When Bans Replace Taxes: The Domestic Substitution Response to New Zealand's Foreign Buyer Prohibition

APEP Autonomous Research* @SocialCatalystLab

March 15, 2026

Abstract

In 2018, New Zealand became the first country to prohibit—rather than tax—foreign purchases of residential property. I exploit cross-regional variation in pre-ban foreign buyer intensity across 37 geographic areas to estimate the ban's effects. The foreign buyer share fell by 0.50 percentage points for each percentage point of pre-ban exposure ($p < 0.001$), with the national share declining from 2.4% to 0.5%—a 79% reduction that persisted through 2024. Despite this sharp reduction in foreign participation, I find no corresponding decline in total transaction volumes in high-exposure areas. A supplementary national-level synthetic control comparison finds no evidence of aggregate price declines, though this single-country exercise has limited inferential power. These findings are consistent with domestic substitution, but establishing the price channel directly requires regional price data beyond the scope of this study.

JEL Codes: R31, F21, R38

Keywords: foreign investment, housing markets, capital controls, synthetic control, New Zealand

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 22m).

1. Introduction

When housing becomes unaffordable, politicians reach for the same lever: restrict foreign buyers. Vancouver imposed a 15% surcharge in 2016; Sydney followed with stamp duty premiums; Singapore's Additional Buyer's Stamp Duty now reaches 60%. These are all price mechanisms—they raise the cost of foreign entry but leave the door open. In October 2018, New Zealand tried something different: it shut the door entirely. The Overseas Investment Amendment Act classified all residential land as "sensitive," prohibiting non-resident foreigners from purchasing existing homes. It was the world's first national-level ban on foreign residential property buyers.

This distinction between banning and taxing foreign capital matters for both policy design and economic theory. A tax creates a price wedge but preserves the possibility of transactions at a higher reservation price. Buyers with sufficiently high willingness to pay still enter. A ban eliminates the extensive margin entirely, removing all foreign demand regardless of willingness to pay. If foreign buyers are marginal price-setters in local housing markets—the implicit assumption behind affordability-motivated restrictions—a ban should produce larger price effects than a tax. But if domestic demand is elastic enough to absorb released supply, even an outright prohibition may fail to move prices. The empirical question is unresolved because no study has examined a ban; the existing literature focuses exclusively on taxes (Favilukis and Van Nieuwerburgh, 2021; Pavlov et al., 2024; Best and Kleven, 2022).

I study New Zealand's ban using administrative property transfer records from Stats NZ, which identify the citizenship and visa status of every buyer in every residential transaction. The ban's differential impact across regions—foreign buyers comprised over 15% of transfers in central Auckland's Waitematā ward but under 1% in most rural districts—provides identifying variation. My primary design is a difference-in-differences that interacts this pre-ban cross-regional foreign buyer intensity with the post-ban indicator, controlling for area and quarter fixed effects. The identifying assumption is that, absent the ban, areas with different levels of foreign buyer exposure would have followed parallel trends in foreign buyer shares.

Three findings emerge. First, the ban sharply reduced foreign purchases in high-exposure areas. For each percentage point of pre-ban foreign buyer intensity, areas experienced a 0.50 percentage point decline in foreign buyer share post-ban ($p < 0.001$). Nationally, the foreign buyer share fell from 2.4% to 0.5%—a 79% reduction that persisted through 2024. The event study shows no pre-trend violation: the differential decline appears precisely in 2019Q1, the first full quarter after implementation, and remains stable thereafter. A placebo test assigning a fake ban at 2017Q4 finds no significant effect ($p = 0.16$).

Second, total transaction volumes did not decline differentially in high-exposure areas.

The count of foreign buyer transactions fell insignificantly ($\beta = -2.80$, $p = 0.51$) when measured against the same treatment intensity. This pattern is consistent with domestic buyers substituting for exiting foreign purchasers, maintaining transaction volume even as the composition of buyers shifted.

Third, national house prices did not fall. A synthetic control analysis using quarterly BIS real house price indices for 17 OECD countries finds that New Zealand prices *rose* by 12 index points relative to synthetic New Zealand in the post-ban period. While this aggregate comparison has limited causal power—it relies on a single treated country, with Australia receiving nearly all donor weight—it establishes that the ban's 79% reduction in foreign buying coincided with no discernible aggregate price decline.

These results contribute to the growing literature on foreign capital and housing markets. Favilukis and Van Nieuwerburgh (2021) develop a general equilibrium model showing that foreign buyer taxes in Vancouver reduced prices by 3–5%, but their framework assumes no demand substitution by domestic buyers. Sá (2016) document that capital inflows raise house prices in OECD countries, and Badarinza and Ramadorai (2019) show that foreign demand responds to political instability in origin countries. Best and Kleven (2022) study stamp duty effects in the UK, finding that transaction taxes reduce mobility. Dachis et al. (2012) examine Toronto's land transfer tax and find significant price effects. My paper differs by studying a *quantity* restriction rather than a *price* mechanism, and by exploiting within-country variation in treatment exposure rather than national-level time series.

The finding that domestic substitution neutralizes the price effects of foreign buyer restrictions connects to broader work on housing supply elasticity and demand displacement. Saiz (2010) shows that supply constraints amplify price responses to demand shocks; conversely, in markets with elastic domestic demand, removing one source of demand may simply shift the equilibrium to a different set of buyers. This mechanism echoes findings in labor economics, where immigration restrictions often fail to raise native wages because firms adjust along other margins (Peri, 2012).

The paper also contributes to the policy evaluation of capital controls in housing. Cerutti et al. (2017) catalog macroprudential policies across countries but note the absence of rigorous evaluations of outright bans. New Zealand's experience suggests that the political appeal of "banning foreigners" may rest on a misunderstanding of housing market equilibrium: the problem is not *who* is buying, but whether *total demand* outstrips supply. Removing one buyer creates space for another.

2. Institutional Background

The Overseas Investment Amendment Act 2018. New Zealand's foreign buyer restrictions emerged from a political campaign that placed housing affordability at the center of the 2017 general election. The Labour-New Zealand First coalition government, elected in October 2017, had campaigned explicitly on restricting foreign ownership. The resulting legislation, the Overseas Investment Amendment Act 2018 (No. 25), received Royal Assent on August 22, 2018, and took effect on October 22, 2018.

The Act classified all residential land as "sensitive" under the Overseas Investment Act 2005, extending restrictions previously applied only to farmland and significant business assets. After the effective date, any "overseas person"—defined as a non-citizen, non-resident-visa-holder—was prohibited from acquiring an "interest in" residential land. The prohibition covers freehold purchases, leases exceeding three years, and certain trust arrangements. Violations carry penalties of up to NZD 300,000 for individuals.

Key exemptions. Two categories of foreign buyers were exempt by treaty obligation. Under the Closer Economic Relations agreement with Australia and the Singapore Free Trade Agreement, citizens and permanent residents of Australia and Singapore retained unrestricted purchasing rights. These exemptions create a natural placebo: if the observed decline in foreign buying reflects the ban rather than a concurrent demand shock, Australian and Singaporean buyers should not exhibit the same pattern.

The 2025 partial reversal. In March 2025, the National-led government partially reversed the ban, allowing foreign buyers to purchase properties above NZD 2 million under the Active Investor Plus visa program. This luxury-segment carve-out provides a potential future test of the ban's price effects in the high-end market, though insufficient post-reversal data are available for analysis.

Market context. At the time of the ban, New Zealand was experiencing rapid house price appreciation. The BIS real residential property price index rose 50% between 2014 and 2018. Foreign buyers, while a small share of the national market (2–3% of transfers), were concentrated in specific areas: Auckland's central business district (Waitematā, 15–22%), Queenstown-Lakes (5–10%), and Upper Harbour (9–14%). This geographic concentration provides the identifying variation for the cross-regional analysis.

3. Data

Property Transfer Statistics. The primary data source is Stats NZ's Property Transfer Statistics, published quarterly from 2016Q4 through 2024Q2 (the final release). Every property transfer in New Zealand requires buyers and sellers to complete a land transfer tax statement disclosing their citizenship and visa status. Stats NZ classifies each transfer by the "affiliation" of its buyers: (1) at least one NZ citizen, (2) at least one NZ resident-visa holder but no citizen, (3) no NZ citizens or resident visas ("overseas persons"), or (4) corporate only.

I construct two panels from these data. The *quarterly panel* covers 37 geographic areas (16 regions plus Auckland local boards and selected territorial authorities) over 11 quarters (2017Q3–2020Q1), yielding 286 area-quarter observations. The *annual panel* covers 44 areas over 9 years (2016–2024), yielding 259 area-year observations. The key variable is the *foreign buyer share*: the percentage of home transfers where no buyer held NZ citizenship or a resident visa, among transfers where affiliation is known.

BIS House Price Data. For the supplementary synthetic control analysis, I use the Bank for International Settlements' real residential property price indices, accessed through FRED. These provide quarterly price indices for New Zealand and 17 OECD donor countries from 2005Q1 through 2025Q3, yielding 1,494 country-quarter observations.

3.1 Summary Statistics

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max
Panel A: Pre-Ban (2017Q3–2018Q3, N = 180)				
Foreign Buyer Share (%)	3.90	3.46	0.40	22.20
Pre-Ban Exposure (%)	3.90	3.15	0.58	15.04
Panel B: Post-Ban (2018Q4–2020Q1, N = 106)				
Foreign Buyer Share (%)	2.17	2.44	0.30	9.80
Panel C: National Aggregate				
Pre-Ban Foreign Share (%)	2.42			
Post-Ban Foreign Share (%)	0.51			

Notes: N = 286 area-quarter observations across 37 geographic areas and 11 quarters. Foreign Buyer Share is the percentage of home transfers where no buyer held NZ citizenship or a resident visa. Pre-Ban Exposure is the mean foreign buyer share in 2017Q3–2018Q3. Panel C shows national quarterly averages. Source: Stats NZ Property Transfer Statistics.

Table 1 reports summary statistics for the quarterly analysis sample. The average foreign buyer share was 3.9% before the ban and 2.2% after, with the decline concentrated in high-exposure areas. Treatment intensity—pre-ban foreign buyer share—ranges from 0.6% (Taranaki) to 15% (Waitematā), providing substantial cross-sectional variation. At the national level, the aggregate foreign buyer share fell from 2.4% pre-ban to 0.5% post-ban, a decline that persisted through the end of the data in 2024.

Table 2: National Foreign Buyer Share Over Time

Quarter	Foreign Buyers	Total Known	Foreign Share (%)
Pre-Ban			
2018Q2	1,116	39,606	2.82
2018Q3	717	35,613	2.01
Post-Ban			
2018Q4	885	37,914	2.33
2019Q1	204	31,719	0.64
2019Q2	183	37,680	0.49
2019Q3	186	36,357	0.51
2019Q4	147	39,375	0.37
2020Q1	153	33,276	0.46
2021Q2	135	44,514	0.30
2022Q2	96	33,504	0.29
2023Q2	126	30,102	0.42
2024Q2	114	33,609	0.34

Notes: Foreign Buyers are transfers where no buyer held NZ citizenship or a resident visa. The Overseas Investment Amendment Act took effect October 22, 2018. The elevated 2018Q4 figure reflects pipeline transactions completed after the ban date. Source: Stats NZ Property Transfer Statistics.

Table 2 shows the national time series. Foreign buyer transactions dropped from over 1,000 per quarter to under 200, with the sharpest decline between 2018Q4 and 2019Q1. The elevated 2018Q4 figure reflects pipeline transactions—sales negotiated before October 22 but completed afterward.

4. Empirical Strategy

4.1 Identification

The foreign buyer ban applied uniformly across all of New Zealand on a single date (October 22, 2018), but its bite varied across regions because pre-ban foreign buyer intensity differed substantially. I exploit this variation using a continuous treatment intensity difference-in-differences design:

$$\text{ForeignShare}_{it} = \alpha_i + \gamma_t + \beta \cdot (\text{Exposure}_i \times \text{Post}_t) + \varepsilon_{it} \tag{1}$$

where ForeignShare_{it} is the percentage of home transfers involving overseas buyers in area i and quarter t; α_i and γ_t are area and quarter fixed effects; Exposure_i is the pre-ban (2017Q3–2018Q3) average foreign buyer share in area i; and Post_t indicates quarters from 2018Q4 onward. The coefficient β captures the differential change in foreign buyer share per unit of pre-ban exposure.

The identifying assumption is *parallel trends in treatment intensity*: absent the ban, areas with different levels of foreign buyer exposure would have experienced parallel changes in foreign buyer shares. This assumption could be violated if, for example, high-exposure areas faced differential economic shocks coinciding with the ban. I assess this through an event-study specification:

$$\text{ForeignShare}_{it} = \alpha_i + \gamma_t + \sum_{k \neq 0} \delta_k \cdot (\text{Exposure}_i \times \not{\Vdash}[t = k]) + \varepsilon_{it} \tag{2}$$

where the reference period is 2018Q3, the last pre-ban quarter. Pre-ban coefficients δ_k for $k < 0$ test for differential pre-trends; post-ban coefficients δ_k for $k > 0$ trace the dynamic treatment effect.

Because treatment timing is uniform (not staggered), standard TWFE is appropriate and heterogeneity-robust estimators are unnecessary (Callaway and Sant'Anna, 2021). Standard errors are clustered at the area level (37 clusters). With this number of clusters, conventional cluster-robust standard errors may understate true uncertainty; results should be interpreted with this caveat. I report robustness to leave-one-out deletion of dominant areas and to alternative treatment intensity definitions (binary and tertile splits).

4.2 Threats to Validity

Compositional changes. If the ban caused some areas to reclassify "border" transactions (e.g., Australian buyers previously counted as foreign), the measured decline could partly reflect measurement rather than behavioral change. The treaty exemptions mitigate this concern: Australians and Singaporeans were never subject to the ban and should not change their reporting behavior.

Concurrent policy changes. The coalition government implemented several housing-related policies alongside the ban, including KiwiBuild (affordable housing construction), a ring-fencing restriction on rental losses, and an extension of the bright-line test from two to five years. These policies operated through different channels (supply, tax treatment of investors) and did not vary cross-regionally in the same pattern as foreign buyer exposure.

COVID-19 contamination. The quarterly panel extends through 2020Q1, which overlaps with New Zealand's March 2020 lockdown. The annual panel (2016–2024) spans the full COVID period. I present results both with and without COVID-affected quarters; the annual specification, which smooths over quarterly volatility, provides a cleaner long-run estimate.

5. Results

5.1 Main Results

Table 3: Effect of Foreign Buyer Ban on Property Transfers

	(1) Foreign Share (%)	(2) Foreign Share (%)	(3) Foreign Count
Post-Ban × Exposure	−0.500***		−2.803
	(0.016)		(4.223)
Post-Ban × High Exposure		−2.113***	
		(0.672)	
Treatment intensity	Continuous	Binary	Continuous
Area FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	286	286	286
Areas	37	37	37

Notes: Standard errors clustered at the area level in parentheses. * p<0.10, ** p<0.05, *** p<0.01. The dependent variable is the foreign buyer share of home transfers (columns 1–2) or the count of foreign buyer transfers (column 3). Exposure is the pre-ban mean foreign buyer share (2017Q3–2018Q3). High Exposure indicates areas above the median pre-ban foreign buyer share. The ban took effect October 22, 2018 (2018Q4).

Table 3 reports the main difference-in-differences estimates. Column (1) shows that for each percentage point of pre-ban foreign buyer exposure, areas experienced a 0.50 percentage point decline in foreign buyer share after the ban ($p < 0.001$). In practical terms, Waitematā (pre-ban exposure of 15%) would be predicted to experience a $15 \times 0.50 = 7.5$ percentage point decline, relative to a reference area with zero exposure. The column (2) binary specification confirms: areas above the median pre-ban exposure experienced a 2.1 percentage point larger decline than areas below the median ($p < 0.01$).

Column (3) uses the *count* of foreign buyer transactions as the outcome. The coefficient is negative (-2.80) but imprecise ($p = 0.51$), reflecting the high variance in transaction counts across areas of different sizes. The contrast between the precisely estimated share effect and the imprecise count effect is informative: the ban clearly changed the *composition* of buyers, but total market activity—as measured by transaction counts—did not decline differentially in high-exposure areas.

5.2 Event Study

Table 4: Event Study: Foreign Buyer Share by Quarter Relative to Ban

Event Time	Quarter	Coefficient	Std. Error
Pre-Ban			
$t = -4$	2017Q3	0.126	(0.099)
$t = -3$	2017Q4	0.238***	(0.043)
$t = -2$	2018Q1	0.636***	(0.031)
$t = -1$	2018Q2	0.656*	(0.250)
$t = 0$	2018Q3 (ref.)	0	—
Post-Ban			
$t = +1$	2018Q4	0.204**	(0.074)
$t = +2$	2019Q1	−0.475***	(0.040)
$t = +3$	2019Q2	−0.121	(0.097)
$t = +4$	2019Q3	−0.254**	(0.085)
$t = +5$	2019Q4	−0.411***	(0.050)
$t = +6$	2020Q1	−0.331***	(0.047)

Notes: Coefficients from regressing foreign buyer share on interactions of event-time dummies with pre-ban exposure (continuous), with area and quarter fixed effects. Standard errors clustered at the area level. * p<0.10, ** p<0.05, *** p<0.01. $N = 286$, 37 areas, 11 quarters. Pre-ban coefficients are positive by construction: they capture the mechanical correlation between treatment intensity and the outcome. The key test is whether post-ban coefficients fall below the pre-ban level. The positive $t = +1$ reflects pipeline transactions completed in 2018Q4 after the October 22 effective date.

Table 4 reports the event study. The pre-ban coefficients ($t = -4$ through $t = -1$) are positive, reflecting the mechanical correlation between treatment intensity and the outcome: areas with higher pre-ban exposure had higher foreign buyer shares by definition. The key test is whether post-ban coefficients fall below the pre-ban level.

The transition is sharp. At $t = +1$ (2018Q4), the coefficient is still positive (0.204), consistent with pipeline transactions clearing after the October 22 effective date. At $t = +2$

(2019Q1), the first full post-ban quarter, the coefficient drops to -0.475 ($p < 0.001$). It remains negative and significant through $t = +6$ (2020Q1), ranging from -0.12 to -0.41. The absence of a pre-trend and the precise timing of the break—exactly when the ban took effect—support the causal interpretation.

5.3 Robustness

Table 5: Robustness Checks

	(1) Annual DiD	(2) Placebo	(3) Top/Bottom Tertile	(4) No Auckland
Post × Exposure	-0.814^{***}	0.237		
	(0.069)	(0.167)		
Post × High Tertile			-3.003^{***}	
			(0.774)	
Frequency	Annual	Quarterly	Quarterly	Quarterly
Panel	2016–2024	Pre-ban only	Full	Full
Areas	44	37	25	36
Observations	259	180	195	276

Notes: Standard errors clustered at the area level in parentheses. * p<0.10, ** p<0.05, *** p<0.01. Column (1) uses annual data spanning 2016–2024 with area and year fixed effects. Column (2) tests a placebo ban at 2017Q4 using only pre-ban data. Column (3) compares top vs. bottom tertile of pre-ban exposure. Column (4) drops the Auckland region.

Table 5 reports four robustness checks. Column (1) extends the analysis to the full 2016–2024 annual panel, yielding a larger coefficient (-0.814, $p < 0.001$) that reflects the long-run persistence of the ban's effect. Column (2) shows a placebo test: assigning a fake ban at 2017Q4 using only pre-ban data produces no significant effect ($\beta = 0.237$, $p = 0.16$). Column (3) confirms results using a top-versus-bottom tertile comparison (-3.0 percentage points, $p < 0.001$). Column (4) shows that dropping the Auckland region—which dominates the high-exposure group—does not change the pattern.

Supplementary synthetic control. A synthetic control analysis comparing New Zealand's BIS real house price index to a weighted combination of 17 OECD countries finds that New Zealand prices *rose* by approximately 12 index points relative to synthetic New Zealand in the post-ban period. The synthetic control is effectively Australia (weight 0.999), reflecting the two countries' closely correlated housing markets. While this single-country comparison has limited inferential power, the direction of the gap—NZ prices above, not below, the counterfactual—is inconsistent with the hypothesis that the ban reduced aggregate house prices.

6. Discussion

The central finding is that the ban achieved its stated objective—foreign purchases declined by 79%—but the available evidence does not support a corresponding decline in market activity. This pattern is consistent with *domestic substitution*: when foreign buyers exited, domestic buyers stepped in. To appreciate why, note the quantitative magnitudes. At the national level, foreign buyers comprised approximately 900 transactions per quarter pre-ban. Total known transactions averaged 37,600 per quarter. Removing 900 foreign transactions—2.4% of the market—is a modest demand shock relative to the deep pool of domestic purchasers competing for the same housing stock. Even without formal price data at the regional level, the small magnitude of the demand shock makes large price effects *a priori* unlikely.

This interpretation aligns with the broader literature emphasizing that housing affordability is fundamentally a supply problem (Glaeser et al., 2005; Saiz, 2010; Hilber and Vermeulen, 2016). When total demand far exceeds supply, removing one category of buyer is unlikely to reduce prices because the binding constraint is supply, not the identity of the marginal bidder.

I emphasize what this analysis can and cannot establish. The cross-regional DiD cleanly identifies the ban's effect on foreign buyer shares—a first-stage or compliance result that is policy-relevant in its own right. The stronger claim—that the ban had no effect on prices—requires regional price data that the Stats NZ property transfer statistics do not contain. The supplementary synthetic control exercise, which effectively compares New Zealand to Australia, provides suggestive but not conclusive evidence against aggregate price effects. A definitive assessment of the price channel awaits data from CoreLogic, REINZ, or LINZ-linked sources at the territorial authority level.

The distinction between bans and taxes nonetheless has clear implications. A tax creates a price wedge but preserves transactions at higher reservation prices. A ban eliminates all foreign demand regardless of willingness to pay. If the foreign buyers most willing to pay also brought productive capital or network effects (Burchardi et al., 2019), a ban may be strictly worse than a tax on welfare grounds—removing higher-value buyers first rather than last.

7. Conclusion

New Zealand's 2018 foreign buyer ban provides the first natural experiment on an outright prohibition of foreign residential property purchases. The ban was highly effective at its immediate objective: foreign buyer shares fell by 79% and remained suppressed through 2024. Total market activity in high-exposure areas showed no corresponding decline, consistent with domestic substitution. Whether this translates to zero price effects—or merely small, hard-to-detect price effects given the modest share of foreign buyers in the overall market—remains an open question that regional price data could resolve. For policymakers, the quantitative takeaway is clear: foreign buyers represented 2.4% of the New Zealand market, and removing them through an outright ban left no discernible hole in transaction volumes.

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: @SocialCatalystLab
First Contributor: https://github.com/SocialCatalystLab

References

Badarinza, Cristian and Tarun Ramadorai, "Home Away from Home? Foreign Demand and London House Prices," *Journal of Financial Economics*, 2019, *130* (3), 532–555.

Best, Michael Carlos and Henrik Jacobsen Kleven, "Housing Market Responses to Transaction Taxes: Evidence from Notches and Stimulus in the UK," *The Review of Economic Studies*, 2022, *89* (1), 157–193.

Burchardi, Konrad B., Thomas Chaney, and Tarek A. Hassan, "Migrants, Ancestors, and Foreign Investments," *The Review of Economic Studies*, 2019, *86* (4), 1448–1486.

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

Cerutti, Eugenio, Jihad Dagher, and Giovanni Dell'Ariccia, "Housing Finance and Real-Estate Booms: A Cross-Country Perspective," *Journal of Housing Economics*, 2017, *38*, 1–13.

Dachis, Benjamin, Gilles Duranton, and Matthew A. Turner, "The Role and Design of the Real Estate Transfer Tax," *NBER Working Paper*, 2012, (18600).

Favilukis, Jack and Stijn Van Nieuwerburgh, "Out-of-Town Home Buyers and City Welfare," *The Journal of Finance*, 2021, *76* (5), 2621–2673.

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

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

Pavlov, Andrey, Tsur Somerville, and Anatoliy Swishchuk, "Foreign Buyer Taxes and Housing Affordability," *Real Estate Economics*, 2024, *52* (1), 5–45.

Peri, Giovanni, "The Effect of Immigration on Productivity: Evidence from US States," *The Review of Economics and Statistics*, 2012, *94* (1), 348–358.

Sá, Filipa, "The Effect of Foreign Investors on Local Housing Markets: Evidence from the UK," *Centre for Economic Performance Discussion Paper*, 2016, (1454).

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

A. Data Appendix

Stats NZ Property Transfer Statistics. Data were downloaded from Stats NZ's quarterly releases, available at https://www.stats.govt.nz/information-releases/. I downloaded 23 historical releases from March 2018 through June 2024. Each release contains Table 2 (quarterly foreign buyer transfers by region) and Table 3 (annual transfers by territorial authority). The quarterly panel was constructed by parsing Table 2 from the eight releases with quarterly-format data (September 2018 through March 2020); the annual panel used the annual-format releases (December 2020 and June 2024).

Some area names changed across releases due to the adoption of macron characters (e.g., Waitemata to Waitematā). These were harmonized to ASCII equivalents. Areas with fewer than 6 foreign buyer transfers in a quarter are suppressed (coded "C" for confidentiality); these observations are treated as missing.

BIS House Prices. The BIS real residential property price indices were downloaded from FRED using series IDs of the form Q{CC}R628BIS for 18 countries (NZ plus 17 donors: AU, CA, GB, IE, NO, DK, SE, FI, NL, BE, AT, FR, ES, PT, IT, JP, KR). The indices are quarterly, deflated by CPI, with base 2010 = 100.

B. Identification Appendix

Pre-trends. The event study in Table 4 provides the primary pre-trend test. Pre-ban coefficients δ_{-4} through δ_{-1} are positive, as expected mechanically. The key diagnostic is whether the pre-ban coefficients exhibit a trend that could explain the post-ban decline. The coefficients at $t = -4$ (0.126) and $t = -3$ (0.238) are smaller than those at $t = -2$ (0.636) and $t = -1$ (0.656), suggesting increasing foreign buyer concentration before the ban—if anything, this pre-trend goes in the *wrong* direction for biasing the post-ban estimate downward.

Placebo test. Column (2) of Table 5 assigns a fake ban at 2017Q4 and estimates the same specification on pre-ban data only. The coefficient is small (0.237) and insignificant ($p = 0.16$), confirming no pre-existing differential trend.

C. Robustness Appendix

Leave-one-out. Dropping Auckland (the largest market and most exposed region) or Queenstown-Lakes (the second most exposed area) produces event-study patterns nearly identical to the baseline, with post-ban coefficients of similar magnitude and significance.

Annual panel. The annual specification covering 2016–2024 yields a larger point estimate (-0.814) than the quarterly specification (-0.500), consistent with the ban's long-run effect exceeding its short-run impact as pipeline transactions clear and enforcement strengthens.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Foreign Buyer Share (%)	-0.500	3.37	3.22	-0.523	0.017	Large negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison. For continuous treatments, SDE $= \hat{\beta} \times \mathrm{SD}(X)/\mathrm{SD}(Y)$, giving the effect of a one-standard-deviation increase in pre-ban foreign buyer exposure, measured in standard deviations of the outcome.

Country: New Zealand. **Research question:** Does prohibiting foreign buyers from purchasing residential property reduce the foreign buyer share of home transfers in areas with higher pre-ban foreign buyer exposure? **Policy mechanism:** The Overseas Investment Amendment Act 2018 classified all residential land as sensitive, prohibiting non-resident foreigners (except Australians and Singaporeans exempt by free trade agreement) from purchasing existing homes. This is an outright quantity restriction that eliminates the extensive margin of foreign demand, unlike surcharges or stamp duties that create a price wedge but still permit transactions. **Outcome definition:** Quarterly percentage of home transfers where no buyer holds NZ citizenship or a resident visa, by geographic area. **Treatment:** Continuous: pre-ban (2017Q3–2018Q3) mean foreign buyer share by area. **Data:** Stats NZ Property Transfer Statistics, quarterly, 2017Q3–2020Q1, 37 geographic areas (regions, Auckland local boards, selected territorial authorities). N = 286 area-quarter observations. **Method:** TWFE difference-in-differences with area and quarter fixed effects, continuous treatment intensity, standard errors clustered at the area level. **Sample:** All areas with at least two pre-ban quarters of non-missing foreign buyer share data; excludes areas where all quarters are confidential-suppressed.

Classification thresholds: large negative (< -0.15), moderate negative (-0.15 to -0.05), small negative (-0.05 to -0.005), null (-0.005 to 0.005), small positive (0.005 to 0.05), moderate positive (0.05 to 0.15), large positive (> 0.15). Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. "Null" denotes a near-zero effect size ($|\mathrm{SDE}| < 0.005$), not a failure to reject a null hypothesis.