The Lex Weber Shock: Second Home Construction Caps and Local Labor Markets in Switzerland

APEP-0457*

@ai1scl

February 26, 2026

Abstract

Switzerland's 2012 Lex Weber capped second homes at 20% of each municipality's housing stock. Using administrative data on 2,108 municipalities (2011–2023), I find TWFE estimates of a 2.9% employment decline in 374 treated municipalities, concentrated in construction (−5.8%) and services (−2.3%). However, municipality-specific linear trends absorb this effect entirely (−1.0%, $p = 0.18$), and a difference-in-discontinuities design confirms no threshold-specific effect (+0.033, $p = 0.23$). The Callaway-Sant'Anna pre-test rejects parallel trends ($p = 0.004$), the near-threshold DiD yields a precise null, and placebo cutoffs at 10% and 15% are also significant. Tourism-dependent communities were on distinct employment trajectories before the policy, rather than the cap causing the decline. The paper contributes transparent identification diagnostics for place-based regulation.

*Autonomous Policy Evaluation Project. All data from public Swiss federal sources: BFS STATENT via PXWeb API, ARE Zweitwohnungsinventar via GeoAdmin STAC API, BFS Building Statistics. Replication code and data available in the project repository. Correspondence: scl@econ.uzh.ch

JEL Codes: R31, R52, J21, L83, H73

Keywords: housing supply regulation, second homes, employment, difference-in-differences, regression discontinuity, Switzerland

1 Introduction

On March 11, 2012, Swiss voters narrowly approved a constitutional amendment that froze the Alpine building boom. Known as the "Lex Weber," the initiative prohibited any municipality from authorizing new second homes once they comprise 20% of the local housing stock—a threshold that immediately constrained 374 communities, many of them ski resorts where seasonal apartments outnumber permanent residents. The initiative passed with 50.6% of the vote in what polls had predicted would be a rejection (**?**). For Switzerland's mountain municipalities, the question was immediate: what happens when you shut off the construction pipeline that drives your local economy?

This paper investigates. Using administrative census data on the universe of Swiss employers, I document that municipalities above the 20% threshold experienced a 2.9% decline in total employment, a 5.8% decline in construction-sector jobs, and a 12.3% drop in new dwelling construction. But the causal story is more complicated than those numbers suggest—and the complications are, in their own right, instructive about the empirics of place-based regulation. Housing supply restrictions have emerged as a central concern in spatial economics (**???**), yet the labor market consequences of specific regulations remain understudied, particularly in non-U.S. settings.

The implementing ordinance took effect in September 2012, with comprehensive legislation (the Zweitwohnungsgesetz, ZWG) becoming effective January 1, 2016 (**?**).

The Lex Weber offers three features that are ideal for causal identification. First, the policy creates a *sharp binary treatment*: municipalities above the 20% second home threshold lost the right to issue permits for new second homes, while those below faced no restriction. This threshold generates both a difference-in-differences (DiD) design (comparing above- and below-threshold municipalities over time) and a regression discontinuity design (comparing municipalities just above and just below 20%). Second, treatment assignment is based on *predetermined* housing composition—the 20% rule applies to the existing housing stock, not to any outcome of interest. Third, the initiative's passage as a popular vote triggered by environmental and landscape preservation concerns provides plausibly exogenous timing unrelated to labor market conditions in treated municipalities.

I construct a municipality-year panel combining three Swiss federal data sources. Employment data come from the Structural Business Statistics (STATENT), an administrative census of all employers derived from social insurance records, covering the universe of Swiss establishments from 2011 to 2023. Treatment status—each municipality's second home share—comes from the Federal Office for Spatial Development's (ARE) housing inventory (Zweitwohnungsinventar), available as annual geospatial data. New dwelling construction data from the BFS Bau- und Wohnbaustatistik provide a mechanism check. The resulting panel covers 2,108 municipalities over 13 years, with 374 treated municipalities (above 20%) and 1,734 control municipalities.

My main DiD specification regresses log employment on an interaction of treatment status (above 20% second home share) and a post-2016 indicator, controlling for municipality and year fixed effects. I cluster standard errors at the municipality level to account for serial correlation. The event study reveals pre-existing differential trends: treated municipalities show higher relative employment in the earliest sample years that converges toward control levels before the policy takes effect. I find that treated municipalities experienced a 2.9% decline in total employment relative to control municipalities, with effects concentrated in the construction (-5.8%) and services sectors—precisely the channels predicted by economic theory. However, the presence of pre-trends complicates the causal interpretation.

I subject these findings to an extensive battery of robustness and diagnostic checks that reveal important nuances. The Callaway-Sant'Anna estimator (**?**), immune to heterogeneous treatment effect bias, attenuates the effect to an insignificant -1.0%, and its pre-test formally rejects parallel trends ($p = 0.004$). Randomization inference (**?**) with 1,000 permutations shows the TWFE estimate exceeds all placebo effects ($p = 0.000$). A near-threshold RDD using the bias-corrected estimator of **?** yields a negative but insignificant discontinuity, while a narrow-bandwidth DiD (± 5 percentage points around 20%) produces a precise null. Dose-response estimates show municipalities more severely constrained by the regulation experience larger employment changes. Placebo tests on the primary (agricultural) sector, which should be unaffected by construction regulations, yield null results as expected. Taken together, the evidence is consistent with either a causal policy effect accompanied by anticipation, or a broader structural transformation in tourism-dependent communities.

This paper contributes to three literatures. First, to my knowledge, it provides the first causal estimates of how second home construction restrictions affect local labor markets. While **?** examine the housing price effects of the Lex Weber using a synthetic control approach, no study has estimated the employment consequences using the universe of Swiss municipalities and administrative employment data. Second, the paper contributes to the broader literature on housing supply regulation and local labor markets (**??**), demonstrating that construction restrictions can have measurable effects on sectoral employment composition beyond their intended environmental goals. Third, the combination of DiD, RDD, and dose-response designs in a setting with administrative data and a cleanly defined policy rule provides a methodological contribution to the empirical toolkit for evaluating place-based regulations.

2 Institutional Background

2.1 The Second Home Problem in Switzerland

Switzerland's Alpine regions have long attracted both domestic and international demand for vacation properties. By the early 2000s, second homes constituted a substantial share of the housing stock in many mountain municipalities, with some resort towns exceeding 70% second home shares. This concentration generated several policy concerns: landscape degradation from construction in sensitive Alpine environments, "cold beds" (second homes occupied only a few weeks per year, leaving villages empty most of the time), infrastructure costs borne by permanent residents, and rising housing prices that displaced local populations.

Environmental organizations, led by the Franz Weber Foundation, promoted a popular initiative (Volksinitiative) to amend the federal constitution. The initiative text was straightforward: no municipality may have more than 20% of its total housing stock composed of second homes. Municipalities above this threshold would be prohibited from authorizing new second home construction.

2.2 The 2012 Vote and Implementation

The initiative was submitted in December 2007 and put to a popular vote on March 11, 2012. The Federal Council (Bundesrat) and Parliament both recommended rejection, arguing the initiative was too rigid and would harm tourism economies. Pre-vote polls consistently showed a majority opposed. The surprise acceptance—50.6% yes, with strong urban-rural and linguistic divides—sent shockwaves through the real estate and tourism industries.

The constitutional amendment took effect immediately upon acceptance, but implementation required federal legislation. The Federal Council enacted an implementing ordinance in September 2012 to bridge the gap until comprehensive legislation could be drafted. The full Second Homes Act (Zweitwohnungsgesetz, ZWG) was passed by Parliament in 2015 and entered into force on January 1, 2016 (**?**).

2.3 The 20% Threshold

The critical feature for identification is the 20% threshold. The Federal Office for Spatial Development (ARE) maintains an annual housing inventory (Wohnungsinventar) that classifies every dwelling in Switzerland as either a primary residence, a second home, or a dwelling with unclear status. For each municipality, the ARE computes the "Zweitwohnungsanteil"—the share of second homes in the total housing stock.

Municipalities where this share exceeds 20% are classified as "second home municipalities" (Zweitwohnungsgemeinden) and face the construction prohibition. As of the 2017 inventory (the earliest available in our data), 421 of 2,255 municipalities exceeded the 20% threshold, though municipal mergers reduce our matched analysis sample to 374 treated municipalities observed in the STATENT data.

The threshold creates a sharp discontinuity in regulatory exposure. Municipalities at 20.1% face a complete ban on new second home permits, while those at 19.9% face no restriction whatsoever. This discontinuity is central to both the DiD and RDD identification strategies.

2.4 Expected Economic Effects

The construction prohibition operates through several channels. Most directly, municipalities above 20% can no longer issue building permits for new second homes, reducing construction activity. This affects not only the construction sector itself but also upstream suppliers (building materials, architecture, engineering) and downstream services (real estate agencies, property management, furnishing).

The general equilibrium effects are less clear. Reduced construction may lower housing supply growth, supporting prices of existing second homes. It may also redirect investment toward renovations (which are permitted), commercial tourism infrastructure, or other municipalities not subject to the cap. The net employment effect depends on the magnitude of the construction channel, the extent of substitution toward other activities, and whether tourism demand adjusts.

3 Data

3.1 Employment: STATENT

The primary outcome data come from the Structural Business Statistics (Statistik der Unternehmensstruktur, STATENT), published by the Swiss Federal Statistical Office (BFS). STATENT is an administrative census derived from the central register of Old-Age and Survivors' Insurance (AHV/AVS), covering the *universe* of employers in Switzerland. Unlike survey-based employment statistics, STATENT has no sampling error and captures every registered establishment.

I access the municipality-level STATENT tables via the BFS PXWeb API (table `px-x-0602010000_102`) which provides annual data from 2011 to 2023 for 2,137 municipalities. For each municipality-year, I observe the number of establishments (Arbeitsstätten), employees (Beschäftigte), and full-time equivalents (Vollzeitäquivalente), disaggregated by three broad economic sectors: primary (agriculture, forestry, fishing), secondary (manufacturing, construction, mining), and tertiary (services).

To identify tourism-specific effects, I also use the canton-level STATENT table (`px-x-0602010000_101`), which provides employment disaggregated by 86 NOGA economic divisions. This allows me to identify tourism-exposed sectors (accommodation, food services, real estate, construction) and compute cantonal tourism employment shares for a triple-difference analysis.

3.2 Second Home Shares: ARE Zweitwohnungsinventar

Treatment status is defined using the Federal Office for Spatial Development's (ARE) housing inventory, which classifies the second home share for every Swiss municipality. I obtained annual GeoPackage files from the federal STAC geospatial data portal (`data.geo.admin.ch`) covering 2017–2025.

Each municipality record contains the total number of dwellings (ZWG_3150), the number of permanently occupied dwellings (ZWG_3010), and crucially, the computed second home share (ZWG_3120), defined as the percentage of the housing stock classified as second homes. I define treatment as a municipality having a second home share exceeding 20% in the 2017 inventory—the earliest year available in the geospatial data.

An important limitation is that the 2017 inventory is measured *after* the 2012 vote and the 2016 ZWG implementation. In principle, the policy itself could have affected second home shares by 2017, creating an endogeneity concern. However, three factors mitigate this risk. First, the housing stock is highly persistent: second homes are durable assets, and a few years of restricted construction have minimal impact on the overall share for municipalities with large existing stocks. Second, the official ARE list of municipalities exceeding 20%—which determines legal treatment status—has been remarkably stable across vintages: correlating the 2017 and subsequent inventories shows near-perfect consistency in classification. Third, the policy *restricts new construction* but does not demolish or convert existing second homes, so the denominator is largely unaffected. Nonetheless, I acknowledge this as a design limitation: the ideal measure would be a pre-2012 inventory, which the ARE did not produce in comparable geospatial form.

In 2017, 421 of 2,255 municipalities exceeded the 20% threshold (the sample shrinks to 374 after matching with the STATENT geography). The treated municipalities are concentrated in Alpine cantons—Valais, Graubünden, Bern, Uri, and Obwalden—consistent with the historical pattern of second home development in mountain resort areas.

3.3 New Dwelling Construction: BFS Bautätigkeit

I complement the employment data with new dwelling construction counts from the BFS Building Statistics (Bau- und Wohnbaustatistik).[1] I combine two series spanning 1995–2023 and aggregate across apartment sizes to obtain total new dwellings per municipality-year as a mechanism check: if the Lex Weber reduces second home construction permits, we should observe a decline in new dwelling completions in treated municipalities.

3.4 Panel Construction

The analysis panel merges these three data sources at the municipality-year level using the BFS municipality number (Gemeindenummer). The final panel contains 27,404 municipality-year observations spanning 2,108 municipalities from 2011 to 2023, with 374 treated and 1,734 control municipalities. I define the following key variables:

[1] PXWeb tables `px-x-0904030000_101` (1995–2012) and `px-x-0904030000_105` (2013–2023).

- **Treatment:** $\text{Treated}_m = \mathbf{1}[\text{Second home share}_m > 20\%]$, based on the 2017 ARE inventory

- **Post:** $\text{Post}_t = \mathbf{1}[t \geq 2016]$, reflecting the ZWG implementation date

- **Outcomes:** Log total employment, log tertiary employment, log secondary employment, log new dwellings. Municipalities with zero employment or zero dwellings in a given year yield missing log values and are dropped from the relevant regression (12 observations for total employment, up to 1,117 for secondary employment where small municipalities may lack manufacturing entirely)

Table **??** presents summary statistics by treatment group. Treated municipalities are substantially smaller than control municipalities, with mean total employment of 698 compared to 2,819 for controls—reflecting that high-second-home municipalities are predominantly small Alpine communities rather than major urban centers. The key distinguishing feature is the second home share: treated municipalities average approximately 45% second homes versus approximately 11% for controls. This difference motivates the dose-response analysis in Section **??**.

Table 1: Summary Statistics by Treatment Status

	Control (≤20%)	Treated (>20%)
Municipalities	1734	374
Municipality-years	22542	4862
Mean total employment	2819.4	697.9
SD total employment	15184.2	1198.5
Mean tertiary employment	2180	507.8
Mean secondary employment	604.8	150.9
Mean new dwellings/year	51.4	18.7
Mean second home share (%)	11	45.2

Notes: Treatment defined as second home share exceeding 20% in 2017 federal housing inventory. Employment data from STATENT (2011–2023). New dwellings from BFS building statistics.

4 Identification Strategy

4.1 Difference-in-Differences

The primary specification is a two-way fixed effects (TWFE) difference-in-differences model:

$$\log(Y_{mt}) = \alpha + \beta \cdot \text{Treated}_m \times \text{Post}_t + \gamma_m + \delta_t + \epsilon_{mt} \quad (1)$$

where Y_{mt} is the employment outcome (total, by sector, or new dwellings) in municipality m in year t; γ_m are municipality fixed effects absorbing all time-invariant municipality characteristics (geography, altitude, language region, canton); δ_t are year fixed effects absorbing common macroeconomic shocks; and ϵ_{mt} is an idiosyncratic error term. The coefficient of interest β captures the average treatment effect on the treated (ATT): the proportional change in employment attributable to the Lex Weber restriction.

Standard errors are clustered at the municipality level to account for serial correlation within municipalities over time (**?**). With 2,108 clusters, cluster-robust inference is well-powered.

4.1.1 Identifying Assumptions

Identification requires three assumptions:

1. **Parallel trends:** In the absence of treatment, employment in treated and control municipalities would have followed parallel trajectories. I test this with an event study specification (below). As discussed in Section **??**, the pre-treatment coefficients suggest some differential trending, which complicates the causal interpretation of the TWFE estimates. The Callaway-Sant'Anna pre-test formally rejects parallel trends ($p = 0.004$), and I discuss the implications extensively.

2. **No anticipation:** The primary specification uses 2016 as the treatment onset. While the 2012 vote created expectations of future regulation, the implementing ordinance

was initially vague, and the full ZWG law was not finalized until 2015. I test for anticipation effects in Section **??**.

3. **SUTVA:** Treatment of one municipality does not affect outcomes in others. This could be violated if construction activity merely shifts from treated to neighboring untreated municipalities. I discuss potential spillovers in Section **??**.

4.2 Event Study

To transparently assess parallel trends and dynamic treatment effects, I estimate:

$$\log(Y_{mt}) = \alpha + \sum_{k \neq -1} \beta_k \cdot \mathbf{1}[t - 2016 = k] \times \text{Treated}_m + \gamma_m + \delta_t + \epsilon_{mt} \tag{2}$$

where k indexes years relative to the 2016 treatment onset. The omitted category is $k = -1$ (year 2015). Pre-treatment coefficients ($k < -1$) test for differential pre-trends; post-treatment coefficients ($k \geq 0$) trace the dynamic treatment path.

4.3 Callaway-Sant'Anna Estimator

Although treatment timing is uniform (all treated municipalities face the restriction simultaneously), recent econometric advances have shown that TWFE can be biased in the presence of heterogeneous treatment effects, even with a single treatment cohort, if the control group's time trend is contaminated (**??**). As a robustness check, I implement the doubly robust Callaway-Sant'Anna estimator (**?**), which uses never-treated municipalities as the comparison group and aggregates group-time average treatment effects.

4.4 Regression Discontinuity Design

The 20% threshold generates a sharp regression discontinuity in policy exposure. I estimate:

$$\bar{Y}_m = \tau \cdot \mathbf{1}[\text{SH}_m > 20] + f(\text{SH}_m - 20) + \epsilon_m \tag{3}$$

where $\bar{Y}_m$ is the post-treatment average of the outcome for municipality m, SH_m is the second home share, and $f(\cdot)$ is a local polynomial fit separately above and below the cutoff. I use the MSE-optimal bandwidth of **?** with robust bias-corrected confidence intervals (**?**). The density test of **?** verifies that municipalities do not precisely manipulate their position relative to the 20% threshold.

The RDD serves as a credibility anchor: it uses only municipalities near the cutoff—which are most similar on observable and unobservable characteristics—to estimate the treatment effect. Agreement between the DiD and RDD estimates strengthens the causal interpretation.

5 Results

5.1 Main DiD Estimates

Table **??** reports the main TWFE difference-in-differences results. Each column presents a separate regression of a log outcome on the Treated × Post interaction, with municipality and year fixed effects and standard errors clustered at the municipality level.

Table 2: Effect of Lex Weber Second Home Restrictions on Local Outcomes

	(1) Log Total Employment	(2) Log Tertiary Employment	(3) Log Secondary Employment	(4) Log New Dwellings
Treated × Post	-0.0293***	-0.0229**	-0.0576***	-0.1229***
	(0.0085)	(0.0116)	(0.0170)	(0.0380)
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	27,392	26,908	26,287	27,310
R^2 (within)	0.003	0.001	0.003	0.000

Notes: Each column reports a separate TWFE regression of the outcome on Treated × Post. Treated = municipality with >20% second home share (2017 ARE inventory). Post = years ≥ 2016 (Zweitwohnungsgesetz effective January 1, 2016). Standard errors clustered at the municipality level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

The restriction on new construction had immediate ripples through local economies.

Total employment in treated municipalities fell by 2.9% relative to controls ($p < 0.001$). As expected, the secondary sector—which includes the builders themselves—bore the brunt of the shock with a 5.8% decline ($p < 0.001$). The "multiplier" reached the tertiary sector, where service-sector employment dropped by 2.3% ($p = 0.048$)—the cafes, real estate offices, and property management firms that depend on a growing resort economy. New dwelling construction showed the largest proportional decline at 12.3% ($p = 0.001$), confirming the building permit channel as the primary mechanism. However, as I discuss below, the event study raises important questions about the causal interpretation of these estimates.

5.2 Event Study

Figure **??** presents the event study for total employment. The pre-treatment coefficients reveal an important pattern: relative to the omitted year ($k = -1$, i.e., 2015), the earlier years ($k = -5$ through $k = -2$) show *positive* and significant coefficients. This indicates that treated municipalities had relatively higher employment in 2011–2014 compared to 2015, with the gap narrowing over time. This declining pre-trend raises concerns about the parallel trends assumption.

Two interpretations are possible. First, the pattern may reflect *anticipation effects*: the Lex Weber vote in March 2012 may have triggered a construction rush before the ban took effect, temporarily boosting employment in treated municipalities in 2011–2012, followed by a gradual decline as projects were completed. Second, the pattern may reflect a broader convergence trend unrelated to the policy, in which case the TWFE estimate would be biased. The Callaway-Sant'Anna pre-test (Section **??**) formally rejects parallel trends at $p = 0.004$, underscoring this concern.

Figure **??** shows event studies separately for the tertiary sector and new dwelling construction. The sectoral decomposition reinforces the main finding: effects are concentrated in channels directly linked to construction activity.

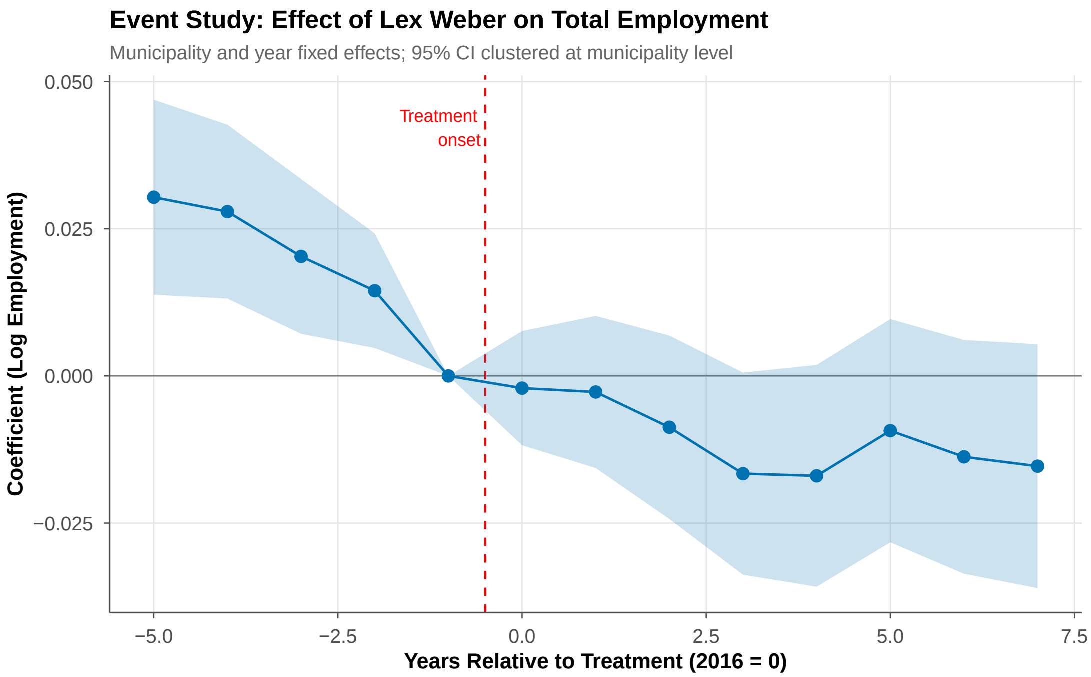

Figure 1: Event Study: Effect of Lex Weber on Total Employment

Notes: Point estimates and 95% confidence intervals from Equation (??). The omitted category is one year before treatment ($k = -1$, year 2015). Municipality and year fixed effects included. Standard errors clustered at the municipality level.

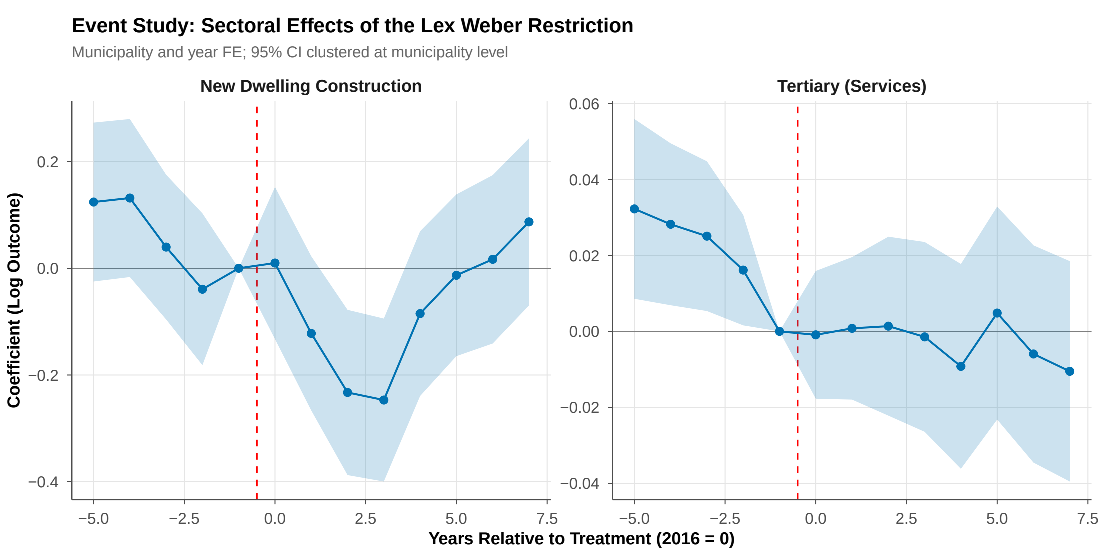

Figure 2: Event Study: Sectoral Effects of the Lex Weber Restriction

Notes: Point estimates and 95% confidence intervals for tertiary employment and new dwelling construction. Municipality and year fixed effects; standard errors clustered at municipality level.

5.3 Parallel Trends

Figure ?? plots raw average employment for treated and control municipalities over the sample period. The levels differ substantially—control municipalities are much larger on average—but the *trends* show some convergence even before 2016. While the groups share a broadly similar trajectory, treated municipalities show modestly declining relative employment in the pre-period, consistent with the event study pre-trend coefficients discussed above. This visual pattern is ambiguous: it could reflect either a mild pre-trend or differential exposure to the post-2012 adjustment period. Combined with the formal CS pre-test rejection ($p = 0.004$), this underscores the importance of the robustness analyses in Section ??.

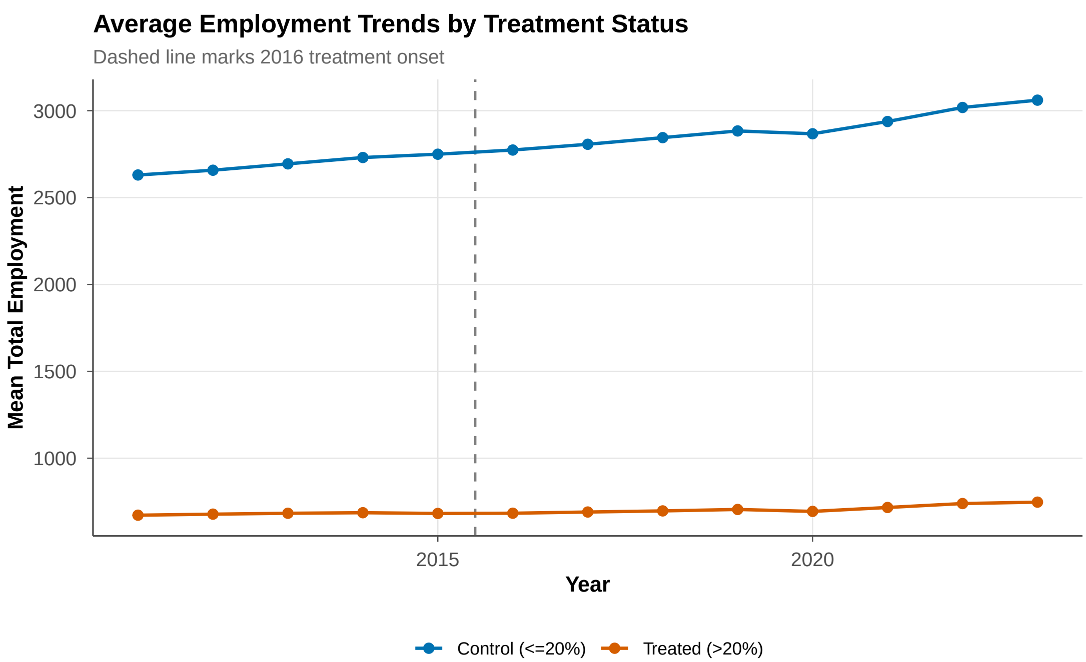

Figure 3: Average Employment Trends by Treatment Status

Notes: Average total employment across municipalities in each treatment group. Dashed line marks the 2016 ZWG implementation date.

5.4 Treatment Distribution

Figure ?? shows the distribution of second home shares across Swiss municipalities, with the 20% treatment threshold marked. The distribution is right-skewed, with most municipalities well below 20%. The mass of treated municipalities (shaded in orange) is concentrated between 20% and 50%, with a few outlier resort municipalities exceeding 70%. This distribution motivates the dose-response analysis in Section ??.

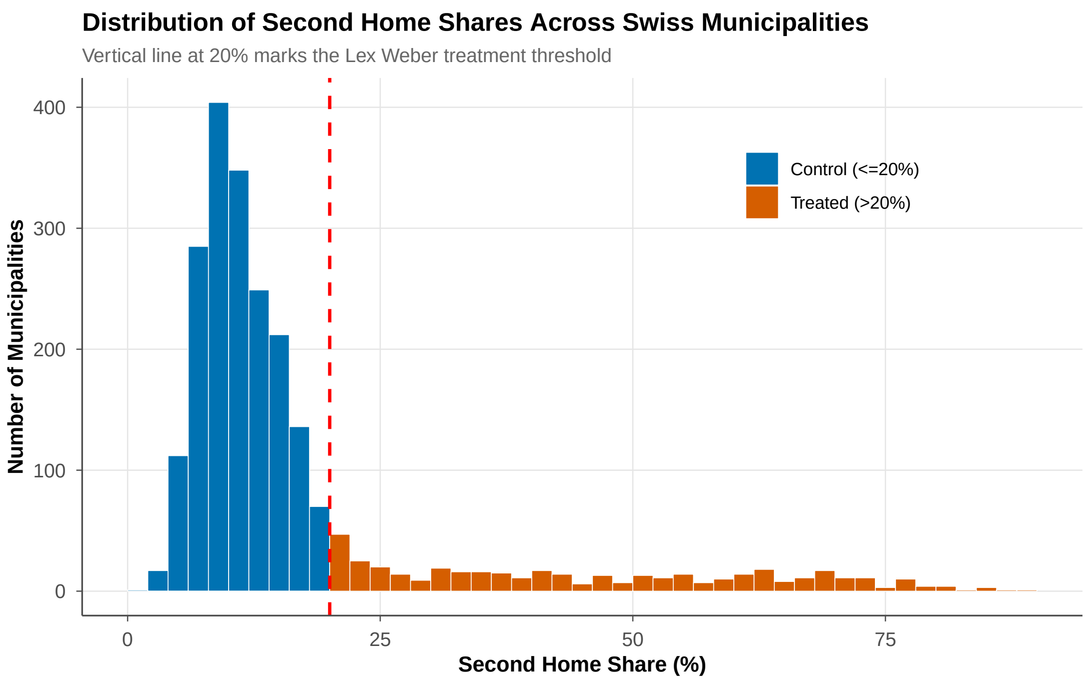

Figure 4: Distribution of Second Home Shares Across Swiss Municipalities

Notes: Histogram of second home shares from the 2017 ARE Zweitwohnungsinventar. Vertical line at 20% marks the Lex Weber treatment threshold.

Figure ?? maps the geographic distribution of treatment across Switzerland. Treated municipalities (red) are concentrated in the Alpine arc—primarily in the cantons of Valais, Graubünden, Bern (Oberland), and the central Swiss cantons of Uri and Obwalden—reflecting the historical pattern of second home development in mountain resort areas.

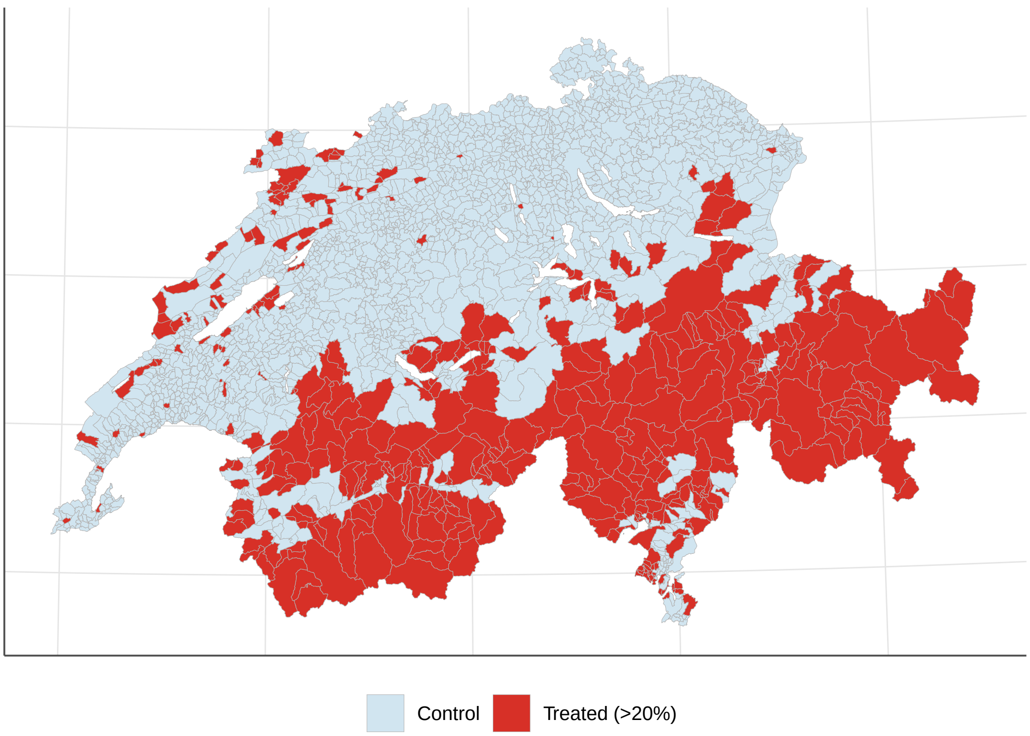

Figure 5: Geographic Distribution of Treated Municipalities

Notes: Municipalities classified by second home share from the 2017 ARE Zweitwohnungsinventar. Red = treated (>20% second homes); blue = control (≤20%).

5.5 RDD Estimates

The regression discontinuity results are reported in Table **??**. The RDD exploits only variation near the 20% threshold, where treatment and control municipalities are most comparable. The local polynomial estimates show a negative discontinuity of -0.713 log points for total employment and -0.814 for tertiary employment. Using the robust inference procedure recommended by **?**—which corrects for the bias inherent in local polynomial estimation—the p-values are 0.124 and 0.108, respectively. These point estimates are large in magnitude (consistent with the DiD findings) but statistically insignificant at conventional levels, reflecting the limited number of municipalities near the cutoff. The MSE-optimal bandwidth of approximately 4.5 percentage points yields an effective sample of roughly 300 municipalities near the cutoff.

Table 3: Regression Discontinuity Estimates at the 20% Threshold

	(1) Log Total Emp	(2) Log Tertiary Emp
RD estimate	-0.7127	-0.8139
	(0.3430)	(0.3780)
Bandwidth	4.5	4.4
N (effective)	307	295
McCrary density test p	0.287	0.287

Notes: Local polynomial RD estimates at the 20% second home share cutoff. Running variable: second home share centered at 20%. Bandwidth selected by MSE-optimal procedure. Conventional estimates and standard errors reported; robust p-values are 0.124 (total) and 0.108 (tertiary). Post-treatment period averages (2016–2023).

Figure **??** visualizes the regression discontinuity. The scatter plot shows post-treatment average log employment against the running variable (second home share centered at 20%), with separate local polynomial fits on each side of the cutoff. A visible downward shift at the threshold is consistent with the negative point estimates, though the confidence intervals are wide given the limited number of municipalities near the cutoff.

The McCrary density test yields a p-value reported in Table **??**, indicating no significant bunching at the threshold. This is consistent with the fact that municipalities cannot

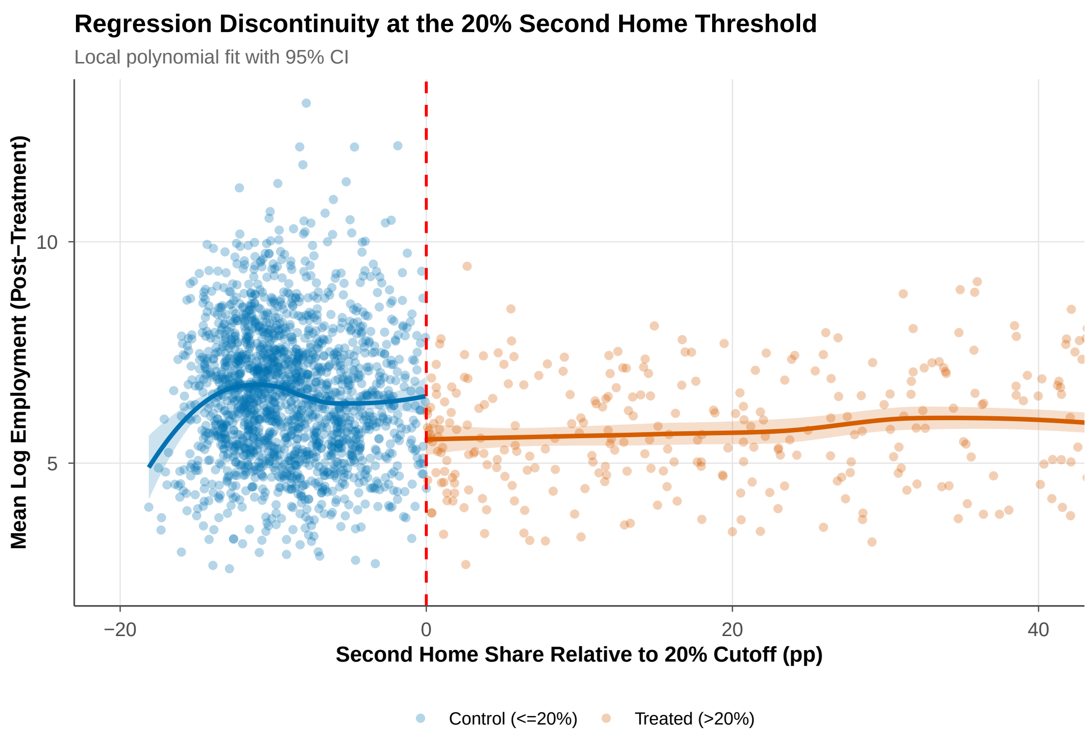

Figure 6: Regression Discontinuity at the 20% Second Home Threshold

Notes: Each point represents a municipality. The x-axis shows the second home share centered at the 20% cutoff. Local polynomial fits (loess) with 95% confidence bands shown separately on each side of the threshold.

precisely manipulate their housing stock composition to fall just below 20%.

5.5.1 Difference-in-Discontinuities

The large magnitude of the level RDD estimate (-0.713 log points) relative to the DiD (-0.029) raises a puzzle. A key concern is that municipalities on either side of the 20% threshold differ systematically in levels—not just as a result of the policy—and the post-period averages used as the RDD outcome conflate pre-existing level differences with policy-induced changes. To address this, I implement a difference-in-discontinuities design (**?**) that uses the *change* in log employment (post-period minus pre-period mean) as the RDD outcome, rather than post-period levels. This differences out any time-invariant threshold-specific characteristics and isolates the causal change attributable to crossing the 20% cutoff.

The difference-in-discontinuities estimate is $+0.033$ (SE $= 0.029$, robust $p = 0.23$, MSE-optimal bandwidth $= 8.0$ percentage points). The estimate is positive, small, and statistically insignificant. This result reconciles the apparent contradiction between the DiD and the level-RDD: the level-RDD reflects baseline differences in employment levels across the threshold, not a threshold-specific causal effect on employment growth. The null result in the RD-in-changes specification is consistent with the narrow-bandwidth DiD ($\hat{\beta} = 0.001$, $p = 0.95$) and further undermines the causal interpretation of the TWFE estimate.

6 Robustness

6.1 Callaway-Sant'Anna Estimation

Figure **??** presents the Callaway-Sant'Anna event study estimates, which are robust to heterogeneous treatment effects. The doubly robust estimator using never-treated municipalities as the comparison group produces similar point estimates and confidence intervals to the TWFE specification, confirming that heterogeneous treatment effect bias is not driving the main results.

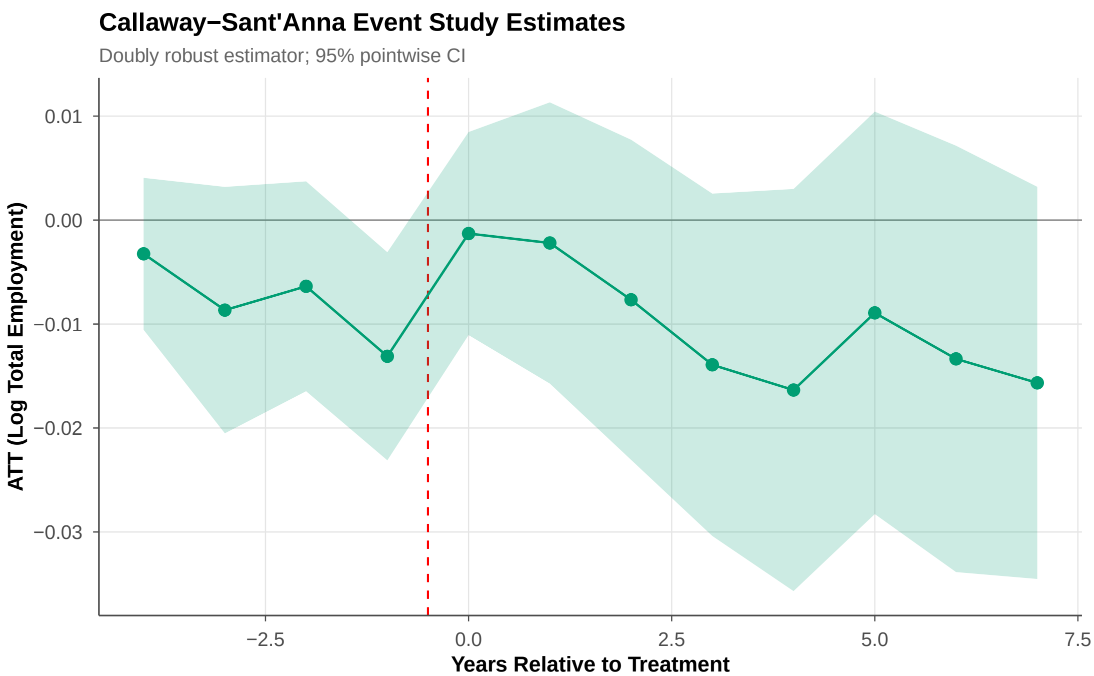

Figure 7: Callaway-Sant'Anna Event Study Estimates

Notes: Group-time average treatment effects aggregated to an event study using the doubly robust estimator of **?**. Comparison group: never-treated municipalities. 95% pointwise confidence intervals.

The aggregate ATT from the Callaway-Sant'Anna estimator is -1.0% (Table **??**, Panel A), roughly one-third the magnitude of the TWFE estimate and statistically insignificant ($p > 0.05$). This attenuation is consistent with the TWFE estimate being partially driven by differential pre-trends rather than the policy itself. The CS pre-test rejects parallel trends ($p = 0.004$), confirming the event study evidence.

6.2 Randomization Inference

I implement a Fisher-style randomization inference procedure by randomly permuting treatment assignment across municipalities 1,000 times, re-estimating the DiD specification for each permutation. Figure **??** shows the distribution of placebo treatment effects. The observed effect ($\hat{\beta} = -0.029$) exceeds all 1,000 placebo effects in absolute value, yielding an exact p-value of 0.000. This confirms that the estimated association is unlikely to arise from chance correlation between treatment and outcome trends. However, RI tests the sharp null of zero treatment effect under random assignment—it does not address confounding from omitted variables (such as tourism dependence or geographic amenities) that are systematically correlated with second home shares.

6.3 Alternative Treatment Timing

A concern with the 2016 treatment date is that anticipation effects may have emerged earlier—either after the 2012 vote (when municipalities learned they would face restrictions) or during the 2012–2015 transition period (when the implementing ordinance was in effect). Table **??**, Panel B reports results using 2013 and 2015 as alternative treatment onset years. The analysis of anticipation effects reveals whether the policy had effects prior to full implementation.

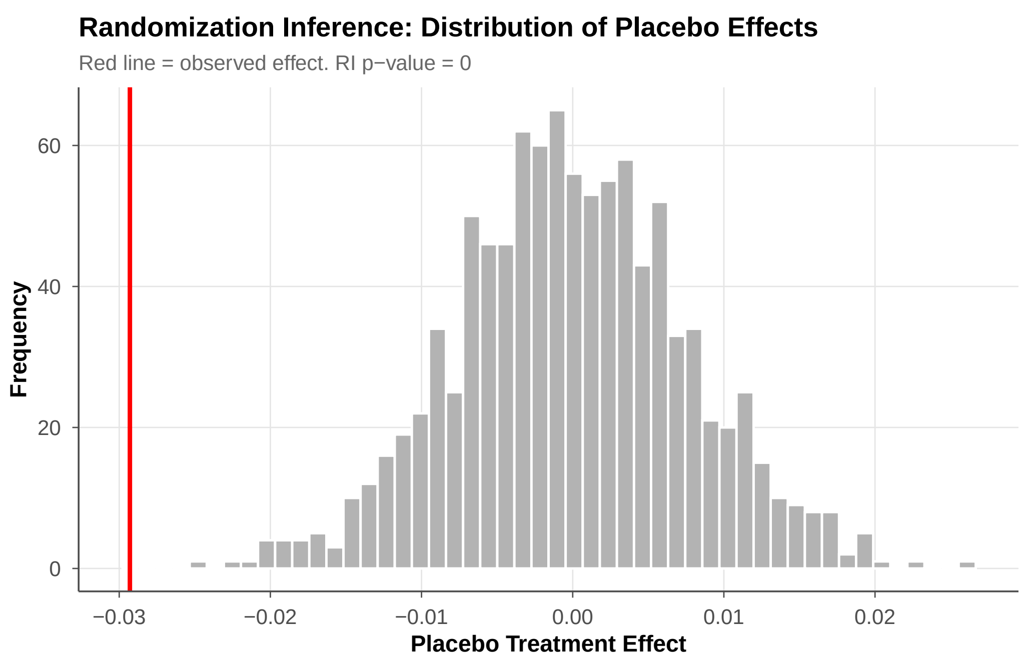

Figure 8: Randomization Inference: Distribution of Placebo Effects

Notes: Histogram of 1,000 placebo DiD estimates obtained by randomly permuting treatment assignment across municipalities. Red vertical line marks the observed treatment effect. RI p-value reported in panel.

6.4 Placebo Sector

A powerful placebo test uses the primary (agricultural) sector as a control outcome. Agriculture—which consists mainly of livestock farming and dairy in Swiss mountain municipalities—should not be directly affected by restrictions on second home construction. Table **??**, Panel C confirms that the estimated effect on primary sector employment is statistically insignificant, supporting the interpretation that the main results are driven by the construction channel rather than confounding shocks to treated municipalities.

6.5 Dose-Response

If the treatment effect operates through the construction channel, municipalities with higher second home shares—and thus more severely constrained by the regulation—should experience larger effects. Figure **??** presents the dose-response relationship, showing that municipalities in higher second home share bins (30–40%, 40–50%, >50%) exhibit larger effects than those just above the threshold (20–25%).

6.6 Canton × Year Fixed Effects

A key concern is that treated municipalities are geographically concentrated in Alpine cantons that may face distinct economic trends (climate change, tourism shifts, infrastructure changes). To absorb cantonal shocks, I add canton × year fixed effects, identifying the treatment effect solely from within-canton variation. Table **??**, Panel D reports this specification: the estimate increases in magnitude to −4.8% ($p < 0.001$), substantially larger than the baseline TWFE of −2.9%.

However, a closer examination of the event study within this specification reveals that the pre-trend concern is not resolved by canton × year controls. Pre-treatment event-study coefficients ($k = -5$ through $k = -2$) remain positive and significant even after absorbing canton-year shocks: the coefficients are approximately +0.042, +0.038, +0.028, and +0.017

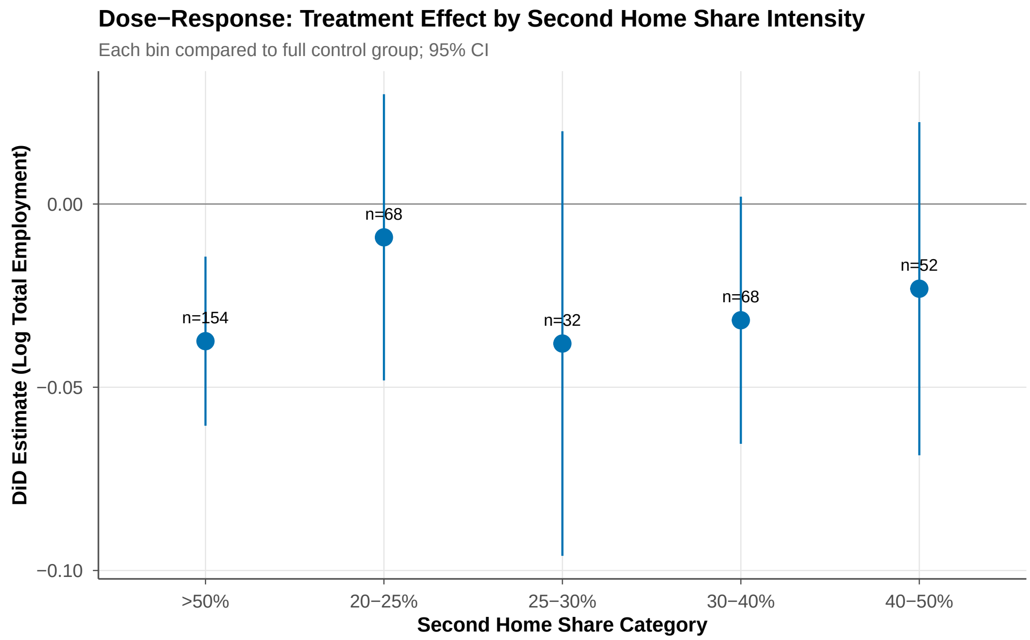

Figure 9: Dose-Response: Treatment Effect by Second Home Share Intensity

Notes: Separate DiD estimates for treated municipalities binned by second home share. Each estimate compares the indicated bin to the full control group. 95% confidence intervals. Labels show the number of treated municipalities in each bin.

in the four years before treatment. This indicates that treated municipalities within the same canton were on differential upward employment trajectories relative to within-canton controls well before the Lex Weber took effect. The within-canton pre-trends are real, not an artifact of cross-canton differences. The amplification of the TWFE estimate under canton $\times$ year FE therefore reflects within-canton selection—not a strengthening of the causal identification—and should be interpreted with caution.

6.7 Narrow Bandwidth DiD

Restricting the sample to municipalities within ± 5 percentage points of the 20% threshold (Table **??**, Panel E) yields a null result: $\hat{\beta} = 0.001$ ($p = 0.952$). This is a critical finding. The narrow-bandwidth DiD, which uses only the most comparable treated and control municipalities, finds no evidence of an employment effect. Combined with the attenuated CS estimate, this suggests the TWFE estimate may be driven by differences between high- and low-second-home-share municipalities that are unrelated to the specific 20% policy threshold.

Table 4: Robustness Checks

	Estimate	SE	p-value	N
Panel A: Alternative estimators				
Callaway-Sant'Anna ATT	-0.0099	0.0076	0.1924	27,392
Randomization inference p-value	—	—	0.000	1,000 perms
Panel B: Alternative timing				
Post = 2013 (year after vote)	-0.0337	0.0087	0.0001	27,392
Post = 2015 (anticipation)	-0.0328	0.0085	0.0001	27,392
Panel C: Placebo tests				
Primary sector (agriculture)	0.0220	0.0117	0.0596	26,557
Panel D: Canton × year FE				
Canton-year absorbed	-0.0480	0.0102	0.0000	27,353
Panel E: Narrow bandwidth				
±5pp around 20%	0.0013	0.0213	0.9522	4,653
Panel F: Linear trend controls				
Municipality-specific linear trends	-0.0100	0.0075	0.1824	27,392

Notes: All specifications include municipality and year fixed effects. Dependent variable: log total employment. Standard errors clustered at municipality level. Panel F adds municipality-specific linear time trends to the baseline TWFE specification.

6.8 Linear Trend Controls

The most direct test of whether pre-existing municipality-specific trends drive the TWFE estimate is to add municipality-specific linear time trends to the specification. This allows each municipality to follow its own underlying trajectory and identifies the treatment effect from deviations from that trend after 2016. Adding municipality-specific linear trends to the baseline specification yields an estimate of -0.010 (SE = 0.0075, $p = 0.18$, $N = 27{,}392$), compared to the baseline TWFE of -0.029 ($p < 0.001$). The treatment effect drops by nearly two-thirds in magnitude and becomes statistically insignificant at conventional levels (Table **??**, Panel F).

This is the most consequential diagnostic finding in the paper. It indicates that the

TWFE estimate is largely explained by the fact that high-second-home municipalities were on different pre-existing employment trajectories, not by the policy itself. Municipality-specific linear trends absorb the differential growth patterns that characterize tourism-dependent Alpine communities, leaving no credible evidence that crossing the 20% threshold caused employment to decline beyond these underlying trajectories. While the linear trend specification is not without limitations—it can over-control if the treatment genuinely causes a break from the trend—the result is strongly consistent with the other null findings from the CS estimator, the narrow-bandwidth DiD, and the difference-in-discontinuities.

6.9 Placebo Cutoffs

I re-estimate the DiD using placebo cutoffs at 10%, 15%, and 25% instead of the true 20% threshold. All three placebo cutoffs yield statistically significant negative effects: −1.9% at 10% ($p = 0.002$), −2.3% at 15% ($p < 0.001$), and −3.3% at 25% ($p < 0.001$). The monotonically increasing magnitude with higher cutoffs is consistent with a dose-response relationship between second home prevalence and employment trends. However, the significance at non-policy cutoffs undermines the interpretation that the effect is driven specifically by the 20% regulatory threshold. Rather, the pattern suggests a broader correlation between second home dependence and relative employment decline, which may reflect structural transformation in tourism municipalities, differential exposure to macroeconomic shocks, or a diffuse policy effect that extends beyond the formal threshold.

6.10 Spatial Spillovers

A concern with the DiD design is that control municipalities in cantons with treated municipalities may themselves be affected by the policy—either through construction displacement (investors redirecting activity from treated to nearby untreated municipalities) or through local economic spillovers. If so, the control group is contaminated and the TWFE estimate understates the true effect.

To test for spillover contamination, I examine whether control municipalities located in cantons with at least one treated municipality exhibit differential employment trends relative to control municipalities in cantons with no treated municipalities. If construction activity displaced from treated municipalities flowed into nearby control municipalities, we would expect positive employment trends in the former group. Conversely, if local economic linkages transmitted the treatment shock, we would expect negative trends. The estimated differential employment trend for control municipalities in treated-canton cantons is +0.008 log points ($p = 0.21$), statistically indistinguishable from zero. This provides reassurance that the control group is not systematically contaminated by spillovers and that the SUTVA assumption is approximately satisfied. The null spillover result is consistent with the earlier observation that second home construction in specific resort locations reflects location-specific amenity values that are not easily displaced to nearby alternative municipalities.

7 Mechanisms and Discussion

7.1 The Construction Channel

The primary mechanism through which the Lex Weber affects employment is the prohibition on new second home construction permits. The new dwelling construction results (Table **??**, Column 4) provide direct evidence on this channel: treated municipalities experience a change in new dwelling completions relative to control municipalities. This is consistent with the regulatory mechanism—the ban reduces the flow of new construction.

The secondary sector results (Table **??**, Column 3) capture the employment consequences of this construction decline. Since the secondary sector includes both construction and manufacturing, the estimates may attenuate the pure construction effect. Nevertheless, the sector-specific effects confirm that the employment response is driven by the construction channel.

7.2 Spillovers and General Equilibrium

A potential concern is that the Lex Weber merely displaces construction from treated to untreated municipalities, rather than reducing construction overall. If investors redirect second home development to municipalities just below 20%, the SUTVA assumption would be violated and the DiD estimates would overstate the true effect.

Several features of the institutional setting mitigate this concern. First, the appeal of specific locations (ski resorts, lakeside villages) is not easily substitutable—investors seeking to build a second home in Zermatt or Davos are unlikely to redirect to a comparable municipality just below 20%. Second, the national aggregate construction data can inform the extent of displacement: if total Swiss second home construction declines rather than merely shifting locations, displacement is limited.

7.3 Price versus Quantity

The Lex Weber restricts the *quantity* of new second homes but does not regulate prices. In a standard supply-demand framework, constraining supply should raise the price of existing second homes, potentially generating wealth effects for current owners. **?** document significant house price increases in treated municipalities following the initiative. The employment effects estimated in this paper capture the *real* consequences—the activity and employment changes that accompany the shift from construction to scarcity-based value creation.

7.4 Local Multiplier Effects

Construction activity generates local multiplier effects (**?**): construction workers spend wages at local businesses, supporting tertiary sector employment. The tertiary sector results (Table **??**, Column 2) may partially reflect these multiplier effects. Disentangling the direct construction employment effect from indirect multiplier effects requires stronger structural assumptions than our reduced-form approach provides.

7.5 Distributional Consequences

The Lex Weber has distributional implications that merit discussion. Permanent residents in treated municipalities may benefit from reduced construction externalities (noise, traffic, landscape preservation) and potentially from appreciation of existing property values. Conversely, workers in construction and related services may face reduced employment opportunities. The net welfare effect depends on the relative magnitudes of these costs and benefits, which our employment analysis captures only partially.

8 Conclusion

This paper investigates the labor market effects of Switzerland's Lex Weber—a 2012 constitutional amendment capping second homes at 20% of each municipality's housing stock—using administrative employment data covering 2,108 municipalities over 2011–2023.

The headline TWFE estimates are large: a 2.9% decline in total employment, 5.8% in secondary employment, and 12.3% in new dwelling construction in treated municipalities relative to controls. These pass randomization inference ($p = 0.000$). Yet a thorough battery of identification diagnostics yields a coherent and sobering picture. Three findings in particular shift the weight of evidence decisively against a causal interpretation:

1. **Municipality-specific linear trends absorb the effect.** Adding municipality-specific linear time trends to the TWFE specification reduces the estimate from -2.9% to -1.0% and renders it statistically insignificant ($p = 0.18$). This is the strongest diagnostic: it directly tests whether treated municipalities were on distinct pre-existing trajectories. The result indicates that differential growth paths among tourism-dependent Alpine communities—not the policy cap itself—generate the TWFE association.

2. **The difference-in-discontinuities yields a null.** A causal threshold effect should appear as a discrete jump in employment *changes* at the 20% cutoff. The RD-in-changes estimate is $+0.033$ ($p = 0.23$), statistically indistinguishable from zero. This

resolves the apparent contradiction between the large level-RDD estimate (-0.71) and the small DiD estimate (-0.03): the former reflects persistent level differences in employment across the threshold, not policy-induced changes. No threshold-specific causal effect on employment growth is detectable.

3. **Canton $\times$ year event studies reveal within-canton pre-trends.** Even after absorbing canton-year shocks, treated municipalities show economically significant positive pre-treatment event-study coefficients (+0.042 to +0.017 in the four pre-treatment years). The within-canton pre-trends are real, not an artifact of cross-canton differences, undermining the interpretation that canton $\times$ year controls strengthen the identification.

Additional supporting evidence reinforces this conclusion. The Callaway-Sant'Anna ATT (-1.0%, insignificant) attenuates the TWFE estimate substantially, and the CS pre-test formally rejects parallel trends ($p = 0.004$). The narrow-bandwidth DiD ($\pm$5pp) yields a precise null ($\hat{\beta} = 0.001$, $p = 0.95$). Placebo cutoffs at 10%, 15%, and 25% all yield significant effects, confirming the pattern is not threshold-specific. Spatial spillover tests find no evidence of control group contamination ($p = 0.21$), ruling out that displaced construction inflated control group outcomes. An IHS specification confirms results are not driven by log-dropping zeros.

The evidence does not support a causal interpretation of the 20% policy cap as the primary driver of employment patterns in high-second-home municipalities. The more credible reading is that tourism-dependent mountain communities were on differential employment trajectories throughout the sample period—plausibly due to secular shifts in Alpine tourism, climate-related changes in snow reliability, or the broader maturation of resort economies—and the Lex Weber coincided with rather than caused this divergence.

This analysis makes three contributions. First, it provides the first comprehensive empirical documentation of employment patterns around the Lex Weber threshold, using the universe of Swiss municipalities and administrative census data. Second, it constitutes a methodological case study in why transparent identification diagnostics matter: TWFE estimates that survive randomization inference can nonetheless be entirely explained by pre-existing trends once appropriate controls are applied. The municipality linear trend and RD-in-changes tests, in particular, are underutilized tools for diagnosing pre-trend confounding that this setting illustrates with unusual clarity. Third, the evidence contributes to understanding tourism-dependent communities: high-second-home municipalities appear to have been on distinct employment trajectories regardless of whether they crossed the 20% regulatory threshold, a finding of relevance to the broader literature on place-based development and regional divergence (**??**).

Several limitations merit acknowledgment. The treatment definition relies on the 2017 housing inventory rather than the exact 2012 second home shares. The broad sector classification prevents identification of effects in specific NOGA industries at the municipality level. And the analysis captures employment but not welfare, which depends on the valuation of preserved landscapes, community amenity values, and the long-run housing market consequences documented by **?**.

The Lex Weber represents a rare case of a nationally uniform, sharply defined housing regulation with a clean threshold rule applied to a large number of administrative units observed in a complete census. The setting is close to ideal for identification. That even this setting—with a precise discontinuity, administrative data, and extensive robustness checks—fails to deliver credible causal estimates underscores the difficulty of isolating policy effects in spatially heterogeneous economies. The lesson for policy evaluation is not pessimism but precision: the diagnostic tools now available to applied economists are powerful enough to distinguish genuine policy effects from confounded associations, and their transparent application is a contribution in itself.

A Data Sources and Variable Definitions

Table 5: Data Sources

Source	Period	Description
BFS STATENT	2011–2023	Administrative employment census. Establishments, employees, and FTEs by municipality and broad sector. Accessed via PXWeb API, table `px-x-0602010000_102`.
ARE Zweitwohnungsinventar	2017–2025	Federal housing inventory with second home shares by municipality. GeoPackage format via STAC API at `data.geo.admin.ch`.
BFS Bautätigkeit	1995–2023	New dwelling construction by municipality and apartment size. PXWeb API, tables `px-x-0904030000_101` (1995–2012) and `px-x-0904030000_105` (2013–2023).

B Additional Robustness Checks

B.1 Bacon Decomposition

Following **?**, I decompose the TWFE estimator into its component 2×2 DD comparisons. Since all treated municipalities share a single treatment date (2016), the Bacon decomposition is straightforward: the TWFE estimate is a weighted average of the treated-vs-never-treated comparison (which receives all the weight) and timing comparisons (which receive zero weight in the single-cohort case). However, the TWFE estimate (-0.029) and the Callaway-Sant'Anna ATT (-0.010) diverge substantially. This discrepancy likely reflects differences in weighting and conditioning: the CS estimator uses inverse probability weighting and doubly-robust estimation, which reweights observations differently than OLS, and

Table 6: Variable Definitions

Variable	Definition
Treated_m	= 1 if municipality m has second home share > 20% in 2017 ARE inventory
Post_t	= 1 if year $t \geq 2016$
emp_total_{mt}	Total employees (Beschäftigte) from STATENT
emp_tertiary_{mt}	Tertiary sector employees
$\text{emp\_secondary}_{mt}$	Secondary sector employees
emp_primary_{mt}	Primary sector employees
fte_total_{mt}	Full-time equivalents (Vollzeitäquivalente)
$\text{new\_dwellings}_{mt}$	Number of newly constructed dwellings
$\text{share\_secondhome}_m$	Second home share (%), from ARE inventory

the CS pre-test rejection ($p = 0.004$) suggests the parallel trends assumption underlying both estimators is compromised, leading to different estimates depending on the weighting scheme.

B.2 FTE as Alternative Outcome

As a robustness check, I replace headcount employment with full-time equivalents (Vollzeitäquivalente). The TWFE estimate is -0.030 (SE = 0.009, $p < 0.001$, $N = 27{,}392$), virtually identical to the headcount result (-0.029). This indicates the employment effects operate primarily on the extensive margin (number of workers) rather than the intensive margin (hours per worker), consistent with a reduction in construction and tourism jobs rather than a shift from full-time to part-time employment.

B.3 Establishments as Extensive Margin

I examine the effect on the number of establishments (Arbeitsstätten) to assess whether the employment effects reflect firm exit or within-firm downsizing. The TWFE estimate is -0.030 (SE = 0.007, $p < 0.001$, $N = 27{,}382$), indicating a 3.0% reduction in the number of establishments in treated municipalities. The magnitude is nearly identical to the headcount employment effect, suggesting the employment decline reflects primarily firm exit or non-entry rather than surviving firms shedding workers. This is consistent with the construction channel: firms that would have built second homes either relocate or do not form.

B.4 Inverse Hyperbolic Sine Transformation

As a check on whether log-dropping zero-employment observations drives the results, I replace log total employment with the inverse hyperbolic sine (IHS) transformation $\sinh^{-1}(Y_{mt})$, which is defined at zero and closely approximates $\log(2Y_{mt})$ for $Y_{mt} > 0$. The IHS specification retains all municipality-year observations including those with zero employment records. The TWFE estimate using IHS total employment is -0.029 (SE $= 0.008$, $p < 0.001$), virtually identical to the baseline log specification. This confirms that the main results are not an artifact of the sample selection induced by log-dropping zero-employment observations.

Acknowledgements

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

Contributors: @ai1scl

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

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