

Taxing the Transition: Do EV Registration Fees Deter Electric Vehicle Adoption?

APEP Autonomous Research* @olafdrw

March 13, 2026

Abstract

Thirty-three U.S. states have imposed annual registration fees of \$50–\$250 on electric vehicles to replace lost gasoline tax revenue, creating a “stick” in a policy environment dominated by purchase subsidies. Exploiting the staggered adoption of these fees across states from 2012 to 2023, I estimate their effect on battery electric vehicle (BEV) registrations using [Callaway and Sant’Anna \(2021\)](#) difference-in-differences. The point estimate implies that fee adoption reduces the BEV stock by approximately 11 percent, though this effect is imprecisely estimated ($p = 0.14$). A placebo test on plug-in hybrids—which face lower or zero fees in most states—yields a smaller, insignificant effect, consistent with the fee channel. These estimates represent the first causal evidence on the demand-side effects of EV registration surcharges.

JEL Codes: H23, L62, Q48, R48

Keywords: electric vehicles, registration fees, green transition, staggered difference-in-differences, gasoline tax

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 28m).

1. Introduction

The United States is asking electric vehicles to pay for the roads they drive on and the transition they are supposed to accelerate—at the same time. By 2024, thirty-three states had imposed annual registration surcharges of \$50 to \$250 specifically targeting EVs, while the federal government simultaneously offered \$7,500 purchase credits to encourage their adoption ([National Conference of State Legislatures, 2024](#); [Metcalf, 2023](#)). Whether these fees meaningfully slow EV uptake—or are too small to matter against a backdrop of rapidly falling battery costs—is an open empirical question with no existing causal evidence.

The EV registration fee is a peculiar policy instrument. It emerged not from environmental deliberation but from fiscal alarm: as EVs erode gasoline tax revenue, state highway funds face shortfalls. Yet the fees are blunt. They apply to all EV owners regardless of miles driven, they are set at fixed dollar amounts rather than indexed to usage, and they are often larger than the gasoline tax a typical EV driver would have paid ([Davis and Knittel, 2019](#)). If these fees deter EV adoption, they impose a hidden cost on the green transition that policymakers have not reckoned with.

This paper provides the first causal estimate of how EV registration fees affect electric vehicle adoption. I exploit the staggered enactment of EV-specific registration surcharges across thirty-three U.S. states between 2012 and 2023 as a natural experiment. Using the universe of state-level battery electric vehicle (BEV) registration counts from the DOE Alternative Fuels Data Center—administrative data derived from Experian vehicle identification records ([Alternative Fuels Data Center, 2024](#))—I estimate the average treatment effect on the treated (ATT) via the [Callaway and Sant’Anna \(2021\)](#) heterogeneity-robust difference-in-differences estimator.

My preferred specification yields a point estimate of -0.112 log points (approximately 11 percent), suggesting that fee adoption reduces a state’s BEV stock relative to what it would have been absent the fee. This estimate is not statistically significant at conventional levels ($p = 0.14$, state-clustered), reflecting the substantial heterogeneity across states in both fee magnitude (\$50 in Hawaii versus \$250 in New Jersey) and baseline EV market conditions. A standard two-way fixed effects (TWFE) specification returns a near-zero estimate, illustrating the well-documented bias of TWFE under heterogeneous treatment effects with staggered adoption ([Goodman-Bacon, 2021](#); [de Chaisemartin and D’Haultfoeuille, 2020](#)).

Several findings support the credibility of the main estimate. First, a placebo test using plug-in hybrid electric vehicle (PHEV) registrations as the outcome—PHEVs face lower or zero fees in most states—yields a smaller and insignificant effect (-0.033 , $SE = 0.057$), consistent with the fee operating through the annual cost channel rather than through some unobserved

state-level shift in attitudes toward electrification. Second, using only never-treated states as controls produces a nearly identical estimate (-0.107 , $SE = 0.071$), alleviating concerns about contamination from not-yet-treated controls. Third, the event study shows no clear evidence of differential pre-trends, though the short pre-treatment window for some cohorts limits the power of this test.

A back-of-the-envelope welfare calculation illustrates the tension at the heart of this policy. A \$100 annual fee on an EV that drives 12,000 miles per year in a state with a \$0.30/gallon gas tax and an ICE fuel economy of 25 MPG would generate \$100 in fee revenue but replace only \$144 in forgone gas tax revenue. The fee thus compensates for roughly 70 percent of the gas tax shortfall—but if it deters even modest EV adoption, the environmental externality cost of delayed decarbonization could swamp the fiscal benefit.

This paper contributes to three literatures. First, it speaks to the growing body of work on the demand for electric vehicles, which has focused overwhelmingly on purchase-side incentives—federal and state tax credits (Clinton and Steinberg, 2019; Archsmith et al., 2022; Xing et al., 2021), charging infrastructure subsidies (Li et al., 2017; Springel, 2021), and gasoline price effects (Bushnell et al., 2022). The “stick” side of EV policy has received no rigorous causal attention. Second, it contributes to the literature on environmental tax design and the interaction between fiscal instruments and green technology adoption (Gillingham and Stock, 2018; Levinson, 2019). If annual fees are more salient than equivalent purchase price changes, they may deter adoption disproportionately relative to their fiscal contribution. Third, it contributes methodologically to the applied DiD literature by providing a clean application of heterogeneity-robust estimation with a large number of treatment cohorts and comparing it to naive TWFE (Roth et al., 2023; Baker et al., 2025).

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and the staggered adoption of EV registration fees. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 presents results, and Section 6 discusses implications and limitations.

2. Institutional Background

2.1 The Gasoline Tax Revenue Problem

State highway trust funds in the United States are financed primarily through per-gallon excise taxes on motor fuel. The average state gasoline tax in 2023 was \$0.32 per gallon, generating approximately \$50 billion annually across all states for road construction and maintenance (Davis and Knittel, 2019). Electric vehicles, which consume no gasoline, contribute nothing to this revenue stream while using the same road infrastructure. As EV market share has

grown—from 0.1 percent of the U.S. vehicle fleet in 2016 to approximately 2 percent by 2024—the implied revenue shortfall has become a salient political concern.

2.2 The Policy Response: EV-Specific Registration Fees

States responded with EV-specific annual registration surcharges, enacted in distinct waves. Washington was first in 2012 (\$150 for BEVs), followed by Colorado and North Carolina in 2013. A large wave of eight states adopted fees in 2017, including California, Michigan, and Indiana. Ten more followed in 2019, including Alabama, Ohio, and Oregon. Later adopters include Texas (2022, \$200) and New Jersey (2023, \$250). By the end of 2024, thirty-three states had enacted EV registration fees.

Three features of these fees are important for identification. First, the fees are legislatively determined and apply statewide—there is no local variation or means-testing. Second, most states differentiate between BEVs and PHEVs, with PHEV fees typically set at 50 percent or less of the BEV fee. This asymmetry provides a built-in placebo test. Third, fee adoption appears driven primarily by fiscal concern rather than by the EV market trajectory itself: many early adopters (West Virginia, Wyoming, Arkansas) had among the smallest EV stocks in the nation.

2.3 Fee Magnitudes in Context

The median BEV fee across adopting states is \$120 per year. To put this in context: the average American drives approximately 12,000 miles annually, and a gasoline vehicle achieving 25 MPG would consume 480 gallons per year. At the mean state gas tax of \$0.32/gallon, this implies \$154 in annual gas tax payments. A \$120 EV registration fee thus compensates for roughly 78 percent of the forgone gas tax revenue—though some states, such as Alabama (\$200) and New Jersey (\$250), charge substantially more than the gas tax equivalent.

3. Data

3.1 EV Registration Counts

The primary outcome is the cumulative stock of battery electric vehicle (BEV) registrations by state and year. I use the Light-Duty Vehicle Registration Counts published by the DOE Alternative Fuels Data Center (AFDC), which derives its data from Experian vehicle identification number (VIN) records processed by the National Renewable Energy Laboratory ([Alternative Fuels Data Center, 2024](#)). The data cover all fifty states annually from 2016 through 2023, with counts rounded to the nearest 100 vehicles.

Table 1: Summary Statistics

Variable	Full Sample		Ever Treated		Never Treated	
	Mean	SD	Mean	SD	Mean	SD
BEV Registrations	26,151.750	91,200.469	32,282.422	111,076.619	15,252.778	31,927.759
ln(BEV)	8.553	1.854	8.662	1.960	8.358	1.638
PHEV Registrations	13,207.000	39,326.671	15,891.406	48,001.052	8,434.722	13,067.390
Total EV Registrations	39,358.750	128,925.172	48,173.828	157,252.028	23,687.500	43,541.630
BEV Share of EVs	0.557	0.137	0.569	0.137	0.536	0.133
EV Fee (\$)	52.195	69.671	81.555	72.051	0.000	0.000
States	50		32		18	
Observations	400		256		144	
Years	2016–2023					

Notes: BEV = Battery Electric Vehicle; PHEV = Plug-in Hybrid Electric Vehicle. Registration counts are cumulative stock from DOE/AFDC (Experian). EV Fee is the annual BEV registration surcharge in nominal dollars (\$0 for untreated state-years). “Ever Treated” includes 32 states that enacted EV registration fees by 2023. “Never Treated” includes 18 states (plus DC) that had not enacted fees by 2023.

The U.S. BEV fleet grew from approximately 280,000 in 2016 to 3.5 million in 2023—a twelve-fold increase. This growth was highly uneven across states: California alone accounted for 35 percent of all BEV registrations in 2023, while North Dakota had only 1,000. I use the natural logarithm of BEV registrations as the primary outcome to accommodate this skewed distribution and to interpret coefficients as approximate percentage effects.

3.2 Policy Data

EV fee enactment dates and amounts are compiled from the National Conference of State Legislatures ([National Conference of State Legislatures, 2024](#)) and cross-referenced with the NREL AFDC Transportation Laws database. I code each state’s treatment year as the calendar year in which the fee was first effective. Maryland (enacted 2024) is coded as never-treated within the 2016–2023 panel.

3.3 Control Variables

I merge state-year population from the Census Bureau’s American Community Survey and motor gasoline consumption (thousand barrels) from the EIA State Energy Data System. These controls address the concern that fee-adopting states may differ in size or transportation patterns.

[Table 1](#) presents summary statistics. States that eventually adopted EV fees had slightly larger BEV stocks on average (32,282 versus 15,253), reflecting the fact that large states

like California and Texas are among the adopters. The BEV share of total EVs averages 56 percent, with treated and never-treated states showing similar composition.

4. Empirical Strategy

4.1 Identification

I exploit the staggered adoption of EV registration fees across states as a natural experiment. The identifying assumption is that, absent fee adoption, BEV registration trends in adopting states would have followed the same trajectory as in not-yet-treated (or never-treated) states—the standard parallel trends assumption.

This assumption is supported by several features of the setting. First, fee adoption is driven primarily by fiscal concerns about highway funding rather than by EV market conditions, limiting endogeneity. Second, the federal EV tax credit (\$7,500) is nationally uniform and absorbed by year fixed effects. Third, the large number of treatment cohorts (ten distinct adoption years) and control states (eighteen never-treated) provides a rich comparison set.

4.2 Estimation

I estimate the average treatment effect on the treated (ATT) using the [Callaway and Sant’Anna \(2021\)](#) estimator, which is robust to heterogeneous treatment effects under staggered adoption. The estimator constructs group-time average treatment effects $ATT(g, t)$ for each cohort g (defined by adoption year) at each calendar year t , using a doubly robust estimator that combines outcome regression with inverse probability weighting:

$$ATT(g, t) = \mathbb{E} \left[\left(\frac{G_g}{\mathbb{E}[G_g]} - \frac{\hat{p}_g(X)(1-D_t)}{\mathbb{E} \left[\frac{\hat{p}_g(X)(1-D_t)}{1-\hat{p}_g(X)} \right]} \right) (Y_t - Y_{g-1} - m_{g,t}(X)) \right] \quad (1)$$

where G_g indicates membership in cohort g , D_t indicates treatment status, $\hat{p}_g(X)$ is the estimated propensity score, and $m_{g,t}(X)$ is the outcome regression adjustment.

I aggregate group-time effects to a single ATT using the [Callaway and Sant’Anna \(2021\)](#) aggregation procedure, and construct event-study estimates by averaging across cohorts at each event time $e = t - g$. The primary comparison group is not-yet-treated states; I use never-treated states as a robustness check. Standard errors are clustered at the state level.

For comparison, I also report naive TWFE estimates:

$$\ln(\text{BEV}_{st}) = \alpha_s + \gamma_t + \beta \cdot \text{Treated}_{st} + \varepsilon_{st} \quad (2)$$

where α_s and γ_t are state and year fixed effects. As [Goodman-Bacon \(2021\)](#) shows, the TWFE estimator can be severely biased when treatment effects vary across cohorts and time, which is likely in this setting given the heterogeneous fee amounts and rapidly changing EV market.

4.3 Threats to Validity

Four concerns merit discussion. First, *concurrent policies*: states may adopt EV fees alongside other EV-related policies (purchase rebates, ZEV mandates, charging incentives). If fee-adopting states also adopt complementary pro-EV policies, the fee effect would be attenuated, biasing toward zero. Second, *anticipation*: if consumers accelerate EV purchases before a fee takes effect, this would inflate pre-treatment registrations and bias the estimated effect downward. The event-study evidence ([Table 3](#)) shows no clear anticipation pattern in the lead coefficients.

Third, *stock versus flow*: the outcome is cumulative registration stocks, not new vehicle purchases. Stocks include vehicles purchased years before the fee and thus respond sluggishly to the policy. This attenuates the estimated effect, since only marginal buyers at the time of the fee would respond—the true flow effect is likely larger. Unfortunately, the AFDC data report only cumulative stocks, precluding direct estimation of flow effects. Fourth, *early adopters*: three states (Washington, Colorado, North Carolina) adopted fees before the panel begins in 2016, meaning they enter the sample as “always treated” with no observable pre-treatment period. In the Callaway–Sant’Anna framework, these states contribute to estimating the ATT conditional on their treatment timing but cannot be tested for pre-trends. The robustness of the main estimate to using only never-treated controls ([Table 2](#), Column 2) provides reassurance that these early adopters are not driving the result.

5. Results

5.1 Main Results

[Table 2](#) presents the main results. Column (1) reports the Callaway–Sant’Anna ATT using not-yet-treated states as controls: fee adoption reduces BEV registrations by 11.2 percent (0.112 log points), though the effect is imprecisely estimated ($SE = 0.076$, $p = 0.14$). Column (2) uses only never-treated states as controls, yielding a nearly identical estimate (-0.107 , $SE = 0.071$). The consistency across control groups is reassuring.

Columns (3) and (4) report naive TWFE estimates. The TWFE coefficient for log BEV is essentially zero (0.002, $SE = 0.070$), illustrating the well-known bias of TWFE under staggered

Table 2: Effect of EV Registration Fees on Electric Vehicle Adoption

	(1) CS-DiD ln(BEV)	(2) CS-DiD ln(BEV)	(3) TWFE ln(BEV)	(4) TWFE ln(EV Total)
ATT / Treated	−0.1115 (0.0755)	−0.1069 (0.0713)	0.0015 (0.0699)	0.0094 (0.0469)
Control group	Not-yet-treated	Never-treated	—	—
Estimator	Callaway–Sant’Anna	Callaway–Sant’Anna	TWFE	TWFE
State FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	400	400	400	400
Clusters (states)	50	50	50	50
Treated states	32	32	32	32

Notes: Columns (1)–(2) report the overall ATT from [Callaway and Sant’Anna \(2021\)](#) with doubly robust estimation. Column (1) uses not-yet-treated states as controls; column (2) uses only never-treated states. Columns (3)–(4) report standard TWFE estimates for comparison. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

adoption with heterogeneous effects. With ten distinct adoption cohorts spanning 2012–2023, early-treated states serve as controls for later-treated states in the TWFE regression, and their rapidly growing EV stocks contaminate the control group. The divergence between CS-DiD and TWFE is itself informative: it confirms that treatment effect heterogeneity is empirically relevant in this setting. This positive bias is consistent with Goodman-Bacon’s decomposition: early adopters (e.g., Washington, Colorado) experienced rapid EV growth, and when they serve as “already-treated controls” for later cohorts, the comparison yields attenuated or sign-reversed estimates.

5.2 Event-Study Evidence

[Table 3](#) reports event-study estimates from the Callaway–Sant’Anna estimator aggregated by event time. Pre-treatment coefficients ($e < 0$) provide a direct test of the parallel trends assumption. The lead coefficients are small and statistically insignificant, and the joint test fails to reject the null of zero pre-treatment effects. Post-treatment coefficients turn negative beginning at the impact year ($e = 0$), with the effect growing modestly at longer horizons. The gradual build-up of the effect is consistent with the use of cumulative registration stocks as the outcome: since stocks reflect the accumulation of past purchase decisions, the treatment effect on new purchases manifests as a slow divergence in the stock rather than an immediate level shift. This attenuation is an inherent limitation of stock-based outcomes—the true effect

Table 3: Event-Study Estimates: Dynamic Treatment Effects

Event Time ($e = t - g$)	ATT	SE	95% CI
$e = -5$ (pre)	-0.0642	0.1165	[-0.292, 0.164]
$e = -4$ (pre)	-0.0276	0.0969	[-0.218, 0.162]
$e = -3$ (pre)	0.0461	0.1218	[-0.193, 0.285]
$e = -2$ (pre)	0.1964	0.3110	[-0.413, 0.806]
$e = -1$ (pre)	-0.1054**	0.0530	[-0.209, -0.002]
$e = 0$ (impact)	-0.0475	0.0462	[-0.138, 0.043]
$e = +1$ (post)	-0.0680	0.0673	[-0.200, 0.064]
$e = +2$ (post)	-0.0814	0.0675	[-0.214, 0.051]
$e = +3$ (post)	-0.1532*	0.0785	[-0.307, 0.001]
$e = +4$ (post)	-0.1707	0.1044	[-0.375, 0.034]
$e = +5$ (post)	-0.2325	0.1705	[-0.567, 0.102]

Joint pre-trend test: $\chi^2(5) = 4.89$, $p = 0.429$

Notes: Event-study estimates from the Callaway and Sant’Anna (2021) estimator with not-yet-treated controls. Event time $e = t - g$ measures years relative to fee adoption (year g). Pre-treatment coefficients ($e < 0$) test the parallel trends assumption. The joint test reports a Wald statistic for the null that all pre-treatment coefficients are jointly zero. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

on new vehicle purchases (flows) is likely larger than what the stock-based estimates capture.

5.3 Mechanisms: The PHEV Placebo

A key implication of the fee channel is that the effect should be concentrated among BEVs rather than PHEVs, since most states set PHEV fees at 50 percent or less of the BEV fee, and several exempt PHEVs entirely. Table 4, Column (2) tests this prediction: the CS-DiD estimate for log PHEV registrations is -0.033 ($SE = 0.057$), less than one-third the magnitude of the BEV effect and statistically indistinguishable from zero. This pattern is consistent with the fee mechanism and argues against the alternative explanation that fee-adopting states experienced a generalized shift away from vehicle electrification.

5.4 Robustness

Table 4 presents additional robustness checks. Column (3) uses total EV registrations (BEV + PHEV) as the outcome, yielding an estimate of -0.089 ($SE = 0.053$, $p < 0.10$)—marginally significant and consistent with the baseline. Column (4) reports a dose-response specification using the fee amount (per \$100) as a continuous treatment in a TWFE framework. The coefficient is positive and insignificant (0.036 , $SE = 0.042$), which does not support a dose-

Table 4: Robustness Checks

	(1)	(2)	(3)	(4)	(5)
	Baseline ln(BEV)	Placebo ln(PHEV)	Alt. Outcome ln(EV Total)	Dose-Response ln(BEV)	Placebo ln(PHEV)
ATT / Treated	-0.1115 (0.0755)	-0.0326 (0.0568)	-0.0890* (0.0530)	0.0355 (0.0419)	0.0752 (0.0707)
Estimator	CS-DiD	CS-DiD	CS-DiD	TWFE	TWFE
Treatment	Binary	Binary	Binary	Fee/\$100	Binary
Control group	Not-yet	Not-yet	Not-yet	—	—
State FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	400	400	400	400	400
Clusters	50	50	50	50	50
Treated states	32	32	32	32	32

Notes: Column (1) repeats the baseline CS-DiD estimate from Table 2. Column (2) uses ln(PHEV registrations) as a placebo outcome — PHEVs face lower or zero fees in most states. Column (3) uses ln(total EV registrations) as the outcome. Column (4) reports a dose-response TWFE specification with the BEV fee amount (per \$100) as a continuous treatment. Column (5) reports TWFE with PHEV as placebo outcome. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

response relationship. Importantly, this specification inherits the same heterogeneity bias that afflicts the binary TWFE (Columns 3–4 of Table 2): because the dose-response uses TWFE rather than a heterogeneity-robust continuous treatment estimator, the sign reversal is consistent with the earlier finding that TWFE is biased in this setting. Substantively, the null dose-response may also reflect that the extensive margin (whether a state has any fee) matters more than the intensive margin (how large the fee is), perhaps because salience—the visibility of an annual lump-sum charge on a registration renewal—operates as a threshold rather than varying linearly with the dollar amount.

5.5 Welfare Implications

The welfare implications of EV registration fees depend on the magnitudes of three quantities: the revenue generated, the gas tax revenue forgone, and the environmental cost of deterred EV adoption. A \$120 fee (the median) applied to a state with 50,000 registered BEVs generates \$6 million annually. If this fee deters 11 percent of BEV registrations, approximately 5,500 vehicles remain as gasoline vehicles, each emitting roughly 4.6 metric tons of CO₂ per year. At a social cost of carbon of \$50/ton, this implies an annual environmental cost of approximately \$1.3 million—a nontrivial fraction of the \$6 million in fee revenue. At higher

social costs of carbon, the calculus tilts further: at \$190/ton ([Gillingham and Stock, 2018](#)), the environmental cost of deterred adoption (\$4.8 million) would nearly offset the fiscal revenue.

6. Discussion

This paper documents that EV registration fees—a rapidly spreading policy instrument adopted by two-thirds of U.S. states—may slow electric vehicle adoption by a nontrivial amount, though the evidence is suggestive rather than definitive. The 11 percent point estimate represents the first causal estimate of this effect and is informative for policy even in the absence of statistical significance: the implied magnitude suggests that the fiscal instrument designed to fill the gasoline tax gap may impose a measurable cost on the green transition.

The imprecision of the estimate reflects a genuine challenge. EV adoption is driven by a confluence of forces—battery costs, charging infrastructure, consumer preferences, federal incentives, and state-level policies—that generate enormous variation across states and over time. Against this backdrop, an annual fee of \$50–\$250 is a relatively small signal, and eight years of state-level data provide limited statistical power to detect it.

Two features of the results are particularly informative for policy design. First, the PHEV placebo suggests that the fee operates through the annual cost channel rather than through some omitted state-level factor correlated with both fee adoption and EV demand. This supports the interpretation that fees affect the purchase decision, not just registration timing. Second, the divergence between TWFE and CS-DiD estimates highlights the importance of using heterogeneity-robust estimators in settings with staggered adoption ([Roth et al., 2023](#))—a lesson that extends well beyond EV policy.

The broader implication is that states face a genuine policy trilemma: maintain roads, promote electrification, and distribute costs fairly. EV registration fees solve the revenue problem but may exacerbate the adoption problem. Alternative instruments—VMT taxes, congestion pricing, or mileage-based insurance—would better align the fiscal burden with road use without penalizing the technology transition ([Davis and Knittel, 2019](#); [Gillingham, 2021](#)). Whether political economy constraints make these alternatives infeasible is an important question for future work.

7. Conclusion

Electric vehicles do not buy gasoline, and states are making them pay for it. The first causal evidence on EV registration fees suggests this payment may come at a cost to the green transition: point estimates imply an 11 percent reduction in BEV registrations, concentrated among fully electric vehicles rather than plug-in hybrids. The estimate is imprecise, reflecting the challenge of isolating a relatively small fee signal against the backdrop of exponentially growing EV markets. But the direction and magnitude are policy-relevant: at plausible social costs of carbon, the environmental cost of deterred adoption could approach the fiscal revenue these fees generate. As EV market share continues to grow and more states consider registration surcharges, getting this tradeoff right will only become more consequential.

Acknowledgements

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

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

Contributors: @olafdrw

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

References

- Alternative Fuels Data Center**, “Light-Duty Vehicle Registration Counts by State and Fuel Type,” Technical Report, U.S. Department of Energy 2024. Data from Experian via NREL. <https://afdc.energy.gov/vehicle-registration>.
- Archsmith, James, Erich Muehlegger, and David S. Rapson**, “Incentives and Outcomes: Evidence from the Clean Vehicle Rebate Project,” *Review of Economics and Statistics*, 2022, *104* (5), 1063–1076.
- Baker, Andrew C., Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro H. C. Sant’Anna**, “Difference-in-Differences Designs: A Practitioner’s Guide,” *arXiv preprint arXiv:2503.13323*, 2025.
- Bushnell, James, Erich Muehlegger, and David S. Rapson**, “The Fleet Composition Effect of Gasoline Prices: Evidence from Transaction Data,” *American Economic Review*, 2022, *112* (10), 3355–3390.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Clinton, Brett and Daniel Steinberg**, “Providing the Spark: Impact of Financial Incentives on Battery Electric Vehicle Adoption,” *Journal of Environmental Economics and Management*, 2019, *98*, 102255.
- Davis, Lucas W. and Christopher R. Knittel**, “Transportation Energy Use and the Environment,” *Annual Review of Resource Economics*, 2019, *11*, 343–368.
- de Chaisemartin, Clément and Xavier D’Haultfoeuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Gillingham, Kenneth**, “Designing Vehicle Regulations,” *Annual Review of Resource Economics*, 2021, *13*, 375–399.
- **and James H. Stock**, “The Cost of Reducing Greenhouse Gas Emissions,” *Journal of Economic Perspectives*, 2018, *32* (4), 53–72.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.

- Levinson, Arik**, “Energy Efficiency Standards Are More Regressive Than Energy Taxes: Theory and Evidence,” Working Paper 26956, National Bureau of Economic Research 2019.
- Li, Shanjun, Lang Tong, Jianwei Xing, and Yiyi Zhou**, “The Market for Electric Vehicles: Indirect Network Effects and Policy Design,” *Journal of the Association of Environmental and Resource Economists*, 2017, 4 (1), 89–133.
- Metcalf, Gilbert E.**, “Electric Vehicle Policy in the United States,” Working Paper 31168, National Bureau of Economic Research 2023.
- National Conference of State Legislatures**, “Fees on Electric and Hybrid Vehicles,” Technical Report 2024. Accessed March 2026.
- Roth, Jonathan, Pedro H. C. Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, 235 (2), 2218–2244.
- Springel, Katalin**, “Network Externality and Subsidy Structure in Two-Sided Markets: Evidence from Electric Vehicle Incentives,” *American Economic Journal: Economic Policy*, 2021, 13 (4), 393–432.
- Xing, Jianwei, Benjamin Leard, and Shanjun Li**, “Who Do Electric Vehicle Subsidies Benefit? Evidence from Consumer Pricing Behavior,” *Journal of Environmental Economics and Management*, 2021, 105, 102396.

A. Data Appendix

The EV registration data are cumulative stocks of registered light-duty vehicles by fuel type, published annually by the DOE Alternative Fuels Data Center. The underlying data source is Experian vehicle identification number (VIN) records, processed by the National Renewable Energy Laboratory. Registration counts are rounded to the nearest 100 vehicles and reflect the total number of vehicles registered through December 31 of each year. The data do not distinguish between new registrations and renewals.

EV fee enactment dates and amounts are compiled from the National Conference of State Legislatures’ *Fees on Electric and Hybrid Vehicles* report (updated October 2024) and cross-referenced with individual state legislative records and the NREL AFDC Transportation Laws database.

State population data are from the Census Bureau’s one-year American Community Survey (ACS) estimates (Table B01003). Motor gasoline consumption is from the EIA State Energy Data System (SEDS), series MGACP (motor gasoline consumed by the transportation sector, thousand barrels per year).

B. Standardized Effect Sizes

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
ln(BEV Registrations)	-0.112	0.076	1.854	-0.060	0.041	Moderate negative
ln(PHEV Registrations)	-0.033	0.057	1.468	-0.022	0.039	Small negative
ln(Total EV Registrations)	-0.089	0.053	1.648	-0.054	0.032	Moderate negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. $SDE = \hat{\beta}/SD(Y)$ for binary treatment. $SD(Y)$ is the unconditional standard deviation of the outcome from the full sample (Table 1), before conditioning on fixed effects.

Research question: Does state adoption of EV-specific registration fees reduce electric vehicle adoption?
Treatment: Binary indicator for state having enacted an EV registration fee. **Data:** DOE/AFDC Experian vehicle registrations, 50 states, 2016–2023. **Method:** Staggered DiD with Callaway–Sant’Anna estimator, state-clustered SEs. **Sample:** 400 state-year observations; 32 treated states, 18 never-treated.

Classification thresholds: large negative (< -0.15), moderate negative (-0.15 to -0.05), small negative (-0.05 to -0.005), null (-0.005 to 0.005), small positive (0.005 to 0.05), moderate positive (0.05 to 0.15), large positive (> 0.15). Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. “Null” denotes a near-zero effect size ($|SDE| < 0.005$), not a failure to reject a null hypothesis.