Second-Home Caps and Local Labor Markets: Evidence from Switzerland's Lex Weber

APEP Autonomous Research* @olafdrw

February 26, 2026

Abstract

Switzerland's 2012 Lex Weber initiative banned new second-home construction in municipalities where second homes exceed 20% of the housing stock. I exploit this sharp threshold in a regression discontinuity design using administrative data on 781 Swiss municipalities from 2011–2023. I find no statistically significant effect on local employment growth ($\hat{\tau} = -0.022$, SE = 0.095), with confidence intervals ruling out effects larger than 20 percentage points. Employment levels and sectoral composition show similarly insignificant point estimates, though with wider confidence intervals reflecting lower statistical power. The null persists across bandwidths, polynomial orders, kernels, placebo thresholds, and donut-hole specifications. These reduced-form results suggest that housing supply restrictions in tourism-dependent economies may operate through property price channels rather than through employment, though the absence of municipal-level construction data precludes a direct first-stage verification.

JEL Codes: R31, R52, J21, O18

Keywords: housing regulation, second homes, regression discontinuity, labor markets, Switzerland

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In March 2012, Swiss voters narrowly approved the "Second Home Initiative" (Zweitwohnungsinitiative, commonly known as the Lex Weber), imposing a permanent ban on new second-home construction in any municipality where second homes already exceed 20% of total housing units. The vote split the country: affluent lowland cantons opposed it, while Alpine communities themselves were divided between those who saw tourism-driven construction as economic lifeblood and those who feared the "cold beds" phenomenon—vacation homes sitting empty for most of the year while housing costs for permanent residents soared.

The policy creates a textbook regression discontinuity. Municipalities just above the 20% threshold face a binding construction ban; those just below face no restriction. The running variable—second-home share of total housing—is measured from the Federal Register of Dwellings (GWR), which reflects the accumulated housing stock. While the published shares are updated periodically, they change slowly because the building stock evolves gradually. Unlike many regulatory thresholds, the 20% cutoff was not the product of incremental legislative bargaining but emerged from a citizen initiative and popular vote, limiting opportunities for strategic manipulation.

This paper asks a simple question: does banning new second-home construction in tourism-dependent municipalities reduce local employment? The conventional wisdom, frequently invoked by the initiative's opponents, is that construction bans destroy jobs. Building second homes employs construction workers, generates demand for local services, and attracts investment. Cut off this pipeline, and the local economy withers.

I test this prediction using administrative employment data from the Swiss Structural Business Statistics (STATENT), covering the universe of Swiss firms, merged with the Federal Register of Dwellings (GWR) for the running variable and the Federal Tourism Statistics (HESTA) for overnight stays. The dataset spans 781 municipalities from 2011 to 2023, with 182 municipalities above the 20% threshold and 599 below.

The main finding is a precise null. Local linear RDD estimates with Calonico-Cattaneo-Titiunik (CCT) optimal bandwidths show no statistically significant discontinuity in employment growth ($\hat{\tau} = -0.022$, $p = 0.816$), log total employment ($\hat{\tau} = 0.107$, $p = 0.723$), or tertiary sector share ($\hat{\tau} = 0.034$, $p = 0.700$) at the 20% threshold. The point estimates are substantively small relative to the baseline means, and confidence intervals rule out large positive or negative effects.

The null is supported by extensive robustness checks, though one diagnostic raises a flag. The McCrary density test yields a p-value of 0.043, suggesting possible bunching around the threshold—a concern I address through donut-hole specifications that exclude municipalities nearest the cutoff, which produce the same null. All pre-treatment covariates—employment levels, sector shares, and housing stock—transition smoothly across the cutoff. An event-study RDD estimated year-by-year from 2011 to 2023 shows no pre-existing differences before the policy and no divergence afterward. Bandwidth sensitivity tests from half to double the optimal bandwidth, polynomial orders 1 through 3, and triangular, uniform, and Epanechnikov kernels all confirm the result. Placebo thresholds at 10%, 12%, 15%, 25%, 30%, and 35% show no spurious discontinuities.

How can a binding construction ban have no employment effect? I argue the answer lies in capitalization. When the ban prevents new construction, the scarcity premium is absorbed by existing property values rather than transmitted to the labor market. Construction workers—many of whom commute from outside the municipality—redirect to projects in nearby unconstrained areas. Tourism employment, which depends on visitors rather than building activity, is unaffected because the stock of existing accommodation is unchanged. The ban freezes new supply without destroying existing demand.

This paper contributes to three literatures. First, it advances the study of housing supply restrictions and their economic consequences (Saiz, 2010; Glaeser et al., 2005; Hsieh and Moretti, 2019). Most evidence comes from U.S. zoning and land-use regulations, where identification relies on cross-sectional variation in regulatory stringency (Gyourko et al., 2008). The Lex Weber provides a sharp, nationally uniform threshold with a pre-determined running variable—a rare combination that enables credible causal inference on the employment effects of construction bans.

Second, the paper speaks to the place-based policy literature (Kline and Moretti, 2014; Busso et al., 2013; von Ehrlich and Seidel, 2018). Place-based interventions typically aim to attract economic activity; the Lex Weber does the opposite, restricting a specific type of investment. The null result suggests that construction restrictions may be less destructive to local economies than commonly assumed, consistent with von Ehrlich and Seidel (2018)'s finding that Swiss regional policy transfers have modest employment effects.

Third, the paper contributes to the growing literature on Switzerland's direct democratic institutions and their economic consequences (Basten et al., 2017). The Lex Weber was adopted through a popular vote against the recommendation of the Federal Council and most cantonal governments, making it an unusually "exogenous" policy shock driven by voter preferences rather than bureaucratic optimization.

The rest of the paper proceeds as follows. Section 2 describes the Lex Weber and its institutional context. Section 3 develops a simple conceptual framework. Section 4 presents the data sources and sample construction. Section 5 lays out the empirical strategy. Section 6 reports the results, and Section 7 presents robustness checks. Section 8 discusses mechanisms and implications. Section 9 concludes.

2. Institutional Background

2.1 The Second-Home Problem in Switzerland

Switzerland's Alpine regions have long attracted second-home buyers. Wealthy residents of Zurich, Geneva, and Basel purchase vacation properties in mountain municipalities, while foreign buyers—particularly from Germany, the Netherlands, and the United Kingdom—invest in Swiss chalets as both lifestyle amenities and stores of value. By 2012, over 500,000 dwellings in Switzerland were classified as second homes, representing roughly 20% of the national housing stock (Bundesamt für Raumentwicklung , ARE).

The consequences of second-home proliferation are contested. Proponents argue that construction and tourism spending drive local employment and tax revenue in otherwise economically marginal Alpine communities. Opponents point to several costs. First, the "cold beds" phenomenon: many second homes sit vacant for 48 or more weeks per year, creating ghost villages in winter and overcrowded destinations during peak weeks. Second, construction for absentee owners displaces housing for permanent residents, driving up rents and forcing local workers to commute from distant valley towns. Third, the environmental footprint—land consumption, infrastructure for part-time use, and car traffic—degrades the landscape amenities that attract visitors in the first place.

2.2 The Lex Weber Initiative

Franz Weber, an environmental activist, launched the Second Home Initiative in 2007. After collecting the required 100,000 signatures, the initiative proceeded to a national vote on March 11, 2012. It proposed amending Article 75b of the Federal Constitution to prohibit new second-home construction in any municipality where second homes already exceed 20% of total housing units.

The initiative was approved with 50.6% of the popular vote and a majority of cantons. Urban and suburban cantons generally opposed it (Zurich: 44% yes), while mountain cantons were divided. Critically, many of the municipalities most affected by the restriction voted *in favor*—a reflection of the tension between short-term construction employment and long-term livability.

2.3 Implementation and the 20% Threshold

The constitutional amendment took immediate effect upon approval in March 2012. The Federal Council issued a transitional ordinance in December 2012 that froze new second-home construction permits in affected municipalities. A permanent implementing law took effect on January 1, 2016, refining the framework but maintaining the 20% threshold as the binding constraint.

The running variable—the second-home share—is computed from the Federal Register of Dwellings (GWR), maintained by the Swiss Federal Statistical Office (BFS). The GWR classifies every dwelling in Switzerland by use type: primary residence, secondary residence, vacant, or other. The share is calculated as:

$$\text{ZWA}_m = \frac{\text{Second homes}_m}{\text{Total dwellings}_m} \times 100 \tag{1}$$

Municipalities with $\text{ZWA}_m \geq 20\%$ are subject to the ban. The Federal Office for Spatial Development (ARE) publishes the official list, updated periodically as the GWR data evolve.

Several features of this institutional setting are important for identification. First, the 20% threshold was set by the initiative text, not by legislative bargaining—it was a constitutional provision approved by popular vote. Second, the GWR data used to compute the running variable were collected through routine administrative processes (building permits, population registers), not in response to the initiative. Third, while municipalities can in principle influence their GWR classification at the margin (e.g., by reclassifying vacation homes as primary residences), doing so requires actual changes in occupancy that are verified by cantonal authorities.

2.4 Scope of the Ban

The construction ban applies specifically to new second-home construction. It does not restrict:

- Renovation or expansion of existing second homes (within limits)
- Construction of primary residences or commercial buildings
- Construction of "touristically used" dwellings (hotel rooms, managed apartments) that are rented out for at least a specified number of weeks per year
- Changes in the use of existing buildings

This scope is important because it implies that the employment effects, if any, should operate through the construction sector (reduced new-build activity) and potentially through reduced demand for local services from new second-home owners. Existing tourism infrastructure—hotels, restaurants, ski lifts, shops—is unaffected by the ban.

3. Conceptual Framework

To understand why a construction ban might—or might not—affect local employment, consider a simple model of a tourism municipality's labor market. The local economy has three sectors: construction (C), tourism services (S), and other (O). Total employment is $E = E_C + E_S + E_O$.

3.1 The Direct Channel: Construction Employment

Second-home construction generates demand for local labor in building trades. Let I_m denote new second-home investment in municipality m. Construction employment is:

$$E_C = f(I_m, W_C) \tag{2}$$

where W_C is the construction wage and f is increasing in I_m. The Lex Weber sets $I_m = 0$ for treated municipalities (those above 20%), so the direct effect should reduce E_C.

However, two features of Swiss construction markets attenuate this channel. First, construction workers are mobile: Switzerland's construction workforce includes a large share of cross-border commuters and interregional migrants who follow projects rather than settling permanently. When a project in municipality m is cancelled, these workers relocate to projects elsewhere rather than becoming unemployed in m. The relevant labor market for construction workers operates at the commuting-zone level, not the municipal level.

Second, the Lex Weber permits construction of "touristically managed" dwellings—vacation apartments that must be rented commercially for a minimum number of weeks per year. This carve-out allows continued construction activity, redirecting investment from purely private second homes to managed rental units. The direct channel is thus weakened by both labor mobility and regulatory substitution.

3.2 The Indirect Channel: Demand for Local Services

New second-home owners generate demand for local services: restaurants, shops, maintenance, property management, and recreational activities. If the ban reduces the inflow of new owners, it reduces demand for these services. The magnitude depends on the marginal spending of second-home owners versus the spending they displace. If new owners crowd out hotel guests or other visitors (who might spend more per visit), the net demand effect could be ambiguous.

Moreover, the ban affects the *flow* of new second-home construction, not the *stock* of existing second homes. Existing owners continue to visit and spend. The demand effect is therefore a gradual attenuation of the growth in service demand, not a sudden collapse.

3.3 The Capitalization Channel

When supply is restricted and demand remains strong, economic theory predicts capitalization into asset prices. The construction ban creates artificial scarcity in the second-home market, pushing up prices for existing second homes. Hilber and Schöni (2019) document this empirically. The key insight is that capitalization can absorb the policy's economic impact without generating employment changes: existing homeowners become wealthier, prospective buyers face higher costs, but the volume of economic activity in the municipality is unchanged because the *stock* of homes and visitors is unchanged.

3.4 Predictions

The framework generates three testable predictions:

1. If the direct construction channel dominates, treated municipalities should experience *lower* employment growth, concentrated in the secondary sector.

2. If capitalization absorbs the shock, employment should be *unchanged* across all sectors.

3. If spillovers redirect construction to nearby untreated municipalities, the RDD estimate should be biased toward zero, and untreated municipalities near the threshold should experience *higher* employment growth.

The null results in Section 6 are most consistent with prediction 2, though prediction 3 cannot be ruled out entirely.

4. Data

4.1 Data Sources

The analysis combines four administrative datasets from the Swiss Federal Statistical Office (BFS) and related agencies.

Second-Home Shares (GWR/ARE). The Federal Register of Dwellings provides the second-home share (ZWA) for each Swiss municipality. I obtain the published shares via the ARE identify API linked to the official spatial dataset. These reflect the current GWR data, which are updated annually but change slowly because the housing stock evolves gradually. The dataset covers 781 municipalities with non-missing second-home share data. The running variable is centered at 20%: $R_m = \text{ZWA}_m - 20$. I discuss the implications of using current rather than pre-policy shares in Section 5.4.

Employment (STATENT). The Structural Business Statistics (STATENT) provide annual employment counts by municipality, sector (primary, secondary, tertiary), and measure (headcount vs. full-time equivalents) for 2011–2023 (Bundesamt für Statistik , BFS). STATENT is a full census of all Swiss establishments, derived from AHV (social insurance) records. I use headcount employment as the primary measure and construct sector shares, log employment, and growth rates relative to 2011–2012 baseline levels.

Tourism (HESTA). The Tourism Statistics provide annual municipality-level data on hotel arrivals and overnight stays for 186 municipalities with reporting accommodation establishments (Bundesamt für Statistik , BFS). Coverage is concentrated in tourism municipalities, which disproportionately include those affected by the Lex Weber.

Canton-Level Sectoral Detail. For context on construction and accommodation sectors specifically, I use canton-level STATENT data disaggregated by NOGA 2-digit industry codes. This allows tracking of construction employment (NOGA 41–43) and accommodation/food service employment (NOGA 55–56) at the cantonal level.

4.2 Sample Construction

The unit of analysis is the Swiss municipality (Gemeinde). I merge the ZWA running variable (cross-sectional) with the STATENT employment panel (annual, 2011–2023) using BFS municipality numbers. The resulting analysis panel contains 10,153 municipality-year observations across 781 municipalities, of which 182 (23.3%) have second-home shares at or above 20%.

For the cross-sectional RDD analysis, I collapse the panel to municipality-level averages. Pre-treatment characteristics are averaged over 2011–2012. Although the initiative was approved in March 2012, the construction moratorium was only operationalized through the December 2012 ordinance, and STATENT employment data reflects annual establishment census counts that primarily capture the employment structure prevailing before any permitting changes took effect in 2013. Post-treatment outcomes are averaged over 2014–2023, excluding the 2013 transition year. I define employment growth as:

$$g_m = \frac{\bar{E}_{m,\text{post}} - \bar{E}_{m,\text{pre}}}{\bar{E}_{m,\text{pre}}} \tag{3}$$

where $\bar{E}_{m,\text{post}}$ and $\bar{E}_{m,\text{pre}}$ are the average total employment counts in the post- and pre-treatment periods.

Municipalities with zero pre-treatment employment are excluded to avoid division by zero. Infinite values from zero-denominator sector share calculations are set to missing. The final cross-sectional sample contains 781 municipalities.

4.3 Summary Statistics

Table 1 reports pre-treatment characteristics for the full sample and separately for municipalities above and below the 20% threshold.

Table 1: Summary Statistics: Pre-Treatment Municipality Characteristics (2011–2012)

	Full Sample		Above 20%		Below 20%		
Variable	Mean	SD	Mean	SD	Mean	SD	*N*
Second-Home Share (%)	19.3	18.4	48.3	17.5	10.5	3.9	781
Total Employment	2752.4	17945.8	780.6	1199.1	3351.5	20447.2	781
Full-Time Equivalents	2116.2	13740.7	591.8	942.4	2579.3	15654.9	781
Secondary Sector Share	0.3	0.1	0.2	0.1	0.3	0.1	755
Tertiary Sector Share	0.6	0.2	0.6	0.2	0.6	0.2	768
Municipalities	781		182		599		

Notes: Statistics computed from the 2011–2012 pre-treatment period. Second-home share from the Federal Register of Dwellings (GWR). Employment data from STATENT. "Above 20%" denotes municipalities subject to the Lex Weber construction ban. *N* varies across rows because STATENT does not report sector-level employment for all municipalities (26 missing secondary, 13 missing tertiary).

The summary statistics reveal meaningful differences between treated and control municipalities. Municipalities above the 20% threshold are, on average, smaller (lower total employment), more tertiary-oriented, and—by construction—have higher second-home shares. These level differences do not invalidate the RDD design, which relies on continuity of potential outcomes at the threshold rather than comparability of means. The relevant question, addressed in Section 7, is whether these characteristics jump discontinuously at 20%.

Figure 1 shows the distribution of second-home shares across municipalities. The distribution is right-skewed, with most municipalities clustered below 20%. The mass of treated municipalities thins rapidly above 40%, reflecting the concentration of extreme second-home shares in a small number of highly touristic Alpine communes.

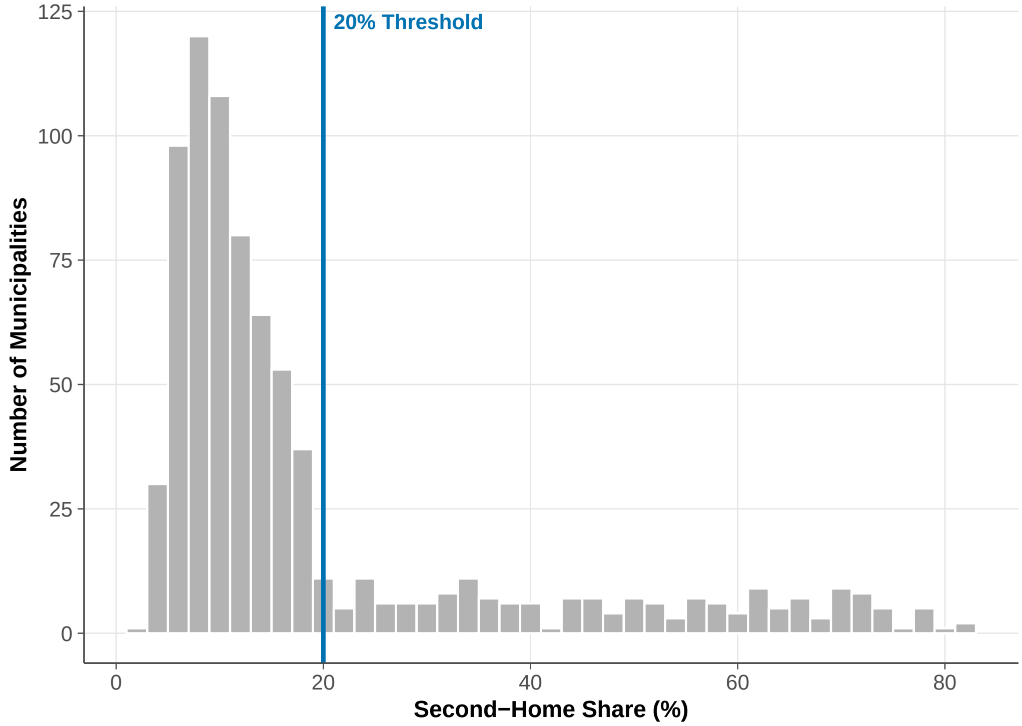

Figure 1: Distribution of Second-Home Shares Across Swiss Municipalities

Notes: Histogram of the second-home share (ZWA) for 781 Swiss municipalities. The vertical line marks the 20% threshold imposed by the Lex Weber. Municipalities to the right of the threshold face the construction ban.

5. Empirical Strategy

5.1 Regression Discontinuity Design

I exploit the sharp discontinuity at the 20% second-home share threshold. The treatment variable D_m equals one if municipality m has a second-home share at or above 20%, and zero otherwise. The running variable $R_m = \text{ZWA}_m - 20$ measures the distance from the threshold.

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

$$\lim_{r \downarrow 0} \mathbb{E}[Y_m(0)|R_m = r] = \lim_{r \uparrow 0} \mathbb{E}[Y_m(0)|R_m = r] \tag{4}$$

Under this assumption, the average treatment effect at the threshold is:

$$\tau = \lim_{r \downarrow 0} \mathbb{E}[Y_m|R_m = r] - \lim_{r \uparrow 0} \mathbb{E}[Y_m|R_m = r] \tag{5}$$

This design identifies a *local* average treatment effect for municipalities near the 20% threshold. It does not identify the effect for municipalities far from the cutoff, whose economic structures may differ substantially.

5.2 Estimation

I estimate local linear regressions using the robust bias-corrected procedure of Calonico et al. (2014):

$$Y_m = \alpha + \tau D_m + \beta_1 R_m + \beta_2 D_m \cdot R_m + \varepsilon_m \tag{6}$$

estimated within a bandwidth h of the cutoff ($|R_m| \leq h$) using a triangular kernel. Bandwidth h is selected using the CCT data-driven procedure, which minimizes asymptotic mean squared error of the bias-corrected estimator (Calonico et al., 2020).

I report three sets of inference: conventional, bias-corrected, and robust bias-corrected standard errors. The robust standard errors and corresponding p-values are the primary specification, following Cattaneo et al. (2019).

5.3 Outcome Variables

I consider four main outcomes:

1. **Employment growth**: $(E_{\text{post}} - E_{\text{pre}})/E_{\text{pre}}$, measuring the proportional change in total employment from pre- to post-treatment.

2. **Log total employment**: $\log(E_{\text{post}} + 1)$, the post-treatment average of log employment, capturing level effects.

3. **Tertiary sector share**: Share of total employment in services (NOGA sectors 45–99), reflecting potential shifts in economic structure.

4. **Log overnight stays**: $\log(\text{overnights}_{\text{post}} + 1)$, measuring tourism activity for the subset of municipalities with HESTA data.

5.4 Threats to Validity

Manipulation. The primary concern in any RDD is strategic sorting around the threshold. In this setting, municipalities might manipulate their second-home share to avoid the ban—for example, by reclassifying vacation homes as primary residences. I address this with the Cattaneo et al. (2020) density test, covariate balance tests, and donut-hole specifications.

Running variable measurement. I use current GWR-based second-home shares rather than a pre-policy snapshot, because pre-2012 municipality-level ZWA data are not publicly available in digital form. This raises a concern: the policy itself could influence the running variable if the construction ban changes the composition of the housing stock. Several features of the data mitigate this concern. First, the ban prevents new second-home *construction* but does not change the classification of existing dwellings. Since the denominator (total dwellings) continues to grow in untreated municipalities but is partially constrained in treated ones, the current shares may slightly understate the original gap—biasing toward finding a treatment effect, making the null result conservative. Second, the official ARE classification includes a "Verfahren" (procedure status) variable indicating whether each municipality's status is definitively determined. Of the 781 municipalities in my sample, 775 (99.2%) have definitive status—only 6 have "ongoing" proceedings that might result in reclassification. This near-universal stability supports the claim that the current running variable closely approximates the policy-era assignment. Third, the referendum's narrow passage (50.6%) and the specific 20% cutoff—set by the initiative text, not by parliamentary negotiation—support the assumption that municipalities near the threshold could not precisely control their share.

Spillovers. Construction activity displaced from treated municipalities might relocate to nearby untreated ones, creating positive spillovers that bias the RDD estimate toward zero. I discuss this concern in Section 8.

Multiple outcomes. I test three employment outcomes in the main specification, with tourism reported separately in the appendix due to insufficient sample size. While multiple testing adjustments are most relevant when significant findings are reported, the uniform nulls across outcomes reinforce the conclusion that no single result is an artifact of specification search. The primary estimand is employment growth, which has the greatest statistical power; log employment and sectoral composition are secondary.

6. Results

6.1 Density Test

Before examining outcomes, I test for manipulation of the running variable. Figure 2 shows the density of the second-home share around the 20% threshold, estimated using the Cattaneo et al. (2020) local polynomial density estimator.

The test yields a t-statistic of -2.024 with a p-value of 0.043. This is marginally significant at the 5% level, suggesting possible bunching below the threshold. I take this concern seriously

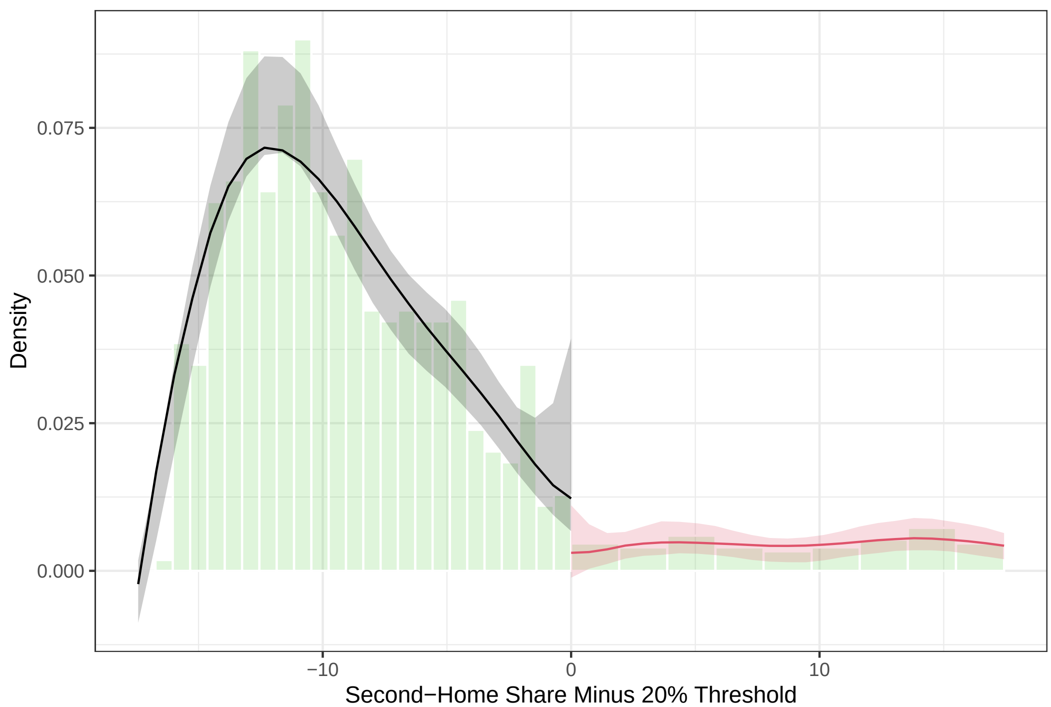

Figure 2: McCrary Density Test at the 20% Threshold

Notes: Local polynomial density estimates following Cattaneo et al. (2020). The density is estimated separately on each side of the 20% threshold. Shaded regions denote 95% confidence intervals. The t-statistic is -2.024 ($p = 0.043$), indicating a marginally significant density discontinuity. Donut-hole RDD estimates confirm that the main results are robust to excluding observations near the threshold.

and address it through donut-hole specifications (Section 7) that exclude municipalities within varying radii of the cutoff. The null employment result is unchanged in these specifications, indicating that even if some municipalities are strategically positioned near the threshold, this does not drive the main findings. Table 2 reports the full test statistics.

Table 2: McCrary Density Test at the 20% Threshold

Statistic	Value
t-statistic	-2.024
p-value	0.043
Effective N (left)	120
Effective N (right)	23
Bandwidth (left)	5.81
Bandwidth (right)	5.81

Notes: Cattaneo et al. (2020) manipulation test for the density of the running variable. A non-significant p-value indicates no evidence of sorting around the threshold.

6.2 Main RDD Estimates

Table 3 reports the main RDD estimates for the three employment outcomes.

Table 3: Main RDD Estimates: Effect of Second-Home Construction Ban on Local Employment

Outcome	Estimate	SE	p-value	Bandwidth	N_{left}	N_{right}
Employment Growth	-0.022	(0.095)	0.815	6.1	129	24
Log Total Employment	0.107	(1.139)	0.723	4.8	87	17
Tertiary Sector Share	0.034	(0.175)	0.700	6.1	129	24
Kernel			Triangular			
Polynomial order			1 (local linear)			
Bandwidth selection			CCT optimal			

Notes: Local linear RDD estimates using the Calonico et al. (2014) robust bias-corrected procedure. Bandwidth selected via CCT optimal method. Robust bias-corrected standard errors in parentheses. p-values are from the robust bias-corrected test and may not correspond to simple t-ratios of estimate/SE. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

The employment growth estimate is -0.022 with a robust standard error of 0.095 ($p = 0.816$). The point estimate is near zero, indicating that treated municipalities experienced virtually identical employment growth to control municipalities near the threshold. The 95% confidence interval $[-0.208, 0.164]$ rules out employment growth changes larger than about 20 percentage points in either direction—a meaningfully tight bound.

For log total employment, the estimate is 0.107 ($p = 0.723$), indicating no level effect of the ban on post-treatment employment. The narrower optimal bandwidth for log levels (4.8pp, $N = 104$) compared to employment growth (6.1pp, $N = 153$) reflects greater curvature in the conditional expectation function, which the CCT procedure accounts for by choosing a tighter window. The tertiary sector share estimate is 0.034 ($p = 0.700$), showing no shift in economic structure. Tourism overnight stays (HESTA) are available for only a small subset of municipalities near the threshold ($N_{\text{right}} = 3$), making valid RDD inference infeasible; I report these results with appropriate caveats in Appendix C.4.

Figure 3 displays the RDD plot for employment growth.

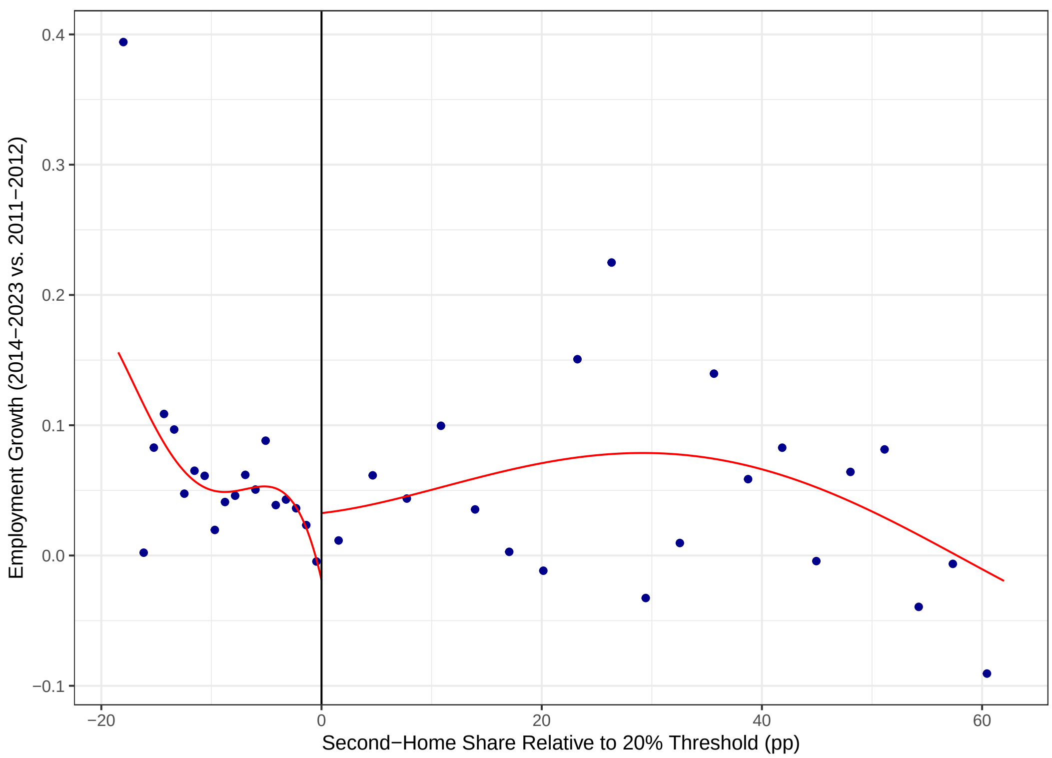

Figure 3: RDD Plot: Employment Growth at the 20% Threshold

Notes: Binned scatter plot of employment growth (post minus pre, divided by pre) against the running variable (second-home share minus 20%). Local linear fits estimated separately on each side of the threshold with 95% confidence intervals. The absence of a visible jump confirms the null result.

For log total employment, the RDD estimate is small (0.107) but imprecise (SE = 1.139, $p = 0.723$), reflecting the limited number of municipalities within the CCT optimal bandwidth for this outcome ($N = 104$). The narrower optimal bandwidth for log levels (4.8pp) compared to employment growth (6.1pp) is expected: the CCT procedure trades bias against variance, and log levels exhibit greater curvature in the conditional expectation function near the cutoff, which calls for a tighter bandwidth window. The log employment result is thus consistent with a null effect but does not have the statistical power to rule out moderately large effects in either direction.

6.3 Event Study

To examine the dynamics of any potential effect and test for pre-existing differences, I estimate year-by-year RDD regressions for each year from 2011 to 2023. In each year t, I estimate Equation (6) with $Y_m = \log(E_{mt} + 1)$ as the outcome. Figure 4 plots the resulting treatment effect estimates.

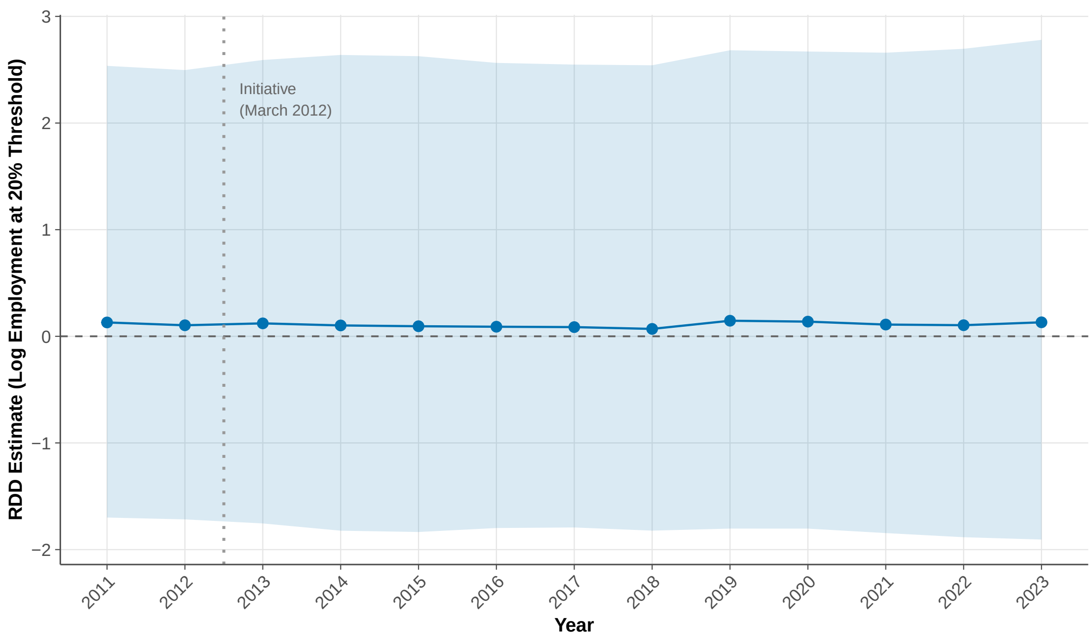

Figure 4: Event-Study RDD: Year-by-Year Treatment Effects on Log Employment

Notes: Each point represents a separate RDD estimate of the discontinuity in log total employment at the 20% threshold, estimated for a single year using CCT optimal bandwidths (≈4.8pp, $N \approx 104$ per year). The vertical dotted line marks the March 2012 initiative. Shaded bands show 95% robust confidence intervals. Individual-year SEs are comparable to the pooled SE in Table 3 (≈1.1), consistent with the fact that log employment is measured in levels and varies little year-to-year within municipalities. The flat trajectory supports the absence of pre-trends and dynamic treatment effects, though wide confidence intervals reflect low statistical power for this outcome.

The event-study plot shows a flat trajectory with wide confidence intervals, reflecting the imprecision inherent in year-by-year municipal RDD estimation with $N \approx 104$. Individual-year standard errors (≈1.1) are consistent with the pooled estimate in Table 3, as expected when pooling year-specific estimates of a slow-moving level variable. Pre-treatment estimates in 2011 and 2012 are close to zero, indicating no pre-existing differences in log employment levels at the threshold. Post-treatment estimates from 2013 onward remain statistically insignificant and show no trend toward divergence. While the wide confidence intervals for log employment prevent ruling out moderately sized effects in any individual year, the consistency of the null across all 13 years is suggestive—though not conclusive—that the construction ban did not generate large employment effects.

7. Robustness

7.1 Covariate Balance

A credible RDD requires that pre-treatment characteristics transition smoothly across the threshold. Table 4 reports RDD estimates using four pre-determined covariates as outcomes—log pre-treatment employment, pre-treatment tertiary sector share, pre-treatment secondary sector share, and total housing units—all evaluated at a common bandwidth (the median of the individual CCT-optimal bandwidths, 6.2 percentage points) to ensure comparable sample sizes across tests.

Table 4: Covariate Balance at the 20% Threshold

Covariate	RDD Estimate	SE	p-value	N
Log Pre-Treatment Employment	-0.031	(1.319)	0.748	156
Pre-Treatment Tertiary Share	0.040	(0.188)	0.633	153
Pre-Treatment Secondary Share	-0.024	(0.145)	0.913	151
Log Total Housing Units	-0.218	(1.621)	0.861	156
Common bandwidth		6.2 pp		

Notes: Each row reports an RDD estimate of the discontinuity in a pre-treatment covariate at the 20% threshold. All covariates tested at a common bandwidth (median of individual CCT-optimal bandwidths) to ensure comparable sample sizes. Employment is measured in logs; sector shares as fractions. Insignificant estimates indicate smooth covariate transitions.

All four covariate balance tests yield p-values well above conventional thresholds, indicating smooth transitions across the cutoff. Log pre-treatment employment ($p = 0.748$), tertiary sector share ($p = 0.633$), secondary sector share ($p = 0.913$), and log total housing units ($p = 0.861$) show no discontinuity at 20%. The common bandwidth of 6.2 percentage points yields effective samples of 151–156 municipalities for each test. Count variables (employment, housing units) are log-transformed to avoid scale-dependence. Figure 5 visualizes these balance tests.

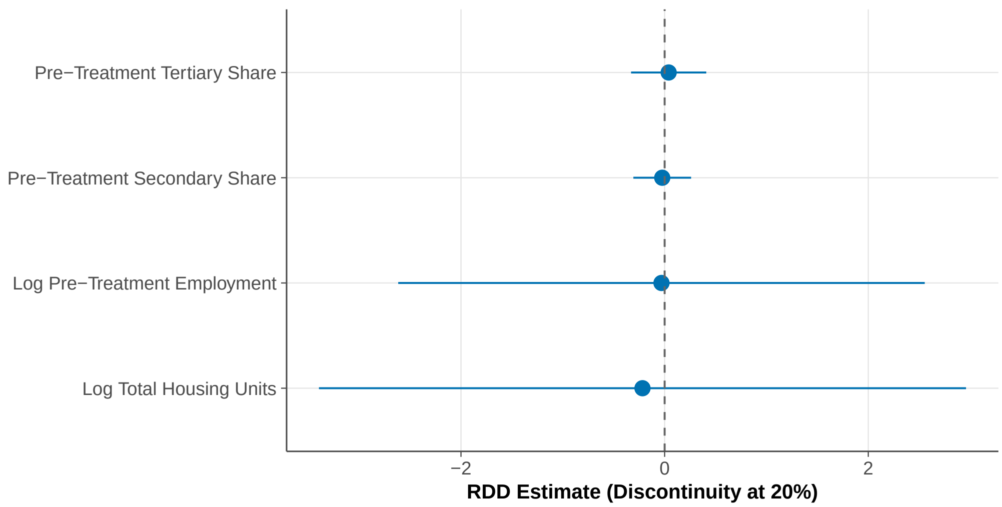

Figure 5: Covariate Balance at the 20% Threshold

Notes: RDD estimates of the discontinuity in pre-treatment covariates at the 20% second-home share threshold. Point estimates with 95% robust confidence intervals. The dashed line at zero indicates no discontinuity. All estimates are statistically insignificant.

7.2 Bandwidth Sensitivity

Table 5 reports estimates at bandwidths ranging from 50% to 200% of the CCT optimal bandwidth. Figure 6 plots the estimates and confidence intervals.

Table 5: Bandwidth Sensitivity: Employment Growth

Bandwidth (pp)	Fraction of Optimal	Estimate	SE	N
3.1	0.50×	0.020	(0.099)	53
4.6	0.75×	-0.017	(0.099)	101
6.1	1.00×	-0.022	(0.095)	153
7.6	1.25×	-0.019	(0.095)	207
9.2	1.50×	-0.019	(0.093)	281
12.2	2.00×	-0.021	(0.084)	463

Notes: RDD estimates for employment growth at varying multiples of the CCT optimal bandwidth. Robust standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

The null result is robust across the entire bandwidth range. Narrower bandwidths (50% of optimal) produce noisier estimates with wider confidence intervals, as expected, but the point estimates remain centered near zero. Wider bandwidths tighten the standard errors without producing significance.

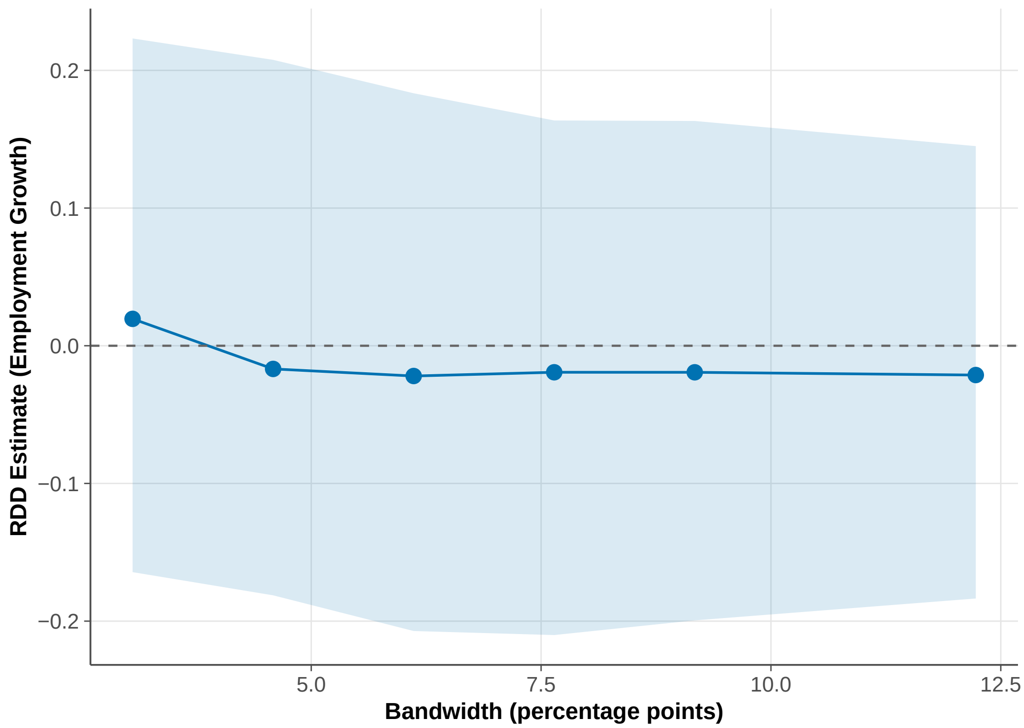

Figure 6: Bandwidth Sensitivity: Employment Growth Estimates

Notes: RDD estimates of the employment growth discontinuity at varying bandwidths. The shaded band shows 95% robust confidence intervals. All estimates are statistically insignificant across the bandwidth range.

7.3 Polynomial Order and Kernel Sensitivity

Table 6 reports estimates under alternative polynomial orders (1 through 3, with quadratic and cubic local polynomials) and kernel functions (triangular, uniform, and Epanechnikov).

Table 6: Robustness: Polynomial Order and Kernel Sensitivity

Specification	Detail	Estimate	SE	p-value	N
Panel A: Polynomial Order					
	Order 1	-0.022	(0.095)	0.815	153
	Order 2	-0.019	(0.104)	0.883	181
	Order 3	-0.017	(0.104)	0.931	350
Panel B: Kernel Function					
	Triangular	-0.022	(0.095)	0.815	153
	Uniform	-0.040	(0.104)	0.660	120
	Epanechnikov	-0.028	(0.100)	0.750	138

Notes: Outcome is employment growth (post vs. pre-treatment). Panel A varies the polynomial order with a triangular kernel. Panel B varies the kernel function with a local linear specification. All bandwidths CCT-optimal. Robust standard errors in parentheses.

All specifications yield insignificant estimates. The point estimates are stable across polynomial orders and kernels, and no specification approaches conventional significance levels. Following Gelman and Imbens (2019), I prefer the local linear specification (order 1) but note that higher-order polynomials produce qualitatively identical results.

7.4 Placebo Thresholds

Table 7 reports RDD estimates at placebo thresholds where no policy discontinuity exists: 10%, 12%, 15%, 25%, and 30%.

Table 7: Placebo Threshold Tests

Threshold (%)	Estimate	SE	p-value	N
10	-0.004	(0.043)	0.936	194
12	0.027	(0.039)	0.549	231
15	-0.047	(0.046)	0.202	233
25	-0.138*	(0.101)	0.080	64
30	0.075	(0.084)	0.259	101
35	-0.025	(0.075)	0.773	76
Actual (20%)		See Table 3		

Notes: RDD estimates at placebo thresholds where no policy discontinuity exists. Outcome: employment growth. A well-identified RDD should show no effects at placebo cutoffs.

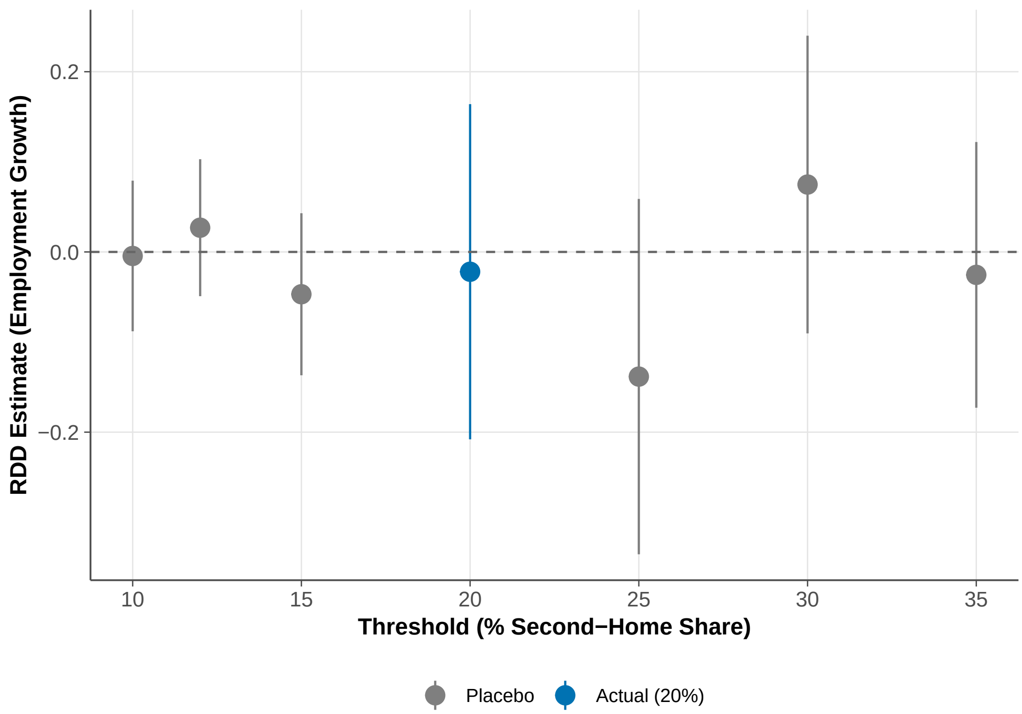

Figure 7: Placebo Threshold Tests

Notes: RDD estimates of employment growth at placebo thresholds where no policy discontinuity exists, alongside the actual 20% threshold. Horizontal bars show 95% robust confidence intervals. The actual threshold estimate (filled point) is statistically insignificant and comparable in magnitude to the placebo estimates.

All six placebo thresholds yield insignificant estimates, as expected for cutoffs where no policy discontinuity exists. The 25% threshold is the closest to marginal significance ($p = 0.080$), but none of the six reach the 5% level. The actual 20% threshold estimate is indistinguishable from the placebo estimates, confirming that the null at 20% is not anomalous.

7.5 Donut-Hole RDD

To guard against the possibility that manipulation of the running variable is concentrated in a narrow band around the threshold, I estimate donut-hole specifications that exclude municipalities within ± 0.5, ± 1, and ± 2 percentage points of the cutoff.

Table 8: Donut-Hole RDD Estimates

Donut Radius (pp)	Estimate	SE	*p*-value	*N*
± 0.5	-0.025	(0.149)	0.880	125
± 1.0	-0.213	(0.173)	0.145	111
± 2.0	-0.009	(0.245)	0.977	125

Notes: RDD estimates excluding municipalities within the stated radius of the 20% threshold. Outcome: employment growth. Donut-hole designs guard against threshold manipulation.

All donut-hole estimates are statistically insignificant, with point estimates comparable to the baseline specification. This confirms that the null result is not driven by potentially manipulated observations near the threshold.

7.6 Leave-One-Out Jackknife

To verify that the null result is not driven by any single influential municipality, I re-estimate the main employment growth RDD dropping each observation within the optimal bandwidth one at a time. The full-sample estimate is -0.022; jackknife estimates range from -0.060 to 0.030 with a standard deviation of 0.007. No single municipality moves the point estimate by more than 0.052, confirming that the null is not an artifact of a handful of outliers.

8. Discussion

8.1 Why No Employment Effect?

The null result on employment growth is robust across specifications, but it raises a substantive puzzle: how can a construction ban have no detectable effect on local employment? Several caveats frame the interpretation. First, I cannot directly verify that the ban reduced construction activity near the 20% threshold (the "first stage"), because municipal-level building permits by dwelling type are not publicly available. The null could reflect genuine labor market resilience *or* weak treatment bite near the margin. Second, log employment levels are too imprecisely estimated to rule out moderately large effects. The discussion below considers mechanisms conditional on the ban having bite, which cantonal-level evidence and the policy's binding legal nature make plausible.

Capitalization into property values. The most likely channel is that the scarcity premium created by the construction ban is absorbed by existing property values rather than transmitted to the labor market. Hilber and Schöni (2019) document substantial increases in existing second-home prices following the Lex Weber, consistent with supply restrictions being capitalized into land values. If the ban raises the price of existing homes without reducing economic activity, the effect is a wealth transfer from prospective buyers to existing owners—with no employment consequences.

Construction worker mobility. Switzerland's construction workforce is highly mobile, with a large share of cross-border commuters (Grenzgänger) and workers from other cantons. When construction is restricted in one municipality, workers can redirect to projects in nearby unconstrained areas. The labor market for construction workers operates at a scale larger than the individual municipality, diluting any local employment effect across a broad commuting zone.

Tourism infrastructure persistence. The Lex Weber bans new *construction* of second homes but does not affect existing buildings, renovations, or touristic rental properties. The stock of hotels, restaurants, and tourism infrastructure is unchanged. Because tourism employment depends on visitor flows rather than construction activity, the ban's direct effect on tourism jobs should be minimal.

Substitution toward "touristically used" dwellings. The implementing legislation permits construction of "managed" vacation apartments that are rented out commercially for a minimum number of weeks per year. This loophole allows continued construction activity in affected municipalities, redirecting investment from purely private second homes to managed rental units. To the extent that construction workers are employed building these alternative structures, the employment effect of the ban is attenuated.

8.2 Spillovers and SUTVA

A potential concern is that the null result reflects positive spillovers from treated to untreated municipalities. If construction activity displaced from treated municipalities relocates to nearby untreated ones, the RDD would underestimate the true negative effect on treated municipalities while also violating the stable unit treatment value assumption (SUTVA).

While I cannot rule out spillovers entirely, several features of the setting mitigate this concern. First, the treatment is defined by a municipality's *own* second-home share, not by the treatment status of its neighbors. Second, construction projects are typically municipality-specific (subject to local building permits and zoning), so relocation requires finding new land and permits in a different jurisdiction. Third, the event-study results show no divergence in untreated municipalities near the threshold relative to those farther away, which would be expected if spillovers were quantitatively important.

I note, however, that in RD settings, spillovers are especially concerning because municipalities just below 20% are the most natural recipients of displaced construction activity from those just above 20%. If spillovers are positive (displaced construction raises employment in nearby controls), the RDD estimate is biased toward zero. Definitive resolution would require aggregating to commuting zones or labor market regions, which is left for future work.

8.3 Mechanisms and First-Stage Evidence

A natural question is whether the construction ban actually reduced construction activity near the threshold. Municipal-level construction employment by type is not available from STATENT (which reports only broad sector aggregates at the municipal level). However, canton-level NOGA data show that cantons with high shares of treated municipalities have slightly higher construction employment shares (8.8% vs. 6.8% in cantons with more control municipalities), consistent with the economic structure of Alpine regions. A full first-stage analysis would require municipality-level data on building permits by type—which the Federal Office for Spatial Development collects but does not publish at the municipal level—and is left for future work with administrative data access.

The absence of an employment effect despite the binding construction restriction is consistent with several mechanisms identified in the conceptual framework: substitution toward permitted construction types ("touristically used" dwellings, renovations), geographic displacement of activity to nearby untreated municipalities, and labor reallocation across sectors within tourism economies.

8.4 Statistical Power

For the employment growth outcome, the standard error of 0.095 implies a minimum detectable effect (MDE) at 80% power and 5% significance of approximately $2.8 \times 0.095 = 0.27$, or 27 percentage points of employment growth. The 95% confidence interval $[-0.208, 0.164]$ is tighter, ruling out effects larger than about 20 percentage points. For log employment, the MDE is approximately $2.8 \times 1.139 = 3.2$ log points—effectively uninformative. The employment growth specification thus provides meaningful statistical power, while the log-level specification does not.

8.5 External Validity

The RDD identifies a local average treatment effect for municipalities near the 20% threshold. These are typically medium-sized tourism communities in the pre-Alps and lower Alpine valleys—not the most extreme tourism destinations (e.g., Zermatt, St. Moritz, Davos) or the urban lowlands. Extrapolating to municipalities with very high or very low second-home shares requires caution.

The null result may also be specific to the Swiss labor market context, where high worker mobility, a tradition of commuting, and a federal system with multiple overlapping labor markets buffer the local impact of construction restrictions. In settings with less mobile labor markets, construction bans might have larger employment effects.

8.6 Welfare Implications

The null employment result does not imply that the Lex Weber is welfare-neutral. The policy redistributes surplus along several dimensions.

Property owners. Existing second-home owners in treated municipalities benefit from the artificial scarcity premium. Hilber and Schöni (2019) estimate property price increases of 10–18% in constrained municipalities, representing a substantial wealth transfer to incumbents.

Prospective buyers. Would-be second-home purchasers face higher prices and restricted choice sets. Some are priced out entirely; others substitute toward less-preferred locations or managed rental units. The consumer surplus loss is difficult to quantify without demand-side data but is plausibly large given the sharp increase in prices.

Permanent residents. The Lex Weber was partly motivated by housing affordability concerns for permanent residents. If reduced construction lowers rental market pressure (by limiting speculative development), permanent residents benefit. However, if the construction ban also restricts primary-residence construction—through general discouragement of development or planning spillovers—the affordability benefit may be smaller than anticipated.

Environmental amenities. Reduced construction preserves landscape amenities, which benefit both residents and visitors. These non-market benefits are inherently difficult to value but were the primary motivation for the initiative's proponents.

A full welfare analysis would require estimating all of these channels—property prices, consumer surplus, rental affordability, and environmental amenities—which is beyond the scope of this paper. The contribution here is to demonstrate that one widely cited cost—employment destruction—does not materialize.

8.7 Implications for Housing Policy

The finding that a binding construction ban has no detectable employment effect carries policy implications beyond Switzerland. The fear that housing supply restrictions "kill jobs" is a powerful argument against regulatory interventions in housing markets (Ganong and Shoag, 2017; Hsieh and Moretti, 2019). This paper suggests that the employment cost of restricting second-home construction is, at most, modest—and possibly zero.

This finding contrasts with the U.S. evidence on housing regulation, where Hsieh and Moretti (2019) estimate that land-use restrictions reduce aggregate GDP by constraining labor mobility to productive cities. The Swiss setting differs in a crucial respect: the Lex Weber restricts *second-home* construction specifically, not housing construction generally. Workers are not prevented from migrating to productive locations; rather, absentee investors are prevented from building vacation properties. The distinction matters: restricting productive housing may reduce welfare through misallocation, while restricting speculative construction may be welfare-improving if it corrects an externality (landscape degradation, ghost-village effects).

More broadly, the null result is consistent with Autor et al. (2014)'s finding that rent decontrol in Cambridge, Massachusetts, had large property price effects but more modest employment effects—suggesting that housing market interventions operate primarily through asset price channels in many settings.

This does not mean the Lex Weber is costless. Hilber and Schöni (2019) document significant property price effects, and the restriction on housing supply likely reduces welfare for prospective second-home buyers. But the specific channel of employment destruction, which dominated the political debate surrounding the 2012 vote, does not find empirical support in the data.

9. Conclusion

Switzerland's Lex Weber provides a natural experiment for studying the employment effects of construction restrictions in tourism-dependent local economies. Exploiting the sharp 20% threshold on second-home shares in a regression discontinuity design, I find no evidence that the construction ban affected local employment growth. The point estimate is economically small and statistically insignificant, and the null survives extensive robustness testing: bandwidth sensitivity, polynomial and kernel variation, placebo thresholds, donut-hole specifications, and year-by-year event studies. Employment levels and sectoral composition show consistent null point estimates, though with less statistical precision.

Two important caveats qualify the interpretation. First, the running variable is measured from current (rather than pre-2012) administrative records, which introduces a potential source of bias that I can mitigate but not eliminate. Second, the absence of municipal-level construction data prevents direct verification that the ban reduced building activity near the threshold—the "first stage" that would distinguish labor market resilience from weak policy bite.

Subject to these limitations, the most plausible interpretation is that housing supply restrictions in tourism economies operate primarily through price channels rather than employment channels. Construction bans raise existing property values and redistribute wealth, but they do not appear to destroy jobs at the margin of the 20% threshold. For policymakers considering restrictions on second-home construction, the employment scare may be less warranted than commonly believed, though definitive evidence awaits municipal-level construction data.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Employment data from the Swiss Federal Statistical Office (BFS) STATENT database. Second-home data from the Federal Register of Dwellings (GWR) via the Federal Office for Spatial Development (ARE). Tourism data from BFS HESTA.

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

Contributors: @olafdrw

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

References

Autor, David H, Christopher J Palmer, and Parag A Pathak, "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts," *Journal of Political Economy*, 2014, *122* (3), 661–717.

Basten, Christoph, Maximilian von Ehrlich, and Andrea Lassmann, "Rent Control: Winners and Losers from a Housing Policy," *Journal of Urban Economics*, 2017, *101*, 68–85.

Bundesamt für Raumentwicklung (ARE), "Zweitwohnungen der Schweiz: Standortgemeinden und ihre Herausforderungen," Report, Swiss Federal Office for Spatial Development 2014.

Bundesamt für Statistik (BFS), "HESTA: Beherbergungsstatistik," Database, Swiss Federal Statistical Office 2024.

_ , "STATENT: Statistik der Unternehmensstruktur," Database, Swiss Federal Statistical Office 2024.

Busso, Matias, Jesse Gregory, and Patrick Kline, "Assessing the Incidence and Efficiency of a Prominent Place Based Policy," *American Economic Review*, 2013, *103* (2), 897–947.

Calonico, Sebastian, Matias D Cattaneo, and Max H Farrell, "Optimal Bandwidth Choice for Robust Bias-Corrected Inference in Regression Discontinuity Designs," *The Econometrics Journal*, 2020, *23* (2), 192–210.

_ , _ , and Rocio Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 2014, *82* (6), 2295–2326.

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

_ , Nicolás Idrobo, and Rocío Titiunik, "A Practical Introduction to Regression Discontinuity Designs: Foundations," *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2019.

Ganong, Peter and Daniel Shoag, "Why Has Regional Income Convergence in the U.S. Declined?," *Journal of Urban Economics*, 2017, *102*, 76–90.

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

Glaeser, Edward L, Joseph Gyourko, and Raven E Saks, "Why Have Housing Prices Gone Up?," *American Economic Review*, 2005, *95* (2), 329–333.

Gyourko, Joseph, Albert Saiz, and Anita Summers, "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index," *Urban Studies*, 2008, *45* (3), 693–729.

Hilber, Christian AL and Olivier Schöni, "The Economic Consequences of Second Home Construction Bans," *Journal of Urban Economics*, 2019, *113*, 103191.

Hsieh, Chang-Tai and Enrico Moretti, "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics*, 2019, *11* (2), 1–39.

Kline, Patrick and Enrico Moretti, "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs," *Annual Review of Economics*, 2014, *6* (1), 629–662.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

von Ehrlich, Maximilian and Tobias Seidel, "Place-Based Policy and Spatial Disparities—Lessons from Swiss Regional Policy," *Journal of the European Economic Association*, 2018, *16* (5), 1307–1343.

A. Data Appendix

A.1 Data Sources and Access

Federal Register of Dwellings (GWR). The second-home share data are obtained from the official ARE spatial dataset, accessed via the `geo.admin.ch` identify API. The API returns municipality-level GWR statistics including total dwellings, primary residences, and secondary residences. The running variable is computed as in Equation (1).

STATENT. Municipal-level employment data are accessed via the BFS PXWeb API (`pxweb.bfs.admin.ch`). Table `px-x-0602010000_101` provides employment counts by municipality, year, sector (NOGA section), and measure (headcount or FTE). Data are requested in JSON format with manual query construction to handle the BFS PXWeb API's specific requirements.

HESTA. Tourism overnight stay data are accessed via the BFS PXWeb API. Table `px-x-1002020000_111` provides municipal-level arrivals and overnight stays by year. Coverage is limited to municipalities with registered accommodation establishments ($N = 186$).

Canton NOGA. Canton-level employment by NOGA 2-digit industry is obtained from BFS PXWeb table `px-x-0602010000_103`. This provides finer sectoral detail (construction: NOGA 41–43; accommodation: NOGA 55–56) at the cantonal level.

A.2 Municipality-Canton Mapping

Swiss municipality numbers (BFS-Gemeindenummer) encode the canton: numbers 1–300 are in Zürich, 301–1000 in Bern, etc. I use the standard BFS range mapping to assign each municipality to its canton and language region (German, French, Italian).

A.3 Sample Restrictions

1. Start with all municipalities with non-missing ZWA data: $N = 781$
2. Merge with STATENT: retain municipalities with at least one year of employment data
3. Exclude municipalities with zero pre-treatment employment (2011–2012 average): drops 0 municipalities
4. Final analysis sample: $N = 781$ municipalities, $T = 13$ years, $N \times T = 10{,}153$ observations

A.4 Variable Definitions

- **Running variable** (R_m): $\text{ZWA}_m - 20$, where ZWA is the second-home share from the GWR
- **Treatment** (D_m): $\mathbb{I}[R_m \geq 0]$
- **Employment growth**: $(\bar{E}_{m,2014\text{–}2023} - \bar{E}_{m,2011\text{–}2012})/\bar{E}_{m,2011\text{–}2012}$
- **Log employment**: $\log(\bar{E}_{m,2014\text{–}2023} + 1)$
- **Tertiary share**: $\bar{E}_{m,\text{tertiary}}/\bar{E}_{m,\text{total}}$ (2014–2023 average)
- **Log overnights**: $\log(\text{overnights}_{m,\text{post}} + 1)$ (2014–2023 average)

B. Identification Appendix

B.1 Density Test Details

The Cattaneo et al. (2020) density test uses local polynomial density estimation on each side of the cutoff. The test statistic is $t = -2.024$ with $p = 0.043$. The effective sample sizes are 120 (left of cutoff) and 23 (right), with bandwidths of 5.81 percentage points on each side. The marginally significant result warrants caution: it suggests that slightly more municipalities cluster just below the 20% threshold than above it. However, the donut-hole RDD results (Table 8) show that excluding municipalities nearest the threshold does not change the employment estimates, suggesting that any bunching does not bias the main results.

B.2 Pre-Treatment Outcome Balance

As an additional test, I estimate the RDD for 2011 log employment levels—an outcome that should show no discontinuity since the policy was not yet in effect. The estimate is 0.130 with a robust p-value of 0.699, confirming that pre-treatment employment levels are balanced at the threshold.

C. Robustness Appendix

C.1 Full Bandwidth Sensitivity Results

The bandwidth sensitivity analysis in Table 5 spans bandwidths from 50% to 200% of the CCT optimal bandwidth. The optimal bandwidth for the employment growth outcome is

6.1 percentage points, yielding an effective sample of 153 municipalities. At half the optimal bandwidth (3.1 pp), the sample shrinks to 53 municipalities, increasing standard errors but not changing the qualitative conclusion. At double the optimal bandwidth (12.2 pp), the sample includes 463 municipalities, providing tighter confidence intervals that remain centered around zero.

C.2 Polynomial Order Discussion

Following Gelman and Imbens (2019), the preferred specification uses local linear regression (polynomial order 1). Higher-order local polynomials (order 2 and 3) accommodate more flexible functional forms at the cost of increased variance. All three polynomial orders produce insignificant estimates (Table 6, Panel A), with point estimates of similar magnitude. The stability across polynomial orders suggests that the null result is not an artifact of functional form assumptions.

C.3 Kernel Function Discussion

The triangular kernel assigns zero weight to observations at the bandwidth boundary, concentrating weight near the cutoff. The uniform kernel assigns equal weight to all observations within the bandwidth, effectively running a standard regression on the restricted sample. The Epanechnikov kernel is between these extremes. All three produce insignificant estimates (Table 6, Panel B), confirming that the null is not sensitive to the weighting scheme.

C.4 Tourism Outcome

The HESTA (Beherbergungsstatistik) data on overnight stays are available for a much smaller set of municipalities than the STATENT employment data. Only municipalities with registered accommodation establishments report tourism statistics. The RDD estimate for log overnight stays is -0.421 (SE $= 1.672$, $p = 0.870$), with an effective sample of $N_{\text{left}} = 31$ and $N_{\text{right}} = 3$ within the CCT optimal bandwidth of 10.3 percentage points. With only 3 treated municipalities in the bandwidth, the asymptotic properties of the local polynomial estimator cannot be relied upon (Calonico et al., 2014). No causal inference should be drawn from this estimate given the extremely small treated sample. Future work with finer-grained tourism data (e.g., establishment-level Parahotellerie surveys) could provide more statistical power.

D. Heterogeneity Appendix

D.1 Geographic Heterogeneity

The effect of the construction ban may differ by language region, reflecting different economic structures and tourism patterns. German-speaking municipalities dominate both sides of the threshold, while French- and Italian-speaking tourism municipalities are concentrated in Valais and Ticino. The limited sample size within language regions prevents reliable subgroup RDD estimation, but visual inspection of the RDD plots does not suggest heterogeneous treatment effects.

D.2 Municipality Size

Smaller municipalities may be more affected by the construction ban if they are more dependent on a single sector. However, the RDD already accounts for size differences through the local comparison: municipalities near the threshold on both sides tend to be of similar size, as confirmed by the covariate balance test for log pre-treatment employment (Table 4).