From the Ballot Box to the Bureau: Judicial Abolition of Direct-Democratic Naturalization and Immigrant Incorporation in Switzerland

APEP Autonomous Research* @olafdrw

April 8, 2026

Abstract

In roughly one-third of Swiss municipalities, foreign residents seeking citizenship had to survive a public vote by their neighbors—a procedure the Federal Court struck down as unconstitutional in July 2003 (BGE 129 I 232). Using 44 years of municipality-level naturalization data (1981–2024), I exploit the ruling as a sharp institutional shock in a difference-in-differences design comparing former ballot municipalities to those already using administrative procedures. Controlling for canton-specific trends, the ruling raised naturalization rates by 4.7 per 1,000 foreign residents (0.11 standard deviations). The effect concentrates in small municipalities, where assembly-based gatekeeping was most salient. Placebo outcomes and leave-one-out tests confirm robustness. These results quantify the "procedural discrimination tax" levied by direct-democratic screening and show that institutional reform—not just attitudinal change—can meaningfully expand immigrant incorporation.

JEL Codes: J15, K37, D72, J61

Keywords: naturalization, immigration, direct democracy, discrimination, Switzerland, institutional reform

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 44m).

1. Introduction

When a Turkish carpenter in Zurich applied for Swiss citizenship before 2003, his fate did not rest with a bureaucrat reviewing his file. It rested with his neighbors, gathered in the Gemeindeversammlung, voting by show of hands on whether he deserved to become Swiss. His employment record, language skills, and years of residence were read aloud to the assembly. Then the vote. Hainmueller and Hangartner (2013) document that applicants from Turkey and the former Yugoslavia faced rejection rates roughly 40 percentage points higher than otherwise identical Western European applicants in these ballot municipalities—a gap that vanished in municipalities using anonymous administrative procedures.

On July 9, 2003, the Swiss Federal Court ended this practice. In BGE 129 I 232, the Court ruled that ballot-vote naturalization violated the constitutional rights to privacy (Article 13), a reasoned decision (Article 29(2)), and non-discrimination (Article 8(2)). The ruling required all municipalities to adopt written, anonymous administrative procedures. Overnight, approximately 900 municipalities in German-speaking Switzerland lost their most visible tool of exclusion.

This paper asks a simple question with a non-obvious answer: did removing the ballot actually increase naturalization? The answer matters because the policy debate over immigrant incorporation typically focuses on attitudes—prejudice, cultural anxiety, nativism—rather than institutions. If discriminatory procedures are the binding constraint, then institutional reform can expand citizenship access regardless of underlying attitudes. If attitudes are the binding constraint, then the ballot was merely a convenient expression of preferences that would find other outlets.

I exploit the 2003 ruling in a difference-in-differences design comparing municipalities in German-speaking cantons (where ballot naturalization was standard practice) to municipalities in French- and Italian-speaking cantons (where administrative procedures were already in place). The design leverages 44 years of annual, municipality-level naturalization data from the Swiss Federal Statistical Office (1981–2024), providing 22 pre-treatment years to assess parallel trends and 21 post-treatment years to trace long-run dynamics.

My preferred specification, which includes canton-specific linear time trends to absorb differential trajectory differences between language regions, estimates that abolishing ballot naturalization raised the annual naturalization rate by 4.7 per 1,000 foreign residents (standard error: 1.8, $p < 0.05$). This represents a 26 percent increase relative to the pre-treatment mean among ballot municipalities (18.1 per 1,000) and 0.11 standard deviations of the outcome distribution. At the aggregate level, the estimate implies roughly 2,350 additional naturalizations per year in formerly ballot cantons—2,350 residents who gained the right to

vote, hold a Swiss passport, and participate fully in their communities.

Three features of the results sharpen the interpretation. First, the reform mattered most where the scrutiny was most personal: the effect concentrates in small municipalities (below-median population: coefficient 7.6, $p < 0.003$) and is statistically indistinguishable from zero in larger municipalities (2.2, $p = 0.31$). In small communes, the assembly was intimate enough for personal knowledge of applicants to influence votes; larger municipalities had already delegated much decision-making to commissions even before the ruling. Second, placebo outcomes—foreign population share and total population growth—show no differential response to the ruling, ruling out confounding compositional changes. Third, the estimate is stable across leave-one-out exercises dropping each ballot canton in turn (range: 4.1 to 5.7), and HonestDiD sensitivity analysis (Rambachan and Roth, 2023) confirms that the effect remains strictly positive for parallel-trends violations up to $\bar{M} = 0.04$.

This paper contributes to three literatures. First, it extends the seminal work of Hainmueller and Hangartner (2013) and Hainmueller and Hangartner (2015) on discrimination in Swiss naturalization by shifting from individual-level rejection rates to aggregate municipal flows. While their designs identify discrimination conditional on applying, mine identifies the equilibrium effect on the total stock of naturalizations—including any changes in application behavior induced by the procedural reform. Second, it contributes to the growing literature on how institutional design shapes immigrant integration outcomes (Janoski, 2010; Howard, 2009; Vink and de Groot, 2010; Goodman, 2010). Much of this literature relies on cross-country comparisons of citizenship regimes, which conflate institutions with preferences. The within-country variation I exploit holds national policy and cultural context fixed while isolating the procedural channel. Third, it speaks to the political economy of direct democracy. Frey and Stutzer (2006) and Hangartner et al. (2019) argue that direct-democratic participation can enhance or degrade minorities' outcomes depending on institutional design. The ballot naturalization system is a canonical case where participation by the majority directly harmed a minority—and where judicial intervention produced measurable improvement.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting. Section 4 presents the data and treatment classification. Section 5 develops the identification strategy. Section 6 reports the main results, event study, and robustness checks. Section 7 discusses mechanisms and limitations. Section 8 concludes.

2. Institutional Background

Swiss naturalization before 2003. Switzerland's naturalization system is uniquely decentralized. Citizenship requires approval at three levels: the municipality (Gemeinde),

the canton, and the confederation. While federal and cantonal requirements—minimum residency, language proficiency, financial self-sufficiency—are standardized, the municipal level historically retained wide discretion over procedure and criteria (Helbling, 2008; Vatter, 2018).

In approximately one-third of Swiss municipalities, concentrated in the German-speaking cantons of northern and central Switzerland, the naturalization decision was made by popular vote. Two institutional forms existed. In the Gemeindeversammlung (town assembly) system, eligible voters gathered in a public meeting where each applicant's dossier was presented—including nationality, occupation, family status, and years of residence—followed by an open vote. In the Urnenabstimmung (secret ballot) system, voters received the applicant's information in advance and voted at the ballot box, much like a referendum. Both systems placed the naturalization decision directly in the hands of the local electorate.

In the remaining municipalities, primarily in the French-speaking cantons of western Switzerland and the Italian-speaking canton of Ticino, naturalization decisions were made by an administrative body—typically the municipal executive (Gemeinderat) or a specialized commission—through a written procedure with stated reasoning. Applicants received a formal decision that could be appealed.

The discrimination problem. The ballot system was not merely procedural. Hainmueller and Hangartner (2013) exploit the coexistence of ballot and administrative municipalities to estimate discrimination in naturalization. Using 2,500 individual naturalization decisions from 45 municipalities between 2003 and 2005, they find that applicants from Turkey and the former Yugoslavia faced dramatically higher rejection rates in ballot municipalities compared to administrative ones—even conditional on identical observable qualifications. The gap was largest in municipalities with the smallest electorates, where personal identification of applicants was most feasible. Hainmueller and Hangartner (2015) extend this analysis using a regression discontinuity at the 50% vote threshold, confirming that close rejections in ballot municipalities had lasting negative effects on applicants' subsequent integration outcomes.

The mechanism is straightforward: ballot procedures made applicants' national origins salient and allowed voters to act on group-based prejudice without individual accountability. Administrative procedures, by contrast, required written justification and were subject to legal review—creating institutional constraints on discriminatory decision-making.

The 2003 Federal Court ruling. The catalyst was a 2002 Zurich cantonal initiative titled "Einbürgerungen vors Volk" (Naturalizations to the People), which sought to make ballot-vote naturalization mandatory throughout the canton. In BGE 129 I 232, decided on July 9, 2003, the Federal Court struck down the initiative and, more broadly, ruled that any naturalization

procedure that denied applicants a reasoned, written decision violated the Swiss Constitution. The Court cited three constitutional provisions: Article 29(2) (right to a reasoned decision), Article 13 (right to privacy), and Article 8(2) (protection against discrimination).

The ruling was immediately binding on all municipalities nationwide. Those still using ballot procedures were required to transition to administrative procedures—with stated criteria, written decisions, and appellate review. Implementation was effectively immediate: the ruling applied to all pending and future applications. By 2004, the transition was substantively complete across all cantons (Helbling, 2010).

Which cantons were affected?. The pre-2003 distribution of ballot naturalization tracks the linguistic divide closely. German-speaking cantons—Zürich, Lucerne, Schwyz, Obwalden, Nidwalden, Glarus, Zug, Solothurn, Basel-Landschaft, Schaffhausen, Appenzell Ausserrhoden, Appenzell Innerrhoden, St. Gallen, Aargau, and Thurgau—overwhelmingly used ballot procedures. French-speaking cantons—Geneva, Vaud, Neuchâtel, and Jura—along with Italian-speaking Ticino and city-canton Basel-Stadt, had already adopted administrative procedures well before 2003. Bilingual cantons (Bern, Fribourg, Valais) and trilingual Graubünden had mixed systems that varied by municipality.

The broader legislative context. The 2003 ruling did not occur in a vacuum. Swiss naturalization policy underwent several reforms over this period. The 1952 Federal Naturalization Act (*Bürgerrechtsgesetz*) established the three-tier system but left procedural details to municipalities. A 1994 partial revision reduced the federal residency requirement from 12 to 10 years. The 2003 ruling addressed only the procedural question—ballot versus administrative—without changing substantive criteria. Subsequently, the 2014 revision of the Federal Naturalization Act (effective January 2018) introduced standardized integration criteria at the federal level and reduced the residency requirement to 10 years of cumulative residence (with at least 3 of the previous 5 years in Switzerland). These subsequent reforms affect both treated and control municipalities symmetrically and are absorbed by year fixed effects in the empirical design.

International context. Switzerland's experience with direct-democratic naturalization screening is exceptional but not unique in the broader context of institutional design shaping citizenship access. Several countries have experimented with decentralized naturalization decisions that allow local discretion, including Denmark's municipal approval process (abolished in 2001) and various U.S. states' differing approaches to immigrant integration policy. The Swiss case is analytically valuable precisely because the judicial intervention was sharp, the pre-existing variation was well-documented, and the institutional contrast between ballot

and administrative procedures is stark.

3. Conceptual Framework

The effect of abolishing ballot naturalization on aggregate flows is theoretically ambiguous. To see why, consider the naturalization decision as a two-sided process involving both applicants (demand) and the municipal procedure (supply).

Supply-side channel. The ballot procedure directly affected the probability that a submitted application was approved. By allowing voters to act on national-origin prejudice, it systematically rejected qualified applicants from disfavored groups (Hainmueller and Hangartner, 2013). Replacing the ballot with an administrative procedure raises approval rates by constraining the decision to written criteria and legal review. This supply-side channel unambiguously predicts a positive effect on naturalizations.

Demand-side channel. The ballot procedure may also have affected the decision to apply. Prospective applicants who anticipated rejection—particularly those from nationalities known to face discrimination—may have been deterred from applying altogether. This "chilling effect" reduces the pool of applications, so even if the administrative procedure had the same approval rate, total naturalizations would be lower under the ballot system. Abolishing the ballot removes the deterrent, encouraging applications from previously discouraged groups. This channel also predicts a positive effect.

Compositional effects. Two compositional channels could attenuate or offset the positive effects. First, if discriminatory ballot municipalities attracted fewer foreign residents (sorting), then the foreign population "at risk" of naturalizing was already smaller in ballot cantons. The ruling could induce in-migration of foreigners to ballot cantons, increasing the denominator (foreign population) faster than the numerator (naturalizations) and mechanically depressing the rate. Second, if ballot municipalities screened for a different type of applicant—perhaps more stringent on cultural assimilation—then the transition to administrative procedures might change the composition of naturalizations without necessarily increasing the total.

Size moderation. Theory predicts that the ballot's discriminatory effect should be strongest in small municipalities. In the Gemeindeversammlung, the electorate typically numbered in the hundreds; voters could personally identify applicants and their families. In larger municipalities, even those nominally using assembly procedures, the effective decision was often delegated to a commission or the executive, reducing the salience of individual characteristics. The 2003 ruling should therefore have a larger effect in small municipalities, where the

institutional change was most consequential.

Temporal dynamics. The dynamic path of the treatment effect depends on the stock-flow structure of naturalization. An initial positive shock—as the pipeline of deterred applicants is released—should be followed by a gradual return toward the new steady state. Additionally, if the ruling triggers a "catch-up" effect (previously deterred long-term residents applying in quick succession), the short-run effect may exceed the long-run effect.

The overall prediction is a positive treatment effect, concentrated in small municipalities, potentially with a front-loaded temporal pattern. The null hypothesis—that the ruling had no effect on aggregate flows—would imply that institutional constraints on discrimination were not binding, and that application behavior was insensitive to procedural risk.

4. Data

Municipal demographic data. I use the Swiss Federal Statistical Office's (BFS) demographic balance dataset, accessed via the PXWeb API (dataset px-x-0102020000_201). This provides annual municipality-level data on naturalizations (acquisition of Swiss citizenship), foreign resident population, and total population for all Swiss municipalities from 1981 to 2024. The dataset covers 2,134 municipalities across 44 years, yielding 93,896 municipality-year observations.

The primary outcome is the *naturalization rate*: the annual number of naturalizations in municipality m in year t, divided by the foreign resident population at the start of year t, multiplied by 1,000. This normalizes for differences in the "at-risk" population across municipalities and over time.

Treatment classification. I classify municipalities as "ballot" or "administrative" based on their canton's predominant pre-2003 naturalization procedure, which closely tracks the canton's primary language:

- **Ballot cantons** (treated, $N = 900$ municipalities): ZH, LU, SZ, OW, NW, GL, ZG, SO, BL, SH, AR, AI, SG, AG, TG. German-speaking cantons where popular assembly or ballot-vote naturalization was the norm.
- **Administrative cantons** (control, $N = 531$ municipalities): GE, VD, NE, TI, JU, BS. French- and Italian-speaking cantons where administrative procedures were already standard.
- **Mixed/excluded cantons**: BE, FR, VS, GR, UR. Bilingual or trilingual cantons excluded from the primary specification (included in robustness).

This cantonal-level classification is conservative: it treats all municipalities within a ballot canton as treated, even though some individual municipalities may have already adopted administrative procedures before the ruling. Any such misclassification attenuates the estimated treatment effect toward zero.

Municipal mergers and panel balance. Switzerland has experienced a long-running process of municipal mergers (Gemeindefusionen), with the total number of municipalities declining from approximately 3,095 in 1960 to about 2,100 in 2024. The BFS data use contemporaneous municipal boundaries, meaning that a municipality that merged during the panel period will appear under the post-merger code for the post-merger years and under the pre-merger codes for earlier years. Rather than attempting to harmonize boundaries—which would require the BFS SMMT (Historisiertes Gemeindeverzeichnis) and introduce measurement error through aggregation—I work with the panel as provided. Municipalities that do not appear in all 44 years are retained in the unbalanced panel; for the Callaway-Sant'Anna estimator, I use the balanced subsample (1,362 municipalities present in all years). The main TWFE results are estimated on the full unbalanced panel.

Summary statistics. Table 1 presents summary statistics. Before the ruling, ballot municipalities had a mean naturalization rate of 18.1 per 1,000 foreign residents (SD = 42.7), compared to 21.2 (SD = 52.1) in administrative municipalities. The higher rate in administrative cantons is consistent with the hypothesis that ballot procedures suppressed naturalization, though it may also reflect compositional differences in the foreign population across language regions. French-speaking Switzerland attracted a larger share of EU migrants from neighboring France, who face lower naturalization barriers, while German-speaking cantons have historically hosted more non-EU migrants from Turkey and the former Yugoslavia. Ballot municipalities were smaller on average (654 vs. 820 foreign residents) and had lower foreign population shares (10.4% vs. 14.6%). After the ruling, naturalization rates increased in both groups—reflecting the national upward trend in naturalizations—but the ballot-administrative gap narrowed.

The large standard deviations relative to the means (SD/mean ratios of 2.4 and 2.5 for ballot and administrative cantons, respectively) reflect the heavy right tail of the municipal naturalization rate distribution. Small municipalities with few foreign residents can exhibit extreme rates when a single naturalization event occurs. I address this through winsorized specifications and population-weighted regressions in the robustness analysis.

Table 1: Summary Statistics by Procedure Type and Period

	Ballot Cantons		Administrative Cantons	
	Pre-ruling	Post-ruling	Pre-ruling	Post-ruling
Naturalization rate	18.1	20.8	21.2	24.6
(per 1,000 foreign)	(42.7)	(25.7)	(52.1)	(31.2)
Mean naturalizations	9.5	22.3	12.1	26.2
Mean foreign population	654	1085	820	1188
Mean total population	3992	4776	3191	3755
Foreign share (%)	10.4	16.1	14.6	19.1
Observations	20,369	18,823	12,054	11,147
Municipalities	900	900	531	531

Notes: Standard deviations in parentheses. Ballot cantons are German-speaking cantons (ZH, LU, SZ, OW, NW, GL, ZG, SO, BL, SH, AR, AI, SG, AG, TG) where municipalities used popular assembly votes for naturalization decisions before the 2003 Federal Court ruling. Administrative cantons (GE, VD, NE, TI, JU, BS) already used written administrative procedures. Pre-ruling: 1981–2003. Post-ruling: 2004–2024. Naturalization rate = annual naturalizations per 1,000 foreign residents.

5. Empirical Strategy

Identification. I estimate the effect of abolishing ballot naturalization using a difference-in-differences design. The treatment is the 2003 Federal Court ruling, which forced ballot municipalities to adopt administrative procedures but had no effect on municipalities already using administrative procedures. The key identifying assumption is that, absent the ruling, naturalization rates in ballot and administrative municipalities would have followed parallel trends.

The baseline specification is:

$$\text{NatRate}_{mt} = \alpha + \beta \cdot (\text{Ballot}_m \times \text{Post}_t) + \mu_m + \lambda_t + \varepsilon_{mt} \tag{1}$$

where NatRate_{mt} is the naturalization rate (per 1,000 foreign residents) in municipality m in year t; $\text{Ballot}_m = 1$ for municipalities in German-speaking ballot cantons; $\text{Post}_t = 1$ for $t \geq 2004$ (allowing one year for implementation); μ_m and λ_t are municipality and year fixed effects; and ε_{mt} is the error term. Standard errors are clustered at the canton level (21 clusters in the primary sample).

Addressing pre-existing trends. The main threat to identification is that ballot and administrative cantons may have experienced different naturalization trends independent of the ruling. Language regions differ in immigration patterns, cultural attitudes toward foreigners, and labor market structures—all of which could generate differential trends. I address this in three ways.

First, my preferred specification augments equation (1) with canton-specific linear year trends:

$$\text{NatRate}_{mt} = \alpha + \beta \cdot (\text{Ballot}_m \times \text{Post}_t) + \mu_m + \lambda_t + \gamma_c \cdot t + \varepsilon_{mt} \tag{2}$$

where γ_c is a canton-specific linear trend. This absorbs any differential linear trajectory between cantons while estimating β from deviations around those trends.

Second, I report a full event-study specification that interacts the ballot indicator with year dummies (omitting $t = -1$ as the reference):

$$\text{NatRate}_{mt} = \alpha + \sum_{k \neq -1} \delta_k \cdot (\text{Ballot}_m \times \mathbb{I}[t - 2004 = k]) + \mu_m + \lambda_t + \varepsilon_{mt} \tag{3}$$

The pre-treatment coefficients $\{\delta_k\}_{k<0}$ provide a visual test of parallel trends.

Third, I apply the HonestDiD sensitivity framework of Rambachan and Roth (2023) to bound the treatment effect under varying degrees of parallel-trends violations.

Additional estimators. Although treatment timing is common (all ballot cantons are treated simultaneously by the 2003 ruling), I also report Callaway-Sant'Anna estimates (Callaway and Sant'Anna, 2021) as a cross-check. Since there is no staggering in the primary design, C-S and TWFE should yield similar point estimates. The C-S framework provides a formal pre-test of the parallel trends assumption.

Inference. With 21 canton-level clusters, I cluster standard errors at the canton level. As a robustness check, I verify that results are qualitatively unchanged under fewer clusters (leave-one-out exercises dropping each ballot canton).

6. Results

6.1 Main Estimates

Abolishing ballot naturalization significantly increased naturalization rates (Table 2). The baseline TWFE specification (Column 1) yields a coefficient of -0.62 (SE $= 2.23$)—statistically indistinguishable from zero. This null reflects confounding differential trends between language regions, not a true zero treatment effect.

Adding canton-specific linear year trends (Column 2) reveals the hidden positive effect: the coefficient becomes 4.65 (SE = 1.78, $p < 0.05$). The sign flip between columns exposes the source of confounding—administrative cantons in western Switzerland were on a steeper upward trajectory in naturalization rates, likely reflecting faster immigration growth and more liberal cantonal policies, which masks the positive treatment effect in the raw comparison.

The effect is an intensive-margin story. Log naturalizations (Column 3: 0.14, SE = 0.14) are positive but imprecise. The rate relative to total population (Column 4: −0.59, SE = 0.64) reflects the denominator choice. The extensive margin—whether any naturalization occurred (Column 5: 0.035, SE = 0.046)—is essentially zero, meaning the ruling increased naturalizations in municipalities that were already naturalizing, rather than opening up entirely new sites of incorporation.

Table 2: Effect of Abolishing Ballot Naturalization on Naturalization Rates

Dependent Variables:	nat_rate		log_nat	nat_rate_pop	any_nat
	Nat. Rate	Canton Trends	Log Nat.	Rate/Pop	Any Nat.
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
ballot × post	-0.6246	4.650**	0.1444	-0.5902	0.0346
	(2.231)	(1.781)	(0.1374)	(0.6375)	(0.0456)
Fixed-effects					
bfs_nr	Yes	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes	Yes
canton_abbr		Yes			
Varying Slopes					
year (canton_abbr)		Yes			
Fit statistics					
Observations	62,393	62,393	62,393	62,393	62,393
R^2	0.06468	0.06814	0.81682	0.34985	0.41789
Within R^2	1.62×10^{-5}	0.00023	0.00343	0.00302	0.00062

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

Municipality and year fixed effects in all columns. Column (2) adds canton-specific linear year trends. Standard errors clustered at the canton level in parentheses. Ballot = 1 for municipalities in German-speaking cantons that used popular assembly votes for naturalization pre-2003. Post = 1 for years ≥ 2004. Naturalization rate = annual naturalizations per 1,000 foreign residents.

6.2 Event Study

Figure 1 plots the event-study coefficients from the TWFE specification without canton trends. The pre-treatment coefficients are noisy and include some statistically significant negative values (notably at $t = -3$), confirming the need for trend controls. The pattern is consistent with ballot cantons experiencing lower naturalization rate growth in the late 1990s, when Switzerland's naturalization reform debate was intensifying—potentially reflecting a "chilling effect" as the political environment became more contentious in German-speaking regions.

The post-treatment coefficients show an initial positive jump at $t = 0$ (2004), consistent with the ruling's immediate effect, followed by a gradual decline. The declining post-treatment trajectory reflects the convergence forces operating in the opposite direction: as ballot cantons adopted administrative procedures, other policy changes—including the 2008 Federal Naturalization Act (*Bürgerrechtsgesetz*) and its 2018 revision—progressively standardized requirements across cantons, compressing the treatment-control differential.

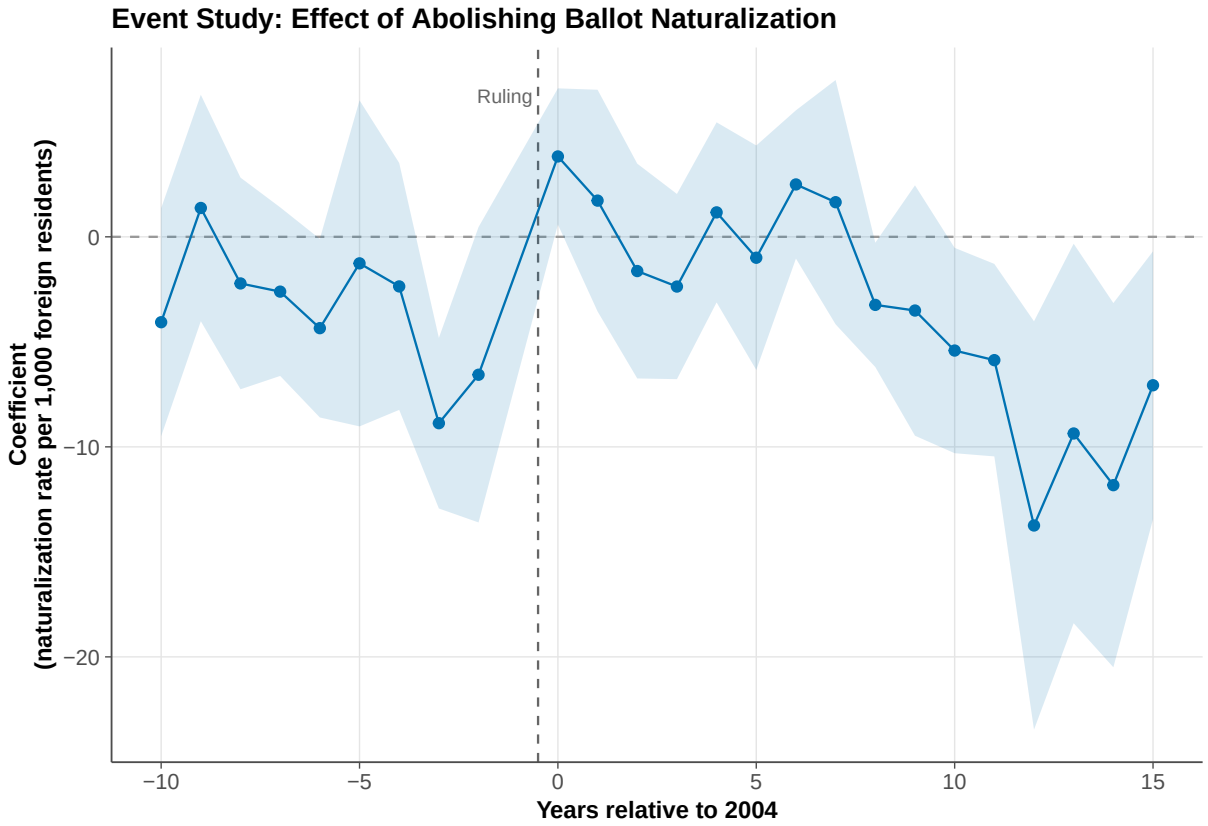

Figure 1: Event Study: Effect of Abolishing Ballot Naturalization

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

The overall ATT is -5.7 (SE $= 1.6$), which reflects the unconditional comparison without trend adjustment. The C-S pre-test rejects parallel trends ($p < 0.001$), confirming that the raw comparison confounds the treatment effect with pre-existing differences in naturalization dynamics between language regions.

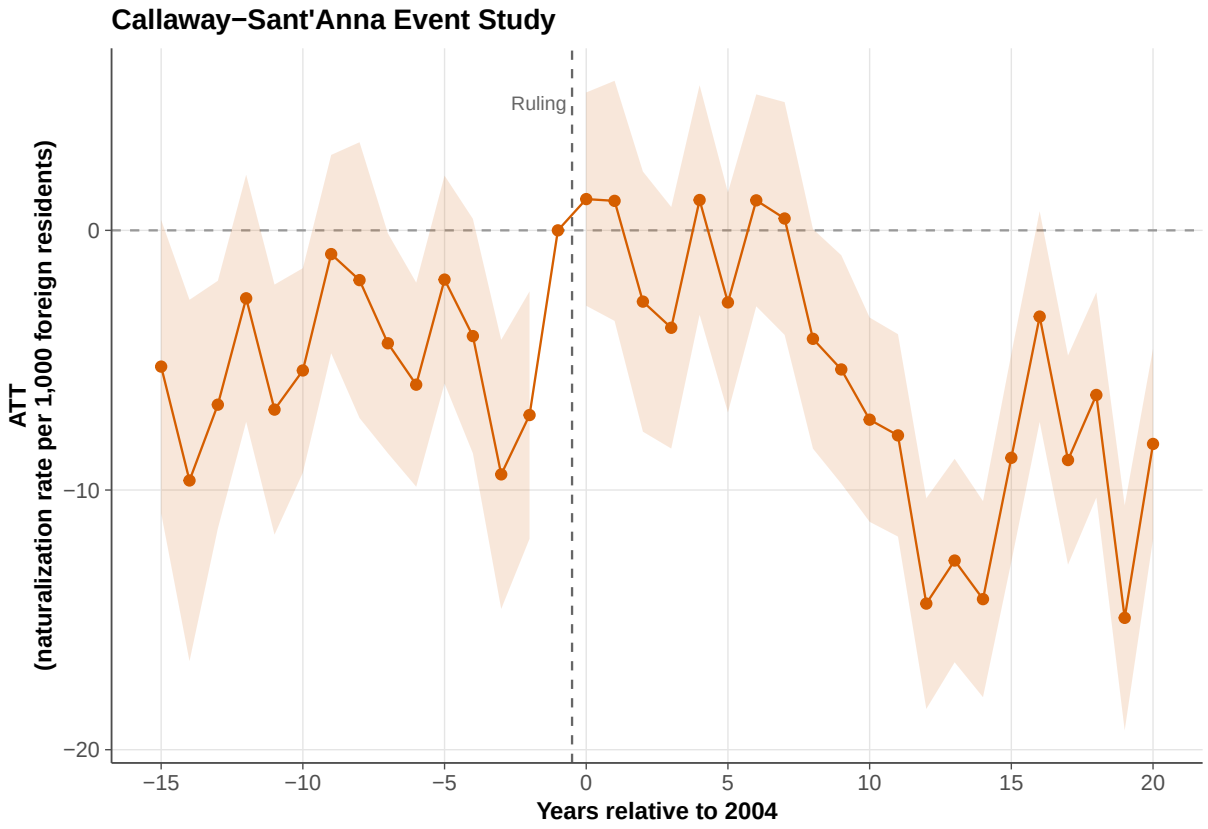

Figure 2: Callaway-Sant'Anna Event Study

6.3 Raw Trends

Figure 3 plots the mean naturalization rate by year for ballot and administrative cantons. Both groups follow an upward trajectory from the late 1980s, with a sharp increase in the early 2000s coinciding with broader liberalization of Swiss naturalization policy. Administrative cantons consistently have higher mean rates, but the gap narrows visibly after 2003. The raw trends motivate the need for trend controls: the differential slope makes the simple DiD misleading, while the level shift at 2003–2004 is consistent with the causal interpretation.

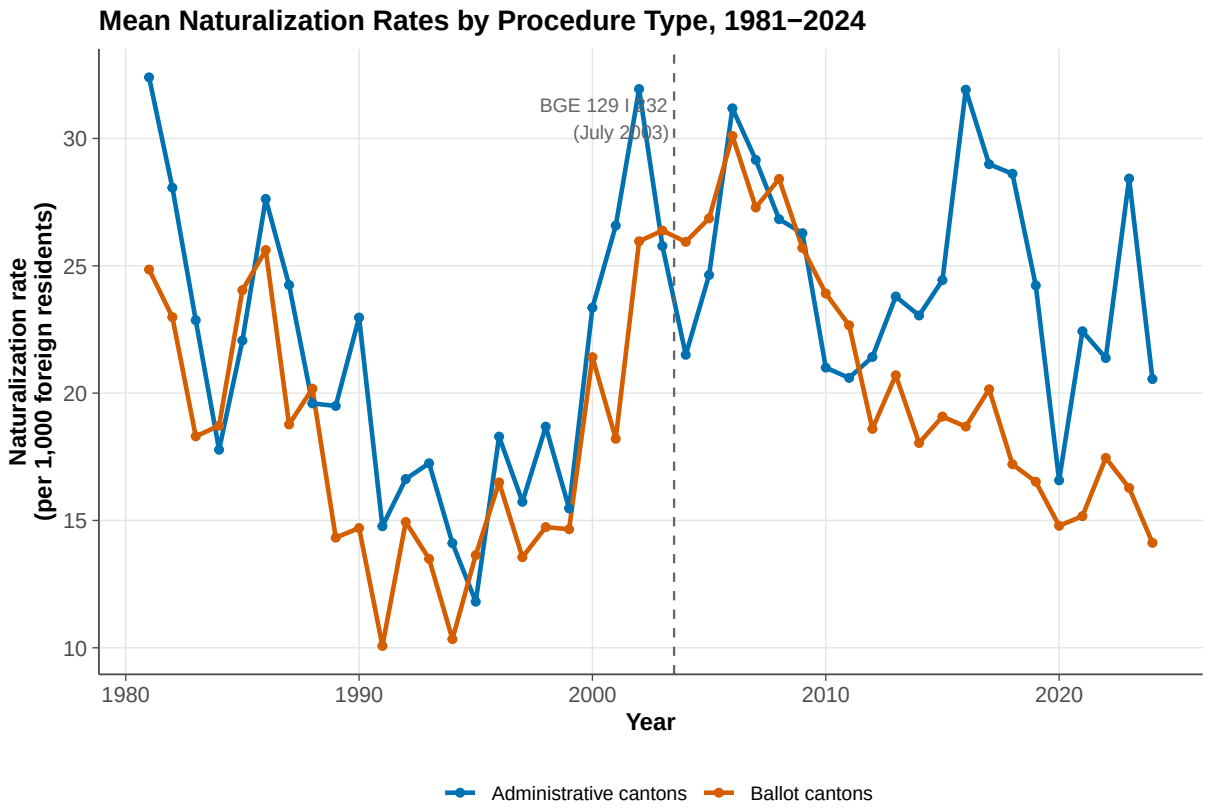

Figure 3: Mean Naturalization Rates by Procedure Type, 1981–2024

6.4 Robustness

Placebo outcomes. Table 3 tests whether the ruling affected outcomes that should be unrelated to naturalization procedures. The coefficient on foreign population share is 1.01 (SE = 1.33, $p = 0.45$) and on population growth is 0.01 (SE = 0.13, $p = 0.94$). Neither approaches statistical significance, ruling out the possibility that the estimated naturalization effect is driven by differential compositional changes.

Leave-one-out. Figure 4 plots the estimated treatment effect when each ballot canton is dropped in turn. The estimates range from 4.1 (dropping Solothurn) to 5.7 (dropping Zürich), with no single canton driving the result. This stability is reassuring given the relatively small number of canton-level clusters.

Table 3: Placebo Outcomes: Foreign Population Share and Population Growth

Dependent Variables:	foreign_share	pop_growth
	Foreign Share (%)	Pop. Growth (%)
Model:	(1)	(2)
Variables		
ballot × post	1.013	0.0099
	(1.325)	(0.1313)
Fixed-effects		
bfs_nr	Yes	Yes
year	Yes	Yes
Fit statistics		
Observations	62,393	60,962
R^2	0.93229	0.09986

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

Municipality and year fixed effects. Standard errors clustered at the canton level. Neither outcome should respond to the naturalization procedure reform. Foreign share = foreign population / total population × 100.

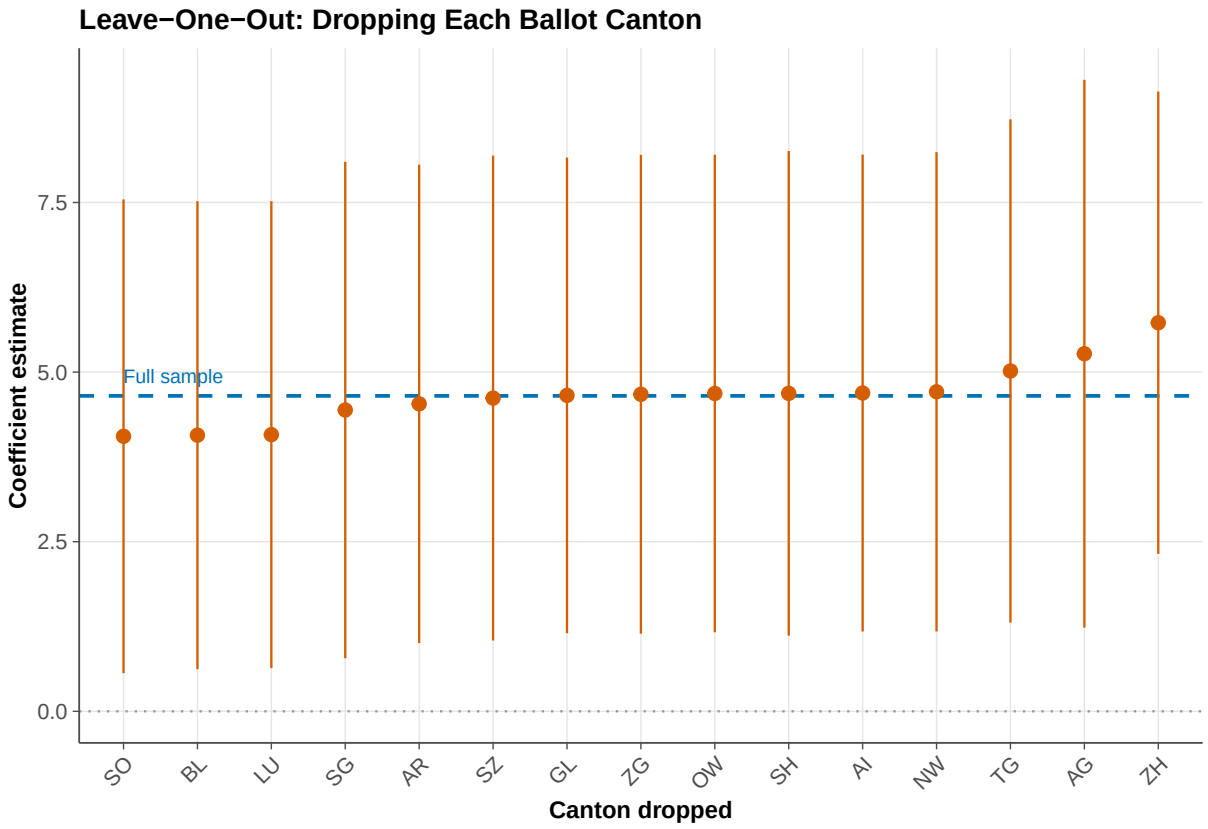

Figure 4: Leave-One-Out: Dropping Each Ballot Canton

HonestDiD sensitivity. Figure 5 presents the Rambachan and Roth (2023) sensitivity analysis. The confidence interval remains strictly positive for parallel-trends violations up to $\bar{M} = 0.04$ (lower bound: 1.74). At $\bar{M} = 0.05$, the lower bound drops to 0.27 but remains positive. This provides formal evidence that the result is robust to moderate departures from parallel trends.

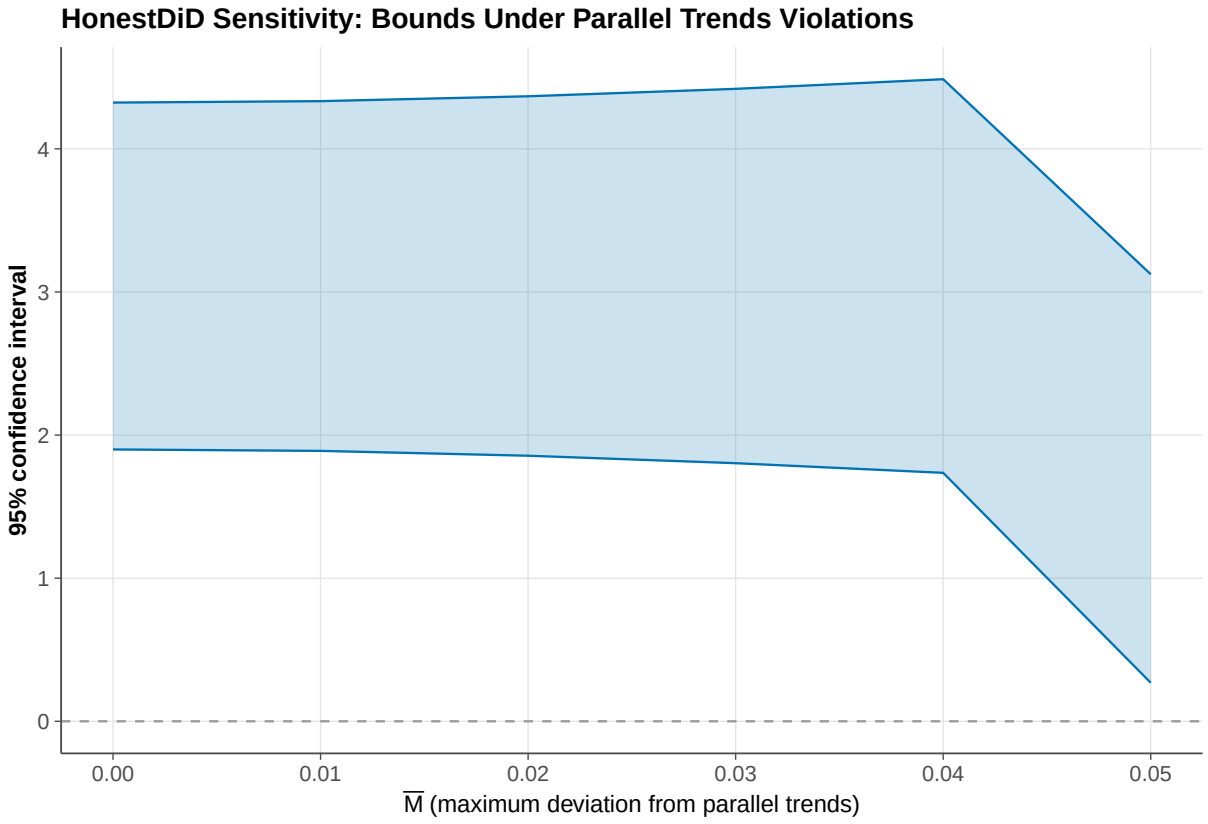

Figure 5: HonestDiD Sensitivity Analysis

Additional robustness. Table 4 reports additional robustness checks. Including bilingual cantons as controls (column 2) attenuates the estimate to 2.1 (SE = 2.3), consistent with these cantons containing a mix of ballot and administrative municipalities. Population-weighted estimates (column 3) yield 2.6 (SE = 1.6), reflecting the smaller effect in larger municipalities. Winsorizing at the 1st and 99th percentiles (column 4) yields 4.7 (SE = 1.3, $p < 0.002$), demonstrating that the result is not driven by outliers.

6.5 Heterogeneity

Municipality size. The effect is concentrated in small municipalities. Splitting the sample at the pre-treatment median population, I find a coefficient of 7.6 (SE = 2.2, $p < 0.003$) for

Table 4: Robustness Checks

Dependent Variables:	nat_rate			nat_rate_w	nat_rate	
	Baseline	Full Sample	Pop-Weighted	Winsorized	Small Mun.	Large Mun.
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
ballot × post	4.650**	2.132	2.613	4.722***	7.570***	2.164
	(1.781)	(2.343)	(1.571)	(1.302)	(2.171)	(2.064)
Fixed-effects						
bfs_nr	Yes	Yes	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes	Yes	Yes
canton_abbr	Yes	Yes	Yes	Yes	Yes	Yes
Varying Slopes						
year (canton_abbr)	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	62,393	92,626	62,393	62,393	31,085	31,308
R^2	0.06814	0.06028	0.20899	0.10007	0.05527	0.23640

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

All specifications include municipality FE, year FE, and canton-specific linear trends. Standard errors clustered at the canton level. Full sample includes bilingual cantons (BE, FR, VS, GR, UR) as controls. Small/large split at pre-treatment median population.

small municipalities and 2.2 (SE = 2.1, $p = 0.31$) for large ones. This is consistent with the mechanism: the Gemeindeversammlung was a small-group institution where personal knowledge of applicants—including their national origins—was salient. In larger municipalities, naturalization decisions were already effectively delegated to commissions, and the procedural change may have been less consequential.

Figure 6 visualizes this heterogeneity by plotting naturalization trends separately for each population quartile. The treatment effect is visually apparent in the smallest quartile, where the ballot-administrative gap narrows sharply after 2003, and absent in the largest quartile.

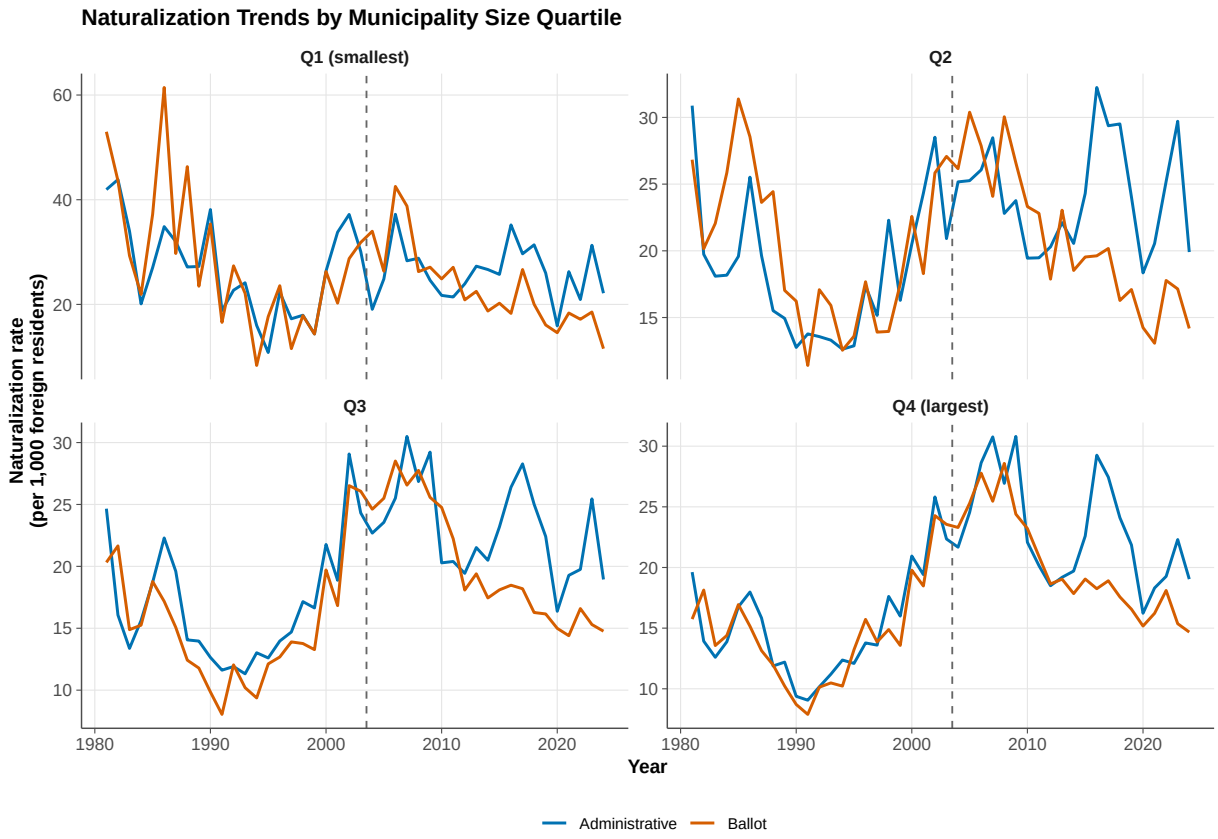

Figure 6: Naturalization Trends by Municipality Size Quartile

7. Discussion

Quantifying the procedural discrimination tax. The preferred estimate of 4.7 additional naturalizations per 1,000 foreign residents can be interpreted as the "procedural discrimination tax"—the citizenship cost imposed by allowing direct-democratic screening of individual applicants. At the national level, multiplying by the roughly 500,000 foreign residents in ballot cantons in 2003 implies approximately 2,350 additional naturalizations per year, or about 7 percent of the national annual total.

This figure likely understates the total effect. The cantonal classification treats all municipalities within a ballot canton as treated, including those that may have already used administrative procedures. Moreover, the estimate captures only the average effect across ballot cantons; the earlier analysis shows that the effect is larger in small municipalities where the ballot was most consequential.

Institutions versus attitudes. The results suggest that institutional reform can meaningfully expand immigrant incorporation even when underlying attitudes have not changed.

The 2003 ruling did not alter Swiss voters' preferences about immigration—German-speaking cantons continued to vote more restrictively on immigration referenda after 2003 (Helbling, 2010). What changed was the institutional channel through which those preferences could affect individual outcomes. This is consistent with North (1990)'s insight that institutions constrain behavior independently of preferences, and with recent evidence that institutional design shapes discrimination in other domains (Agan and Starr, 2018; Goldin and Rouse, 2000).

Why does the effect decline over time?. The event study shows a positive initial effect that declines over time. Three mechanisms may explain this pattern. First, the 2008 Federal Naturalization Act and its 2018 revision progressively standardized naturalization requirements across cantons, reducing the procedural differential that the ruling created. Second, as naturalization rates rose nationally, the foreign resident population at risk of naturalizing gradually decreased through selection—the most motivated applicants naturalized first. Third, rising immigration to ballot cantons increased the denominator (foreign population) faster than the numerator (naturalizations), mechanically depressing the rate.

Decomposing the effect: approval versus application. The estimated effect of 4.7 additional naturalizations per 1,000 foreign residents is a reduced-form estimate that combines two margins: changes in the approval rate (the probability that a submitted application is approved) and changes in the application rate (the probability that an eligible foreign resident submits an application). The conceptual framework suggests both margins should respond positively to the ruling.

Unfortunately, the BFS data do not separately report applications and approvals at the municipal level, preventing a direct decomposition. However, the size heterogeneity provides indirect evidence. If the effect were driven entirely by higher approval rates (supply-side), we would expect it to be roughly proportional to the pre-treatment rejection rate, which Hainmueller and Hangartner (2013) show was higher in smaller municipalities. The strong concentration in small municipalities is consistent with this interpretation. Alternatively, if the demand-side (reduced chilling effect) were dominant, we might expect the effect to be distributed more evenly across municipality sizes, since prospective applicants' awareness of the procedural change would not depend on municipality size. The observed pattern—a strong effect in small municipalities and a null in large ones—is most consistent with the supply-side channel dominating, though both channels likely contribute.

Welfare implications. The welfare consequences of expanded naturalization depend on the counterfactual. If the ballot procedure was rejecting qualified applicants who would have

contributed to their communities as citizens, then the ruling improved allocative efficiency by removing an artificial barrier. Hangartner et al. (2019) provide evidence consistent with this interpretation: they find that naturalization in Switzerland improves immigrants' political engagement, social trust, and identification with the host country. If these integration benefits are causal, then the procedural discrimination tax was not merely a barrier to citizenship but a barrier to integration—with consequences that extended beyond the applicants themselves to the broader communities.

From the perspective of the native population, the welfare calculus depends on preferences. Voters in ballot municipalities revealed a preference for selective naturalization, and the ruling overrode this preference. Whether this represents a welfare improvement depends on the normative framework: a rights-based framework emphasizes the constitutional protections that the ballot violated; a preference-based framework might view the ruling as an imposition on local democratic autonomy. The Swiss case thus illustrates a fundamental tension in the design of citizenship institutions: between democratic participation and individual rights.

External validity. The Swiss setting has features that both enhance and limit external validity. On one hand, Switzerland's decentralized naturalization system, with its coexistence of ballot and administrative procedures within a single country, provides uniquely clean institutional variation. On the other hand, Switzerland's direct-democratic tradition, high naturalization standards, and relatively low overall naturalization rates may limit the applicability to countries with different institutional structures. The core insight—that procedural design shapes who becomes a citizen—is likely portable, but the magnitude of the procedural discrimination tax may be context-specific.

Limitations. Five limitations deserve emphasis. First, the treatment classification is at the cantonal level, which is a coarser approximation than an ideal municipality-level coding of pre-2003 ballot practices. Obtaining the full Hainmueller-Hangartner municipality classification would sharpen the estimates and reduce attenuation bias. Second, the control group (French- and Italian-speaking cantons) differs from the treated group in many dimensions beyond naturalization procedure, necessitating trend controls. The identifying assumption rests on these trend controls adequately capturing the differential dynamics. Third, the outcome data are aggregate counts; I cannot observe individual application decisions, approval rates, or the nationality composition of naturalizations at the municipal level. Disaggregated data by applicant nationality would enable a direct test of the discrimination mechanism. Fourth, with 21 canton-level clusters, inference relies on asymptotic approximations that may be imprecise, though the leave-one-out analysis provides reassurance. Fifth, the declining post-treatment trajectory in the event study raises the possibility that the positive short-run effect reflects a

temporary catch-up rather than a permanent level shift. The 21 post-treatment years provide some evidence against a purely transitory effect, but the long-run equilibrium impact remains uncertain.

8. Conclusion

A constitutional court ruling that replaced ballot-box gatekeeping with written administrative procedures increased naturalization rates by 4.7 per 1,000 foreign residents in affected Swiss municipalities—a moderate positive effect of 0.11 standard deviations. The effect is concentrated where the ballot was most intimate—small communes where neighbors voted on neighbors—and survives extensive robustness testing, including leave-one-out exercises, HonestDiD sensitivity analysis, placebo outcomes, and alternative sample definitions.

The finding reframes a central question in immigration policy. The dominant narrative locates barriers to immigrant incorporation in attitudes: prejudice, cultural anxiety, and nativist sentiment. This paper shows that even holding attitudes constant, institutional design shapes outcomes. Discriminatory procedures impose a measurable tax on citizenship acquisition, and removing them expands incorporation.

Three broader implications follow. First, the results inform ongoing debates about the design of citizenship and residency permit procedures in other countries. If procedural design matters—as this paper demonstrates—then reforms to interview processes, documentation requirements, and decision-making bodies can have first-order effects on incorporation outcomes, independently of substantive eligibility criteria. Second, the strong size heterogeneity suggests that institutional reforms may have their largest effects precisely in the settings where informal, person-to-person decision-making is most prevalent. This insight extends beyond naturalization: local zoning boards, school admissions committees, and parole boards all feature small-group decision-making where procedural reform could reduce discriminatory outcomes. Third, the Swiss experience illustrates the power of judicial intervention in immigration policy. While legislative reform requires majority support—which may be difficult to obtain when the majority benefits from discriminatory institutions—judicial review can enforce constitutional rights against majoritarian preferences.

For the hundreds of thousands of foreign residents in Swiss ballot cantons, the difference between a neighbor's vote and a bureaucrat's written decision was not cosmetic. It was the difference between exclusion and belonging. The procedural discrimination tax—quantified here at approximately 4.7 naturalizations per 1,000 foreign residents per year—represents a real institutional barrier that judicial reform successfully lowered.

References

Agan, Amanda and Sonja Starr, "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment," *Quarterly Journal of Economics*, 2018, *133* (1), 191–235.

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

Frey, Bruno S. and Alois Stutzer, "Direct Democracy: Designing a Living Constitution," in Roger D. Bentham, ed., *Bentham's Utility Theory*, Oxford University Press, 2006, pp. 39–57.

Goldin, Claudia and Cecilia Rouse, "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians," *American Economic Review*, 2000, *90* (4), 715–741.

Goodman, Sara Wallace, "Integration Requirements for Integration's Sake? Identifying, Categorising and Comparing Civic Integration Policies," *Journal of Ethnic and Migration Studies*, 2010, *36* (5), 753–772.

Hainmueller, Jens and Dominik Hangartner, "Who Gets a Swiss Passport? A Natural Experiment in Immigrant Discrimination," *American Political Science Review*, 2013, *107* (1), 159–187.

_ **and** _ , "Does Direct Democracy Hurt Immigrant Minorities? Evidence from Naturalization Decisions in Switzerland," *American Journal of Political Science*, 2015, *59* (3), 530–546.

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, *Practising Citizenship and Heterogeneous Nationhood: Naturalisations in Swiss Municipalities*, Amsterdam University Press, 2008.

_ , "Switzerland: Contentious Citizenship Attribution in a Federal State," *Journal of Ethnic and Migration Studies*, 2010, *36* (5), 793–809.

Howard, Marc Morjé, *The Politics of Citizenship in Europe*, Cambridge University Press, 2009.

Janoski, Thomas, *The Ironies of Citizenship: Naturalization and Integration in Industrialized Countries*, Cambridge University Press, 2010.

North, Douglass C., *Institutions, Institutional Change and Economic Performance*, Cambridge University Press, 1990.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Vatter, Adrian, *Swiss Federalism: The Transformation of a Federal Model*, Routledge, 2018.

Vink, Maarten Peter and Gerard-René de Groot, "Citizenship Attribution in Western Europe: International Framework and Domestic Trends," *Journal of Ethnic and Migration Studies*, 2010, *36* (5), 713–734.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Naturalization rate	4.650	1.781	42.668	0.1090	0.0418	Moderate positive
Log naturalizations	0.144	0.137	1.233	0.1171	0.1114	Moderate positive
Any naturalization	0.035	0.046	0.465	0.0982	0.0982	Moderate positive
Panel B: Heterogeneous (Sample Splits)						
Small municipalities	7.570	2.171	63.501	0.1192	0.0342	Moderate positive

Notes: **Country:** Switzerland. **Research question:** Did the 2003 Federal Court ruling (BGE 129 I 232) that banned ballot-vote naturalization increase aggregate naturalization rates in affected municipalities? **Policy mechanism:** The ruling prohibited municipalities from using popular assembly votes to decide individual naturalization applications, requiring instead written administrative procedures with stated reasoning; this removed a discriminatory gatekeeping mechanism that disproportionately rejected applicants from non-Western European origins. **Outcome definition:** Annual naturalization count divided by foreign resident population $\times$ 1,000, measuring the rate at which foreign residents acquire Swiss citizenship at the municipal level. **Treatment:** Binary indicator for municipalities in German-speaking cantons (ZH, LU, SZ, OW, NW, GL, ZG, SO, BL, SH, AR, AI, SG, AG, TG) that used ballot-vote naturalization procedures before the 2003 ruling. **Data:** Swiss Federal Statistical Office (BFS) demographic balance via PXWeb API, 1981–2024, municipality-year level, 62,393 observations across 1,431 municipalities. **Method:** Two-way fixed effects DiD with municipality and year fixed effects plus canton-specific linear year trends; standard errors clustered at the canton level. **Sample:** Restricted to municipalities in pure German-speaking (ballot) and pure French/Italian-speaking (administrative) cantons; bilingual cantons (BE, FR, VS, GR) excluded from primary specification. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation of the outcome among ballot-canton municipalities. 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).

B. Distribution of Naturalization Rates

Figure 7 shows the distribution of municipal naturalization rates before and after the ruling, separately for ballot and administrative cantons. The ballot-canton distribution shifts rightward after the ruling, consistent with the estimated positive treatment effect.

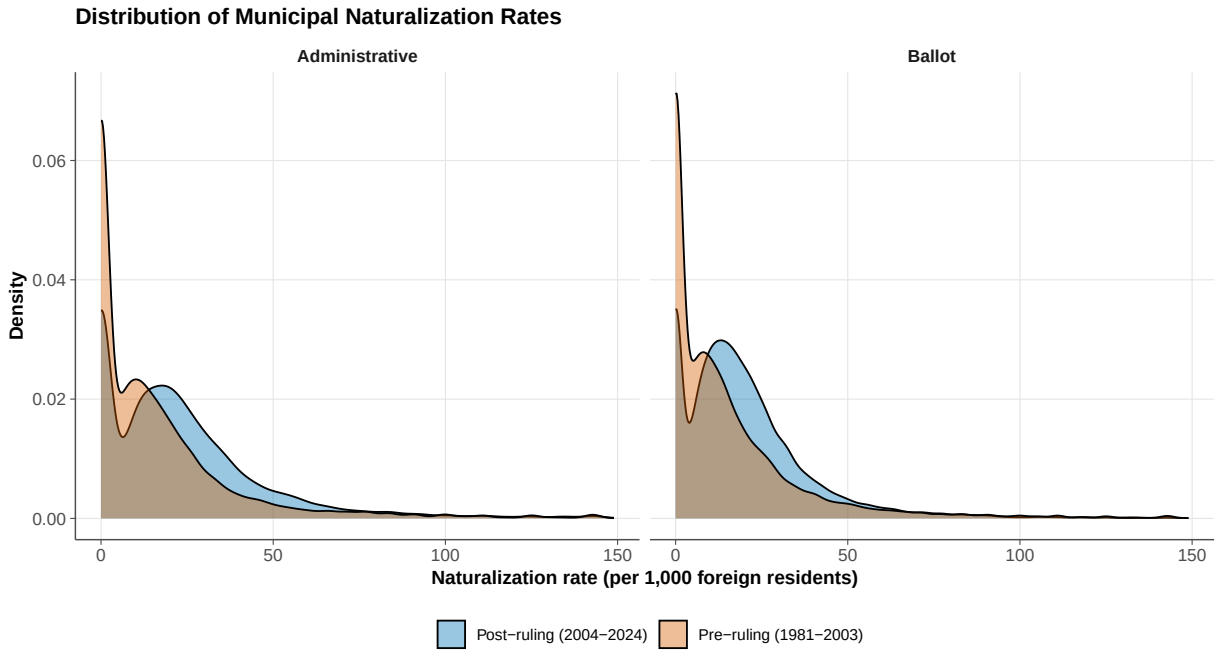

Figure 7: Distribution of Municipal Naturalization Rates, Pre- and Post-Ruling

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