

The Merger That Didn't Move: Municipal Consolidation and Foreign Residential Sorting in Switzerland

APEP Autonomous Research* @olafdrw

April 1, 2026

Abstract

Municipal mergers are often justified by administrative efficiency, but they may also alter the local tax-service bundle and the social character of a jurisdiction. I study whether Swiss municipal consolidation changed where foreign residents live. Using 47 mergers between 2015 and 2020, harmonized to current municipal boundaries, and 1,798 never-merged municipalities as controls, I estimate staggered difference-in-differences models on annual population-by-citizenship data from 2010 to 2024. The preferred Sun-Abraham estimate for the foreign population share is 0.002 in share terms (s.e. 0.002), and event-study coefficients show no visible post-merger break. Log foreign population, foreign population growth, and subgroup estimates are similarly small or imprecise. Swiss mergers appear to matter for local political participation, but in this sample they do not materially reallocate foreign residents.

JEL Codes: H70, J15, R23

Keywords: municipal mergers, immigration, residential sorting, Switzerland, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 35m).

1. Introduction

Governments consolidate municipalities to save money, professionalize administration, and stabilize service delivery. Those reforms are usually evaluated in fiscal or political terms: do mergers reduce spending, and do they weaken local democratic participation? Yet mergers also change the *location choice problem* facing households. A merger can flatten tax differences across neighboring jurisdictions, repackage local public services, and dilute the identity of a small community inside a larger one. If those changes matter for residential sorting, the effects should appear in local demographic composition.

This paper asks whether Swiss municipal mergers changed the residential location of foreign residents. Switzerland is a useful setting because mergers are common, precisely dated, and mostly voluntary. The same institutional feature that makes Switzerland a canonical laboratory for direct democracy also makes it a clean laboratory for local consolidation: municipalities are numerous, boundaries matter, and reform timing is recorded in the official commune mutations registry. Between 2015 and 2020, 47 current-boundary municipalities in my sample experienced a first merger, while 1,798 never merged. I combine these treatment dates with annual municipality-level population data by citizenship from the Swiss Federal Statistical Office.

The headline result is a disciplined null. In the preferred Sun-Abraham specification, mergers increase the foreign population share by 0.002 in share terms with a standard error of 0.002. Relative to a pre-merger foreign share of roughly 15 percent in both treated and control municipalities, the estimate is economically tiny. The log foreign population effect is similarly close to zero, and the growth-rate estimates are positive but imprecise. Event-study coefficients around the merger date are also flat: there is no visible break at treatment and no evidence of a delayed demographic adjustment. Weighted specifications, canton-by-year fixed effects, and sample restrictions all tell the same story.

That null is informative. The political economy of municipal reform often assumes a trade-off between administrative scale and local attachment. Existing Swiss evidence shows that mergers can depress direct-democratic participation ([Koch and Rochat, 2017](#); [Frey, 2023](#)). If the same reforms also reshaped local amenities in ways salient to migrants, one would expect at least modest changes in foreign population shares or foreign population growth. I do not find them. In this setting, mergers appear to affect *voice* more than *exit*.

The paper contributes to three literatures. First, it adds a demographic and urban-sorting margin to the municipal amalgamation literature, which focuses mainly on spending, turnout, and administrative capacity ([Allers and Geertsema, 2016](#); [Blom-Hansen et al., 2016](#); [Ladner and Buhlmann, 2009](#); [Steiner and Kaiser, 2017](#)). Second, it links local-government structure

to Tiebout-style sorting and fiscal federalism (Tiebout, 1956; Oates, 1972). Third, it applies modern staggered-adoption estimators to a voluntary reform setting where treatment timing is heterogeneous and naive two-way fixed effects are potentially misleading (Callaway and Sant’Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021; Rambachan and Roth, 2023; Roth et al., 2023).

The paper is deliberately modest about identification. Swiss mergers are voluntary, and merged municipalities may differ from never-merged municipalities in slow-moving ways. I therefore emphasize event-study evidence and a narrow empirical claim: I do not detect a meaningful post-merger break in foreign residential composition. That is a more defensible contribution than claiming that mergers causally raise or lower immigrant settlement by a precisely estimated amount.

2. Institutional Background

Swiss municipalities are small by international standards and historically numerous. Consolidation has been a recurrent policy response to administrative fragmentation, especially in cantons that face difficulty staffing local offices or financing public services in tiny communes. Unlike forced territorial reforms, Swiss mergers typically proceed through feasibility studies, cantonal encouragement, and local approval. This creates substantial heterogeneity in timing and geography, but it also means merger decisions are closely tied to local conditions (Tavares, 2018; Swianiewicz, 2010).

Three features of the Swiss setting matter for this paper.

First, municipality boundaries are meaningful. Tax rates, local public services, and political participation all have a strong municipal component, so consolidation is not a purely administrative relabeling.

Second, the reform is common enough to study with a panel. The BFS AGVCH/SMMT mutations registry records dissolved municipalities, successors, and effective dates.

Third, the country maintains high-quality annual local population statistics by citizenship, allowing me to observe whether demographic composition changes after merger.

Why might mergers change foreign residential sorting? One mechanism is fiscal. If neighboring municipalities differed in tax burden or service provision before consolidation, merger harmonization can compress those differences, weakening the gains from sorting across very small jurisdictions. Another is social. A merger can preserve administrative continuity while diluting a strong local identity, potentially changing how welcome outsiders perceive a

place to be. A third mechanism runs the other way: larger successor municipalities may be administratively stronger and therefore better able to provide services valued by both Swiss and foreign residents. The sign is ambiguous *ex ante*.

The ambiguity makes a null result scientifically useful. If political participation falls after merger while demographic composition does not move, then the core social response to consolidation is political disengagement rather than household relocation.

3. Data

The first input is the BFS AGVCH commune mutations API, which provides the merger timeline. I use the 2024 commune snapshot and the mutations registry to map each historical BFS code recursively into its eventual current-boundary successor municipality. Pre-merger population counts are then summed across all predecessor communes that map into the same current municipality-year, so the pre-merger foreign share is a population-weighted aggregation on the same boundary that appears after consolidation. The treatment unit in the paper is therefore the *current-boundary successor municipality*, not the dissolved predecessor commune. This avoids the standard boundary-break problem in municipal panels.

The second input is the BFS population-by-citizenship table `px-x-0102010000_104`, which reports annual total and Swiss population by municipality. I derive the foreign population residually as total minus Swiss population, so the foreign-resident measure corresponds to non-Swiss citizens rather than broader migration-background categories. The available sample runs from 2010 through 2024. To satisfy the minimum pre-period requirement for a V1 DiD design, I restrict treated municipalities to those with a first merger between 2015 and 2020. That yields five clean pre-treatment years for the earliest treated cohort, at the cost of excluding earlier mergers from this short-format version of the project.

After harmonization, the analysis panel contains 27,675 municipality-year observations across 1,845 municipalities. Of these, 47 are treated successor municipalities and 1,798 are never-merged controls. The outcome variables are: foreign population share, log foreign population, foreign population growth, change in foreign population share, Swiss population growth, and total population growth.

Two facts from the summary statistics are immediately useful. First, treated and control municipalities look similar on the main outcome before merger: the foreign-share mean is 15.4 percent in treated municipalities and 15.3 percent in controls. Second, treated municipalities are larger in the sense of foreign population counts, which is why weighted and unweighted specifications are both worth reporting.

Table 1: Pre-Merger Summary Statistics (2010–2014)

	Foreign share	Log foreign pop.	Foreign growth (%)	Total growth (%)	Municipalities
Treated	0.154	6.021	5.144	0.860	47
Control	0.153	5.177	5.698	1.110	1798

Notes: Summary statistics use the 2010–2014 pre-merger window. Treated municipalities are current-boundary successor communes that first merge between 2015 and 2020. Control municipalities never merge during the sample. Foreign growth and total growth are year-over-year percentage changes.

4. Empirical Strategy

The baseline estimating equation is

$$Y_{it} = \beta \text{PostMerger}_{it} + \alpha_i + \gamma_t + \varepsilon_{it}, \quad (1)$$

where i indexes current-boundary municipalities and t indexes years. α_i are municipality fixed effects and γ_t are year fixed effects. Standard errors are clustered at the municipality level.

Because merger timing is staggered, I do not rely exclusively on TWFE. The preferred estimates come from a Sun-Abraham event-study specification that allows treatment effects to vary across cohorts:

$$Y_{it} = \sum_{k \neq -1} \beta_k 1\{t - G_i = k\} + \alpha_i + \gamma_t + \varepsilon_{it}, \quad (2)$$

where G_i is the first merger year for treated municipalities. I report the aggregated ATT and the event-study path. Never-merged municipalities form the control group.

The identifying assumption is weaker than literal random assignment but stronger than pure descriptive comparison: absent merger, treated municipalities would not have experienced a sharp *break* in foreign residential composition relative to controls. The event-study helps assess that claim. If pre-treatment coefficients are flat and the post-treatment coefficients remain small, the design is informative even if one remains skeptical that mergers are fully exogenous. Because Swiss mergers are often planned through feasibility studies and local votes, any anticipation during the run-up to the legal merger date would tend to attenuate the measured post-merger break.

5. Results

The main results show no meaningful demographic reallocation after merger. The preferred Sun-Abraham estimate for foreign population share is 0.002, with a standard error of 0.002.

Table 2: Municipal Mergers and Foreign Residential Composition

	Foreign share	Log foreign pop.	Foreign growth (%)	Δ foreign share
Sun-Abraham ATT	0.002 (0.002)	-0.007 (0.017)	1.799 (1.221)	0.001 (0.001)
TWFE	0.003 (0.003)	-0.012 (0.023)	0.716 (0.614)	0.000 (0.001)

Notes: Each column reports the post-merger effect of municipal consolidation on a foreign-population outcome. The preferred estimate is the aggregated Sun-Abraham ATT with municipality and year fixed effects. TWFE is shown for comparison. Standard errors are clustered at the municipality level. Treated municipalities are those first merging between 2015 and 2020; never-merged municipalities are the control group. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The sign is slightly positive, not negative, but the effect is too small to support a substantive sorting interpretation. The associated 95 percent confidence interval runs from roughly -0.003 to 0.006 , which rules out only fairly large changes in average foreign share. The log foreign population estimate is similarly close to zero. Foreign population growth is more volatile and the point estimate is positive, but it is estimated too imprecisely to pin down whether the true effect is close to zero or modestly positive.

These magnitudes are economically unremarkable. A 0.002 ATT on foreign share is only a small fraction of the pre-merger foreign share in the panel. The weighted robustness estimate is actually negative. This is not a setting where one specification finds a large effect and another washes it out; rather, the entire reasonable specification set points toward demographic stability.

Table 3: Event Study: Foreign Population Share

Relative year	Estimate	SE
-5	-0.001	(0.002)
-4	-0.000	(0.002)
-3	0.000	(0.001)
-2	0.000	(0.001)
0	-0.000	(0.001)
1	0.001	(0.002)
2	0.003	(0.002)
3	0.002	(0.002)
4	0.001	(0.002)
5	-0.000	(0.003)

Notes: Coefficients come from a Sun-Abraham event-study specification with municipality and year fixed effects. Relative year -1 is omitted. Pre-treatment coefficients are small, and post-treatment coefficients remain close to zero. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The event study reinforces that interpretation. Pre-treatment coefficients are small, and

the post-treatment path never shows a discrete or persistent jump. If mergers were materially changing the attractiveness of municipalities to foreign residents, one would expect either an immediate break in the share series or a steadily accumulating post-merger wedge. Neither is present.

6. Robustness and Comparison Outcomes

Table 4: Robustness for Foreign Population Share

Specification	Estimate	SE	N
Baseline TWFE	0.003	(0.003)	27675
Canton x year FE	0.004	(0.003)	27675
Population-weighted Drop top 1	-0.002	(0.004)	27675

Notes: The dependent variable is the foreign population share. Column 2 absorbs canton-specific shocks with canton-by-year fixed effects. Column 3 weights observations by baseline 2014 population. Column 4 drops the largest 1% of municipalities by baseline size. The point estimates remain economically small across all specifications.

The foreign-share result is stable across four checks. Adding canton-by-year fixed effects increases the estimate modestly but leaves it small. Weighting by baseline municipal population flips the sign. Dropping the largest one percent of municipalities changes almost nothing. The natural conclusion is that the foreign-share effect is close to zero, not that it is hiding behind one fragile modeling choice.

Table 5: Comparison Outcomes

Outcome	Estimate	SE	N
Log foreign population	-0.012	(0.023)	27675
Foreign population growth	0.716	(0.614)	25790
Swiss population growth	0.077	(0.153)	25830
Total population growth	0.122	(0.173)	25830
High baseline foreign share	0.001	(0.004)	13845
Low baseline foreign share	0.005	(0.004)	13830

Notes: The upper panel reports TWFE estimates for additional demographic outcomes. The lower panel splits the foreign-share specification by whether the municipality's 2014 foreign share was above or below the sample median. None of the comparison or subgroup estimates indicate large post-merger re-sorting effects.

Comparison outcomes tell the same story. Log foreign population does not show a precise decline after merger. Swiss population growth and total population growth are also modest

and imprecise, suggesting that consolidation did not trigger a broader demographic shake-up that happened to leave the foreign share unchanged. Split-sample estimates by baseline foreign share are larger in low-foreign-share municipalities, but they remain imprecise and are best read as exploratory rather than as evidence of systematic heterogeneity.

This pattern matters for interpretation. If mergers were pushing foreigners out while attracting Swiss households, or vice versa, one would expect some non-zero signal in these companion outcomes. Instead, the demographic consequences of consolidation appear muted across the board.

7. Discussion

The paper’s empirical message is narrower than “mergers do nothing.” Swiss municipal mergers have documented political consequences, and the fiscal literature still debates their efficiency effects (Allers and Geertsema, 2016; Ladner and Buhlmann, 2009; Steiner and Kaiser, 2017). My result is instead that the margin of *foreign residential sorting* seems largely unaffected in the medium run.

That finding is consistent with a simple interpretation: the local changes induced by Swiss mergers are salient enough to matter for political participation but not large enough to overturn the fundamentals driving residential location. Jobs, housing, commuting patterns, and social networks may dominate any merger-induced change in local identity or administration. If so, reformers evaluating consolidation should not expect much demographic re-sorting, even when political participation changes measurably.

The paper also provides a caution about what a credible null looks like. The design does not prove that no municipality ever became more or less attractive after merger. It does show that, in a well-powered panel with 47 treated successor municipalities and a large control group, there is no evidence of an economically meaningful average break in foreign residential composition. That is a useful factual constraint on theories that assign a large allocative role to municipal boundaries in this context.

8. Conclusion

Swiss municipal mergers between 2015 and 2020 do not appear to have materially reallocated foreign residents. Across foreign population share, log foreign population, and growth-based outcomes, the estimates are small, imprecise, and dynamically flat. The most natural interpretation is that municipal consolidation changes political participation more than it changes where people live.

That conclusion has two implications. First, the demographic stakes of Swiss mergers may be lower than both supporters and critics sometimes imply. Second, future work on local consolidation should distinguish more sharply between *voice* margins and *exit* margins. The former may react quickly; the latter, at least here, barely move.

Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Foreign population share	0.002	0.002	0.098	0.017	0.024	Small positive
Log foreign population	-0.007	0.017	1.766	-0.004	0.010	Null
Foreign population growth	1.799	1.221	16.692	0.108	0.073	Moderate positive
<i>Panel B: Heterogeneous</i>						
Foreign share, high baseline foreign share	0.001	0.004	0.082	0.011	0.053	Small positive
Foreign share, low baseline foreign share	0.005	0.004	0.034	0.135	0.104	Moderate positive

Notes: **Country:** Switzerland. **Research question:** Do voluntary municipal mergers change foreign residential composition in Swiss municipalities? **Policy mechanism:** Mergers replace multiple local jurisdictions with a single successor municipality, potentially altering tax-service bundles and local identity without directly targeting immigration policy. **Outcome definition:** Municipal foreign population share, log foreign population, and annual foreign population growth derived from BFS population-by-citizenship statistics. **Treatment:** Binary; indicator for years after the municipality’s first merger between 2015 and 2020. **Data:** BFS AGVCH merger registry and BFS PXWeb population by citizenship table, 2010–2024, harmonized to current boundaries, $N = 27,675$ municipality-year observations across 1845 municipalities. **Method:** Difference-in-differences with municipality and year fixed effects; pooled rows use the preferred Sun-Abraham ATT, heterogeneous rows use split-sample TWFE. **Sample:** Municipalities first merging between 2015 and 2020 and never-merged controls; heterogeneity splits municipalities at the median 2014 foreign share. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

References

- Allers, Maarten A. and J. Bieuwe Geertsema, “The Effects of Local Government Amalgamation on Public Spending, Taxation, and Service Levels: Evidence from 15 Years of Municipal Consolidation,” *Journal of Regional Science*, 2016, 56 (4), 659–682.
- Blom-Hansen, Jens, Kurt Houlberg, Soren Serritzlew, and Daniel Treisman, “Municipal Amalgamations and the Repercussions for Local Civic Participation: A Danish Study,” *Political Studies*, 2016, 64 (4), 796–814.
- Callaway, Brantly and Pedro H. C. Sant’Anna, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Frey, Timothee, “Do Municipal Mergers Reduce Voter Turnout? Evidence from Switzerland,” *Swiss Political Science Review*, 2023, 29 (2), 175–195.
- Goodman-Bacon, Andrew, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, 225 (2), 254–277.

- Koch, Philippe and Philippe E. Rochat**, “The Effects of Municipal Mergers on Turnout: Evidence from the Swiss Canton of Glarus,” *Swiss Political Science Review*, 2017, 23 (4), 398–420.
- Ladner, Andreas and Marc Buhlmann**, “Municipal Amalgamations and Their Effects: A Review of Evidence from Switzerland,” *Jahrbuch für Regionalwissenschaft*, 2009, 29, 155–175.
- Oates, Wallace E.**, *Fiscal Federalism*, New York: Harcourt Brace Jovanovich, 1972.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Roth, Jonathan, Pedro H. C. Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, 235 (2), 2218–2244.
- Steiner, Reto and Claire Kaiser**, “Assessing Municipal Amalgamation: A Focus on Switzerland,” *European Urban and Regional Studies*, 2017, 24 (2), 167–186.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.
- Swianiewicz, Pawel**, “If Territorial Fragmentation Is a Problem, Is Amalgamation a Solution? An East European Perspective,” *Local Government Studies*, 2010, 36 (2), 183–203.
- Tavares, Antonio F.**, “Joining Forces? Factors Behind Consensual and Non-Consensual Municipal Amalgamations,” *Territory, Politics, Governance*, 2018, 6 (4), 494–514.
- Tiebout, Charles M.**, “A Pure Theory of Local Expenditures,” *Journal of Political Economy*, 1956, 64 (5), 416–424.

Acknowledgements

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

Contributors: @olafdrw

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

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