The Ballot Box Rejection: Does Local Anti-Immigration Sentiment Alter Municipal Demographics? Evidence from Switzerland's 2014 Mass Immigration Initiative

APEP Autonomous Research* @olafdrw

March 31, 2026

Abstract

Switzerland's 2014 Mass Immigration Initiative passed with 50.3% of the national vote, creating sharp variation in municipal vote margins. I exploit this variation in a regression discontinuity design across 2,098 municipalities to test whether local expression of anti-immigration sentiment affects subsequent demographic composition. Municipalities narrowly voting Yes experienced 0.79 percentage points less foreign population share growth over 2015–2018 compared to those narrowly voting No (bias-corrected, $p = 0.078$). Total population growth shows no discontinuity. The effect survives alternative bandwidth, kernel, and covariate specifications. However, a pre-treatment placebo test shows a marginally insignificant baseline differential ($p = 0.103$), and placebo cutoffs at 40% and 45% produce effects, suggesting a gradient rather than a sharp discontinuity. The evidence is suggestive that democratic expression of anti-immigration preferences generates sorting of foreign residents across municipalities.

JEL Codes: J15, J61, D72, R23

Keywords: immigration, regression discontinuity, direct democracy, residential sorting, Switzerland

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 32m).

1. Introduction

On February 9, 2014, Swiss voters approved the Mass Immigration Initiative by 50.3% to 49.7%—a margin of roughly 19,500 votes out of 2.9 million cast. The initiative, launched by the Swiss People's Party (SVP), mandated annual quotas on immigration and threatened Switzerland's bilateral agreements with the European Union on free movement of persons. It was the sharpest democratic rebuke of open borders in modern European history. But the vote was not merely a national event. Across Switzerland's roughly 2,100 municipalities, yes-shares ranged from 19% to 94%, creating enormous local variation in the intensity of anti-immigration sentiment revealed by the ballot.

This paper asks a simple question: did this variation matter? Specifically, did municipalities that narrowly voted in favor of the initiative—thereby revealing majority anti-immigration sentiment at the most local level of government—subsequently experience different patterns of foreign population growth than otherwise similar municipalities that narrowly voted against it?

The question sits at the intersection of two literatures that have developed largely in parallel. A large body of work documents what shapes attitudes toward immigration (Dustmann and Preston, 2007; Hainmueller and Hopkins, 2014; Mayda, 2006; Card et al., 2012; Alesina et al., 2023), while a separate literature estimates the labor market and demographic consequences of immigration itself (Borjas, 2003; Card, 2001; Peri and Sparber, 2009; Foged and Peri, 2016; Dustmann et al., 2013). What remains underexplored is whether the *expression* of anti-immigration sentiment—through formal democratic channels—generates real demographic consequences, even when the vote outcome is determined at a higher level of government. If it does, the mechanism must operate through local sorting: immigrants choosing to avoid places that voted against them, or employers and landlords in high-yes municipalities behaving differently toward foreign workers and residents.

I exploit the 50% vote-share threshold in a sharp regression discontinuity design (RDD). The running variable is the municipal yes-share on the MII, centered at 50%. Municipalities just above the threshold revealed a local anti-immigration majority; those just below did not. Because no individual municipality's vote determines federal policy—the quota mandate applied nationally—the local vote margin is a signal of sentiment, not a direct policy instrument. The RDD identifies whether crossing the threshold of majority anti-immigration sentiment causes differential changes in foreign population composition.

The main finding is suggestive: municipalities narrowly voting Yes experienced approximately 0.79 percentage points less foreign population share growth over 2015–2018 relative to the pre-period (2010–2013), compared to municipalities narrowly voting No (bias-corrected

estimate, robust $p = 0.078$). The conventional estimate is -0.71 pp ($p = 0.059$). Total population growth shows no corresponding discontinuity ($p = 0.70$), indicating that the effect is specific to the foreign-born population rather than reflecting general demographic divergence.

The design passes standard validity checks. A McCrary density test finds no evidence of manipulation at the 50% threshold ($p = 0.40$), as expected given that municipal vote shares result from aggregating thousands of individual secret ballots. Covariate balance is clean: pre-treatment population, foreign share, voter turnout, and electorate size all show statistically insignificant differences at the threshold (all $p > 0.1$). The point estimate is robust to alternative kernels (Epanechnikov, uniform), bandwidth choices, and the inclusion of pre-treatment covariates.

I present these results with appropriate caution for two reasons. First, a pre-treatment placebo test—estimating the RDD on the change in foreign share from 2010–2011 to 2012–2013, before the vote occurred—yields a point estimate of -0.35 with $p = 0.103$. While not statistically significant, this baseline differential raises the possibility that municipalities trending toward Yes votes were already experiencing slower foreign population growth. Second, placebo cutoffs at 40% and 45% produce marginally significant effects, while those at 55% and 60% do not. This pattern suggests a gradient in the relationship between anti-immigration sentiment and demographic change, rather than a sharp discontinuity precisely at 50%. Whether the gradient is causal or reflects pre-existing sorting is a question this design cannot fully resolve.

This paper contributes to several literatures. First, it adds to the growing body of work on the real-world consequences of anti-immigration voting (Hangartner et al., 2019; Barone et al., 2016; Halla et al., 2017; Steinmayr, 2021; Tabellini, 2020). Most prior work examines how immigration affects voting; I study the reverse channel. Second, it contributes to the close-vote RDD literature (Lee, 2008; Lee and Lemieux, 2010), applying the design to a referendum setting where municipal vote shares provide unusually clean running variables. Third, it speaks to the literature on immigration and sorting (Card, 2001; Amuedo-Dorantes and Arenas-Arroyo, 2021), documenting that the revealed preferences of native communities may influence where immigrants settle, even absent direct local policy authority over immigration.

The Swiss setting offers several advantages. Switzerland's system of direct democracy produces frequent, binding referenda with granular municipal-level data on vote shares (Helbling and Kriesi, 2014). The free-movement agreement with the EU meant that, prior to the MII, EU citizens could settle freely across municipalities—making any post-vote sorting response especially telling. And the near-50/50 national result created substantial variation around the threshold, with 674 municipalities voting below 50% and 1,424 above.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of the MII. Section 3 presents the data. Section 4 lays out the empirical strategy. Section 5 reports results. Section 6 discusses implications and limitations.

2. Institutional Background

2.1 The Mass Immigration Initiative

Switzerland's relationship with the European Union is governed by a series of bilateral agreements, the most significant of which—the Agreement on the Free Movement of Persons (AFMP)—took effect in stages between 2002 and 2007. Under the AFMP, EU and EFTA nationals gained the right to live and work in Switzerland, subject to transitional safeguards that expired for EU-15 and EFTA states in 2007 and for EU-8 states in 2011 (Beerli et al., 2021). By 2013, foreign nationals constituted roughly 23% of Switzerland's permanent resident population, one of the highest shares in Europe.

The Mass Immigration Initiative (*Masseneinwanderungsinitiative*), launched by the Swiss People's Party (SVP), proposed a constitutional amendment requiring the Swiss government to set annual quantitative limits on all categories of immigration, including EU free movement. Critically, the initiative's text included a "guillotine clause": if the immigration caps could not be reconciled with the AFMP within three years, the Federal Council was required to renegotiate or abrogate the agreement. The initiative thus threatened not only immigration flows but Switzerland's entire bilateral framework with the EU (Helbling and Kriesi, 2014).

The vote took place on February 9, 2014. Turnout was 55.8%—above the Swiss average for federal votes. The initiative passed by the thinnest possible margin: 50.3% yes, with 19,526 votes separating the camps. The canton-by-canton pattern was stark. Urban, French-speaking, and border cantons voted heavily against; rural, German-speaking, and interior cantons voted in favor. But within cantons, substantial municipal variation existed, driven by local demographic composition, economic structure, and political orientation.

2.2 Implementation and the Three-Year Window

The MII's passage created a three-year implementation window (2014–2017) during which Switzerland was constitutionally obligated to impose immigration quotas but had not yet legislated the details. This period was characterized by intense policy uncertainty. The EU refused to renegotiate the AFMP, and the Federal Council faced the prospect of invoking the guillotine clause—which would have terminated not only free movement but also six other bilateral agreements governing trade, research, and transport.

In practice, the quotas were never implemented as the initiative envisioned. After extensive parliamentary debate, Switzerland adopted a "light" implementation in December 2016: a domestic hiring preference (*Inländervorrang light*) requiring employers to register vacancies with public employment services in occupations with unemployment rates above 8%, later lowered to 5%. The EU accepted this compromise, and the AFMP was preserved. But for the three years following the vote, the policy signal was ambiguous: Switzerland had voted to restrict immigration, the mechanism was undefined, and the outcome was uncertain.

This uncertainty is central to the paper's mechanism. If the MII vote affected municipal demographics, it did so not through direct policy implementation—which was uniform across municipalities—but through the *signal* the local vote sent about community attitudes. A municipality that voted 55% yes sent a different signal about its receptiveness to foreigners than one that voted 45% yes, even though both were subject to identical federal policies.

3. Data

I combine two data sources at the municipal level. Vote data come from the Federal Statistical Office (BFS) via the `swissdd` R package, which provides municipal-level results for all federal votes. For the MII, this includes the number of yes votes, no votes, valid ballots, eligible voters, and turnout for each municipality. Population data come from the BFS *Statistik der Bevölkerung und der Haushalte* (STATPOP), accessed through the PXWeb API. STATPOP provides annual counts of permanent residents by citizenship (Swiss vs. foreign) at the municipal level from 2010 to 2019.

I match municipalities across the two datasets using official BFS municipality identifiers. After dropping municipalities with missing vote or population data, the analysis sample comprises 2,098 municipalities—approximately 85% of all municipalities in 2014. The primary running variable is the municipal yes-share, centered at 50%. The primary outcome is the change in the municipal foreign population share (as a percentage of total population) from the pre-period (mean of 2010–2013) to the post-period (mean of 2015–2018).

3.1 Summary Statistics

Table 1 presents summary statistics separately for municipalities above and below the 50% threshold. The two groups differ substantially in raw averages: yes-share municipalities are smaller (mean population 2,808 vs. 5,758), have lower pre-treatment foreign shares (13.1% vs. 17.7%), and lower turnout (56.9% vs. 60.5%). These differences are expected—larger, more urban, and more cosmopolitan municipalities tend to vote against immigration restrictions—and underscore the importance of the RDD design, which compares only municipalities near

Table 1: Summary Statistics by MII Vote Outcome

	Below 50%		Above 50%		
	Mean	SD	Mean	SD	Diff.
Yes share (%)	42.4	5.8	61.4	7.4	19.06
Turnout (%)	60.5	6.4	56.9	6.2	-3.62
Eligible voters	3555.1	11882.1	1914.5	2565.9	-1640.63
Population (pre)	5758.2	20083.9	2808.0	4022.6	-2950.21
Foreign share, pre (%)	17.7	10.2	13.1	8.6	-4.58
Δ Foreign share (pp)	2.0	2.0	2.0	1.8	-0.00
Δ Population (%)	6.8	7.4	5.0	6.1	-1.77
Municipalities	674		1424		

Notes: Summary statistics by whether the municipality voted above or below 50% Yes on the Mass Immigration Initiative (February 9, 2014). "Pre" statistics are averaged over 2010–2013. Δ Foreign share is the change in the municipal foreign population share (percentage points) from the pre-period (2010–2013) to the post-period (2015–2018). Population data from BFS STATPOP; vote data from the Federal Statistical Office via swissdd.

the threshold.

Critically, the mean change in foreign share from pre to post is nearly identical across the two groups: 2.0 percentage points for both. The identifying variation comes not from this unconditional comparison but from the local discontinuity at 50%.

4. Empirical Strategy

4.1 Identification

I estimate a reduced-form sharp regression discontinuity design with the municipal yes-share as the running variable and 50% as the cutoff.[1] The treatment is not a policy change—the federal policy outcome does not depend on any individual municipality's vote—but rather the public revelation of majority anti-immigration sentiment:

$$Y_m = \alpha + \tau \cdot \mathbb{I}[S_m \geq 0.5] + f(S_m - 0.5) + \varepsilon_m \tag{1}$$

where Y_m is the outcome for municipality m, S_m is the yes-share, $\mathbb{I}[S_m \geq 0.5]$ indicates a local yes-majority, $f(\cdot)$ is a local polynomial estimated separately on each side of the cutoff, and τ is the parameter of interest: the effect of crossing the majority-yes threshold on subsequent demographic change.

[1]The original research design also envisioned a canton-level panel analysis of cross-border commuter (Grenzgänger) flows using quarterly BFS data. I defer this complementary analysis to future work, focusing here on the municipal-level RDD where treatment and outcome are measured at the same geographic unit.

The identifying assumption is continuity of potential outcomes at the cutoff:

$$\lim_{s \downarrow 0.5} \mathbb{E}[Y_m(0)|S_m = s] = \lim_{s \uparrow 0.5} \mathbb{E}[Y_m(0)|S_m = s] \tag{2}$$

This requires that municipalities just above and just below 50% would have experienced the same demographic trends in the absence of crossing the threshold. The assumption is plausible here because municipal vote shares aggregate thousands of individual secret ballots; precise manipulation of the aggregate share to land just above or below 50% is infeasible (McCrary, 2008; Lee, 2008).

4.2 Estimation

All specifications use the `rdrobust` package (Calonico et al., 2014; Cattaneo et al., 2020b) with MSE-optimal bandwidth selection (Imbens and Kalyanaraman, 2012). The baseline specification uses a local linear polynomial with a triangular kernel. I report both conventional and bias-corrected estimates with robust standard errors, following Cattaneo et al. (2020b). Standard errors are clustered at the canton level (26 cantons) to account for spatial correlation in both vote patterns and demographic trends. I follow Gelman and Imbens (2019) in avoiding high-order polynomials.

4.3 Interpretation

A key feature of this design is that crossing the 50% threshold does not change any policy at the municipal level. Federal immigration policy is determined by the national vote, not by individual municipalities. The local vote margin therefore captures revealed sentiment: it is the municipality's public, democratic statement about immigration. Any effect of crossing the threshold must operate through this signal—through the behavior of immigrants, employers, landlords, or municipal authorities responding to the knowledge that their community voted for or against restricting immigration.

This interpretation places the design closer to the "close election" RDD of Lee (2008) than to a standard policy discontinuity. The treatment is the *revelation of majority anti-immigration sentiment*, not the implementation of a policy. This means the effect, if it exists, is a reduced-form composite: it could reflect immigrant sorting (avoiding hostile communities), employer behavior (differential hiring or sponsorship), informational effects (salience of local attitudes), or some combination.

4.4 Threats to Validity

Manipulation. The McCrary density test, reported in Table 4, yields $p = 0.40$, providing no evidence of bunching at the threshold. Strategic manipulation of municipal vote shares would require coordination among thousands of voters, which is neither feasible nor incentivized: no municipal-level policy hinges on crossing 50%.

Covariate balance. Table 4 Panel A tests whether pre-determined covariates—population, foreign share, turnout, and electorate size—are smooth through the threshold. All four tests are statistically insignificant at the 10% level, consistent with local quasi-random assignment.

Pre-treatment trends. The placebo outcome test in Table 4 Panel C estimates the RDD on the change in foreign share from 2010–2011 to 2012–2013—before the vote occurred. The point estimate is -0.35 ($p = 0.103$). While not significant at conventional levels, this pre-treatment differential is large enough relative to the main effect (-0.79) to warrant caution. It suggests that municipalities that would go on to vote Yes may have already been experiencing somewhat slower foreign population growth, potentially reflecting the same underlying attitudes that drove the vote.

5. Results

5.1 Main Results

Table 2 presents the main RDD estimates for three outcomes: change in foreign population share (percentage points), change in total population (percent), and change in foreign population (percent). The preferred specification—bias-corrected with robust standard errors (Panel B)—estimates that municipalities narrowly voting Yes experienced a 0.79 pp reduction in foreign share growth relative to narrowly-No municipalities, significant at the 10% level ($p = 0.078$). The conventional estimate (Panel A) is -0.71 pp ($p = 0.059$).

To put this in context, the mean change in foreign share across all municipalities was approximately 2.0 pp. The estimated discontinuity thus represents roughly 40% of the average change—a large effect, though imprecisely estimated. The MSE-optimal bandwidth is 9.0 pp, yielding an effective sample of 1,028 municipalities.

The second column of Table 2 tests for effects on total population growth. The estimate is small (-0.74%) and statistically insignificant ($p = 0.70$ in the robust specification). This null result is informative: it rules out the possibility that the foreign share effect is driven by differential native in-migration or overall demographic divergence. The composition of the population shifts, but its size does not.

Table 2: RDD Estimates: Effect of Local MII Majority on Demographic Change

	Δ Foreign Share (pp)	Δ Total Population (%)	Δ Foreign Population (%)
Panel A: Conventional			
Above 50% Yes	-0.707*	-0.434	-7.444
	(0.375)	(1.741)	(6.508)
Panel B: Bias-Corrected, Robust SE			
Above 50% Yes	-0.788*	-0.740	-8.474
	(0.448)	(1.938)	(7.866)
Bandwidth (pp)	9.0	10.9	9.8
Effective N	1028	1240	1116
Kernel	Triangular	Triangular	Triangular
Polynomial	Linear	Linear	Linear
Clustered SE	Canton	Canton	Canton

Notes: Local polynomial RDD estimates using `rdrobust` with MSE-optimal bandwidth selection. The running variable is the municipal yes-share on the February 2014 Mass Immigration Initiative, centered at 50%. Panel A reports conventional estimates; Panel B reports bias-corrected point estimates with robust standard errors (Cattaneo, Idrobo, and Titiunik 2020). Standard errors clustered at the canton level in parentheses. Δ Foreign share is the change in the municipal foreign population share (pp) from 2010–2013 to 2015–2018. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The third column examines the change in the foreign population itself (in percent). The point estimate (−8.47%) is large but imprecisely estimated ($p = 0.28$), reflecting the high variance of percentage changes in small municipalities.

5.2 Robustness

Table 3 examines the sensitivity of the main result to specification choices. Panel A compares the baseline estimate to a specification including pre-treatment covariates (log population, pre-period foreign share, and turnout). The covariate-adjusted estimate (−0.77 pp, $p = 0.082$) is nearly identical to the baseline, consistent with the clean covariate balance documented in Table 4.

Panel B varies the kernel. The Epanechnikov kernel produces a similar estimate (−0.78 pp, $p = 0.073$), while the uniform kernel produces a somewhat larger estimate (−0.87 pp, $p = 0.065$). The consistency across kernels is reassuring, as each weights observations differently within the bandwidth.

Panel C tests a quadratic polynomial. The point estimate increases in absolute value (−0.87 pp) but becomes less precise ($p = 0.107$), as expected with added flexibility (Gelman and Imbens, 2019). The sign and magnitude remain consistent.

Table 3: Robustness of the Foreign Share RDD Estimate

Specification	Coef.	Robust SE	p-value	BW (pp)
Panel A: Baseline and Covariates				
Baseline (triangular, $p = 1$)	-0.788*	0.448	0.078	9.0
With covariates	-0.767*	0.441	0.082	—
Panel B: Alternative Kernels				
Triangular	-0.788*	0.448	0.078	9.0
Epanechnikov	-0.776*	0.433	0.073	8.5
Uniform	-0.868*	0.470	0.065	6.2
Panel C: Polynomial Order				
Order $p = 1$	-0.788*	0.448	0.078	9.0
Order $p = 2$	-0.866	0.537	0.107	12.8

Notes: All specifications report bias-corrected RDD estimates with robust standard errors. Covariates include log pre-period population, pre-period foreign share, and voter turnout. MSE-optimal bandwidth selection throughout. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5.3 Validity and Placebo Tests

Table 4 consolidates the design validation exercises. Panel A confirms covariate balance at the threshold: pre-treatment population ($p = 0.52$), foreign share ($p = 0.20$), turnout ($p = 0.13$), and electorate size ($p = 0.67$) all pass smoothness tests. Panel D reports the McCrary density test ($p = 0.40$).

Panel B reports placebo cutoff tests at 40%, 45%, 55%, and 60%. The results at 55% and 60% show no effect ($p = 0.40$ and $p = 0.88$, respectively), consistent with the absence of a discontinuity away from the true threshold. However, the tests at 40% ($p = 0.084$) and 45% ($p = 0.051$) are marginally significant. These results suggest that the relationship between anti-immigration sentiment and foreign population growth may be a gradient across the vote-share distribution rather than a sharp break at 50%. The effect at 45% is positive—municipalities just above 45% grew faster than those just below—which may reflect the same underlying sorting pattern viewed from a different point on the curve.

Panel C reports the pre-treatment placebo outcome. As discussed, the estimate of -0.35 ($p = 0.103$) is not significant but cannot be dismissed. I interpret this as a caution flag rather than a fatal flaw: the pre-treatment differential is less than half the post-treatment effect, and it is not statistically distinguishable from zero. But it prevents a clean causal interpretation. The main result should be read as suggestive evidence of a sorting response, not a definitive estimate.

Table 4: Validity: Covariate Balance, Placebo Cutoffs, and Placebo Outcome

	Coef.	Robust SE	p-value	Test
Panel A: Covariate Balance at Threshold				
Population (pre)	-581.179	1125.953	0.523	Balance
Foreign share (pre, %)	-1.747	1.601	0.200	Balance
Turnout (%)	-1.388	0.940	0.128	Balance
Eligible voters	-218.474	705.255	0.673	Balance
Panel B: Placebo Cutoffs				
Cutoff at 40%	-0.890*	0.514	0.084	Placebo
Cutoff at 45%	0.659*	0.338	0.051	Placebo
Cutoff at 55%	0.246	0.291	0.398	Placebo
Cutoff at 60%	0.034	0.221	0.876	Placebo
Panel C: Placebo Outcome				
Δ Foreign share (2010–2013)	-0.346	0.212	0.103	Pre-trend
Panel D: Manipulation Test				
McCrary density	—	—	0.398	Density

Notes: Panel A tests whether pre-determined covariates are smooth through the 50% threshold using RDD estimation. Panel B tests for effects at false cutoffs (40%, 45%, 55%, 60%). Panel C uses the pre-treatment change in foreign share (2010–2011 to 2012–2013) as a placebo outcome. Panel D reports the McCrary (2008)/Cattaneo, Jansson, and Ma (2020) density manipulation test. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5.4 Mechanisms

The finding that foreign share declines at the threshold while total population does not constrains the set of plausible mechanisms. Three channels are consistent with the evidence.

First, *immigrant sorting*: foreign residents and prospective immigrants may avoid municipalities that voted Yes, either because the vote outcome is directly observable (all municipal results are public) or because the underlying attitudes manifest in everyday interactions. This is the most direct interpretation and is consistent with the broader literature on residential sorting and social attitudes (Card, 2001; Amuedo-Dorantes and Arenas-Arroyo, 2021).

Second, *employer behavior*: in municipalities with revealed anti-immigration majorities, employers may be less willing to hire or sponsor foreign workers, whether out of shared preferences, concern about local political dynamics, or anticipation of future restrictions. Beerli et al. (2021) document that the AFMP's liberalization substantially affected employer behavior; the MII vote may have partially reversed these effects locally.

Third, *information revelation*: the vote may have updated immigrants' beliefs about the

local environment. Prior to the vote, a foreign resident in a given municipality may not have known the distribution of attitudes toward immigration. The vote made this information public, precise, and salient. If immigrants value living in communities where they feel welcome (Dustmann and Preston, 2007; Facchini and Mayda, 2008), the vote could trigger selective out-migration from high-yes municipalities.

These channels are not mutually exclusive, and the data do not permit definitive adjudication among them. All three predict the observed pattern—reduced foreign share with no effect on total population—and all three are consistent with the marginal significance of the estimate, which suggests a modest rather than dramatic behavioral response.

6. Discussion

The central finding of this paper is that municipalities narrowly crossing the threshold of majority support for Switzerland's Mass Immigration Initiative subsequently experienced slower growth in their foreign population share. The effect is economically meaningful—roughly 0.8 percentage points, or 40% of the mean change—but statistically marginal, with p-values between 0.06 and 0.08 depending on specification. The pre-treatment placebo and placebo-cutoff results further temper causal claims.

What should we make of suggestive evidence? The result is consistent with a growing literature documenting that anti-immigration politics has real consequences beyond the policies it produces (Hangartner et al., 2019; Tabellini, 2020; Steinmayr, 2021). Even when the MII's federal policy consequences were ultimately modest—the "light" implementation preserved free movement—the local vote may have served as a coordination device, making anti-immigration sentiment legible and salient in ways that affected individual location decisions.

The null result on total population deserves emphasis. It would be natural to expect that if foreign residents leave (or avoid) high-yes municipalities, native residents might flow in to fill vacancies. That this does not happen suggests the effect operates on the margin of new arrivals rather than through population replacement—a finding more consistent with sorting of incoming immigrants than with displacement or flight.

The placebo cutoff results raise a broader methodological point. If the relationship between anti-immigration sentiment and demographic change is a gradient rather than a discontinuity, the RDD estimates a local effect that may not be qualitatively distinct from the pattern observed across the full distribution. This does not invalidate the design—the continuity assumption can hold even when the treatment effect is part of a broader gradient (Lee and Lemieux, 2010)—but it changes the interpretation. The "treatment" is not a discrete

regime change but rather a point along a continuous dose-response relationship between community sentiment and immigrant sorting. Indeed, the gradient interpretation may be the more natural one: there is no institutional reason why 50% should mark a discrete behavioral break at the municipal level, since municipal vote outcomes carry no local policy consequences. The 50% threshold is analytically convenient but not institutionally privileged, and readers should interpret the estimates accordingly.

A related limitation is that the foreign population share outcome confounds several margins of adjustment: in-migration, out-migration, naturalization, differential fertility, and denominator effects from Swiss population changes. The null result on total population growth rules out some of these channels, but a full decomposition—separating foreign arrivals from departures and naturalizations—would require administrative data on migration flows that the BFS STATPOP stock data cannot provide.

The Swiss setting is both an advantage and a limitation. Switzerland's direct democratic institutions, granular municipal data, and pre-MII regime of free movement make the identification strategy feasible and the mechanism interpretable. But these same features limit external validity. Countries without binding referenda on immigration, without free movement agreements, or without Switzerland's particular combination of economic prosperity and cultural diversity may exhibit different patterns. The finding that democratic expression of anti-immigration sentiment affects sorting is most directly relevant to other contexts where local attitudes are publicly revealed—through elections, protests, or policy statements—and where immigrants have freedom to choose among locations.

7. Conclusion

When 50.3% of Swiss voters said they wanted fewer immigrants, they did not merely register a preference. They created a public, municipality-by-municipality map of anti-immigration sentiment. The evidence in this paper suggests—tentatively—that this map mattered: municipalities on the anti-immigration side of the threshold subsequently attracted fewer foreign residents than those on the other side.

The deeper implication concerns the relationship between democratic expression and social sorting. Referenda do not merely implement policy; they produce information. A municipality's vote on the Mass Immigration Initiative told prospective immigrants something about the community's character that may not have been knowable before. If this information affects where people choose to live, then direct democracy has demographic consequences that extend beyond the policies it enacts. The ballot box does not just decide; it sorts.

Whether this sorting is efficient or equitable is a question this paper cannot answer. But

it is a question that matters increasingly as immigration becomes the dominant cleavage in democratic politics across Europe and beyond (Hainmueller and Hopkins, 2014; Rydgren, 2008). Understanding how democratic expression of immigration preferences feeds back into the spatial distribution of immigrants is essential for evaluating both the politics and the economics of immigration in open societies.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data were obtained from the Swiss Federal Statistical Office (BFS) via the PXWeb API and the `swissdd` R package.

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

Contributors: @olafdrw

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

References

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva, "Immigration and Redistribution," *Review of Economic Studies*, 2023, *90* (1), 1–39.

Amuedo-Dorantes, Catalina and Esther Arenas-Arroyo, "Immigration Enforcement and Children's Living Arrangements," *Journal of Policy Analysis and Management*, 2021, *40* (4), 1125–1153.

Barone, Guglielmo, Alessio D'Ignazio, Guido de Blasio, and Paolo Naticchioni, "Mr. Rossi, Mr. Hu and Politics: The Role of Immigration in Shaping Natives' Voting Behavior," *Journal of Public Economics*, 2016, *136*, 1–13.

Beerli, Andreas, Jens Rüfli, Michael Siegenthaler, and Giovanni Peri, "The Abolition of Immigration Restrictions and the Performance of Firms and Workers: Evidence from Switzerland," *American Economic Review*, 2021, *111* (3), 976–1012.

Borjas, George J., "The Labor Demand Curve Is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market," *Quarterly Journal of Economics*, 2003, *118* (4), 1335–1374.

Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 2014, *82* (6), 2295–2326.

Card, David, "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration," *Journal of Labor Economics*, 2001, *19* (1), 22–64.

_ , Christian Dustmann, and Ian Preston, "Immigration, Wages, and Compositional Amenities," *Journal of the European Economic Association*, 2012, *10* (1), 78–119.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma, "Simple Local Polynomial Density Estimators," *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.

_ , Nicolás Idrobo, and Rocío Titiunik, *A Practical Introduction to Regression Discontinuity Designs: Foundations*, Cambridge University Press, 2020.

Dustmann, Christian and Ian P. Preston, "Racial and Economic Factors in Attitudes to Immigration," *The B.E. Journal of Economic Analysis & Policy*, 2007, *7* (1), 1–41.

_ , **Tommaso Frattini, and Ian P. Preston**, "The Effect of Immigration along the Distribution of Wages," *Review of Economic Studies*, 2013, *80* (1), 145–173.

Facchini, Giovanni and Anna Maria Mayda, "From Individual Attitudes towards Migrants to Migration Policy Outcomes: Theory and Evidence," *Economic Policy*, 2008, *23* (56), 651–713.

Foged, Mette and Giovanni Peri, "Immigrants' Effect on Native Workers: New Analysis on Longitudinal Data," *American Economic Journal: Applied Economics*, 2016, *8* (2), 1–34.

Gelman, Andrew and Guido Imbens, "Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs," *Journal of Business & Economic Statistics*, 2019, *37* (3), 447–456.

Hainmueller, Jens and Daniel J. Hopkins, "Public Attitudes Toward Immigration," *Annual Review of Political Science*, 2014, *17*, 225–249.

Halla, Martin, Alexander F. Wagner, and Josef Zweimüller, "Immigration and Voting for the Far Right," *Journal of the European Economic Association*, 2017, *15* (6), 1341–1385.

Hangartner, Dominik, Elias Dinas, Moritz Marbach, Konstantinos Matakos, and Dimitrios Xefteris, "Does Exposure to the Refugee Crisis Make Natives More Hostile?," *American Political Science Review*, 2019, *113* (2), 442–455.

Helbling, Marc and Hanspeter Kriesi, "Why Citizens Prefer High- over Low-Skilled Immigrants: Labor Market Competition, Welfare State, and Deservingness," *European Societies*, 2014, *16* (2), 207–226.

Imbens, Guido and Karthik Kalyanaraman, "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *Review of Economic Studies*, 2012, *79* (3), 933–959.

Lee, David S., "Randomized Experiments from Non-Random Selection in U.S. House Elections," *Journal of Econometrics*, 2008, *142* (2), 675–697.

_ **and Thomas Lemieux**, "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 2010, *48* (2), 281–355.

Mayda, Anna Maria, "Who Is Against Immigration? A Cross-Country Investigation of Individual Attitudes Toward Immigrants," *Review of Economics and Statistics*, 2006, *88* (3), 510–530.

McCrary, Justin, "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 2008, *142* (2), 698–714.

Peri, Giovanni and Chad Sparber, "Task Specialization, Immigration, and Wages," *American Economic Journal: Applied Economics*, 2009, *1* (3), 135–169.

Rydgren, Jens, "Immigration Sceptics, Xenophobes or Racists? Radical Right-Wing Voting in Six West European Countries," *European Journal of Political Research*, 2008, *47* (6), 737–765.

Steinmayr, Andreas, "Contact versus Exposure: Refugee Presence and Voting for the Far Right," *Review of Economics and Statistics*, 2021, *103* (2), 310–327.

Tabellini, Marco, "Gifts of the Immigrants, Woes of the Natives: Lessons from the Age of Mass Migration," *Review of Economic Studies*, 2020, *87* (1), 454–486.

A. Data Appendix

Vote data. Municipal-level results for the Mass Immigration Initiative (February 9, 2014) were obtained via the `swissdd` R package (see Helbling and Kriesi, 2014), which queries the Federal Statistical Office's (BFS) open data portal. The dataset includes, for each municipality: the number of yes votes, no votes, valid and invalid ballots, eligible voters, and voter turnout. The running variable is computed as $S_m = \text{Yes}_m/(\text{Yes}_m + \text{No}_m)$, centered at 0.5.

Population data. Annual municipal population by citizenship (Swiss nationals vs. foreign nationals) comes from the BFS *Statistik der Bevölkerung und der Haushalte* (STATPOP), accessed via the PXWeb API for years 2010–2019. The unit of observation is the municipality-year. I use the permanent resident population (*ständige Wohnbevölkerung*).

Sample construction. I begin with 2,212 municipalities in the BFS vote records for the MII. I match these to STATPOP population data using official BFS municipality identifiers, yielding 2,098 matched municipalities (94.8%). The 114 unmatched municipalities are predominantly very small communes that either merged during the period or had incomplete population records. The analysis file is a municipality-level cross-section, with pre-period variables averaged over 2010–2013 and post-period variables averaged over 2015–2018.

Key variable definitions.

- **Yes margin** ($S_m - 0.5$): the running variable, measuring the municipal yes-share minus 50%. Positive values indicate a local yes-majority.
- **Δ Foreign share**: the change in the municipal share of foreign nationals in total population, from the mean of 2010–2013 to the mean of 2015–2018, in percentage points.
- **Δ Total population**: the percentage change in total municipal population from the mean of 2010–2013 to the mean of 2015–2018.
- **Δ Foreign population**: the percentage change in the number of foreign nationals in the municipality from the mean of 2010–2013 to the mean of 2015–2018.

B. Identification Appendix

McCrary density test. I implement the density test of McCrary (2008) using the `rddensity` package of Cattaneo et al. (2020a). The test examines whether the density of the running

variable (municipal yes-share) exhibits a discontinuity at 50%, which would suggest strategic manipulation. The test statistic has a p-value of 0.40, providing no evidence of manipulation. This is expected: individual voters cast secret ballots, and no single voter or small group can precisely control the aggregate municipal share.

Covariate balance. I test for smoothness of four pre-determined covariates at the threshold using separate RDD regressions with each covariate as the outcome: pre-period population ($p = 0.52$), pre-period foreign share ($p = 0.20$), voter turnout ($p = 0.13$), and number of eligible voters ($p = 0.67$). All four are statistically insignificant, supporting the identifying assumption.

Placebo cutoffs. I re-estimate the main specification at four placebo cutoffs (40%, 45%, 55%, 60%). Effects at 55% and 60% are insignificant ($p = 0.40$ and $p = 0.88$). Effects at 40% and 45% are marginally significant ($p = 0.084$ and $p = 0.051$), though with opposite signs. This pattern is consistent with a gradient rather than a sharp break and is discussed in the main text.

C. Robustness Appendix

Bandwidth sensitivity. The MSE-optimal bandwidth for the main specification is 9.0 pp. I re-estimate using bandwidths at 50%, 75%, 100%, 125%, 150%, and 200% of the optimal. The coefficient remains negative across all bandwidths, ranging from approximately -0.5 to -0.9 pp. Narrower bandwidths yield larger point estimates but wider confidence intervals, consistent with the bias-variance tradeoff inherent in nonparametric estimation (Imbens and Kalyanaraman, 2012).

Polynomial order. Following Gelman and Imbens (2019), I restrict attention to linear and quadratic local polynomials. The quadratic specification yields a larger point estimate (-0.87 pp) but is less precise ($p = 0.107$). Higher-order polynomials are avoided due to well-documented over-fitting concerns.

Kernel choice. Results are robust across triangular (-0.79 pp), Epanechnikov (-0.78 pp), and uniform (-0.87 pp) kernels, with robust p-values between 0.065 and 0.078. The uniform kernel produces the largest estimate, consistent with its equal weighting of all observations within the bandwidth.

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Foreign share change (pp)	-0.79	0.45	1.86	-0.423	0.241	Large negative
Total pop. change (%)	-0.74	1.94	6.59	-0.112	0.294	Moderate negative
Foreign pop. change (%)	-8.47	7.87	37.06	-0.229	0.212	Large negative
Panel B: Heterogeneous (urban vs. rural)						
Foreign share (urban Q4)	-0.24	0.40	1.21	-0.202	0.334	Large negative
Foreign share (rural Q1)	-1.88	0.94	2.48	-0.761	0.379	Large negative

Notes: **Country:** Switzerland. **Research question:** Does a municipality narrowly voting in favor of the 2014 Mass Immigration Initiative (which mandated caps on EU immigration) cause a differential reduction in its foreign population share relative to a municipality narrowly voting against? **Policy mechanism:** The February 2014 initiative threatened to end the bilateral Free Movement of Persons Agreement with the EU, creating a three-year policy uncertainty window (2014–2017) during which municipalities with stronger anti-immigration mandates may have become less attractive to foreign residents and cross-border workers through local signaling, employer caution, or migrant sorting. **Outcome definition:** Change in the municipal foreign population share (percentage points), computed as the difference between the mean share in 2015–2018 and the mean share in 2010–2013, using BFS STATPOP permanent resident population by citizenship. **Treatment:** Binary indicator for the municipality's yes-share exceeding 50% on the Mass Immigration Initiative, with the municipal yes-share as a continuous forcing variable in a regression discontinuity design. **Data:** BFS STATPOP municipal population by citizenship (2010–2019) merged with municipal vote results from the Federal Statistical Office (2,098 municipalities), via the PXWeb API and swissdd R package. **Method:** Local polynomial RDD (rdrobust) with MSE-optimal bandwidth, triangular kernel, bias-corrected estimates with robust standard errors (Cattaneo, Idrobo, and Titiunik 2020), clustered at the canton level. **Sample:** 2,098 Swiss municipalities with matched vote and population data; effective sample within optimal bandwidth varies by specification (approximately 600–1,400 municipalities). SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the cross-sectional standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

D. Standardized Effect Sizes