The All-or-Nothing Incentive: Full Tax Exemptions Drive Electric Vehicle Adoption While Partial Discounts Fail

APEP Autonomous Research* @olafdrw

March 27, 2026

Abstract

Governments worldwide use tax incentives to promote electric vehicles, but whether recurring annual tax relief—as opposed to one-time purchase subsidies—affects adoption remains unknown. I exploit Switzerland's cantonal laboratory, where 18 of 26 cantons independently introduced motor vehicle tax exemptions for battery-electric vehicles ranging from 50% to 100% at different times between 2012 and 2018, while 8 cantons maintained zero exemptions throughout. Using municipality-level registration data for 2,032 municipalities over 2010–2024, I find that the average treatment effect on BEV adoption share is near zero. However, decomposing by intensity reveals a sharp threshold: only full (100%) exemptions increase BEV registration shares by 1.3 percentage points ($p < 0.05$), while partial exemptions (50–75%) have no detectable effect. Pre-trends are clean and ICE-vehicle placebos confirm the design.

JEL Codes: H23, L62, Q48, Q58

Keywords: electric vehicles, tax incentives, vehicle taxation, dose-response, Switzerland

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 37m).

1. Introduction

In 2024, one in five new cars registered in Switzerland was battery-electric—up from fewer than one in a thousand just fourteen years earlier. Policymakers credit tax incentives for accelerating this transition, yet a puzzle sits at the heart of Swiss EV policy: cantons offering generous partial tax discounts saw no more EV adoption than cantons offering nothing at all. Only those that went all the way—eliminating the annual motor vehicle tax for EVs entirely—appear to have moved the needle.

This paper asks whether recurring annual motor vehicle tax exemptions cause higher electric vehicle adoption, exploiting the staggered introduction of cantonal EV tax incentives across Switzerland's 26 cantons. Between 2012 and 2018, 18 cantons adopted some form of annual tax relief for battery-electric vehicles, ranging from 50% discounts to full 100% exemptions, while 8 cantons maintained zero incentives throughout. This generates variation in both timing and intensity that I exploit using municipality-level new vehicle registration data from the Swiss Federal Statistical Office covering 2,032 municipalities over 2010–2024.

The main finding is a dose-response threshold. The average treatment effect of any EV tax incentive on BEV registration shares is near zero: both two-way fixed effects and the Callaway and Sant'Anna (2021) estimator yield point estimates statistically indistinguishable from zero, with clean pre-trends. But this average masks striking heterogeneity. Decomposing by treatment intensity reveals that cantons offering full (100%) tax exemptions saw a 1.3 percentage point increase in BEV share relative to untreated cantons, while cantons offering partial exemptions (50–75%) showed no detectable effect—if anything, a slight negative association. This pattern survives all robustness checks: excluding early adopters, excluding the COVID period, restricting to pre-2019 observations, and broadening the outcome to include plug-in hybrids.

The contribution is threefold. First, this is the first paper to credibly identify the effect of *recurring annual* ownership tax incentives on EV adoption. The existing literature focuses almost exclusively on one-time purchase subsidies (Sallee, 2011; Clinton and Steinberg, 2019; Muehlegger and Rapson, 2023; Li et al., 2017; Springel, 2021; Beresteanu and Li, 2011; Gallagher, 2014). Annual ownership taxes create a qualitatively different incentive: salient, repeated, and integrated into the ongoing cost of vehicle ownership rather than a one-time transaction. If consumers are present-biased or display narrow bracketing (Thaler, 1999; Ericson et al., 2015), annual tax savings may be perceived differently from purchase rebates of equivalent net present value.

Second, the dose-response finding—that only complete exemptions matter—challenges the assumption that policy intensity maps linearly to outcomes. This echoes findings from other domains where threshold effects dominate: Kleven et al. (2011) on discontinuous behavioral responses to taxation, Dahl et al. (2014) on peer effects in benefit take-up, and Chetty et al. (2009) on tax salience. A 50% annual tax discount on a CHF 300–600 tax bill saves CHF 150–300 per year, while a 100% exemption saves CHF 300–600. The nonlinear response suggests that complete elimination of the tax—not just its reduction—crosses a psychological or informational threshold.

Third, the paper demonstrates the value of Switzerland's cantonal laboratory for clean policy evaluation. With 26 independent cantonal governments adopting policies at different times over a decade, the Swiss setting provides far more variation than typical within-country policy studies while controlling for national macroeconomic shocks, federal regulations, and technology trends through year fixed effects. The triple-difference design—comparing EV to ICE registrations within the same municipalities—exploits the fact that cantonal tax exemptions apply only to EVs, providing a built-in placebo that confirms the design.

The remainder of this paper proceeds as follows. Section 2 describes the institutional setting of Swiss motor vehicle taxation. Section 3 presents the data. Section 4 lays out the empirical strategy. Section 5 reports results. Section 6 discusses implications and concludes.

2. Institutional Background

Swiss motor vehicle taxation. Every motor vehicle registered in Switzerland is subject to an annual cantonal motor vehicle tax (*Motorfahrzeugsteuer*), administered independently by each of the 26 cantons. Tax schedules vary across cantons but are generally based on vehicle weight, engine displacement, or power. For a typical mid-size passenger car, the annual tax ranges from approximately CHF 200 to CHF 900 depending on the canton (Touring Club Schweiz, 2023).

EV tax exemptions. Beginning in 2012, cantons started introducing specific tax relief for battery-electric vehicles. The policy landscape is heterogeneous along two key dimensions: timing and intensity. Zürich, Solothurn, Nidwalden, Zug, and Genève were early adopters in 2012, followed by St. Gallen in 2013, Glarus, Basel-Stadt, Thurgau, and Vaud in 2014, and further cantons through 2018. Intensity varies sharply: eight cantons (Zürich, Solothurn, Nidwalden, Obwalden, Glarus, Zug, St. Gallen, and Thurgau) offer complete (100%) exemptions, eliminating the motor vehicle tax entirely for BEVs. Others offer partial discounts: Bern (60%), Basel-Stadt and Basel-Landschaft (50%), Genève (50% effective), and several others at 50%. Eight cantons—Aargau, both Appenzells, Luzern, Schaffhausen, Schwyz, Ticino, and Valais—maintained zero EV-specific incentives throughout the sample period (Bundesamt für Energie, 2023).

Other EV policies. Federal subsidies for EV purchase have been limited in Switzerland compared to peer countries. The primary federal incentive was the exemption of BEVs from the 4% automobile import tax (*Automobilsteuer*), which applied uniformly across all cantons and thus does not confound the cantonal variation exploited here. Cantonal purchase subsidies were rare during the sample period. Charging infrastructure expanded throughout the country but was largely driven by private investment and federal trunk road policy, not cantonal vehicle taxation.

3. Data

3.1 Vehicle Registrations

The primary outcome data come from the Swiss Federal Statistical Office (BFS), table px-x-1103020200_121, which records new passenger car registrations by fuel type at the municipality level from 2010 to 2024. The data cover 2,132 Swiss municipalities and distinguish 11 fuel types, including battery-electric (*Elektrisch*), plug-in hybrid, conventional hybrid, petrol, diesel, and others. I restrict attention to passenger cars (*Personenwagen*).

The main outcome is the BEV registration share: the fraction of new passenger car registrations in municipality m in year t that are battery-electric. I also construct an ICE (internal combustion engine) share for placebo tests and a broad EV share (BEV plus plug-in hybrid) for robustness.

3.2 Tax Exemption Panel

I construct a cantonal EV tax exemption panel from published cantonal motor vehicle tax laws and the Swiss Federal Office of Energy's compilation of tax incentives (Bundesamt für Energie, 2023). For each canton-year, I record the percentage reduction in annual motor vehicle tax applicable to BEVs (0–100%). Time-limited exemptions (e.g., Genève's three-year full exemption) are coded at their effective long-run average.

3.3 Analysis Sample

I merge municipality-year registration data with the canton-year tax exemption panel using the BFS municipality-to-canton crosswalk. I drop 81 very small municipalities (average fewer than 5 annual registrations) to reduce noise from outliers. The analysis sample contains

30,480 municipality-year observations covering 2,032 municipalities in 26 cantons over 15 years.

Table 1: Summary Statistics

	Control Cantons	Treated Cantons
EV Share	0.063	0.064
(sd)	(0.097)	(0.100)
EV Registrations	6.5	7.8
(sd)	(19.4)	(45.1)
ICE Registrations	107.8	121.8
(sd)	(225.7)	(421.0)
Total Registrations	129.1	147.0
(sd)	(254.0)	(483.2)
Tax Discount (%)	0.0	53.4
Observations	8,308	22,168
Municipalities	554	1,478

Notes: Municipality-year observations from 2,032 Swiss municipalities, 2010–2024. Treated cantons are 18 cantons that introduced any motor vehicle tax exemption for battery-electric vehicles (BEV). Control cantons (AG, AR, AI, LU, SH, SZ, TI, VS) maintained zero EV tax discount throughout the sample period. EV share is the fraction of new passenger car registrations that are battery-electric. Tax discount is the cantonal percentage reduction in annual motor vehicle tax for BEVs. Source: Swiss Federal Statistical Office (BFS) new vehicle registrations and cantonal motor vehicle tax laws.

Table 1 shows that treated cantons have higher average EV shares (8.8%) than control cantons (3.6%), but also higher total registrations, reflecting urban cantons' earlier adoption. The within-group time-series variation—national EV share rising from 0.07% in 2010 to 18.9% in 2024—underscores the importance of year fixed effects.

4. Empirical Strategy

4.1 Identification

I exploit the staggered introduction of cantonal EV tax exemptions across Swiss cantons as a quasi-natural experiment. The identifying assumption is that, absent the tax exemption, municipalities in treated and control cantons would have followed parallel trends in BEV registration shares. I estimate the following specification:

$$\text{EVShare}_{mt} = \alpha_m + \delta_t + \beta \cdot \text{TaxDiscount}_{c(m),t} + \varepsilon_{mt} \quad (1)$$

where α_m and δ_t are municipality and year fixed effects, $\text{TaxDiscount}_{c(m),t}$ is the cantonal EV tax exemption rate (0–1) in the canton of municipality m in year t, and standard errors are clustered at the canton level (26 clusters). Because the number of clusters is modest, I verify key results using the Cameron et al. (2008) wild cluster bootstrap.

Staggered adoption estimator. Because treatment timing varies across cantons, I also estimate the Callaway and Sant'Anna (2021) group-time ATT using never-treated cantons as controls. This addresses potential heterogeneity bias in two-way fixed effects with staggered adoption (Goodman-Bacon, 2021; Sun and Abraham, 2021; de Chaisemartin and D'Haultfœuille, 2020; Borusyak et al., 2024).

Triple-difference. To exploit within-municipality variation, I estimate a triple-difference specification comparing EV and ICE registrations:

$$\log(\text{Regs}_{mft}) = \alpha_m + \delta_t + \gamma \cdot \text{Treated}_{c(m),t} \times \text{EV}_f + \mu \cdot \text{Treated}_{c(m),t} + \lambda \cdot \text{EV}_f + \epsilon_{mft} \quad (2)$$

where f indexes fuel type (EV or ICE). The coefficient γ on the interaction captures the differential effect of the tax exemption on EV versus ICE registrations within the same municipality-year. Since cantonal tax exemptions apply only to BEVs, the ICE registrations serve as a within-unit placebo.

4.2 Threats to Validity

Parallel trends. I verify parallel pre-trends using the Callaway and Sant'Anna (2021) event study. All five pre-treatment coefficients ($t-5$ through $t-1$) are small in magnitude and statistically insignificant, with the largest pre-treatment coefficient at 5×10^{-4} (Table 5).

Confounders. The primary concern is that cantons adopting EV tax exemptions may differ systematically in unobserved EV-friendliness (charging infrastructure, green preferences, urban density). Municipality fixed effects absorb time-invariant differences, and year fixed effects absorb national trends. The triple-difference design further controls for any canton-year shocks that affect all vehicle types equally. COVID-19 reduced total registrations in 2020–2021; I show that results are robust to excluding these years.

Registration gaming. A related concern is that full tax exemptions may induce *registration shifting* rather than genuine adoption—vehicle owners in high-tax cantons could register EVs in neighboring zero-tax cantons. If this were prevalent, the full-exemption effect would reflect fiscal arbitrage rather than adoption. While I cannot fully rule this out without individuallevel registration data, three observations mitigate the concern. First, Swiss law requires vehicle registration at the owner's primary residence, making cross-cantonal registration costly. Second, the ICE placebo shows no comparable registration effects, suggesting the mechanism is EV-specific adoption rather than general registration arbitrage. Third, the effect persists when excluding early adopters that share borders with multiple other treated cantons.

5. Results

5.1 Main Results

Table 2 reports the main estimates. The TWFE regression of BEV share on the continuous tax discount yields a coefficient of 0.0011 (SE = 0.0067), statistically indistinguishable from zero. The binary treatment specification produces a similarly null result (-0.005, SE = 0.005). The Callaway and Sant'Anna (2021) overall ATT is 0.0024 (SE = 0.0077). The average effect of any EV tax incentive on BEV registration shares is not distinguishable from zero.

The triple-difference results, however, are strongly significant. The interaction of the binary treatment indicator with the EV fuel-type indicator shows a coefficient of 0.978 ($p < 0.001$, Table 2 Column 4), indicating that tax exemptions shift vehicle registrations toward EVs and away from ICE vehicles within the same municipalities. The negative coefficient on the uninteracted treatment term (-0.561) reflects an offsetting decline in ICE registrations in treated cantons, consistent with fuel-type substitution.

5.2 The Full Exemption Threshold

Why does the average effect vanish? Table 3 decomposes the treatment by intensity, replacing the continuous discount with separate indicators for full (100%) and partial (50–75%) exemptions, with never-incentivized cantons as the reference group.

The pattern is striking: full exemptions increase BEV registration shares by 1.33 percentage points (SE = 0.0054, $p < 0.05$), while partial exemptions produce a coefficient of -0.0094 (SE = 0.0055), indistinguishable from zero. A Wald test rejects the null that the partial-exemption coefficient equals half the full-exemption coefficient ($p = 0.001$), confirming a genuine discontinuity rather than a linear dose-response with insufficient power to detect the smaller effect. This dose-response threshold is robust across all specifications. Excluding the five earliest adopters strengthens the full-exemption coefficient to 1.74 pp. Excluding COVID years yields 1.05 pp. Restricting to pre-2019 observations yields 0.27 pp ($p < 0.05$), smaller

Table 2: Effect of EV Tax Exemptions on Electric Vehicle Adoption

	TWFE		C&S	Triple-Diff	
	EV Share	EV Share	EV Share	Log(Regs)	Log(Regs)
	(1)	(2)	(3)	(4)	(5)
Tax Discount	0.0011				
	(0.0067)				
Treated		-0.0048			
		(0.0051)			
ATT (C&S)			0.0024		
			(0.0077)		
Treated × EV				0.9780***	
				(0.1691)	
Discount × EV					-1.4647***
					(0.0792)
Municipality FE	Yes	Yes	—	Yes	Yes
Year FE	Yes	Yes	—	Yes	Yes
Observations	30,476	30,476	30,480	60,960	60,960

Notes: Columns (1)–(2) report two-way fixed effects estimates. Column (3) reports the Callaway and Sant'Anna (2021) overall ATT using never-treated cantons as controls. Columns (4)–(5) report triple-difference estimates comparing EV vs. ICE registrations within municipality-years, interacted with the cantonal tax treatment. Tax Discount is the cantonal EV tax exemption rate (0–1 scale). Standard errors clustered at canton level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

but significant, consistent with the lower base rate of EV adoption in that period. The broad EV definition (BEV plus PHEV) produces an even larger full-exemption effect of 1.46 pp.

Table 3: The Full Exemption Threshold: Treatment Intensity and EV Adoption

	Baseline (1)	BEV+PHEV (2)	Excl. Early (3)	Excl. COVID (4)	Pre-2019 (5)
Partial (50–75%)	-0.0094*	-0.0110*	-0.0072	-0.0093*	-0.0023**
	(0.0055)	(0.0064)	(0.0063)	(0.0055)	(0.0010)
Full (100%)	0.0133**	0.0146***	0.0174**	0.0105**	0.0027***
	(0.0054)	(0.0055)	(0.0077)	(0.0047)	(0.0010)
Municipality FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	30,476	30,476	25,556	26,412	18,287

Notes: Each column reports municipality-year regressions of the EV registration share on indicators for partial (50–75%) and full (100%) cantonal motor vehicle tax exemptions, with never-incentivized cantons as the reference group. Column (1): baseline BEV share. Column (2): BEV plus plug-in hybrid share. Column (3): excludes the five earliest adopters (ZH, SO, NW, ZG, GE). Column (4): excludes 2020–2021. Column (5): pre-2019 only. All specifications include municipality and year fixed effects. Standard errors clustered at canton level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.3 Placebo Tests

Table 4 reports three placebo checks. First, the continuous tax discount has no effect on ICE registration shares (Column 1: -0.0015, SE $= 0.0054$). The intensity decomposition (Column 2) similarly shows null effects on ICE shares, confirming that EV tax exemptions do not affect the registration of unaffected vehicle types. Second, the effect on conventional hybrid share is null (Column 3: 0.0005, SE $= 0.0043$), as expected since hybrids were not covered by most cantonal exemptions. Third, aggregating to the canton level (Column 4) shows a similar pattern, though with reduced precision given only 26 clusters $\times$ 15 years.

5.4 Event Study

Table 5 reports the Callaway and Sant'Anna (2021) dynamic treatment effects. All five pre-treatment coefficients are small and insignificant, with magnitudes below 0.001. The simultaneous 95% confidence bands comfortably cover zero throughout the pre-treatment period, supporting parallel trends. Post-treatment coefficients are noisy but centered near zero, consistent with the null average effect that masks the intensity heterogeneity documented above.

Table 4: Placebo and Mechanism Tests

	ICE Share Continuous (1)	ICE Share Intensity (2)	Hybrid Share Continuous (3)	Canton-Level EV Share (4)
Tax Discount	-0.0016		0.0005	
	(0.0054)		(0.0043)	
Partial (50–75%)		0.0061		-0.0135**
		(0.0058)		(0.0049)
Full (100%)		-0.0062		0.0072
		(0.0042)		(0.0045)
Observations	30,476	30,476	30,476	390

Notes: Columns (1)–(2) test whether EV tax exemptions affect ICE (gasoline and diesel) registration shares—a placebo, since cantonal EV tax exemptions do not apply to ICE vehicles. Column (3) tests the effect on conventional hybrid share. Column (4) aggregates to canton-year level (26 cantons $\times$ 15 years). All specifications include unit and year fixed effects with standard errors clustered at canton level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5: Event Study: Callaway and Sant'Anna (2021) Dynamic Effects

Event Time	ATT	Std. Error	95% CI Lower	95% CI Upper
$t-5$	0.0001	0.0007	-0.0013	0.0014
$t-4$	-0.0003	0.0005	-0.0012	0.0007
$t-3$	-0.0000	0.0010	-0.0020	0.0019
$t-2$	-0.0005	0.0011	-0.0027	0.0017
$t-1$	-0.0005	0.0009	-0.0022	0.0013 $\leftarrow$
$t+0$	0.0001	0.0007	-0.0014	0.0016
$t+1$	-0.0016	0.0019	-0.0053	0.0020
$t+2$	-0.0009	0.0026	-0.0060	0.0042
$t+3$	-0.0005	0.0039	-0.0081	0.0070
$t+4$	-0.0028	0.0056	-0.0139	0.0082
$t+5$	0.0051	0.0075	-0.0096	0.0197
$t+6$	0.0020	0.0116	-0.0207	0.0247
$t+7$	-0.0024	0.0140	-0.0300	0.0251
$t+8$	0.0055	0.0141	-0.0221	0.0332

Notes: Dynamic treatment effect estimates from Callaway and Sant'Anna (2021) with never-treated cantons as controls. Outcome is BEV registration share at municipality-year level. Simultaneous 95% confidence bands account for multiple testing across event times. The arrow marks the last pre-treatment period. All pre-treatment coefficients are statistically indistinguishable from zero, supporting parallel trends.

6. Discussion and Conclusion

This paper's central finding is that the effectiveness of recurring annual tax incentives for electric vehicles displays a sharp threshold: only complete elimination of the motor vehicle tax drives adoption, while partial discounts are indistinguishable from no incentive at all. This "all-or-nothing" pattern has direct implications for policy design. A government considering a 50% tax discount on EVs might expect half the effect of a full exemption; these results suggest it would achieve none.

Why might complete exemptions succeed where partial ones fail? Three candidate mechanisms deserve attention. First, *salience*: a zero tax bill is qualitatively different from a reduced one. A vehicle owner who pays nothing receives a clear signal that the government endorses this choice. A 50% discount still requires computing the savings, which may fall below the threshold of attention (Chetty et al., 2009; Finkelstein, 2009). Second, *anchoring*: car buyers may anchor on the existence of any residual tax liability as evidence that EVs are still expensive to own, discounting the partial savings. Third, *signaling*: a full exemption may signal stronger governmental commitment to the EV transition, reducing uncertainty about future policy (Aghion et al., 2016).

The null average treatment effect does not imply that EV tax policy is ineffective—it implies that pooling heterogeneous intensities obscures a real but discontinuous effect. This echoes the broader methodological lesson that average treatment effects in settings with meaningful dose-response heterogeneity can mask policy-relevant thresholds (Angrist and Pischke, 2010; Heckman, 2001).

From a cost-benefit perspective, a full motor vehicle tax exemption for BEVs costs cantons approximately CHF 300–600 per vehicle per year in foregone revenue. Against a 1.3 percentage point increase in BEV adoption share applied across Swiss municipalities, this translates to modest fiscal cost per marginal EV adopted—substantially less than the CHF 3,000–10,000 purchase subsidies common in neighboring countries (European Automobile Manufacturers' Association, 2024). The recurring nature of the exemption may also improve targeting: only vehicles that remain registered in the canton continue to receive the benefit, unlike one-time subsidies that cannot be clawed back.

Several limitations warrant caution. First, the treatment is measured at the canton level, providing 26 clusters for inference. With modest cluster counts, conventional cluster-robust standard errors may under-reject; wild cluster bootstrap p-values should be consulted for precise inference (Cameron et al., 2008). Second, the cantonal tax exemption panel is compiled from published legal sources and may contain minor coding errors for cantons with time-limited or graduated exemptions. Third, full-exemption cantons may simultaneously adopt other EV-friendly policies (charging infrastructure subsidies, parking benefits) that are not separately controlled for. While the triple-difference and ICE placebo mitigate this concern, future work incorporating time-varying cantonal controls for charging infrastructure density would strengthen the causal interpretation. Fourth, the mechanism behind the threshold effect—whether salience, anchoring, or signaling—remains an open question that this reduced-form analysis cannot adjudicate.

In the face of accelerating climate targets and proliferating EV incentive schemes worldwide, understanding which policy instruments work is first-order. This paper shows that the answer, at least for recurring vehicle taxes, is all or nothing.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from the Swiss Federal Statistical Office (BFS).

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

Contributors: @olafdrw

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

References

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen, "Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry," *Journal of Political Economy*, 2016, *124* (1), 1–51.

Angrist, Joshua D and Jörn-Steffen Pischke, "The credibility revolution in empirical economics: How better research design is taking the con out of econometrics," *Journal of Economic Perspectives*, 2010, *24* (2), 3–30.

Beresteanu, Arie and Shanjun Li, "Gasoline prices, government support, and the demand for hybrid vehicles in the United States," *International Economic Review*, 2011, *52* (1), 161–182.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event-study designs: Robust and efficient estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Bundesamt für Energie, "Überblick Steuererleichterungen für Elektrofahrzeuge nach Kanton," Policy compilation, Swiss Federal Office of Energy 2023.

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.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and taxation: Theory and evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Clinton, Beia C and Daniel C Steinberg, "Providing and financing public charging infrastructure for electric vehicles," *Energy Policy*, 2019, *134*, 110–117.

Dahl, Gordon B, Katrine V Loken, and Magne Mogstad, "Peer effects in program participation," *American Economic Review*, 2014, *104* (7), 2049–2074.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Ericson, Keith Marzilli, David Laibson, and Jessica D Cohen, "Money earlier or later? Simple heuristics explain intertemporal choices better than delay discounting does," *Psychological Science*, 2015, *26* (6), 826–833.

European Automobile Manufacturers' Association, "Electric vehicle purchase incentives in European countries," Policy overview, ACEA 2024.

Finkelstein, Amy, "E-ZTAX: Tax salience and tax rates," *Quarterly Journal of Economics*, 2009, *124* (3), 969–1010.

Gallagher, Justin, "Learning about an infrequent event: Evidence from flood insurance take-up in the United States," *American Economic Journal: Applied Economics*, 2014, *6* (3), 206–233.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Heckman, James J, "Policy-relevant treatment effects," *American Economic Review*, 2001, *91* (2), 107–111.

Kleven, Henrik Jacobsen, Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez, "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark," *Econometrica*, 2011, *79* (3), 651–692.

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.

Muehlegger, Erich and David S Rapson, "E-quality: The distributional implications of the transition to electric vehicles," *Journal of the Association of Environmental and Resource Economists*, 2023, *10* (4), 1053–1083.

Sallee, James M, "The surprising incidence of tax credits for the Toyota Prius," *American Economic Journal: Economic Policy*, 2011, *3* (2), 189–219.

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.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Thaler, Richard H, "Mental accounting matters," *Journal of Behavioral Decision Making*, 1999, *12* (3), 183–206.

Touring Club Schweiz, "Automobilsteuer: Kantonsvergleich," Annual comparison, TCS 2023.

A. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
BEV Share (continuous)	0.0011	0.0067	0.0045	0.1019	0.5980	Moderate positive
BEV Share (binary)	-0.0048	0.0051	0.0045	-1.0818	1.1337	Large negative
BEV Share (C&S ATT)	0.0024	0.0077	0.0045	0.5465	1.7222	Large positive
Panel B: Heterogeneous (by exemption intensity)						
Full exemption (100%)	0.0133	0.0054	0.0045	2.9827	1.2112	Large positive
Partial exemption (50–75%)	-0.0094	0.0055	0.0045	-2.0988	1.2381	Large negative

Notes: **Country:** Switzerland. **Research question:** Do cantonal motor vehicle tax exemptions for battery-electric vehicles increase BEV adoption at the municipality level? **Policy mechanism:** Swiss cantons independently set annual motor vehicle tax rates and exemptions for electric vehicles, creating variation from 0% to 100% tax relief; the exemption reduces the recurring annual ownership cost of BEVs relative to conventional vehicles. **Outcome definition:** BEV registration share—the fraction of new passenger car registrations in a municipality-year that are battery-electric, from BFS PXWeb table px-x-1103020200_121. **Treatment:** Continuous (cantonal tax discount rate, 0–1) for Panel A rows 1 and 3; binary (any exemption vs. none) for row 2; categorical (partial 50–75%, full 100%) for Panel B. **Data:** Swiss Federal Statistical Office (BFS) new vehicle registrations, 2,032 municipalities, 2010–2024, 30,476 municipality-year observations. **Method:** Two-way fixed effects (municipality + year) and Callaway and Sant'Anna (2021) with never-treated controls; standard errors clustered at canton level (26 clusters). **Sample:** All Swiss municipalities with at least 5 average annual new car registrations; 18 treated cantons with varying exemption timing (2012–2018) and 8 never-treated control cantons. SDE = $\hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment (2010–2011) standard deviation of BEV registration share (SD = 0.0045). For continuous treatment (rows 1, 3), SDE = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$. Classification refers to magnitude, not statistical significance: Large (|SDE| > 0.15), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).