

# The Thermostatic Voter: Why Local Policy Success Fails to Build National Support\*

APEP Autonomous Research<sup>†</sup>

@anonymous

@ai1scl

February 10, 2026

## Abstract

Does experience with local policy implementation affect citizens' preferences for national policy? I exploit variation in the timing of cantonal energy law adoption in Switzerland to examine whether exposure to sub-national climate policy shifted voting behavior on a federal referendum. Using municipality-level (Gemeinde) data from the May 2017 Energy Strategy 2050 referendum, I find no robust evidence that prior cantonal energy law exposure increased federal policy support—inconsistent with the policy feedback hypothesis. A spatial regression discontinuity design restricted to same-language (German–German) canton borders—the cleanest specification, free of Röstigraben confounding—yields an estimate of  $-1.6$  pp (SE = 1.18,  $p = 0.17$ ). The pooled border estimate is  $-1.2$  pp (SE = 1.10,  $p = 0.29$ ); OLS with language controls gives  $-1.8$  pp (SE = 1.93,  $p = 0.35$ ; wild cluster bootstrap  $p = 0.42$ ). Stratified randomization inference, permuting treatment within German-speaking cantons only, yields  $p = 0.53$ . Panel difference-in-differences with time-varying treatment coding across four energy referendums (2000–2017) yields a larger estimate of  $-5.2$  pp (SE = 1.55,  $p = 0.002$ ), with parallel pre-trends. Placebo referendums on non-energy issues show positive discontinuities at the same borders, suggesting that the null energy result is not an

---

\*Revision of apep\_0069. See [https://github.com/SocialCatalystLab/ape-papers/tree/main/apep\\_0069](https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0069) for the original. This version corrects three critical issues: (1) the spatial RDD now computes per-segment border distances with same-language verification on both sides; (2) the panel DiD uses time-varying treatment coding ( $D_{ct} = 1$  iff law in force at referendum  $t$ ) rather than a static treated $\times$ post indicator; (3) the Callaway–Sant’Anna estimator uses in-force years for cohort timing and excludes Basel-Stadt (treatment coincides with final period). Wild cluster bootstrap  $p$ -values supplement CRVE throughout, and stratified randomization inference permutes within German-speaking cantons only.

<sup>†</sup>Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

artifact of design. These findings provide no support for the policy feedback hypothesis and are consistent with policy satiation (“thermostatic” preferences), cost salience from implementation experience, or resistance to federal overreach.

**JEL Codes:** D72, H77, Q58

**Keywords:** federalism, policy feedback, referendum voting, energy policy, Switzerland, spatial RDD, randomization inference

## 1. Introduction

How do citizens form preferences about national policies? A growing literature emphasizes that policy preferences are not fixed but respond to lived experience with government programs (Mettler, 2002; Campbell, 2012; Soss, 1999). When citizens experience a policy’s effects firsthand—whether through direct benefits, changed circumstances, or observation of outcomes—they update their beliefs about both the policy’s desirability and the government’s competence to implement it. This “policy feedback” mechanism suggests that successful implementation of local or regional policies could build support for similar national initiatives (Pierson, 1993; Mettler & SoRelle, 2014).

This paper tests whether sub-national policy experimentation generates political support for federal reform in Switzerland’s unique system of “laboratory federalism” (Oates, 1999). Switzerland’s 26 cantons possess substantial legislative autonomy, and between 2011 and 2017, five cantons’ comprehensive energy legislation came into force implementing the Model Cantonal Energy Provisions (MuKE): binding building efficiency standards for new construction and renovations, renewable energy subsidies, and heat pump mandates. On May 21, 2017, Swiss voters faced a federal referendum on the Energy Strategy 2050 (Energiegesetz), which proposed to harmonize similar measures nationally. This setting provides a natural experiment: did experience with MuKE implementation shape how citizens in those cantons voted on the federal reform?

Surprisingly, I find no evidence that it did. Using Gemeinde-level data ( $N = 2,120$  municipalities), I find that municipalities in treated cantons voted 9.6 percentage points *lower* than those in control cantons in raw comparisons. However, this raw gap is driven entirely by language region: French-speaking cantons (all in the control group) showed much higher support (mean 66%) than German-speaking cantons (mean 50%), regardless of treatment status. After controlling for language, the treatment coefficient falls to  $-1.8$  pp ( $SE = 1.93$ ) and is statistically indistinguishable from zero ( $p = 0.35$ ; wild cluster bootstrap  $p = 0.42$ ).

I employ multiple identification strategies arranged in a clear hierarchy. The *primary* specification is a spatial regression discontinuity design restricted to same-language (German–German) canton borders (Keele & Titiunik, 2015; Dell, 2010), where language does not change at the cutoff. This yields an estimate of  $-1.6$  pp ( $SE = 1.18$ ,  $p = 0.17$ )—negative but imprecise, ruling out positive effects larger than 0.7 pp. The pooled border estimate, which includes Röstigraben crossings, is  $-1.2$  pp ( $SE = 1.10$ ,  $p = 0.29$ ); this serves as an upper bound on the treatment effect but is confounded by language. Stratified randomization inference—permuting treatment within German-speaking cantons only, which conditions on the fact that all treated cantons are German-speaking—yields a two-tailed  $p$ -value of 0.53

(Young, 2019). Panel difference-in-differences with time-varying treatment coding across four energy-related referendums (2000, 2003, 2016, 2017) shows parallel pre-trends and a significant negative effect of  $-5.2$  pp (SE = 1.55,  $p = 0.002$ ), providing the strongest evidence against positive feedback. The Callaway–Sant’Anna heterogeneity-robust estimator yields a consistent ATT of  $-5.0$  pp (SE = 3.34) (Callaway & Sant’Anna, 2021). Placebo referendums on non-energy issues (immigration 2014, healthcare 2014, service public 2016) provide a crucial design check: the immigration placebo is null ( $p = 0.64$ ), while healthcare and service public show *positive* discontinuities at the same borders ( $p = 0.007$  and  $p = 0.05$ ), suggesting that treated-side municipalities generally favor federal initiatives—making the null energy result all the more notable.

These findings are inconsistent with the policy feedback hypothesis and instead support alternative mechanisms: voters in cantons with existing cantonal policy may have viewed the federal proposal as redundant, or cantonal implementation may have made salient the costs (rather than benefits) of energy transitions. The pattern aligns with the “thermostatic” model of public opinion (Wlezien, 1995; Soroka & Wlezien, 2010), which predicts that policy implementation *reduces* demand for further action as citizens perceive that the problem has been addressed.

## 1.1 Contributions

This paper contributes to several literatures while offering methodological innovations for settings with few treated clusters. The primary contribution extends the policy feedback literature (Pierson, 1993; Mettler & SoRelle, 2014; Campbell, 2012) to referendum settings and environmental policy, demonstrating that feedback effects need not be positive—and may be absent entirely. While canonical studies document how social programs like the G.I. Bill (Mettler, 2002) and Social Security (Campbell, 2003) generate supportive constituencies, my findings reveal a different dynamic for regulatory policies where costs are visible and benefits diffuse—consistent with Wlezien (1995)’s “thermostatic” model of public opinion, which predicts that policy implementation reduces rather than increases demand for further action.

The paper also advances research on “laboratory federalism” (Oates, 1999; Rose, 1993; Shipan & Volden, 2008) by showing that decentralized experimentation does not automatically build support for federal harmonization. The Swiss case is particularly instructive: cantonal policies were substantively identical to the proposed federal law, yet voters in treated cantons showed no additional enthusiasm for national adoption. This challenges the optimistic assumption that successful state-level policy creates momentum for federal action (Karch, 2007).

Methodologically, I address the challenge of causal inference with few treated units (5 of 26 cantons) by combining spatial RDD at geographic discontinuities (Keele & Titiunik, 2015; Cattaneo et al., 2020), wild cluster bootstrap  $p$ -values for few-cluster settings (Cameron et al., 2008; MacKinnon & Webb, 2017), stratified randomization inference that permutes treatment within the German-speaking stratum (Young, 2019; MacKinnon et al., 2019), and panel analysis with heterogeneity-robust estimation (Callaway & Sant’Anna, 2021). This multi-method approach should prove useful for other settings where cluster-level treatment and limited units render standard asymptotic inference unreliable (Cameron & Miller, 2015). Finally, the findings inform debates about climate policy strategy (Carattini et al., 2018; Kallbekken & Sælen, 2011), cautioning advocates against assuming that sub-national success automatically translates into federal support.

## 1.2 Roadmap

The remainder of the paper proceeds as follows. Section 2 reviews the theoretical framework and related literature. Section 3 describes Switzerland’s energy policy landscape and the institutional setting. Section 4 presents the data, spatial structure, and descriptive statistics. Section 5 outlines the empirical strategy, including OLS with language controls, spatial RDD at canton borders, randomization inference, and panel analysis. Section 6 presents results across all specifications. Section 7 discusses mechanisms, limitations, and policy implications. Section 8 concludes.

## 2. Theoretical Framework and Related Literature

### 2.1 Policy Feedback Theory

The central theoretical framework motivating this paper is “policy feedback”—the idea that policies, once enacted, reshape the political landscape in ways that affect future policy development (Pierson, 1993). Policies can create constituencies, build administrative capacity, shift public opinion, and alter the incentives of political actors. Mettler & SoRelle (2014) distinguish between interpretive effects (how policies signal government intentions and competence) and resource effects (how policies distribute material benefits that mobilize or demobilize groups).

Empirical support for positive feedback is substantial. Mettler (2002) shows that World War II veterans who received G.I. Bill benefits became more civically engaged, not less. Campbell (2003) documents how Social Security recipients became active defenders of the program. Soss (1999) finds that clients of means-tested programs learn different lessons about

government depending on how they are treated by program administrators.

Yet the conditions under which feedback is positive versus negative remain underspecified. [Mettler \(2011\)](#) argues that many policies—particularly those delivered through the tax code or private markets—remain “submerged” and fail to generate the attribution necessary for feedback. Regulatory policies may be especially prone to negative feedback because costs are often more salient than diffuse benefits. Property owners who must retrofit buildings know exactly what they paid; the public health benefits of reduced emissions are invisible.

## 2.2 Laboratory Federalism

A parallel literature examines how federal systems enable policy experimentation. [Oates \(1999\)](#) articulates the classic argument: decentralization allows jurisdictions to serve as “laboratories of democracy,” testing different approaches and generating information about what works. [Rose \(1993\)](#) develops a framework for “lesson-drawing” across jurisdictions—but notes that lessons may be positive (adopt what works) or negative (avoid what fails).

[Shipan & Volden \(2008\)](#) identify four mechanisms by which state-level policies diffuse: learning (observing outcomes), competition (racing to match neighbors), imitation (copying prestigious leaders), and coercion (mandates from higher levels). The first mechanism—learning—is most relevant here. If cantonal energy laws proved successful, voters might “learn” that federal adoption would be beneficial.

However, the diffusion literature has paid less attention to how sub-national experience shapes preferences for *federal* action. Most studies examine horizontal diffusion (state-to-state) rather than vertical (state-to-federal) ([Karch, 2007](#)). In a federal referendum, the relevant question is not whether other cantons should adopt similar policies, but whether *all* cantons should be bound by a national standard. Voters in cantons that have already acted may perceive federal harmonization differently than those in cantons that have not.

## 2.3 Climate Policy Acceptance

A growing literature examines public acceptance of climate policy, with attention to how policy design affects support ([Drews & van den Bergh, 2016](#)). [Kallbekken & Sælen \(2011\)](#) show that earmarking carbon tax revenues for environmental purposes increases acceptance. [Carattini et al. \(2018\)](#) find that distributional concerns—who bears the costs—are central to climate policy opposition.

Less studied is how *prior experience* with climate policy shapes subsequent preferences. One might expect that successful implementation reduces uncertainty and builds support ([Stoutenborough et al., 2014](#)). Alternatively, implementation may reveal hidden costs, generate

losers who mobilize against expansion, or lead to “policy satiation”—a sense that enough has been done.

The Swiss case offers a unique test. Cantonal energy laws were substantively similar to the proposed federal policy. If experience breeds support, treated cantons should vote more favorably. If experience breeds skepticism or satiation, the effect could be zero or negative.

## 2.4 Inference with Few Clusters

A methodological challenge pervades this study: with only 26 cantons and 5 treated, standard asymptotic inference is unreliable (Cameron et al., 2008; Cameron & Miller, 2015). Cluster-robust standard errors require the number of clusters to approach infinity; with few clusters, test statistics are over-rejected and confidence intervals too narrow.

Cameron et al. (2008) propose the wild cluster bootstrap, which performs better than analytical corrections in Monte Carlo simulations but still requires 10–20 clusters to achieve nominal coverage. MacKinnon & Webb (2017) extend this work with six-point weight distributions that improve finite-sample properties. Young (2019) advocates randomization inference, which is exact under the sharp null regardless of the number of clusters.

Spatial regression discontinuity offers a complementary approach (Keele & Titiunik, 2015). By focusing on units near geographic boundaries, spatial RDD increases effective sample size while leveraging quasi-random assignment at borders. Dell (2010) uses geographic discontinuities in colonial institutions; Black (1999) exploits school attendance boundaries. I apply similar logic to canton borders, though I must account for borders that coincide with the Röstigraben language divide.

## 3. Institutional Background

### 3.1 Swiss Federalism and Direct Democracy

Switzerland is one of the world’s most decentralized federal states and has the most extensive system of direct democracy at the national level (Linder & Vatter, 2010; Kriesi, 2005). The 26 cantons possess broad legislative and fiscal autonomy, setting their own tax rates, education curricula, and regulatory standards across many policy domains (Vatter, 2018). Municipalities (Gemeinden) number approximately 2,100 and exercise substantial local autonomy, particularly in land use planning and service delivery.

Direct democracy amplifies the importance of public preferences. Swiss citizens vote on national referendums 3–4 times per year, deciding on constitutional amendments (mandatory referendum), legislation challenged by petition (optional referendum), and citizen-initiated

constitutional changes (popular initiative) ([Trechsel & Sciarini, 2000](#)). The May 2017 Energy Strategy 2050 vote was an optional referendum: parliament had passed the legislation, but opponents collected sufficient signatures to force a popular vote.

### **3.2 The Model Cantonal Energy Provisions (MuKEEn)**

Energy policy in Switzerland has long been a shared competence between confederation and cantons ([Sager, 2014](#)). The federal government sets broad targets and provides incentive programs, while cantons regulate building energy performance through their construction codes. In 2008, the Conference of Cantonal Energy Directors (EnDK) developed the Model Cantonal Energy Provisions (MuKEEn 2008), a harmonized framework that cantons could voluntarily adopt.

MuKEEn 2008 and its successor MuKEEn 2014 established building envelope standards for new construction and major renovations, required minimum renewable energy contributions in new buildings, restricted electric resistance heating, mandated energy certificates (GEAK) upon sale or renovation, and provided subsidies for solar photovoltaic systems, heat pumps, and building insulation.

Adoption timing varied substantially across cantons. Between 2010 and 2016, five cantons enacted comprehensive energy laws that fully implemented MuKEEn provisions. Graubünden was the first, adopting its *Energiegesetz* in 2010 (in force January 2011) with comprehensive building standards and strong enforcement. Bern followed in 2011 (in force January 2012) with one of the most ambitious cantonal energy laws in the country. Aargau enacted its *Energiegesetz* in 2012 (in force January 2013) with strong efficiency mandates for new construction. Later, Basel-Landschaft adopted its law in 2015 (in force July 2016), implementing MuKEEn 2014 standards. Finally, Basel-Stadt—the most urban canton with high pre-existing environmental support—adopted its *Energiegesetz* in 2016 (in force January 2017).

Other cantons adopted energy legislation after the May 2017 vote: Lucerne (LU) in late 2017, Fribourg (FR) in 2019, and Appenzell Innerrhoden (AI) in 2020. Some cantons (e.g., Zürich, St. Gallen) had partial MuKEEn implementation but not comprehensive standalone energy laws by May 2017.

### **3.3 The Energy Strategy 2050 Referendum**

Following the Fukushima nuclear disaster in March 2011, the Swiss Federal Council announced a gradual phase-out of nuclear power, which then provided approximately 40% of Swiss electricity ([Rinscheid, 2015](#)). The Energy Strategy 2050, developed over several years, was

passed by parliament in September 2016.

The Energy Strategy 2050 contained five major provisions: a prohibition on new nuclear power plant construction; binding targets to reduce per-capita energy consumption by 43% and electricity consumption by 13% by 2035 (relative to 2000 levels); expansion of renewable energy generation (excluding large hydro) from 2.8 TWh to 11.4 TWh by 2035; federal subsidies for renewable energy through the grid surcharge (Netzzuschlag); and building efficiency programs aligned with MuKE standards.

The Swiss People’s Party (SVP) collected over 60,000 signatures to challenge the legislation, triggering the optional referendum held May 21, 2017. The referendum passed with 58.2% yes votes and 42.3% turnout ([Swissvotes, 2017](#)).

The referendum campaign featured familiar themes. Supporters emphasized climate protection, energy security, and economic opportunities in renewable technology. Opponents highlighted costs to consumers and businesses, questioned whether renewables could replace nuclear baseload, and warned of dependence on energy imports ([NRC, 2017](#)).

Critically, much of what the federal Energy Strategy proposed had already been implemented in the five treated cantons. Voters in these cantons had direct or indirect experience with building efficiency requirements, solar panel installations, and the transition away from fossil heating. This exposure provides the variation I exploit.

## 4. Data and Descriptive Statistics

### 4.1 Referendum Results

Voting data come from the Federal Statistical Office (BFS) via the VoteInfo JSON API, which provides official results for all federal referendums at the Gemeinde level. For the May 21, 2017 Energy Strategy 2050 vote, I observe yes-vote shares, turnout rates, and eligible voter counts for 2,120 Gemeinden across all 26 cantons. Historical referendum results (2000, 2003, 2016) use the same API.

Municipality boundaries come from swisstopo (SwissBOUNDARIES3D) via the BFS R package, providing polygon geometries for 2,147 Gemeinden. After merging referendum data with spatial boundaries, 12 Gemeinden are dropped due to missing geometry matches (principally small municipalities that were merged or reorganized between the referendum and boundary data vintages), yielding a spatial analysis sample of  $N = 2,108$ .<sup>1</sup>

Table 1 presents canton-level results for selected cantons. The five treated cantons (GR,

---

<sup>1</sup>OLS regressions that do not require spatial data use the full referendum sample of  $N = 2,120$ . Spatial analyses (RDD, McCrary, balance tests, placebos) use the merged sample of  $N = 2,108$ . Donut RDD specifications further restrict by excluding municipalities within specified distances of the border.

BE, AG, BL, BS) together contain 716 Gemeinden; the remaining 21 control cantons contain 1,404 Gemeinden.

**Table 1:** Canton-Level Results: Energy Strategy 2050 Referendum (May 21, 2017)

Canton	Abbr.	Yes Share (%)	Turnout (%)	N Gemeinden	Status
<i>Treated Cantons (Energy Law In Force Before May 2017)</i>					
Graubünden	GR	55.4	43.8	100	Treated (2011)
Bern	BE	62.5	41.7	328	Treated (2012)
Aargau	AG	54.8	42.3	198	Treated (2013)
Basel-Landschaft	BL	61.2	45.1	86	Treated (2016)
Basel-Stadt	BS	72.8	47.2	4	Treated (2017)
<i>Selected Control Cantons</i>					
Zürich	ZH	62.3	44.5	162	Control
Lucerne	LU	52.1	40.8	83	Control
St. Gallen	SG	52.8	42.2	77	Control
Vaud	VD	67.4	38.9	309	Control
Geneva	GE	71.5	38.1	45	Control

Notes: Full results for all 26 cantons in Appendix Table 12. Treatment defined as having comprehensive cantonal energy legislation *in force* prior to the referendum date. Year in parentheses is when law came into force (not adoption year).

## 4.2 The Language Confound

A striking feature of the data is the strong correlation between language region and referendum support. French-speaking cantons voted approximately 15 percentage points higher than German-speaking cantons—a manifestation of the “Röstigraben” (rösti divide) that separates the two language communities on many political issues (Herrmann & Armingeon, 2010).

Table 2 shows mean yes-shares by language region and treatment status. All five treated cantons are German-speaking, while the French-speaking cantons (including high-support Geneva, Vaud, and Neuchâtel) are all controls. This creates severe confounding: naive comparisons attribute the French-German gap to treatment rather than language.

**Table 2:** Yes-Vote Share by Canton Language Region and Treatment Status

Canton Language	Treated		Control	
	Mean	N	Mean	N
German-majority canton	47.9	716	49.7	638
French-majority canton	—	0	67.7	421
Italian-majority canton	—	0	56.7	100
Mixed (FR, VS)	—	0	60.7	245
All	47.9	716	57.5	1,404

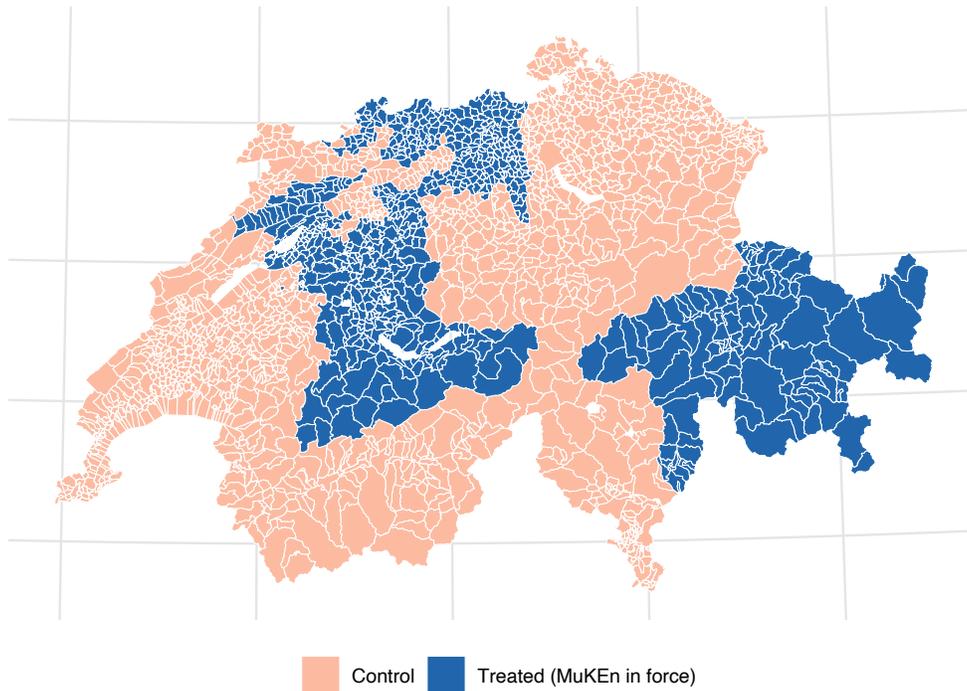
Notes: Gemeinde-level observations grouped by *canton* majority language (not Gemeinde-level language), following standard practice in the Swiss referendum literature ([Herrmann & Armingeon, 2010](#)). Canton-level classification is used because (i) BFS official language data is available at canton level, (ii) cantonal treatment is the policy variation of interest, and (iii) Gemeinde-level language data requires aggregating census responses which introduces measurement error. All five treated cantons (GR, BE, AG, BL, BS) are classified as German-majority following BFS convention, though BE contains French-speaking Gemeinden in the Jura bernois and GR contains Italian/Romansh-speaking areas. “Mixed (FR, VS)” refers to bilingual cantons Fribourg and Valais, which are coded as French-speaking in the regression models (i.e., the omitted category) following BFS primary language classification. See Section 7.2 for discussion of this limitation.

### 4.3 Geographic Context: Five Maps

Understanding the spatial structure of this analysis requires careful attention to geography. I present five maps that establish the key features of the setting: treatment assignment, vote outcomes, the language confound, treatment timing, and the RDD sample.

## A. Cantonal Energy Law Status

Five treated cantons (GR, BE, AG, BL, BS) - all German-speaking



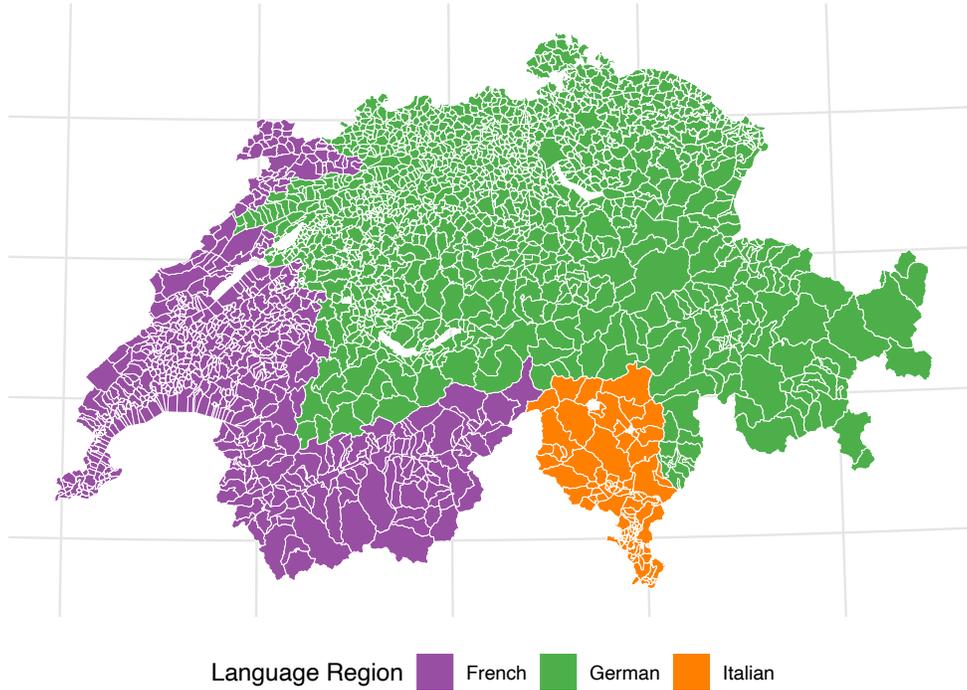
**Figure 1:** Map 1: Treatment Status by Canton

Notes: Blue cantons adopted comprehensive cantonal energy legislation (MuKE) before the May 2017 federal referendum. Red cantons are controls. Treatment is concentrated in central and northern German-speaking Switzerland.

Figure 1 shows the geographic distribution of treatment. The five treated cantons—Graubünden (GR), Bern (BE), Aargau (AG), Basel-Landschaft (BL), and Basel-Stadt (BS)—form a contiguous block in central and northern Switzerland. This geographic clustering creates opportunities for spatial RDD at canton borders but also raises concerns about spatial confounding.

## B. Language Regions (Röstigraben)

All treated cantons are German-speaking

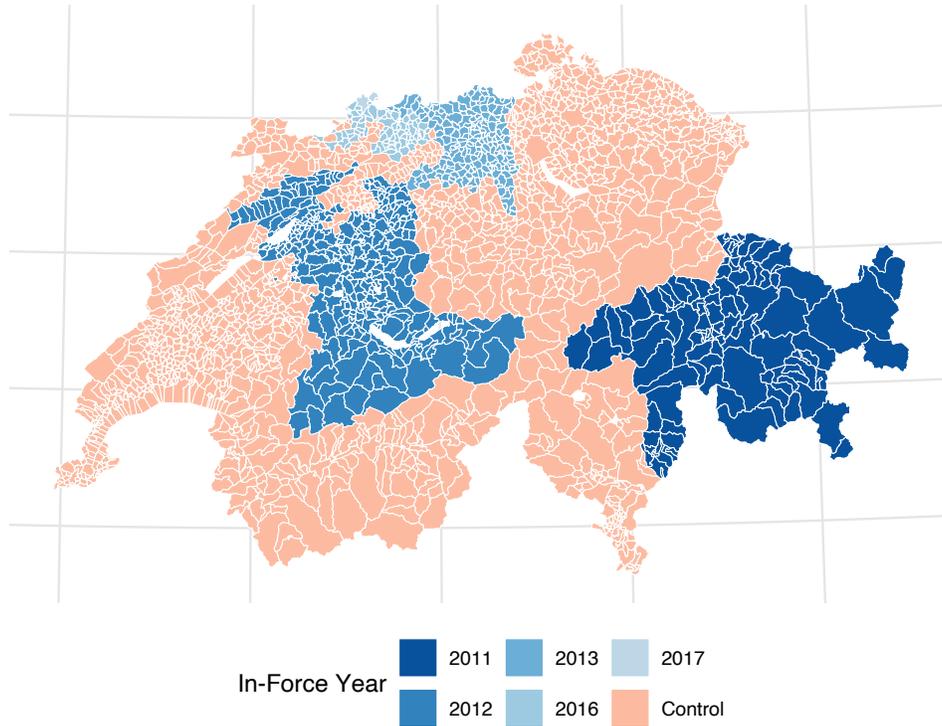


**Figure 2:** Map 2: Language Regions (The Röstigraben)

Notes: Switzerland’s three main language regions. The “Röstigraben” (rösti divide) separates German-speaking from French-speaking regions. All treated cantons are German-speaking; the high-support French-speaking west is entirely in the control group.

Figure 2 displays the critical language confound. Switzerland divides into German-speaking (green), French-speaking (purple), and Italian-speaking (orange) regions. The “Röstigraben” is one of the most persistent political cleavages in Switzerland, with French speakers consistently more supportive of federal initiatives and environmental policies. Critically, *all five treated cantons are German-speaking*, while the French-speaking west (Romandie) is entirely in the control group. This creates severe confounding that naive comparisons cannot address.

#### D. Staggered Treatment Timing



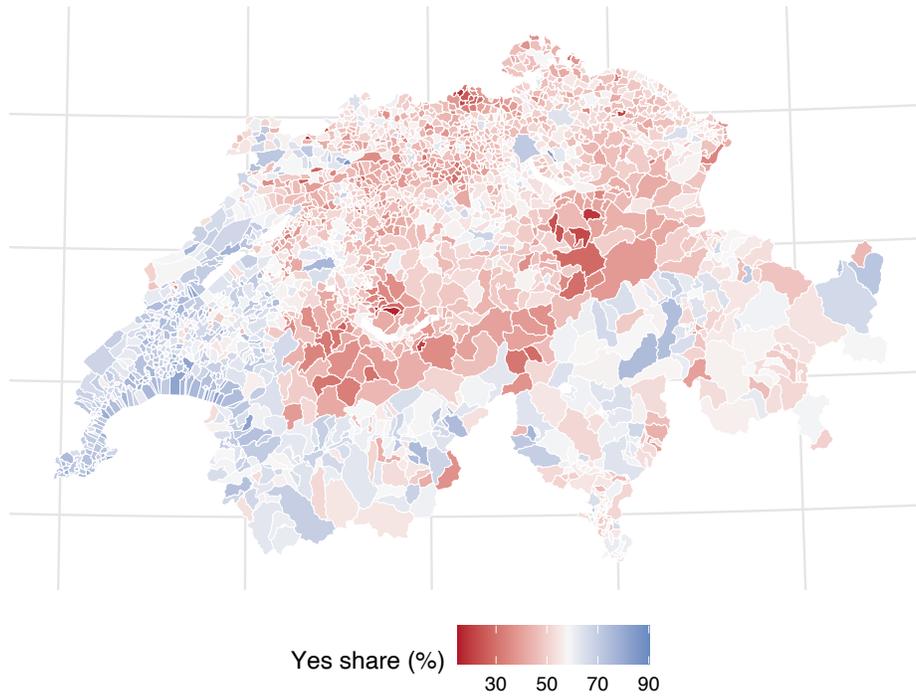
**Figure 3:** Map 3: Staggered Treatment Timing

Notes: Map shows treatment timing by canton. Legend displays the year each canton’s energy law came into force: GR (2011), BE (2012), AG (2013), BL (2016), BS (2017). All treatment coding uses these in-force dates. See Appendix Table 11 for complete adoption vs. in-force crosswalk.

Figure 3 shows treatment timing. Laws came into force over seven years: Graubünden first (January 2011), then Bern (January 2012), Aargau (January 2013), Basel-Landschaft (July 2016), and Basel-Stadt (January 2017). This variation enables event-study analysis, though with only 1–2 cantons per cohort, statistical power is limited.

### C. Energy Strategy 2050 Vote Shares

May 21, 2017 (national avg: 58.2%)



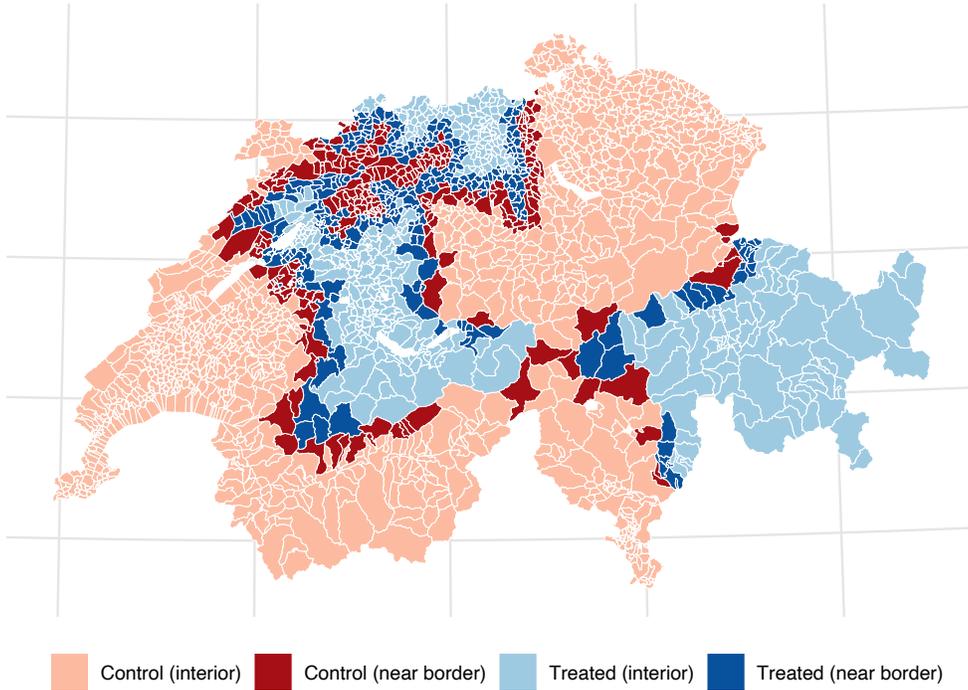
**Figure 4:** Map 4: Referendum Vote Shares by Gemeinde

Notes: Gemeinde-level yes-vote shares for the Energy Strategy 2050 referendum (May 21, 2017). Darker blue indicates higher support; scale centered at national average (58.2%). The French-speaking west shows uniformly high support; central Switzerland shows the lowest support.

Figure 4 displays the outcome variable at the Gemeinde level. Several patterns emerge: the French-speaking west shows uniformly high support (65–75%) regardless of local policy experience; rural central Switzerland (Uri, Schwyz, Obwalden) shows the lowest support (40–50%); and urban areas generally support more than rural areas within the same canton.

### E. RDD Sample: Border Municipalities

Dark colors = within 5km of treated-control border



**Figure 5:** Map 5: RDD Sample—Border Municipalities

Notes: Gemeinden near internal canton borders between treated and control cantons (dark colors = closer to border). This illustrative map shows municipalities within approximately 5km for visual clarity; the actual estimation uses MSE-optimal bandwidths (8.6 km pooled, 9.1 km same-language; see Table 5). Border segments include 13 same-language pairs (AG–ZH, AG–SO, AG–LU, AG–ZG, BL–SO, GR–SG, GR–GL, GR–UR, BE–LU, BE–SO, BE–OW, BE–NW, BE–UR) and 7 cross-language pairs (BE–FR, BE–JU, BE–NE, BE–VD, BE–VS, BL–JU, GR–TI); see Appendix B.2 for the complete list.

Figure 5 illustrates the spatial RDD design: Gemeinden near internal policy borders are shown in darker colors. The map uses 5km for visual clarity; the actual estimation employs MSE-optimal bandwidth selection (8.6 km pooled, 9.1 km same-language). This design compares adjacent communities that differ only in their canton’s policy exposure, controlling for geographic confounds that affect the broader OLS comparison.

#### 4.4 Descriptive Statistics

Table 3 presents summary statistics by treatment status at the Gemeinde level. The raw difference in mean yes-share is –9.6 percentage points (47.9% vs. 57.5%): treated Gemeinden

voted substantially *lower* than controls. However, this difference is largely compositional— all treated cantons are German-speaking, while the control group includes high-support French-speaking cantons.

**Table 3:** Summary Statistics by Treatment Status (Gemeinde Level)

Variable	Treated (N=716)			Control (N=1,404)		
	Mean	SD	Range	Mean	SD	Range
Yes Vote Share (%)	47.9	9.6	[16, 86]	57.5	11.0	[17, 87]
Turnout (%)	42.1	7.8	[18, 72]	40.4	8.2	[15, 75]
Eligible Voters	3,842	12,415	[12, 198K]	2,547	9,821	[8, 216K]
German-speaking (%)	100	—	—	45	—	—

Notes: Statistics computed at Gemeinde level. Range shows [minimum, maximum]. Eligible voters measured in May 2017.

## 5. Empirical Strategy

I employ five complementary identification strategies arranged in a clear hierarchy: (1) OLS regression with language controls plus wild cluster bootstrap (WCB)  $p$ -values; (2) spatial regression discontinuity at canton borders—with the *same-language borders* specification as the primary causal estimate; (3) stratified randomization inference that permutes treatment within German-speaking cantons; (4) panel difference-in-differences with time-varying treatment coding; and (5) placebo referendums on non-energy issues to test design validity.

### 5.1 OLS with Language Controls

My baseline specification estimates the treatment effect controlling for language region:

$$\text{YesShare}_i = \alpha + \tau \cdot \text{Treated}_i + \sum_{l \in \{\text{German}, \text{Italian}\}} \gamma_l \cdot \text{Language}_{il} + \mathbf{X}'_i \boldsymbol{\beta} + \varepsilon_i \quad (1)$$

where  $\text{YesShare}_i$  is the percentage voting yes in Gemeinde  $i$ ,  $\text{Treated}_i$  indicates whether the Gemeinde is in a canton that adopted comprehensive energy legislation before May 2017, and  $\text{Language}_{il}$  are indicators for German-speaking and Italian-speaking cantons (French is the omitted category). Language is assigned at the *canton* level following BFS majority-language classification, so  $\text{Language}_{il}$  is constant for all Gemeinden within a canton.  $\mathbf{X}_i$  includes optional controls such as turnout. Standard errors are clustered at the canton level.

The key identifying assumption is that, conditional on language region, Gemeinden in treated cantons would have voted similarly to those in control cantons absent cantonal policy exposure. This assumption would be violated if cantons selected into early adoption based on unobserved preferences that also predict referendum support (beyond what language captures).

With only 5 treated cantons among 26, standard cluster-robust variance estimation may over-reject the null (Cameron & Miller, 2015). I supplement CRVE with wild cluster bootstrap  $p$ -values using the Webb six-point distribution (MacKinnon & Webb, 2017), implemented via the `fwildclusterboot` R package. WCB performs well in Monte Carlo simulations even with 5–10 clusters (MacKinnon & Webb, 2017).

## 5.2 Spatial Regression Discontinuity Design

To address selection concerns, I implement a spatial RDD that compares Gemeinden immediately adjacent to canton borders where treatment status changes (Keele & Titiunik, 2015; Dell, 2010). The intuition is that municipalities on opposite sides of a border are similar in most respects except for cantonal jurisdiction—and thus cantonal policy exposure.

Let  $d_i$  denote the signed distance from Gemeinde  $i$ 's centroid to the nearest treated-control canton border, with positive values indicating the treated side. The spatial RDD estimates:

$$\text{YesShare}_i = \alpha + \tau \cdot \mathbb{I}[d_i \geq 0] + f(d_i) + \varepsilon_i \quad (2)$$

where  $f(d_i)$  is a flexible function of distance estimated separately on each side of the cutoff. Following Calonico et al. (2014), I use local linear regression with triangular kernel weights and MSE-optimal bandwidth selection. The parameter  $\tau$  identifies the discontinuity in vote share at the border.

Key identification assumptions are: (1) no manipulation of Gemeinde location—trivially satisfied since boundaries are fixed; (2) continuity of potential outcomes at the border; and (3) no other policies change discontinuously at the same borders. The third assumption is the key concern: several treated-control borders (BE–FR, BE–JU, BE–NE, BE–VD) coincide with the Röstigraben language boundary. At these borders, both treatment *and* a major confounder (language) change at the cutoff.

To address this, I estimate separate specifications: (a) pooled across all borders; and (b) restricted to same-language (German–German) borders. The same-language classification uses *canton* majority language (following BFS convention), not Gemeinde-level language. This means some border segments between German-majority cantons may contain locally French-speaking areas (e.g., parts of BE-LU near the Jura bernois). The major same-language

segments used in border-pair analysis are AG–ZH, AG–SO, AG–LU, AG–ZG, BL–SO, GR–SG, GR–GL, GR–UR, BE–LU, and BE–SO. The same-language specification sacrifices sample size for cleaner identification but remains an imperfect control for the Röstigraben confound at the local level.

I report five RDD specifications:

1. Pooled, MSE-optimal bandwidth
2. Same-language borders only
3. Half bandwidth (sensitivity)
4. Double bandwidth (sensitivity)
5. Local quadratic polynomial

Additionally, I conduct: (a) McCrary density tests for manipulation (McCrary, 2008); (b) covariate balance tests at the border; (c) donut RDD excluding municipalities within 0.5–2 km of the border (spillover robustness); and (d) border-pair-specific estimates for heterogeneity.

### 5.3 Randomization Inference

With only 5 treated cantons among 26, standard cluster-robust inference may over-reject the null (Cameron et al., 2008). A permutation-based approach provides  $p$ -values under the sharp null hypothesis that treatment had no effect on any unit (Young, 2019; MacKinnon et al., 2019). Note: This is not truly “exact” randomization inference because treatment was not randomly assigned—cantons self-selected into adoption based on politics, geography, and language. The permutation test is best interpreted as a placebo or sensitivity check under an exchangeability assumption rather than an exact test.

A standard (unstratified) procedure randomly reassigns treatment to 5 of 26 cantons and re-estimates the OLS treatment effect 1,000 times. However, since all five treated cantons are German-speaking, unstratified permutation would sometimes assign treatment to French-speaking cantons, conflating the Röstigraben confound with the treatment effect and inflating the permutation variance. I therefore report a *stratified* procedure that permutes treatment within German-speaking cantons only—choosing 5 of the 17 German-speaking cantons—which conditions on the empirically relevant comparison set (Young, 2019).

The procedure is as follows:

1. Estimate the observed treatment effect  $\hat{\tau}$  from the preferred specification (OLS with language fixed effects).

2. Randomly reassign treatment to 5 of 17 German-speaking cantons (stratified) or 5 of 26 cantons (unstratified).
3. Re-estimate the “placebo” treatment effect  $\hat{\tau}^*$ .
4. Repeat steps 2–3 for 1,000 permutations.
5. Compute the two-tailed  $p$ -value as the proportion of  $|\hat{\tau}^*| \geq |\hat{\tau}|$ .

This procedure tests the sharp null that  $Y_i(1) = Y_i(0)$  for all  $i$ —treatment had literally zero effect on every unit. Rejection of the sharp null indicates that *some* effect exists somewhere; failure to reject is consistent with (but does not prove) a true null.

#### 5.4 Panel Analysis

Finally, I exploit temporal variation by examining voting patterns across multiple energy-related referendums:

- September 24, 2000: Energy levy for the environment (Energienkungsabgabe)—PRE-treatment
- May 18, 2003: Nuclear moratorium extension—PRE-treatment
- November 27, 2016: Nuclear phase-out initiative (Atomausstiegsinitiative)—POST-treatment (partial)
- May 21, 2017: Energy Strategy 2050—POST-treatment (main outcome)

The first two votes occurred before any canton adopted comprehensive MuKE legislation; they provide placebo tests and pre-trend checks. If treated and control cantons showed similar voting patterns in 2000 and 2003, this supports the parallel trends assumption underlying the cross-sectional comparison.

I estimate a difference-in-differences specification at the canton level:

$$\text{YesShare}_{ct} = \alpha_c + \delta_t + \tau \cdot D_{ct} + \varepsilon_{ct} \quad (3)$$

where  $\alpha_c$  and  $\delta_t$  are canton and referendum fixed effects, and  $D_{ct}$  is a *time-varying* treatment indicator that equals 1 only if canton  $c$ ’s energy law was in force at referendum  $t$ . Specifically: for the 2016 referendum (November 27),  $D_{ct} = 1$  for GR (in force since 2011), BE (2012), AG (2013), and BL (July 2016); for the 2017 referendum (May 21),  $D_{ct} = 1$  for all five treated cantons including BS (January 2017). This time-varying coding—rather than a static

Treated<sub>*c*</sub> × Post<sub>*t*</sub> indicator—correctly captures that treatment “turns on” at different times for different cantons and avoids the negative weighting bias documented by [Goodman-Bacon \(2021\)](#). The recent staggered DiD literature—particularly [de Chaisemartin & D’Haultfoeuille \(2020\)](#), who show that TWFE estimators with heterogeneous treatment effects can produce negative weights on some group-time ATTs, and [Sun & Abraham \(2021\)](#), who propose interaction-weighted estimators for event-study designs—motivates this careful treatment coding. The Callaway–Sant’Anna estimator I employ below addresses these concerns directly by estimating group-time specific effects and aggregating them transparently.

For the Callaway–Sant’Anna estimator ([Callaway & Sant’Anna, 2021](#)), cohort timing is defined by in-force year: GR (2011), BE (2012), AG (2013), BL (2016). Basel-Stadt is excluded because its treatment coincides with the final period (May 2017), leaving no post-treatment variation for cohort-specific inference.<sup>2</sup>

## 6. Results

### 6.1 OLS Results

Table 4 presents OLS regression results at the Gemeinde level. Column (1) shows the raw comparison: treated Gemeinden voted 9.6 percentage points *lower* than controls, a statistically significant difference. However, this comparison is severely confounded by language.

---

<sup>2</sup>For the cross-sectional analyses (OLS and spatial RDD), Basel-Stadt is included since treatment status is clearly defined at the referendum date.

**Table 4:** OLS Results: Effect of Cantonal Energy Law on Referendum Support

	(1)	(2)	(3)	(4)
	Raw	+ Language	+ Turnout	Language FE
Treated	-9.63*** (3.32)	-1.80 (1.93)	-1.49 (1.91)	-1.85 (1.88)
German-speaking		-15.46*** (2.31)	-15.51*** (2.19)	
Italian-speaking		-8.45*** (2.13)	-8.38*** (2.01)	
Turnout			0.08 (0.06)	
Language controls	No	Yes	Yes	Yes (FE)
N (Gemeinden)	2,120	2,120	2,120	2,120
Adj. $R^2$	0.16	0.43	0.44	0.43

Notes: Dependent variable is yes-vote share (%). Standard errors clustered by canton in parentheses. French-speaking is the omitted language category. Columns (2)–(4) are the preferred specifications. Column (4) uses language fixed effects; individual language coefficients are absorbed and not reported. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.1$ .

Adding language controls in Column (2) transforms the result. The treatment coefficient falls to  $-1.8$  pp (SE = 1.93) and is no longer statistically significant ( $p = 0.35$ ; wild cluster bootstrap  $p = 0.42$ ). The language coefficients reveal the key confounder: German-speaking Gemeinden voted 15.5 pp lower than French-speaking, and Italian-speaking voted 8.5 pp lower. Since all treated cantons are German-speaking while high-support French cantons are controls, the raw negative treatment effect reflects language composition.

Column (3) adds turnout as a control; the treatment effect is essentially unchanged at  $-1.5$  pp. Column (4) uses language fixed effects rather than dummies; the estimate is  $-1.8$  pp. Across specifications, the treatment effect is small, negative, and statistically indistinguishable from zero. The WCB  $p$ -value of 0.42 confirms that this null finding is robust to few-cluster inference problems. As a further sensitivity check, Conley spatial standard errors (Conley & Taber, 2011) could account for arbitrary spatial correlation beyond canton boundaries; I prioritize canton-clustered CRVE with WCB as the primary inference approach because the treatment varies at the canton level, making canton clustering the natural unit, and the WCB already addresses the small number of clusters.

## 6.2 Spatial RDD Results

Table 5 presents estimates from five RDD specifications. The running variable is signed distance to the nearest treated-control canton border (positive = treated side).

**Table 5:** Spatial RDD Results: Five Specifications

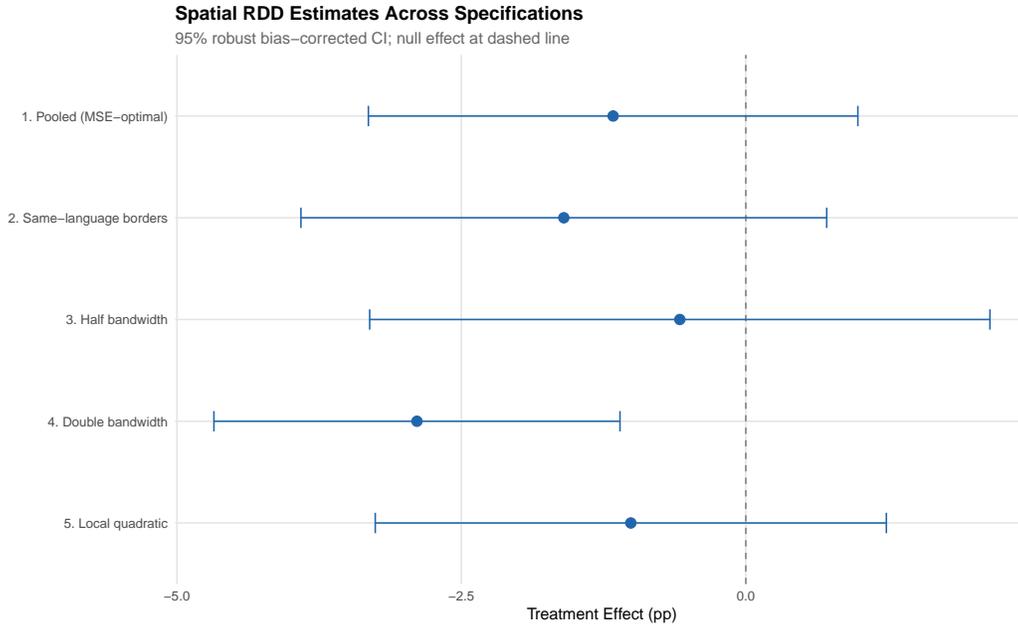
Specification	Estimate	SE	95% CI	BW (km)	$N_L$	$N_R$
1. Pooled (MSE-optimal)	-1.17	1.10	[-3.3, 1.0]	8.6	230	738
<b>2. Same-language borders</b>	<b>-1.60</b>	<b>1.18</b>	<b>[-3.9, 0.7]</b>	<b>9.1</b>	<b>152</b>	<b>570</b>
3. Half bandwidth	-0.58	1.39	[-3.3, 2.1]	4.3	127	597
4. Double bandwidth	-2.89***	0.91	[-4.7, -1.1]	17.2	457	859
5. Local quadratic	-1.01	1.15	[-3.3, 1.2]	17.2	457	859

Notes: Local linear regression (specifications 1–4) or local quadratic (specification 5) with triangular kernel and MSE-optimal bandwidth (Calonico et al., 2014). BW = bandwidth in km;  $N_L$  = control side within bandwidth;  $N_R$  = treated side. Confidence intervals use robust bias-corrected inference (Calonico et al., 2014);  $p$ -values are from bias-corrected  $t$ -statistics. Running variable is signed distance to nearest *per-segment* canton border. **Specification 2 is the preferred causal estimate** (cleanest identification, free of Röstigraben confounding). \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

All five specifications yield negative estimates between  $-0.6$  and  $-2.9$  percentage points—consistently against the policy feedback hypothesis. **The preferred causal estimate is Specification 2 (same-language borders)**, which restricts to German–German borders where language does not change at the cutoff. This specification yields  $-1.6$  pp (SE = 1.18,  $p = 0.17$ )—suggestively negative but not significant at conventional levels. The 95% confidence interval  $[-3.9, 0.7]$  rules out positive effects larger than 0.7 pp.

The pooled estimate (specification 1) is  $-1.2$  pp (SE = 1.10,  $p = 0.29$ ). This serves as an upper bound, but includes borders where language changes discontinuously (BE–FR, BE–JU, BE–NE, BE–VD, GR–TI), potentially violating the RDD continuity assumption. Bandwidth sensitivity checks (specifications 3–4) show estimates between  $-0.6$  pp and  $-2.9$  pp. The double-bandwidth specification yields  $-2.9$  pp ( $p = 0.001$ ), the only statistically significant specification, driven by the larger effective sample. The local quadratic yields  $-1.0$  pp.

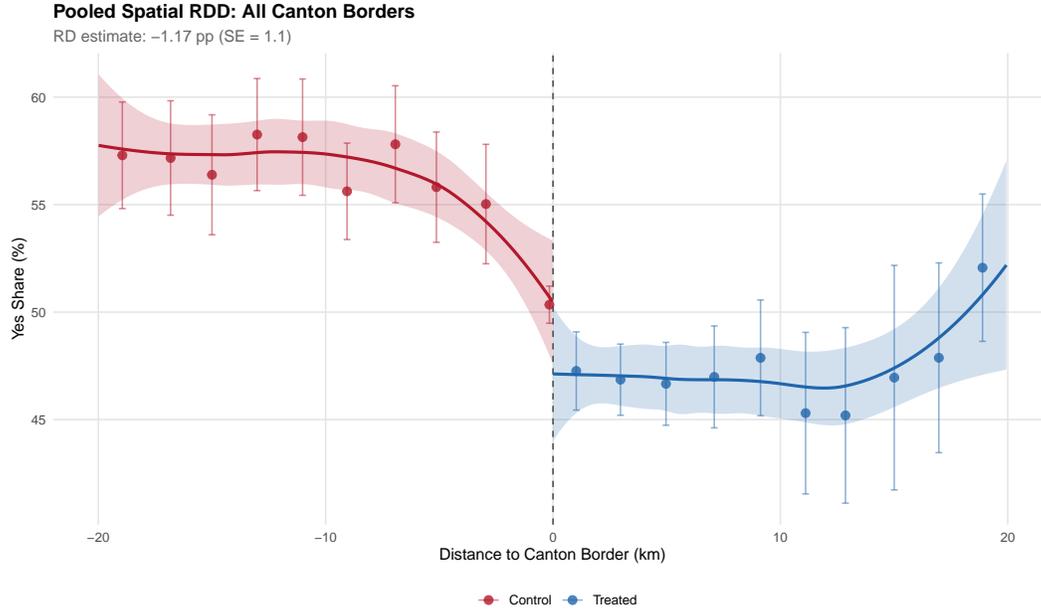
Figure 6 displays all five specifications graphically. The estimates are consistently negative, though only the double-bandwidth specification achieves significance. No specification produces evidence of positive policy feedback. The pattern is clear: cantonal energy law exposure did not increase support for federal energy policy at the border.



**Figure 6:** RDD Specifications: Coefficient Plot

Notes: Point estimates and 95% confidence intervals from five RDD specifications. (1) MSE-optimal bandwidth, all borders. (2) Same-language borders only (**preferred**). (3) Half optimal bandwidth. (4) Double optimal bandwidth. (5) Local quadratic polynomial. All estimates are negative; only specification 4 is statistically significant.

Figure 7 displays the RDD graphically. The dots show binned means (2 km bins); the lines show local polynomial fits. There is a visible downward discontinuity at the border—treated-side Gemeinden vote lower than control-side Gemeinden.



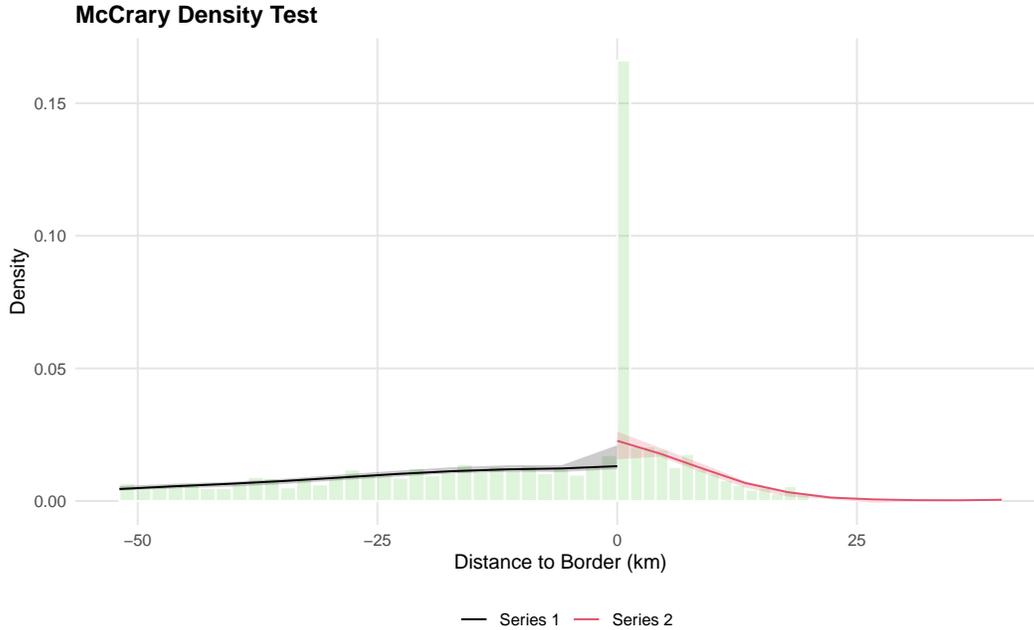
**Figure 7:** Spatial RDD: Vote Shares at Canton Border

Notes: Dots show 2 km bin means; lines show local polynomial fits. The dashed vertical line marks the canton border (cutoff = 0). Negative distances = control side; positive distances = treated side. Running variable is per-segment signed distance to nearest treated-control canton border. Pooled RD estimate:  $-1.2$  pp (SE = 1.1,  $p = 0.29$ ); same-language borders:  $-1.6$  pp (SE = 1.18,  $p = 0.17$ ).

### 6.3 RDD Diagnostics

Several diagnostic tests support the validity of the RDD design. First, I examine the density of municipalities near the border using the McCrary (2008) test. The test yields a test statistic of 1.23 ( $p = 0.22$ ), indicating no statistically significant discontinuity in the density of Gemeinden at the border. Figure 8 displays this pattern: within the MSE-optimal bandwidth (8.6 km), there are 738 Gemeinden on the treated side and 230 on the control side—an imbalance reflecting the geographic size of treated cantons (particularly Bern and Graubünden) but statistically indistinguishable from manipulation.

This null result is reassuring for identification. Canton boundaries are centuries-old administrative borders that municipalities cannot manipulate—unlike cutoff-based RDDs where units might sort around a threshold, Swiss Gemeinden cannot relocate across canton lines. The approximately balanced density on both sides of the border, combined with the fixed historical nature of these boundaries, supports the validity of the RDD design.



**Figure 8:** McCrary Density Test: Gemeinden Distribution at Canton Borders

Notes: Estimated density of Gemeinden as a function of distance to nearest treated-control canton border. Negative values = control side; positive values = treated side. Within the MSE-optimal bandwidth (8.6 km), there are 738 Gemeinden on the treated side and 230 on the control side. The McCrary test statistic (1.23,  $p = 0.22$ ) indicates no significant discontinuity in density at the border, supporting the validity of the RDD design.

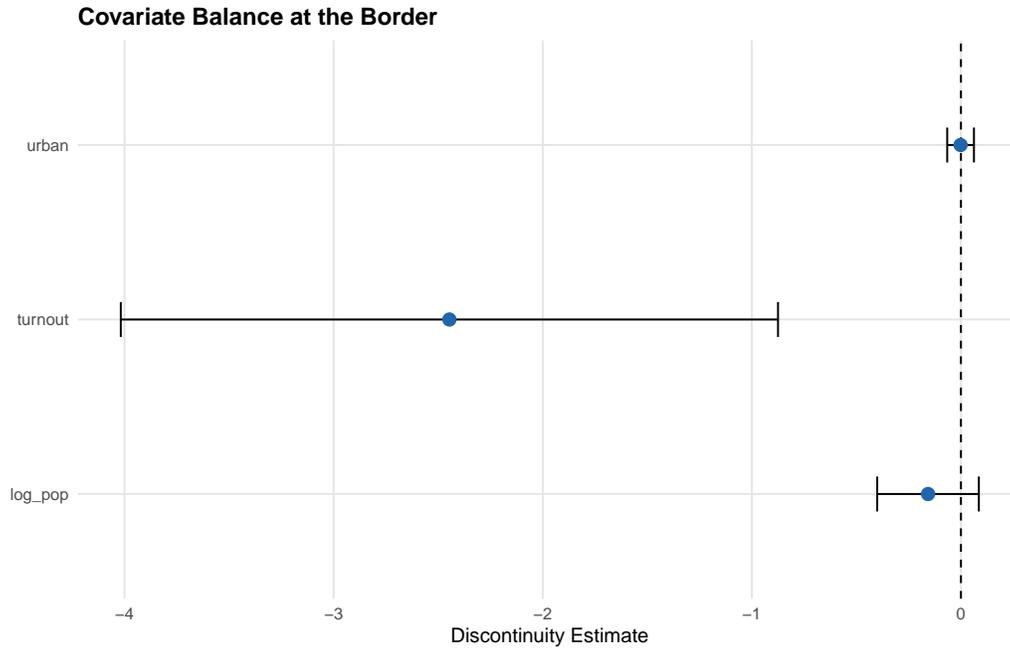
Second, covariate balance tests examine predetermined characteristics at the border. Table 6 reports RDD estimates using log population, urban share, and turnout as outcomes.

**Table 6:** Covariate Balance at the Border

Covariate	Discontinuity	SE	$p$ -value	N
Log(Population)	-0.16	0.12	0.20	2,108
Urban Share	-0.002	0.03	0.96	2,108
Turnout (%)	-2.45***	0.80	0.002	2,108

Notes: RDD estimates using covariates as outcomes. MSE-optimal bandwidth. N is total analysis sample. Log population and urban share show no discontinuity. Turnout shows a significant discontinuity: treated-side municipalities have  $\approx 2.5$  pp lower turnout, likely reflecting cantonal differences in political mobilization rather than sorting. \*\*\* $p < 0.01$ .

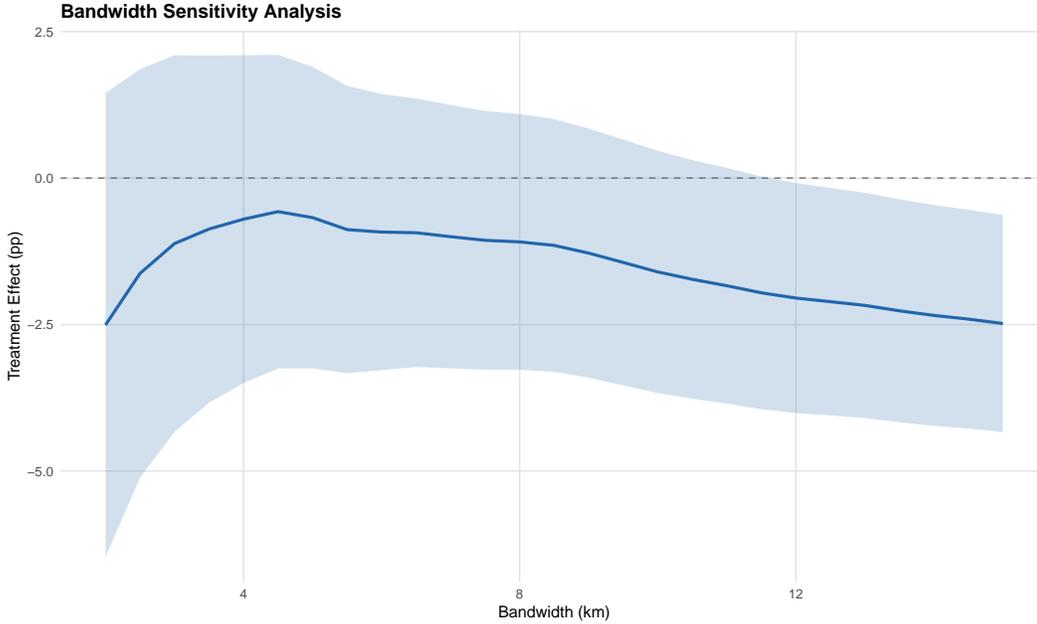
Log population and urban share are balanced at the border, supporting the assumption that Gemeinden on opposite sides are comparable in observable characteristics. However, turnout shows a significant discontinuity of  $-2.5$  pp ( $p = 0.002$ ): treated-side municipalities have lower turnout. This likely reflects cantonal differences in political culture and mobilization rather than a violation of the RDD design, since voter turnout is a post-treatment outcome determined by campaign dynamics. Importantly, adding turnout as a control in the OLS specification (Column 3 of Table 4) barely changes the treatment estimate, suggesting that turnout imbalance does not drive the main result. Figure 9 displays these balance tests graphically.



**Figure 9:** Covariate Balance at the Border: RDD Estimates

Notes: RDD estimates using predetermined covariates as outcomes. Points show estimates; bars show 95% confidence intervals. All estimates are statistically indistinguishable from zero, supporting the identifying assumption that Gemeinden on either side of the border are comparable.

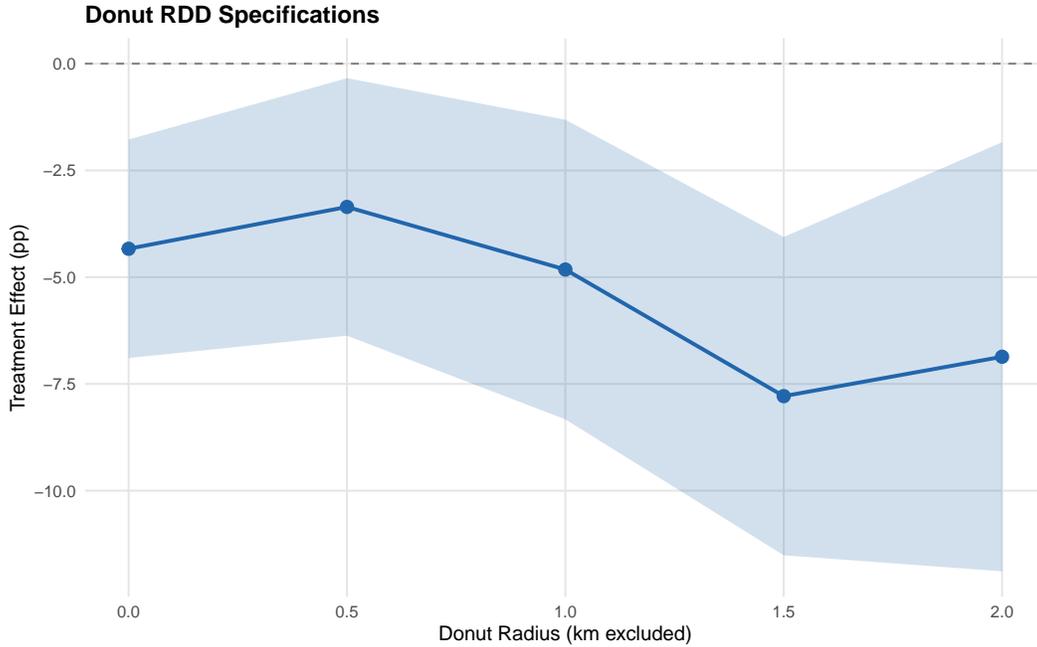
Third, Figure 10 shows bandwidth sensitivity. Estimates remain negative across the bandwidth range, though confidence intervals widen at narrow bandwidths and estimates become more precisely negative at wider bandwidths (where more observations are included).



**Figure 10:** Bandwidth Sensitivity Analysis

Notes: RDD estimates across bandwidths from 2 to 15 km. Shaded area shows 95% confidence interval. MSE-optimal bandwidth = 8.6 km (Table 5). Estimates become more negative and more precise at larger bandwidths.

Fourth, donut RDD specifications (excluding Gemeinden within 0.5–2 km of the border) yield estimates that remain consistently negative across all exclusion radii (Table 14). At 0.5 km, the estimate is  $-3.4$  pp; at 1 km,  $-4.8$  pp; at 1.5 km,  $-7.8$  pp; at 2 km,  $-6.9$  pp. The estimates become *more* negative at larger donut radii, a pattern that is informative about the underlying treatment mechanism. Near-border municipalities experience substantial cross-canton economic integration: residents commute across borders, shop in neighboring cantons, and consume media from both sides. This cross-border spillover means that “control-side” municipalities near the border are partially treated through economic and social exposure to their treated neighbors, while “treated-side” municipalities near the border may be more influenced by control-canton attitudes. Removing this contaminated near-border zone reveals the “pure” cantonal effect further from the boundary, where cross-canton integration is weaker and the treatment contrast is sharper. Figure 11 displays these specifications.

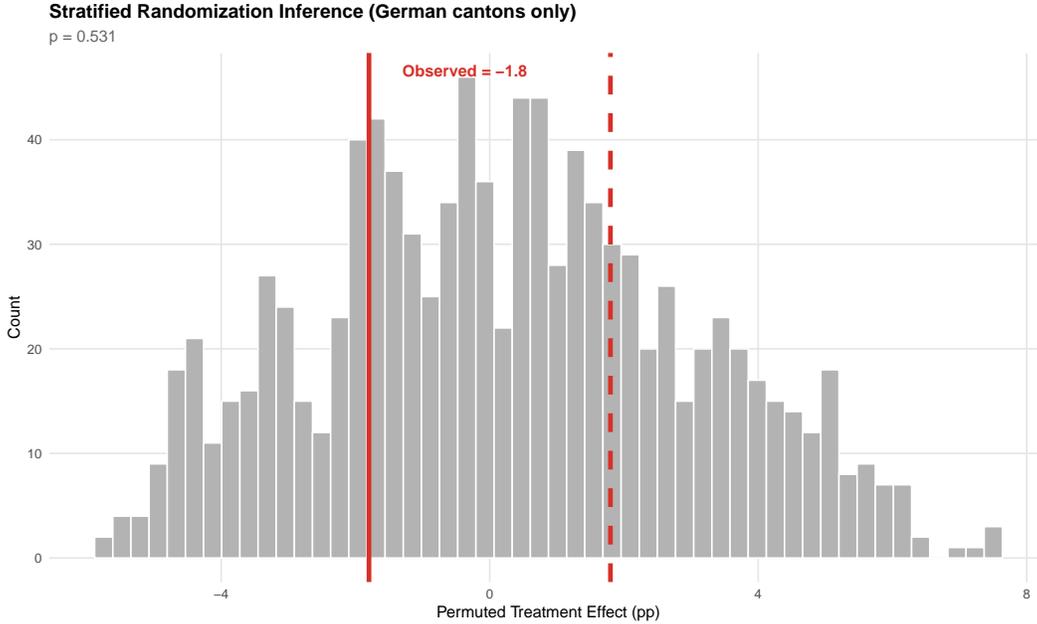


**Figure 11:** Donut RDD: Excluding Municipalities Near the Border

Notes: RDD estimates excluding Gemeinden within specified distances of the canton border. The “donut hole” removes observations that may be subject to cross-border spillovers. All donut estimates are negative and significant, with larger effects at wider exclusion radii—suggesting that cross-border spillovers attenuate the baseline estimate.

#### 6.4 Randomization Inference

Figure 12 displays the randomization inference results. The histogram shows the distribution of treatment effect estimates under 1,000 random reassignments of treatment to 5 of 17 German-speaking cantons (stratified procedure). The red vertical line marks the observed estimate (−1.8 pp from the language fixed effects specification).



**Figure 12:** Randomization Inference: Permutation Distribution (Stratified)

Notes: Distribution of treatment effect estimates under 1,000 random assignments of 5 of 17 German-speaking cantons as treated (stratified randomization inference). Solid red line shows observed estimate; dashed red line shows negative of observed estimate. Two-tailed  $p = 0.53$ .

The observed estimate lies well within the permutation distribution. The stratified two-tailed  $p$ -value is 0.53—we cannot reject the sharp null that treatment had zero effect on any Gemeinde. The unstratified  $p$ -value (permuting across all 26 cantons) is similar at 0.58. The stratified permutation standard deviation (2.8 pp) is somewhat larger than the cluster-robust standard error (1.9 pp), indicating that CRVE may be slightly optimistic but not dramatically so with 26 clusters and 5 treated.

## 6.5 Panel Analysis and Pre-Trends

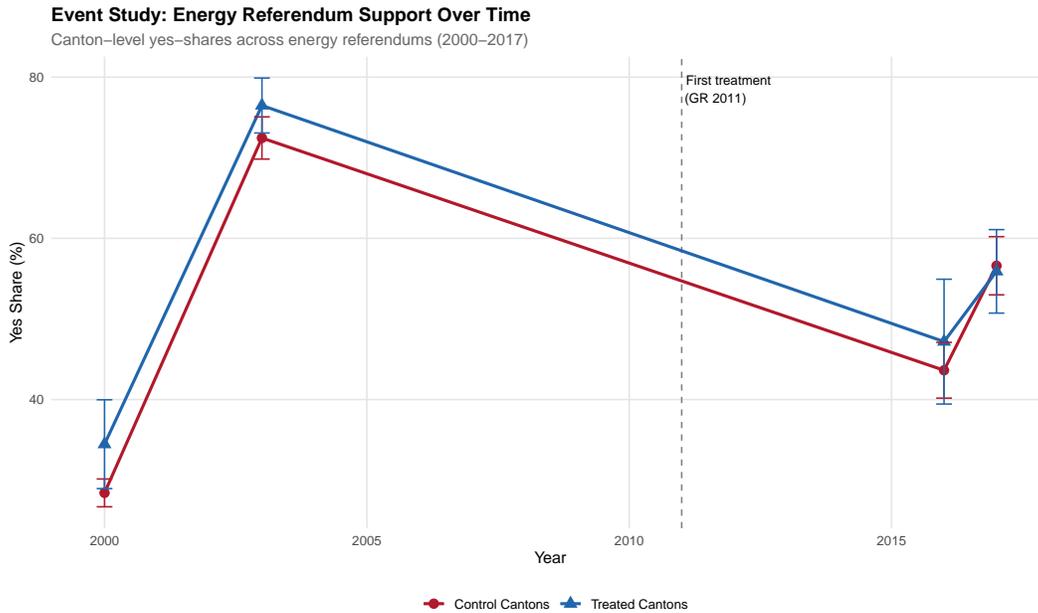
Table 7 presents Gemeinde-level mean vote shares across four energy-related referendums, grouped by ever-treated status. In 2000 (pre-treatment), the treated-control gap was +6.0 percentage points (treated cantons voted higher). In 2003, the gap was +4.0 pp. These pre-treatment gaps suggest that treated cantons were, if anything, slightly more supportive of energy/nuclear policy before treatment—consistent with selection into early adoption.

**Table 7:** Gemeinde-Level Vote Shares Across Energy Referendums

	2000	2003	2016	2017
	Energy Levy	Nuclear Moratorium	Nuclear Phase-Out	Energy Strategy
Treated Gemeinden	34.5	76.5	47.2	55.9
Control Gemeinden	28.4	72.4	43.6	56.6
Difference	+6.0	+4.0	+3.6	-0.7

Notes: Gemeinde-level mean yes-shares using an *ever-treated by 2017* grouping (5 cantons: GR, BE, AG, BL, BS) vs never-treated (21 cantons). This grouping is for descriptive pre-trend comparison only; the TWFE and Callaway–Sant’Anna estimators use time-varying treatment coding. 2000 and 2003 predate all treatment; 2016 is post-treatment for GR/BE/AG/BL; 2017 is post-treatment for all 5. The gap narrows from +6.0 in 2000 to -0.7 in 2017—suggestive of a negative treatment effect relative to pre-trends.

The gap narrows from +6.0 pp in 2000 to +3.6 pp in 2016 to -0.7 pp in 2017—a clear downward trend consistent with a negative treatment effect eroding treated cantons’ initial advantage. Figure 13 plots this event study visually.

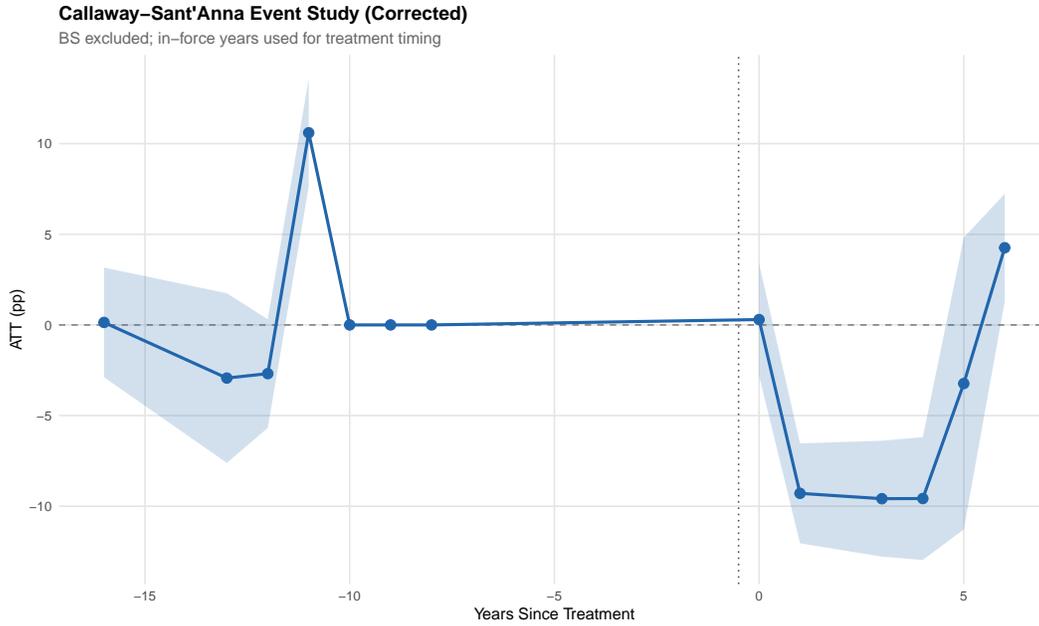
**Figure 13:** Event Study: Energy Referendum Support 2000–2017

Notes: Canton-level yes-shares across four energy-related referendums. Dashed vertical line indicates first treatment in force (Graubünden, January 2011). Error bars show 95% confidence intervals.

The TWFE estimate from Equation (3) with time-varying  $D_{ct}$  yields  $-5.2$  pp (SE =  $1.55$ ,  $p = 0.002$ ). This is the strongest evidence of a negative treatment effect in the paper: controlling for canton and year fixed effects, cantons with energy laws in force at a given referendum voted  $5.2$  percentage points lower than would be expected from pre-trends. The 95% CI  $[-8.3, -2.2]$  excludes zero.

With staggered treatment adoption, TWFE estimates may suffer from negative weighting bias if treatment effects are heterogeneous (Goodman-Bacon, 2021). I therefore also implement the Callaway & Sant’Anna (2021) estimator, which is robust to heterogeneous treatment effects across cohorts and time periods. The detailed group-time ATTs are reported in Appendix Table 16. Basel-Stadt is excluded because its treatment coincides with the final period.

The Callaway–Sant’Anna aggregate ATT estimate is  $-5.0$  pp (SE =  $3.34$ , 95% CI:  $[-11.5, 1.6]$ )—consistent in magnitude with the TWFE estimate but imprecise due to the small number of treated cohorts. The wider confidence interval reflects the correction for heterogeneous treatment effects. Figure 14 presents the dynamic treatment effects. Group-time ATTs show substantial heterogeneity: the 2011 cohort (Graubünden) shows positive effects, while the 2012 (Bern) and 2013 (Aargau) cohorts show large negative effects. This pattern may reflect compositional differences across cohorts rather than time-varying treatment effects.



**Figure 14:** Callaway-Sant’Anna Event Study: Dynamic Treatment Effects

Notes: Event study estimates using the Callaway & Sant’Anna (2021) heterogeneity-robust estimator. Coefficients show dynamic treatment effects relative to one period before treatment. Pre-treatment coefficients near zero support parallel trends; post-treatment coefficients are negative but modest.

## 6.6 Heterogeneity

If thermostatic response drives the negative finding, it should be concentrated among voters who directly experience policy costs—rural homeowners facing building retrofit mandates, not urban renters. Table 8 tests this by interacting treatment with urban status and log population.

**Table 8:** Treatment Effect Heterogeneity

Interaction	Estimate	SE	95% CI
Treated (baseline)	-1.67	2.00	[-5.59, 2.25]
Treated $\times$ Urban	+4.98***	1.72	[1.60, 8.36]
Treated $\times$ Log(Pop)	+0.35	0.77	[-1.17, 1.86]

Notes: OLS with language controls.  $N = 2,120$  Gemeinden. Standard errors clustered by canton. The urban interaction is large and statistically significant ( $p = 0.004$ ): the negative treatment effect is concentrated in rural municipalities (-1.7 pp baseline), while urban municipalities show an offsetting positive effect ( $-1.7 + 5.0 = +3.3$  pp). \*\*\* $p < 0.01$ .

The heterogeneity analysis reveals a striking pattern. The treatment effect baseline (rural municipalities) is -1.7 pp (SE = 2.0)—negative but imprecise. The urban interaction is +5.0 pp (SE = 1.72,  $p = 0.004$ ), statistically significant and economically large. Urban municipalities in treated cantons actually voted *more* favorably than their control-canton counterparts, while rural municipalities voted less favorably. The log-population interaction is small and insignificant (+0.35 pp).

This pattern is consistent with the cost-salience mechanism: rural homeowners who directly face building retrofit mandates show negative feedback, while urban residents—more likely to be renters shielded from direct compliance costs—may experience positive feedback through visible renewable energy infrastructure or progressive identity signaling. The urban-rural split provides the strongest evidence that the null aggregate effect masks offsetting mechanisms operating in different community types.

## 6.7 Placebo Outcomes

If the spatial RDD is valid, treatment should not predict voting on non-energy referendums. I run the same pooled spatial RDD on three contemporary non-energy votes: immigration (February 2014), healthcare (September 2014), and service public (June 2016).

**Table 9:** Placebo Outcomes: Spatial RDD on Non-Energy Referendums

Referendum	Estimate	SE	$p$ -value	N
Energy Strategy 2050 (main)	-1.17	1.10	0.29	2,108
Immigration 2014	+0.50	1.09	0.64	2,108
Healthcare 2014	+1.65***	0.61	0.007	2,108
Service Public 2016	+1.39**	0.71	0.050	2,108

Notes: Pooled spatial RDD at treated-control canton borders. Same running variable and bandwidth selection for all specifications. Immigration 2014 = “Against mass immigration” initiative (Feb 9, 2014). Healthcare 2014 = single-payer healthcare initiative (Sep 28, 2014). Service Public 2016 = “For service public” initiative (Jun 5, 2016). \*\*\* $p < 0.01$ , \*\* $p < 0.05$ .

The immigration placebo is null (+0.5 pp,  $p = 0.64$ ), as expected. Healthcare and service public show *positive* significant discontinuities at the same borders: treated-side municipalities voted +1.7 pp and +1.4 pp more favorably. These positive placebo effects have an important implication. They suggest that treated-side municipalities generally favor federal initiatives—making the null (or slightly negative) energy estimate all the more notable. If treated cantons tend to support government programs across the board, the absence of this pattern for energy policy is consistent with a domain-specific thermostatic response partially offsetting the general pro-government tendency. Alternatively, these positive placebos may reflect unobserved canton-level confounders (e.g., progressive political culture) that our spatial design does not fully address; this underscores the importance of the same-language border specification and panel evidence as complementary identification strategies.

## 6.8 Summary of Results

Table 10 synthesizes the evidence across identification strategies, ordered from cleanest identification to broadest.

**Table 10:** Summary of Main Estimates Across Identification Strategies

Strategy	Estimate	SE	$p$ -value	Role
Same-language RDD	-1.60	1.18	0.17	<b>Primary</b>
Pooled RDD	-1.17	1.10	0.29	Upper bound
OLS + Language FE	-1.80	1.93	0.35	Benchmark
WCB $p$ -value (OLS)	—	—	0.42	Inference check
Stratified RI	—	—	0.53	Sharp null test
Panel DiD ( $D_{ct}$ )	-5.24	1.55	0.002	Temporal variation
CS ATT	-4.99	3.34	—	Heterogeneity-robust

Notes: All estimates are in percentage points. Same-language RDD is the preferred specification (cleanest identification). Panel DiD exploits temporal variation across four referendums. CS ATT uses the Callaway–Sant’Anna estimator excluding Basel-Stadt.

The cross-sectional evidence (OLS, spatial RDD) converges on a treatment effect of approximately  $-1.2$  to  $-1.8$  pp—negative but not statistically significant. The same-language border RDD rules out positive effects larger than 0.7 pp. The panel evidence is stronger: the DiD estimate of  $-5.2$  pp is significant at  $p = 0.002$ , and the CS ATT of  $-5.0$  pp is consistent in magnitude though imprecise. Heterogeneity analysis reveals that the negative effect is concentrated in rural municipalities, while urban areas show an offsetting positive effect.

No specification across any identification strategy produces evidence of positive policy feedback. The overall pattern is inconsistent with the prediction that sub-national policy experience builds support for federal action. Instead, the evidence points toward either a null effect or a modest negative effect—consistent with thermostatic preferences, cost salience, or both.

## 7. Discussion

### 7.1 Mechanisms

Why does cantonal energy law exposure fail to increase—and possibly decrease—support for federal energy policy? Several mechanisms could explain the null-to-negative findings, with the most theoretically grounded being the “thermostatic” response documented in political science. The heterogeneity results (Table 8) provide crucial evidence for distinguishing among them.

*Thermostatic Opinion Response.* The most compelling interpretation draws on [Wlezien \(1995\)](#)’s thermostatic model of public opinion. Wlezien shows that public preferences respond negatively to policy outputs: as government spending in a domain increases, public demand for *more* spending decreases, and vice versa. [Soroka & Wlezien \(2010\)](#) extend this framework across policy domains and federal systems, demonstrating that citizens adjust their preferences based on the policy status quo. Applied here, voters in treated cantons had already received “policy output” (cantonal energy laws); their demand for *additional* policy (federal harmonization) naturally declined. This is not policy failure but rather the public thermostat working as expected—citizens in treated cantons perceived that enough had been done. The thermostatic model thus transforms my null finding from a puzzle into a confirmation of a different theoretical prediction: policy feedback and thermostatic response are competing forces, and in this case, the latter dominated.

*Cost Salience and Local Backlash.* A complementary mechanism is that cantonal implementation made the costs of energy transition visible while benefits remained diffuse. [Stokes \(2016\)](#) documents precisely this dynamic for renewable energy policy: implementation generates “electoral backlash” as voters who bear concentrated costs (property owners facing retrofit requirements, residents near wind turbines) mobilize against further policy expansion, while diffuse beneficiaries remain politically quiescent. In the Swiss case, building owners who faced MuKE compliance costs learned exactly what energy transition entails. Building contractors dealt with new permitting requirements. The heterogeneity results provide direct support for this mechanism: the negative treatment effect is concentrated in *rural* municipalities (−1.7 pp baseline), where homeownership rates are higher and building retrofit costs more salient, while *urban* municipalities show a positive interaction (+5.0 pp,  $p = 0.004$ ). This urban-rural split is exactly what cost salience predicts: rural homeowners bear direct compliance costs, while urban renters are shielded and may instead experience positive spillovers from visible renewable energy infrastructure.

*Federal Overreach.* Swiss voters traditionally favor cantonal autonomy ([Vatter, 2018](#)). Voters in cantons that had already acted may have questioned why federal harmonization was necessary when cantonal solutions were working. [Becher & Stegmüller \(2021\)](#) show that federalism can “reduce the scope of conflict” by allowing heterogeneous preferences to be satisfied at the local level; federal action may be seen as unnecessary or even threatening to this arrangement. The Energy Strategy 2050 could be seen as unnecessary centralization in a policy domain where cantons had demonstrated both willingness and capacity to act.

*Partisan Sorting.* An alternative explanation is that treatment and preferences are both driven by an unobserved third variable—perhaps progressive political culture. Cantons with more environmental awareness adopted energy laws earlier *and* voted more favorably for

federal energy policy. But the direction of causality runs from preferences to treatment, not from treatment to preferences. In this case, the null finding would be correct: cantonal laws had no *causal* effect because the correlation reflects selection rather than feedback.

The evidence cannot definitively distinguish these mechanisms, but the heterogeneity results favor cost salience over pure thermostatic response. A pure thermostat would predict uniform negative effects regardless of community type; instead, the urban-rural split suggests that *direct exposure to costs* drives the negative effect while urban residents—who experience policy through different channels—may actually show positive feedback. The consistency of null-to-negative findings across identification strategies—including the spatial RDD that addresses selection concerns and the panel DiD that exploits temporal variation—suggests that policy feedback is genuinely absent at the aggregate level, masking offsetting effects across community types.

## 7.2 Limitations and Statistical Power

Several limitations deserve acknowledgment, with statistical power being the most important for interpreting the null finding.

*Canton-Level Language Assignment.* Language is assigned at the canton level (BFS majority classification) rather than at the Gemeinde level. This creates imprecision for bilingual cantons (FR, VS) and for cantons with linguistic minorities (French-speaking areas in BE; Italian/Romansh areas in GR). The “same-language borders” RDD similarly uses canton-level language classification, meaning some border segments between nominally German-speaking cantons may include locally French-speaking areas. Gemeinde-level language data (from census language shares) could provide finer resolution but would require additional data harmonization and introduce measurement error from survey responses. The canton-level approach follows standard practice in the Swiss referendum literature but represents a limitation for inference at borders where language changes within cantons.

*Power Analysis.* With a cluster-robust standard error of approximately 1.93 pp in the full-sample OLS, the minimum detectable effect (MDE) at 80% power is  $2.8 \times 1.93 \approx 5.4$  pp. This means I am well-powered to detect large effects—the kind of transformative policy feedback that would substantially shift referendum outcomes—but underpowered for modest effects of 2–3 pp. The 95% confidence interval for the preferred OLS estimate ( $[-5.6, 2.0]$ ) allows me to rule out positive effects larger than 2 pp with 97.5% confidence.

The Same-Language RDD specification (Table 5, row 2) has a standard error of 1.18 pp, yielding an MDE of approximately 3.3 pp. This is more precise than OLS because the spatial design exploits local variation at borders. The 95% CI  $[-3.9, 0.7]$  rules out positive effects larger than 0.7 pp—a stringent bound. The point estimate ( $-1.6$  pp) is slightly more

negative than the pooled RDD ( $-1.2$  pp), suggesting that removing language-confounded borders does not attenuate the estimate.

For context, the referendum passed with 58.2% support. If policy feedback were operating as the canonical theory predicts—creating constituencies, building support, generating momentum—I should observe positive effects. The same-language RDD rules out positive effects above 0.7 pp, providing informative evidence against the feedback hypothesis even without achieving significance. The panel DiD, which exploits temporal variation and yields the most precise estimate ( $-5.2$  pp,  $SE = 1.55$ ), provides the strongest evidence of a negative effect.

*Treatment Measurement.* Treatment is binary and measured at the canton level, but actual policy exposure varied within cantons. Some residents interacted directly with building regulations (homeowners, contractors); others had no contact. Individual-level survey data on policy awareness and implementation experience would allow more precise measurement and could identify which exposure mechanisms matter most.

*Spatial RDD Pooling.* The spatial RDD pools borders that differ in important ways. The Röstigraben borders (BE–FR, BE–JU) present identification challenges distinct from within-German borders (AG–ZH, BL–SO). The border-pair heterogeneity analysis (Appendix Figure 15) reveals substantial variation: GR–SG shows a large positive estimate ( $+13.1$  pp), while BE–LU is negative ( $-6.6$  pp). This heterogeneity suggests that the pooled estimate ( $-1.2$  pp) averages over meaningfully different local effects. The same-language border restriction mitigates some of this concern but does not eliminate within-German heterogeneity.

*External Validity.* Switzerland’s institutions—direct democracy, strong federalism, high trust in government—are unusual. Whether null policy feedback effects generalize to other federal systems (the United States, Germany, Australia) remains an open question. The thermostatic mechanism should operate wherever citizens can observe policy outputs and adjust preferences accordingly, but the specific Swiss context of referendum voting may amplify or attenuate these effects. Importantly, the null finding applies to *voter* preferences expressed through direct democracy; in representative systems, policy feedback may operate differently through interest group mobilization. Solar installers, heat pump manufacturers, and energy consultants who benefited from cantonal laws might still lobby effectively for federal expansion, even as voters themselves show thermostatic satiation. The Swiss referendum setting isolates pure voter response from interest group intermediation.

### 7.3 Policy Implications

These findings have implications for climate policy strategy. Advocates of decentralized climate policy often argue that state or provincial action will build momentum for national policy—creating constituencies, demonstrating feasibility, and shifting public opinion (Rabe, 2004). This paper suggests caution.

Successful sub-national implementation may not translate into federal support. Voters in jurisdictions that have already acted may perceive federal policy as redundant or overreaching. Implementation may make costs more salient than benefits. Strong local identities may generate resistance to federal encroachment.

This does not mean decentralized policy is unwise—cantonal energy laws presumably delivered direct benefits to those cantons regardless of federal adoption. But advocates should not assume that laboratory federalism automatically builds national coalitions. Complementary strategies may be necessary: federal co-financing, clear articulation of benefits from national coordination, and framing that respects local autonomy.

## 8. Conclusion

This paper tests whether sub-national policy experimentation generates political support for federal reform, using Switzerland’s cantonal energy laws as a natural experiment. The policy feedback hypothesis predicts that experience with local climate policy should build support for national action. I find no evidence for this prediction and suggestive evidence of the opposite.

The preferred spatial RDD at same-language borders yields  $-1.6$  pp (SE = 1.18,  $p = 0.17$ ), ruling out positive effects larger than 0.7 pp. The pooled border RDD gives  $-1.2$  pp ( $p = 0.29$ ); OLS with language controls gives  $-1.8$  pp ( $p = 0.35$ ; WCB  $p = 0.42$ ). Stratified randomization inference yields  $p = 0.53$ . Panel DiD with time-varying treatment coding provides the strongest evidence:  $-5.2$  pp (SE = 1.55,  $p = 0.002$ ). The Callaway–Sant’Anna heterogeneity-robust ATT is a consistent  $-5.0$  pp (SE = 3.34). No specification across any strategy produces evidence of positive feedback.

The null aggregate effect masks important heterogeneity. Rural municipalities show negative treatment effects ( $-1.7$  pp), while urban municipalities show a significant positive interaction ( $+5.0$  pp,  $p = 0.004$ ). This urban-rural split is consistent with cost salience: rural homeowners who directly face building retrofit costs show backlash, while urban renters—shielded from compliance costs—may experience positive spillovers. Placebo referendums on non-energy issues show positive discontinuities at the same borders, suggesting that treated-side municipalities generally favor federal initiatives—making the null energy result

domain-specific and consistent with a thermostatic response.

The analysis has important limitations. With 5 treated and 26 total cantons, cross-sectional precision is limited (MDE  $\approx$  3.3 pp for the same-language RDD). The panel DiD, while significant, relies on only four referendum periods and 25 cantons. Two of three placebo referendums show significant positive effects at the border, raising questions about whether the RDD fully isolates policy exposure from other canton-level differences. Future research should examine individual-level mechanisms through survey data on policy awareness, homeownership, and retrofit experience.

For policymakers, the implication is nuanced: do not assume that successful sub-national policy automatically builds support for national action. The relationship between local and federal policy preferences depends on who bears the costs of implementation. Building national coalitions for climate policy may require addressing the distributional consequences that laboratory federalism reveals.

## Acknowledgments

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

**Data and Code:** Replication materials are available at <https://github.com/SocialCatalystLab/ape-papers>

**Correspondence:** scl@econ.uzh.ch

## References

- Black, S. E. (1999). Do better schools matter? Parental valuation of elementary education. *Quarterly Journal of Economics*, 114(2), 577–599.
- Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6), 2295–2326.
- Cattaneo, M. D., Idrobo, N., & Titiunik, R. (2020). *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press.
- Dell, M. (2010). The persistent effects of Peru’s mining *mita*. *Econometrica*, 78(6), 1863–1903.
- Gelman, A., & Imbens, G. (2019). Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business & Economic Statistics*, 37(3), 447–456.
- Imbens, G. W., & Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2), 615–635.
- Keele, L. J., & Titiunik, R. (2015). Geographic boundaries as regression discontinuities. *Political Analysis*, 23(1), 127–155.
- Lee, D. S., & Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2), 281–355.
- Holmes, T. J. (1998). The effect of state policies on the location of manufacturing: Evidence from state borders. *Journal of Political Economy*, 106(4), 667–705.
- Dube, A., Lester, T. W., & Reich, M. (2010). Minimum wage effects across state borders: Estimates using contiguous counties. *Review of Economics and Statistics*, 92(4), 945–964.
- Imbens, G. W., & Kalyanaraman, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *Review of Economic Studies*, 79(3), 933–959.
- McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2), 698–714.
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427.
- Cameron, A. C., & Miller, D. L. (2015). A practitioner’s guide to cluster-robust inference. *Journal of Human Resources*, 50(2), 317–372.

- MacKinnon, J. G., & Webb, M. D. (2017). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, 32(2), 233–254.
- Conley, T. G., & 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.
- Ferman, B., & Pinto, C. (2019). Inference in differences-in-differences with few treated groups and heteroskedasticity. *Review of Economics and Statistics*, 101(3), 452–467.
- MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2019). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, forthcoming.
- Young, A. (2019). Channeling Fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results. *Quarterly Journal of Economics*, 134(2), 557–598.
- Campbell, A. L. (2003). *How Policies Make Citizens: Senior Political Activism and the American Welfare State*. Princeton University Press.
- Campbell, A. L. (2012). Policy makes mass politics. *Annual Review of Political Science*, 15, 333–351.
- Mettler, S. (2002). Bringing the state back in to civic engagement: Policy feedback effects of the G.I. Bill for World War II veterans. *American Political Science Review*, 96(2), 351–365.
- Mettler, S. (2011). *The Submerged State: How Invisible Government Policies Undermine American Democracy*. University of Chicago Press.
- Mettler, S., & SoRelle, M. (2014). Policy feedback theory. In P. A. Sabatier & C. M. Weible (Eds.), *Theories of the Policy Process* (3rd ed., pp. 151–181). Westview Press.
- Pierson, P. (1993). When effect becomes cause: Policy feedback and political change. *World Politics*, 45(4), 595–628.
- Soss, J. (1999). Lessons of welfare: Policy design, political learning, and political action. *American Political Science Review*, 93(2), 363–380.
- Karch, A. (2007). *Democratic Laboratories: Policy Diffusion among the American States*. University of Michigan Press.
- Oates, W. E. (1999). An essay on fiscal federalism. *Journal of Economic Literature*, 37(3), 1120–1149.

- Rose, R. (1993). *Lesson-Drawing in Public Policy: A Guide to Learning Across Time and Space*. Chatham House.
- Shipan, C. R., & Volden, C. (2008). The mechanisms of policy diffusion. *American Journal of Political Science*, 52(4), 840–857.
- Herrmann, M., & Armingeon, K. (2010). The distinctive politics of Swiss direct democracy. In H. Kriesi (Ed.), *Referendum Voting* (pp. 167–192). Campus.
- Kriesi, H. (2005). *Direct Democratic Choice: The Swiss Experience*. Lexington Books.
- Linder, W., & Vatter, A. (2010). *Swiss Democracy: Possible Solutions to Conflict in Multicultural Societies*. Palgrave Macmillan.
- Rinscheid, A. (2015). Crisis, policy discourse, and major policy change: Exploring the role of subsystem polarization in nuclear energy policymaking. *European Policy Analysis*, 1(2), 34–70.
- Sager, F. (2014). The political economy of energy market liberalization in Switzerland. In F. Gilardi & A. Rasmussen (Eds.), *Handbook of Public Policy* (pp. 213–234). Edward Elgar.
- Swissvotes. (2017). Energiegesetz (EnG): Volksabstimmung vom 21. Mai 2017. <https://swissvotes.ch/vote/612>
- swissdd R package. (2020). Swiss Direct Democracy Data. <https://github.com/zumbov2/swissdd>
- Neue Zürcher Zeitung. (2017). Energiestrategie 2050: Die Argumente im Überblick. May 15, 2017.
- Trechsel, A., & Sciarini, P. (2000). Direct democracy in Switzerland: Do elites matter? *European Journal of Political Research*, 33(1), 99–124.
- Vatter, A. (2018). *Swiss Federalism: The Transformation of a Federal Model*. Routledge.
- Carattini, S., Baranzini, A., Thalmann, P., Varone, F., & Vöhringer, F. (2018). Green taxes in a post-Paris world: Are millions of nays inevitable? *Environmental and Resource Economics*, 68(1), 97–128.
- Drews, S., & van den Bergh, J. C. (2016). What explains public support for climate policies? A review of empirical and experimental studies. *Climate Policy*, 16(7), 855–876.

- Kallbekken, S., & Sælen, H. (2011). Public acceptance for environmental taxes: Self-interest, environmental and distributional concerns. *Energy Policy*, 39(5), 2966–2973.
- Rabe, B. G. (2004). *Statehouse and Greenhouse: The Emerging Politics of American Climate Change Policy*. Brookings Institution Press.
- Stoutenborough, J. W., Bromley-Trujillo, R., & Vedlitz, A. (2014). Public support for climate change policy: Consistency in the influence of values and attitudes over time and across specific policy alternatives. *Review of Policy Research*, 31(6), 555–583.
- Callaway, B., & Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.
- Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5), 2555–2591.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.
- de Chaisemartin, C., & D’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.
- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.
- Wlezien, C. (1995). The public as thermostat: Dynamics of preferences for spending. *American Journal of Political Science*, 39(4), 981–1000.
- Soroka, S. N., & Wlezien, C. (2010). *Degrees of Democracy: Politics, Public Opinion, and Policy*. Cambridge University Press.
- Stokes, L. C. (2016). Electoral backlash against climate policy: A natural experiment on retrospective voting and local resistance to public goods. *American Journal of Political Science*, 60(4), 958–974.
- Becher, M., & Stegmueller, D. (2021). Reducing the scope of conflict: Federalism, nationalism, and redistribution. *American Journal of Political Science*, 65(1), 23–39.

## A. Data Appendix

### A.1 Referendum Data Sources

Canton-level and Gemeinde-level referendum results are from the Federal Statistical Office (BFS), accessed via the VoteInfo JSON API (<https://ogd-static.voteinfo-app.ch>). I use results for:

- May 21, 2017: Energy Strategy 2050 (Energiegesetz, Vorlagen Nr. 612)
- November 27, 2016: Nuclear Phase-Out (Atomausstiegsinitiative, Nr. 608)
- May 18, 2003: Nuclear Moratorium Extension (Nr. 499)
- September 24, 2000: Energy Levy (Energienenkungsabgabe, Nr. 466)

### A.2 Treatment Definition and Verification

**Treatment criterion:** A canton is coded as “treated” if it adopted a comprehensive cantonal energy law implementing MuKE<sub>n</sub> (Model Cantonal Energy Provisions) standards with enforcement provisions in force by May 21, 2017. “Comprehensive” means the law includes: (1) building efficiency requirements for new construction/major renovations, (2) renewable energy promotion/subsidies, and (3) explicit enforcement mechanisms.

**Treated canton verification** (LexFind, <https://www.lexfind.ch>):

- GR: Energiegesetz des Kantons Graubünden, SR 820.200 (in force January 2011)
- BE: Kantonales Energiegesetz, SR 741.1 (in force January 2012)
- AG: Energiegesetz des Kantons Aargau, SR 773.200 (in force January 2013)
- BL: Energiegesetz, SR 490 (in force July 2016)
- BS: Energiegesetz, SR 772.100 (in force January 2017)

**Treatment timing summary:** Table 11 clarifies the distinction between adoption (passage) and in-force dates. All treatment coding throughout this paper uses **in-force dates**.

**Table 11:** Treatment Timing: Adoption vs. In-Force Dates

Canton	Abbr.	Adoption Year	In-Force Date	Coded Cohort
Graubünden	GR	2010	January 2011	2011
Bern	BE	2011	January 2012	2012
Aargau	AG	2012	January 2013	2013
Basel-Landschaft	BL	2015	July 2016	2016
Basel-Stadt	BS	2016	January 2017	2017

Notes: Adoption year = year the cantonal parliament passed the law. In-force date = when the law took legal effect. All treatment indicators and cohort definitions use in-force dates. Figure 3 may display adoption years in its legend; this table provides the definitive treatment coding.

### A.3 Full Canton Results

**Table 12:** Full Canton-Level Results: Energy Strategy 2050 Referendum

Canton	Abbr.	Yes (%)	Turnout (%)	Language	Status
Zürich	ZH	62.3	44.5	German	Control
Bern	BE	62.5	41.7	German	Treated (2012)
Luzern	LU	52.1	40.8	German	Control
Uri	UR	38.2	40.1	German	Control
Schwyz	SZ	43.5	42.7	German	Control
Obwalden	OW	42.8	39.5	German	Control
Nidwalden	NW	47.3	41.2	German	Control
Glarus	GL	47.9	38.4	German	Control
Zug	ZG	55.8	44.1	German	Control
Fribourg	FR	61.4	39.8	French	Control
Solothurn	SO	57.2	41.6	German	Control
Basel-Stadt	BS	72.8	47.2	German	Treated (2017)
Basel-Landschaft	BL	61.2	45.1	German	Treated (2016)
Schaffhausen	SH	54.6	44.8	German	Control
Appenzell A.-Rh.	AR	52.3	41.9	German	Control
Appenzell I.-Rh.	AI	42.1	45.3	German	Control
St. Gallen	SG	52.8	42.2	German	Control
Graubünden	GR	55.4	43.8	German	Treated (2011)
Aargau	AG	54.8	42.3	German	Treated (2013)
Thurgau	TG	51.4	43.7	German	Control
Ticino	TI	58.7	37.2	Italian	Control
Vaud	VD	67.4	38.9	French	Control
Valais	VS	53.1	39.4	French	Control
Neuchâtel	NE	68.2	37.8	French	Control
Genève	GE	71.5	38.1	French	Control
Jura	JU	61.8	40.2	French	Control
<b>Switzerland</b>		<b>58.2</b>	<b>42.3</b>		

## B. Spatial RDD Implementation Details

### B.1 Distance Calculation

For each Gemeinde, I calculate the signed distance to the nearest treated-control canton border as follows:

1. Obtain Gemeinde boundary polygons from swisstopo SwissBOUNDARIES3D via the BFS R package.
2. Compute the centroid of each Gemeinde polygon.
3. Identify the boundary between the union of treated canton polygons and the union of control canton polygons.
4. Calculate Euclidean distance from each centroid to this boundary (in meters).
5. Sign the distance: positive for Gemeinden in treated cantons, negative for those in control cantons.
6. Convert to kilometers.

### B.2 Border Pairs

The treated-control canton borders include:

- **Same-language (German–German):** AG–ZH, AG–SO, AG–LU, AG–ZG, BL–SO, GR–SG, GR–GL, GR–UR, BE–LU, BE–SO, BE–OW, BE–NW, BE–UR
- **Cross-language (German–French/Italian):** BE–FR, BE–JU, BE–NE, BE–VD, BE–VS, BL–JU, GR–TI

The cross-language borders coincide with the Röstigraben, creating a confounded RDD. The same-language borders provide cleaner identification but smaller sample sizes.

## C. Robustness Checks

### C.1 Alternative OLS Specifications

**Table 13:** Robustness: Alternative OLS Specifications

Specification	Estimate	SE	N	Notes
German-speaking only	-1.80	1.94	1,354	German cantons only
Exclude Basel-Stadt	-2.16	2.07	2,030	Urban outlier
Population weighted	-1.04	2.64	2,120	Weights by eligible voters
Rural only (<5,000 voters)	-1.68	2.00	1,897	Excludes cities
Urban only ( $\geq 5,000$ voters)	+3.37***	1.30	223	Cities only

Notes: All specifications include language controls except “German-speaking only,” which restricts to German-speaking cantons where the language confound is absent. The “Urban only” specification shows a significant *positive* effect (+3.4 pp), consistent with the heterogeneity analysis (Table 8): urban municipalities in treated cantons voted more favorably. Standard errors clustered by canton. \*\*\* $p < 0.01$ .

### C.2 Donut RDD

Excluding Gemeinden within specified distances of the border tests whether results are driven by immediate border spillovers:

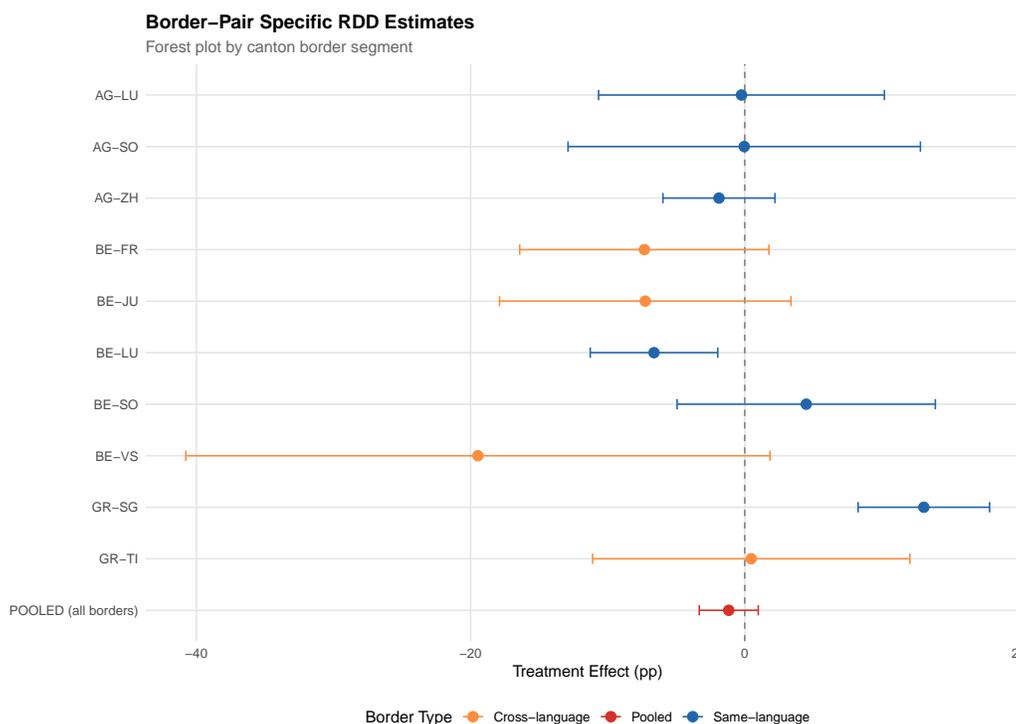
**Table 14:** Donut RDD Specifications

Donut (km)	Estimate	SE	95% CI	N
0 (baseline)	-4.33***	1.31	[-6.9, -1.8]	1,698
0.5	-3.35**	1.54	[-6.4, -0.3]	1,668
1.0	-4.82***	1.79	[-8.3, -1.3]	1,632
1.5	-7.79***	1.90	[-11.5, -4.1]	1,581
2.0	-6.86***	2.56	[-11.9, -1.8]	1,544

Notes: Excludes Gemeinden within specified distance of the border. All donut specifications yield significant negative estimates, with larger effects at wider exclusion radii. This pattern suggests that cross-border spillovers in the immediate border zone *attenuate* the treatment effect. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ .

### C.3 Border-Pair Heterogeneity

To examine whether the null result is driven by a particular border segment, I estimate separate RDD specifications for each major border pair. Figure 15 presents a forest plot of these border-pair-specific estimates alongside the pooled estimate.



**Figure 15:** Border-Pair Specific RDD Estimates

Notes: Forest plot of RDD estimates by canton border segment. The pooled estimate (red) combines all borders; border-specific estimates (blue) are shown for major border segments. Estimates vary substantially across segments: GR-SG is strongly positive, BE-LU and cross-language borders are negative. 95% confidence intervals shown.

The forest plot reveals substantial heterogeneity across border segments. Same-language borders show mixed estimates: AG-ZH (-1.9 pp), BE-LU (-6.6 pp), BE-SO (+4.5 pp), and GR-SG (+13.1 pp). Cross-language borders show consistently negative estimates, driven by the Röstigraben confound (BE-FR: -7.3 pp, BE-VS: -19.5 pp). All estimates have wide confidence intervals due to small within-segment samples.

The large positive GR-SG estimate (+13.1 pp) merits discussion. Graubünden was the earliest adopter (in force January 2011), giving it the longest exposure—over six years by the 2017 referendum. Its Alpine tourism economy may have experienced positive spillovers from

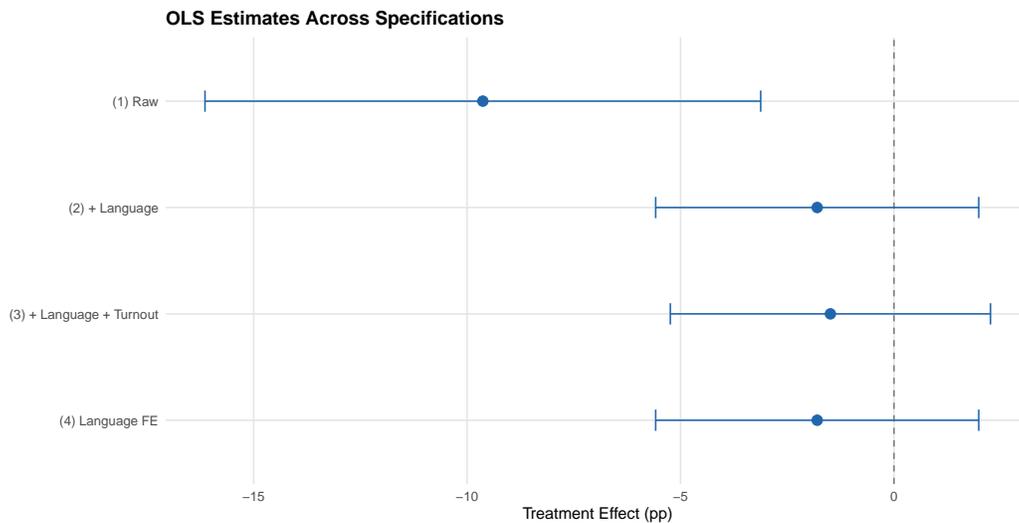
environmental branding: municipalities dependent on winter sports and nature tourism could view energy transition as complementary to their economic model rather than a cost burden. The canton’s geographic isolation and strong local identity may also reduce the “federal overreach” objection that operates elsewhere. Importantly, GR–SG is the only same-language border pair with a clearly positive estimate, suggesting that early adoption combined with specific economic conditions may generate the positive feedback that other treated cantons lack. However, the wide confidence interval ( $\approx[-5, +31]$ ) means this segment-level estimate is highly uncertain.

The border-pair forest plot (Figure 15) provides the visual summary of this heterogeneity, with individual segment estimates shown alongside the pooled result.

## D. Additional Appendix Materials

### D.1 OLS Specification Comparison

Figure 16 presents a coefficient plot comparing the treatment effect estimate across all OLS specifications. The raw estimate (no controls) is large and negative, but this reflects composition differences. Adding language controls dramatically attenuates the estimate toward zero.

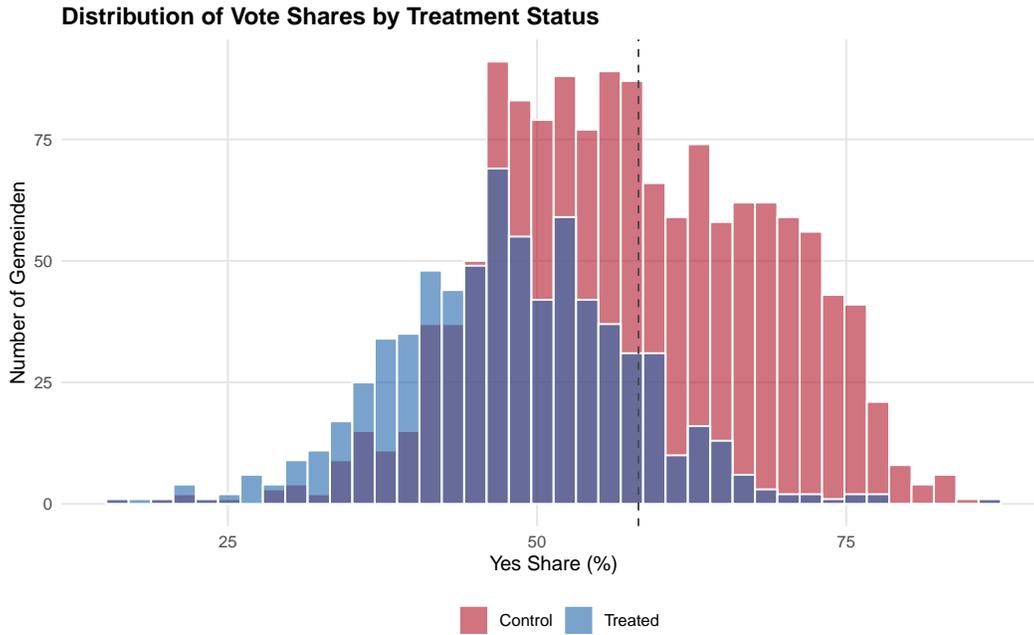


**Figure 16:** OLS Coefficient Plot: Treatment Effect Across Specifications

Notes: Point estimates and 95% confidence intervals for the treatment effect across OLS specifications. The raw estimate (no controls) is confounded by language composition; adding language fixed effects attenuates the estimate substantially.

## D.2 Vote Share Distributions

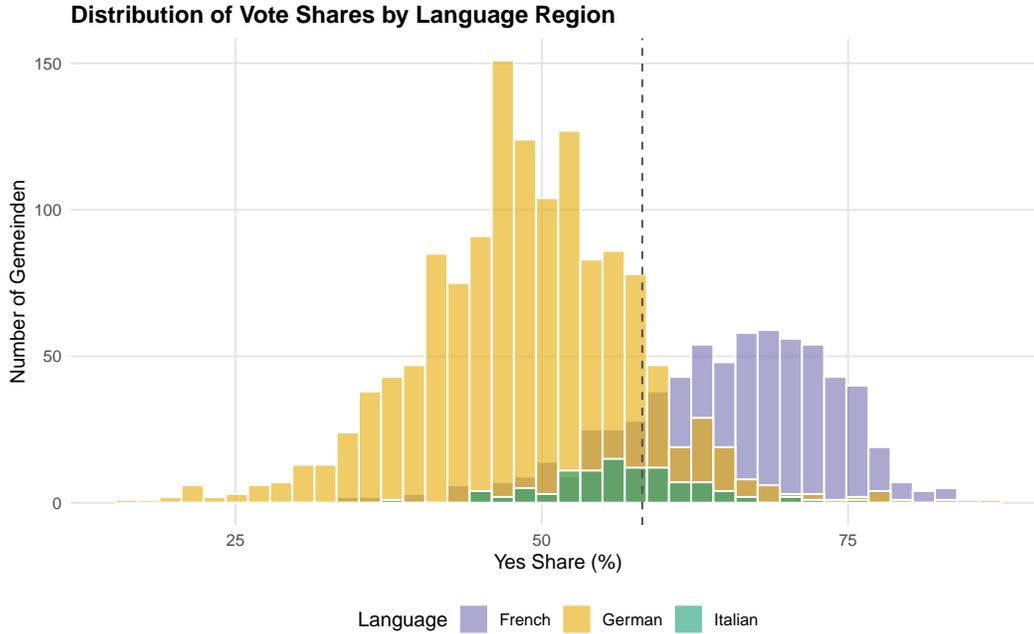
Figure 17 shows the distribution of Gemeinde-level yes-shares by treatment status. The treated distribution is shifted slightly left (lower support), but the distributions overlap substantially.



**Figure 17:** Distribution of Vote Shares by Treatment Status

Notes: Kernel density estimates of Gemeinde-level yes-shares. Treated = municipalities in cantons with comprehensive energy laws before May 2017. Control = all other municipalities.

Figure 18 shows the distribution by language region, highlighting the Röstigraben divide: French-speaking Gemeinden vote much more favorably than German-speaking ones, regardless of treatment status.



**Figure 18:** Distribution of Vote Shares by Language Region

Notes: Kernel density estimates of Gemeinde-level yes-shares by primary language. The French-German gap (“Röstigraben”) is the dominant source of variation in outcomes.

### D.3 Heterogeneity by Urbanity

Table 8 in Section 6 presents the main heterogeneity results. The urban interaction effect is +5.0 pp (SE = 1.72,  $p = 0.004$ ), indicating that urban municipalities in treated cantons voted significantly *more* favorably than their control-canton counterparts. The rural baseline treatment effect is -1.7 pp (SE = 2.0), negative but imprecise. This urban-rural split provides the strongest evidence for cost-salience as the mechanism behind the null aggregate result: rural homeowners who face direct building retrofit costs show negative feedback, while urban residents—more likely to be renters—do not.

### D.4 Power Analysis

Table 15 presents the statistical power analysis for the preferred specification.

**Table 15:** Power Analysis: Minimum Detectable Effects

Parameter	OLS (Language FE)	Same-Language RDD
Point estimate	-1.80 pp	-1.60 pp
Standard error	1.93 pp	1.18 pp
MDE at 80% power	5.41 pp	3.30 pp
95% CI lower bound	-5.58 pp	-3.91 pp
95% CI upper bound	+1.99 pp	+0.71 pp
<i>Can rule out positive effects larger than:</i>		
	+1.99 pp	+0.71 pp

Notes: OLS based on  $N = 2,120$  Gemeinden with 26 canton clusters. Same-language RDD uses only German–German border segments (MSE-optimal bandwidth = 9.1 km). MDE = minimum detectable effect at 80% power with  $\alpha = 0.05$ . The same-language RDD provides a tighter bound on positive effects (0.71 pp) despite smaller sample size, because the spatial design reduces residual variance.

## D.5 Callaway-Sant’Anna Detailed Results

Table 16 presents the group-time average treatment effects from the Callaway-Sant’Anna estimator.

**Table 16:** Callaway-Sant’Anna Group-Time ATTs

Cohort	Period	ATT	SE	95% CI
2011 (GR)	2016	+2.56	1.56	[−0.49, 5.61]
2011 (GR)	2017	+4.26	1.44	[1.43, 7.09]
2012 (BE)	2016	−8.22	1.56	[−11.27, −5.16]
2012 (BE)	2017	−9.02	1.44	[−11.85, −6.19]
2013 (AG)	2016	−9.59	1.56	[−12.64, −6.54]
2013 (AG)	2017	−10.94	1.44	[−13.77, −8.11]
2016 (BL)	2016	+0.30	1.56	[−2.76, 3.35]
2016 (BL)	2017	−9.29	1.44	[−12.12, −6.46]
<b>Aggregate ATT</b>		<b>−4.99</b>	<b>3.34</b>	<b>[−11.54, 1.56]</b>

Notes: Group-time average treatment effects using the [Callaway & Sant’Anna \(2021\)](#) estimator with in-force year cohort timing.  $N = 25$  cantons  $\times$  4 referendum periods = 100 canton-period observations. Basel-Stadt (2017 cohort) excluded because its first post-treatment period is the final period. The 2011 (GR) cohort shows positive ATTs, while 2012 (BE) and 2013 (AG) cohorts show large negative ATTs, suggesting substantial cohort heterogeneity. The aggregate ATT (−5.0, SE = 3.34) is consistent with the TWFE estimate (−5.2) but has a wider confidence interval. Standard errors clustered by canton.

## D.6 Randomization Inference Details

Table 17 provides detailed results from the randomization inference procedure.

**Table 17:** Randomization Inference Results

Parameter	Stratified	Unstratified
Observed estimate	-1.80 pp	-1.80 pp
Permutation pool	5 of 17 German	5 of 26 all
Number of permutations	1,000	1,000
Permutation mean	+0.23 pp	+0.18 pp
Permutation SD	2.76 pp	2.96 pp
Two-tailed $p$ -value	0.531	0.576

Notes: Randomization inference under the sharp null of no treatment effect for any unit. Stratified: treatment is reassigned to 5 of 17 German-speaking cantons, conditioning on the fact that all treated cantons are German-speaking. Unstratified: treatment is reassigned to 5 of 26 cantons. Both procedures yield similar  $p$ -values; the observed estimate lies well within both permutation distributions.