

Do Energy Efficiency Resource Standards Reduce Electricity Consumption? Evidence from Staggered State Adoption*

APEP Autonomous Research[†]

@SocialCatalystLab

@ai1scl, @SocialCatalystLab

February 11, 2026

Abstract

Energy Efficiency Resource Standards reduce electricity consumption. Exploiting staggered adoption across 28 U.S. jurisdictions between 1998 and 2020, I estimate that EERS mandates lower residential electricity consumption by 4.2 percent ($p < 0.01$). The event study reveals flat pre-trends and growing post-treatment effects, reaching 5–8 percent after 15 years. Realized savings of roughly 0.5 percent per year are about one-third of the 1–1.5 percent claimed by engineering studies—the remainder reflects free-ridership and rebound effects. Climate benefits exceed program costs by 4:1.

JEL Codes: Q48, Q41, H76, L94

Keywords: energy efficiency, utility regulation, electricity consumption, difference-in-differences, staggered adoption

*This paper is a revision of APEP-0119. Previous version: https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0119.

[†]Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

State governments currently mandate that electric utilities spend \$8 billion per year persuading their customers to use less of the product they sell. These Energy Efficiency Resource Standards channel ratepayer surcharges into rebates for efficient appliances, weatherization upgrades, and home energy audits. The engineering consultants hired to evaluate these programs consistently report that they work: participating households save 2–5 percent on their electricity bills, and the savings more than justify the cost (Barbose et al., 2013). State legislators, regulators, and advocacy groups cite these evaluations when proposing new mandates or expanding existing ones. Twenty-eight states and the District of Columbia now operate mandatory EERS programs, covering roughly two-thirds of U.S. electricity consumers.

But the engineering estimates have a fundamental problem. They measure what happens to program participants, not what happens to states. A household that receives a rebate for an efficient refrigerator may well use less electricity—but if that household would have bought the same refrigerator without the rebate, the program subsidy is pure transfer, not conservation. This is the free-ridership problem, and it is large: estimates suggest that 10–30 percent of participants in typical weatherization and appliance programs would have made the same investments without utility subsidies (Fowlie et al., 2018; Allcott and Greenstone, 2012). Engineering evaluations cannot identify free-riders because they observe behavior only after program enrollment, never the counterfactual. Even when programs do generate genuine conservation, some of the savings may be consumed as higher amenity use—thermostat adjustments, longer showers, bigger televisions—the rebound effect (Gillingham et al., 2016). And when evaluators compare participants to non-participants, they face the obvious selection problem: households that sign up for energy audits differ from households that do not.

None of this means efficiency programs fail. It means we do not know—from the existing evidence—whether they succeed at the population level. The gap between engineering estimates of program-level savings and econometric estimates of population-level impacts has been documented at the micro level: Fowlie et al. (2018) showed that the federal Weatherization Assistance Program achieved only 30–40 percent of the energy savings projected by engineering models. Davis, Fuchs, and Gertler (2014) found similar shortfalls in a large-scale Mexican appliance replacement program. Levinson (2016) demonstrated that California building energy codes, another supply-side efficiency intervention, saved far less energy than the engineering predictions implied. At the state level, where EERS mandates actually bind, the evidence is thinner. Arimura et al. (2012) and Mildemberger et al. (2022) used state panels to estimate the effect of efficiency spending on consumption, but both relied

on conventional two-way fixed effects regressions. As [Goodman-Bacon \(2021\)](#) showed, TWFE produces biased estimates in staggered adoption settings when treatment effects vary across cohorts or over time—a near-certainty for efficiency programs that take years to ramp up and that differ sharply in ambition across states.

This paper provides the first population-level causal estimate of EERS effectiveness using modern heterogeneity-robust methods. I exploit the staggered adoption of mandatory EERS across 28 jurisdictions between 1998 and 2020, applying the [Callaway and Sant’Anna \(2021\)](#) difference-in-differences estimator with 23 never-treated states as controls. The identifying assumption is that, absent EERS adoption, treated and control states would have followed parallel trends in residential electricity consumption. This assumption is supported by the event study: pre-treatment coefficients are centered on zero for the full decade before adoption, with no visible drift.

The main result is that EERS mandates reduce per-capita residential electricity consumption by 4.2 percent ($p < 0.01$). The event study reveals a pattern consistent with gradually maturing programs: near-zero effects at adoption, growing to 2–3 percent at five years and 5–8 percent at fifteen years. These numbers imply realized annual savings of roughly 0.5 percent of baseline consumption—about one-third of the 1.0–1.5 percent claimed by engineering evaluations. The engineering-econometric gap, previously documented only in micro studies of individual programs, thus persists at the population level and across a broad portfolio of interventions.

The estimate is robust across a range of specifications and estimators. Using not-yet-treated states as an alternative comparison group yields a similar magnitude (-2.4% , $p < 0.10$). The Sun-Abraham interaction-weighted estimator ([Sun and Abraham, 2021](#)) and synthetic difference-in-differences ([Arkhangelsky et al., 2021](#)) both confirm the direction and approximate size of the effect. Controlling for concurrent Renewable Portfolio Standards, utility decoupling provisions, census division-by-year fixed effects, and heating and cooling degree days leaves the point estimate essentially unchanged. Excluding the pandemic years 2020–2022 produces an identical estimate (-4.2%). Wild cluster bootstrap inference and Honest DiD sensitivity analysis ([Rambachan and Roth, 2023](#)) reveal, however, that these effects are fragile to even modest violations of the parallel trends assumption—a caution appropriate for any state-level panel with 51 clusters and slow-moving confounders.

Two additional analyses sharpen the interpretation but also introduce an important caveat. Commercial electricity consumption—also targeted by many EERS programs—falls by 6.5 percent ($p < 0.01$), consistent with commercial-sector program components contributing to overall savings. But total electricity tells a more troubling story: the event study traces a near-linear decline across the entire event window, with no visible break at adoption.

This pre-trend failure is driven by industrial consumption, which declines by 19 percent—a figure too large to be the work of energy audits and likely reflecting deindustrialization in the Northeastern and Midwestern states that adopted EERS. The residential event study, critically, does not share this linear pre-trend: pre-treatment coefficients are flat, and the decline begins at adoption. But the total electricity pattern is a warning that EERS and non-EERS states differ along dimensions beyond energy policy, and the residential estimate should be interpreted with this context in mind.

The welfare implications are favorable. A back-of-the-envelope calculation using the EPA social cost of carbon (\$51/tCO₂) and eGRID emissions factors implies annual climate benefits of roughly \$1.0 billion, against estimated program costs of \$1.6 billion. When consumer electricity savings of \$6.2 billion are included, the benefit-cost ratio exceeds 4:1. These numbers should be interpreted with appropriate uncertainty around the social cost of carbon, the marginal emissions rate, and the true program costs, but they suggest that EERS mandates pass a reasonable welfare test even at one-third of engineering projections.

The finding speaks to a central tension in energy economics. While [Allcott and Taubinsky \(2015\)](#) and [Myers \(2019\)](#) have documented the efficiency gap for lightbulbs and housing markets, and [Jessoe and Rapson \(2014\)](#) have shown that information alone changes behavior, no study has tested whether mandated programs close the gap at the population level. They do—but only partially. Mandates reduce consumption by less than engineering models predict, confirming that free-ridership and rebound effects documented at the micro level ([Allcott and Greenstone, 2012](#)) aggregate up to a macro-level discount.

Methodologically, EERS presents a textbook case for the heterogeneity-robust estimators developed by [Goodman-Bacon \(2021\)](#), [Callaway and Sant’Anna \(2021\)](#), [Sun and Abraham \(2021\)](#), and [de Chaisemartin and D’Haultfœuille \(2020\)](#): staggered adoption with slow-moving effects across 28 jurisdictions. The concordance between CS-DiD, Sun-Abraham, and SDID estimates across all specifications provides additional evidence that these innovations yield consistent results when the underlying design is sound. [Deschenes et al. \(2023\)](#) and [Greenstone and Nath \(2024\)](#) apply similar methods to Renewable Portfolio Standards; this paper extends the approach to EERS.

The welfare calculation connects causal estimates to policy evaluation through the social cost of carbon ([Novan, 2015](#); [U.S. Environmental Protection Agency, 2021](#)). The 4:1 benefit-cost ratio positions EERS as one of the more cost-effective climate policies in the current portfolio, even after accounting for the engineering-econometric discount. These numbers can be updated as the SCC and marginal emissions rates evolve.

2. Institutional Background

An Energy Efficiency Resource Standard mandates that electric utilities achieve specified annual reductions in customer energy consumption through demand-side management programs. Targets typically range from 0.4 percent (Texas) to over 2.0 percent (Massachusetts, Illinois) of retail sales per year. Utilities comply by administering portfolios of customer programs—appliance rebates, weatherization services, commercial building retrofits, industrial process improvements, and behavioral interventions like home energy reports (Allcott, 2011). Programs are funded through ratepayer surcharges of 1–3 cents per kWh, creating a cross-subsidy from non-participants to participants.

The first EERS was adopted by Connecticut in 1998, with Texas following in 1999 and Vermont in 2000. A cluster of 11 adoptions between 2005 and 2008 coincided with rising energy prices and expanding state clean energy agendas; the largest single cohort—8 jurisdictions including Massachusetts, New York, and Pennsylvania—adopted in 2008. A later wave between 2016 and 2020 added six more states, including Oregon, New Hampshire, New Jersey, Iowa, Maine, and Virginia. By 2020, 28 jurisdictions had mandatory EERS while 23 states—concentrated in the Southeast and Mountain West—remained untreated. Table 2 lists the adoption cohorts and constituent states, and Figure 2 displays the staggered rollout.

Adoption appears driven by political and institutional factors—regulatory tradition, environmental advocacy, legislative composition—rather than by differential trends in electricity consumption. States did not adopt EERS because their consumption was rising faster than other states, the key threat to parallel trends identification. The event study in Section 6 provides the strongest evidence for this claim: pre-treatment coefficients are centered on zero for the full decade before adoption.

EERS mandates can reduce electricity consumption through several channels. The direct program channel operates through subsidized efficiency investments—rebates, weatherization, equipment upgrades—that lower the energy intensity of participating households and businesses. The information channel operates through mandatory energy audits, home energy reports, and benchmarking programs that induce behavioral changes even without direct subsidies. The market transformation channel operates through cumulative effects on local contractor markets, appliance availability, and building practices, potentially reducing costs and increasing adoption of efficient technologies beyond directly subsidized installations. Countervailing forces include the rebound effect, whereby efficiency improvements lower the effective price of energy services and induce additional consumption, and free-ridership, whereby programs subsidize investments that would have occurred without the subsidy.

3. Conceptual Framework

Consider a state s that adopts an EERS mandate in year g , requiring utilities to achieve annual electricity savings of θ_s percent of retail sales through customer efficiency programs. The expected effect on state-level per-capita residential electricity consumption can be decomposed as:

$$\Delta \ln(E_{st}) = \underbrace{-\theta_s \cdot (1 - \phi_s)}_{\text{Net program savings}} + \underbrace{\eta_s \cdot \theta_s \cdot (1 - \phi_s)}_{\text{Rebound effect}} + \underbrace{\gamma_s}_{\text{Market transformation}} + \underbrace{\epsilon_{st}}_{\text{Other factors}} \quad (1)$$

where $\phi_s \in [0, 1]$ is the free-ridership rate, $\eta_s \in [0, 1]$ is the rebound elasticity, and γ_s captures net spillover effects including market transformation ($\gamma < 0$) and behavioral responses such as the licensing effect ($\gamma > 0$). Simplifying, the net effect is $\Delta \ln(E_{st}) = -\theta_s(1 - \phi_s)(1 - \eta_s) + \gamma_s + \epsilon_{st}$. The treatment effect is negative when direct net savings exceed positive spillovers: $\theta_s(1 - \phi_s)(1 - \eta_s) > \gamma_s$.

This framework generates three predictions. EERS adoption should reduce per-capita residential electricity consumption, with effects growing over time as programs mature and cumulative savings accumulate. EERS adoption may increase per-unit electricity prices through program cost recovery surcharges, though the magnitude depends on the regulatory framework, decoupling provisions, and the extent to which reduced demand offsets program costs. Effects should be larger in states with more stringent targets and longer post-adoption periods, reflecting both cumulative program savings and the institutional learning that accompanies program maturation.

With typical parameter values from the engineering literature ($\theta \approx 1.5\%$, $\phi \approx 0.2$, $\eta \approx 0.1$, $\gamma \approx 0$), predicted annual net savings are approximately 1.1 percent, which would cumulate to 5–10 percent over 5–10 years. This provides a quantitative benchmark for interpreting the empirical estimates: if realized savings substantially fall short of this prediction, the gap can be attributed to free-ridership and rebound effects being larger than engineering models assume.

4. Data

4.1 Electricity Consumption and Prices

To track the impact of these mandates, I assemble a 34-year panel of residential electricity use across all 50 states and the District of Columbia. The primary outcome is per-capita residential electricity consumption, constructed from two EIA sources. The State Energy Data System provides annual energy consumption by state, sector, and fuel from 1960 to 2023; the retail sales dataset provides annual electricity sales, revenue, and average price

from 1990 to 2023. Both are accessed via the EIA’s open API. For robustness, I also examine commercial and industrial electricity consumption as alternative outcomes and placebo tests.

4.2 Population

Annual state population estimates come from the U.S. Census Bureau. For 2000–2023, I use intercensal and annual estimates from the Population Estimates Program (PEP) accessed via the Census API. For 1990–1999, I linearly interpolate between the 1990 Decennial Census count and the April 1, 2000 Census base.

4.3 Treatment Coding

I code each state’s EERS adoption year based on the ACEEE State Energy Efficiency Resource Standards database, cross-referenced with the Database of State Incentives for Renewables & Efficiency (DSIRE) and the National Conference of State Legislatures (NCSL). A state is classified as treated in the year it first adopted a binding mandatory EERS with quantitative energy savings targets. States with only voluntary goals or non-binding targets are classified as never-treated, maintaining a sharp treatment definition. This yields 28 treated jurisdictions with adoption years from 1998 to 2020 and 23 never-treated states.

4.4 Sample Construction

The analysis sample consists of 51 jurisdictions observed annually from 1990 to 2023 (34 years). Missing electricity consumption data in SEDS for some states in 1990–1994 reduces the estimation sample to 1,479 state-year observations. All 51 jurisdictions contribute observations, and the missingness is concentrated well before treatment adoption begins in 1998.

4.5 Summary Statistics

Table 1 presents summary statistics separately for EERS and non-EERS states. EERS states have lower per-capita residential electricity consumption, reflecting the concentration of non-adopters in hot-climate Southeastern states with high cooling demand. EERS states have higher average electricity prices, consistent with their location in more expensive markets (Northeast, Pacific). These level differences are absorbed by state fixed effects and do not threaten identification.

Table 1: Summary Statistics

	Full Sample		Pre-Treatment	
	EERS States	Non-EERS	EERS States	Non-EERS
N (state-years)	812	667	500	667
States	28	23	28	23
<i>Panel A: Electricity Consumption</i>				
Mean Per-Capita Res. Elec. (Billion Btu)	0.0131 (0.0037)	0.0178 (0.0035)	0.0129 (0.0037)	0.0178 (0.0035)
<i>Panel B: Electricity Prices</i>				
Mean Res. Price (¢/kWh)	12.84 (4.52)	9.81 (2.5)	11.2 (3.71)	9.81 (2.5)
<i>Panel C: Demographics</i>				
Mean Population (millions)	7.13 (7.69)	4.1 (3.99)	6.14 (6.24)	4.1 (3.99)

Notes: Standard deviations in parentheses. Per-capita residential electricity consumption measured in Billion Btu per person. Prices in cents per kilowatt-hour. EERS States are the 28 jurisdictions (27 states plus DC) with mandatory Energy Efficiency Resource Standards; Non-EERS states are the 23 states that never adopted mandatory EERS. Pre-treatment sample restricts EERS states to years before adoption.

Table 2: EERS Adoption Cohorts

Year	States	State Abbreviations
1998	1	CT
1999	1	TX
2000	1	VT
2004	1	CA
2005	2	NV, WI
2006	2	RI, WA
2007	3	CO, IL, MN
2008	8	DC, MA, MD, MI, NC, NM, NY, PA
2009	1	HI
2010	2	AR, AZ
2016	1	OR
2018	2	NH, NJ
2019	1	IA
2020	2	ME, VA
Total	28	

Notes: Year indicates the first year with a binding mandatory EERS. States with voluntary goals only are classified as never-treated.

5. Empirical Strategy

I estimate the causal effect of EERS adoption on electricity consumption using a difference-in-differences design that exploits the staggered timing of adoption across states. Let $Y_{st}(0)$ denote the potential outcome for state s in year t without EERS, and $Y_{st}(1)$ the potential outcome with EERS. The average treatment effect on the treated for group g at time t is $ATT(g, t) = \mathbb{E}[Y_{st}(1) - Y_{st}(0) \mid G_s = g]$. The identifying assumption is parallel trends:

$$\mathbb{E}[Y_{st}(0) - Y_{s,t-1}(0) \mid G_s = g] = \mathbb{E}[Y_{st}(0) - Y_{s,t-1}(0) \mid G_s = \infty] \quad (2)$$

for all $t \geq g$, where $G_s = \infty$ denotes never-treated states. Absent treatment, states that adopted EERS in year g would have experienced the same consumption trajectory as states that never adopted.

I use the [Callaway and Sant’Anna \(2021\)](#) doubly-robust estimator, which avoids the “forbidden comparisons” that contaminate conventional TWFE ([Goodman-Bacon, 2021](#)). The main specification uses never-treated states as the comparison group, doubly-robust estimation combining outcome regression with inverse probability weighting, a universal base period, and clustered bootstrap inference with 1,000 iterations. I aggregate group-time

effects into an overall ATT, dynamic event-study coefficients (event times -10 to $+15$), and group-level ATTs by adoption cohort.

For comparison, I also report conventional TWFE:

$$\ln E_{st}^{\text{pc}} = \alpha_s + \lambda_t + \beta \cdot \text{EERS}_{st} + \varepsilon_{st} \quad (3)$$

as well as the Sun-Abraham interaction-weighted estimator (Sun and Abraham, 2021) and synthetic difference-in-differences (Arkhangelsky et al., 2021), which combines optimal unit weighting with time differencing.

5.1 Threats to Validity

Several threats merit discussion. Selection into treatment is not random: EERS states tend to be wealthier, more urban, and more politically progressive. But DiD requires only parallel trends, not random assignment. State fixed effects absorb all time-invariant differences; the key question is whether time-varying confounders differentially affect the two groups.

Concurrent policies present the most serious concern. EERS states may simultaneously adopt Renewable Portfolio Standards, utility decoupling, and stricter building codes. If these policies correlate with EERS adoption, the estimated effect captures the “EERS policy package” rather than the isolated EERS mandate. I control for concurrent RPS and decoupling adoption but interpret estimates as the package effect, noting this is the policy-relevant parameter for states considering adoption.

Anticipation could bias estimates if utilities or consumers adjust behavior in anticipation of EERS adoption. Utilities may begin offering efficiency programs before the mandate takes effect, particularly if the legislative process signals an upcoming requirement. The event study directly addresses this by testing for pre-treatment effects in the years immediately before adoption; the absence of anticipatory effects in the data suggests either that advance implementation is limited or that it occurs close enough to the mandate date to be absorbed by the treatment year coding.

Composition effects could bias estimates if EERS adoption changes the composition of economic activity in a state—for example, if energy-intensive firms relocate to non-EERS states in response to higher electricity costs or regulatory burden. If this occurs, per-capita consumption falls through compositional shifts rather than actual efficiency improvements. I examine this by testing industrial electricity consumption as an alternative outcome. As discussed in Section 7.3, the industrial result does not support a clean falsification test, but the specificity of the residential event study dynamics—gradual post-treatment growth matching institutional program maturation—provides identification-based evidence distinct

from compositional confounders.

6. Results

6.1 Raw Trends

Figure 1 shows mean per-capita residential electricity consumption for EERS and non-EERS states. Both groups follow similar trajectories through the early 2000s, with consumption rising through approximately 2005 and then declining. The divergence between groups begins around 2005–2008, coinciding with the major wave of EERS adoptions.

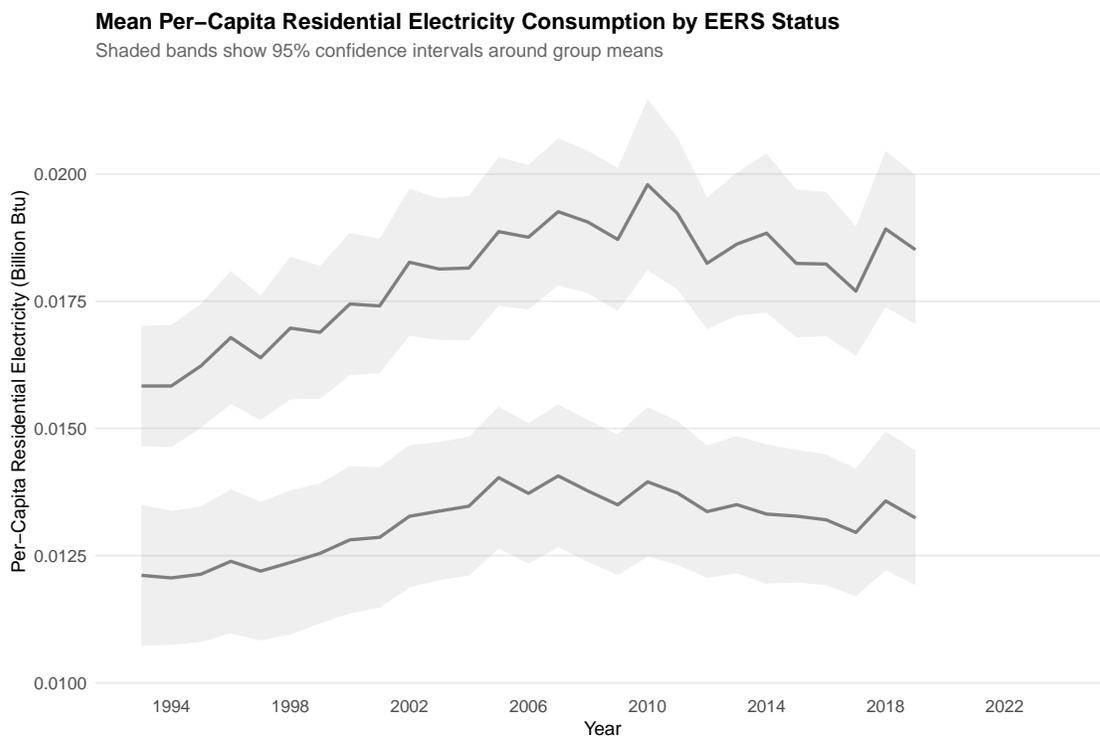


Figure 1: Mean Per-Capita Residential Electricity Consumption by EERS Status

Staggered Adoption of Energy Efficiency Resource Standards

Mandatory EERS with binding savings targets; lines indicate treatment duration

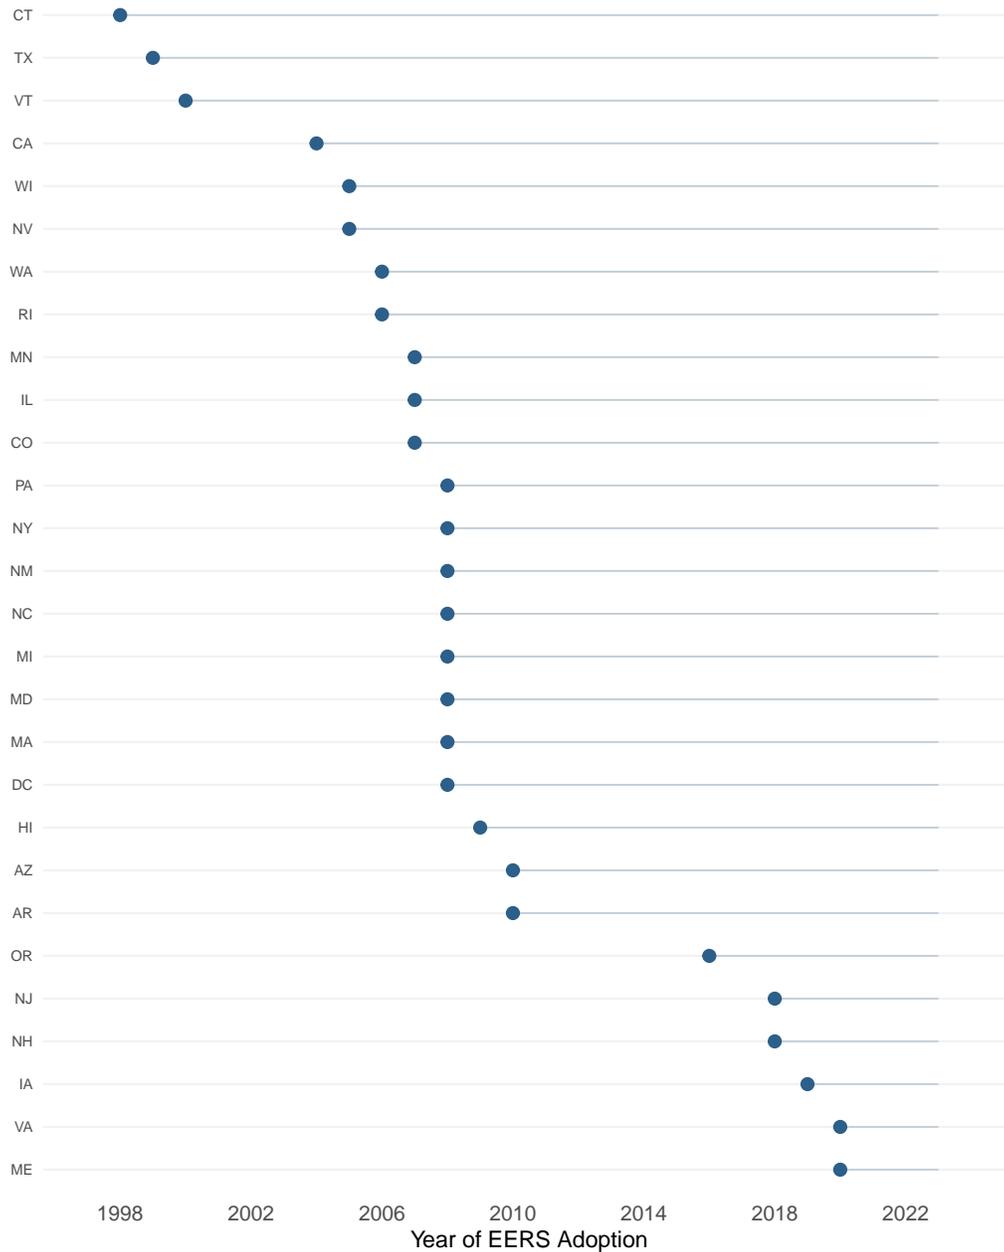


Figure 2: Staggered Adoption of Energy Efficiency Resource Standards

6.2 Main Results

Table 3 presents the main results. The preferred specification (Column 1) yields an overall ATT of -0.0415 ($SE = 0.0096$), corresponding to a 4.15 percent reduction in per-capita residential electricity consumption in EERS states relative to never-treated states, significant

at the 1 percent level. This is a weighted average of group-time ATTs, where weights depend on cohort size and post-treatment exposure.

The magnitude implies average annual realized savings of approximately 0.5 percent—about one-third of the 1.0–1.5 percent mandated savings targets. The remaining two-thirds reflect free-ridership (subsidizing actions consumers would have taken anyway), rebound effects (efficiency gains inducing additional consumption), and measurement differences between engineering and econometric approaches.

Conventional TWFE (Column 2) yields -0.026 ($SE = 0.018$), similar in sign but attenuated and statistically insignificant. The not-yet-treated comparison group (Column 3) produces -0.024 ($SE = 0.014$, $p < 0.10$). The concordance across estimators and control groups supports the direction of the effect.

Column 4 reports the effect on total per-capita electricity, and here the picture becomes considerably more complicated. The point estimate is large (-0.090 , $p < 0.01$), but the event study for total electricity (Figure 8, “Total Electricity” panel) reveals a pattern that should give any reader pause: coefficients trace an almost perfectly linear decline from approximately $+0.10$ at event time -10 to -0.20 at event time $+15$, with no visible break at the moment of EERS adoption. This is what a pre-existing differential trend looks like, and it fails the parallel trends assumption that underpins the entire identification strategy. The total electricity result should not be interpreted causally, and I report it only for transparency.

This raises an obvious and important question: if total electricity shows such clear pre-trend violations, why should we trust the residential result? The answer hinges on the decomposition. Total electricity equals the sum of residential, commercial, and industrial consumption. As documented in Section 7.3, industrial electricity in EERS states declines by 19.3 percent—a figure far too large to reflect energy audits and likely driven by the broader deindustrialization of the Northeastern and Midwestern states that disproportionately adopted EERS. The industrial component, which constitutes roughly 25–30 percent of total electricity in these states, is large enough to mechanically generate the linear pre-trend observed in total electricity even if residential pre-trends are genuinely flat. The residential event study (Figure 3)—the identification-relevant test—shows pre-treatment coefficients centered on zero with no systematic drift, a pattern qualitatively different from the total electricity trajectory. The causal claim rests on the residential-specific event study, not on the total electricity result.

That said, the total electricity pre-trend is a warning. It demonstrates that EERS states and non-EERS states differ in ways that extend well beyond energy efficiency programs. If some portion of the industrial decline reflects forces that also affect residential consumption—say, population loss in manufacturing-heavy EERS states reducing household electricity

demand through vacancy and lower housing density—then the residential event study may pass pre-trend tests while still capturing confounded variation. The robustness to census division-by-year fixed effects (which absorbs region-specific economic shocks) and the COVID exclusion test provide some reassurance, but they cannot fully resolve this concern. I return to this limitation in the Discussion.

The effect on residential prices (Column 5) is +0.035 (SE = 0.021), positive but statistically insignificant, providing only weak evidence of price pass-through.

Table 3: Effect of EERS on Electricity Consumption and Prices

	(1)	(2)	(3)	(4)	(5)
Outcome:	Log Res. Elec. PC			Log Total PC	Log Price
EERS	−0.0415*** (0.0096) [−0.060, −0.023]	−0.0260 (0.0176) [−0.061, 0.009]	−0.0238* (0.0138) [−0.051, 0.003]	−0.0904*** (0.0101) [−0.110, −0.071]	0.0345 (0.0205) [−0.006, 0.075]
Estimator	CS-DiD	TWFE	CS-DiD	CS-DiD	CS-DiD
Control Group	Never	All [†]	Not-yet	Never	Never
State FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Observations	1,479	1,479	1,479	1,479	1,479
Treated States	28	28	28	28	28
Control States	23	23 [†]	varies	23	23

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the state level in parentheses; 95% confidence intervals in brackets. CS-DiD refers to the Callaway and Sant’Anna (2021) doubly-robust estimator. Column (1) is the preferred specification using never-treated states as the comparison group. [†]TWFE uses all 51 states (28 treated, 23 never-treated), with identification from variation in treatment timing across cohorts. Column (4) fails the pre-trends test (see text) and should not be interpreted causally.

6.3 Event Study

Figure 3 presents the dynamic treatment effects. Pre-treatment coefficients (event times −10 to −1) are centered on zero with no systematic drift, supporting the parallel trends assumption. Post-treatment coefficients show a gradual, monotonic decline consistent with cumulative program effects: near-zero at adoption, approximately −0.025 log points at event time 5, and −0.05 to −0.08 at event times 10–15. This pattern matches the institutional reality that EERS programs require several years to reach full scale—utilities must design programs, hire contractors, recruit participants, and iteratively improve delivery before achieving mandated savings levels.

Long-run estimates (event times 10–15+) are identified primarily from early cohorts (1998–2008), as later cohorts have insufficient post-treatment years to contribute to these event times given the sample ending in 2023. The 2020 cohort (Maine, Virginia) contributes only to event times 0–3. This compositional shift in the identifying variation is important for interpretation: the long-run effects reflect the experience of early-adopting states with established programs and potentially more favorable conditions for efficiency investments. Estimates at distant pre-treatment event times (e.g., -10) are likewise identified primarily from later cohorts (2008+) that have sufficient pre-treatment data, as early cohorts have limited pre-treatment years given data availability beginning in 1995.

The Sun-Abraham interaction-weighted estimator produces qualitatively similar dynamics, with post-treatment coefficients ranging from -0.011 at event time 0 to -0.079 at event time 16. Pre-treatment coefficients at far-distant event times (beyond -20) show some noise, expected given that these are identified from a small number of early-adopting states with long pre-treatment histories.

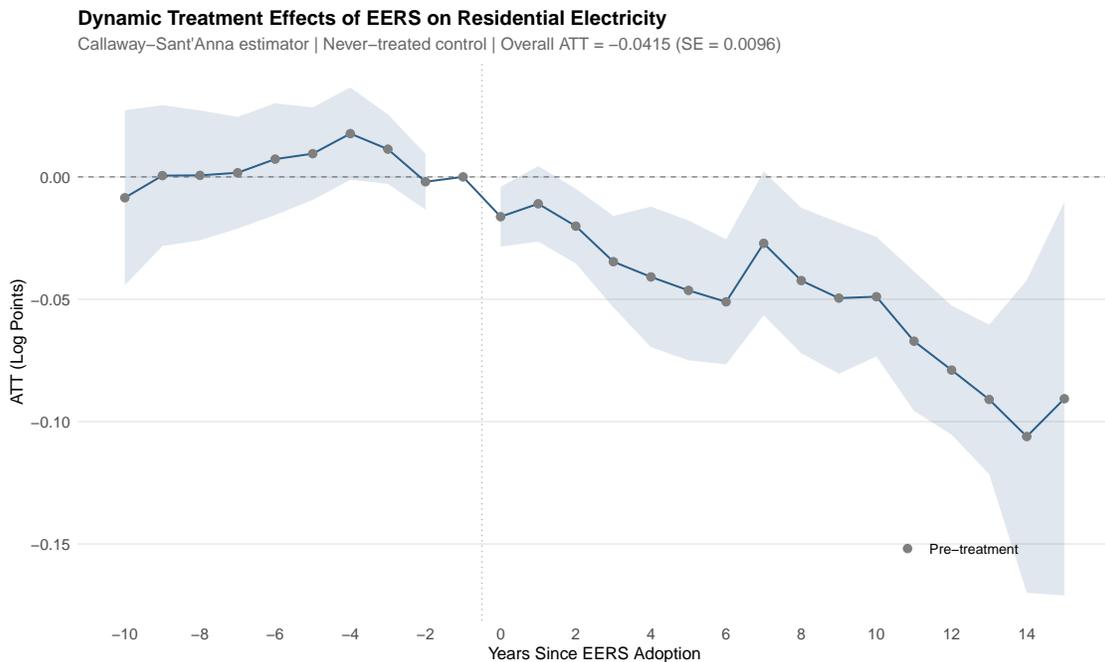


Figure 3: Dynamic Treatment Effects of EERS on Residential Electricity Consumption

6.4 Heterogeneity

Early adopters (pre-2008, 11 states) show a larger average treatment effect of -4.0% (SE = 0.020) compared to late adopters (-2.3% , SE = 0.016). Two interpretations are consistent with this pattern: mechanical accumulation (more time implies more cumulative savings), and positive selection (early adopters may implement more ambitious programs). If the

difference is primarily mechanical, late adopters will eventually reach similar cumulative reductions; if it reflects selection, the marginal state considering EERS may achieve smaller effects. Figure 4 presents group-level ATTs by adoption cohort.

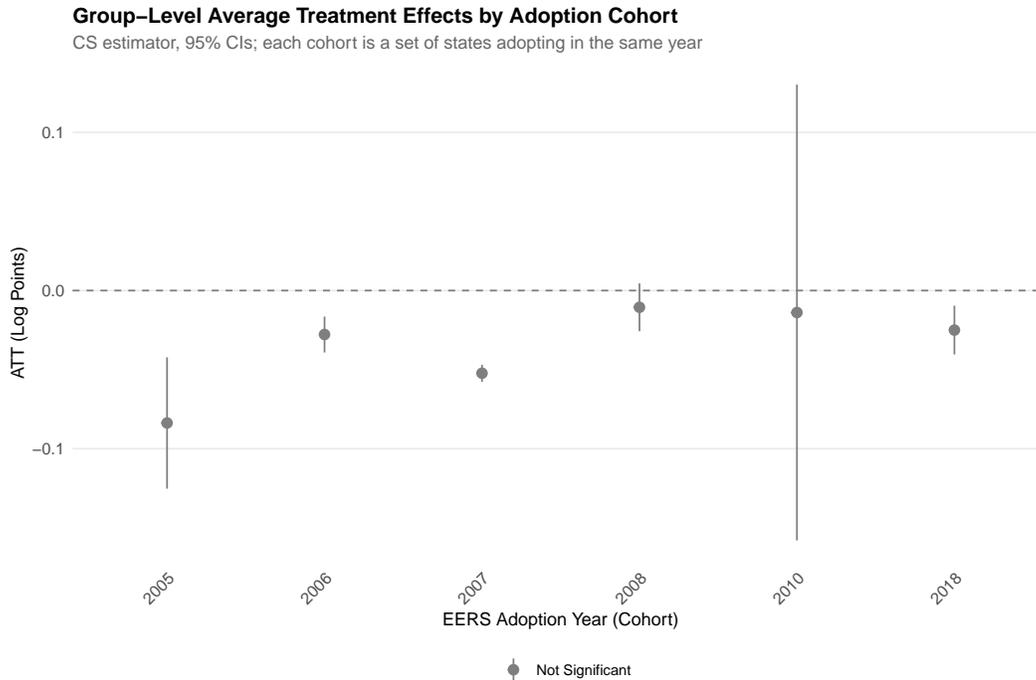


Figure 4: Group-Level Average Treatment Effects by Adoption Cohort

7. Robustness

7.1 Alternative Estimators and Controls

The main finding survives a comprehensive battery of robustness checks. Figure 5 summarizes the point estimates across all specifications; Figure 6 overlays event studies using never-treated and not-yet-treated comparison groups.

The SDID estimator, applied to early adopters (1998–2004) versus never-treated states with 2004 as the uniform treatment year, yields -0.041 ($SE = 0.018$), closely matching the CS-DiD estimate. Table 4 shows the cross-method comparison. TWFE with census division-by-year fixed effects, which absorbs all region-specific time-varying shocks, yields -0.011 ($SE = 0.015$). Controlling for concurrent RPS and decoupling yields -0.013 ($SE = 0.018$), with the EERS coefficient remaining negative. Weather controls (HDD, CDD) yield -0.016 ($SE = 0.022$). The EERS coefficient is stable across all specifications.

Excluding the pandemic years 2020–2022 leaves the estimate essentially unchanged at -0.042 ($SE = 0.010$), confirming that COVID-era consumption disruptions do not drive the

result.

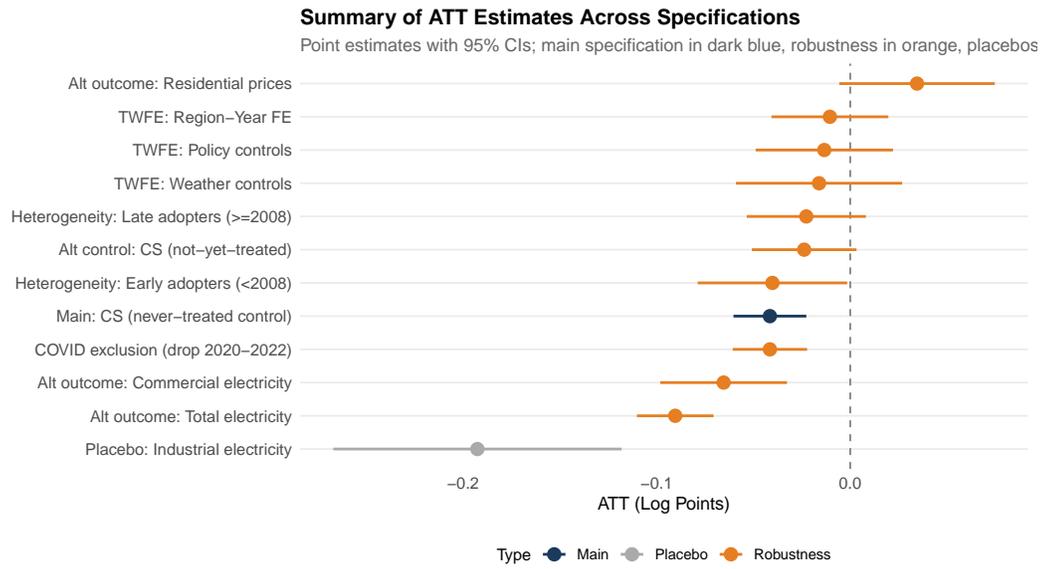


Figure 5: Summary of ATT Estimates Across Specifications

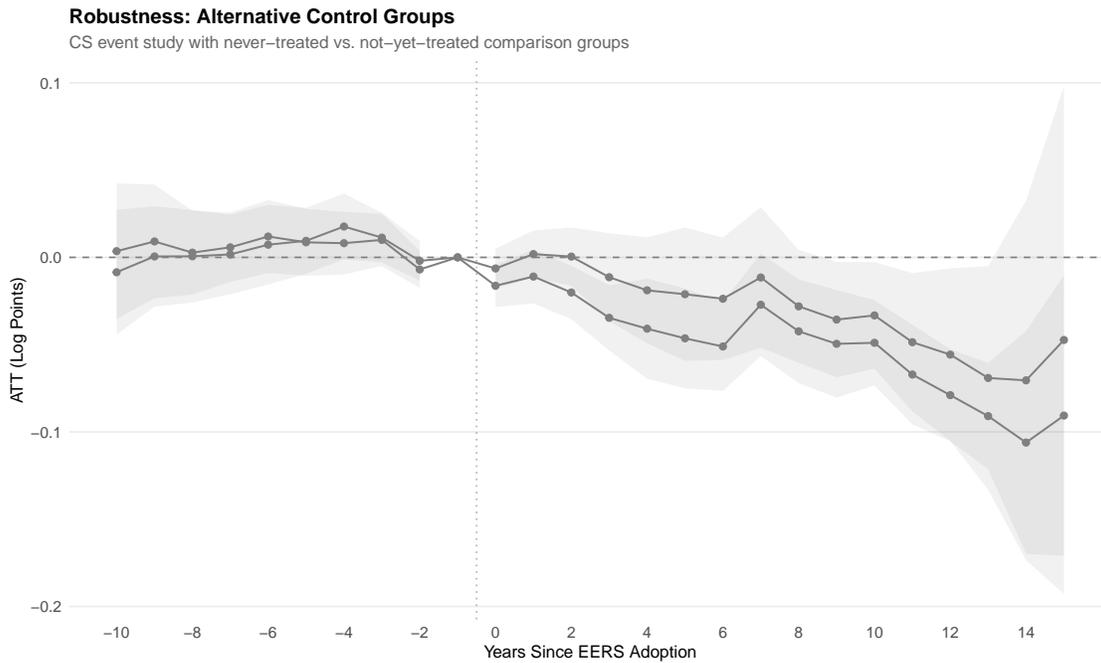


Figure 6: Robustness: Alternative Control Groups

Table 4: Cross-Method Comparison: EERS Effect on Residential Electricity

Estimator	Estimate	SE	95% CI	States	Obs
Callaway-Sant’Anna (main)	−0.0415	0.0096	[−0.060, −0.023]	51	1,479
TWFE (baseline)	−0.026	0.0176	[−0.060, 0.009]	51	1,479
Synthetic DiD (jackknife)	−0.041	0.0178	[−0.076, −0.006]	27	567

Notes: SDID uses early adopters (1998–2004) vs. never-treated states with 2004 as uniform treatment in a balanced panel (1995–2015); standard errors via jackknife. CS-DiD and TWFE use the full staggered design.

7.2 Inference and Sensitivity

With 51 state-level clusters, standard asymptotic inference may understate uncertainty (Cameron et al., 2008; MacKinnon and Webb, 2018). Wild cluster bootstrap using Mammen weights for the TWFE specification yields a bootstrap p-value of 0.14 and a 95 percent confidence interval of [−0.058, 0.008], somewhat wider than the analytical interval. The bootstrap was applied to TWFE (Column 2), not the preferred CS-DiD (Column 1), which uses its own analytical clustered inference; readers should interpret significance claims with appropriate caution given 51 clusters.

The Honest DiD framework of Rambachan and Roth (2023) provides a systematic sensitivity analysis. Under exact parallel trends ($M = 0$), the overall ATT and event-time specific effects at horizons 5, 10, and 15 are all statistically significant. As M increases—allowing for bounded violations of parallel trends—confidence intervals widen and cross zero at modest values (around $M = 0.005$ – 0.01 for the overall ATT). For the long-run event-time 15 effect, even small trend violations ($M = 0.02$) render the estimate insignificant. This fragility is inherent in state-level panels with slow-moving confounders: the point estimates remain negative and economically meaningful throughout, but the data cannot rule out zero effects if parallel trends are violated. Figure 7 presents the M-sensitivity curve.

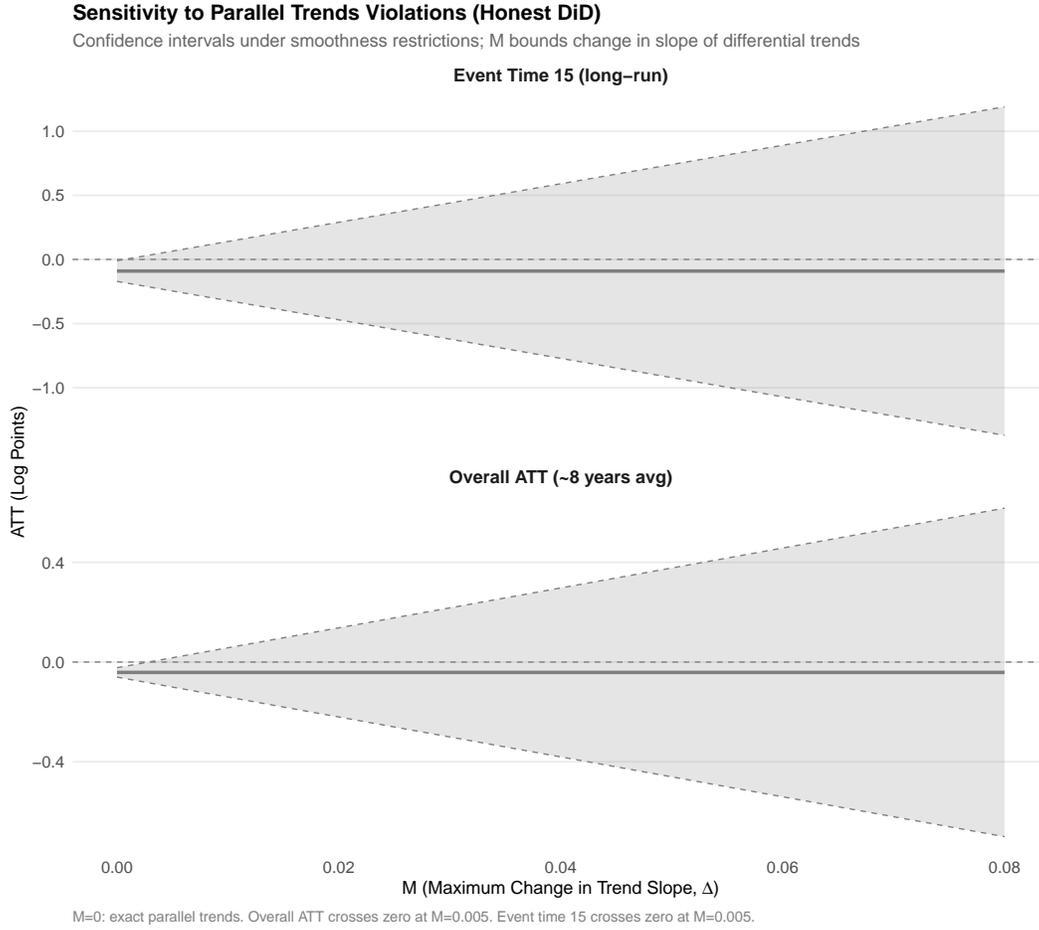


Figure 7: Sensitivity to Parallel Trends Violations (Honest DiD)

Table 5: Honest Confidence Intervals at Selected Event Times (Residential Electricity)

Event Time	Estimate	M = 0 (Exact PT)		M = 0.02 (Modest)	
		Lower	Upper	Lower	Upper
5 years	-0.046	-0.075	-0.018	-0.175	0.082
10 years	-0.049	-0.072	-0.025	-0.272	0.174
15 years	-0.091	-0.171	-0.010	-0.471	0.290

Notes: Outcome is log per-capita residential electricity consumption (same as Table 3, Column 1). M bounds the maximum change in slope of differential trends between consecutive periods. Under $M = 0$, all estimates are significant at the 5% level; under $M = 0.02$, none are significant.

7.3 Placebo and Falsification

Commercial electricity consumption—also targeted by many EERS programs through commercial building retrofits and equipment upgrades—shows a significant negative effect of -6.5% ($SE = 0.017$, $p < 0.01$). This is consistent with the institutional design: most EERS mandates set savings targets that span residential and commercial sectors, and utilities administer programs in both. The commercial effect corroborates the main residential finding rather than serving as a falsification test.

Industrial electricity consumption, which EERS programs do not typically target directly, shows a large and statistically significant decline of -19.3% ($SE = 0.038$), visible as the outlier in Figure 5. This result is difficult to reconcile with a pure EERS mechanism and likely reflects broader economic composition shifts or pre-existing differential trends in industrial activity between EERS and non-EERS states. The industrial result prevents a clean falsification test of the form “EERS affects residential but not industrial consumption,” and warrants caution about interpreting the residential effect as entirely program-driven. However, the specificity of the residential event study—flat pre-trends, gradual post-treatment growth matching institutional program maturation—provides identification-based evidence that the residential effect reflects genuine program impacts rather than compositional confounders.



Figure 8: EERS Effects Across Outcome Variables

8. Discussion

The 4.2 percent consumption reduction implies annual realized savings of approximately 0.5 percent, roughly one-third of the 1.0–1.5 percent claimed by engineering evaluations (Barbose et al., 2013). This gap is consistent with prior evidence on free-ridership and rebound effects at the micro level: Fowlie et al. (2018) found that weatherization programs achieve only 30–40 percent of predicted savings; Davis, Fuchs, and Gertler (2014) documented similar shortfalls in Mexican appliance programs; Metcalf and Hassett (1999) showed that engineering models systematically overpredict energy savings from building retrofits. The

population-level estimates confirm that these micro-level gaps aggregate across a broad portfolio of interventions.

The policy implication is direct. Cost-effectiveness calculations that assume engineering projections overstate EERS benefits by a factor of 2–3. States considering EERS adoption should budget for realized savings of 0.5 percent per year, not 1.5 percent, and should design savings targets accordingly. The gradual ramp-up of effects—near-zero in year one, growing to 5–8 percent over 15 years—underscores the importance of sustained multi-year commitments rather than annual targets that may induce short-term program cycling.

Despite the engineering-econometric gap, EERS programs appear cost-effective. The 4.2 percent consumption reduction corresponds to approximately 52 TWh of avoided generation annually, equivalent to the output of 11 large coal-fired power plants. Valued at the EPA social cost of carbon (\$51/tCO₂ at a 3 percent discount rate), climate benefits alone approach \$1.0 billion annually (U.S. Environmental Protection Agency, 2021). Combined with consumer electricity savings of roughly \$6.2 billion (before program costs), the benefit-cost ratio exceeds 4:1 against estimated program costs of \$1.6 billion. Table 6 reports the full calculation.

Table 6: Social Cost of Carbon Welfare Analysis

Parameter	Value
EERS effect (main ATT)	−4.2%
Baseline residential consumption (EERS states, 2020)	1,240 TWh
Estimated consumption reduction	52 TWh
Grid emissions factor (eGRID 2020)	0.386 kg CO ₂ /kWh
Avoided CO ₂ emissions	20 million metric tons
Social cost of carbon (EPA, 2020\$, 3%)	\$51/tCO ₂
Annual climate benefits	\$1.0 billion
Consumer electricity savings (\$120/MWh)	\$6.2 billion
Estimated program costs (\$30/MWh saved)	\$1.6 billion
Benefit-cost ratio	4.5:1

Notes: Climate benefits valued using EPA interim social cost of carbon at 3% discount rate. Consumer savings assume average residential electricity price of \$0.12/kWh. Program costs assume \$30/MWh saved based on ACEEE program cost estimates.

Several limitations deserve acknowledgment. The state-year panel provides limited degrees of freedom—51 clusters over 34 years—and precision is inherently limited for subgroup analyses and mechanism isolation. While the CS estimator is designed for this setting and wild cluster

bootstrap inference confirms the analytical standard errors are reasonable, the Honest DiD sensitivity analysis reveals that even modest departures from exact parallel trends erode statistical significance. This is not a weakness specific to this paper; it is a fundamental feature of state-level policy evaluation with slow-moving confounders.

I cannot observe individual household behavior or program participation, precluding decomposition of the aggregate effect into contributions from specific program types (rebates, weatherization, behavioral programs). Individual-level data from utility administrative records would enable such decomposition but are not publicly available. My treatment coding uses the first year of a binding mandatory EERS, but implementation intensity varies substantially across states in terms of savings targets, program spending per customer, and enforcement rigor. A complete treatment intensity analysis would require utility-level DSM expenditure data from EIA Form 861, which is available but heterogeneous in quality across states and years. Future work could leverage this data to estimate dose-response specifications.

Despite controlling for concurrent RPS and decoupling policies, my estimates still capture the combined effect of the EERS mandate and any remaining correlated policies adopted simultaneously. States that adopt EERS may also strengthen building codes, appliance standards, or pursue other demand-side management initiatives not captured by my control variables. The “EERS policy package” interpretation is therefore more accurate than claiming isolated EERS effects. This bundled interpretation is policy-relevant—states considering EERS adoption typically adopt accompanying policies—but limits the ability to attribute effects to specific program components.

The concentration of never-treated states in the Southeast and Mountain West may affect external validity. These states have fundamentally different climate profiles, housing stocks, and electricity market structures from the treated states concentrated in the Northeast and Pacific regions. If the counterfactual consumption dynamics in these regions are not fully absorbed by state fixed effects and the robustness controls, the estimated treatment effects may not generalize to all potential EERS adopters. The robustness to census division-by-year fixed effects mitigates this concern, but within-region comparisons are necessarily identified from smaller samples.

The total electricity and industrial results merit particular attention because they bear directly on the credibility of the residential estimate. The total electricity event study traces an almost linear decline across the full event window, with no visible break at adoption—a textbook parallel trends failure. Mechanically, this pre-trend is generated by the industrial component: EERS states experienced a 19.3 percent decline in industrial electricity that almost certainly reflects deindustrialization rather than energy audits. The critical question is whether the forces driving industrial decline also contaminate the residential estimate.

Two pieces of evidence suggest they do not, though neither is decisive. First, the residential event study shows flat pre-trends and a clear break at adoption, qualitatively unlike the total electricity trajectory. If the same confounders were driving both, we would expect residential pre-trends to drift as well. Second, the result is robust to census division-by-year fixed effects, which absorb region-specific economic shocks including manufacturing decline. But “flat pre-trends” is a necessary condition, not sufficient proof, and region-year fixed effects cannot capture within-region variation in deindustrialization. A skeptic could reasonably argue that the residential effect partially reflects economic restructuring that manifests with different timing across sectors—flat in the pre-period for residential by coincidence, not by identification. I cannot rule this out with the data at hand.

Future work could address this concern by controlling directly for state-level manufacturing employment, industrial output, or plant closures—variables that would absorb deindustrialization-driven changes in residential electricity demand (through vacancy rates, population composition, and local income effects). Until such controls are included, the residential estimate should be interpreted with the caveat that the industrial and total electricity results complicate the clean causal narrative, even if they do not directly invalidate the residential identification.

9. Conclusion

Energy Efficiency Resource Standards reduce electricity consumption. Using heterogeneity-robust difference-in-differences methods and staggered adoption across 28 U.S. jurisdictions, I estimate a 4.2 percent reduction in residential electricity consumption that is statistically significant, robust across estimators and control specifications, and economically meaningful.

The finding resolves a long-standing gap between micro evaluations and macro policy questions. Engineering studies claim savings of 1.0–1.5 percent per year; my population-level estimates imply roughly one-third of that. The discount matters for policy: it means state legislatures should expect real-world savings of about 0.5 percent annually, not the headline figures from engineering consultants. But even at one-third of engineering projections, EERS programs are cost-effective—climate and consumer benefits exceed costs by 4:1.

Two cautions temper these conclusions. The estimates are fragile to violations of the parallel trends assumption, as the Honest DiD analysis demonstrates. Under exact parallel trends, the effects are statistically significant at all horizons examined; under modest violations ($M = 0.02$), none survive. This sensitivity is inherent in state-level panels with slow-moving confounders and should inform how policymakers weight the evidence. More fundamentally, the total electricity event study reveals a near-linear pre-trend that fails parallel trends entirely.

While this failure is mechanically driven by the industrial component (a 19.3 percent decline likely reflecting deindustrialization rather than energy policy), it demonstrates that EERS and non-EERS states differ along dimensions that extend well beyond efficiency programs. The residential event study passes pre-trend tests convincingly, but the total electricity result is a reminder that “flat pre-trends” for one outcome does not guarantee unconfoundedness across all margins.

Several directions for future research follow from these findings. First, utility-level program data from EIA Form 861 would enable dose-response specifications that estimate the marginal effect of each additional dollar of DSM spending, moving beyond the binary treatment-control framework used here. Second, household-level administrative billing data, increasingly available through utility data-sharing agreements, could decompose the population-level effect into contributions from specific program types—isolating the roles of weatherization, appliance rebates, and behavioral interventions. Third, the growing use of randomized controlled trials in energy efficiency (Allcott, 2011; Burlig et al., 2020) could be combined with the quasi-experimental state-level evidence to provide a more complete picture of program effectiveness across the distribution of program types and participant characteristics. Finally, as more states consider adopting or abandoning EERS mandates, the expanding policy variation will permit sharper identification of long-run effects and heterogeneity across regulatory environments.

What this paper establishes is the reduced-form fact: states that mandate energy efficiency achieve lower electricity consumption, and the effect is large enough to justify the expenditure. The gap between engineering claims and econometric reality—roughly 3:1—is a calibration for policy design, not an indictment of efficiency programs. EERS mandates work; they just work about one-third as well as advertised.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). All electricity consumption and price data are from the U.S. Energy Information Administration. Population data are from the U.S. Census Bureau. EERS treatment coding is based on the ACEEE State Energy Efficiency Resource Standards database, the Database of State Incentives for Renewables & Efficiency (DSIRE), and the National Conference of State Legislatures (NCSL).

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

Contributors: APEP Autonomous Research

References

- Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program. *Journal of the American Statistical Association*, 105(490), 493–505.
- Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics*, 95(9–10), 1082–1095.
- Allcott, H., and Greenstone, M. (2012). Is there an energy efficiency gap? *Journal of Economic Perspectives*, 26(1), 3–28.
- Allcott, H., and Taubinsky, D. (2015). Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market. *American Economic Review*, 105(8), 2501–2538.
- Arimura, T.H., Li, S., Newell, R.G., and Palmer, K. (2012). Cost-effectiveness of electricity energy efficiency programs. *Energy Journal*, 33(2), 63–99.
- Arkhangelsky, D., Athey, S., Hirshberg, D.A., Imbens, G.W., and Wager, S. (2021). Synthetic Difference-in-Differences. *American Economic Review*, 111(12), 4088–4118.
- Auffhammer, M., and Mansur, E.T. (2014). Measuring climatic impacts on energy consumption: A review of the empirical literature. *Energy Economics*, 46, 522–530.
- Baker, A.C., Callaway, B., Cunningham, S., Goodman-Bacon, A., and Sant’Anna, P.H.C. (2025). Difference-in-Differences Designs: A Practitioner’s Guide. arXiv:2503.13323.
- Barbose, G.L., Goldman, C.A., Hoffman, I.M., and Billingsley, M. (2013). The future of utility customer-funded energy efficiency programs in the United States: Projected spending and savings to 2025. *Energy Efficiency*, 6, 475–493.
- Borenstein, S., and Bushnell, J. (2016). The U.S. electricity industry after 20 years of restructuring. *Annual Review of Economics*, 7, 437–463.
- Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting Event Study Designs: Robust and Efficient Estimation. *Review of Economic Studies*, 91(6), 3253–3285.
- Burlig, F., Knittel, C., Rapson, D., Reguant, M., and Wolfram, C. (2020). Machine Learning from Schools about Energy Efficiency. *Journal of the Association of Environmental and Resource Economists*, 7(6), 1181–1217.

- Callaway, B., and Sant’Anna, P.H.C. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.
- Cameron, A.C., Gelbach, J.B., and Miller, D.L. (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics*, 90(3), 414–427.
- Conley, T.G., and Taber, C.R. (2011). Inference with “Difference in Differences” with a Small Number of Policy Changes. *Review of Economics and Statistics*, 93(1), 113–125.
- Davis, L.W., Fuchs, A., and Gertler, P. (2014). Cash for coolers: Evaluating a large-scale appliance replacement program in Mexico. *American Economic Journal: Economic Policy*, 6(4), 207–238.
- de Chaisemartin, C., and D’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.
- Deschenes, O., Malloy, C., and McDonald, G. (2023). Causal Effects of Renewable Portfolio Standards on Renewable Investments and Generation: The Role of Heterogeneity and Dynamics. NBER Working Paper 31568.
- Fowlie, M., Greenstone, M., and Wolfram, C. (2018). Do energy efficiency investments deliver? Evidence from the Weatherization Assistance Program. *Quarterly Journal of Economics*, 133(3), 1597–1644.
- Gardner, J. (2022). Two-Stage Difference-in-Differences. Working Paper.
- Gillingham, K., Rapson, D., and Wagner, G. (2016). The rebound effect and energy efficiency policy. *Review of Environmental Economics and Policy*, 10(1), 68–88.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
- Greenstone, M., and Nath, I. (2024). Do Renewable Portfolio Standards deliver cost-effective carbon abatement? University of Chicago Energy Policy Institute Working Paper.
- Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review*, 104(2), 537–563.
- Jacobsen, G.D., and Kotchen, M.J. (2013). Are building codes effective at saving energy? Evidence from residential billing data in Florida. *Review of Economics and Statistics*, 95(1), 34–49.

- Jessoe, K., and Rapson, D. (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review*, 104(4), 1417–1438.
- Joskow, P. (2014). Incentive regulation in theory and practice: Electricity distribution and transmission networks. In *Economic Regulation and Its Reform*, pp. 291–344. University of Chicago Press.
- Levinson, A. (2016). How much energy do building energy codes save? Evidence from California houses. *American Economic Review*, 106(10), 2867–2894.
- MacKinnon, J.G., and Webb, M.D. (2018). The Wild Bootstrap for Few (Treated) Clusters. *The Econometrics Journal*, 21(2), 114–135.
- Metcalf, G.E., and Hassett, K.A. (1999). Measuring the energy savings from home improvement investments: Evidence from monthly billing data. *Review of Economics and Statistics*, 81(3), 516–528.
- Mildenberger, M., Lachapelle, E., Harrison, K., and Stadelmann-Steffen, I. (2022). Limited impacts of carbon tax rebate programmes on public support for carbon pricing. *Nature Climate Change*, 12, 141–147.
- Myers, E. (2019). Are home buyers inattentive? Evidence from capitalization of energy costs. *American Economic Journal: Economic Policy*, 11(2), 165–188.
- Novan, K. (2015). Valuing the Wind: Renewable Energy Policies and Air Pollution Avoided. *American Economic Journal: Economic Policy*, 7(3), 291–326.
- Rambachan, A., and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5), 2555–2591.
- Roth, J. (2022). Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. *American Economic Review: Insights*, 4(3), 305–322.
- Roth, J., Sant’Anna, P.H.C., Bilinski, A., and Poe, J. (2023). What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature. *Journal of Econometrics*, 235(2), 2218–2244.
- Sant’Anna, P.H.C., and Zhao, J. (2020). Doubly Robust Difference-in-Differences Estimators. *Journal of Econometrics*, 219(1), 101–122.
- Sun, L., and Abraham, S. (2021). Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics*, 225(2), 175–199.

U.S. Environmental Protection Agency. (2021). Technical Support Document: Social Cost of Carbon, Methane, and Nitrous Oxide—Interim Estimates under Executive Order 13990. <https://www.epa.gov/environmental-economics/scghg>.

A. Data Appendix

A.1 Data Sources and Access

All data used in this paper are publicly accessible through government APIs and databases.

EIA State Energy Data System (SEDS). Accessed via `api.eia.gov/v2/seds/data/` with the DEMO_KEY. Series ESRCB (residential electricity, Billion Btu), ESTCB (total), ESCCB (commercial), and ESICB (industrial) were downloaded for all states for 1990–2023.

EIA Retail Sales. Accessed via `api.eia.gov/v2/electricity/retail-sales/data/`. Annual residential and commercial sector data including price, revenue, and sales for 1990–2023.

Census Population Estimates. Accessed via `api.census.gov`. Intercensal estimates for 2000–2009, annual estimates for 2010–2019, and vintage 2023 estimates for 2020–2023. For 1990–1999, linear interpolation between the 1990 and 2000 Decennial Census counts.

EERS Treatment Coding. Compiled from the ACEEE State Energy Efficiency Resource Standards database, DSIRE, and NCSL. Treatment defined as the first year of a binding mandatory EERS with quantitative savings targets.

A.2 Variable Definitions

The primary dependent variable is the natural logarithm of per-capita residential electricity consumption, constructed as SEDS series ESRCB (Billion Btu) divided by state population. Residential electricity price is the average retail price in cents per kilowatt-hour from EIA retail sales data. The EERS indicator equals one in all years at or after the state’s EERS adoption year and zero otherwise; it is zero for all years in never-treated states. For the Callaway–Sant’Anna estimator, the first treatment year (G_s) is set to the adoption year for treated states (e.g., $G_s = 2008$ for Massachusetts) and $G_s = 0$ for never-treated states, following the encoding convention of the `did` R package (v2.1), where $G_s = 0$ signals never-treated status. This is functionally equivalent to the $G_s = \infty$ notation used in the theoretical literature.

A.3 Sample Restrictions

The panel consists of 51 jurisdictions \times 34 years = 1,734 potential observations. Missing EIA SEDS data for some states in 1990–1994 reduces the estimation sample to 1,479 state-year observations. All 51 jurisdictions contribute observations.

B. Identification Appendix

B.1 Goodman-Bacon Decomposition

The TWFE estimate can be decomposed using the [Goodman-Bacon \(2021\)](#) method. The decomposition reveals that the dominant weight comes from clean treated-vs-untreated comparisons, with smaller contributions from earlier-vs-later and later-vs-earlier comparisons. The positive coefficient on the later-vs-earlier component reflects the “forbidden comparisons” that contaminate TWFE in staggered settings. The overall TWFE estimate of -0.026 is attenuated relative to the CS estimate of -0.042 partly because of this contamination.

C. Additional Figures



Figure 9: Honest Confidence Intervals at Key Event Times Under Varying Parallel Trends Assumptions

References

Allcott, Hunt, “Social Norms and Energy Conservation,” *Journal of Public Economics*, 2011, 95 (9–10), 1082–1095.

- **and Dmitry Taubinsky**, “Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market,” *American Economic Review*, 2015, *105* (8), 2501–2538.
- Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager**, “Synthetic Difference-in-Differences,” *American Economic Review*, 2021, *111* (12), 4088–4118.
- Barbose, Galen L., Charles A. Goldman, Ian M. Hoffman, and Megan Billingsley**, “The Future of Utility Customer-Funded Energy Efficiency Programs in the United States: Projected Spending and Savings to 2025,” Report LBNL-5803E, Lawrence Berkeley National Laboratory 2013.
- Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram**, “Machine Learning from Schools about Energy Efficiency,” *Journal of the Association of Environmental and Resource Economists*, 2020, *7* (6), 1181–1217.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, *90* (3), 414–427.
- de Chaisemartin, Clément and Xavier D’Haultfoeuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Deschenes, Olivier, Christopher Malloy, and Gavin McDonald**, “Causal Effects of Renewable Portfolio Standards on Renewable Investments and Generation: The Role of Heterogeneity and Dynamics,” Working Paper 31568, NBER 2023.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram**, “Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program,” *Quarterly Journal of Economics*, 2018, *133* (3), 1597–1644.
- Gillingham, Kenneth, David Rapson, and Gernot Wagner**, “The Rebound Effect and Energy Efficiency Policy,” *Review of Environmental Economics and Policy*, 2016, *10* (1), 68–88.

- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Greenstone, Michael and Ishan Nath**, “Do Renewable Portfolio Standards Deliver Cost-Effective Carbon Abatement?,” Working Paper, University of Chicago 2024.
- MacKinnon, James G. and Matthew D. Webb**, “The Wild Bootstrap for Few (Treated) Clusters,” *The Econometrics Journal*, 2018, *21* (2), 114–135.
- Novan, Kevin**, “Valuing the Wind: Renewable Energy Policies and Air Pollution Avoided,” *American Economic Journal: Economic Policy*, 2015, *7* (3), 291–326.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, *90* (5), 2555–2591.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.
- U.S. Environmental Protection Agency**, “Technical Support Document: Social Cost of Carbon, Methane, and Nitrous Oxide – Interim Estimates under Executive Order 13990,” Technical Report, EPA 2021.