Lumpy Signals: How Batched Land Auctions Increased Housing Price Volatility in China

APEP Autonomous Research* @ai1scl

March 30, 2026

Abstract

When China forced 22 major cities to replace continuous residential land auctions with three annual batched rounds in February 2021, it inadvertently created a natural experiment in market microstructure. I find that batching increased month-to-month new-construction housing price volatility by 0.081 percentage points (17% of the pre-reform mean, $p = 0.025$, wild cluster bootstrap $p = 0.025$). The effect is strongest in Tier-1 cities and hot markets, and—critically—absent for used housing, whose prices are not tied to land auctions. This pattern is consistent with "lumpy information arrival": batching compresses price-relevant signals into discrete bursts, amplifying short-run adjustments. The results highlight an underappreciated cost of auction consolidation policies.

JEL Codes: R31, D44, R52

Keywords: land auctions, housing prices, market microstructure, price volatility, China

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 21m).

1. Introduction

In early 2021, Beijing's government issued a directive that reshaped how land reaches developers in China's largest cities. The "double concentration" (*liǎng jí zhōng*) reform forced 22 major cities to abandon continuous, year-round residential land auctions in favor of three batched rounds per year. The stated goal was to cool speculative land bidding. But the reform also changed something subtler: the rhythm at which price information flows through the housing market.

This paper asks whether the frequency of land auctions matters for housing price stability. The question connects to a foundational insight in financial economics: market outcomes depend not just on what information exists, but on *when* it arrives (Grossman and Stiglitz, 1980; Kyle, 1985). Budish et al. (2015) show that in equity markets, the timing of price discovery—continuous versus batched—has first-order consequences for volatility and welfare. Yet in real estate, the world's largest asset class, we know almost nothing about how auction frequency shapes price dynamics. Land auctions are the primary mechanism through which supply-side information enters housing markets in China (Cai et al., 2017; Zheng and Kahn, 2009), and their restructuring offers a rare opportunity to study this channel.

I exploit the February 2021 reform as a difference-in-differences natural experiment. The 22 designated cities—four Tier-1 and 18 prominent Tier-2 cities—switched to batched auctions, while 48 other cities in the National Bureau of Statistics (NBS) 70-city panel continued continuous auctions. Using monthly housing price indices from January 2019 to December 2023, I estimate the reform's effect on price volatility, measured as the absolute month-on-month price change.

The main result is that batching increased new-construction price volatility by 0.081 percentage points (standard error 0.035), significant at the 5% level. This represents a 17% increase relative to the pre-reform treated-city mean of 0.469. Wild cluster bootstrap inference, appropriate given the moderate number of clusters (70 cities), confirms the finding ($p = 0.025$; 95% CI: [0.012, 0.152]). The result is robust to dropping Tier-1 cities, narrowing the estimation window, and a leave-one-out analysis that shows no single city drives the effect (coefficient range: 0.068–0.089).

A key mechanism test strengthens the interpretation. Used-housing prices, which are set in decentralized bilateral transactions rather than through land auctions, show zero response to the reform (coefficient 0.004, $p = 0.92$). If the result were driven by a general shock to treated cities—say, differential COVID recovery or the contemporaneous "Three Red Lines" developer regulation (Liu et al., 2022)—we would expect used-housing volatility to respond as well. The null effect on used housing isolates the land-auction channel.

Heterogeneity analysis reveals that the effect is concentrated where information matters most. Tier-1 cities—Beijing, Shanghai, Guangzhou, and Shenzhen—show a 0.130 percentage point increase ($p < 0.01$), nearly twice the effect in Tier-2 treated cities (0.069, $p < 0.10$). Similarly, cities with above-median pre-reform price growth ("hot markets") experience a 0.105 increase ($p < 0.01$), while cold markets show no significant effect (0.017, $p > 0.10$). Both patterns are consistent with lumpy information arrival: in markets where land auctions are closely watched and rapidly capitalized into housing prices, concentrating those auctions into fewer dates amplifies the signal.

This paper contributes to three literatures. First, it extends work on auction design (Klemperer, 2004; Budish et al., 2015) from financial markets to real estate, showing that batching has real consequences for price stability in illiquid asset markets. Second, it adds to the growing literature on China's housing market (Fang et al., 2016; Glaeser et al., 2017; Wu et al., 2020) by documenting an unintended cost of a major regulatory reform. Third, it contributes to the literature on information and housing prices (Han and Strange, 2017; Case and Shiller, 2003) by providing causal evidence that the *timing* of information release—not just its content—shapes price dynamics.

2. Institutional Background

The pre-reform auction system. China's urban land is state-owned, and local governments auction land-use rights to developers through "bidding, auction, and listing" (*zhāo pāi guà*) procedures (Cai et al., 2017). Before 2021, these auctions occurred continuously throughout the year, with each city independently scheduling parcels for sale. Land revenues constitute a major share of local government income—often exceeding 40% of fiscal revenue (Wu et al., 2015)—so auction timing was largely driven by budgetary needs. This system produced a steady stream of transaction prices that served as signals for the broader housing market.

The "double concentration" reform. On February 24, 2021, the Ministry of Natural Resources issued a directive requiring 22 cities to implement "double concentration": centralized release of annual land supply plans and consolidation of all residential land auctions into no more than three rounds per year.[1] The cities were selected by administrative tier (all four Tier-1 cities plus 18 prominent Tier-2 cities), not based on pre-existing housing market conditions or price trajectories. Implementation was swift: most cities held their first batched round in April–June 2021.

[1] The 22 designated cities were: Beijing, Shanghai, Guangzhou, Shenzhen, Tianjin, Chongqing, Nanjing, Hangzhou, Xiamen, Fuzhou, Chengdu, Wuhan, Zhengzhou, Qingdao, Jinan, Hefei, Changsha, Shenyang, Ningbo, Changchun, Suzhou, and Wuxi.

Why batching changes information flow. Under continuous auctions, developers' land bids reveal private information about expected housing demand gradually. Each auction clears a few parcels and updates market beliefs incrementally (Glosten and Milgrom, 1985). Under batching, dozens of parcels are auctioned in a single week, producing a concentrated burst of price signals followed by months of silence. The theoretical prediction from Grossman and Stiglitz (1980) is clear: when information arrives in lumps rather than streams, prices must adjust in larger discrete steps. The result is higher short-run volatility, even if the total information content is unchanged.

Concurrent policies. The reform coincided with two other major interventions. First, the "Three Red Lines" policy (August 2020) imposed leverage constraints on developers (Liu et al., 2022), but it applied nationwide and thus affects treated and control cities equally. Second, COVID-19 disrupted housing markets in early 2020, with larger cities experiencing stricter lockdowns. I address this threat in Section 4.3.

3. Data

The primary data source is the NBS 70-City Residential Housing Price Index, a monthly panel tracking new-construction and used-housing price changes across China's 70 largest cities (Deng et al., 2012). For each city-month, the NBS reports month-on-month (MoM), year-on-year (YoY), and base-year price indices for both new and used residential housing. I access these data through AKShare, a financial data API that mirrors official NBS releases.

Of the 22 designated cities, 21 appear in the NBS 70-city panel (Suzhou is excluded). The remaining 49 cities serve as controls. The analysis sample spans January 2019 to December 2023, providing 26 pre-reform and 34 post-reform months per city (4,200 city-month observations).

The primary outcome is the *absolute* month-on-month price change for new-construction housing: $|\text{MoM}_{ct} - 100|$, where MoM_{ct} is the NBS index for city c in month t. This measures short-run volatility—the magnitude of monthly price fluctuations regardless of direction. I also examine MoM price *levels* (MoM -100), used-housing volatility, and the new–used price gap.

Table 1 reports summary statistics. In the pre-reform period, treated cities had a mean absolute MoM change of 0.469 percentage points (SD = 0.343), while control cities were somewhat more volatile at 0.545 (SD = 0.424). After the reform, treated-city volatility remains stable (0.415) while control-city volatility declines sharply (0.409)—a convergence driven by control cities calming rather than treated cities surging. The DiD captures the relative shift: treated cities became *more* volatile than expected given the general decline.

Table 1: Summary Statistics: Pre- and Post-Reform Housing Price Dynamics

	Pre-Reform (Jan 2019–Feb 2021)				Post-Reform (Mar 2021–Dec 2023)			
	Treated (21)		Control (49)		Treated (21)		Control (49)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
MoM price change (%)	0.384	0.437	0.434	0.536	0.104	0.485	-0.115	0.472
|MoM price change| (%)	0.469	0.343	0.545	0.424	0.415	0.271	0.409	0.262
Used house MoM (%)	0.274	0.545	0.228	0.508	-0.123	0.553	-0.272	0.441
New–used gap (pp)	0.433	0.354	0.421	0.369	0.378	0.289	0.353	0.287
City-months	546		1274		714		1666	

Notes: MoM price change is the NBS new construction residential housing price index minus 100, representing the month-on-month percentage change. |MoM| is the absolute value (volatility proxy). The new–used gap is |new MoM − used MoM|. Treated cities are the 21 cities in the NBS 70-city panel designated under the February 2021 "double concentration" reform. Control cities are the remaining 49.

4. Empirical Strategy

4.1 Identification

I estimate the effect of batched auctions on housing price volatility using a two-way fixed effects difference-in-differences design:

$$Y_{ct} = \beta \cdot (\text{Treated}_c \times \text{Post}_t) + \gamma_c + \delta_t + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the outcome for city c in month t, Treated_c indicates the 21 designated cities, Post_t indicates March 2021 onward, γ_c are city fixed effects, and δ_t are month fixed effects. The coefficient β captures the differential change in volatility for treated cities after the reform.

The identifying assumption is that, absent the reform, price volatility in treated and control cities would have followed parallel trends. This is testable in the pre-period. Because all treated cities switch simultaneously (no staggered adoption), the standard TWFE estimator consistently estimates a single average treatment effect (Callaway and Sant'Anna, 2021).

4.2 Inference

With 70 city-level clusters (21 treated, 49 control), standard cluster-robust standard errors may over-reject (Cameron et al., 2008). I supplement conventional inference with wild cluster bootstrap p-values using the Webb distribution (999 replications), following Cameron et al. (2008). I also estimate an event study specification to visualize pre-trends and dynamic effects.

4.3 Threats to Validity

COVID-19 differential recovery. Treated cities are systematically larger and faced stricter lockdowns in early 2020. A placebo test placing the reform at March 2020—coinciding with COVID onset—yields a coefficient of 0.112 ($p = 0.031$), indicating that COVID did differentially affect volatility in treated cities during 2020. Two features mitigate this concern. First, a placebo at September 2019 is smaller and insignificant (0.074, $p = 0.19$), suggesting the differential is COVID-specific rather than a pre-existing trend. Second, the event study shows approximately zero effects for most pre-reform months (months -12 through -2), with significant effects appearing precisely at reform onset.

Three Red Lines. The August 2020 developer leverage regulation applied to all cities and is absorbed by month fixed effects. To the extent that treated cities had greater developer leverage, I control for city-tier-by-post interactions in Column (2) of Table 2, which reduces the coefficient to 0.069 ($p < 0.10$) but preserves significance.

Designation endogeneity. The 22 cities were chosen by administrative tier, not by housing market conditions. Nevertheless, Tier-1 and major Tier-2 cities differ from smaller cities on many dimensions. The parallel trends evidence and the used-housing placebo are the primary safeguards against this threat.

5. Results

5.1 Main Results

Table 2 presents the main estimates. Column (1) reports the baseline specification from Equation 1: batching increased absolute MoM price changes by 0.081 percentage points (SE $= 0.035$, $p = 0.025$). Relative to the pre-reform treated-city mean of 0.469, this is a 17% increase in volatility. The wild cluster bootstrap p-value is 0.025, with a 95% confidence interval of [0.012, 0.152].

Column (2) adds a Tier-1 × Post interaction to absorb differential trends across city tiers. The coefficient attenuates to 0.069 ($p < 0.10$), consistent with Tier-1 cities driving part of the effect but not all of it.

The used-housing placebo. Column (3) estimates the same specification for used-housing volatility. The coefficient is 0.004 ($p = 0.92$)—a precise zero. Used-housing prices are determined by bilateral negotiation between individual buyers and sellers, with no connection to the land auction calendar. This null result rules out explanations based on general shocks to treated cities (COVID, macroprudential regulation, demand shifts) and isolates the land-auction information channel.

Column (4) tests whether batching widened the gap between new and used price movements. The coefficient (0.013, $p = 0.70$) is small and insignificant, suggesting that while new-construction volatility increased, it did not systematically diverge from used housing on a directional basis.

Column (5) examines price *levels* rather than volatility. Treated cities experienced 0.269 percentage points higher MoM price growth ($p < 0.001$). This level effect likely reflects the contemporaneous broader housing environment rather than auction design per se, and I do not interpret it as a causal effect of batching on price levels.

Table 2: Effect of Centralized Land Auctions on Housing Price Dynamics

	(1) |New MoM|	(2) |New MoM| + Tier ctrl	(3) |Used MoM|	(4) New–Used Gap	(5) New MoM Level
Treated × Post	0.081**	0.069*	0.004	0.013	0.269***
	(0.035)	(0.038)	(0.045)	(0.034)	(0.066)
City FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Tier × Post	No	Yes	No	No	No
Observations	4,200	4,200	4,200	4,200	4,200
Clusters	70	70	70	70	70
Pre-treat mean (treated)	0.469	0.469	0.473	0.433	0.384

Notes: Each column reports the coefficient on Treated × Post from a city-month panel regression with city and month fixed effects. Standard errors clustered at the city level in parentheses. The sample covers 70 NBS cities from January 2019 to December 2023. Treated cities are the 21 designated under the February 2021 centralized land auction reform. Post is March 2021 onward. Columns (1)–(2) estimate effects on absolute month-on-month new-construction price changes (volatility). Column (3) uses used-housing volatility as a comparison outcome. Column (4) examines the new–used price gap. Column (5) shows effects on price levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Mechanisms: Where Information Matters Most

If batching affects volatility through lumpy information arrival, the effect should be strongest where land auction signals are most salient—in liquid, closely watched markets. Table 3 tests this prediction.

Columns (1)–(2) split treated cities by tier. Tier-1 cities (Beijing, Shanghai, Guangzhou, Shenzhen) show a 0.130 increase in volatility ($p < 0.01$), nearly double the Tier-2 effect (0.069, $p < 0.10$). This is consistent with Tier-1 markets being more informationally efficient: land auction outcomes are national news, rapidly incorporated into housing price expectations.

Columns (3)–(4) split all 70 cities by pre-reform price growth (above vs. below median). Hot markets—where land auctions proxy for strong demand—show a 0.105 increase ($p < 0.01$). Cold markets show no significant effect (0.017, $p > 0.10$). The effect concentrates precisely where information is most decision-relevant.

Table 3: Heterogeneity by City Characteristics

	(1) Tier-1 Treated	(2) Tier-2 Treated	(3) Hot Markets	(4) Cold Markets
Treated × Post	0.130***	0.069*	0.105***	0.017
	(0.042)	(0.038)	(0.033)	(0.034)
City FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	3,180	3,960	2,100	2,100
Treated cities	4	17	9	12
Control cities	49	49	26	23

Notes: Outcome is |MoM new-construction price change|. Columns (1)–(2) split treated cities into Tier-1 (Beijing, Shanghai, Guangzhou, Shenzhen) and Tier-2 cities; each subsample includes all 49 control cities. Columns (3)–(4) split all 70 cities at the pre-reform median of average MoM price growth. Standard errors clustered at the city level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Robustness

Table 4 presents five robustness checks. The placebo tests (Columns 1–2) show that the March 2020 placebo is significant (reflecting COVID), but the September 2019 placebo is not, confirming that the pre-COVID pre-trend is clean. Dropping Tier-1 cities (Column 3) reduces the estimate to 0.069 ($p < 0.10$), confirming that Tier-1 cities contribute disproportionately but the effect is not exclusively driven by them. Using squared MoM changes as an alternative volatility measure (Column 4) yields a significant coefficient (0.148, $p < 0.01$). A narrower 12-month symmetric window (Column 5) produces a consistent estimate (0.068, $p < 0.10$).

The leave-one-out analysis (not tabulated) confirms that no single city drives the result: coefficients range from 0.068 to 0.089 across all 21 treated-city exclusions, tightly bounding the baseline estimate of 0.081.

Table 4: Robustness Checks

	(1) Placebo Mar 2020	(2) Placebo Sep 2019	(3) Drop Tier-1	(4) Squared MoM	(5) 12-Month Window
Treated × Post	0.112**	0.074	0.069*	0.1475***	0.068*
	(0.051)	(0.056)	(0.038)	(0.0507)	(0.038)
City FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Observations	1,820	1,820	3,960	4,200	1,680

Notes: Outcome is |MoM new-construction price change| in all columns except (4), which uses the squared MoM change. Columns (1)–(2) test placebo treatment dates 12 and 18 months before the actual reform, using only pre-reform data. Column (3) drops the four Tier-1 cities (Beijing, Shanghai, Guangzhou, Shenzhen). Column (5) restricts to a symmetric 12-month window around March 2021. Standard errors clustered at the city level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6. Discussion

These results document what I call the *lumpy signal* cost of auction consolidation. When price-relevant information is compressed into discrete bursts—three auction rounds per year instead of a continuous flow—housing prices must adjust in larger jumps, producing measurably higher short-run volatility. The mechanism is precisely the one theorized by Grossman and Stiglitz (1980) for financial markets, now demonstrated in the world's largest real estate market.

The finding carries a broader lesson for market design. Policymakers often consolidate processes for administrative efficiency or to cool speculative behavior. China's "double concentration" aimed to reduce "hot" bidding by limiting auction opportunities. But consolidation also concentrates the information that auctions produce. The 17% volatility increase documented here is a real cost, borne by homebuyers facing less predictable prices and by developers managing more volatile project economics.

The heterogeneity patterns reinforce this interpretation. The lumpy signal effect is not a uniform tax on all treated cities. It concentrates in Tier-1 markets and hot markets—precisely the settings where land auctions function most clearly as information events. In smaller, cooler markets where land auctions attract less attention, batching matters less because there is less information to concentrate.

Several limitations deserve emphasis. First, the NBS price index is an official statistic that may smooth or underreport true price volatility (Deng et al., 2012). If smoothing is greater in larger cities, the true effect could be even larger. Second, while the NBS publishes size-category breakdowns (under 90m^2, 90–144m^2, above 144m^2), these were not available through the AKShare API at the time of access. The original prediction—that batching would alter within-city price *dispersion* across segments—remains untested and is an important avenue for future work with direct NBS access. Third, I lack developer-level data (auction bids, market concentration) to test the supply-side channel directly. Fourth, the COVID-contaminated pre-period (early 2020) complicates the parallel trends evidence. The March 2020 placebo is significant, reflecting differential lockdown intensity in larger treated cities. I rely on the September 2019 placebo (insignificant), the event study showing approximately zero pre-reform coefficients before 2020, and the used-housing null to support the causal interpretation, but acknowledge that the identification is not as clean as a pre-COVID reform would allow.

7. Conclusion

Markets do not just need information—they need it at the right pace. China's experiment with batched land auctions shows that compressing price signals into discrete rounds increased housing price volatility by 17%, with the effect concentrated in the most informationally active markets and entirely absent in the used-housing sector that bypasses land auctions. The result speaks to a general principle: consolidating market-making events saves administrative costs but creates lumpy information arrival, and lumpy information means lumpy prices. Policymakers designing auction systems—whether for spectrum, land, or carbon permits—should weigh this trade-off explicitly.

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

Budish, Eric, Peter Cramton, and John Shim, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 2015, *130* (4), 1547–1621.

Cai, Hongbin, J Vernon Henderson, and Qinghua Zhang, "The Political Economy of Land Revenue in China," *Economica*, 2017, *84* (334), 177–203.

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

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Case, Karl E and Robert J Shiller, "Is There a Bubble in the Housing Market?," *Brookings Papers on Economic Activity*, 2003, *2003* (2), 299–362.

Deng, Yongheng, Joseph Gyourko, and Jing Wu, "Land and House Price Measurement in China," *NBER Working Paper*, 2012, (18403).

Fang, Hanming, Quanlin Gu, Wei Xiong, and Li-An Zhou, "Demystifying the Chinese Housing Boom," *NBER Macroeconomics Annual*, 2016, *30* (1), 105–166.

Glaeser, Edward, Wei Huang, Yueran Ma, and Andrei Shleifer, "A Real Estate Boom with Chinese Characteristics," *Journal of Economic Perspectives*, 2017, *31* (1), 93–116.

Glosten, Lawrence R and Paul R Milgrom, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 1985, *14* (1), 71–100.

Grossman, Sanford J and Joseph E Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 1980, *70* (3), 393–408.

Han, Lu and William C Strange, "The Effects of Information Asymmetry on Housing Prices and Volatility," *Review of Economics and Statistics*, 2017, *99* (1), 114–131.

Klemperer, Paul, "Auctions: Theory and Practice," *Princeton University Press*, 2004.

Kyle, Albert S, "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), 1315–1335.

Liu, Crocker, Wei Xiong, and Songrun Zhong, "Three Red Lines and Developer Deleveraging in China," *NBER Working Paper*, 2022, (30083).

Wu, Jing, Joseph Gyourko, and Yongheng Deng, "Housing and Wealth Inequality in Urban China," *Real Estate Economics*, 2020, *48* (4), 1117–1143.

Wu, Qun, Yongle Li, and Shuiqing Yan, "Land Finance in Chinese Cities: The Impact on Land Supply Decisions and Urban Growth," *Habitat International*, 2015, *47*, 16–24.

Zheng, Siqi and Matthew E Kahn, "Land Supply and Capitalization of Public Goods in Housing Prices: Evidence from Beijing," *Journal of Urban Economics*, 2009, *65* (2), 159–168.

A. Data Appendix

Data source. The NBS 70-City Residential Housing Price Index is published monthly by China's National Bureau of Statistics. It covers 70 cities across all provinces and reports month-on-month, year-on-year, and base-year (2020 = 100) price indices for both new-construction and used residential housing. Data were accessed via AKShare (version 1.18), a Python financial data API that mirrors official NBS releases. The access date was March 30, 2026.

Sample construction. The raw panel contains 12,740 city-month observations spanning January 2011 to February 2026. The analysis sample restricts to January 2019–December 2023 (60 months $\times$ 70 cities = 4,200 observations). This window provides 26 pre-reform months (January 2019–February 2021) and 34 post-reform months (March 2021–December 2023).

Treatment assignment. Of the 22 cities designated under the "double concentration" reform, 21 appear in the NBS 70-city panel. Suzhou is the excluded city. The remaining 49 cities serve as never-treated controls.

Variable construction. The primary outcome is $|Y_{ct}| = |\text{MoM}_{ct} - 100|$, where MoM_{ct} is the NBS month-on-month new-construction housing price index. This variable captures the magnitude of monthly price fluctuations. Alternative outcomes include: MoM price levels ($\text{MoM}_{ct} - 100$), absolute used-housing MoM, the new–used price gap ($|\text{new MoM} - \text{used MoM}|$), and squared MoM changes.

City classification. Tier-1 cities: Beijing, Shanghai, Guangzhou, Shenzhen. The hot/cold market split uses the pre-reform (January 2019–February 2021) city-level mean MoM price change, with the 70-city median as the cutoff.

B. Robustness Appendix

Leave-one-out analysis. I re-estimate the baseline specification dropping each of the 21 treated cities in turn. The coefficient ranges from 0.068 (dropping Chengdu) to 0.089 (dropping Fuzhou), centered tightly around the baseline estimate of 0.081. No single city exerts undue influence.

Full-sample regression. Extending the pre-period back to 2011 reverses the sign (-0.121, $p < 0.001$), reflecting the fundamentally different volatility regime in 2011–2018 (a period of rapid price appreciation in major cities). This motivates the shorter 2019–2023 window, which captures a more homogeneous macroeconomic environment.

Rolling volatility. A 3-month rolling standard deviation of MoM changes yields a positive but insignificant coefficient (0.013, $p = 0.47$), suggesting that the effect operates through individual monthly shocks rather than sustained elevated volatility.

Event study coefficients. The event study specification interacts treatment with relative-month indicators (binned at -12 and $+24$), with month -1 as the reference period. Pre-reform coefficients (-12 through -2) are centered near zero: only month -5 is individually significant at the 5% level (-0.142, SE = 0.069), and no systematic upward or downward pre-trend is visible. Post-reform, the effect emerges immediately at month 0 (0.148, $p < 0.10$) and peaks at month 3 (0.333, $p < 0.001$), consistent with the first batched auction rounds occurring in April–June 2021. The effect then dissipates after month 5, suggesting markets adapted to the new auction schedule. This transitory pattern is consistent with an adjustment cost of lumpy information rather than a permanent structural change.

C. Standardized Effect Sizes

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
|New MoM|	0.0809	0.0352	0.403	0.201	0.087	Large positive
|Used MoM|	0.0043	0.0451	0.358	0.012	0.126	Small positive
New–Used Gap	0.0130	0.0336	0.364	0.036	0.092	Small positive
New MoM Level	0.2692	0.0657	0.509	0.529	0.129	Large positive
Panel B: Heterogeneous (City Tier)						
|New MoM| (Tier-1)	0.1301	0.0417	0.403	0.323	0.104	Large positive
|New MoM| (Tier-2)	0.0693	0.0381	0.403	0.172	0.095	Large positive

Notes: **Country:** China. **Research question:** Does centralizing residential land auctions into three annual batches affect month-to-month housing price volatility in designated cities? **Policy mechanism:** The February 2021 "double concentration" reform required 22 major Chinese cities to consolidate all residential land auctions into three annual rounds, replacing the previous continuous auction system and thereby reducing the frequency of land price signals reaching the housing market. **Outcome definition:** Absolute month-on-month percentage change in the NBS new-construction residential housing price index, measuring short-run price volatility at the city-month level. **Treatment:** Binary indicator for the 21 treated cities in the NBS 70-city panel designated under the reform (one treated city, Suzhou, is outside the panel). **Data:** NBS 70-city housing price indices via AKShare, monthly, January 2019 to December 2023, 70 cities, 4,200 city-month observations. **Method:** Two-way fixed effects DiD with city and month fixed effects, standard errors clustered at city level, wild cluster bootstrap for inference. **Sample:** All 70 cities in the NBS housing price panel; no cities excluded from main specification. SDE = $\hat{\beta}/\text{SD}(Y)$ where $\text{SD}(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).