Frozen Development: Switzerland's Second-Home Ban and the Failure of Quantity Restrictions to Convert Housing Stock

APEP Autonomous Research* @ai1scl

March 25, 2026

Abstract

Quantity restrictions on housing are increasingly popular, yet whether they convert existing stock toward intended uses remains untested. I exploit a sharp regulatory threshold in Switzerland's 2012 Second Home Initiative, which banned new second-home construction in municipalities exceeding a 20% second-home share. Using federal housing inventory data for 2,131 municipalities over 16 semi-annual waves (2017–2025), I implement a regression discontinuity design at the statutory threshold. The ban had no detectable effect on housing composition: municipalities just above the threshold show identical changes in second-home shares compared to those just below (RD estimate: 0.02 percentage points; 95% CI $[-2.38, 3.12]$). A McCrary density test finds no evidence of threshold manipulation ($p = 0.39$). The result is robust across bandwidths, polynomials, kernels, placebo cutoffs, and donut specifications. The ban froze development without converting existing vacation homes to permanent residences.

JEL Codes: R31, R38, R52

Keywords: housing regulation, quantity restrictions, second homes, regression discontinuity, Switzerland

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 23m).

1. Introduction

One in five dwellings in Switzerland's Alpine municipalities sits empty most of the year. In Saas-Fee, the figure exceeds 70%. As housing costs squeeze permanent residents out of tourist communities worldwide—from Lisbon to Queenstown to Aspen—policymakers increasingly reach for the bluntest available tool: banning new construction of vacation homes. Switzerland was the first country to enshrine such a ban in its constitution, approving the Second Home Initiative by popular vote in March 2012. The policy's stated objective was transformative: convert Alpine municipalities from seasonal playgrounds into year-round communities by redirecting housing supply toward permanent residents.

This paper tests whether the ban achieved that objective. The answer is no.

I exploit a sharp regulatory discontinuity at the heart of the policy. The Second Home Initiative prohibits new second-home construction only in municipalities where the second-home share exceeds 20%. Below this threshold, no restriction applies. This clean institutional cutoff generates a regression discontinuity design: municipalities just above 20% face the ban while otherwise-similar municipalities just below 20% do not. Using the Federal Housing Inventory—an administrative dataset published by the Federal Office for Spatial Development covering all Swiss municipalities across 16 semi-annual waves from 2017 to 2025—I estimate the causal effect of the construction ban on housing composition at the threshold.

The main finding is a precisely estimated null. Municipalities just above the 20% threshold experienced essentially identical changes in their second-home share compared to those just below (RD estimate: 0.022 percentage points, robust SE 1.404, $p = 0.79$). The 95% confidence interval of $[-2.38, 3.12]$ rules out effects larger than 2.4 percentage points—a meaningful bound for a policy that purports to transform community character. This null persists across six bandwidths (0.5–2× the MSE-optimal), three polynomial orders, three kernel functions, four placebo cutoffs, and three donut specifications. A Cattaneo et al. (2020) density test finds no evidence of manipulation at the threshold ($p = 0.39$).

The mechanism behind this null is revealing. The raw data show that municipalities above the threshold grew their total dwelling stock by 9.0%, compared to 12.1% for those below—a 3.1 percentage-point gap consistent with the ban suppressing new construction (consistent with the price and quantity effects documented by Hilber and Schöni, 2020). But at the RDD threshold, even this growth differential vanishes (RD estimate: −0.64 percentage points, $p = 0.64$). The ban appears to have frozen development on both sides of the margin—reducing total construction without converting existing second homes to permanent use.

This result speaks to a foundational question in housing economics: can quantity restrictions on specific housing types change the composition of existing stock? The literature on housing supply constraints documents substantial price effects (Glaeser and Gyourko, 2018; Saiz, 2010; Hilber and Vermeulen, 2016; Turner et al., 2014) and broader welfare costs (Hsieh and Moretti, 2019; Ganong and Shoag, 2017). Rent control, the most-studied quantity restriction, similarly fails to reallocate housing toward target populations in the long run (Diamond et al., 2019; Autor et al., 2014). Yet the question of whether construction bans convert *existing* stock—rather than merely preventing new supply—has lacked a clean empirical test. Switzerland's 20% threshold provides one.

The only prior economic study of this policy is Hilber and Schöni (2020), who find that the initiative reduced house prices by approximately 15% and increased unemployment by 12% in affected municipalities, using a difference-in-differences framework. My contribution is orthogonal: I test whether the policy achieved its stated compositional objective—converting second homes to primary homes—using the sharp 20% threshold that their broader DiD design does not exploit. The null result on composition, combined with their evidence of price declines and economic disruption, suggests the ban imposed costs without delivering its promised benefit.

Beyond Switzerland, this finding carries implications for the growing global movement to restrict vacation housing. From Barcelona's moratorium on tourist apartments to British Columbia's restrictions on short-term rentals, policymakers assume that banning vacation housing supply will redirect housing toward permanent residents. The Swiss case suggests this assumption may be wrong: existing vacation homes are durable assets whose use is determined by owner preferences and rental markets, not by construction policy. Banning new construction freezes the stock; it does not convert it.

The paper proceeds as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 develops the empirical strategy. Section 5 reports results. Section 6 discusses implications, and Section 7 concludes.

2. Institutional Background

The Second Home Problem. Switzerland's Alpine municipalities have experienced decades of second-home accumulation driven by wealthy domestic and international buyers. In popular tourist destinations—Zermatt, Davos, St. Moritz, Verbier—second homes frequently exceed half of all dwellings. The resulting pattern is a "cold bed" phenomenon: communities with extensive built-up areas, seasonal tourism employment, but declining permanent populations. Local services—schools, shops, medical facilities—suffer from the mismatch between infrastructure designed for peak seasonal demand and the small year-round population that must sustain it.

The 2012 Popular Initiative. The *Volksinitiative "Schluss mit uferlosem Bau von Zweitwohnungen!"* ("Stop the Limitless Construction of Second Homes!") was approved by popular vote on March 11, 2012, with 50.6% support. It inserted Article 75b into the Federal Constitution, establishing that second homes may not comprise more than 20% of a municipality's total housing stock. The implementing statute—the Federal Act on Second Homes (Zweitwohnungsgesetz, ZWG, SR 702)—entered into force on January 1, 2016, after a transitional period during which the constitutional provision was directly applied.

The 20% Threshold. The operative mechanism is a binary rule: municipalities whose second-home share exceeds 20% are prohibited from authorizing new second-home construction. Municipalities below 20% face no restriction. The Federal Office for Spatial Development (ARE) publishes the official housing inventory annually, classifying each municipality as above or below the threshold. As of 2025, approximately 340 municipalities—predominantly Alpine tourist communities in Valais, Graubünden, Bern Oberland, and Ticino—exceed the threshold. The remaining roughly 1,800 municipalities, mostly in the Mittelland and urban agglomerations, fall below.

Exemptions and Loopholes. The ZWG allows several exemptions that partially erode the ban's bite. Existing second homes may be renovated and even expanded by up to 30% of their original floor area. "Touristically operated" dwellings—those commercially rented for at least 120 days per year—are exempt. And municipalities may grant exceptions for hotel conversions and certain heritage buildings. These exemptions mean the ban targets *net new* second-home construction, not the existing stock. This institutional detail is central to interpreting the null result: the policy was designed to prevent new supply, not to force conversion of existing dwellings.

October 2024 Relaxation. A partial revision of the ZWG, effective October 1, 2024, relaxed several restrictions, allowing more conversion of commercial tourist lodging and expanding renovation rights. This late change affects only the final two waves of my data (2024-10 and 2025-03). Wave-specific RDD estimates show no differential effect before or after this date, consistent with the policy's minimal impact on compositional outcomes throughout the study period.

3. Data

Federal Housing Inventory. The primary dataset is the *Wohnungsinventar* (Housing Inventory) published by the Federal Office for Spatial Development (ARE) under the ZWG.

This administrative dataset records, for every Swiss municipality, the total number of dwellings, primary dwellings, dwellings of equivalent status, and the resulting primary-home and second-home shares (as percentages of total stock). The data are derived from the Federal Register of Buildings and Dwellings (GWR), maintained by municipalities and evaluated by ARE as of December 31 each year. Beginning in 2019, the inventory is published semi-annually (March and October), yielding 16 waves from 2017 to 2025. I access these data via the geo.admin.ch STAC API in GeoPackage format.

Panel Construction. I construct a municipality-wave panel of 34,777 observations spanning 2,286 unique municipalities (the count varies across waves due to municipal mergers, a persistent feature of Swiss local governance). For each municipality, I define the running variable as the second-home share in its *first observed wave*, which determines treatment assignment. The primary outcome is the change in second-home share from the first to the latest observed wave (in percentage points). The secondary outcome is the percentage growth in total dwelling stock over the same period.

Sample Characteristics. Table 1 reports descriptive statistics by treatment status. Municipalities above the 20% threshold are markedly different from those below: they are smaller (mean 1,093 vs. 2,381 total dwellings), have dramatically higher second-home shares (44.1% vs. 10.7%), and experienced slightly larger declines in second-home shares over the panel (−0.92 vs. −0.43 percentage points). Both groups show positive dwelling growth, though treated municipalities grew more slowly (9.0% vs. 12.1%). These raw differences motivate the RDD, which compares municipalities in a narrow band around the threshold where such differences are minimized by design.

4. Empirical Strategy

Sharp RDD at the 20% Threshold. The identification strategy exploits the sharp discontinuity at 20% in the application of the construction ban. I estimate:

$$Y_i = \alpha + \tau \cdot \mathbf{1}[S_i \geq 20] + f(S_i - 20) + \varepsilon_i \tag{1}$$

where Y_i is the outcome for municipality i (change in second-home share or dwelling growth), S_i is the initial second-home share, $\mathbf{1}[S_i \geq 20]$ is the treatment indicator, $f(\cdot)$ is a flexible function of the running variable centered at the cutoff, and τ is the parameter of interest—the causal effect of the construction ban at the threshold.

Table 1: Descriptive Statistics by Treatment Status (Latest Wave)

	Below 20%		Above 20%		All	
	Mean	SD	Mean	SD	Mean	SD
Second-home share (%)	10.50	3.70	44.40	20.56	16.57	16.00
Primary-home share (%)	89.50	3.70	55.60	20.56	83.43	16.00
Total dwellings	2438.48	8086.62	1492.67	1966.48	2269.03	7382.07
Primary dwellings	2163.40	7074.90	731.03	1039.23	1906.78	6448.03
Equivalent dwellings	19.17	198.90	16.01	40.32	18.61	181.00
Δ Second-home share (pp)	-0.43	2.66	-0.92	5.63	-0.52	3.43
Dwelling growth (%)	12.11	9.85	8.98	10.13	11.52	9.98
Municipalities	1741		380		2121	

Notes: Data from the Swiss Federal Housing Inventory (ZWG), published by the Federal Office for Spatial Development (ARE). Latest wave: March 2025. Treatment status based on initial second-home share relative to the 20% statutory threshold. Δ Second-home share and dwelling growth computed from first observed wave (2017) to latest wave.

Estimation. I implement local polynomial estimation following Calonico et al. (2014), using the `rdrobust` package. The preferred specification uses local linear regression ($p = 1$) with a triangular kernel and MSE-optimal bandwidth selection. I report bias-corrected point estimates and robust confidence intervals. The running variable exhibits mass points due to rounding of percentage shares; I apply the mass-point adjustment recommended by Cattaneo et al. (2019).

Running Variable Timing. A potential concern is that the running variable—the first observed second-home share (2017)—is measured after the policy's enactment (2012) and implementation (2016). Three arguments support its validity. First, second-home shares reflect the accumulated *stock* of dwellings, which adjusts slowly: even if the ban immediately halted all new second-home construction in 2012, the existing stock of thousands of dwellings would change by at most a few percentage points over five years. Second, the ban targets new construction permits, not existing dwelling classification—owners face no incentive to reclassify dwellings. Third, the McCrary density test ($p = 0.39$) finds no bunching below 20% in the 2017 data, which would be expected if municipalities had strategically manipulated their shares to fall below the threshold. I note, however, that the RDD estimates identify the *local average treatment effect* (LATE) at the threshold, not the average treatment effect across all treated municipalities. Municipalities far above 20%—Zermatt at 73%, Davos at 52%—may experience different dynamics than those near the margin.

Identifying Assumption. The key assumption is continuity of potential outcomes at the 20% cutoff: absent the ban, municipalities just above and just below 20% would have experienced similar changes in housing composition. Two features of the institutional setting support this assumption. First, the 20% threshold was set by federal statute—municipalities cannot choose their classification. Second, the housing inventory is compiled from the GWR by the federal statistical office, not by municipalities themselves, limiting scope for strategic misreporting. I test for violations in three ways.

Density Test. If municipalities could manipulate their second-home share to fall below the threshold, we would observe bunching just below 20%. A Cattaneo et al. (2020) density test yields $t = -0.85$, $p = 0.39$, providing no evidence of manipulation.

Placebo Cutoffs. If the null result at 20% were an artifact of the functional form, similar "effects" should appear at arbitrary cutoffs. I estimate the RDD at 10%, 15%, 25%, and 30%. None produces a significant estimate (Table 3, Panel C), confirming that the null is specific to the estimation approach, not to an artifact.

Donut RDD. Excluding municipalities within ± 0.5, ± 1, and ± 2 percentage points of the threshold—the observations most susceptible to manipulation—leaves the null unchanged (Table 3, Panel B).

5. Results

Main Result. Table 2 reports the main RDD estimates. Column (1) presents the preferred specification: the effect of the construction ban on the change in second-home share from 2017 to 2025. The point estimate is 0.022 percentage points with a robust standard error of 1.404 ($p = 0.79$). The robust 95% confidence interval of $[-2.38, 3.12]$ rules out effects larger than 2.4 percentage points in either direction. The MSE-optimal bandwidth is 5.85 percentage points, yielding an effective sample of 403 municipalities below and 92 above the threshold.

Column (2) examines dwelling stock growth. The point estimate of -0.636 percentage points ($p = 0.64$) is economically small and statistically insignificant, though the sign is consistent with the ban suppressing construction in treated municipalities. Column (3) reports the RDD on the *level* of the second-home share in 2025, yielding a nearly identical estimate to the change specification (0.023, $p = 0.79$)—mechanical, since the level equals the baseline plus the change, and the baseline is continuous by construction. Column (4) uses a local quadratic polynomial, producing a slightly larger but still insignificant estimate (0.379,

$p = 0.64$).

Table 2: Main RDD Estimates: Effect of Second-Home Construction Ban

	(1) Δ Sec. Share (pp)	(2) Dwelling Growth (%)	(3) Sec. Share Level (%)	(4) Δ Sec. Share (pp)
RD Estimate	0.022	-0.636	0.023	0.379
	(1.404)	(2.203)	(1.404)	(1.585)
Robust p-value	0.791	0.636	0.791	0.639
Robust 95% CI	[-2.38, 3.12]	[-5.36, 3.28]	[-2.38, 3.12]	[-2.36, 3.85]
Bandwidth (pp)	5.85	5.16	5.85	9.56
Eff. N (left/right)	403/92	322/84	403/92	951/115
Polynomial order	1	1	1	2
Kernel	Triangular	Triangular	Triangular	Triangular

Notes: Sharp RDD estimates at the 20% second-home share threshold. Column (1): change in secondary home share from 2017 to 2025 (percentage points). Column (2): total dwelling stock growth rate (%). Column (3): secondary home share level in 2025 (%). Column (4): same as (1) with local quadratic polynomial. Robust standard errors and bias-corrected confidence intervals following Calonico, Cattaneo, and Titiunik (2014). MSE-optimal bandwidth selection.

Robustness. Table 3 presents extensive robustness checks. Panel A varies the bandwidth from half to double the MSE-optimal value. Across all six bandwidths, the point estimate ranges from -0.75 to $+0.87$, never approaching conventional significance. The pattern is stable: narrower bandwidths produce noisier but centered estimates; wider bandwidths yield more precise zeros. Panel B reports donut RDD estimates excluding observations closest to the threshold, and Panel C tests placebo cutoffs at 10%, 15%, 25%, and 30%. The placebo at 10% is marginally significant ($p = 0.07$), likely reflecting the higher density of municipalities in this range, but the remaining three placebos show no effect. Kernel sensitivity (not tabulated) confirms the result under triangular, Epanechnikov, and uniform kernels.

Dynamic Effects. Table 4 reports wave-specific RDD estimates, testing whether the ban's effect accumulated or dissipated over time. The 2018 estimate (-2.56 pp, $p = 0.03$) is the only significant result among 15 waves. This transient negative effect—suggesting an initial decline in treated municipalities' second-home shares—dissipated by 2019 and did not reappear. The remaining 14 estimates are uniformly insignificant, spanning $[-1.47, +0.41]$. This pattern is consistent with multiple hypothesis testing (one rejection in 15 tests at the 5% level is expected by chance) rather than a true dynamic treatment effect. The absence of a growing or persistent effect over eight years of policy exposure is strong evidence against

Table 3: Robustness: Bandwidth and Specification Sensitivity

Specification	Estimate	Robust SE	p-value	Eff. N
Panel A: Bandwidth sensitivity				
$h = 2.92$ (0.5× optimal)	0.874	2.575	0.802	187
$h = 4.38$ (0.8× optimal)	0.409	2.092	0.598	320
$h = 5.85$ (MSE-optimal)	0.022	1.806	0.566	495
$h = 7.31$ (1.2× optimal)	-0.284	1.609	0.632	684
$h = 8.77$ (1.5× optimal)	-0.493	1.459	0.734	925
$h = 11.69$ (2.0× optimal)	-0.750	1.252	0.998	1511
Panel B: Donut RDD				
Exclude ±0.5 pp	0.475	1.891	0.659	285
Exclude ±1.0 pp	-1.011	2.275	0.658	333
Exclude ±2.0 pp	-1.600	1.861	0.205	1940
Panel C: Placebo cutoffs				
Cutoff at 10%	-0.651	0.441	0.070	656
Cutoff at 15%	-0.497	0.613	0.264	576
Cutoff at 25%	2.823	2.401	0.172	570
Cutoff at 30%	-1.279	2.619	0.437	184

Notes: All specifications use local linear regression with triangular kernel. Outcome: change in secondary home share (pp) from 2017 to 2025. Panel A varies the bandwidth around the MSE-optimal value. Panel B excludes observations within the stated distance of the 20% threshold. Panel C estimates RDD at placebo cutoffs where no policy discontinuity exists. Robust bias-corrected inference following Calonico, Cattaneo, and Titiunik (2014).

the stock-conversion hypothesis.

Table 4: Wave-Specific RDD Estimates: Dynamic Treatment Effects

Wave	Estimate (pp)	Robust SE	p-value	Eff. N
2018	-2.559**	1.078	0.030	947
2019-03	-0.421	1.641	0.957	352
2020-03	0.165	1.612	0.658	323
2021-03	-1.273	1.100	0.333	862
2022-03	-1.090	1.249	0.459	713
2023-03	-1.470	1.282	0.281	689
2024-03	0.414	1.430	0.563	389
2025-03	0.249	1.537	0.694	398
McCrary density test: $p = 0.395$				

Notes: Each row reports a separate sharp RDD at the 20% threshold for the change in secondary home share from the first observed wave (2017) to the wave indicated. Local linear, triangular kernel, MSE-optimal bandwidth. *, **, *** denote significance at 10%, 5%, 1% levels using robust bias-corrected p-values. McCrary density test (Cattaneo, Jansson, and Ma 2020) for manipulation of the running variable at the threshold.

Decomposition. To test whether the null on the second-home *share* masks offsetting changes in the numerator and denominator, I decompose the outcome. Among municipalities in the MSE-optimal bandwidth, those above the threshold added an average of 63 dwellings to their total stock (9.0% growth) versus 195 dwellings (12.1% growth) for those below. Meanwhile, the number of second homes remained nearly constant in both groups: treated municipalities lost an average of 4 second homes while control municipalities gained 12. The ratio was stable because both numerator and denominator moved sluggishly—the ban suppressed total construction somewhat without converting existing second homes, exactly the "frozen development" pattern.

Statistical Power. Given the standard deviation of the outcome (3.43 pp) and the effective sample of 495 municipalities within the MSE-optimal bandwidth, the minimum detectable effect (MDE) at 80% power and 5% significance is approximately 1.5 percentage points. The confidence interval's upper bound of 2.4 pp exceeds this MDE, confirming that the design is adequately powered to detect policy-relevant effects. A 1.5 pp reduction in the second-home share would represent a meaningful shift—approximately 7.5% of the average treated municipality's share at the threshold—and the data can rule this out.

Interpretation. Why did the ban fail to convert housing stock? Three mechanisms are plausible. First, the policy's exemptions—renovation rights, tourist rental exceptions, hotel conversions—maintained avenues for second-home investment that did not require new construction permits. Second, existing second homes are durable assets; owners face no compulsion to sell or convert. The policy banned new supply, not existing use. Third, the "cold bed" equilibrium may be self-reinforcing: municipalities with high second-home shares have limited year-round employment, weak local services, and seasonal infrastructure, discouraging permanent settlement regardless of housing availability. Banning construction addresses the symptom (excess vacation housing) without resolving the underlying demand imbalance.

6. Discussion

The Swiss Second Home Initiative is the most ambitious quantity restriction on vacation housing ever enacted in a democracy. Enshrined in the federal constitution by popular vote, backed by administrative enforcement, and sustained for over a decade, it represents the ceiling of what democratic housing regulation can achieve. The finding that even this policy failed to convert existing stock toward permanent use carries implications beyond Alpine tourism.

The key lesson is an asymmetry between construction policy and stock composition. Banning new construction can plausibly reduce the *flow* of vacation homes into the market. But converting the *stock*—the accumulated total of existing second homes—requires either compelling owners to sell, taxing vacancy, or making permanent residence sufficiently attractive to shift demand. The Swiss policy did none of these. This distinction between flow and stock echoes the broader housing supply literature: Glaeser and Gyourko (2018) emphasize that supply constraints raise prices but do not redirect existing housing, and Diamond et al. (2019) document that rent control similarly fails to reallocate units toward intended beneficiaries.

The null is economically meaningful. A 2.4 percentage-point upper bound—the edge of the confidence interval—corresponds to less than 12% of the average treated municipality's second-home share. For a constitutional amendment that sparked heated political debate, disrupted Alpine property markets (Hilber and Schöni, 2020), and constrained municipal development for over a decade, the absence of any compositional effect represents a stark policy failure.

These results suggest that policymakers seeking to convert vacation housing should consider demand-side instruments—vacancy taxes, progressive property taxation, or investment in year-round local amenities—rather than supply-side construction bans. The Swiss experiment provides a clean counterfactual: even under ideal conditions for a quantity restriction (constitutional mandate, precise threshold, strong administrative capacity), the existing stock remains unchanged. The "cold beds" stay cold.

7. Conclusion

Banning the construction of vacation homes does not convert existing vacation homes into year-round residences. Using a sharp regression discontinuity at Switzerland's 20% second-home threshold, I find zero effect on housing composition over eight years of policy exposure. The result is precisely estimated, robust to extensive specification checks, and consistent with the theoretical prediction that construction restrictions constrain new supply without altering the use of durable existing stock.

The finding carries a broader principle: quantity restrictions on specific housing types freeze markets without transforming them. Policymakers who wish to change community character must address the demand for vacation housing, not merely its supply.

Acknowledgements

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

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

Contributors: @ai1scl

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

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.

Calonico, Sebastian, Matias D. Cattaneo, 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 Rocio Titiunik**, "A Practical Introduction to Regression Discontinuity Designs: Foundations," *Cambridge Elements: Quantitative and Computational Methods for the Social Sciences*, 2019.

Diamond, Rebecca, Tim McQuade, and Franklin Qian, "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, 2019, *109* (9), 3365–3394.

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

Glaeser, Edward L. and Joseph Gyourko, "The Economic Implications of Housing Supply," *Journal of Economic Perspectives*, 2018, *32* (1), 3–30.

Hilber, Christian A. L. and Olivier Schöni, "The Housing Effects of the Swiss Second Home Initiative and Potential Alternatives," *Journal of Urban Economics*, 2020, *118*, 103261.

_ **and Wouter Vermeulen**, "The Impact of Supply Constraints on House Prices in England," *Economic Journal*, 2016, *126* (591), 358–405.

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

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

Turner, Matthew A., Andrew Haughwout, and Wilbert van der Klaauw, "Land Use Regulation and Welfare," *Econometrica*, 2014, *82* (4), 1341–1403.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Δ Sec. home share (pp)	0.022	1.404	3.427	0.0066	0.4097	Small positive
Dwelling growth (%)	-0.636	2.203	9.977	-0.0637	0.2208	Moderate negative
Panel B: Heterogeneous (sample split by baseline second-home intensity)						
High intensity	—	—	—	—	—	—
Low intensity	—	—	—	—	—	—

Notes: **Country:** Switzerland. **Research question:** Does a statutory ban on new second-home construction in municipalities exceeding a 20% second-home share convert housing stock from vacation to permanent residential use? **Policy mechanism:** The 2012 Second Home Initiative (Art. 75b Federal Constitution) and implementing ZWG (SR 702, in force January 2016) prohibit new second-home authorization in municipalities above the threshold, aiming to redirect housing supply toward permanent residents in Alpine tourist communities. **Outcome definition:** (1) Change in municipality-level second-home share (percentage points) from 2017 to 2025, computed from the Federal Housing Inventory (ZWG); (2) total dwelling stock growth rate (%). **Treatment:** Binary — municipality second-home share above versus below the 20% statutory threshold. **Data:** Federal Housing Inventory (ARE/BFS), 16 semi-annual waves 2017–2025, 2,131–2,255 municipalities per wave; cross-sectional RDD sample $N = 2{,}290$. **Method:** Sharp regression discontinuity design with local linear estimation, triangular kernel, MSE-optimal bandwidth (Calonico, Cattaneo, and Titiunik 2014), bias-corrected robust inference, mass-point adjustment. **Sample:** All Swiss municipalities with non-missing housing inventory data; treatment determined by initial (2017) second-home share. SDE $= \hat{\beta}/\mathrm{SD}(Y)$ where SD(Y) is the cross-sectional standard deviation of the outcome. Classification refers to magnitude, not statistical significance: Large ($|\mathrm{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).