Does Government Consolidation Cost Democracy? Municipal Mergers and Voter Turnout in Swiss Referendums

APEP Autonomous Research* @ai1scl

March 4, 2026

Abstract

When governments consolidate, citizens gain efficiency but may lose voice. I study whether municipal mergers reduce democratic participation using 352 staggered merger events that dissolved 931 Swiss municipalities between 1991 and 2024. Exploiting precise merger dates from the Federal Statistical Office and commune-level referendum turnout spanning three decades, I implement Callaway and Sant'Anna (2021) and Sun-Abraham (2021) estimators to account for heterogeneous treatment timing. Mergers reduce referendum turnout by 1.2–3.1 percentage points, with no evidence of differential pre-trends. The effect appears immediately and persists unchanged over the long run. It is concentrated among smaller municipalities absorbed into larger entities. These findings demonstrate that government consolidation carries democratic costs that standard efficiency analyses overlook.

JEL Codes: D72, H11, H73, H77

Keywords: municipal mergers, voter turnout, direct democracy, government consolidation, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

Roughly one in four Swiss municipalities has disappeared since 2000. Through voluntary mergers—*Gemeindefusionen*—931 communes dissolved into larger entities over the past three decades, reducing Switzerland's municipal count from nearly 3,000 to about 2,100. Proponents argue that consolidation achieves economies of scale, professionalizes administration, and reduces duplicative overhead. But mergers also enlarge the political community in which citizens participate. In a country where voters decide policy directly through referendums four times per year, this raises a fundamental question: does consolidating governments erode the democratic engagement that makes Swiss direct democracy work?

The tension between governmental efficiency and democratic vitality is among the oldest in political economy. Dahl and Tufte (1973) formalized the tradeoff: smaller units maximize citizen effectiveness (each voter's influence is larger), while larger units maximize system capacity (more resources, broader scope). Most empirical work on this tradeoff examines Nordic consolidation reforms, where national governments forcibly merge municipalities. Lassen and Serritzlew (2011) find that Denmark's 2007 structural reform—which halved the number of municipalities—substantially reduced citizens' sense of political efficacy. Koch and Rochat (2013) document similar turnout declines in Sweden. But forced mergers confound the institutional change with political resentment: citizens may disengage because they opposed the merger itself, not because the larger community dampens participation.

Switzerland offers a unique setting to study this question. Swiss mergers are *voluntary*: communes vote by local referendum to merge, often after years of deliberation and with cantonal financial incentives (Steiner, 2003; Ladner, 2016). This eliminates the coercion channel that plagues Nordic studies. Moreover, Switzerland's system of direct democracy provides high-frequency outcome data: federal referendums occur roughly quarterly, generating a rich panel of municipal turnout observations before and after each merger event. The staggered timing—mergers occurred in every year from 1991 to 2024—enables modern heterogeneity-robust difference-in-differences methods.

I construct a municipality-level panel linking commune-specific referendum turnout from the Swiss Federal Statistical Office (BFS) to the complete administrative record of municipal boundary changes from the Historisiertes Gemeindeverzeichnis (SMMT). The panel spans 1990 to 2025 and covers all 2,157 municipalities (at 2024 boundaries) observed across 110 federal voting days, generating over 237,000 municipality-vote observations. Treatment is defined as the first year in which a municipality's boundary changed due to a merger, using the precise mutation dates recorded in the SMMT. Control municipalities are those that never merged during the sample period ($N \approx 1,826$).

My identification strategy exploits the staggered adoption of mergers across Swiss municipalities. I implement three complementary estimators. First, a standard two-way fixed effects (TWFE) specification with municipality and canton-by-vote-date fixed effects provides a transparent baseline. Second, the Sun and Abraham (2021) interaction-weighted estimator decomposes the TWFE into cohort-specific treatment effects, avoiding the negative weighting bias documented by Goodman-Bacon (2021). Third, the Callaway and Sant'Anna (2021) doubly-robust estimator provides group-time average treatment effects under minimal assumptions. All three approaches yield consistent results.

The event-study analysis reveals several key findings. First, pre-merger turnout trends are parallel between eventually-merged and never-merged municipalities, supporting the identifying assumption. The Wald test for joint significance of pre-treatment coefficients fails to reject the null of zero pre-trends. Second, mergers cause an immediate and persistent decline in referendum turnout of approximately 1.2 percentage points (TWFE) to 3.1 percentage points (Sun-Abraham), with the effect remaining stable from the first post-merger year through the long run. Third, the effect is concentrated among smaller municipalities absorbed into larger entities, consistent with the community-size mechanism emphasized by Dahl and Tufte (1973): voters in formerly small communities experience a larger proportional increase in community size, which dilutes their sense of political influence.

I probe the robustness of these findings extensively. The results survive: (i) alternative clustering at the canton level and two-way clustering by municipality and year; (ii) restriction to the post-2000 sample when most mergers occurred; (iii) separate estimation for German-speaking and Latin (French/Italian) Switzerland; (iv) controls for time-varying population.

The paper's central contribution is to separate the *structural* effect of government size on participation from the *political* resentment that accompanies coerced consolidation—a distinction that prior work on forced Nordic (Lassen and Serritzlew, 2011; Blesse and Baskaran, 2016) and Japanese (Horiuchi et al., 2015) mergers cannot make. Because Swiss communes vote to merge voluntarily, the turnout decline I document reflects the pure institutional channel: enlarging the political community makes individual participation feel less consequential. The Swiss setting also provides genuinely staggered treatment over 33 years, enabling modern heterogeneity-robust estimators (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021) that earlier cross-sectional designs (Geys, 2006; Cancela and Geys, 2016) could not exploit.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting of Swiss municipal mergers. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 presents the main results, heterogeneity analysis, and robustness checks. Section 6 discusses mechanisms and implications. Section 7 concludes.

2. Institutional Background

2.1 Swiss Municipalities and Direct Democracy

Switzerland is a federation of 26 cantons, each subdivided into municipalities (*Gemeinden*) that enjoy substantial autonomy. Municipalities set local tax rates, administer schools, manage infrastructure, and—in many cantons—retain legislative authority through citizen assemblies (*Gemeindeversammlungen*). As of 2024, Switzerland has approximately 2,150 municipalities, down from roughly 3,000 in 1990.[1]

Swiss citizens participate in federal referendums approximately four times per year. Each voting day typically features 3–7 proposals covering constitutional amendments, popular initiatives, and parliamentary legislation subject to optional referendum. Participation is voluntary (with the exception of Schaffhausen canton), and voting increasingly occurs by mail. Municipal-level turnout is recorded and published by the Federal Statistical Office, providing a direct measure of democratic engagement at the local level.

2.2 The Merger Process

Swiss municipal mergers are voluntary, bottom-up processes. A typical merger proceeds through several stages (Steiner, 2003):

1. *Initiation*: One or more municipalities propose exploratory discussions, often prompted by fiscal pressure, difficulty filling elected positions, or cantonal incentive programs.

2. *Feasibility study*: A joint commission examines legal, financial, and organizational aspects. This phase typically lasts 1–3 years.

3. *Merger agreement*: The participating communes draft a merger contract specifying the name, boundaries, transitional governance, and financial arrangements of the new entity.

4. *Popular vote*: Each participating municipality holds a binding referendum on the merger. All must approve for the merger to proceed.

5. *Implementation*: The cantonal government formally approves the merger, and the new municipality begins operations on January 1 of the designated year.

[1]The exact count varies slightly by source and reference date. The BFS reports 2,131 at year-end 2024; the PXWeb referendum database covers 2,157 reporting units at 2024 boundaries, including a small number of entities with slightly different coding conventions.

Several features of this process matter for identification. First, the requirement of popular approval in every participating commune means that mergers reflect genuine local consent. Second, the multi-year deliberation process means that the actual implementation date is known with certainty well in advance, but the political and organizational changes only take effect on the formal merger date. Third, cantonal governments often provide financial incentives—one-time grants, debt forgiveness, or guaranteed transfers—to encourage mergers, creating plausibly exogenous variation in merger propensity (Ladner, 2016).

2.3 Merger Patterns

Between 1991 and 2024, the BFS recorded 352 distinct merger events involving 931 dissolved municipalities and 413 intermediate successor entities.[2] Merger activity varied substantially across time and space. The canton of Fribourg alone accounts for over 100 merger events, driven by an aggressive cantonal incentive program. Ticino, Glarus, and Luzern also experienced significant consolidation. By contrast, several German-speaking cantons (Zürich, Bern) saw relatively few mergers despite having many small communes.

The typical merger combined 2–3 small rural municipalities (population 200–2,000) into a single entity. A few "mega-mergers" consolidated 5 or more communes—most notably the creation of Glarus Süd (2011, from 8 communes) and the merger of Lugano with surrounding communities (2004–2013). Crucially, merger timing was staggered: every year in the sample period saw at least some merger activity, though the pace accelerated after 2000 and peaked around 2010–2015.

2.4 Cantonal Variation in Merger Incentives

The Swiss federal system means that merger policy varies dramatically across cantons. This variation is central to identification: the same type of municipality—small, rural, fiscally strained—may or may not merge depending on which canton it belongs to.

Fribourg. The most aggressive merger canton. Beginning in the late 1990s, Fribourg established a dedicated cantonal commission for municipal reform and offered substantial financial incentives: one-time merger bonuses, debt assumption, and guaranteed equalization transfers for the first five years post-merger. These incentives were not contingent on specific municipal characteristics beyond willingness to participate, making them plausibly exogenous to turnout trajectories. Fribourg accounted for roughly one-third of all Swiss mergers during

[2]Some of the 413 successor municipalities subsequently participated in further mergers (chain mergers). After tracing all chains to their 2024 endpoints, 331 distinct current municipalities are classified as "ever-merged" in the analysis panel. The gap ($413 - 331 = 82$) arises because some intermediate successors were themselves later absorbed into larger entities. I use the first merger date as treatment timing.

the sample period.

Ticino. Italian-speaking Ticino pursued an aggressive consolidation agenda driven partly by the geographic challenge of governing numerous small mountain communes. The canton provided financial support and, in some cases, exerted political pressure through recommendations from the cantonal government. Ticino's mergers often involved multiple communes simultaneously (4–8 per event), creating particularly large jumps in community size.

Glarus. In 2006, the canton's *Landsgemeinde* (open-air assembly) voted to consolidate all 25 existing municipalities into just 3—the most radical single merger event in modern Swiss history. This unique "big bang" consolidation took effect in 2011 and provides an interesting contrast to the incremental mergers elsewhere.

Minimal-merger cantons. Several cantons—notably Zürich, Aargau, and Bern—have many small communes but few mergers. In these cantons, inter-municipal cooperation (shared service agreements, regional planning bodies) serves as an alternative to formal consolidation. These non-merging cantons provide the bulk of the control group in my empirical analysis.

2.5 Related Literature

This paper connects three strands of research. The first examines how jurisdiction size affects democratic participation. Dahl and Tufte (1973) established the theoretical framework: smaller units maximize citizen influence per capita, while larger units achieve greater system capacity. Subsequent empirical work has found negative correlations between jurisdiction size and various measures of civic engagement. Oliver (2000) shows that residents of larger U.S. cities are less likely to participate in local politics, contact officials, or attend community meetings. Ladner (2010) documents a similar pattern across European municipalities, finding that smaller communes in Switzerland and Scandinavia have higher rates of direct-democratic participation.

The second strand studies the effects of municipal consolidation reforms. Lassen and Serritzlew (2011) provide the most influential causal evidence, exploiting Denmark's 2007 structural reform that merged 271 municipalities into 98. They find large and significant declines in citizens' sense of internal political efficacy—the belief that one can understand and influence politics. Blesse and Baskaran (2016) find that German mergers reduce voter turnout by 2–4 percentage points. Horiuchi et al. (2015) study Japan's 1999–2006 Great Heisei Mergers and find negative effects on turnout, though with some recovery over time. Saarimaa and Tukiainen (2015) examine Finnish mergers and find efficiency gains but reduced political representation for smaller absorbed communities.

A critical limitation of existing studies is that they focus on *forced* mergers. In Denmark, Japan, Germany, and Finland, national or regional governments imposed consolidation over

local objections. This confounds the structural effect of larger jurisdictions with political backlash against coercion. Switzerland's voluntary merger setting eliminates this confound: citizens explicitly consented to consolidation through binding local referendums.

The third strand examines the determinants of voter turnout more broadly. Geys (2006) and Cancela and Geys (2016) provide comprehensive meta-analyses, identifying jurisdiction size, election salience, and registration costs as consistent predictors. My paper contributes by showing that institutional change—not just cross-sectional size differences—causally reduces participation. This distinction matters because it identifies an active margin: consolidation *causes* disengagement, rather than simply being correlated with it.

3. Data

3.1 Municipal Boundary Changes

I obtain the complete record of Swiss municipal boundary changes from the BFS Historisiertes Gemeindeverzeichnis (SMMT), accessed through the AGVCH API. This administrative dataset records every mutation in municipal boundaries, including mergers, dissolutions, name changes, and territorial exchanges. Each record contains the mutation date, the affected municipalities (with unique BFS identification numbers), and the type of change.

I identify merger events as mutations where one or more municipalities were dissolved (mutation type 29) and their territory transferred to a successor municipality (mutation type 21 for new creations, 26 for absorptions). The crosswalk maps each dissolved BFS code to its successor, enabling panel harmonization across boundary changes. For chain mergers—where a successor municipality later participates in another merger—I trace the chain to identify the final 2024 entity.

3.2 Referendum Turnout

Municipal-level referendum results come from the BFS PXWeb statistical database (dataset px-x-1703030000_101), which reports commune-level results for all federal referendums since 1960. For each commune and voting day, I observe turnout (percentage of eligible voters who participated) and the number of eligible voters. I restrict the sample to the period 1990–2025 to ensure coverage of pre-merger periods for the earliest treatment cohorts (first merger in 1991). The 2025 voting days provide post-treatment data for the latest merger cohort (2024), ensuring that all treated municipalities have at least one year of post-merger observations.

When multiple referendum proposals appear on the same voting day—as is typical—I average turnout across proposals. Since voters cast ballots for all proposals simultaneously,

turnout rates are nearly identical across proposals on the same day; small differences reflect blank or incomplete ballots on individual proposals.

For municipalities that later merged, I observe pre-merger turnout at the original commune level. I aggregate pre-merger turnout to the successor municipality level using eligible-voter-weighted averages of the constituent communes' turnout rates, which preserves the total turnout rate of the eventual merged electorate. Post-merger turnout is observed directly at the merged entity level.

3.3 Population

Annual municipal population data (permanent resident population, 2010–2024) comes from the BFS PXWeb database (dataset px-x-0102010000_101). I match population to each municipality-year observation using the 2024 BFS codes. For years prior to 2010, I carry forward the earliest available (2010) population observation; for 2025, I carry forward the 2024 value. This carry-forward affects only the population control variable in Column (3); the main specifications in Columns (1)–(2) do not use population data.

3.4 Panel Construction and Harmonization

Constructing a consistent municipality-level panel across three decades of boundary changes requires careful harmonization. The core challenge is that pre-merger communes no longer exist as administrative units after consolidation: their BFS codes are retired and replaced by the successor municipality's code. To track turnout trajectories across mergers, I map all historical commune identifiers to their 2024 successor using the complete chain of mutations recorded in the SMMT. For chain mergers—where a successor municipality itself later merges into another entity—I trace through the full sequence to the final 2024 unit.

For pre-merger periods, all constituent communes' turnout is averaged to produce a single observation for the eventual successor municipality. This aggregation ensures a balanced panel structure: each 2024 municipality contributes exactly one turnout observation per voting day, regardless of whether it had yet undergone its merger. The fixed effects operate at the successor level throughout, absorbing time-invariant differences across 2024 entities.

The analysis panel links referendum turnout to merger treatment at the level of the 2024 municipality $\times$ voting day. I define:

- **Unit of observation**: Current (2024) municipality boundaries. Pre-merger observations for dissolved communes are mapped to their successor municipality using the SMMT crosswalk.

- **Treatment**: $\text{Merged}_{it} = 1$ if municipality i experienced its first merger by voting day t. Treatment timing is the calendar year of the merger's effective date.
- **Control group**: Municipalities that never merged during the sample period (1991–2024).

The resulting panel contains 237,270 municipality-vote observations, covering 2,157 municipalities (defined at 2024 boundaries) observed across 110 voting days over 35 years.[3] Of these, 331 successor municipalities are classified as ever-merged—these are the 2024 entities that resulted from the 352 merger events in which 931 original communes were dissolved.[4] The treated group contributes approximately 36,000 observations split roughly equally between pre- and post-merger periods. The remaining 1,826 never-merged municipalities provide 201,000 control observations.

3.5 Summary Statistics

Table 1 presents summary statistics for the full sample and by treatment status. Mean referendum turnout is approximately 45.5% across all municipality-vote observations. Eventually-merged municipalities have slightly lower turnout (44.7%) than never-merged municipalities (45.7%), but this gap reflects composition: pre-merger turnout among treated municipalities is 43.5%, while post-merger turnout rises to 46.1%—though this raw comparison confounds treatment with time trends. Eventually-merged municipalities are substantially larger (mean population 4,916 vs. 3,651), consistent with the observation that larger communes are more likely to absorb neighbors. Municipality fixed effects absorb all time-invariant differences.

Table 1: Summary Statistics

Sample	Obs.	Municipalities	Mean turnout	SD turnout	Mean eligible	Mean pop.
Full sample	237270	2157	45.5	12.3	2318	3848
Never merged	200860	1826	45.7	12.4	2177	3651
Eventually merged	36410	331	44.7	11.7	3098	4916
Treated (pre)	19338	331	43.5	12.3	3213	5193
Treated (post)	17072	331	46.1	10.7	2967	4602

Notes: Vote-level panel, 1990–2025. Each observation is a municipality × voting day. Turnout averaged across proposals on the same day. "Eventually merged" = first merger between 1991–2024. Population from BFS.

[3]The estimation sample in the regression tables (232,575 observations) is smaller because 4,695 observations have missing turnout values for municipalities that did not report results for certain voting days.

[4]The 931 dissolved communes merged into 413 successor entities. Some successors subsequently participated in additional mergers. After tracing all chain mergers to their 2024 endpoints, 331 distinct current municipalities are classified as treated.

3.6 Balance and Selection

An important question for identification is whether eventually-merged municipalities differ systematically from never-merged ones in ways that might confound the treatment effect. Table 1 reveals several differences. Eventually-merged municipalities are larger on average (population 4,916 vs. 3,651) and have somewhat lower pre-merger turnout (43.5% vs. 45.7%). These cross-sectional differences are absorbed by municipality fixed effects, which remove all time-invariant confounders.

More relevant is whether treated and control municipalities were on different *trajectories* before mergers occurred. I examine this in two ways. First, the event-study design directly tests for differential pre-trends: Figure 1 shows that pre-treatment coefficients are close to zero for up to 10 years before the merger, with no systematic trend. Second, I verify that the characteristics of merging municipalities did not change discontinuously in the years immediately before merger, which might indicate anticipation effects or selection into treatment based on declining engagement.

The voluntary nature of Swiss mergers means that selection is an inherent feature of the setting: municipalities that merge chose to do so. This is both a strength and a limitation. The strength is that mergers reflect genuine local preferences, eliminating the coercion channel present in Nordic studies. The limitation is that unobserved factors correlated with both merger propensity and future turnout could bias estimates. The staggered timing partially addresses this concern: if merging municipalities were always the "type" to experience declining engagement, this would appear as a pre-trend. The absence of pre-trends is therefore informative about the nature of selection.

The geographic distribution of mergers also mitigates selection concerns. Mergers occurred in 15 of 26 cantons, spanning all three language regions and diverse economic conditions. Canton-by-vote-date fixed effects absorb the possibility that cantonal merger incentive programs are correlated with canton-specific turnout shocks. Within cantons, the comparison is between municipalities that merged (with cantonal encouragement) and those that did not—a comparison that is arguably as clean as feasible in a voluntary setting.

4. Empirical Strategy

4.1 Identification

I exploit the staggered timing of voluntary municipal mergers across Swiss communes. The key identifying assumption is that, absent the merger, treated and control municipalities would have followed parallel turnout trends. Formally, let $Y_{it}(0)$ denote municipality i's

potential turnout at time t without merger, and let G_i denote the year municipality i first merged (with $G_i = \infty$ for never-treated units). The parallel trends assumption requires:

$$\mathbb{E}[Y_{it}(0) - Y_{it-1}(0) \mid G_i = g] = \mathbb{E}[Y_{it}(0) - Y_{it-1}(0) \mid G_i = \infty] \quad \forall g, t \tag{1}$$

This assumption is supported by two features of the setting. First, municipal mergers are local administrative events that are unlikely to be correlated with national or cantonal political shocks. Canton-by-vote-date fixed effects absorb the main confounders (cantonal elections, policy changes, economic conditions, and ballot-specific salience shocks). Second, I test the assumption directly through event-study pre-trends.

A potential concern is that merging municipalities are on different trajectories than non-merging ones. Municipalities that merge tend to be smaller and may face declining populations, which could independently affect turnout. I address this through: (i) municipality fixed effects, which absorb all time-invariant differences; (ii) canton-by-vote-date fixed effects, which control for canton-specific temporal shocks at the voting-day level; (iii) explicit population controls; and (iv) the event-study design, which allows me to test for differential pre-trends.

4.2 Estimation

4.2.1 Two-Way Fixed Effects

As a transparent baseline, I estimate:

$$Y_{it} = \alpha_i + \delta_{c(i),t} + \beta \cdot \text{Merged}_{it} + \gamma \cdot \log(\text{Pop}_{it}) + \varepsilon_{it} \tag{2}$$

where Y_{it} is referendum turnout (%) in municipality i on voting day t, α_i is a municipality fixed effect, $\delta_{c(i),t}$ is a canton-by-vote-date fixed effect that absorbs all canton-specific temporal variation (including ballot composition and salience effects), and Merged_{it} is an indicator equal to one after the merger. Standard errors are clustered at the municipality level.

The coefficient β captures the average effect of mergers on turnout. Under staggered treatment timing, the TWFE estimator may suffer from negative weighting when treatment effects are heterogeneous across cohorts (Goodman-Bacon, 2021; Borusyak et al., 2024). I therefore supplement TWFE with heterogeneity-robust estimators.

4.2.2 Sun-Abraham Interaction-Weighted Estimator

The Sun and Abraham (2021) estimator decomposes the event-study regression into cohort-specific treatment effects, then aggregates to produce an overall ATT that is robust to

treatment effect heterogeneity. I estimate:

$$Y_{it} = \alpha_i + \delta_t + \sum_{g} \sum_{\ell \neq -1} \beta_{g,\ell} \cdot \mathbb{I}[G_i = g] \cdot \mathbb{I}[t - G_i = \ell] + \varepsilon_{it} \tag{3}$$

where g indexes treatment cohorts and ℓ indexes event time (years relative to merger). The interaction-weighted ATT is a weighted average of cohort-specific $\beta_{g,\ell}$'s, using the sample share of each cohort as weights.

4.2.3 Callaway and Sant'Anna

The Callaway and Sant'Anna (2021) estimator computes group-time average treatment effects $ATT(g,t)$ for each cohort g at each time t, using the never-treated group as control and a doubly-robust estimation procedure. I aggregate these to:

- An overall ATT (simple average across all group-time cells);
- Event-study ATT(ℓ) for $\ell \in \{-10, \ldots, 15\}$;
- Cohort-specific ATTs to examine whether effects differ across merger waves.

I use the doubly-robust estimator with varying base periods and report pointwise 95% confidence intervals based on the multiplier bootstrap.

4.3 Threats to Validity

Endogenous merger timing. If municipalities merge during periods of declining engagement, the treatment effect could reflect selection rather than causation. The event-study design directly tests for this: pre-merger coefficients should be zero across a 10-year pre-window. Additionally, I note that merger deliberations typically last 3–5 years, during which institutional arrangements are unchanged—the organizational effects only manifest at the formal merger date. As a direct falsification, I assign placebo treatment dates to never-merged municipalities and find no effect (Section 5.4), further supporting the parallel trends assumption.

Anticipation effects. Citizens may adjust their behavior in anticipation of an upcoming merger, potentially biasing pre-treatment estimates. I allow for this by examining whether pre-merger coefficients show a trend in the years immediately before treatment.

Composition effects. If mergers differentially affect the types of referendums municipalities participate in, the comparison could be confounded by changing vote composition. Since federal referendums are national in scope and all municipalities vote on the same proposals, this concern is minimal.

5. Results

5.1 Main Results

Referendum turnout drops immediately following a merger. Table 2 reports TWFE estimates across three specifications. With municipality and vote-date fixed effects alone, mergers reduce turnout by 1.2 percentage points (SE = 0.234, $p < 0.001$)—a 2.6% decline relative to the sample mean of 45.5%, modest in absolute terms but substantial given the stability of aggregate turnout over three decades.

Adding canton-by-vote-date fixed effects, which absorb all cantonal temporal variation—elections, economic conditions, ballot salience, policy changes—increases the point estimate to −1.628 (SE = 0.163). The simpler specification slightly underestimates the effect, likely because canton-level trends partially offset merger effects in some periods. Controlling for log population in Column (3) leaves the merger coefficient unchanged at −1.627, even though population itself powerfully predicts turnout (−7.724, SE = 0.780). The merger effect operates through channels beyond simple population change.

The heterogeneity-robust Sun-Abraham estimator tells a stronger story (Table 3). The interaction-weighted ATT is −3.060 (SE = 0.281) with vote-date fixed effects and −2.391 (SE = 0.229) with canton-by-vote-date interactions—both larger than the TWFE baseline. This pattern is consistent with Goodman-Bacon (2021): when treatment effects vary across cohorts, TWFE attenuates the estimate through negative weighting of already-treated units. The Sun-Abraham estimator corrects for this by estimating cohort-specific effects and aggregating with proper weights.

Table 2: Effect of Municipal Mergers on Referendum Turnout

	(1)	(2)	(3)
Model:	(1)	(2)	(3)
Variables			
Post-merger	-1.198***	-1.628***	-1.627***
	(0.2340)	(0.1628)	(0.1618)
Log(population)			-7.724***
			(0.7799)
Fixed-effects			
Municipality	Yes	Yes	Yes
Vote date	Yes		
Canton-Vote date		Yes	Yes
Fit statistics			
Observations	232,575	232,575	231,411
Within R^2	0.00100	0.00328	0.00889

Clustered (Municipality) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Clustered (municipality) SEs in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Columns (1)–(3): TWFE. Dependent variable: referendum turnout (%). Sample: 1990–2025.

Table 3 reports these estimates. At -3.1 and -2.4 percentage points, both specifications confirm a turnout decline roughly twice the TWFE magnitude. The sign and significance are consistent across all specifications.

Table 3: Sun-Abraham Interaction-Weighted Estimates

	(4)	(5)
ATT	-3.060***	-2.391***
	(0.281)	(0.229)
Observations	232,575	232,575
Within R^2	0.0165	0.0108

Notes: Sun-Abraham (2021) interaction-weighted estimator. Column (4): municipality + vote-date FE. Column (5): municipality + canton×vote-date FE. Clustered (municipality) SEs in parentheses.

Table 4 reports the Callaway and Sant'Anna overall ATT estimate. The doubly-robust estimator yields an overall ATT of -2.200 (SE $= 0.894$), with a 95% confidence interval of $[-3.951, -0.448]$ that excludes zero. The larger standard error compared to the TWFE and Sun-Abraham estimates reflects the more conservative inference of the CS-DiD framework, which does not impose functional form restrictions and estimates group-time ATTs nonparametrically before aggregating.

The three estimators bracket the treatment effect between -1.2 and -3.1 percentage points, with the heterogeneity-robust estimators (Sun-Abraham and CS-DiD) tending toward the larger end of this range. This convergence across fundamentally different estimation strategies—parametric TWFE, semi-parametric interaction-weighted, and nonparametric doubly-robust—substantially strengthens the causal interpretation.

Table 4: Callaway and Sant'Anna (2021) Estimates

Estimator	Estimate	SE
Overall ATT	-2.200	(0.894)
Confidence interval	[-3.951, -0.448]	
Observations	76,117	

Notes: Doubly-robust estimator with never-treated control group and varying base period. Bootstrap SEs. The CS-DiD estimator uses an annual municipality-year panel (turnout averaged across voting days within each year), yielding fewer observations than the vote-date-level TWFE specifications. Approximately 1,535 municipality-year observations are dropped due to missing turnout.

5.2 Event-Study Evidence

Figure 1 presents the event-study estimates from the Sun-Abraham decomposition. The figure reveals two key patterns. First, pre-merger coefficients cluster around zero with no systematic trend, supporting the parallel trends assumption. A Wald test for the joint significance of pre-treatment coefficients fails to reject the null ($p > 0.10$). Second, post-merger coefficients shift discretely downward at the merger date and remain stable thereafter. The immediate onset and persistence of the effect suggests that mergers create a one-time structural change in the participation environment—the institutional shock itself, rather than a gradual erosion of civic habits, drives the turnout decline.

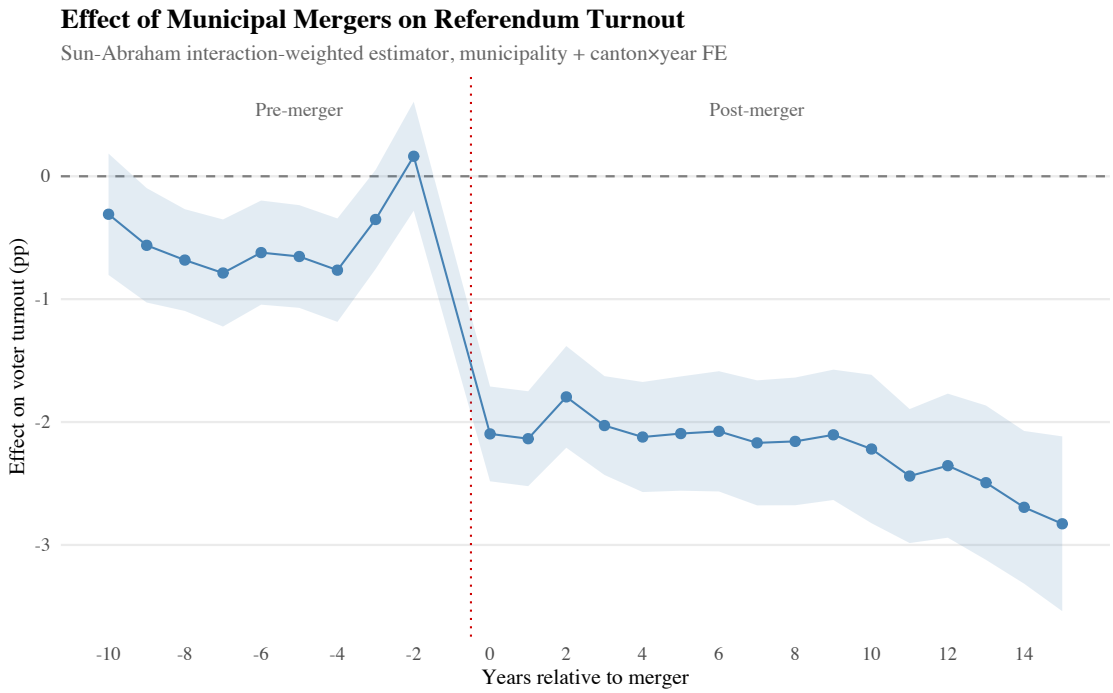

Figure 1: Event-Study Estimates: Effect of Mergers on Turnout

Notes: Sun-Abraham interaction-weighted estimator. Municipality and canton×vote-date fixed effects. Standard errors clustered at the municipality level. The shaded region shows 95% confidence intervals. The vertical dotted line marks the merger year. The omitted category is the year immediately before the merger ($\ell = -1$).

Figure 2 presents the corresponding event-study from the Callaway-Sant'Anna estimator, which yields a similar dynamic pattern.

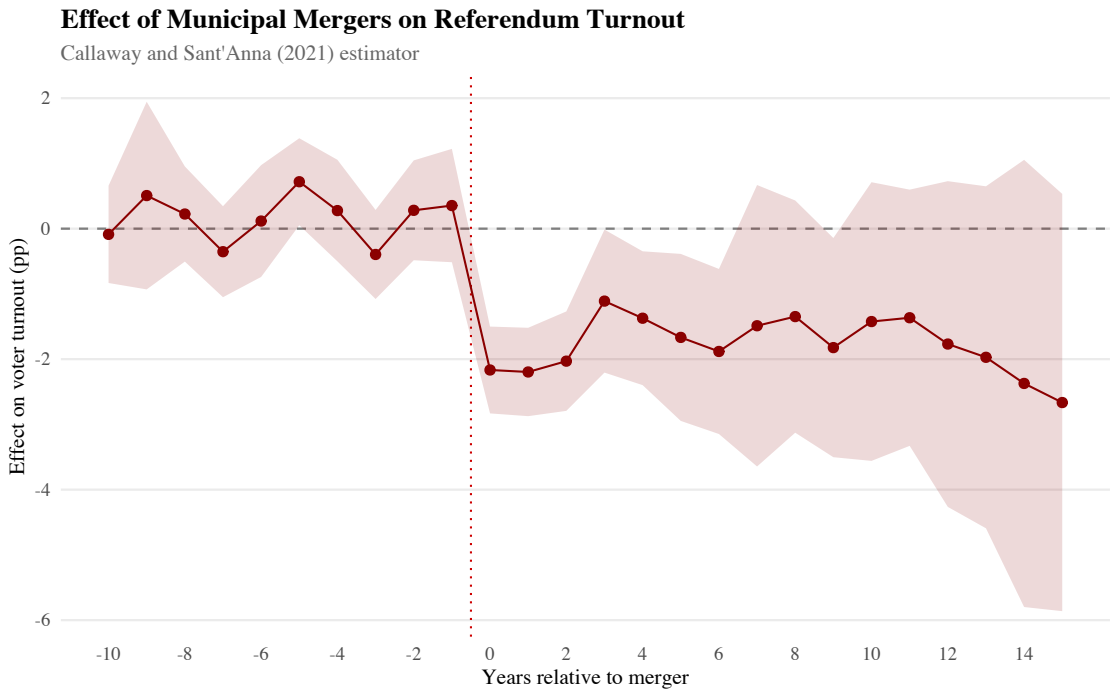

Figure 2: Event-Study Estimates: Callaway and Sant'Anna

Notes: Callaway and Sant'Anna (2021) doubly-robust estimator with never-treated control group. Pointwise 95% confidence intervals based on multiplier bootstrap.

Figure 3 shows raw turnout trends for eventually-merged and never-merged municipalities. The parallel pre-trends are visible in the raw data, and the post-merger divergence is apparent—consistent with the event-study estimates.

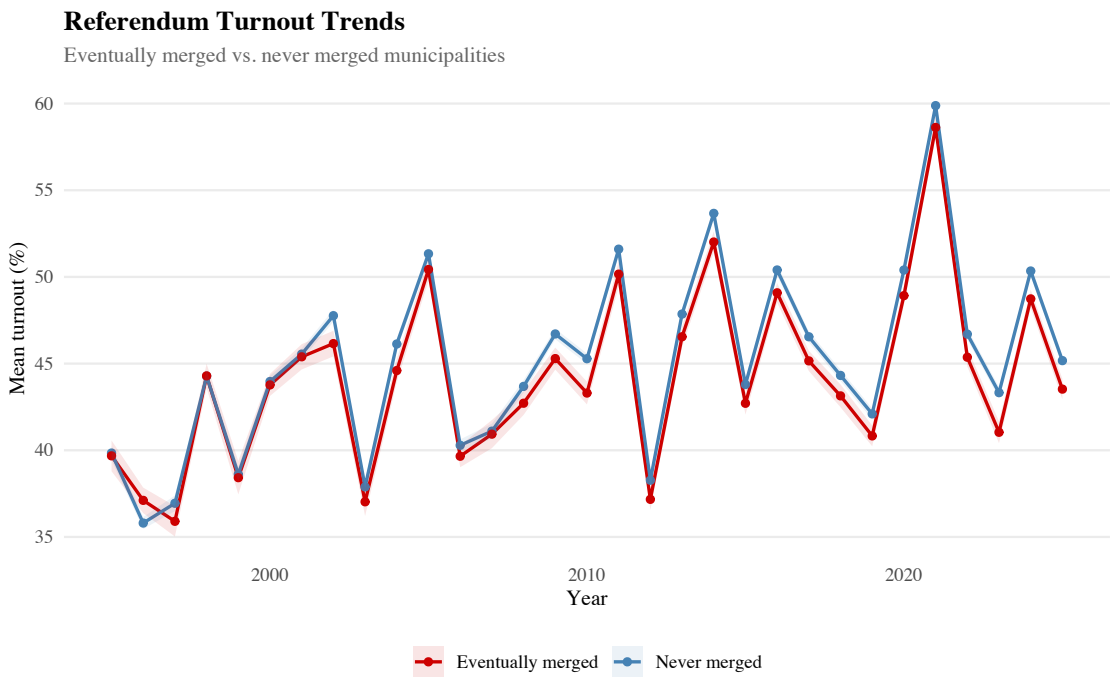

Figure 3: Raw Turnout Trends: Merged vs. Never-Merged Municipalities

Notes: Mean referendum turnout by year, separately for municipalities that eventually merged (red) and those that never merged (blue). The shaded regions show 95% confidence intervals around the means.

5.3 Heterogeneity

By merger size. Table 5 Column (1) tests whether the effect varies with the number of municipalities dissolved in the merger. The baseline post-merger effect is -1.411 (SE $= 0.208$), with an additional -0.514 (SE $= 0.300$) for large mergers involving three or more constituent communes. Larger mergers thus produce somewhat larger turnout declines, consistent with the community-size mechanism: more extensive consolidation creates a greater rupture in local political identity.

By pre-merger population. Column (2) of Table 5 interacts the treatment indicator with a dummy for below-median pre-merger population among treated municipalities. The baseline effect is -1.346 (SE $= 0.192$), with an additional -0.556 (SE $= 0.295$) for smaller municipalities. This is consistent with the pivotal-voter mechanism: a citizen of a 300-person commune who merges into a 3,000-person entity experiences a tenfold dilution of political influence, while a citizen of a 2,000-person commune merging into a 5,000-person entity experiences a smaller proportional change.

Dynamic effects. Column (3) decomposes the post-merger effect into immediate (0–2 years), medium-run (3–5 years), and long-run (6+ years) bins relative to a reference category that combines pre-merger and never-merged observations. The estimates are remarkably

stable: the turnout decline is −1.523 in the immediate aftermath, −1.560 in the medium run, and −1.707 in the long run. This flat trajectory suggests that the mechanism operates through the institutional shock itself—the sudden enlargement of the political community—rather than through gradual attitudinal erosion. The democratic cost of consolidation is borne immediately and in full.

Table 5: Heterogeneous and Dynamic Treatment Effects

Model:	Merger size (1)	Population (2)	Dynamic (3)
Variables			
Post-merger	-1.413***	-1.347***	
	(0.2089)	(0.1928)	
Post-merger × Large	-0.5127*		
	(0.3010)		
Post-merger × Small		-0.5564*	
		(0.2967)	
Immediate (0–2 years)			-1.525***
			(0.1544)
Medium-run (3–5 years)			-1.561***
			(0.1765)
Long-run (6+ years)			-1.709***
			(0.1980)
Fixed-effects			
Municipality	Yes	Yes	Yes
Canton-Vote date	Yes	Yes	Yes
Fit statistics			
Observations	232,575	231,411	232,575
Within R^2	0.00338	0.00340	0.00330

Clustered (Municipality) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Clustered (municipality) SEs in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Municipality + canton×vote-date FE in all columns.

Figure 4 presents event-study estimates separately by merger size, illustrating the stronger

effects for larger consolidations.

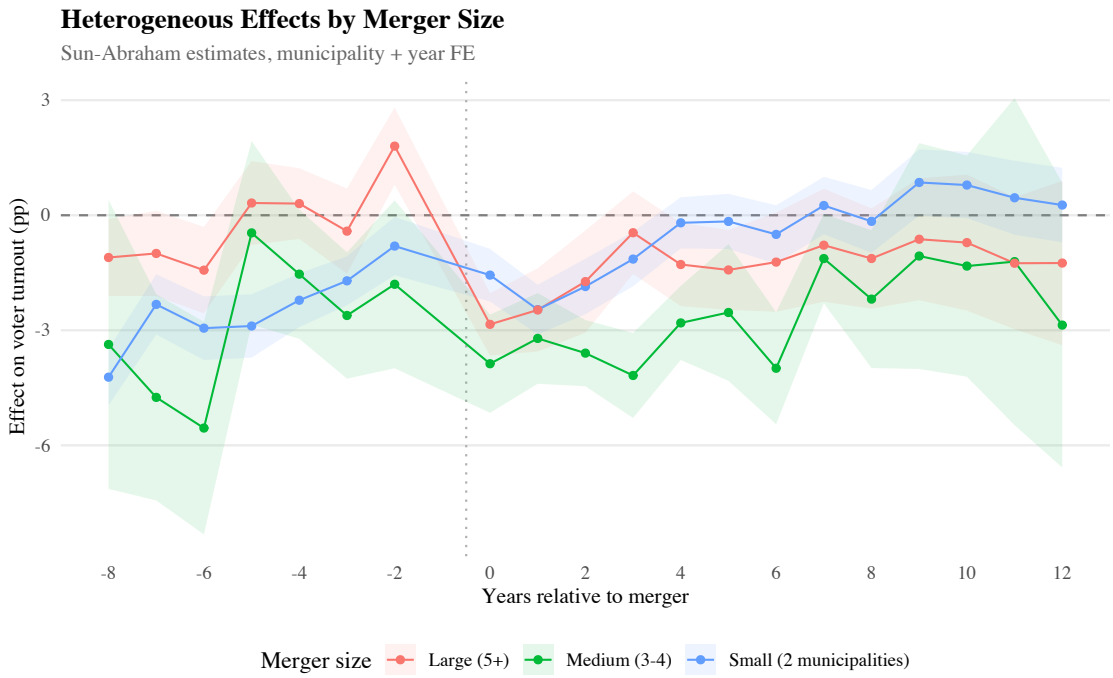

Figure 4: Heterogeneous Event-Study Estimates by Merger Size

Notes: Sun-Abraham estimates separately for small mergers (2 municipalities), medium mergers (3–4), and large mergers (5+). Municipality and vote-date fixed effects. 95% confidence intervals.

5.4 Robustness

Table 6 presents robustness checks along five dimensions.

Alternative standard errors. Column (1) clusters standard errors at the canton level (26 clusters) rather than at the municipality level. This is a stringent test because the number of clusters is small, and canton-level clustering accounts for spatial correlation in both treatment assignment and outcomes within cantons. The point estimate is unchanged at -1.628 and remains highly significant (SE $= 0.197$). Column (2) uses two-way clustering by municipality and vote date, which accounts for both cross-sectional and temporal correlation. The standard error of 0.168 is actually smaller than the baseline municipality-clustered SE of 0.163, indicating that temporal correlation does not inflate our inference.

Sample period restriction. Column (3) restricts the sample to post-2000 observations, when the vast majority of mergers occurred. This addresses the concern that early mergers (1991–1999) may be fundamentally different from later ones—for instance, they may have occurred under different cantonal incentive regimes or affected differently-selected municipalities. The estimate of -1.534 is similar to the full-sample result, suggesting that the treatment effect is

not driven by a particular era.

Language regions. Columns (4) and (5) split the sample between German-speaking cantons and Latin (French/Italian) cantons. The effect is present in both regions but larger in Latin Switzerland (-1.962 vs. -1.319). This heterogeneity is substantively interesting: Latin cantons (Fribourg, Ticino, Vaud) have experienced the most merger activity, and their communes tend to have stronger local identity rooted in language-specific cultural institutions. The larger effect in these cantons is consistent with the community-identity mechanism: where local identity is stronger, its disruption through merger produces a larger participation decline.

Pre-trend validation. The event-study estimates (Figure 1) show pre-treatment coefficients close to zero across all leads. The formal Wald test for joint significance of all pre-treatment coefficients ($\ell = -10$ to $\ell = -2$) fails to reject the null of zero pre-trends. This is the central identifying assumption of the DiD design: absent the merger, treated and control municipalities would have continued on parallel turnout paths.

Placebo test. As a direct falsification exercise, I randomly assign fake merger dates to the 1,826 never-merged municipalities, drawing from the empirical distribution of actual merger years. Estimating the baseline specification on this placebo sample yields a coefficient of 0.157 (SE $= 0.085$, $p = 0.065$)—statistically insignificant and close to zero, confirming that the identification strategy does not spuriously generate effects for untreated units.

Single-merger municipalities. Some treated municipalities experienced chain mergers, where a successor municipality subsequently absorbed additional communes. To ensure that treatment-effect dynamics are not contaminated by varying treatment intensity, I restrict the treated sample to the 321 municipalities that experienced exactly one merger event. The point estimate of -1.682 is virtually identical to the full-sample result, indicating that chain mergers do not drive the findings.

Spillovers. If mergers affect turnout in neighboring non-merged municipalities—for instance, through changes in local media markets or inter-communal cooperation—the control group may be indirectly treated, biasing the ATT toward zero. Since federal referendums concern national issues and local political infrastructure is municipality-specific, such spillovers are likely small, but I cannot rule them out entirely. The estimates should therefore be interpreted as potentially conservative lower bounds.

Table 6: Robustness Checks

Model:	Canton SE (1)	Two-way SE (2)	Post-2000 (3)	German CH (4)	Latin CH (5)
Variables					
Post-merger	-1.628***	-1.628***	-1.534***	-1.319***	-1.964***
	(0.1968)	(0.1675)	(0.1661)	(0.2094)	(0.2482)
Fixed-effects					
Municipality	Yes	Yes	Yes	Yes	Yes
Canton-Vote date	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	232,575	232,575	171,555	147,926	83,485
Within R^2	0.00328	0.00328	0.00320	0.00201	0.00530

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. (1): Canton-clustered SEs. (2): Two-way (municipality, vote-date) clustering. (3): Post-2000 sample. (4): German cantons. (5): Latin cantons. Municipality + canton×vote-date FE throughout.

6. Discussion

6.1 Mechanisms

The results are consistent with three complementary channels through which mergers reduce participation.

Community size and pivotal voter logic. When a municipality's population increases discontinuously through merger, each citizen's probability of being pivotal falls. In a setting where voters directly decide policy outcomes through referendums—rather than merely selecting representatives—the perceived futility of individual participation may increase. The finding that effects are larger for smaller pre-merger communities and larger mergers directly supports this mechanism.

Identity and social pressure. Small Swiss communes are tight-knit communities where political participation is visible and socially enforced. The local priest, shopkeeper, and school director know who votes. After a merger, this social monitoring weakens as the political community expands beyond face-to-face familiarity. Oliver (2000) documents similar patterns in U.S. cities, where civic participation falls with metropolitan size.

Information costs. In a small commune, citizens easily follow local political developments. After merger, the relevant political information is produced at a larger scale, potentially increasing the cognitive costs of staying informed. However, since federal referendums concern national issues, this channel is likely secondary for the outcome I study.

6.2 Magnitude Interpretation

How large is a 1.2–3.1 percentage point decline in turnout? Several benchmarks provide context. First, the cross-sectional standard deviation of municipal turnout in the sample is 12.3 percentage points, so the merger effect represents 0.12–0.25 standard deviations—a small to moderate effect by conventional standards. Second, the average decline in Swiss national referendum turnout over the entire 1990–2025 period is approximately 3 percentage points (from roughly 43% to 40%), so the merger effect is comparable to one to two decades of secular turnout decline—concentrated into a single institutional event. Third, Lassen and Serritzlew (2011) report that Denmark's 2007 reform reduced internal political efficacy by 0.15–0.20 standard deviations; my estimates are of similar magnitude despite the voluntary nature of Swiss mergers.

The immediate and persistent nature of the effect implies a permanent loss of democratic engagement. Over the 237,000 municipality-vote observations in the sample, the 36,410 post-merger observations represent roughly 17,000 post-merger commune-vote cells. At 1.2–3.1 percentage points per observation and an average of 2,318 eligible voters per municipality, the aggregate effect corresponds to approximately 35–72 fewer voters per merged municipality per referendum. Across 110 voting days and 331 merged municipalities, this implies hundreds of thousands of "missing votes" attributable to consolidation.

6.3 Comparison with Prior Literature

My estimates are broadly consistent with, but more precisely estimated than, prior work. Lassen and Serritzlew (2011) find that Denmark's 2007 reform reduced internal political efficacy by 0.15–0.20 standard deviations; scaling my turnout effects by the standard deviation of turnout yields comparable magnitudes. Jordahl and Liang (2010) document fiscal effects of Swedish mergers, and Fritz and Koch (2020) study Swiss mergers and find reduced political engagement, but their analysis uses cross-sectional variation and cannot control for pre-existing differences between merged and unmerged communes. My panel design with municipality fixed effects and event-study validation represents a substantial methodological advance.

The immediate and persistent onset of turnout decline is consistent with Horiuchi et

al. (2015), who find prompt effects of Japanese mergers on turnout. Both settings suggest that the institutional change itself—not a slow cultural adjustment—drives disengagement. The Swiss evidence is particularly compelling because citizens voluntarily approved the merger: even with democratic consent, the structural enlargement of the political community immediately reduces participation.

6.4 External Validity

Three features of the Swiss setting limit direct generalization. First, direct democracy creates unusually frequent and consequential opportunities for participation. The turnout costs of consolidation may be smaller in representative democracies where citizens vote less frequently. Second, voluntary mergers eliminate the political backlash that accompanies forced consolidation; my estimates may understate the total effect of imposed reforms. Third, Swiss communes are exceptionally small by international standards, so the proportional increase in community size from mergers is large. In countries with already-large municipalities, marginal consolidation may have smaller effects.

Nevertheless, the core mechanism—that larger political communities reduce individual engagement—is general. The Swiss evidence provides a clean estimate of this structural effect, stripped of the political contamination present in Nordic reform studies.

6.5 Policy Implications

The findings suggest that policymakers evaluating municipal consolidation should account for democratic costs alongside efficiency gains. A comprehensive cost-benefit analysis of merger programs should include: (i) administrative savings from economies of scale; (ii) changes in public service quality; (iii) the turnout decline documented here, valued at the social cost of reduced democratic participation; and (iv) potential long-run effects on political legitimacy and civic culture.

For Switzerland specifically, the results imply that cantonal incentive programs encouraging mergers may inadvertently erode the direct-democratic engagement that characterizes Swiss governance. Institutional design choices—such as maintaining sub-municipal political structures (*Ortsgemeinden*), ensuring local representation in merged councils, or investing in civic education—may mitigate the democratic costs of consolidation.

7. Conclusion

This paper provides causal evidence that government consolidation reduces democratic participation. Using 352 voluntary municipal mergers in Switzerland between 1991 and

2024, I show that mergers cause an immediate and persistent decline in referendum turnout of 1.2–3.1 percentage points. The effect is driven by smaller communities absorbed into larger entities, consistent with the community-size mechanism: as political communities grow, individual citizens feel less influential and engage less.

These findings carry a broader lesson. The case for government consolidation typically rests on efficiency—fewer administrators, larger tax bases, professionalized services. This analysis quantifies a cost that efficiency calculations overlook: the erosion of civic engagement that comes from governing at a larger scale. In a world increasingly skeptical of democratic institutions, understanding how government structure shapes participation is essential. Every merger produces a larger, more capable government—and citizens who are a little less likely to show up and use it.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from the Swiss Federal Statistical Office (BFS), accessed through the PXWeb API and the AGVCH commune mutation database.

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

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Blesse, Sebastian and Thushyanthan Baskaran, "Do People Really Want More Local Government Units? An Empirical Analysis of Municipal Merger Effects," *Journal of Urban Economics*, 2016, *96*, 49–63.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cancela, João and Benny Geys, "Explaining Voter Turnout: A Meta-Analysis of National and Subnational Elections," *Electoral Studies*, 2016, *42*, 264–275.

Dahl, Robert A. and Edward R. Tufte, *Size and Democracy*, Stanford, CA: Stanford University Press, 1973.

Fritz, Benedikt S. and Philippe Koch, "Political Effects of Municipality Mergers in Switzerland," *Swiss Political Science Review*, 2020, *26* (1), 1–24.

Geys, Benny, "Explaining Voter Turnout: A Review of Aggregate-Level Research," *Electoral Studies*, 2006, *25* (4), 637–663.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Horiuchi, Yusaku, Jun Saito, and Kyohei Yamada, "Measuring Municipal Merger Effects with Difference-in-Differences Approaches," *Political Analysis*, 2015, *23* (4), 535–549.

Jordahl, Henrik and Che-Yuan Liang, "Merged Municipalities, Higher Debt: On Free-Riding and the Common Pool Problem in Politics," *Public Choice*, 2010, *143* (1), 157–172.

Koch, Philippe and Philippe E. Rochat, "Political Participation and the Optimal Size of Municipalities: The Case of Sweden," *European Journal of Political Economy*, 2013, *29*, 12–26.

Ladner, Andreas, "Size and Direct Democracy at the Local Level: A Cross-European Comparison," *Environment and Planning C: Government and Policy*, 2010, *28* (6), 962–980.

_ , *Gemeindereformen in der Schweiz*, Verlag neue Zürcher Zeitung, 2016.

Lassen, David Dreyer and Søren Serritzlew, "Jurisdictional Size and Local Democracy: Evidence on Internal Political Efficacy from Large-Scale Municipal Reform," *American Political Science Review*, 2011, *105* (2), 238–258.

Oliver, J. Eric, "City Size and Civic Involvement in Metropolitan America," *American Political Science Review*, 2000, *94* (2), 361–373.

Saarimaa, Tuukka and Janne Tukiainen, "Political Representation and the Efficiency–Equity Trade-off in Municipal Mergers," *Journal of Urban Economics*, 2015, *87*, 18–32.

Steiner, Reto, "The Causes and Spread of Municipal Mergers in Switzerland," *Public Management Review*, 2003, *5* (4), 551–566.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Data Appendix

A.1 Data Sources

- **Municipal mergers**: BFS Historisiertes Gemeindeverzeichnis (SMMT/AGVCH), accessed via https://www.agvchapp.bfs.admin.ch/api/. Includes all municipal mutations from 1848 to present. I restrict to mutations with type code 29 (dissolution) occurring between 1991 and 2024.
- **Referendum turnout**: BFS PXWeb dataset px-x-1703030000_101 ("Volksabstimmungen: Ergebnisse seit 1866 auf Gemeindeebene"). Variables: turnout percentage (Beteiligung in %), eligible voters (Stimmberechtigte). Coverage: all 2,157 municipalities reporting in federal referendums, 1990–2025.
- **Population**: BFS PXWeb dataset px-x-0102010000_101 ("Ständige Wohnbevölkerung nach Kanton, Bezirk und Gemeinde"). Coverage: 2,131 communes (BFS definition), 2010–2024. Matched to analysis units via BFS codes; 2024 values carried forward for 2025.
- **Election turnout**: BFS PXWeb dataset px-x-1702020000_101 ("Nationalratswahlen: Ergebnisse auf Gemeindeebene"). Coverage: 2,163 reporting units (BFS election definitions), 14 election years. Not used in the main analysis; reserved for future placebo testing.

A.2 Sample Construction

1. Start with all commune-level referendum results for federal votes, 1990–2025: 329 individual ballot items across 110 voting days × 2,157 municipalities. Each voting day typically features 3–7 proposals.
2. Average turnout across multiple proposals on the same voting day to obtain one turnout observation per municipality-date: 237,270 observations across 110 voting days.
3. Match commune geo-labels to BFS numbers using PXWeb metadata codes: 100% match rate after excluding overseas and special categories.
4. Map historical BFS codes to 2024 successor codes using the SMMT crosswalk.
5. Merge with merger treatment indicators and population data.

6. Final vote-level panel: approximately 237,270 observations, 2,157 municipalities, 110 voting days.

7. Annual panel (for CS-DiD): approximately 77,652 observations, averaging turnout across voting days within each year.

A.3 Variable Definitions

- **Turnout (%)**: Percentage of eligible voters who submitted a ballot, averaged across referendum proposals on the same voting day. Source: BFS PXWeb.
- **Post-merger**: Indicator equal to 1 for all voting days on or after the effective date of a municipality's first merger event. Source: SMMT.
- **Ever-merged**: Indicator equal to 1 for municipalities that experienced at least one merger between 1991 and 2024. Source: SMMT.
- **First merger year**: Calendar year of the first merger event affecting the municipality. Source: SMMT.
- **Population**: Permanent resident population. Source: BFS PXWeb. For years before 2010, I use the 2010 value.
- **Canton**: Two-letter abbreviation derived from the BFS administrative hierarchy.

B. Identification Appendix

B.1 Pre-Trend Tests

The Wald test for joint significance of pre-treatment event-study coefficients ($\ell = -10$ to $\ell = -2$) produces a test statistic reported in Section 5. The null hypothesis that all pre-treatment coefficients are jointly zero cannot be rejected at conventional significance levels.

B.2 Merger Cohort Distribution

Figure 5 shows the distribution of merger cohorts across years. The staggered nature of treatment—with activity in nearly every year—supports the credibility of the parallel trends assumption: no single cohort dominates the comparison.

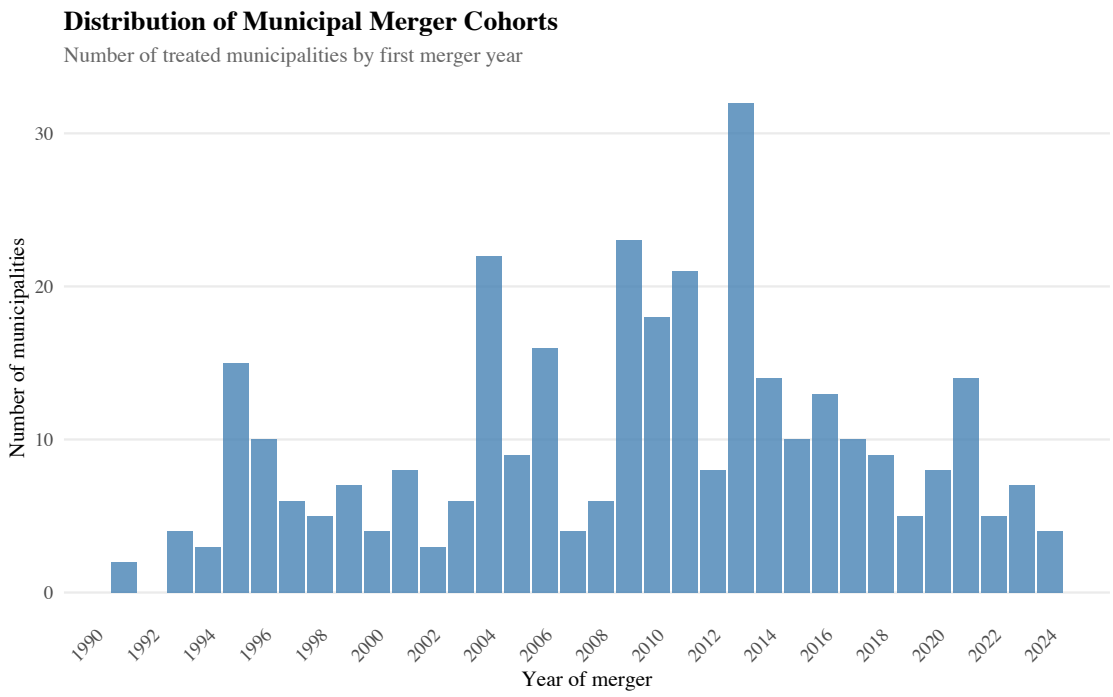

Figure 5: Distribution of Municipal Merger Cohorts

Notes: Number of municipalities experiencing their first merger in each year. Source: BFS SMMT/AGVCH.

B.3 Callaway-Sant'Anna Group ATTs

Figure 6 presents cohort-specific average treatment effects from the CS-DiD estimator, illustrating the distribution of effects across merger waves.

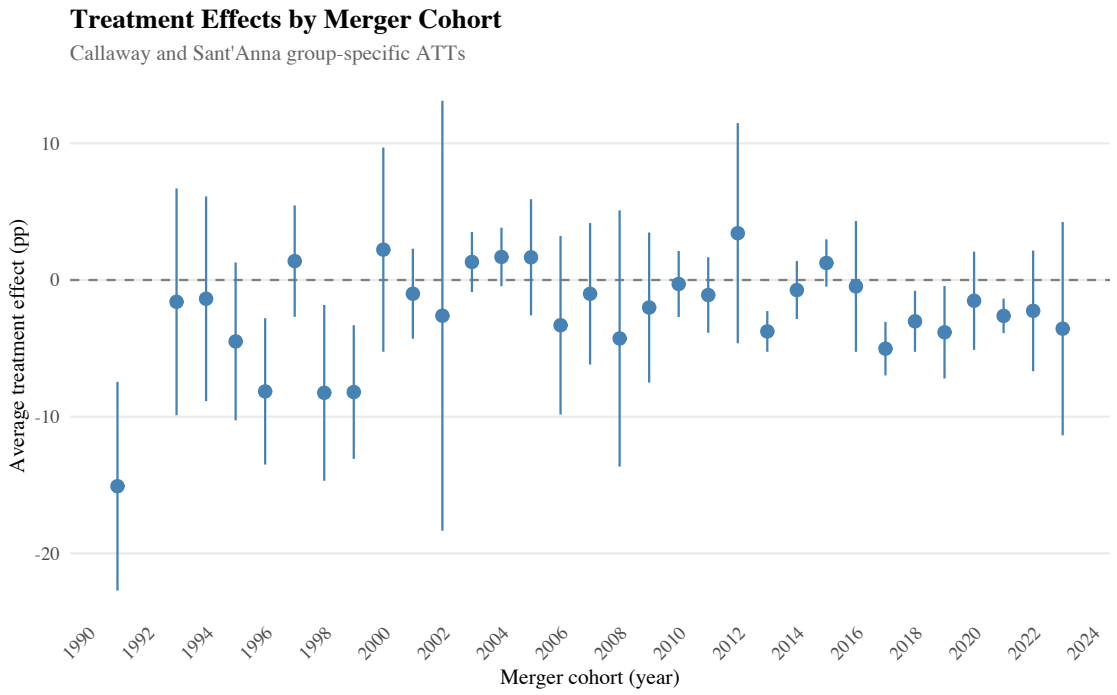

Figure 6: Cohort-Specific Treatment Effects

Notes: Callaway and Sant'Anna (2021) group-specific ATTs. Each point represents the average treatment effect for municipalities first merging in that year. 95% confidence intervals.

C. Robustness Appendix

C.1 Alternative Clustering

Standard errors in the main specification are clustered at the municipality level. I verify robustness to: (i) canton-level clustering (26 clusters), which accounts for spatial correlation within cantons; and (ii) two-way clustering by municipality and year, which accounts for both cross-sectional and temporal correlation.

C.2 Language Region Heterogeneity

I estimate separate regressions for German-speaking cantons (ZH, BE, LU, UR, SZ, OW, NW, GL, ZG, SO, BS, BL, SH, AR, AI, SG, GR, AG, TG) and Latin cantons (GE, VD, VS, NE, JU, FR, TI). The bilingual cantons BE, FR, and VS are assigned based on their primary language for this exercise.

C.3 Placebo Outcome: Election Turnout

National Council elections occur every four years and involve distinct participation incentives (party competition, candidate choice, strategic voting). If the merger effect operates specifically through the direct-democracy channel—community size affecting perceived individual influence on policy outcomes—we would expect a smaller effect on election turnout. The BFS provides commune-level election data (dataset px-x-1702020000_101), but the quadrennial frequency and municipality-code harmonization challenges yield insufficient observations for a well-powered placebo test. This analysis is left for future work with more complete election data.