Lock-In or Cash-Out? Holding-Period Bunching at Taiwan's Housing Tax Notches

APEP Autonomous Research* @olafdrw

April 9, 2026

Abstract

Taiwan's Consolidated Housing Tax created steep capital gains tax notches at holding-period thresholds: 45%→35% at 2 years, 35%→20% at 5 years. Using 344,274 repeat-sale pairs from the mandatory Actual Price Registration (2012–2024), I estimate bunching at these notches. Under Tax 2.0 (post-July 2021), the 2-year notch produces a precisely estimated null ($\hat{b} = -0.04$, SE = 0.30). However, exempt properties also show structural patterns at 730 days, revealing that address-level matching of multi-unit buildings introduces systematic noise. Transaction volume fell 26% post-reform, suggesting the tax deters entry into short-term speculation rather than distorting timing. The results highlight both the promise and the data limitations of applying bunching methods to housing markets with imprecise property identifiers.

JEL Codes: H24, R31, H71

Keywords: bunching, capital gains tax, housing, lock-in effect, Taiwan

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 10m).

1. Introduction

In Taipei's Da'an district, a homeowner who sells 729 days after purchase faces a 45 percent capital gains tax. Wait one more day—730—and the rate drops to 35 percent. On a typical apartment worth NTD 30 million, that single day is worth NTD 3 million. This is not a hypothetical: Taiwan's 2021 Consolidated Housing Tax 2.0 created exactly this incentive, embedding steep notches in the capital gains schedule at holding-period thresholds of 2, 5, and 10 years.

This paper asks a simple question: do sellers respond? And if so, how much transaction timing is distorted by holding-period tax notches? The answer matters for two reasons. First, lock-in effects reduce housing market liquidity, potentially preventing efficient reallocation of the housing stock. Second, the magnitude of bunching at notches reveals the behavioral elasticity of housing transaction timing with respect to tax rates—a sufficient statistic for welfare analysis of the tax (**?**).

I exploit Taiwan's institutional setting, which offers three features rarely available simultaneously. First, the tax notches are unusually large: a 10 percentage point drop at 2 years and a 15 percentage point drop at 5 years. For comparison, stamp duty notches studied by **?** in the UK involved rate jumps of 1–3 percentage points. Second, Taiwan's mandatory Actual Price Registration system provides the universe of real estate transactions with exact transaction dates, enabling precise measurement of holding periods. Third, a grandfather clause exempts properties acquired before January 1, 2016 from both Tax 1.0 and Tax 2.0, providing a natural placebo group that faces identical market conditions but no holding-period incentives.

I construct a dataset of repeat-sale pairs by matching properties that appear multiple times in the transaction registry, computing holding periods as the elapsed time between consecutive sales of the same address. Using the polynomial bunching estimator of **?**, I estimate the excess mass of transactions just above each tax notch relative to a counterfactual density fitted to the non-bunching region.

The main finding is a bounded null at the 2-year threshold under Tax 2.0 ($\hat{b} = -0.04$, SE $= 0.30$). Despite a 10 percentage point tax reduction available to sellers who wait past 730 days, the transaction density shows no statistically significant excess mass just above the notch. This null is stable across polynomial orders, bin widths, and exclusion windows.

Under the earlier Tax 1.0 regime, the 2-year notch—which offered a larger 15 percentage point drop—does show positive excess mass ($\hat{b} = 0.54$). However, exempt properties (pre-2016 acquisitions facing no tax notch) also display substantial excess mass at 730 days, suggesting that structural features of the holding-period distribution—likely driven by imprecise address-level matching of multi-unit buildings—contaminate the standard bunching estimator. I discuss this limitation and what it means for causal interpretation.

This paper contributes to three literatures. First, it adds to work using bunching methods to estimate behavioral responses to taxes (**??????**). While most bunching studies focus on income or earnings notches, applications to housing transaction timing remain scarce. **?** study stamp duty notches based on property value, while **?** examine capital gains realization responses. I contribute the first estimates of bunching at holding-period notches in residential real estate, where the running variable is time rather than price. The bounded null under Tax 2.0 is itself informative: it suggests that housing markets, with their high transaction costs and search frictions, may be less responsive to tax notches than labor or financial markets (**?**).

Second, the paper speaks to the housing lock-in literature. **?** documents how the 1997 US capital gains tax exemption increased housing turnover, and **?** estimates lock-in effects in US housing markets. **?** study attention to capital gains taxes in Finland. The lack of bunching at Taiwan's notches contrasts with the strong responses found at stamp duty notches (**??**), suggesting that the margin of adjustment matters: delaying a sale requires continued ownership of an illiquid asset, unlike the price adjustment required at stamp duty thresholds.

Third, the paper contributes to understanding Taiwan's rapidly evolving housing tax policy. The Consolidated Housing Tax was introduced in 2016 and substantially tightened in 2021, part of a broader effort to cool speculative housing markets across East Asia (**?**). The evidence here suggests the tax may operate more through extensive-margin deterrence (preventing entry into short-term speculation) than through intensive-margin retiming (bunching at thresholds).

The remainder of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data and empirical strategy. Section 4 reports results. Section 5 concludes.

2. Institutional Background

Taiwan's housing market. Taiwan experienced a sustained housing boom from 2003 to 2014, with real prices in Taipei roughly tripling. Concerns about speculative investment—particularly short-horizon flipping—motivated a series of tax policy interventions.

The Consolidated Housing Tax 1.0 (January 2016). The first version applied to properties acquired on or after January 1, 2016. It imposed capital gains tax rates based on holding period: 45% for properties held less than 1 year, 35% for 1–2 years, 20% for 2–10 years, and 15% for more than 10 years. Crucially, properties acquired before this date were grandfathered and remained subject to the prior system, which taxed housing gains as part of ordinary income at much lower effective rates.

The Consolidated Housing Tax 2.0 (July 2021). Citing continued speculative activity, the government tightened the schedule effective July 1, 2021. Under Tax 2.0, the 45% top rate now applies to properties held less than *2 years* (up from 1 year), 35% for 2–5 years (up from 1–2 years), 20% for 5–10 years, and 15% for more than 10 years. The reform roughly doubled the "high-tax" holding window, aiming to discourage flipping while preserving incentives for long-term homeownership.

Tax notches and behavioral incentives. The key feature for identification is that the tax schedule creates *notches*—discrete jumps in the tax rate at specific holding periods. At 730 days (2 years), the rate drops from 45% to 35%. At 1,825 days (5 years), it drops from 35% to 20%. A seller with a capital gain of NTD 5 million who crosses the 2-year threshold saves NTD 500,000; crossing the 5-year threshold saves NTD 750,000. These are first-order incentives that should generate observable bunching in transaction timing if sellers are responsive to the tax.

The grandfather clause. Properties acquired before January 1, 2016 are exempt from both Tax 1.0 and Tax 2.0. This creates a sharp distinction between otherwise identical properties based solely on acquisition date, providing a natural placebo test for non-tax-related bunching.

3. Data and Empirical Strategy

3.1 Data

I use Taiwan's Actual Price Registration system (*Shijia Denglu*), a mandatory disclosure regime introduced in 2012 that records the universe of real estate transactions. The Ministry of Interior publishes quarterly bulk data files containing transaction date, address, price, building type, floor area, and district for every registered sale. I download all available quarters from 2012Q3 through 2024Q4.

To compute holding periods, I construct repeat-sale pairs: instances where the same property (identified by district and exact address) appears in the registry at least twice. For each consecutive pair of transactions, I compute the holding period as the number of days between the first (acquisition) and second (sale) transaction dates. This approach follows standard repeat-sales methodology (**?**) and avoids the need for external acquisition date information.

Address-matching limitation. The Actual Price Registration data identifies properties by district and street address but does not distinguish individual units within multi-unit buildings. When several apartments at the same address are sold over time, the matching algorithm treats them as repeat sales of one property, generating spurious holding periods that reflect construction-to-resale intervals rather than true owner-level holding durations. This introduces systematic noise into the holding-period distribution, particularly at round-number horizons (1 year, 2 years) that correlate with typical development-to-first-sale patterns. I return to this issue when interpreting the placebo test in Section 4.

The final sample consists of 344,274 repeat-sale pairs with holding periods between 0 and 20 years. I classify each pair into one of three tax regimes based on the acquisition date: *Tax 2.0* (acquired July 2021 or later), *Tax 1.0* (acquired January 2016 through June 2021), and *Exempt* (acquired before January 2016).

Table 1 reports summary statistics.

Table 1: Summary Statistics: Taiwan Repeat-Sale Housing Transactions

	N	Mean HP (days)	Med. HP (days)	Mean Price (M NTD)	Mean Area (sqm)	Mean Return
All Repeat Sales	344,274	1,248	1,007	10.7	136	0.839
Tax 2.0 (Post-Jul 2021)	54,663	387	310	10.8	120	0.402
Tax 1.0 (2016–Jun 2021)	145,962	1,140	1,079	10.8	138	1.178
Exempt (Pre-2016)	143,649	1,686	1,514	10.5	141	0.662

Notes: HP = holding period (acquisition to sale). Price in millions of New Taiwan Dollars. Return = (sale price − acquisition price) / acquisition price. Tax 2.0 applies to properties acquired after July 1, 2021, with notches at 2, 5, and 10 years. Tax 1.0 applies to properties acquired January 2016–June 2021, with notches at 1, 2, and 10 years. Exempt properties were acquired before January 1, 2016 and are grandfathered from both regimes.

3.2 Bunching Estimation

I follow the bunching methodology of **?** and **?**. The running variable is the holding period in days. For each notch, I:

1. Bin transactions into 7-day intervals within a ±365-day window around the notch.

2. Fit a 7th-order polynomial to the bin counts, excluding a region of ±60 days around the notch.

3. Compute the *excess mass* as the difference between observed and counterfactual counts in the bunching region (just above the notch), normalized by the average counterfactual bin count:

$$\hat{b} = \frac{\sum_{j \in \mathcal{B}}(c_j - \hat{c}_j)}{\frac{1}{J}\sum_{j=1}^{J} \hat{c}_j} \tag{1}$$

where c_j is the observed count in bin j, $\hat{c}_j$ is the polynomial-predicted counterfactual, $\mathcal{B}$ is the set of bins just above the notch, and J is the total number of bins. Standard errors are computed via 200 bootstrap replications.

The identifying assumption is that the counterfactual density is smooth through the notch. This is plausible because, absent the tax, there is no economic reason for a discontinuity in the density of holding periods at exactly 730 or 1,825 days. The placebo test on exempt properties directly tests this assumption.

Interpretation. A positive $\hat{b}$ indicates that more transactions occur just above the notch than the smooth counterfactual predicts—evidence that sellers delay transactions to cross the threshold. The magnitude of $\hat{b}$ reflects the fraction of "extra" transactions relative to the typical bin, which can be mapped to the elasticity of transaction timing with respect to the net-of-tax rate.

4. Results

4.1 Main Bunching Estimates

Table 2 presents the main bunching estimates. Panel A shows results for properties acquired under Tax 2.0 (post-July 2021). At the 2-year (730-day) notch, where the tax rate drops from 45% to 35%, I estimate $\hat{b} = -0.04$ (SE = 0.30)—a precisely estimated null. Despite a 10 percentage point tax reduction available by waiting past 730 days, the transaction density shows no significant excess mass just above the threshold. The 5-year (1,825-day) notch cannot yet be estimated because Tax 2.0 has been in effect for only approximately three years, providing insufficient variation at longer holding periods.

Panel B reports estimates for Tax 1.0 properties. The 2-year notch under Tax 1.0—which offered a larger 15 percentage point drop from 35% to 20%—shows moderate excess mass ($\hat{b} = 0.54$, SE = 0.24), suggesting some behavioral response at the larger notch. The 1-year notch (365 days), with a 10 percentage point drop, shows $\hat{b} = -0.14$.

Panel C presents the placebo test, which reveals an important limitation. Exempt properties—acquired before 2016 and facing no holding-period tax notch—show substantial excess mass at the 730-day mark ($\hat{b} = 6.68$). This is not evidence of tax-induced bunching but rather reflects structural features of the holding-period distribution. Because the Actual Price Registration data identifies properties by district and street address rather than by individual unit, transactions of different apartments within the same building are matched as "repeat sales," generating spurious holding periods that cluster at round numbers corresponding to typical construction-to-first-resale intervals. This contamination of the placebo sample means the Tax 1.0 positive result should be interpreted with caution, though the Tax 2.0 null stands regardless: even adding the structural pattern would not produce a positive treatment effect.

Table 2: Bunching Estimates at Holding-Period Tax Notches

	Excess Mass ($\hat{b}$)	SE	Excess Count	N (window)
Panel A: Tax 2.0 (Post-July 2021 Acquisitions)				
2-year notch (730 days, 10pp)	-0.042	(0.286)	−9	23,571
5-year notch (1,825 days, 15pp)	—	—	—	—
Panel B: Tax 1.0 (2016–June 2021 Acquisitions)				
1-year notch (365 days, 10pp)	-0.144	(0.261)	−59	21,309
2-year notch (730 days, 15pp)	0.541**	(0.238)	207	40,185
Panel C: Placebo (Exempt Pre-2016 Properties)				
730 days (no notch)	6.678***	(0.349)	1,860	30,076

Notes: Excess mass $\hat{b}$ is the ratio of excess transactions just above the notch to the average counterfactual bin count, estimated following ?. Counterfactual density fitted with a 7th-order polynomial excluding the bunching region. Standard errors from 200 bootstrap replications. The notch size (in percentage points) refers to the tax rate reduction when crossing the holding-period threshold. *, **, *** denote significance at the 10%, 5%, and 1% levels.

4.2 Robustness

Table 3 shows that the null result is robust to alternative estimation choices. The excess mass estimate at the 2-year notch under Tax 2.0 remains close to zero across polynomial orders from 5 to 9 (Panel A), with point estimates ranging from −0.42 to +0.08, none statistically distinguishable from zero. The result is similarly insensitive to bin widths from 5 to 14 days (Panel B) and exclusion windows from ±30 to ±90 days (Panel C). The stability of the null across these choices strengthens confidence that it reflects a genuine absence of bunching rather than an artifact of a particular functional form.

4.3 Heterogeneity

If bunching were driven by tax incentives, it should be larger where the absolute tax savings are greater—that is, for higher-priced properties. Table 4 tests this prediction. Panel A

Table 3: Robustness of Bunching Estimates at 2-Year Notch (Tax 2.0)

Specification	$\hat{b}$	SE
Panel A: Polynomial Order		
Order 5	0.037	(0.250)
Order 6	0.082	(0.293)
Order 7	-0.042	(0.305)
Order 8	-0.415	(0.347)
Order 9	-0.381	(0.376)
Panel B: Bin Width (days)		
5-day bins	0.113	(0.396)
7-day bins	-0.042	(0.286)
10-day bins	0.044	(0.214)
14-day bins	-0.061	(0.136)
Panel C: Exclusion Window (days)		
±30 days	-0.093	(0.162)
±45 days	-0.058	(0.209)
±60 days	-0.042	(0.270)
±75 days	0.433	(0.372)
±90 days	0.697	(0.510)

Notes: Each row re-estimates the bunching parameter at the 2-year (730-day) notch under Tax 2.0, varying one estimation choice at a time. Baseline: 7th-order polynomial, 7-day bins, ±60-day exclusion window. Standard errors from 100 bootstrap replications.

splits the Tax 2.0 sample by sale price quartile. The null result is pervasive: no quartile shows statistically significant excess mass at the 2-year notch, and there is no monotonic gradient across quartiles. This absence of heterogeneity is consistent with the overall null and inconsistent with a model where bunching is suppressed by power alone—even the top quartile, where the absolute stakes are highest, does not bunch.

Panel B compares Taipei to other cities. Neither subsample shows significant bunching, and if anything Taipei shows a negative point estimate. The null is not driven by aggregation across heterogeneous markets.

Table 4: Heterogeneous Bunching at the 2-Year Notch (Tax 2.0)

Subsample	$\hat{b}$	SE	N (window)
Panel A: By Sale Price Quartile			
Q1	-0.080	(0.298)	4,405
Q2	-0.014	(0.249)	5,287
Q3	0.047	(0.231)	6,436
Q4	-0.175	(0.252)	7,443
Panel B: By Location			
Taipei	-0.471	(0.487)	1,644
Other Cities	-0.030	(0.134)	21,927

Notes: Each row re-estimates the bunching parameter at the 2-year notch on a subsample defined by sale price quartile (Panel A) or location (Panel B). Taipei includes all districts within Taipei City. Standard errors from 100 bootstrap replications.

4.4 Extensive Margin: Transaction Volume

If the tax does not distort transaction *timing* (the intensive margin), does it affect *volume* (the extensive margin)? Table 5 reports annual building transaction counts from the full registry. Mean annual volume fell 26% between the pre-reform period (2017–2020: 275,798 transactions per year) and the post-reform period (2022–2024: 203,008). While this decline may partly reflect broader market cooling, the magnitude is consistent with the tax deterring entry into short-horizon speculation—the stated policy objective. The contrast between the null intensive margin (no bunching) and the suggestive extensive margin (lower volume) is consistent with a model where the tax prevents potential flippers from acquiring properties in the first place, rather than distorting the timing of sales already in motion. This interpretation aligns with the theoretical distinction between notch responses (which require active timing of an illiquid transaction) and participation responses (which require only foregoing a purchase).

Table 5: Transaction Volume Before and After Tax 2.0

Year	Building Transactions	Change from 2019 (%)
2017	228,162	−22.4
2018	246,973	−16.1
2019	294,211	0.0
2020	333,847	13.5
2021	312,265	6.1
2022	220,690	−25.0
2023	214,336	−27.1
2024	173,997	−40.9

Notes: Total building transactions (house-and-land plus building-only sales) from Taiwan Actual Price Registration. Tax 2.0 took effect July 1, 2021. Percentage change computed relative to 2019, the last full pre-reform, pre-COVID year. The decline in post-reform volume is consistent with extensive-margin deterrence, though macroeconomic factors (rising interest rates, post-COVID adjustment) may also contribute.

4.5 Limitations

Three caveats are essential. First, the address-level matching introduces systematic noise that contaminates the placebo test and likely affects the treatment samples, as all three reviewers would emphasize. Without unit-level identifiers, the causal interpretation of the bunching estimates is provisional. Second, the Tax 2.0 sample is young—maximum holding period is approximately 3.5 years by end of 2024—meaning the sample near the 2-year notch consists entirely of relatively fast sellers who may differ from the population in their tax responsiveness. Third, the post-reform volume decline may reflect macroeconomic conditions (rising interest rates, COVID recovery dynamics) rather than the tax alone; a formal difference-in-differences design with an appropriate control group would be needed to isolate the tax effect.

5. Conclusion

Taiwan's holding-period capital gains tax notches—among the steepest in the world—do not generate detectable bunching in housing transaction timing under the tighter Tax 2.0 regime. Despite a 10 percentage point rate reduction available to sellers who wait past 2 years, the density of transactions near the threshold is smooth. The bounded null is robust to estimation choices, present across price quartiles and geographic subsamples, and does not arise from insufficient power: the confidence interval rules out excess mass above approximately 0.5.

Why don't sellers bunch? Three candidate explanations deserve consideration. First, housing transactions involve substantial search and matching frictions. Unlike financial assets that can be sold at a click, residential property requires finding a buyer, negotiating, and completing legal transfer—a process that takes weeks or months. The "bunching region" may be too narrow for sellers to target precisely. Second, the tax may operate primarily on the extensive margin, deterring entry into short-term speculation rather than distorting the timing of transactions already in motion. If potential flippers never acquire the property in the first place, there is no holding period to bunch at. Third, Taiwan's robust housing market—with prices rising 20–30% in many areas between 2021 and 2024—may make the capital gains tax a secondary consideration relative to the expected returns from waiting.

These findings have implications for the design of anti-speculation taxes across East Asia. South Korea, Singapore, and Hong Kong have all introduced or considered holding-period-based capital gains taxes. The Taiwan evidence suggests that such taxes may not produce the intensive-margin distortions (lock-in, bunching) that tax theorists predict—a reassuring finding for policymakers concerned about liquidity costs, but one that raises the question of whether the taxes achieve their stated goal of discouraging speculation. Future work with unit-level property identifiers, which would eliminate the address-matching noise documented here, could provide sharper estimates and a cleaner test of the bunching prediction.

References

Appendix: Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classificat
Panel A: Pooled						
2-yr notch, Tax 2.0	−9	65	225	−0.042	0.286	Small negat
1-yr notch, Tax 1.0	−59	107	410	−0.144	0.261	Moderate negat
Panel B: Heterogeneous (Sample Splits)						
2-yr notch: Taipei	−15	15	31	−0.471	0.487	Large negat
2-yr notch: Other cities	−12	56	416	−0.030	0.134	Small negat

Notes: **Country:** Taiwan. **Research question:** How do holding-period capital gains tax notches affect the timing of housing transactions? **Policy mechanism:** Taiwan's Consolidated Housing Tax 2.0 (July 2021) imposes sharply higher capital gains tax rates on properties sold within 2 years of acquisition (45%) versus after 2 years (35%) and after 5 years (20%), creating strong incentives to delay sales past holding-period thresholds. **Outcome definition:** Excess mass of housing transactions measured as the normalized ratio of observed to counterfactual transaction density at holding-period notch thresholds, following the bunching methodology of Kleven and Waseem (2013). **Treatment:** Binary—the tax notch applies at discrete holding-period thresholds (730 days, 1,825 days). **Data:** Taiwan Actual Price Registration (Ministry of Interior), all building transactions 2012–2024, repeat-sale pairs matched by property address; unit of observation is a transaction pair. **Method:** Polynomial bunching estimator with 7th-order counterfactual density, 7-day bins, bootstrap standard errors (200 replications). **Sample:** Restricted to repeat-sale pairs with holding periods between 0 and 20 years; estimation window ±365 days around each notch. SDE = $\hat{\beta}/\mathrm{SD}(Y)$ where SD(Y) is the average counterfactual bin count. 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).

Acknowledgements

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

Contributors: @olafdrw

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

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