

Crowding In: How China's Affordable Rental Housing Program Affected Private Rents

MINGLEI LYU

t330036104@mail.bnbu.edu.cn

Beijing Normal-Hong Kong Baptist University

Research Article

Keywords: social housing, rental markets, crowd-out, difference-indifferences, China

Posted Date: October 13th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11340204/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Crowding In: How China’s Affordable Rental Housing Program Affected Private Rents

Anonymous submission (double-anonymous review)

October 2026

Abstract

Social rental housing is expanding at unprecedented scale in developing countries, yet its effect on the private rental markets that surround it remains poorly documented. This paper assesses how China’s Affordable Rental Housing (ARH) program — 8.7 million planned units between 2021 and 2025, the largest social-rental expansion in any country in recent decades — affected private market rents across 74 Chinese cities, using a monthly panel from 2018 to 2026. The program’s city-level plan targets, negotiated between municipal housing bureaus and the ministry, ranged from 4,700 to 600,000 units, providing quasi-experimental variation in the intensity of committed social-rental supply. Estimates from difference-in-differences designs that exploit this intensity variation, with and without a control group, show that rents rose by roughly 4.0 to 4.3 percent per standard deviation of planned units over the post-period. The effect appears within months of designation and is concentrated in cities with high pre-program rents and large inflows of new urban workers, consistent with demand-side crowding-in rather than the supply-side crowding-out documented in the U.S. literature on subsidized housing.

JEL codes: R21, R31, H53. *Keywords:* social housing; rental markets; crowd-out; difference-in-differences; China.

Every year for the past two decades, roughly twenty million people have moved into China’s largest cities, and most of them rent. Housing these newcomers has become a first-order policy problem, because rents in the major cities are high relative to the wages newcomers earn, and formal rental supply lags far behind demand. Between 2021 and 2025, the central government answered with the Affordable Rental Housing (ARH) program, 8.7 million planned units of small, rent-regulated flats for new urban workers, the largest social-rental expansion in any country in recent decades (Fang et al., 2016), with 6.5 million units concentrated in roughly 40 designated cities. The program is supply-conditional, in that rental land must be listed separately in each city’s annual land-supply plan; tax-privileged, with a value-added tax rate of 1.5 percent and a 4 percent property tax for qualifying landlords; credit-privileged, since ARH loans are exempt from

the real-estate lending caps that bind banks elsewhere, and performance-monitored, with ministry audits feeding into cadre evaluations. Cities therefore competed to commit, but committed to very different degrees. Whether an expansion of committed social-rental supply at this scale helps or harms the private rental markets around it is the question this paper answers, and the benchmark prediction comes from two decades of U.S. evidence pointing to crowd-out of unsubsidized supply.

This paper estimates the causal effect of the ARH program on private rents in 74 Chinese cities, using monthly data from August 2018 through September 2026. The question matters because the two leading theoretical predictions point in opposite directions, and which one wins is not obvious even in principle. On the substitution side, a rent-regulated unit is close to the private units that house the same workers. If the subsidized stock houses people who would otherwise rent privately, the residual private demand falls and market rents fall with it. This is the logic behind the U.S. findings. Low-income housing subsidies largely displace privately provided units (Sinai and Waldfogel, 2005), LIHTC development offsets itself within local markets (Eriksen and Rosenthal, 2010), and the end of rent control triggered conversions that shrank the low-cost stock (Sims, 2007). On the attraction side, subsidized supply lowers the effective cost of living in a city, and if workers are mobile the city gains tenants faster than it gains subsidized units, and private landlords then face steeper demand curves and raise rents. Which force dominates depends on details that mature-market evidence cannot pin down. Is tenancy bundled with local citizenship, how elastic informal supply is, and whether the tenant population is growing. The U.S. literature, built on slow-growing rental markets with an ample regulated stock (Saiz, 2010; Gyourko and Molloy, 2015), has largely answered in favor of crowd-out (Sims, 2007; Eriksen and Rosenthal, 2010; Baum-Snow and Marion, 2009), even as the wider affordability debate has turned toward supply (Been et al., 2019). None of its settings resembles a country injecting 8.7 million subsidized units into rapidly urbanizing cities.

Identification comes from variation in **program intensity across the designated cities themselves**, benchmarked against a control group of non-designated cities. Municipal ARH targets, negotiated between city housing bureaus and the ministry, ranged from 4,700 units in Haikou to 600,000 in Guangzhou: a 128-fold difference. Scaled by census population it spans 1.6 to 48.4 units per thousand residents (s.d. 9.1). In dynamic difference-in-differences designs with city and calendar-month fixed effects and house-price controls, a one-standard-deviation increase in planned intensity raises private rents by 4.0–4.3% over the post-period (within-designated-city dose: +0.44 log points per unit per 1,000 residents; dose against the control group: +0.47).

The identifying assumption is that, conditional on fixed effects, plan intensity is uncorrelated with differential rent shocks at the treatment date. The 28 pre-period coefficients are individually small, none significant, and a linear pre-trend test does not reject; joint pre-trend tests, however, are sensitive to covariance estimation with the small number of treated clusters, so we treat the pre-period evidence as necessary rather than sufficient and lean on trend-adjusted sensitivity bounds, which leave the conclusion unchanged. The simple binary comparison of designated versus non-designated cities fails pre-trend tests because large cities and smaller cities were already on

different rent trajectories before 2021: a level difference we document and discuss. Within a growing body of work that treats Chinese housing policies as quasi-natural experiments, covering purchase restrictions (Jia et al., 2018), housing-reform and labor outcomes (Chen et al., 2023), residence-policy relaxation (Tang et al., 2024), and credit tightening (Zhang and Wan, 2025), our setting adds a supply-side program observed at unprecedented scale. It speaks most directly to recent evidence that institutional landlords raised rents on comparable Irish properties (McCarthy, 2026): there, a demand-side shift by landlords priced rents up; here, a supply-side program addressed to tenants appears to do the same through the workers it attracts.

The paper relates to three literatures. First, the U.S. social housing literature: Sims (2007) shows that the end of Massachusetts rent control triggered condo conversions and shrank the low-cost rental stock; Sinai and Waldfoegel (2005) find that low-income housing subsidies largely crowd out the private occupied stock; Eriksen and Rosenthal (2010)’s instrumented estimates suggest LIHTC development offsets nearly all of itself within a ten-mile radius, though less so at wider geographies; Baum-Snow and Marion (2009) document neighborhood spillovers; Mast (2023) finds new market-rate construction filters down to relieve the low-income segment, and Diamond et al. (2019) show rent control redistributes toward incumbent tenants at the cost of reduced rental supply. Second, the land-use regulation literature links supply restrictions to segregation and renter welfare (Lens and Monkkonen, 2016; Saiz, 2010); this setting isolates a pure supply-side intervention, the opposite instrument, at unprecedented scale. Third, the Chinese housing literature has documented the market’s extraordinary expansion and recent turning point (Fang et al., 2016; Wu et al., 2012; Rogoff and Yang, 2021), but no published work evaluates the ARH program causally. Within the pages of this journal, Chinese housing policies have been treated as quasi-natural experiments repeatedly over the past decade: purchase restrictions and prices, credit tightening and developer debt, residence-policy relaxation and spatial equilibrium. The closest study in spirit (McCarthy, 2026) finds that institutional investors entering Ireland’s rental market raised rents on comparable properties by 4.1 percentage points, a demand-side repricing whose structural echo here is the program attracting the tenants it subsidizes. Across settings, the closest international evidence on social rental supply itself remains Lee (2007) on Korean public rental housing, which finds crowd-out.

This paper makes three contributions. First, it provides causal city-level estimates of the world’s largest active social-rental program, in a market where tenancy is decoupled from local public services and the subsidized stock, small units, occupancy-capped, rented below market, is horizontally differentiated from private supply; these are precisely the conditions under which the crowd-out mechanism has the hardest work to do. Second, it shows that the direction of the effect reverses relative to the U.S. benchmark: rents rise, and the rise concentrates in the high-rent, high-inflow cities where new workers concentrate, which is the placement a demand-side mechanism predicts and a supply-side substitution mechanism does not. Third, it makes a methodological point about evaluating policy rollouts: where designation itself correlates with market size, the within-group dose margin dominates the designated-versus-undesigned comparison, and we show concretely how the latter fails and why the former does not.

Section I describes the program and the data. Section II presents the empirical strategy. Section III reports the main results, the heterogeneity, and the mechanism evidence. Section IV presents the robustness battery, and Section V concludes.

I The Program and the Data

I.1 The program

The State Council’s Opinion No. 22 (June 2021) established ARH as the “inclusive” tier of China’s tripartite housing-security system. Units are capped at 70 m² ($\geq 70\%$ of stock) and rented below market rates to new urban workers without local settlement rights. Three design features matter for the analysis. First, the program is *supply-conditional*: land for rental housing must be listed separately in annual municipal land-supply plans, a meaningful lever in a system where urban land allocation is administratively controlled (Duranton and Puga, 2015). Second, it is *tax- and credit-privileged*: from October 2021, landlords renting through the program pay value-added tax at a reduced 1.5 percent rate rather than the standard 5 percent, roughly a 70 percent cut in effective VAT for rental businesses, plus a 4 percent property tax; from February 2022 (PBOC Notice No. 30), ARH loans are exempted from the real-estate loan concentration caps that bind banks elsewhere. Third, it is *performance-monitored*: the ministry audits the 40 designated cities and feeds results into cadre evaluations, and projects require certification to access benefits.

The scale is first-order. The national plan committed 8.7 million units for 2021–25, 6.5 million of them in the 40 designated cities, a volume comparable to roughly one year of national commodity-housing completions at the peak. Execution was substantial and fast. The first year alone delivered 942,000 units against a 936,000 national target; the cumulative count reached 3.3 million by the end of 2022 and 5.73 million by the end of 2023, about two-thirds of the national plan with two years to spare. Delivery leaned on the program’s distinctive financing architecture: rental-land listings priced below commodity-residential land, conversion of idle commercial and industrial stock at low regulatory cost, and, from February 2022, the exemption of ARH loans from the real-estate lending caps, which freed bank credit precisely as the broader property sector’s credit was being choked off. For the eight cities with auditable cumulative figures, completion rates as of 2024–25 range from 43% (Hefei) to essentially 100% (Beijing), with a median of 84% (Table I). The program was built; the announcements became buildings (Figure 1).

I.2 Designation and plan intensity

The ministry never published the designated list in full; it is reconstructed here from provincial confirmations (Jiangsu: Nanjing, Suzhou, Wuxi, Changzhou, Nantong) and municipal implementation documents. The reconstruction yields 41 candidate cities, essentially all provincial capitals, directly-administered municipalities, and the largest inflow-facing prefectures. The extra city relative to the official count of 40 reflects an arithmetic tension in the reconstruction: the Jiangsu province confirmation includes four non-capital prefectures, which cannot all be accommodated alongside the




Fig. 1: Program execution and the rent-effect platform. Bars: cumulative national ARH units completed (ministry releases). Shaded band: the post-2021 dose-effect platform from Figure 3b.

commonly listed Guangdong prefectures within a strict count of 40. Plan intensity, the Five-Year target scaled by 2020 census population, averages 14.0 units per thousand residents (s.d. 9.1). Variation is wide within the group: Xiamen (48.4), Guangzhou (32.1), and Hangzhou (25.1) anchor the top; Harbin (3.0), Tianjin (3.6), and Haikou (1.6) the bottom. The ARH push coincides with a structural turning point in Chinese housing, the end of two decades of breakneck construction (Rogoff and Yang, 2021), making its market consequences a first-order policy question alongside its construction targets (Fang et al., 2016).

This within-group variation is the paper’s identifying margin; Figure 2 plots the distribution.

Table I: ARH execution: plan vs. completed, eight auditable cities

City	Plan (10k units)	Completed (10k)	As of	Rate (%)
Beijing	40	40.0	2025	100
Wuhan	25	23.7	2024	95
Chengdu	25	22.5	2025	90
Hangzhou	30	26.7	2025	89
Shanghai	47	37.6	2024-06	80
Chongqing	40	30.0	2025	75
Guangzhou	60	34.1	2024	57
Hefei	15	6.5	2025-12	43

Completion rates as disclosed by municipal governments and the ministry; audit dates differ across cities (noted in the “As of” column). Sources: municipal housing bureaus, People’s Daily, city statistical bureaus.




Fig. 2: Plan intensity across designated cities (units per 1,000 residents). Within-group variation spans Haikou (1.6) to Xiamen (48.4).

I.3 The data

Rents. City-level monthly average asking rents (yuan/m²/month) for 74 cities, August 2018 through September 2026 (7,209 city-months), from the Wind EDB series compiled from Beike-system monthly rental reports: the only multi-city rental panel with pre-2021 depth currently available. Two properties of this series matter for reading the estimates. It measures asking, not transacted, rents: asking prices are the market’s posted signal and adjust quickly, which suits a design built around monthly timing, but they likely overstate levels in a softening market and understate the depth of negotiated discounts: the asking-versus-transaction wedge is a known feature of Chinese listing platforms, and we return to it in the conclusion. And it is a consistent index across cities and years: the same platform, the same aggregation, the same city definitions, so the design’s comparisons are within-series throughout. 43 cities have complete 98-month series; gaps (concentrated in 2020:08–2021:08) reflect upstream publication lags. Observations deviating more than $4 \times 1.4826 \cdot MAD$ from the city median are flagged as spikes (0.43% of observations; isolated outliers in Wuhan 2023:01, Chengdu 2023:01, and Nanjing 2023:02 likely reflect upstream transcription).

House prices. Official NBS 70-city second-hand house price indices (year-on-year), available for 60 of 74 cities; the house-price control therefore defines a slightly smaller estimation sample. **Population and GDP.** Annual city population and GDP (2015–2025) from statistical-bureau series. **Program data.** Plan targets from municipal Five-Year housing plans and implementation opinions (39 of 41 candidate cities have published or officially reconstructable totals, each source-verified); executed units from ministry releases and municipal disclosures. 2020 census population provides the intensity denominator. Three sample restrictions matter for what follows. First, five designated cities lack official house-price indices (Dongguan, Foshan, Nantong, Changzhou, and Suzhou, all provincial-priority cities in Guangdong and Jiangsu), so specifications controlling for house-price inflation cover 36 designated cities. Second, the dose-within design further requires a published intensity, leaving 34 cities. Third, 19 cities (11 designated) have a gap at 2021:07 in the upstream rent series: a publication artifact of the source, not a market event, but its timing inside the treatment window demands attention. Two checks bound its influence: dropping the month entirely moves the binary estimate from +6.64 to +6.63, and a balanced panel restricted to the 29 cities with complete in-sample series returns +8.1 ($p = 0.08$, about half the observations), so the missing month neither creates nor carries the result. Table II reports summary statistics for the analysis sample of Section I.3. Three features stand out. First, the rent level spans an order of magnitude across the sample, from under 11 yuan per square meter per month in the weakest markets to above 130 in Beijing at its peak, and the within-city variation around these levels is what identifies the design.

Second, plan intensity among the 39 designated cities with published targets (mean 14.0, median 12.5 units per thousand residents) is right-skewed: a handful of coastal cities planned at more than twice the median, which is the variation the dose design consumes.

Third, the sample years straddle an unusual cycle: house-price inflation averaged across city-

months is slightly negative over the window, reflecting the 2021–25 downturn, so the rent estimates are identified against a backdrop of falling prices rather than a boom that could mechanically explain a rent increase.

Table II: Summary statistics

	Mean	SD	P25	Median	P75	N
Rent (yuan/m ² /month)	28.9	17.0	19.0	24.8	29.5	74
Plan intensity (units/1k res.)	14.0	9.1	7.1	12.5	18.6	39
Plan target (10k units)	16.2	14.2	5.5	12.5	23.0	39
Population (10k, 2020)	858.4	542.2	498.0	763.1	997.5	74
GDP (100M yuan)	11376.5	10609.1	4671.6	8099.2	14133.2	74
House price y/y (%)	-1.0	6.0	-5.6	-2.4	3.3	5784

City-level statistics except house-price inflation (city-month observations). Rents averaged over 2018–2026; plan targets and intensity per the Five-Year plans (2021–25); population from the 2020 census. Spike-filtered rent observations excluded.

II Empirical Strategy

II.1 Dose design

The preferred specification exploits within-designated-city variation in program intensity:

$$\log Rent_{ct} = \beta (Intensity_c \times Post_t) + \alpha_c + \delta_t + \gamma \Delta HP_{ct} + \varepsilon_{ct}, \quad (1)$$

where $Intensity_c$ is city c ’s Five-Year ARH target scaled by its 2020 census population, measured in units per thousand residents, and $Post_t$ switches on from January 2021. The intensity variable deserves a description, because it is the object the whole design rests on. Targets were set city by city in negotiations between municipal housing bureaus and the ministry over 2020–2022; a city’s number reflects its land availability, fiscal capacity, and the local stock of candidates for conversion — collectively a bureaucratic planning decision, not a rental-market forecast. We scale by population so that a 600,000-unit commitment in Guangzhou, a city of nearly nineteen million, is comparable to a 25,000-unit commitment in Xiamen, a city of five million; the raw targets differ by a factor of 24, the scaled intensities by a factor of 30. No designated city planned zero, so the design has no within-group placebo; what it has is a steep dose gradient, from Haikou at 1.6 units per thousand residents to Xiamen at 48.4, and the question is whether rents tracked it.

January 2021 dates the policy information environment: the Central Economic Work Conference of December 2020 elevated “housing problems in big cities” to a top priority and municipal plans began circulating then; the May 2021 MOHURD symposia and the June 2021 State Council Opinion followed. Choosing a treatment date this early is a judgment call, and Table V shows what turns on it: estimates drift only from +6.5 at the earliest date to +5.4 at the latest, and the later, fully-executed dates give the smaller numbers because by 2022 the rents of designated cities were

already carrying the effect. Dating the program at its information moment treats the anticipation, which was real (cities announced plans, landlords repriced), as part of the treatment rather than as contamination. Standard errors are clustered at the city level. Because treatment is simultaneous across designated cities, the two-way fixed-effects specification is immune to the staggered-timing biases documented in the recent literature (Goodman-Bacon, 2021; Callaway and Sant’Anna, 2021; de Chaisemartin and D’Haultfœuille, 2020; Sun and Abraham, 2021), and the event studies are read alongside the imputation approach of Borusyak et al. (2024). Continuous-treatment designs carry their own requirements, strong monotonicity and dose-parallel trends in the terminology of Callaway et al. (2024), so we present the dose estimate twice: within designated cities only (34 cities, no outside control) and against the non-designated control group (60 cities), where the control cities anchor common shocks at dose zero. All event studies use the standard two-way fixed-effects interaction parameterization; an imputation (did2s) variant of the binary design (Borusyak et al., 2024) gives a similar post-period path but a somewhat different pre-period profile, and we use the TWFE version throughout for figure-statistic consistency. Pre-period coefficients are interpreted with the caveats of Roth (2022), and the sensitivity bounds are in the spirit of Rambachan and Roth (2023). City-level clustering follows the standard concerns about serial and spatial correlation in panel policy estimates (Bertrand et al., 2004).

Figure 3 (panel b) shows the dose event study: none of the 28 pre-period coefficients is individually significant, the largest lying within 0.16 log points ($\times 100$) of zero. Formal joint tests over 28 pre-period coefficients are sensitive to covariance estimation when clusters are few (footnote 1). The informative checks are a five-month nearest-window test ($p = 0.027$; a small, marginally significant near-term deviation) and a one-degree-of-freedom linear pre-trend test, which is marginal ($t = 1.87$, $p = 0.06$). We therefore read the pre-period as broadly consistent with identification but not probative, and place the weight on trend-adjusted sensitivity bounds in Section IV.¹ The identifying assumption is that, conditional on city and month fixed effects and house-price appreciation, plan intensity is uncorrelated with differential rent shocks at the treatment date. Three features support it. First, the targets were bureaucratically determined: a city’s number came out of negotiations over how much suitable land it could list, how much subsidy its budget could carry, and how many existing buildings were candidates for conversion. These inputs move with fiscal and land conditions, not with expected rent growth, and the city and month fixed effects absorb the level of exactly those conditions.

Second, the design decouples measurement from treatment. Intensity reflects commitments negotiated over 2020–2022 and enters the regression as a fixed city characteristic interacted with the post period. Identification therefore requires only that, once city and time effects are removed, anticipation did not scale with the plan number. It does not require that cities failed to anticipate rents at all.

¹Under a diagonal covariance approximation the 28-coefficient joint test passes ($\chi^2(28) = 36.2$, $p = 0.14$); under the full clustered covariance it rejects, but the statistic is dominated by an unstable high-dimensional inversion with as few treated clusters as this design has. The five-month window ($\chi^2(5) = 12.7$, $p = 0.027$) and the linear-trend test avoid the inversion problem. Estimates that add a group-specific linear trend remain positive (+4.1, $p = 0.15$).

Third, the program was actually delivered. Execution tracked plans: 942,000 units in the first year, 5.73 million cumulative by end-2023, and a median completion rate of 84 percent across the eight auditable cities in Table I, so β measures the effect of committed-and-delivered supply, not of empty announcements. One caveat on the house-price control: if ARH affects house prices themselves, including ΔHP absorbs part of the total effect; estimates without it are slightly larger (dose-with-control +0.51, binary +7.1), so the controlled estimates are, if anything, conservative.

II.2 The binary benchmark

The binary event study (Figure 3, panel a) shows clear pre-period separation: designated and non-designated cities differed by roughly 1.5 log points in rent trajectories well before 2021, with the designated cities on the softer path: the signature of the weaker large-city rental cycle of 2019–20, not of anticipation, since the separation long predates any policy information and deepens through 2020 when the pandemic hit dense service economies hardest. The pattern persists under earlier treatment dates and pandemic-window exclusions, and the near-window joint test rejects ($p = 0.016$).

The economic point is that China’s large-city rental markets were on a different trajectory from smaller ones for reasons unrelated to the program; any before-after comparison of the two groups would load that trajectory difference onto the policy. We retain the binary specification of Section II.2 as a between-group benchmark — it answers a different question, whether designated status as a package matters, and its estimate bounds the extensive-margin effect the dose design cannot see. It survives linear-trend extrapolation of its own pre-period (+9.5 log points), the clean-control sample excluding the provincially-designated cities present in the sample, and the full specification-curve grid (+3.3 to +8.2). The dose design carries the identification.




Fig. 3: Event studies: binary vs. dose design. Monthly coefficients with 95% confidence intervals (city-clustered), reference period $k = -1$, treatment date January 2021. Panel (a): 36 designated vs. 24 non-designated cities (house-price-control sample). Panel (b): dose with control group, 60 cities; coefficients per additional planned unit per 1,000 residents.

III Results

Section II.1 laid out the two designs; Table III presents the main estimates, with all coefficients in log points ($\times 100$) per additional planned unit per 1,000 residents. Column 1 is the within-designated-city dose: $+0.44$ ($p = 0.021$). Column 2 adds the 24 non-designated cities as a dose-zero control group and lands at $+0.47$ ($p < 0.001$), essentially unchanged, which says that cities outside the program experienced the same common rent shocks. Identification does not lean on the designated group alone. Column 3 adds a city-specific linear trend ($+0.45$). Columns 4 through 6 add province-by-year fixed effects, which absorb every province-level shock — national credit conditions, provincial subsidy allocations, regional market cycles. The dose coefficient survives at $+0.20$ within provinces and $+0.47$ with controls, meaning the identifying variation survives inside provinces, where all cities faced the same provincial plan; cities inside the same province, facing the same provincial policy environment, differ in plan intensity and rents move with that difference. The binary benchmark (columns 7 and 8) is $+6.6$ log points, or $+8.1$ with province-by-year effects — larger than the dose estimate, as it should be: designation bundled the ARH commitments with everything else a designated city did for rental supply in 2021–25, from its own subsidy budgets to its enforcement of landlord registration, while the dose coefficient isolates the margin that varies across otherwise-similar designated cities. The two designs bracket the program’s effect from different sides, and their agreement in sign through every fixed-effects configuration is itself evidence. A common confound would have to survive within-city, within-province, and dose-contrast comparisons simultaneously. Columns 9 and 10 drop the house-price control: the dose estimate rises to $+0.51$ and the binary to $+7.1$, so treating house prices as a pre-treatment control is, if anything, conservative.

How large are these effects economically? A standard deviation of plan intensity is 9.1 units per thousand residents, so columns 1 and 2 imply a rent increase of roughly 4.0–4.3 percent over the post-period. For the median designated city — 12.5 units per thousand residents planned — this is a rent level about 5.9 percent above its no-program counterfactual.

Set against the sample average rent of 29 yuan per square meter per month, a 4 percent difference is roughly 1.2 yuan per square meter each month, close to a tenth of a typical migrant worker’s monthly housing budget in these cities.

The subsidy side of the ledger matters for interpretation: eligible tenants of ARH units pay below-market rents by construction, so the program transfers to its direct beneficiaries while raising the price faced by the unsubsidized renting margin: the group the program does not, or not yet, cover.

In the dose event study (Figure 3b), effects appear immediately at the treatment date, build from $+0.34$ log points per unit per 1,000 residents over the first year to a plateau of $+0.54$ – 0.55 across months 24–47, and then decay to $+0.11$ (imprecisely estimated) after month 56, i.e. from late 2025. Two features of this path deserve note. First, the plateau arrives quickly, by the second year, not the fifth, which is more consistent with a demand response to credible committed supply than with a slow construction channel. The construction margin moves slowly: land listing, permitting, conversion, leasing-out stretches over years, and completions accumulate across the whole Five-Year

window — yet the rent response is at full height by month twelve, which points to the fast-moving margin: information and anticipation. The project certification system made each committed unit an auditable claim on future supply with attached tax and credit privileges, the land listings were public, and landlords on both sides of the market could observe the pipeline. An early plateau is what one expects if renters and landlords reprice against a credible announced pipeline; it is hard to produce if the effect had to wait for physical units.

Second, the tail: after month 56, roughly late 2025, the dose effect decays to +0.11 and is no longer precisely estimated, while the binary path in panel (a) moves the opposite way, rising toward +8. A sample truncated at mid-2025 leaves the static estimates essentially unchanged (binary +6.59, dose +0.51), so the average effect is not an artifact of the tail. For the divergence we offer a reading rather than a test: the large-city rental cycle that produced the binary design’s pre-period separation was still unwinding in 2025–26, and the binary path inherits it; the dose path, anchored by dose-zero cities, nets it out. We read the dose profile as the more reliable description of the program’s dynamic footprint, while flagging that its late-sample decay, market cycle, data quality, or genuine erosion, cannot be settled with these data.

The path is jointly far from zero: the static dose coefficient of Table III is precisely the efficient test of this post-period path ($t = 3.6$, $p < 0.001$), and 41 of 68 post-period coefficients are individually significant at 5%.

Table III: Main results: ARH plan intensity and private rents

	Coef. ($\times 100$)	SE	N
Dose (within, 34 cities)	0.44**	(0.18)	3,268
Dose (with control group, 60 cities)	0.47***	(0.13)	5,760
Dose within + city trend	0.45**	(0.18)	3,268
Binary	6.64***	(2.35)	5,760
Dose within + prov $\times$ year	0.20**	(0.08)	3,268
Dose with control + prov $\times$ year	0.47***	(0.13)	5,760
Binary + prov $\times$ year	8.05***	(2.43)	5,760
Dose with control, NO hpi control	0.51***	(0.14)	5,760
Binary, NO hpi control	7.06***	(2.41)	5,760

Dependent variable: log city-month asking rent. Dose-within sample: 34 designated cities with published targets and house-price controls (no outside control). Dose-with-control and binary samples: 60 cities with house-price controls (36 designated vs. 24 non-designated). All specifications include city and calendar-month FE and the second-hand house-price inflation control. Post = January 2021. City-clustered SEs. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

III.1 Heterogeneity and mechanisms

The U.S. crowd-out benchmark makes a prediction that goes beyond the sign of the average effect: if subsidized units substitute for private ones, the rent decline should be largest where the two stocks compete head-on — cheap cities, formal rental segments, and markets where the tenant base

is stable. A demand-side attraction mechanism makes the opposite placement prediction: the effect should live where new tenants want to settle and can now afford to. These are testably different, and Table IV puts them against each other. The estimates show increases, concentrated exactly where the demand-side mechanism predicts. Interacting dose with pre-period city characteristics (Table IV, with formal difference tests): effects are present in above-median-rent cities (+0.49 per additional unit per 1,000 residents, $p < 0.001$) and small and insignificant in below-median-rent cities (+0.16); effects are more than three times larger in above-median-population cities (+0.65, $p < 0.001$) than below (+0.19), a difference significant at $p < 0.001$. Both patterns are consistent with ARH lowering the effective cost of settling in expensive, inflow-heavy cities, expanding rental demand faster than the subsidized stock expands supply. Reading Table IV group by group: the rent split puts the entire effect above the median — +0.49 per additional unit per 1,000 residents ($p < 0.001$) in above-median-rent cities against +0.16 ($p = 0.46$) below, a difference significant at $p = 0.15$ with this sample’s power. The population split is the sharpest: +0.65 ($p < 0.001$) in above-median-population cities against +0.19 ($p = 0.005$) below, and the difference itself is significant at $p < 0.001$ — the strongest single piece of evidence in the paper for the attraction channel, because population is the most direct proxy for the inflow the mechanism requires.

By contrast, effects are statistically indistinguishable between cities with rapid and slow 2018–20 house-price growth (+0.48 vs. +0.47, difference test $p = 0.98$). This null difference speaks to the third candidate mechanism: a land-channel story in which listing rental land crowds out commodity-housing land, raising house prices and thence rents through the own-or-rent margin. If that channel operated, cities where land scarcity binds hardest should show the largest rent effects; they do not, and the absence of a house-price-growth gradient suggests the program’s rental-market footprint runs through the rental market specifically rather than through land markets. Figure 4 plots the grouped event studies: the above-rent and above-population groups drive the aggregate effect, with small and statistically quiet pre-periods in each group.




Fig. 4: Grouped dose event studies. Each panel re-estimates the dose design (Equation 1) separately for above- and below-median cities, with non-designated cities as controls. Panel (a): 2019 rent level. Panel (b): 2020 population.

The reversal of sign relative to the international benchmark is the paper’s central finding, and

it is worth being precise about what is being reversed. Lee (2007) reports that Korean public rental provision crowds out private rental investment; U.S. evidence on subsidized supply points the same way (Sinai and Waldfogel, 2005; Eriksen and Rosenthal, 2010; Baum-Snow and Marion, 2009). Three features distinguish the Chinese setting. First, tenancy in China carries no local-service entitlement, schooling and social insurance attach to hukou, not to a lease, so a subsidized rental unit does not buy access that a private unit would; the two are horizontally differentiated products, not substitutes on a single margin. Second, the private rental stock in these cities is dominated by small, informal landlords, so the supply response at the margin is elastic rather than the constrained corporate supply of mature markets.

Third, and most important, the program targets a growing population. The cities in our sample are exactly the ones that keep attracting workers on consumption and productivity margins (Glaeser et al., 2001; Moretti, 2011); new urban workers arrive whether or not subsidized units exist, and the program lowers the effective cost of settling in the expensive cities they prefer. Each feature weakens the substitution channel through which crowd-out operates in mature markets.

The pattern in Table IV is what this account predicts: the rent effect lives in the high-rent, high-inflow cities where new workers concentrate, and is absent where the marginal arrival is scarce. Consistent with a demand-side reading, housing-voucher design changes raise rents where recipients concentrate (Collinson and Ganong, 2018), and Diamond et al. (2019) show incumbent tenants capturing large benefits even as rent control contracts supply; here the subsidized stock is new supply rather than a cap on existing supply, and the private market prices up the demand it attracts.

We stress the limits of the claim: with city-level rents we observe the pricing consequence, not the migrating households, and we cannot decompose the demand response between new arrivals and within-city tenure choices. What the data do rule out is the benchmark prediction that committed social-rental supply at this scale would depress private rents. A final measurement caveat belongs here rather than in a footnote: if the program shifted the composition of listings, more subsidized-adjacent units, or formal landlords replacing informal ones, the asking-rent series could move for reasons that are not pure price adjustment. The direction of such composition effects on our estimate is ambiguous, and the series gives us no way to separate composition from price. The reassurance is indirect but real: composition stories do not naturally produce the dose gradient, the heterogeneity placement, or the timing that the estimates show. Two channels cannot be separately measured with available data — internal migration responses and the composition of private supply (formal vs. informal), and are discussed as limitations in Section V.

IV Robustness

We organize the robustness checks around four concerns: the exact start date of treatment, contamination from provincial policies or the pandemic, unreliable inference with few treated clusters, and a linear pre-trend. Table V addresses each, and all rows use the January 2021 treatment date. *Timing*: the binary estimate ranges +5.4 to +6.6 log points across five treatment dates from December 2020

Table IV: Heterogeneity: dose effects by pre-period city characteristics

	Above median		Below median		Diff. p
	Coef.	p	Coef.	p	
High 2019 rent (calendar year)	0.49	(0.001)	0.16	(0.456)	(0.147)
Large population 2020	0.65	(0.000)	0.19	(0.005)	(0.001)
High 2018-20 HPI growth (corrected)	0.48	(0.009)	0.47	(0.004)	(0.979)

Dose-with-control specification of Table III with the dose variable interacted with above/below-median group indicators (city-level, pre-period constructed). Coefficients in log points ($\times 100$) per additional planned unit per 1,000 residents; p -values in parentheses. “Diff. p ” tests the within-regression equality of the two group coefficients. Grouping variables: mean 2019 rent (calendar year); 2020 census population; cumulative 2018–20 house-price inflation.

to January 2022.

Samples: excluding the provincially-designated cities present in the sample raises nothing (+6.3), and excluding lockdown-peak months does the same (+6.7). A balanced panel of the 29 cities with complete in-sample series returns +8.1 ($p = 0.08$, about half the observations, reported in text). Dropping the gap month 2021:07 itself gives +6.6, also in text.

Inference: province-clustered standard errors (+6.6, $p = 0.009$), and randomization inference, which sidesteps asymptotic approximations altogether. Reassigning designated status across the 60 sample cities 300 times places the true binary estimate beyond all but two percent of its placebo distribution. For the dose design the version that matches the estimand reshuffles the intensity values themselves across cities, preserving the intensity distribution while severing its link to any particular city — Figure 5 plots the resulting distribution of placebo dose coefficients, tightly centered below the observed +0.45, with none of the 300 draws reaching it. This test also carries weight against the spatial-spillover concern: because placebo doses are scattered across the same geographic space as the true ones, a spillover-driven estimate would survive random reshuffling only by luck, and it does not.

Trend adjustment: extrapolating the (small, negative) pre-period linear trend of the binary event study into the post-period leaves +9.5; correcting with group-specific linear trends instead leaves +4.1 ($p = 0.15$) — still positive, though imprecisely so, and we flag this as the least favorable specification.

The specification curve (Figure 6) spans +3.3 to +8.2 log points across thirty timing $\times$ fixed-effects $\times$ sample combinations, all positive. Fake treatment dates before 2019 either render the treatment collinear with fixed effects (2018:07) or return a smaller, marginally significant coefficient consistent with the documented pre-trend drift (2019:07). No check moves the sign; the group-trend specification is the one caveat, and even there the point estimate is 70 percent of the baseline. Taking the battery together, the estimate that survives the least friendly specification is still a rent increase of roughly 2 percent per standard deviation of intensity — economically meaningful, and opposite in sign to the crowd-out benchmark. It is worth being explicit about what the battery

does and does not establish. It establishes that the finding is not an artifact of one treatment date, one control set, one sample definition, or one covariance approximation, and that the dose estimate in particular survives the trends that contaminate the binary design. It does not establish that some unobserved shock perfectly correlated with plan intensity and rising rents is impossible; no design of this type can. What it establishes is that such a shock would need to track municipal plan negotiations closely enough to generate a dose gradient, hit exactly at the designation date, and leave no trace in 28 pre-period coefficients: a demanding assignment. Finally, because 20 of the 60 sample cities share provinces with designated cities, and rent shocks may spill across nearby markets, city-clustered standard errors address serial but not spatial dependence; the randomization inference, which reassigns treatment across the full city distribution, is the more robust of the two approaches to such spillovers. With the estimate defended, the remaining questions are where the effect lives and what it means: the heterogeneity and mechanism evidence of Section III.1.




Fig. 5: Randomization inference for the dose design: placebo dose coefficients from randomly reshuffling dose intensities across sample cities (300 draws), preserving the intensity distribution. Red line: observed coefficient; empirical $p < 0.001$.

V Conclusion

This paper examines how the largest social-rental housing expansion in any country in recent decades affected the private rental markets around it, using the variation in city-level commitments within China’s Affordable Rental Housing program as its identifying margin. Exploiting variation in committed program intensity across China’s designated Affordable Rental Housing cities, targets negotiated between city housing bureaus and the ministry that ranged from 4,700 to 600,000 units, we find that private rents rose rather than fell: by roughly 4.0 to 4.3 percent per standard




Fig. 6: Specification curve: the binary estimate across 30 combinations of treatment timing, fixed effects, and sample. All specifications positive.

Table V: Robustness

Check	Coef. ($\times 100$)	SE / note
Timing: Post = Dec 2020 (CEWC)	6.47***	(2.34)
Timing: Post = Jan 2021 (preferred)	6.64***	(2.35)
Timing: Post = May 2021 (symposia)	6.34***	(2.28)
Timing: Post = Jul 2021 (Opinion 22)	5.94***	(2.22)
Timing: Post = Jan 2022	5.44**	(2.07)
Clean control (drop 10 provincial pilots)	6.26**	(2.39)
Excl. lockdown months (2022)	6.74***	(2.37)
Province-clustered SEs	6.64***	(2.36)
Placebo: fake post = Jul 2019	4.40**	(2.14)
Group-specific linear trend	4.14	(2.86)
Randomization inference (300 draws)	6.64***	(2.65)
R10 linear pre-trend extrapolation	9.54	—

Binary benchmark unless noted. Post = January 2021 except timing rows. City-clustered SEs. The binary pre-trend joint test (see text, Section III.B) and the fake date 2018:07 (treatment collinear with fixed effects, unidentified) are omitted from the table and discussed in the text.

deviation of plan intensity over 2021–2025, robust to the inclusion of a non-designated control group, province-by-year shocks, house-price controls, trend adjustment, and randomization inference. The effect appears within months of designation, plateaus by the second year, and is concentrated in the high-rent, high-inflow cities the program most aims to help; by contrast, it does not vary with house-price pressure, indicating that the program operates through the rental market specifically.

These findings reverse the sign the U.S. literature has taught the field to expect, for reasons this setting makes visible. Chinese tenancy carries no local-service entitlement. The private rental stock is informal and elastically supplied. And the tenant population is still growing. In these conditions, committed social-rental supply works less as a substitute for private rentals than as urban attraction infrastructure. It lowers the cost of settling for the workers a city wants, and the private market raises rents on the demand they bring.

For policy, the results caution against reading the ARH program as rent relief. Its affordability benefits accrue to eligible tenants of the subsidized stock; the unsubsidized renting margin, often the near-eligible and the not-yet-eligible, pays the higher rents. If rent stabilization for new urban workers is the objective, the scale of subsidized supply would need to stand against the inflow it attracts rather than against the existing stock alone. The broader stakes go beyond rents: binding urban housing supply depresses aggregate output by keeping workers out of productive places (Hsieh and Moretti, 2019), so a program that lets more workers stay in the big cities — even at higher private rents — has a first-order efficiency claim that a rent-centric reading misses.

The analysis has limits. Rents are asking rather than transacted prices, and asking rents may overstate the slowdown of a weakening market; the house-price control restricts the sample to 36 of the 39 cities with published targets, and the excluded cities concentrate in Guangdong and Jiangsu; effects decay after mid-2025 for reasons these data cannot distinguish, and the demand response cannot be decomposed between migration and within-city tenure choice. Tracing the workers themselves, who comes, who pays the higher rent, and who gets the subsidized unit, is the natural next step, and the one on which the policy verdict ultimately turns.

References

- Baum-Snow, N., & Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics*, 93(3–4), 654–666. <https://doi.org/10.1016/j.jpubeco.2009.01.001>
- Been, V., Ellen, I. G., & O'Regan, K. (2019). Supply skepticism: Housing supply and affordability. *Housing Policy Debate*, 29(1), 25–40. <https://doi.org/10.1080/10511482.2018.1476899>
- Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1), 249–275. <https://doi.org/10.1162/003355304321042836>
- Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, 91(6), 3253–3285. <https://doi.org/10.1093/restud/rdae007>
- Callaway, B., Goodman-Bacon, A., & Sant'Anna, P. H. C. (2024). Difference-in-differences with a continuous treatment (NBER Working Paper No. 32117). National Bureau of Economic Research. <https://doi.org/10.3386/w32117>
- Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230. <https://doi.org/10.1016/j.jeconom.2020.12.001>
- Chen, J., Hu, M., & Lin, Z. (2023). China's housing reform and labor market participation. *The Journal of Real Estate Finance and Economics*, 67(2), 218–242. <https://doi.org/10.1007/s11146-021-09827-3>
- Collinson, R., & Ganong, P. (2018). How do changes in housing voucher design affect rent and neighborhood quality? *American Economic Journal: Economic Policy*, 10(2), 62–89. <https://doi.org/10.1257/pol.20170099>
- de Chaisemartin, C., & D'Haultfoeulle, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996. <https://doi.org/10.1257/aer.20181169>
- Diamond, R., McQuade, T., & Qian, F. (2019). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco. *American Economic Review*, 109(9), 3365–3394. <https://doi.org/10.1257/aer.20181289>
- Duranton, G., & Puga, D. (2015). Urban land use. In G. Duranton, J. V. Henderson, & W. C. Strange (Eds.), *Handbook of regional and urban economics* (Vol. 5, pp. 467–560). Elsevier. <https://doi.org/10.1016/B978-0-444-59531-7.00008-7>
- Eriksen, M. D., & Rosenthal, S. S. (2010). Crowd out effects of place-based subsidized rental housing: New evidence from the LIHTC program. *Journal of Public Economics*, 94(11–12), 953–966. <https://doi.org/10.1016/j.jpubeco.2010.07.002>
- Fang, H., Gu, Q., Xiong, W., & Zhou, L.-A. (2016). Demystifying the Chinese housing boom. *NBER Macroeconomics Annual*, 30, 105–166. <https://doi.org/10.1086/681598>
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277. <https://doi.org/10.1016/j.jeconom.2021.03.014>

- Glaeser, E., Kolko, J., & Saiz, A. (2001). Consumer city. *Journal of Economic Geography*, 1(1), 27–50. <https://doi.org/10.1093/joeg/1.1.27>
- Gyourko, J., & Molloy, R. (2015). Regulation and housing supply. In G. Duranton, J. V. Henderson, & W. C. Strange (Eds.), *Handbook of regional and urban economics* (Vol. 5, pp. 1289–1337). Elsevier. <https://doi.org/10.1016/B978-0-444-59531-7.00019-1>
- Hsieh, C.-T., & Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2), 1–39. <https://doi.org/10.1257/mac.20170388>
- Jia, S., Wang, Y., & Fan, G.-Z. (2018). Home-purchase limits and housing prices: Evidence from China. *The Journal of Real Estate Finance and Economics*, 56(3), 386–409. <https://doi.org/10.1007/s11146-017-9615-2>
- Lee, C.-I. (2007). Does provision of public rental housing crowd out private housing investment? A panel VAR approach. *Journal of Housing Economics*, 16(1), 1–20. <https://doi.org/10.1016/j.jhe.2006.09.004>
- Lens, M. C., & Monkkonen, P. (2016). Do strict land use regulations make metropolitan areas more segregated by income? *Journal of the American Planning Association*, 82(1), 6–21. <https://doi.org/10.1080/01944363.2015.1111163>
- Mast, E. (2023). JUE Insight: The effect of new market-rate housing construction on the low-income housing market. *Journal of Urban Economics*, 133, 103383. <https://doi.org/10.1016/j.jue.2022.103383>
- McCarthy, B. (2026). Institutional investment and residential rental market dynamics. *The Journal of Real Estate Finance and Economics*. Advance online publication. <https://doi.org/10.1007/s11146-025-10037-4>
- Moretti, E. (2011). Local labor markets. In O. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (Vol. 4B, pp. 1237–1313). Elsevier. <https://doi.org/10.1016/B978-0-444-53542-5.00004-9>
- Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *The Review of Economic Studies*, 90(5), 2555–2591. <https://doi.org/10.1093/restud/rdad018>
- Rogoff, K., & Yang, Y. (2021). Has China’s housing production peaked? *China & World Economy*, 29(1), 1–31. <https://doi.org/10.1111/cwe.12360>
- Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3), 305–322. <https://doi.org/10.1257/aeri.20210236>
- Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3), 1253–1296. <https://doi.org/10.1162/qjec.2010.125.3.1253>
- Sims, D. P. (2007). Out of control: What can we learn from the end of Massachusetts rent control? *Journal of Urban Economics*, 61(1), 129–151. <https://doi.org/10.1016/j.jue.2006.09.001>
- Sinai, T., & Waldfoegel, J. (2005). Do low-income housing subsidies increase the occupied housing stock? *Journal of Public Economics*, 89(11–12), 2137–2164. <https://doi.org/10.1016/j.jpubeco.2004.06.015>

- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199. <https://doi.org/10.1016/j.jeconom.2020.09.006>
- Tang, R., Tang, Y., & Zhang, R. (2024). The aggregate and distributional impacts of residence policy relaxation. *The Journal of Real Estate Finance and Economics*, 69(2), 343–376. <https://doi.org/10.1007/s11146-022-09912-1>
- Wu, J., Gyourko, J., & Deng, Y. (2012). Evaluating conditions in major Chinese housing markets. *Regional Science and Urban Economics*, 42(3), 531–543. <https://doi.org/10.1016/j.regsciurbeco.2012.01.004>
- Zhang, B., & Wan, X. (2025). Credit tightening and housing prices: Empirical research on the channel of the housing supply side. *The Journal of Real Estate Finance and Economics*, 71(4), 615–641. <https://doi.org/10.1007/s11146-023-09974-9>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Replicationpackagecodeandcleaneddata.zip](#)