Housing Tenure, Welfare, and Distributional Effects of Short-Term Rental Platforms*

Ertunc Aydogdu† Fabio Braggion‡ Mauro Mastrogiacomo§
Nicola Pavanini¶ Yushi Peng||

November 2026

Abstract

Cities restrict short-term rentals (STR) to protect residents, but which residents benefit is unclear: a unit leaving STR can be let to a tenant or sold to an owner-occupier, and who can buy depends on borrowing limits. We study Amsterdam's 2017 60-night Airbnb cap using listings linked to administrative records on households, transactions and mortgages. Per standard deviation of exposure, the cap reduced reservation days by 14% and house prices by 1.1%. We estimate an equilibrium model with borrowing-constrained households, visitors, and landlords choosing between STR, long-term rental and selling. Tightening the cap to 30 nights sends three in four units leaving STR to owner-occupiers. Rents and prices fall by 5%, and lower-income purchases rise most, a gain tighter loan-to-income limits partly undo. Residents gain €1.6 billion, but landlords and visitors lose €4.9 billion. Per unit removed, a 10% tax delivers similar resident gains at a third of the cost.

JEL Classification: R31, R38, L85, G51, D15
Keywords: short-term rentals, Airbnb, housing tenure, rental supply, borrowing constraints

*We thank Lu Han, Tobias Klein, Matthijs Korevaar, Florian Sniekers for helpful comments. We also thank seminar participants at Tilburg University, De Nederlandsche Bank for valuable feedback. Nicola Pavanini gratefully acknowledges financial support from the European Research Council (Grant Agreement No. 101042069). The views expressed herein are those of the authors and do not necessarily reflect the views of De Nederlandsche Bank or any affiliated institution.

†Nova School of Business and Economics. Email: ertunc.aydogdu@novasbe.pt.
‡Tilburg University and CEPR. Email: f.braggion@tilburguniversity.edu.
§De Nederlandsche Bank and Vrije Universiteit Amsterdam. Email: m.mastrogiacomo@vu.nl.
¶Tilburg University and CEPR. Email: n.pavanini@tilburguniversity.edu.
||Tilburg University. Email: y.peng@tilburguniversity.edu.

1. Introduction

Cities around the world are restricting short-term rentals (STR). New York effectively banned most short-term rentals of entire homes in 2023, Barcelona has announced it will phase out all tourist-apartment licences by 2028, and the European Union has introduced common registration and data-sharing rules for platforms.[1] The stated rationale is almost always the same. When a home is let to tourists on a platform such as Airbnb, it is no longer available to residents. In a city where the housing stock adjusts slowly, this *supply withdrawal* reduces the number of homes available for long-term rental (LTR) tenants and owner-occupiers, and raises both rents and house prices. Restricting STR should reverse it and benefit residents, but which residents benefit most is far from obvious.

When STR becomes less profitable, a landlord can convert the unit to an LTR, but they can also sell it to an owner-occupier. Units converted to LTR lower rents for tenants, while units sold lower house prices for buyers. Which of the two responses dominates therefore determines whether an STR restriction benefits mainly renters or buyers. Who can purchase the units that are sold depends on credit regulations, particularly limits on loan-to-value (LTV) and loan-to-income (LTI) ratios. Credit regulations therefore also determine who benefits from STR restrictions. This interaction is particularly important for households in the free-market rental segment, who earn too much to qualify for social housing but face borrowing constraints that prevent them from becoming homeowners.

In this paper we quantify these trade-offs in Amsterdam, where the municipality began enforcing a 60-night annual limit on entire-home STR in April 2017. Before enforcement, Airbnb listings in Amsterdam nearly tripled between 2015 and their 2017 peak, from about 4,000 to 11,000 per month, while the housing stock grew by less than 1.5% a year. Enforcement began when Airbnb started automatically blocking listings that reached the limit. We link listing-level Airbnb data to administrative registers covering every Amsterdam household (including income, wealth, and tenure), the universe of housing transactions, rents from the national rental survey, loan-level mortgage data from De Nederlandsche Bank, as well as neighborhoods' attributes and tourism flows.

We first use the 2017 enforcement as a natural experiment, comparing districts in Amsterdam with high and low pre-regulation Airbnb exposure. Per standard deviation of exposure, the cap reduced entire-home listings by 6.5% and total reservation days by 14.3%, while nightly Airbnb prices rose by about 2%, consistent with a supply contraction. Quality-adjusted house prices fell by 1.1% and mortgage amounts by 1.7%,

[1] New York City, Local Law 18 of 2022; Barcelona, Catalan Decree-Law 3/2023; European Union, Regulation (EU) 2024/1028.

with no change in LTV, LTI or mortgage rates. A shift-share design gives a house-price elasticity with respect to Airbnb density of 0.17 to 0.31, and shows that free-market rents respond (an elasticity of 0.07) while regulated rents, which are set by an administrative points system, do not. This asymmetry is what supply withdrawal predicts: only rents that are free to respond to fewer available units should rise. These estimates establish that supply withdrawal is at work in Amsterdam, but on their own they cannot tell us where the units went once the cap made STR less profitable, who gained and who lost, or how alternative policies would perform.

We therefore build and estimate an equilibrium model of Amsterdam's private housing market, with housing products defined by 23 areas and three property-size classes. The model has three types of agents. Households make dynamic tenure and location choices subject to LTV and LTI constraints, with moving costs. Visitors choose between Airbnb listings in different areas and hotels. Investors (landlords) decide each year how to split their units between STR and LTR, a decision that we define as the *inner nest*, and how many units to keep as rentals rather than sell to owner-occupiers, which instead we call the *outer nest*. Ownership prices, long-term rents and STR prices clear the three markets jointly. The outer nest is the key departure from existing single-margin STR models (Calder-Wang, 2021; Almagro and Domínguez-Iino, 2025): it lets a restriction shrink the rental sector itself, and the household block determines who absorbs the units investors sell.

We estimate the three model blocks sequentially. First, household preferences come from conditional choice probabilities (Hotz and Miller, 1993), with price sensitivity identified from financial moving costs following Bayer et al. (2016) and Kvaerner et al. (2023). Second, visitor demand is a logit model where STR prices are instrumented with the bite of the day-cap and the characteristics of competing listings in neighboring districts. Third, the two investor margins are identified from a shift-share instrument for global Airbnb demand and from ECB-driven changes in mortgage rates.

Three results stand out from the model estimates and counterfactuals. First, restricting short-term rentals benefits buyers more than tenants. Investors' decision to keep a unit as a rental or sell it to an owner-occupier responds to a change in returns about as strongly as their decision to list it on Airbnb or let it long term. Per euro of return, the sale margin is in fact slightly more responsive. As a result, for every 10 units that leave the STR market, about 3 become long-term rentals and about 7 are sold to owner-occupiers. Tightening the cap from 60 to 30 nights removes about 4,100 units from STR, a quarter of the total. In equilibrium, about 3,100 of these are sold and 1,000 become long-term rentals. Free-sector rents and house prices both fall by about 5%, and exits from the city fall by 22%.

Second, the gains are progressive, but borrowing limits shape them. Households in the bottom income quintile are about twice as sensitive to housing costs as those in the

top quintile and face moving costs exceeding their annual income. Renters in the bottom half of the wealth distribution cannot afford any ownership product at the prevailing LTV and LTI limits. Lower house prices nonetheless bring constrained households into ownership: purchases by households in the two lowest income quintiles rise by 45%, compared with 13% in the two highest. Combining the cap with an LTI limit of 3 leaves its aggregate effects almost unchanged, but cuts the rise in low-income purchases to 28%. The share of the policy's benefit that reaches constrained renters is therefore governed by borrower-based limits.

Third, residents' gains come at a large cost to landlords and visitors, and a tax achieves them more cheaply than a cap. Under the 30-night cap, residents gain €1.6 billion in present value, while landlords lose €3.8 billion and visitors €1.1 billion. Per unit of STR removed, a 10% tax on STR stays delivers similar gains to residents at about a third of the cost. The cap raises nightly STR prices by 17%, so part of the scarcity value of STR goes to the remaining hosts; the tax collects it as revenue instead, and because visitors can switch to hotels, hosts bear nearly 90% of it.

Two features of the setting qualify these results. First, the model does not include the amenity and nuisance effects of STR on neighbors, which Almagro and Domínguez-Iino (2025) show are important in Amsterdam, so our welfare figures leave out a potential benefit of restrictions: the cap raises aggregate welfare only if each STR night removed imposes about €230 or more of external costs on residents. Second, investors' supply is static, so the model gives the long-run reallocation once adjustment is complete; with adjustment frictions, the short-run response through the ownership margin would be smaller.

Related literature. The papers closest to ours are Calder-Wang (2021) and Almagro and Domínguez-Iino (2025), which estimate equilibrium models of residential sorting and STR supply; the latter also studies Amsterdam. Neither has an ownership market: Calder-Wang (2021) models the rental market with a fixed rental stock, and Almagro and Domínguez-Iino (2025) treat homeowners as renting to themselves and take tenure status as given. Restricting STR can therefore only move units between short- and long-term rental, and its benefits for residents run through rents. We add the ownership market, with borrowing-constrained buyers, and the investor's choice between renting and selling. This addition changes the answer quantitatively: of the units that leave STR under a day-cap, about 7 in 10 are sold to owner-occupiers rather than let long-term, so a model without the ownership margin would overstate the long-term rental supply gained by a factor of about 3.7. It also changes who benefits, since most of the gain to residents comes through lower house prices and goes to buyers, and it lets us study STR and borrower-based policies within one framework. Reduced-form evidence on Airbnb and housing costs (Barron et al., 2021; Koster et al., 2021; Garcia-López et al.,

2020), on substitution between Airbnb and hotels (Zervas et al., 2017; Farronato and Fradkin, 2022), and on Airbnb and LTR supply (Horn and Merante, 2017; Li et al., 2022a) motivates our model. Our own identification strategies and estimates for Amsterdam are in line with this work.

We also contribute to quantitative models of housing tenure, in which the joint determination of prices and rents governs how credit conditions and housing policy pass through to households (Sommer et al., 2013; Kaplan et al., 2020; Greenwald and Guren, 2025). Our household block follows the dynamic demand models of neighborhood choice of Bayer et al. (2016), which Almagro and Domínguez-Iino (2025) apply to Amsterdam, and adds tenure choice subject to LTV and LTI limits. A central object in this literature is how segmented the rental and ownership markets are (Greenwald and Guren, 2025; Halket et al., 2020), and investors are the agents who move units across that boundary (Garriga et al., 2023). We introduce STR as a third use of the housing stock, and we estimate rather than calibrate the degree of segmentation, separately by location and size, using quasi-experimental variation and administrative micro-data on household income, wealth and mortgages.

Finally, we relate to work on the unintended effects of rental affordability policies. Rent control in San Francisco led landlords to convert units to condominiums for sale, reducing rental supply (Diamond et al., 2019). The end of rent control in Cambridge, Massachusetts, raised nearby property values (Autor et al., 2014). Favilukis et al. (2023) quantify the trade-off between redistribution and misallocation in New York. We show that the same exit margin that Diamond et al. (2019) document for rent control also operates under STR restrictions, where it shifts the benefits from tenants to buyers. Studies of LTV and LTI limits find that they reduce leverage and house price growth and redirect constrained buyers to cheaper locations (Acharya et al., 2022; Tzur-Ilan, 2023; Peydrò et al., 2024; van Bekkum et al., 2024), and can raise rental demand (Gete and Reher, 2018). We show that these limits also determine who absorbs the units investors sell when STR is restricted, and therefore how the benefits of STR restrictions are distributed among residents.

2. Institutional Background and Data

We study Amsterdam from 2015 to 2019, the window around the city's April 2017 Airbnb day-cap regulation. In this section we present three important institutional features of the Amsterdam housing market that guide our analysis, the details of the 2017 regulation, and the description of our data sources.

2.1. The Dutch Housing Market

Three features of the Dutch housing market shape how Airbnb could affect housing costs and generate the testable predictions we take to the data: a highly inelastic housing supply, three institutionally defined housing tenure categories, and a regulated rental market segmentation.

Inelastic supply. The Netherlands has an estimated long-run housing supply elasticity of 0.40 (Cavalleri et al., 2019), among the lowest in the OECD and less than half the OECD average of 1.01 (Figure J.1). Amsterdam is even more constrained: structural modification of existing buildings is prohibited in the central UNESCO-designated canal district, strict citywide caps limit building height and density, and the city's position on reclaimed land below sea level leaves little scope for outward expansion. Over our sample period, entire-home Airbnb listings rose from roughly 4,000 to a peak of 11,000 per month. Real house prices increased by approximately 60%, substantially outpacing surrounding provinces and the national average (Figure J.2). These trends are consistent with excess demand in a supply-constrained housing market. With the housing stock largely fixed in the short run, reallocating units toward short-term rentals is expected to put upward pressure on house prices and long-term rents.

Tenure composition. The Dutch housing stock consists of three tenure categories: owner-occupied properties, social rentals owned by housing corporations, and private rentals. Housing corporations face legal restrictions that prevent their properties from being converted to short-term rentals. Short-term rental platforms such as Airbnb therefore draw primarily on owner-occupied and privately rented properties. In Amsterdam, both private rentals and Airbnb listings are concentrated in the historic center (Tables 1 and B.1). Given the city's limited housing supply, this overlap suggests that the trade-off between short- and long-term rental use is particularly important in these districts.

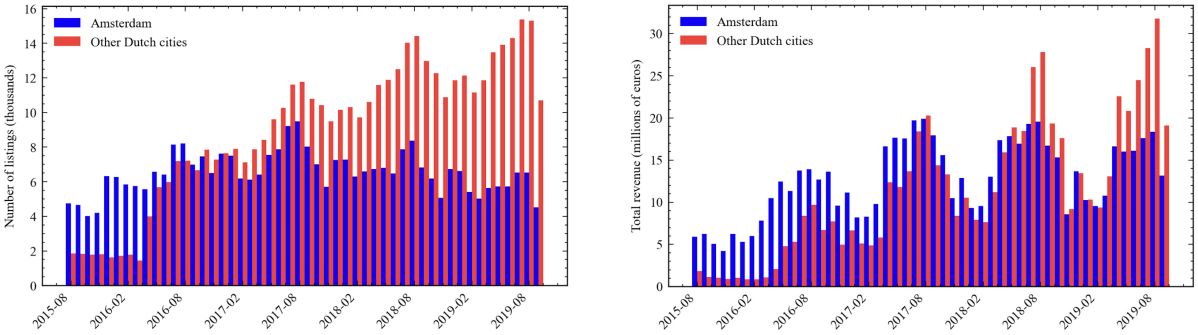
Rental market segmentation. The Dutch rental market is segmented by the *Woningwaarderingsstelsel* (WWS), a national points system. Properties with more than 142 points can be rented at market prices, while those at or below this threshold face statutory rent ceilings.[2] The regulated sector is substantial: housing corporations alone account for 46% of Amsterdam's housing stock (Table B.1).

This distinction provides a useful test of the supply-withdrawal mechanism. If Airbnb reduces the supply of long-term rentals, free-sector rents should rise, while regulated rents should be unaffected because of statutory ceilings. However, this difference alone cannot rule out alternative explanations, such as changes in neighborhood amenities associated with Airbnb (Almagro and Domínguez-Iino, 2025).

[2]The 142-point threshold is the liberalisation cutoff set by the *Besluit huurprijzen woonruimte*; the corresponding maximum regulated rent was €710.68/month in 2016 and is indexed annually.

2.2. Airbnb and the 2017 Regulation

Airbnb entered the Dutch market in 2009, with Amsterdam becoming its largest market by listings and revenue. Before the regulation, entire-home listings averaged 6,908 per month, or 16.8 per 1,000 housing units (Table B.1). Listings were concentrated in the historic center and nearby tourist districts (Figure A.1), providing substantial cross-district variation in Airbnb exposure. Meanwhile, Airbnb activity continued expanding in other Dutch cities, including Rotterdam, The Hague, and Utrecht (Figure 1), providing variation in platform growth outside Amsterdam.



(a) Active listings (thousands) **(b)** Total monthly revenue (millions of euros)

Figure 1. Evolution of Airbnb Activity

Notes: Entire-home listings only. Other Dutch Municipalities includes all listings outside Amsterdam. Sample period August 2015 to September 2019.

Amsterdam introduced a 60-night annual limit on entire-home short-term rentals in 2014, but automated enforcement through Airbnb began only in April 2017. Under an agreement effective January 2017, Airbnb committed to blocking further bookings once a listing reached the annual limit. The system became operational in April 2017, which we use as the effective enforcement date.[3] The municipality also conducted inspections through its Housing Supervision Department (*Toezicht Wonen*), with fines of up to €20,500 for a first violation. Airbnb activity declined sharply following the introduction of automated enforcement (Figure 1).[4] Appendix C discusses other contemporaneous housing policies.

2.3. Data Sources

We combine seven data sources into a unified panel: commercial Airbnb listing data from AirDNA; Statistics Netherlands' (CBS - Centraal Bureau voor de Statistiek) administrative housing registers that include the universe of housing transactions (the BKW transaction register and BAG address register), supplemented by unit-level

[3]The limit was lowered to 30 nights from 1 January 2019, after the end of our reduced-form window (April 2016 to March 2018). We study this tightening as a counterfactual.

[4]Gemeente Amsterdam, *Toeristische verhuur van woonruimte: Rapportage 1 januari 2017 to 1 juli 2017* (2017), https://openresearch.amsterdam/nl/page/61430/rapportage-toeristische-verhuur-van-woonruimte.

property characteristics from the Dutch association of real estate agents (NVM); the CBS *Huurenquête* rental survey with rental prices and property characteristics; household microdata registers from CBS that include households' demographics, location, and financial information (Gbaadresobjectbus, Gbapersoontab, Inhatab, Vehtab, Eigendomtab); district-level data from CBS and the municipality of Amsterdam; and tourism statistics from the Amsterdam Office of Research and Statistics (OIS). To provide further descriptive evidence, we also access mortgage-level data from the De Nederlandsche Bank (DNB), the Dutch Central Bank. The next sections provide a more detailed description of such data.

2.3.1. Airbnb listings

We obtain Airbnb listing data from AirDNA, a commercial provider compiling listing-level records for the Dutch market from August 2015 through September 2019.[5] For each listing-month we observe property type, location coordinates, reservation days, revenue, nightly price, and occupancy rate. We focus on entire-home listings throughout, because they remove a full unit from the long-term housing stock while private and shared rooms monetize spare capacity without withdrawing any unit. Entire-home listings account for approximately 85% of Amsterdam's Airbnb supply during our sample period. Airbnb's share of total short-term rental activity across platforms is a separate quantity: across the six platforms monitored by the Amsterdam municipality, Airbnb accounted for approximately 80% of advertised listings in the first half of 2017 (Gemeente Amsterdam, 2017), with cross-platform duplication from host multi-homing inflating the denominator and increasing Airbnb's unique-listing share.[6]

Panel A of Table B.2 documents the resulting within-city variation at the district-month level, and Panel B reports listing-level characteristics separately for Amsterdam and the other Dutch cities. Figure 1 displays the evolution of Airbnb activity over the sample period for the two groups, with Amsterdam's trajectory diverging after April 2017 from the other Dutch cities not subject to the cap.

2.3.2. Property characteristics and transactions

We combine two property transaction registers: the BKW (*Bestaande Koopwoningen*), a CBS administrative register covering the universe of ownership transactions with sale prices and living area, and the NVM, which adds unit-level characteristics unavailable in BKW (construction year, property type, floor area, number of rooms) for roughly 85%

[5] AirDNA constructs its dataset from platform scrapes and API access; see `https://www.airdna.co` for methodology details. The August 2015 start date reflects the earliest month for which AirDNA covers all four major Dutch metros at the listing level. The dataset is used extensively in academic research on short-term rentals (e.g., Calder-Wang, 2021; Li et al., 2022b).

[6] More details on cross-platform listings are in Appendix D.

of residential sales (Table J.1). For each transaction we observe the sale price, transaction date, and property characteristics.[7] The full sample contains 44,856 transactions across 106 Amsterdam districts over January 2014 to December 2019, yielding 6,654 district-month cells.[8]

Table B.3 documents the systematic differences that motivate district quality adjustment. We distinguish between high- and low-Airbnb-exposure districts, defined respectively as the top and bottom terciles of pre-regulation listings per 1,000 housing units. High-exposure districts transact smaller units, older by 56 years on average, overwhelmingly apartments, and at prices 51% higher in levels and 78% higher per square meter. Airbnb exposure necessarily correlates with these characteristics because the platform concentrated in the compact, historic, tourist-dense center. We therefore use quality-adjusted hedonic residuals as our primary price outcome for the empirical facts in Section 3, constructed via the procedure detailed in Appendix I.2.

2.3.3. Rental Prices

Unlike housing transactions, which are recorded in administrative registries, rents are observed only through surveys. The CBS survey *Huurenquête* provides our primary source: for each sampled dwelling we observe monthly rent, sector status, and detailed property attributes including living area, WOZ value (*Waardering Onroerende Zaken*, the Dutch tax assessment value), construction year, and ownership type.

Table B.4 documents three patterns of the rental market.[9] First, within the free sector, high-exposure districts have an 11% rent premium over low-exposure districts, with slightly larger units, and substantially higher assessed values. Second, regulated rents show the opposite cross-sectional pattern, with lower rents in high-exposure districts, due to the regulated rents' formula that penalizes smaller living area. This shows that regulated rents respond to administrative formulas, not location premia. Third, as in Table B.1, high-exposure districts have a larger share of free-sector rentals, hence in those areas Airbnb will have greater incidence on market-clearing rents.

The survey supports reliable cross-sectional comparisons but cannot directly sustain a district-year panel: stratified sampling yields too few observations per district-year cell across Amsterdam's 110 districts. We therefore impute rental sector status and rent levels for all 412,655 units in the full Amsterdam housing stock via a procedure described in Appendix I.3.

[7] We retain transactions involving housing corporations (*woningcorporaties*) and include a corporation indicator in all specifications, since corporation sale prices reflect regulated or negotiated terms rather than market conditions. Social housing sales account for 21.5% of transactions in peripheral districts (with little Airbnb presence), but only 11.2% in central districts (where Airbnb is most present). Excluding them would differentially thin the sample in peripheral districts and bias treatment effects toward zero.

[8] More details on the transaction registers are in Appendix I.4.

[9] The table draws on the 2015 to 2023 survey (38,382 dwelling-year observations in the high- and low-exposure terciles shown); statistics for the 2016 to 2018 analysis window are nearly identical.

2.3.4. *District characteristics*

District-level data from CBS and the municipality of Amsterdam serve three analytical roles in our descriptive analysis. Demographics, income, and housing stock composition enter as controls absorbing cross-sectional differences correlated with both Airbnb exposure and housing outcomes. Private rental share enters as the split variable for the heterogeneity analysis in Section 3. Tourism infrastructure (hotel beds, cultural facilities, hospitality density, and inverse distance to Dam Square, all measured in 2015) provides the components of the touristiness index used to construct the shift-share instrument in Section 2.4.

We also collect district-level attributes likely relevant to residents, such as transit access, school quality, and availability of green space. Transit access is the average distance by road from residents' homes to the nearest train station measured yearly. School quality is the share of final-year primary pupils in the area's schools who reach the national target level in reading, mathematics and Dutch, averaged over the three subjects, measured yearly except for 2019 when the test was cancelled. Green space is the share of the area's surface covered by public green, from the national topographic register BGT, measured in 2021 and held constant over the sample.

Table 1 reports cross-sectional comparisons between high- and low-exposure districts. High-exposure districts have twice the private rental share of low-exposure districts, while social housing shows the opposite pattern, so the units most able to shift between short-term and long-term use are disproportionately located where Airbnb penetrated deepest. Median household income, however, is statistically indistinguishable across the two groups, and the income quintile shares are equally balanced. This parity is not automatic: Airbnb concentrated in the historic center, where prices are 78% higher per square meter, so income balance across exposure groups directly contradicts a gentrification mechanism in which high-Airbnb districts were already higher-income areas on steeper price trajectories. It does not rule out demand-side confounders operating through channels other than income, but it rules out the simplest affluence-driven explanation for the cross-district price gradient.

2.3.5. *Household panel*

For the structural model we construct an annual panel of every household with an Amsterdam address between 2015 and 2019 from the CBS registers, taking January 1 snapshots to create annual observations.[10] The five-year window captures the 2017 day-cap while the broader institutional environment remained largely stable: social housing eligibility rules were unchanged, and mortgage regulation had no effective bite on the margin the model studies. The statutory LTV cap declined from 103% to

[10] Within-year double-moves (under 1% of households) are considered as single transitions.

Table 1. District Characteristics by Baseline Airbnb Exposure

	High Exposure		Low Exposure			
	Mean	SD	Mean	SD	Diff.	p
Panel A: Socioeconomic						
Median HH income (€)	39,542	11,992	39,015	14,770	527	0.89
HH income share (Q1, %)	30.6	6.2	29.3	10.6	1.4	0.54
HH income share (Q5, %)	19.3	7.6	16.1	11.5	3.2	0.21
Panel B: Housing stock						
Housing units	5,655	2,291	2,946	2,003	2,709	0.00
Single-family (%)	2.5	3.4	30.4	27.8	−27.9	0.00
Multi-family (%)	97.6	3.4	69.6	27.8	27.9	0.00
Social housing (%)	34.4	18.1	52.5	27.6	−18.1	0.00
Owner-occupied (%)	29.0	6.8	29.9	17.1	−0.9	0.77
Private rental (%)	36.6	15.1	17.6	20.9	19.0	0.00
Mean WOZ (€ 000s)	283	94	200	96	83	0.00
Mean residents per unit	1.73	0.16	2.17	0.33	−0.44	0.00
Panel C: Residential mobility (annual)						
Duration of residence (years)	8.1	0.9	8.5	3.2	−0.4	0.52
Arrivals from outside city	837	312	536	502	301	0.00
Arrivals from other Amsterdam area	1,074	378	585	401	490	0.00
Departures from city	765	275	484	339	280	0.00
Departures to other Amsterdam area	1,157	430	541	358	616	0.00
Migration within area	146	84	87	79	58	0.00
Panel D: Demographics						
Population	9,949	4,146	6,632	4,846	3,317	0.00
Households	5,893	2,351	3,225	2,159	2,668	0.00
Age 0 to 17 (%)	13.4	3.2	19.5	5.8	−6.1	0.00
Age 66+ (%)	9.6	2.7	11.9	6.0	−2.3	0.05
Single-person HH (%)	59.0	5.4	48.2	11.3	10.8	0.00
HH with children (%)	18.0	4.9	30.1	10.8	−12.1	0.00
European immigrant (excl. Netherlands, %)	15.1	3.2	10.1	4.9	5.0	0.00
Panel E: Tourism infrastructure						
Hotel beds	1,172	1,882	145	292	1,027	0.00
Cultural facilities	14.0	10.1	1.9	1.8	12.1	0.00
Hospitality establishments	110	96	14	11	96	0.00
Distance to center (km)	2.1	1.0	6.2	2.3	−4.1	0.00

Notes: Pre-regulation period values (2015 to 2016). High/low exposure defined as top/bottom tercile of Airbnb listings per 1,000 housing units; middle tercile omitted. Diff. reports high minus low; p-values from two-sample t-tests with unequal variances (Welch). HH = household. WOZ = assessed property value. Q1 and Q5 refer to bottom and top quintiles of the national income distribution. Distance to center measured from district centroid to Dam Square (raw distance; the instrument uses the inverse).

100% over the period, but because Amsterdam buyers systematically overbid relative to appraised values, average origination LTVs relative to transaction prices hovered around 85% throughout, well below the effective cap used in the model.

For each household-year we observe the residential address (mapped to a housing

product), tenure status, gross income, age of the household head, and household size. From the CBS wealth register we additionally observe household wealth: liquid wealth (financial assets minus non-mortgage debt) for renters, and liquid wealth plus housing equity for homeowners, since sale proceeds are available as a down payment on a subsequent purchase. We discretize income, wealth, age, and household size into 300 household-type cells (10 wealth deciles $\times$ 5 income quintiles $\times$ 3 age groups $\times$ 2 family sizes), with bin boundaries pooled across years.

We exclude social and regulated private renters from the household panel, because allocation in these segments is governed by queuing and administratively capped rents rather than market prices, and their landlords do not participate in the STR market. As reported in Appendix I.3, these segments account for approximately 64% of the stock. This exclusion is therefore not innocuous: social housing acts as an absorbing state, so our sample underrepresents the lowest-income households.

We mitigate this concern in three ways. First, households transitioning to the regulated sector are not dropped, but endogenously choose to exit the private LTR or the ownership markets.Second, binding borrowing constraints ensure the lowest-wealth private-market renters are correctly identified as unable to substitute to ownership, so the distributional findings do not depend on observing social-sector households. Third, regulated-sector rents are capped by the points system and landlords cannot convert regulated units to STR use without first liberalizing the tenancy, so the regulated sector does not participate in the reallocation the model studies. What we do not capture is the indirect effect on social housing waiting times: if regulation reduces private rents and slows transitions from social to private housing, the welfare gain to the lowest-income households is understated.

We also abstract from the following dynamic. Rising Airbnb revenues over the sample may have increased landlords' incentive to liberalize regulated tenancies in order to access the STR market, generating an Airbnb-induced regulated-to-unregulated transition flow that our static segment classification does not separately model. Observing this transition channel would require more disaggregate rental data.

The final sample contains 961,602 household-year observations from 209,921 unique households observed for an average of 4.6 years. Table B.5 reports key statistics, and three patterns motivate our model's structure. First, annual mobility is 14.9%, with owner stay rates far exceeding renter stay rates (91.0% vs. 77.7%), an asymmetry that motivates tenure-specific moving costs and a dynamic stay/move framework. Among movers, 63% leave the private market (transitioning to social housing, other Dutch cities, or abroad), while 37% relocate within Amsterdam's private stock. Second, among within-Amsterdam movers, 47% switch tenure, indicating the rent-versus-own margin is quantitatively active and justifying endogenous tenure choice. Third, movers are systematically younger and less wealthy than stayers, reflecting state dependence:

homeowners who move incur transaction costs proportional to property value, reducing post-move wealth and shrinking future feasible sets.

2.3.6. Tourism statistics

Total visitor flow across all accommodation types averages 8.37 million visitor-nights per month city-wide.[11] Hotels account for the majority of visitor-nights. Airbnb's share peaks at 11.2% in 2016 before declining to 7.4% by 2019, a pattern consistent with the regulation compressing STR capacity. The mean nightly Airbnb price is €161.

2.3.7. Mortgage data

Our mortgage data come from the anonymized loan-level dataset (LLD) collected by De Nederlandsche Bank (DNB), which covers approximately 85% of Dutch mortgages. We use a representative random sample comprising about 10% of this universe, focusing on mortgage originations for home purchases in the four largest Dutch cities.[12] The Amsterdam estimation sample contains 19,605 originations over 2015–2018. The dataset provides information on loan characteristics (lender, origination date, loan amount, maturity, interest rate, and repayment terms), borrower characteristics (income and age), and property characteristics (purchase price and location). Loans, borrowers, and properties have separate anonymized identifiers, but these cannot be directly linked to individual CBS records. Instead, CBS uses observed mortgage characteristics to match records to individuals with similar mortgage profiles. Moreover, property location is recorded at the four-digit postal-code (PC4) level, which does not map one-to-one into the districts used in our housing analysis. We therefore conduct the mortgage analysis separately at the PC4 level. Further details are provided in Appendix G.4.

2.4. Sample Construction and Aggregation

Amsterdam comprises 25 administrative areas (*gebieden*).[13] Excluding uninhabited and purely industrial zones leaves 23 residential areas, indexed by $\mathcal{N} = \{1, \ldots, 23\}$. Within each area $n \in \mathcal{N}$, dwellings are grouped into three size categories: small ($<$50 m^2), medium (50 to 79 m^2), and large ($\geq$80 m^2), so $\mathcal{S} = \{\mathrm{S}, \mathrm{M}, \mathrm{L}\}$.[14] A *housing*

[11] Visitor-night totals and the Airbnb share are from the Amsterdam Office of Research and Statistics annual tourism monitor.

[12] Observations are trimmed at the 1st and 99th percentiles of loan size, LTV, and LTI.

[13] Amsterdam's administrative geography nests as follows: 518 neighborhoods (*buurten*) aggregate into 110 districts (*wijken*), which aggregate into 25 areas (*gebieden*), which in turn fall within 9 boroughs (*stadsdelen*); the ninth, Westpoort, is a port and industrial zone with negligible residential stock and is excluded throughout, leaving 8 boroughs in the analysis sample. Throughout the paper we use the English terms: *neighborhoods* for *buurten*, *districts* for *wijken*, *areas* for *gebieden*, and *boroughs* for *stadsdelen*. We define products at the area level, which balances within-unit price homogeneity against sufficient cell sizes for estimation.

[14] Alternative definitions based on number of rooms yield qualitatively similar results.

product is an area-size pair $j = (n, s) \in \mathcal{J} \equiv \mathcal{N} \times \mathcal{S}$, yielding $|\mathcal{J}| = 23 \times 3 = 69$ products. Time is discrete, with $t = 1, \ldots, T$ indexing years.

Within product-year cells, the median coefficient of variation (standard deviation over mean) is about 0.25 for both ownership prices and rents, and area-size-year fixed effects explain 75.5% and 50.7% of their variation, respectively (Table J.2). We therefore use a single representative price per cell, treating remaining within-cell dispersion as dwelling-level heterogeneity that the product-level model abstracts from. The structural model aggregates to this product-year level for household locations, housing stock counts, prices, rents, and STR activity. The district-month panel is retained in parallel for the empirical facts in Section 3.

2.4.1. Stock allocation

The housing stock registers classify each dwelling's tenure type and ownership category as of January 1. After excluding social and regulated units, for each product j and period t the private dwelling stock H_{jt} is allocated across three uses:

$$O_{jt} + R_{jt} + S_{jt} = H_{jt} \qquad \forall\, j \in \mathcal{J},\ t = 1, \ldots, T, \tag{1}$$

where O_{jt} denotes owner-occupied units identified from the tenure classification, R_{jt} is LTR supply, and S_{jt} is STR capacity, restricted to entire-home listings geocoded to products via district polygon matching and active for at least 30 days during the calendar year (excluding occasional one-off listings that do not represent a sustained withdrawal from residential use). We write $K_{jt} \equiv S_{jt} + R_{jt}$ for total rental units. We treat H_{jt} as exogenous: net additions averaged below 1.5% of the stock annually during 2015 to 2019 (a conservative upper bound from CBS new-construction permits and municipal stock statistics), consistent with the Netherlands' position among the lowest long-run housing supply elasticities in the OECD. At this rate, the fixed-stock assumption is a first-order approximation: even cumulating linearly across the five-year sample, the implied stock change is under 8%, an order of magnitude smaller than the cross-product allocation shifts the model identifies. New-build units that do enter during the sample appear in the data through a higher H_{jt} in subsequent years. We treat these as exogenous additions rather than as reallocations of existing stock, so the model's predictions should be read conditional on the observed path of total dwellings.

2.4.2. Aggregation of prices and visitor data

Ownership prices p^O_{jt} are annual product-level means from the BKW register, averaging approximately 9,000 transactions per year across the 69 products, where the superscript O denotes the owner-occupied use. Rents r^L_{jt} (superscript L for long-term

rental) are imputed as described in Section 2.3.3, and aggregated to the product-year level. For visitor demand, AirDNA listing data are aggregated to the district-month level over the full 2015 to 2019 calendar window (99 residential districts, 60 months, 5,870 district-month cells after excluding those with no recorded listings; Table 5).

Table 2 reports mean allocation shares and the STR revenue premium by borough. Allocation shares exhibit large spatial heterogeneity: in Centrum, STR accounts for 13% of the rental stock versus under 4% in the outermost suburban boroughs (Nieuw-West and Zuidoost), reflecting the concentration of tourist demand.[15] The revenue premium reveals the economic force driving this allocation: in Centrum, STR generates 1.46 times the annual return of LTR, while in several peripheral boroughs the premium falls below parity, meaning STR is less profitable than LTR.[16] This cross-product variation in the revenue differential is the identifying variation for the investor supply parameters.

Table 2. Housing stock allocation and STR revenue premium by borough

	Owner (%)	LTR (%)	STR (%)	Rental (%)	STR/Rental (%)	Revenue premium
Centrum	40.2	52.3	7.5	59.8	13.1	1.46
West	44.1	47.9	8.0	55.9	15.2	1.03
Zuid	41.7	53.8	4.5	58.3	7.7	0.95
Oost	52.9	41.0	6.1	47.1	14.2	0.85
Noord	67.7	26.9	5.4	32.3	19.6	0.83
Nieuw-West	61.7	36.8	1.5	38.3	3.9	0.84
Zuidoost	63.0	36.3	0.7	37.0	2.1	1.74
Weesp	87.4	12.2	0.4	12.6	3.4	1.22

Notes: Figures are 2015 to 2019 mean percentages of liberalized private stock H_{jt}. Rental = LTR + STR; Owner + Rental = 100% by the stock constraint (1). STR/Rental is the mean of product-level STR shares within each borough, not the ratio of borough aggregates; the two differ due to Jensen's inequality. Revenue premium = annualized STR revenue (mean nightly price × occupied nights per year) divided by annual LTR rent. Boroughs ordered by centrality.

2.4.3. District panel, treatment definition, and Bartik instrument

Before presenting our motivating empirical facts, we describe the construction of the underlying district-month panel, the definition of the regulation treatment, and an additional Bartik instrument used for part of the analysis.

District panel. The empirical analysis operates on a district-month panel, with all sources geocoded to CBS district boundaries: Airbnb listings via spatial join of GPS coordinates to CBS boundary polygons, housing transactions via BAG district

[15] Noord's STR/rental share of 19.6% is mechanically inflated by a small denominator: Noord has a very small rental sector overall (32.3% rental share versus 59.8% in Centrum), so even a modest number of STR listings produces a high within-rental share. On an absolute stock basis, Noord's STR share of total dwellings is 5.4%, well below the Centrum figure of 7.5%.

[16] Zuidoost's premium of 1.74 is an outlier driven by a very thin STR market (0.7% of the total stock): a handful of listings with high nightly prices inflate the borough average. At the product level, the premium is estimated with substantial noise in low-STR cells.

identifiers, and rental data via unit-level district identifiers. The Airbnb panel covers 99 districts over 50 months (August 2015 to September 2019), of which 20 months precede the April 2017 regulation and 30 follow it. The housing panel covers 106 districts over 72 months (January 2014 to December 2019), of which 39 precede and 33 follow. DiD specifications use a symmetric 12-month window around the regulation date (April 2016 to March 2018), yielding 99 districts and 2,177 district-month observations.

Treatment definition. The treatment is baseline Airbnb intensity: the average number of active entire-home listings per 1,000 housing units during the pre-regulation period (August 2015 to March 2017), standardized to mean zero and unit standard deviation so that regression coefficients have a per-standard-deviation interpretation. This measure is predetermined: it reflects market-determined location patterns established before the regulation, and averaging over twenty months smooths transitory fluctuations. Figure A.1 displays the distribution. Panel (a) shows pronounced right skew: in raw units the mean is 16.8 listings per 1,000 housing units with standard deviation 15.3, and the 10^{th} to 90^{th} percentile range spans 1.7 to 36.4, a 21-fold range (Table B.2, Panel A). The rightmost histogram bar, at approximately 5.5 standard deviations above the mean, corresponds to the densest canal-district areas with raw intensities exceeding 100 listings per 1,000 units. Panel (b) maps the variation spatially: exposure declines steeply from the historic city center outward, matching the geographic gradient in Table B.1. High- and low-exposure districts, used for the descriptive splits throughout, are the top and bottom terciles of this distribution (Table 1).

Bartik instrument. Our Bartik instrument interacts predetermined tourist infrastructure with Airbnb growth in other Dutch cities. The touristiness index combines four characteristics measured in 2015: hotel bed capacity, cultural facilities, horeca establishments (restaurants, cafés, and bars), and proximity to Dam Square. Its first principal component explains 63.7% of their variation and correlates positively with Airbnb density ($r = 0.53$). Further details are provided in Appendix I.1.

Our preferred measure of platform growth is the number of Airbnb listings in other Dutch cities, indexed to August 2015 and constructed using the same AirDNA data.[17] As an alternative, we use worldwide Google Trends searches for "Airbnb".

The identifying assumption is that, conditional on district and time fixed effects, Airbnb growth outside Amsterdam does not differentially affect housing prices and rents in more tourist-oriented Amsterdam districts except through local Airbnb activity. Measuring tourist infrastructure before the regulation mitigates reverse-causality concerns, although it does not rule out differential housing-demand trends correlated with broader tourism growth.

[17] The sample covers 450,124 listing-months from 30,227 unique listings in other Dutch cities.

3. Empirical Facts

Our empirical facts document how Amsterdam's April 2017 regulation impacted Airbnb listings, house prices, free-sector rents, and mortgages. We find that the day-cap reduced entire-home Airbnb listings by 6–7% per standard deviation of pre-regulation exposure in the year following its introduction. This validates the regulation as an identifying shock on Airbnb activity. The same regulation lowered house prices by 1.1% and reduced mortgage amounts by 1.7%, with no effect on mortgage rates and borrowers' LTV and LTI. Using a Bartik shift-share instrument, we find that a 10% increase in Airbnb density raises house prices by 3.4% and free-sector rents by 0.7%. Regulated rents remain unaffected. In what follows we present the headline estimating equations, the main results in Table 3, and some event-study graphs, while robustness checks, alternative specifications, and most underlying tables are in Appendix G. Last, we also document in Section 3.5 how binding are borrower-based limits across households' income and wealth distribution.

3.1. Effect of Regulation on Airbnb Activity

We begin by examining how the April 2017 day-cap affected Airbnb activity across Amsterdam. Although the regulation applied throughout the city, we expect its impact to be stronger in districts with greater pre-regulation Airbnb activity, where more properties were likely to be affected by the 60-night limit. We exploit this variation using a continuous difference-in-differences design. The main estimating equation is:

$$Y_{jt} = \alpha_j + \gamma_t + \beta_{\text{DiD}} \cdot \text{Exposure}_j \times \text{Post}_t + \varepsilon_{jt}, \tag{2}$$

where Y_{jt} is an Airbnb outcome in district j at month t, α_j and γ_t are district and year-month fixed effects, Exposure_j is standardized baseline Airbnb density (Section 2.4), and Post_t indicates periods from April 2017 onward. The sample is a symmetric 12-month window (April 2016 to March 2018), district-month observations are weighted by mean pre-regulation total listing count, and standard errors are clustered at the district level. The coefficient β_{DiD} is interpreted as the approximate percentage change in outcome per standard deviation of baseline exposure. The corresponding event-study specification replaces Post_t with a full set of period indicators:

$$Y_{jt} = \alpha_j + \gamma_t + \sum_{k \neq -1} \beta_k \cdot \text{Exposure}_j \times \mathbf{1}[t = k] + \varepsilon_{jt}, \tag{3}$$

where k indexes months relative to the regulation (April 2017 $= 0$) and $k = -1$ is the omitted reference period.

The left panel of Figure 2 presents the event-study estimates for log Airbnb listings per 1,000 housing units. Before the regulation, the coefficients fluctuate around zero,

Table 3. Reduced Form Evidence: Regulation DiD and Bartik Instruments

	Airbnb Listings			ln(Rent)	
	ln(List.)	ln(Days)	ln(Price)	Regulated	Free
$\text{Exposure}_j \times \text{Post}_t$	-0.067^{***}	-0.154^{***}	0.019^{**}		
	(0.022)	(0.023)	(0.009)		
ln(Airbnb listings)				-0.018	0.073^{**}
				(0.036)	(0.032)
Kleibergen-Paap F-statistic				15.0	
Segment FE				Yes	Yes
Year FE				Yes	Yes
Year-Month FE	Yes	Yes	Yes		
District FE	Yes	Yes	Yes	Yes	Yes
R^2		0.971			
Observations	2,177	2,177	2,177	586	586

	House	Mortgages			
	Price	ln(Loan)	LTV	LTI	Rate
$\text{Exposure}_j \times \text{Post}_t$	-0.011^{*}	-0.017^{**}	0.587	-0.008	-0.008
	(0.006)	(0.007)	(0.369)	(0.020)	(0.006)
District FE	Yes				
Year-Month FE	Yes				
Area (PC4) FE		Yes	Yes	Yes	Yes
Lender × Month FE		Yes	Yes	Yes	Yes
Product × Lender FE					Yes
R^2	0.854				
Observations	2,177	19,605	19,605	19,605	19,575

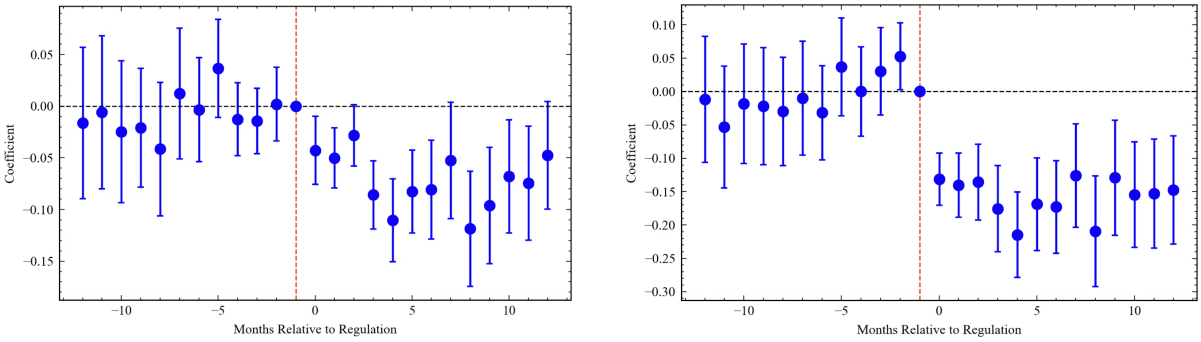
Notes: Unit of observation for Airbnb Listings, Rents, and House Price is district-month, while for Mortgages is an individual mortgage. Sample period is April 2016 to March 2018 (12 months before and after the April 2017 regulation) for Airbnb Listings and House Price, while it is 2016-2018 for Rents and 2015-2018 for Mortgages. Exposure is standardized baseline Airbnb density (listings per 1,000 housing units, pre-regulation mean; standardized to mean zero and unit standard deviation). House price is the mean hedonic residual. Observations for Airbnb Listings are weighted by baseline listing counts. For Rents we show IV estimates from a single regression with segment-specific coefficients. Endogenous variable is log Airbnb listings interacted with segment. Instrument is touristiness × national listing growth in other Dutch cities, interacted with segment. Standard errors in parentheses are clustered at the district level (99 districts) for Airbnb Listings, Rents, House Price and at the PC4 level for Mortgages. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

with no evidence of pre-trends.[18] Following the regulation, the coefficients fall to approximately -13% per standard deviation of baseline exposure.

The intensive margin on the right panel of Figure 2 shows the same pattern with a larger magnitude: total reservation days decline by approximately 14% per standard deviation, and mean reservation days per listing fall by approximately 7.8% (Table G.1), so the regulation reduced both the number of active listings and the intensity of use per continuing host. Nightly Airbnb prices rise by approximately 1.9 to 2.4% in the post-regulation period (Figure A.2), consistent with a contraction in supply rather than

[18] A linear pre-trend test, obtained by regressing the pre-regulation coefficients on a time trend, does not reject a zero slope ($p = 0.41$).

demand.[19] The decline in listings is concentrated in districts with high private-rental share, consistent with regulation displacing private-sector hosts from the short-term market (Appendix G).



(a) Airbnb Listings **(b)** Airbnb Utilization

Figure 2. Event Study: Effect of Regulation on Airbnb Listings and Utilization

Notes: Dependent variable is log listings per 1,000 housing units (left) and log total reservation days (right). Coefficients from equation (3); March 2017 is the omitted reference period and the vertical dashed line marks April 2017, the first post-regulation month. Vertical bars represent 95% confidence intervals with standard errors clustered at district level. Sample is 99 Amsterdam districts over April 2016 to March 2018.

3.2. Effect of Airbnb on House Prices

We estimate Airbnb's effect on house prices using two complementary approaches. First, a difference-in-differences design exploits the April 2017 regulation to estimate its impact across districts with different Airbnb exposure. However, this approach may be affected by spillovers if Airbnb activity shifts from high- to low-exposure districts following the regulation. Second, we use a Bartik shift-share instrument that combines predetermined tourism infrastructure with Airbnb growth in other Dutch cities. This approach addresses potential endogeneity arising from changes in neighborhood attractiveness and provides an alternative source of variation to the regulatory shock. Together, the two strategies offer complementary evidence on Airbnb's effect on house prices.

The DiD specification is identical to equation (2) with Y_{jt} now denoting a quality-adjusted house-price outcome (Section 2.3). We examine three price outcomes: the mean hedonic residual (preferred), the hedonic residual per square meter, and log mean price per square meter, whose inclusion ensures that results do not depend on the LASSO specification.

Figure 3 presents the event-study estimates for quality-adjusted house prices per square meter. Before the regulation, the coefficients show no pre-trend.[20] Following the regulation, the coefficients turn negative, reaching values between -0.04 and -0.08.

[19] Airbnb activity continued expanding in other Dutch cities during this period, suggesting that the contraction in Amsterdam was not driven by a nationwide decline in Airbnb demand.

[20] A linear pre-trend test does not reject a zero slope ($p = 0.43$).

Table 3 reports a decline in house prices of 1.1% per standard deviation of pre-regulation Airbnb exposure. Consistent with this finding, our Bartik IV estimates indicate that a 10% increase in Airbnb density raises house prices by up to 3.4%. Detailed estimates and alternative specifications are reported in Appendix G.2.

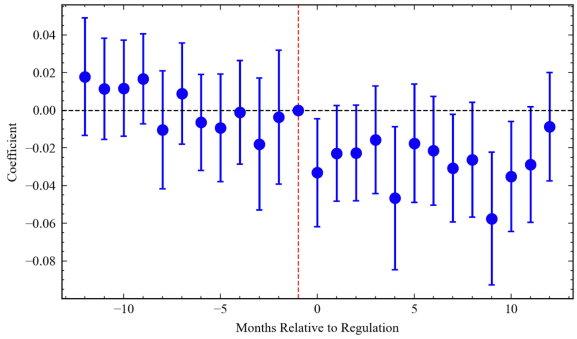


Figure 3. Event Study: Effect of Regulation on House Prices

Notes: Dependent variable is log house prices per square meter (hedonic residual specification). Coefficients from equation (3) with Y_{jt} now denoting the house-price outcome, relative to March 2017 (vertical dashed line marks April 2017). Vertical bars represent 95% confidence intervals with standard errors clustered at district level.

We can compare the economic magnitudes given by two approaches by expressing the DiD results as an elasticity. Dividing the estimated effect on house prices (-0.011) by the corresponding effect on Airbnb listings (-0.065; Table G.1) yields an elasticity of 0.17, compared with 0.31 from the Bartik IV estimation. The difference may reflect the shorter time window of the DiD analysis and the different sources of variation exploited by the two approaches. Nevertheless, the DiD-implied elasticity falls within the confidence interval of the IV estimate, providing reassuring evidence that the two approaches deliver broadly consistent results.[21]

3.3. Effect of Airbnb on Rents

Rents provide a more direct test of the supply-withdrawal mechanism than house prices. If Airbnb reduces long-term rental supply, its presence should raise free-sector rents but have little effect on regulated rents, which are set administratively. However, Airbnb activity may also be correlated with changes in neighborhood attractiveness that independently affect rents. To address this endogeneity concern, we instrument Airbnb density using the Bartik instrument described in Section 2.4.3. We prefer this approach to the difference-in-differences design used for house prices because rental prices are observed annually, leaving us with only three years of data (2016–2018) around the regulation.

We estimate segment-specific effects in a single regression by interacting the endoge-

[21] Airbnb listings per 1,000 units in the top-exposure tercile grew by approximately 40 log points between 2015 and the 2017 peak; at an elasticity of 0.31, the implied price contribution is roughly 12% of the level in those districts, against total real appreciation of around 30 to 35% over the same period.

nous variable and the instrument with segment indicators:

$$\ln r^L_{jst} = \alpha_j + \gamma_t + \mu_s + \beta_s \cdot \ln(\text{Airbnb}_{jt}) + \varepsilon_{jst}, \tag{4}$$

where r^L_{jst} is the long-term rental price in district j, year t, and $s \in \{\text{free}, \text{regulated}\}$ indexes market segments. α_j, γ_t, and μ_s are district, year, and segment fixed effects, and β_s is allowed to differ by segment. $\ln(\text{Airbnb}_{jt}) \times \mathbf{1}[s]$ is instrumented with the touristiness shift-share Bartik instrument described in Section 2.4.3, interacted with the same segment indicator.[22][23] The panel covers 586 district-year-segment observations across 98 districts, three years (2016 to 2018), and two segments.[24] Stacking the two segments in a single regression rather than running separate free-sector and regulated-sector regressions is equivalent to the latter in point estimates but lets us test the cross-segment equality $\beta_{\text{free}} = \beta_{\text{regulated}}$ directly.

Table 3 presents the IV estimates for rents. A 10% increase in Airbnb density raises free-sector rents by approximately 0.7%, while the estimated effect on regulated rents is close to zero and statistically insignificant. This difference is consistent with a reduction in long-term rental supply: free-sector rents respond to market conditions, whereas regulated rents are administratively constrained. The larger effect on house prices than on rents is also consistent with house prices reflecting expected future rental income. Alternative instruments and potential attenuation bias are discussed in Appendix G.3.

3.4. Effect of Airbnb on Mortgage Outcomes

The descriptive analysis so far showed that the April 2017 regulation lowered house prices and raised free-sector rents, with both effects concentrated in the districts with high Airbnb exposure. This section asks whether those effects left a footprint on the mortgage market. If the cap relaxed prices where Airbnb exposure was highest, the loans originated to buy homes in those areas may adjust in size, and the adjustment may show up at the borrowing-constraint margins. We bring the same DiD and shift-instrument empirical designs to loan-level administrative mortgage data.

We defer the details of the empirical analysis to Appendix G.4, and provide here just the overview of the main findings. The DiD results (Table 3) show that a one-standard-deviation increase in Airbnb exposure leads to a 1.7% reduction in loan size, but has no

[22] Equation (4) is shorthand for the fully interacted specification: $\ln(\text{Airbnb}_{jt})$ enters separately for each segment, each instrumented by the corresponding segment-interacted instrument, and the district fixed effects are likewise interacted with segment, yielding district-segment fixed effects that absorb all time-invariant district-specific differences in rent levels between sectors.

[23] We use a single instrument here, rather than the two-instrument set of Section 3.2, because the annual panel spans only three years, so the Google Trends instrument adds little incremental time-series variation beyond the other-Dutch-cities growth instrument.

[24] Standard errors are clustered at the district level. With only three time periods, the effective time-series dimension per cluster is small, but the 98 cross-sectional clusters are well above the threshold at which cluster-robust inference is reliable.

effect on mortgage rates and on borrowers' LTV and LTI. The interpretation of these findings is that the regulation led to a symmetric reduction in house prices and loan sizes, leaving LTV ratio unchanged, and that the reduction in house prices relaxed the affordability constraint for marginal (low-income) households, leaving the LTI ratio unchanged on average. The null effect on mortgage rates shows that lenders did not adjust that margin, and that the primary mechanism was a reduction in mortgage demand. The shift-share instrument analysis delivers very similar results (Table G.11). Last, we find no effect on mortgage prepayment and arrears (Appendix G.4).

Overall, these results justify our choice not to explicitly model mortgage demand, supply, and performance in our structural framework, but rather to just consider borrower-based prudential limits as constraints to the set of properties that households can afford purchasing.

3.5. Borrowing Constraints and Tenure Sorting

Table B.8 shows how LTV and LTI constraints limit households' access to homeownership. Renters in the bottom two wealth quartiles (Q1–Q2) cannot afford any of the 69 housing products defined in Section 2.4. The median renter in Q3 can afford only one product, compared with 57 in Q4. By contrast, even homeowners in the lowest wealth quartile can afford 45 products. These constraints are particularly relevant for renters in central, high-tourism areas, who are concentrated in the bottom wealth quartiles and are more exposed to Airbnb activity. These findings motivate the model's assumption that borrowing constraints determine which households can switch from renting to owning.

Figure A.3 provides additional evidence on residential mobility. Low-wealth households exit Amsterdam's private housing market at roughly twice the rate of high-wealth households, regardless of local Airbnb intensity. This pattern is consistent with homeownership providing greater residential stability, while borrowing-constrained renters have fewer options to adjust to rising housing costs. These differences motivate the model's dynamic tenure and location choices.

Together, these two patterns suggest that the housing market effects of STR regulation cannot be captured by single-margin models, and that borrowing constraints are necessary to quantify ownership demand responses and distributional implications.

3.6. Identification and Why We Need a Model

Identification summary. Our empirical analysis relies on three identifying assumptions. First, the difference-in-differences design requires parallel trends across districts with different Airbnb exposure, consistent with our event-study evidence. Second, the Bartik IV assumes that Airbnb growth in other Dutch cities affects Amsterdam house prices only through local Airbnb activity, conditional on fixed effects. Third, the

rent analysis exploits the fact that regulated rents are set administratively and cannot respond to local housing market conditions.

Several findings support our identification strategies and the supply-withdrawal mechanism. Following the regulation, Airbnb listings declined most in districts with larger private rental sectors (Appendix G), while Airbnb activity continued expanding in other Dutch cities. Nightly Airbnb prices also increased, consistent with a contraction in supply rather than demand. Moreover, household incomes are similar across high- and low-exposure districts (Table 1), and the free-sector rent estimates are robust to alternative instruments (Table G.8).

Why we need a model. The reduced-form evidence cannot tell us how properties are reallocated when STR becomes less profitable. Landlords may convert their units into long-term rentals or sell them to owner-occupiers, with different implications for rental supply and the households that benefit. Moreover, reduced-form estimates cannot quantify the welfare consequences of these reallocations or evaluate alternative policies, such as tighter STR restrictions, borrowing limits and taxes on STR stays. To address these questions, we develop a structural model that jointly determines housing tenure, prices, and household choices. The model distinguishes between the conversion of STR into LTR and the sale of rental properties to owner-occupiers, while incorporating household borrowing constraints to evaluate distributional effects.

4. Model

We model Amsterdam's private housing market, excluding social housing and regulated private rentals, as defined in Section 2. The remaining stock is allocated across the 69 area-size products, with the private stock H_{jt} divided across three uses as in equation (1). Three types of agents interact in this market. *Households* solve a dynamic discrete-choice problem over location-tenure bundles, subject to borrowing constraints. *Short-term visitors* choose between Airbnb listings and a hotel outside option. *Investors* allocate units across the three uses through a two-level nested logit that separates the STR/LTR margin (inner nest) from the rental-versus-ownership margin (outer nest). Prices in all three submarkets (STR, LTR, and ownership) adjust to clear markets simultaneously.

The household demand block builds on the dynamic sorting framework of Bayer et al. (2016), extended to include renters, endogenous tenure choice, and borrowing constraints following Kvaerner et al. (2023). The investor supply side is new: the two-level nested logit that nests STR/LTR allocation within a rental-versus-ownership choice has no direct precedent in the literature. Single-margin STR models (Calder-Wang, 2021) capture only the inner nest; portfolio models following Koijen and Yogo (2019) capture only the outer nest. The nested structure is what generates the portfolio-

reallocation channel.

4.1. Environment

Tenure and choice space. Households choose tenure $h \in \mathcal{H} = \{\text{R}, \text{O}\}$, where R denotes long-term renting and O denotes owner-occupation. A *product-tenure bundle* is a pair $k = (j, h) \in \mathcal{K} = \mathcal{J} \times \mathcal{H}$, yielding $|\mathcal{K}| = 138$ inside alternatives. We augment the choice set with an *outside option* $k = 0$, representing exit from Amsterdam's private housing market.[25] We distinguish the *product choice* (which bundle k to occupy, conditional on moving) from the *mobility choice* (whether to move or stay in the current bundle), where $d_{i,t-1} \in \{0\} \cup \mathcal{K}$ denotes household i's bundle in the previous period and $d_{i,t-1} = 0$ indicates the outside option. The mobility decision introduces moving costs that create state dependence, formalized in the value functions below.

Outside-option composition. The outside option $k = 0$ pools four exits along two dimensions. Two are geographic exits from the city, migrating outside Amsterdam and owner-occupation outside Amsterdam; two remain inside Amsterdam but leave the private market, transition to social housing or to regulated private rental. The model treats these as a single absorbing state because the paper's identification target is supply-side reallocation across the STR, LTR, and ownership margins of the private market, not the decomposition of the exit margin itself. Three pieces of evidence speak to whose exits are bundled here. Figure A.3 shows that low-wealth households exit at roughly twice the rate of high-wealth households throughout the sample period. Table B.5 shows that exits are dominated by renters. Table B.8 shows that the bottom two wealth quartiles cannot afford a single ownership product out of 69. Taken together, the bundled outside option is concentrated among low-wealth, borrowing-constrained renters whose exit margin is plausibly dominated by inside-Amsterdam transitions to social housing or regulated private rental rather than by long-distance migration. Under the current bundling, the model's distributional incidence findings should be read as conservative on the welfare loss to the lowest-income renters, since social-housing transitions are themselves welfare-relevant exits that the bundled outside option treats as a single absorbing state.

Borrowing constraints. Dutch households face a statutory loan-to-value (LTV) cap set by the Ministry of Finance and a debt-service-to-income (DSTI) cap set by Nibud (*Nationaal Instituut voor Budgetvoorlichting*), the independent advisory body whose lending norms are legally binding for mortgage lenders. The LTV cap fell from 106% to 100% in 1-percentage-point annual steps between 2012 and 2018 (Caloia, 2022), so during our 2015 to 2019 sample it ranged from 103% to 100%. The cap applies to the appraised (*taxatie*) value, not the transaction price; in Amsterdam's overheated market, buyers frequently overbid above appraisal, and since mortgage lending is capped

[25] Any subsequent reappearance in the Amsterdam population register is modeled as a new entrant.

at the appraisal, the gap must be covered by a down payment, making the effective constraint tighter in practice than the statutory rate suggests. Average origination LTVs for first-time buyers hovered around 85% during our sample (Eijsink and Van Dijk, 2023). We approximate the Nibud income-dependent DSTI norms with a uniform loan-to-income (LTI) cap of 4 (Caloia, 2022). Together, these two constraints narrow the choice space from $\mathcal{K}$ to a type-specific feasible set $\mathcal{K}_t^\tau \subseteq \mathcal{K}$, with the outside option ($k = 0$) always available.

Household types and state space. Each household i at time t is characterized by a *type* $\tau_{it} := \tau(Z_{it})$, where $\tau(\cdot)$ bins the continuous demographic vector $Z_{it} = (Y_{it}, A_{it}, \text{age}_{it}, \text{fam}_{it})$ comprising gross income Y_{it}, wealth A_{it}, age of the household head age_{it}, and family size fam_{it} (single- versus multi-person). This vector is discretized into 300 type cells (Section 2). The household's vector of state variables is:

$$\Omega_{it} = \left(\tau_{it},\, d_{i,t-1},\, \{r_{jt}^L, p_{jt}^O, X_{jt}, \xi_{jt}\}_{j\in\mathcal{J}}\right),$$

comprising the current type τ_{it}, the lagged choice $d_{i,t-1}$, and market conditions across all products. Here X_{jt} denotes observable product attributes (transit access, school quality, and green space), and ξ_{jt} represents product quality unobserved by the econometrician but observed by all agents.

Timing within period t. At the start of each period, product attributes X_{jt} and unobserved quality ξ_{jt} are realized, and all agents take prices as given and optimize. Households choose product-tenure bundles to maximize expected discounted utility, generating demand for LTR Q_{jt}^L and for ownership Q_{jt}^O. Visitors choose between short-term rentals (Airbnb) and hotels, generating STR demand Q_{jt}^A. Investors determine the rental-versus-ownership split and the STR/LTR allocation within rental. Equilibrium prices $(p_{jt}^A, r_{jt}^L, p_{jt}^O)$ are the values at which all three markets clear subject to the stock constraint (1). All agents are price-takers: households observe the rent and ownership price vectors (r_{jt}^L, p_{jt}^O), visitors observe STR prices p_{jt}^A, and investors observe all three. No agent internalizes the effect of its choices on prices.

4.2. Household Problem

A population of I_t resident households makes housing decisions over discrete time periods $t = 1, \ldots, T$. The population evolves through endogenous exit (choosing the outside option $k = 0$) and exogenous arrival of new entrants.

Feasible choice sets and financial constraints. Each household of type τ faces a feasible set of inside alternatives $\mathcal{K}_t^\tau \subseteq \mathcal{K}$ determined by borrowing constraints that depend on income and wealth. The outside option ($k = 0$) is always available. We treat all rental products as feasible for all household types: $(j, \text{R}) \in \mathcal{K}_t^\tau$ for all $j \in \mathcal{J}$.[26]

[26] In practice, landlords may screen tenants based on rent-to-income ratios (typically requiring rent

Homeownership requires satisfying loan-to-value (LTV) and loan-to-income (LTI) constraints. Let A_τ and Y_τ denote the wealth and income levels associated with type τ. Assuming the household uses all available wealth as a down payment, the required mortgage is at most $L_{\tau t}(j) = \max\{0,\ p^O_{jt} - A_\tau\}$. Ownership of product j is feasible if and only if:

$$\frac{L_{\tau t}(j)}{p^O_{jt}} \leq \overline{\text{LTV}} \quad \text{and} \quad \frac{L_{\tau t}(j)}{Y_\tau} \leq \overline{\text{LTI}}, \tag{5}$$

where $\overline{\text{LTV}}$ and $\overline{\text{LTI}}$ are regulatory caps. We set $\overline{\text{LTV}} = 0.95$ and $\overline{\text{LTI}} = 4$ as round-number approximations to the effective constraints faced by a representative Amsterdam household during our sample window, as described in Section 4.1. Rearranging equations (5) yields a maximum affordable price:

$$\bar{p}^{\text{afford}}_{\tau t} = \min\left(\frac{A_\tau}{1 - \overline{\text{LTV}}},\ A_\tau + \overline{\text{LTI}} \cdot Y_\tau \right),$$

so ownership of product j is feasible if and only if:

$$(j, \text{O}) \in \mathcal{K}^\tau_t \quad \Longleftrightarrow \quad p^O_{jt} \leq \bar{p}^{\text{afford}}_{\tau t}. \tag{6}$$

The $\overline{\text{LTI}}$ cap is not perfectly stable across 2015 to 2019: the Dutch joint-borrower DSTI rule was reformed multiple times during the sample, raising the effective ceiling for two-income households (Thiel and Zaunbrecher, 2023). The implied drift in the maximum mortgage is on the order of 2% over the sample window, with a roughly 3% upper bound on the bias to the household borrowing set; this is small relative to the cross-sectional variation in household income and wealth that drives identification. Treating $\overline{\text{LTI}} = 4$ as constant is therefore a simplification.

Moving costs and type transitions. A move from current choice $d_{i,t-1} = (j_{i,t-1}, h_{i,t-1})$ to a new choice $k = (j, h)$ entails financial and psychological costs. Following the simplifying assumption of Bayer et al. (2016), financial moving costs (FMC) depend only on the property being vacated (origin), not on the destination. Seller-side costs reduce wealth upon departure; buyer-side costs (transfer tax on acquisition) are instead amortized into the user cost of ownership (equation (13)), ensuring they affect the flow utility of the destination rather than the moving decision.[27] As a result, we define financial

below one-third of gross income). We omit this constraint for two reasons. First, unlike LTV/LTI caps, rent-to-income screening is informal and varies across landlords, making it difficult to calibrate uniformly. Second, the flow utility specification penalizes high-rent choices for low-income households through the user cost coefficient γ^τ_{fmc}; at estimated parameter values, $\gamma_{\text{fmc},Y} < 0$ (Section 5.1), so this coefficient is larger for lower-income types and the choice probabilities assign negligible weight to products whose rents are high relative to income, achieving a similar effect to a hard constraint.

[27] This separation is standard in dynamic discrete-choice housing models. It ensures that the move/stay margin identifies the marginal utility of wealth (through seller-side FMC), while the product choice

moving costs as:

$$\text{FMC}_{it} = \begin{cases} 0, & \text{if } h_{i,t-1} = \text{R} \quad \text{(renter moving)}, \\ c^{\text{sell}}\, p^{O}_{j_{i,t-1},t}, & \text{if } h_{i,t-1} = \text{O} \quad \text{(owner moving)}, \end{cases} \tag{7}$$

where $c^{\text{sell}} = 0.02$ is the proportional seller-side transaction-cost rate (realtor commissions and notary fees) and $p^{O}_{j_{i,t-1},t}$ is the current-period price of the origin product.[28]

Post-move wealth accounts for transaction costs:

$$\bar{A}_{it} = A_\tau - \text{FMC}_{it}, \tag{8}$$

and the *post-move type* is $\bar{\tau}_{it} = \tau(\bar{Z}_{it})$ with $\bar{Z}_{it} = (Y_\tau, \bar{A}_{it}, \text{age}_\tau, \text{fam}_\tau)$. Since FMC depends only on origin characteristics, $\bar{\tau}_{it}$ is determined before the household chooses a destination.

Financial moving costs enter utility through the marginal utility of wealth, which we specify as depending on income:

$$\gamma^{\tau}_{\text{fmc}} = \gamma_{\text{fmc},0} + \gamma_{\text{fmc},Y} \cdot \ln Y_\tau, \tag{9}$$

so the direct utility cost of the wealth reduction is $\gamma^{\tau}_{\text{fmc}} \cdot \text{FMC}_{it}$.[29]

Psychological moving costs (PMC) are non-pecuniary and depend linearly on household characteristics:

$$\text{PMC}_\tau = \gamma_{\text{pmc},0} + \gamma_{\text{pmc},Y} \cdot \ln Y_\tau + \gamma_{\text{pmc,age}} \cdot \text{age}_\tau + \gamma_{\text{pmc,fam}} \cdot \text{fam}_\tau, \tag{10}$$

where $\gamma_{\text{pmc},0}$ is a constant and $(\gamma_{\text{pmc},Y}, \gamma_{\text{pmc,age}}, \gamma_{\text{pmc,fam}})$ govern demographic heterogeneity.[30] Entrants to Amsterdam (households with no prior Amsterdam residence) face $\text{FMC}_{it} = 0$ (no origin property). We normalize $\text{PMC}_\tau = 0$ for entrants, since they have no established residence to disrupt; equation (10) applies only to incumbent movers. Thus $\bar{\tau} = \tau$ for entrants.

margin reflects the full annual cost of occupancy (through buyer-side costs embedded in the user cost).

[28] For renters, we assume zero financial moving cost; minor costs such as lease-breaking penalties are negligible relative to ownership transaction costs.

[29] Since income is unchanged by moving, $\gamma^{\tau}_{\text{fmc}} = \gamma^{\bar{\tau}}_{\text{fmc}}$: the marginal utility of wealth is the same before and after the move.

[30] Wealth does not enter the specifications of γ_{fmc} (equation (9)) or PMC (equation (10)) to avoid collinearity with the wealth effect operating through the FMC-induced type transition $\tau \to \bar{\tau}$, since wealth already affects utility through the post-move type $\bar{\tau}$.

Flow utilities. A household of type τ choosing bundle $k = (j, h)$ has flow utility:

$$u_{kt}^{\tau} = \alpha_0 + \alpha_h \cdot \mathbf{1}[h = \mathrm{O}] + \tilde{Z}_{\tau}' \alpha_Z + \alpha_t + \delta_j + (\beta_0 + \tilde{Z}_{\tau}' \beta_Z)' X_{jt} - \gamma_{\mathrm{fmc}}^{\tau} \cdot \mathrm{UC}_{kt} + \xi_{jt}, \tag{11}$$

where j and h denote the product and tenure components of k, and

$$\tilde{Z}_{\tau} = (\ln Y_{\tau},\ \ln A_{\tau},\ \mathrm{age}_{\tau},\ \mathrm{fam}_{\tau})$$

collects type characteristics (income and wealth in EUR thousands, entering in logs).[31] Here α_h is a homeownership premium, α_t and δ_j are year and product fixed effects, X_{jt} are time-varying observed attributes (transit access, school quality, green space), UC_{kt} is the user cost of housing (defined below), and ξ_{jt} is unobserved product quality common to all types. The coefficient on user cost is $\gamma_{\mathrm{fmc}}^{\tau}$ from equation (9), so that the marginal utility of wealth governs both the financial cost of moving and the flow cost of housing.

The user cost of bundle $k = (j, h)$ is:

$$\mathrm{UC}_{kt} = \begin{cases} r_{jt}^{L} & \text{if } h = \mathrm{R} \quad \text{(annual rent)}, \\ p_{jt}^{O} \cdot \mu_{jt}^{O} & \text{if } h = \mathrm{O} \quad \text{(annual ownership cost)}, \end{cases} \tag{12}$$

where r_{jt}^{L} is the annual rental price and μ_{jt}^{O} is the ownership cost rate, defined as:

$$\mu_{jt}^{O} = r_t^{m} + c_t^{\mathrm{prop}} + \delta^{\mathrm{dep}} + \frac{c^{\mathrm{buy}}}{\bar{T}^{\mathrm{hold}}} - \mathbb{E}_t[\pi_{j,t+1}]. \tag{13}$$

The components of equation (13) are as follows. r_t^{m} is the post-tax mortgage interest rate, reflecting the Dutch *hypotheekrenteaftrek* (mortgage interest tax deduction). c_t^{prop} is the property tax rate, approximately 0.3% in Amsterdam, combining the OZB municipal property tax ($\approx 0.04\%$) and the *eigenwoningforfait* imputed rent tax ($\approx 0.25\%$ of assessed value).[32] δ^{dep} is the physical depreciation rate, set to 2% annually following Poterba (1984), reflecting maintenance, wear, and obsolescence of the dwelling structure. $\mathbb{E}_t[\pi_{j,t+1}]$ is the expected capital gain, included following the standard Poterba (1984) user cost formulation; we estimate expected appreciation from product-level price dynamics via the AR(2) process in equation (40) (Section 5). The term $c^{\mathrm{buy}} / \bar{T}^{\mathrm{hold}}$ amortizes buyer-side transaction costs (transfer tax, notary fees; $c^{\mathrm{buy}} = 0.03$) over the expected holding period $\bar{T}^{\mathrm{hold}} = 8$ years, calibrated from the median completed

[31] For the small share of households with zero or negative net financial assets, we set $\ln A_{\tau}$ to the value at the lower bound of the lowest wealth bin; see Section 5.1 for details.

[32] OZB rate from Gemeente Amsterdam's municipal tax schedule (*Belastingtarieven*) for 2015 to 2019; *eigenwoningforfait* rate from Belastingdienst schedule applicable to WOZ values in the sampled range.

ownership spell in our sample.[33] Together with the seller-side rate $c^{\text{sell}} = 0.02$ from equation (7), total round-trip transaction costs are approximately 5% of property value. For the median owner during 2016 to 2019, $\mu_{jt}^{O} \approx 0.04$, reflecting low post-tax Dutch mortgage rates and moderate expected capital gains from historical price trends.

Identifying price sensitivity from moving costs. The coefficient on user cost in equation (11) is $\gamma_{\text{fmc}}^{\tau}$, the same parameter that governs the utility cost of financial moving costs. This is the identification strategy of Bayer et al. (2016): observed mobility patterns identify $\gamma_{\text{fmc}}^{\tau}$ from the move/stay margin, and this parameter then translates the flow cost of housing into utility units. The exclusion restriction is that FMC depends only on the origin property (equation (7)), so it shifts the move/stay decision without affecting relative product attractiveness conditional on moving.

Outside option. The flow utility of the outside option is:

$$u_{0t}^{\tau} = \alpha_{\text{out},0} + \alpha_{\text{out},Y} \cdot \ln Y_{\tau} + \alpha_{\text{out,age}} \cdot \text{age}_{\tau} + \alpha_{\text{out,fam}} \cdot \text{fam}_{\tau} + \alpha_{\text{out},t}, \tag{14}$$

where $(\alpha_{\text{out},Y}, \alpha_{\text{out,age}}, \alpha_{\text{out,fam}})$ capture demographic heterogeneity in outside-option value and $\alpha_{\text{out},t}$ are year fixed effects.

Entry. Each period, a flow of N_t^{entry} new households enters the Amsterdam housing market. These represent young adults forming independent households, inter-city migrants from other Dutch cities, and international arrivals. The model treats all three groups identically, distinguishing them only through their type τ. Entrants choose among all alternatives in $\{0\} \cup \mathcal{K}_t^{\tau}$. For entrants, $k = 0$ represents the alternative of not entering Amsterdam at all, so any implicit cost of entering is captured by the outside-option parameters α_{out} from equation (14). The number and type distribution of entrants $\{N_t^{\tau}\}_{\tau}$ are treated as exogenous to the model. We calibrate them from observed first appearances in the municipal population register.[34]

The household population evolves as $I_{t+1} = I_t - E_t + N_t^{\text{entry}}$, where $E_t = \sum_{i=1}^{I_t} \mathbf{1}[d_{it} = 0]$ is aggregate exit to the outside option. Households that choose the outside option leave the model permanently and do not return. Any subsequent appearance in the Amsterdam population register is treated as a new entry.

Value functions and dynamic programming. Let $\mathbb{V}_{it}(\Omega_{it})$ denote the ex-ante value function before observing the household idiosyncratic taste shock ε_{ikt}, a choice-specific preference residual orthogonal to observed state variables. The household's problem is Markovian: current payoffs and transition probabilities depend on the state Ω_{it} but not

[33] We use completed spells (those that begin and end within 2015 to 2019) rather than the stock-based stay rate of 91% (Table B.5), which includes many spells censored at the sample boundary.

[34] In our sample, annual entry averages approximately 12,400 households (6.5% of the private-market population), with entrants skewing young, lower-wealth, and predominantly renting (Table B.5).

on its history. Define the choice-specific value as:

$$v_{kt}^{\tau} \equiv u_{kt}^{\tau} + \beta \, \bar{V}_{t+1}^{\tau}(k), \tag{15}$$

where $\bar{V}_{t+1}^{\tau}(k) \equiv \mathbb{E}\left[\mathbb{V}_{i,t+1}(\Omega_{i,t+1}) \mid \Omega_{it}, d_{it} = k\right]$ is the expected continuation value and β is the discount factor.[35]

The household faces a two-part decision: whether to move, and where. A stayer receives $v_{d_{i,t-1},t}^{\tau}$, evaluated at the pre-move type since no transaction costs are incurred and wealth is unchanged within the period. Type transitions for stayers due to price appreciation or income changes occur at the period boundary when $\tau_{i,t+1}$ is realized. For owner-stayers, housing equity accrues as $A_{i,t+1} = A_{\tau} + (p^{O}_{j_{i,t-1},t+1} - p^{O}_{j_{i,t-1},t})$, with type updated via $\tau_{i,t+1} = \tau(Z_{i,t+1})$ accordingly; renter-stayer wealth is unchanged within the period. While movers experience type transitions at the moment of moving through FMC reducing wealth (equation (8)), stayers experience them only at the period boundary through capital gains. Income and family size are treated as stable within the five-year estimation window; wealth evolves through capital gains on owned property (for owners) or is unchanged (for renters).

A mover to $k \in \{0\} \cup \mathcal{K}_t^{\bar{\tau}}$ receives $v_{kt}^{\bar{\tau}}$, evaluated at the post-move type $\bar{\tau}$ (reflecting reduced wealth from FMC), minus the psychological moving cost PMC_{τ} (equation (10)). The Bellman equation is:

$$\mathbb{V}_{it}(\Omega_{it}) = \mathbb{E}_{\varepsilon}\left[\max\left\{v_{d_{i,t-1},t}^{\tau} + \varepsilon_{i,\text{stay},t}, \quad \max_{k \in \{0\} \cup \mathcal{K}_t^{\bar{\tau}}} \left\{v_{kt}^{\bar{\tau}} + \varepsilon_{ikt}\right\} - \text{PMC}_{\tau}\right\} \;\middle|\; \Omega_{it}\right], \tag{16}$$

where $\bar{\tau} = \tau(\bar{Z}_{it})$ is the post-move type after FMC reduces wealth (equation (8)) and $(\varepsilon_{i,\text{stay},t}, \varepsilon_{ikt})$ are i.i.d. Type-I extreme value idiosyncratic taste shocks (Rust, 1987), independent across households, alternatives, and periods. For the outside option ($k = 0$), $v_{0t}^{\bar{\tau}}$ includes the flow utility from equation (14) and the continuation value of remaining outside the private market; since departure is an absorbing state, both the flow utility and continuation value are captured by the reduced-form parameters α_{out}.

Conditional on moving, PMC_{τ} is constant across destinations and cancels from the relative comparison, so the mover's destination choice depends only on $v_{kt}^{\bar{\tau}}$. The logit conditional choice probability (CCP) is:

$$\Pr(d_{it} = k \mid \text{move}, \Omega_{it}) = \frac{\exp\left(v_{kt}^{\bar{\tau}}\right)}{\sum_{k' \in \{0\} \cup \mathcal{K}_t^{\bar{\tau}}} \exp\left(v_{k't}^{\bar{\tau}}\right)}, \tag{17}$$

where the outside option ($k = 0$) is included among the mover's alternatives. Entrants

[35] We set $\beta = 0.95$, a standard value in the dynamic discrete choice literature, corresponding to an annual discount rate of approximately 5%.

face the same destination CCP (17) as movers, with FMC = PMC = 0 and hence $\bar{\tau} = \tau$; since they have no lagged choice $d_{i,t-1}$, the stay option does not apply. The unconditional probability that incumbent i occupies bundle k combines the stay probability with the move probability times the conditional destination choice (17). Rather than solving the fixed-point problem (16) directly, we use observed choice probabilities to recover v via the Hotz and Miller (1993) inversion described in Section 5.

Aggregation to demand. Aggregating choice probabilities over incumbent households and new entrants yields product-level demands:

$$Q_{jt}^{L} = \underbrace{\sum_{i=1}^{I_t} \Pr\left(d_{it} = (j,\mathrm{R}) \mid \Omega_{it}\right)}_{\text{incumbents}} + \underbrace{\sum_{\tau} N_t^{\tau} \overset{\tau}{\Pr}\left((j,\mathrm{R}) \mid t\right)}_{\text{entrants}},$$

$$Q_{jt}^{O} = \underbrace{\sum_{i=1}^{I_t} \Pr\left(d_{it} = (j,\mathrm{O}) \mid \Omega_{it}\right)}_{\text{incumbents}} + \underbrace{\sum_{\tau} N_t^{\tau} \overset{\tau}{\Pr}\left((j,\mathrm{O}) \mid t\right)}_{\text{entrants}}.$$

Here Q_{jt}^{L} and Q_{jt}^{O} count the number of households choosing to rent or own product j at time t; the incumbent term includes both stayers and movers.

4.3. Visitor Problem

In each period t, a fixed total of Q_t^V visitor-nights of accommodation demand (from tourists and business travelers) arises in Amsterdam.[36] Each visitor chooses among Airbnb entire-home listings (the relevant substitute for residential housing) across area-size products $j \in \mathcal{J}$ and a composite city-wide hotel outside option with mean utility normalized to zero.

Visitor utility. Visitor i's indirect utility from booking Airbnb product j is:

$$u_{ijt}^{V} = \delta_j^{A} + \beta_A' X_{jt}^{A} - \alpha_A \, p_{jt}^{A} + \xi_{jt}^{A} + \varepsilon_{ijt},$$

where δ_j^{A} is a product fixed effect capturing time-invariant location appeal (centrality, canal proximity), X_{jt}^{A} are time-varying attributes (average review score, seasonal indicators, and event dummies for King's Day, Amsterdam Dance Event, and Pride), p_{jt}^{A} is the per-night Airbnb price, $\alpha_A > 0$ is price sensitivity, ξ_{jt}^{A} is unobserved product quality, and ε_{ijt} is the visitor idiosyncratic taste shock, i.i.d. Type-I extreme value across visitors and products. The utility of the hotel outside option is $u_{it}^{H} = \varepsilon_{i0t}$.[37]

[36] This assumption is conservative for evaluating housing market effects: if regulation deterred some visitors entirely, STR demand would fall further, freeing additional units for long-term use and amplifying the housing-market benefit of regulation. At the model level, visitor flow is defined annually; in estimation, monthly OIS totals Q_m^V are used directly (Section 5.2).

[37] If STR restrictions shift visitors toward hotels and raise hotel prices, the model understates the deterioration of the outside option, so by ignoring hotel price responses it overstates the net benefits

Choice probabilities and demand. Defining mean utility $\bar{\delta}_{jt}^{A} \equiv \delta_{j}^{A} + \beta_{A}' X_{jt}^{A} - \alpha_{A}\, p_{jt}^{A} + \xi_{jt}^{A}$, the logit structure yields the following market shares for Airbnb (A) and hotels (H):

$$s_{jt}^{A} = \frac{\exp(\bar{\delta}_{jt}^{A})}{1 + \sum_{j' \in \mathcal{J}} \exp(\bar{\delta}_{j't}^{A})}, \tag{18}$$

$$s_{t}^{H} = \frac{1}{1 + \sum_{j' \in \mathcal{J}} \exp(\bar{\delta}_{j't}^{A})}. \tag{19}$$

Aggregating over the visitor population yields demand:

$$Q_{jt}^{A} = Q_{t}^{V} \cdot s_{jt}^{A}, \qquad Q_{t}^{H} = Q_{t}^{V} \cdot s_{t}^{H}. \tag{20}$$

We estimate these shares at the district-month level and aggregate them to the product level for the equilibrium system.

Occupancy and market clearing. Not all STR capacity is occupied due to search frictions, seasonal variation, and vacancy between bookings. Let S_{jt} denote the number of units that investors decide to allocate to STR use in product j, and let $\phi_{j}^{A,*} \in (0, 1]$ denote the *natural* occupancy rate absent binding day-caps, estimated from pre-cap booking patterns and held constant across periods.[38] In the model, regulation constrains the intensive margin (nights per listing); the extensive margin (which units are listed) is governed by the investor's inner-nest choice. The realized occupancy rate is:

$$\phi_{jt}^{A} = \min\left\{ \phi_{j}^{A,*},\ \frac{\bar{n}_{t}^{\text{cap}}}{365} \right\}, \tag{21}$$

where $\bar{n}_{t}^{\text{cap}}$ is the citywide maximum nights per year permitted under STR regulations

of day-caps to the housing market. This works in the opposite direction to the exogenous-visitor-flow assumption above, which understates those benefits by ruling out visitor deterrence, so the two outside-option abstractions bound the housing-market benefit from opposite sides. Empirical evidence on STR-hotel substitution supports this direction: Zervas et al. (2017) and Farronato and Fradkin (2022) document that Airbnb expansion exerts measurable downward pressure on hotel revenue and prices, with effects concentrated in periods of peak hotel demand. Symmetric reasoning implies that a binding day-cap would push visitors back toward hotels and raise hotel rates in those same peak periods; we abstract from this channel in the visitor logit and treat hotel utility as the constant outside-option normalization.

[38] We estimate $\phi_{j}^{A,*}$ from average pre-cap (2015 to 2016) booking rates at the product level; see the instrument construction in Section 5.1. The natural occupancy rate is a reduced-form object reflecting both demand-side booking patterns and supply-side listing availability. We treat it as fixed rather than modeling hosts' intensive-margin availability decisions for two reasons. First, in the pre-cap period, hosts faced no regulatory constraint and the modal listing was available year-round; intensive-margin variation was primarily demand-driven. Second, the model's extensive margin (inner nest: whether to list as STR versus LTR) is the first-order supply response to regulation; the day-cap constrains the intensive margin mechanically through $\bar{n}_{t}^{\text{cap}}$, making separate modeling of host availability redundant.

(see Section 2 for Amsterdam's regulatory timeline).[39]

The STR market clears when visitor demand equals occupied capacity:

$$Q_t^V \cdot s_{jt}^A(p_t^A) = 365\,\phi_{jt}^A \cdot S_{jt}. \tag{22}$$

Given ϕ_{jt}^A and S_{jt}, the equilibrium price vector $(p_{1t}^A, \ldots, p_{Jt}^A)$ is the simultaneous solution to the system of J clearing conditions (22) for each period t.

4.4. Investor Problem

We model the equilibrium allocation of the exogenous housing stock H_{jt} across STR, LTR, and owner-occupation using a two-level nested logit. The model represents the supply side through a representative marginal investor. Two features of the Dutch data support this abstraction. First, we observe product-year allocation shares (STR, LTR, and owner-occupied counts) but not the landlord-level sequence of individual reallocation decisions behind them, so the aggregate rental-vs-ownership split is the identified object and a representative-agent model delivers the correct mapping from return differentials to that split. Second, the substantial cross-product variation in the STR revenue premium (Table 2: 1.46 in Centrum versus below parity in peripheral boroughs) provides the identifying variation needed to estimate the nest elasticities, and a representative-investor specification suffices to recover these elasticities from the observed shares.

In practice, the allocation emerges from the decisions of many heterogeneous actors (landlords choosing to sell, owner-occupiers choosing to buy, developers selecting tenure type), but at the market level the equilibrium split between rental and ownership responds to relative returns as if governed by a representative non-owner-occupier landlord at the margin. In the *outer nest*, a unit is either held as a rental investment or allocated to owner-occupation. In the *inner nest*, investors assign each rental unit to either STR or LTR use; the resulting inclusive value (expected per-unit payoff from optimal STR/LTR allocation) enters the *outer nest*, where it is weighed against the opportunity cost of the ownership alternative.

The model treats the total stock H_{jt} as exogenous: no new construction or demolition is determined within the model, so the aggregate number of dwellings in each product is taken from the data and varies only through observed net additions. Within this fixed stock, the logistic share functions allow two-way reallocation between rental and ownership: when rental returns rise, the rental portfolio expands (units move from ownership to rental); when returns fall, it contracts (units move from rental to ownership). Three

[39] With $\bar{n}_t^{\text{cap}} = 60$, the regulatory ceiling is $60/365 \approx 0.16$, well below typical natural occupancy rates, so the cap binds for essentially all products once in effect. The model assumes perfect compliance; to the extent that enforcement is imperfect, estimated policy effects should be interpreted as intended rather than realized impacts.

observations support the exogenous-stock assumption as a reasonable approximation for our sample period. First, net additions averaged below 1.5% of the stock annually during 2015 to 2019 (Section 2.4). Second, the main channels through which new rental supply enters the market (purchases from developers, inheritance, and acquisitions from housing corporations) are plausibly exogenous to the STR regulatory shocks we study. Third, the secular decline in owner-occupation outside Centrum (Table B.6) likely reflects Amsterdam's broader housing boom and new construction rather than investor purchases from owner-occupiers, a trend absorbed by the observed variation in H_{jt} across periods.

The resulting logistic share functions describe the equilibrium split as a smooth function of relative returns.[40] Market clearing reconciles these supply-side allocations with household demand: prices adjust so that the number of households demanding ownership equals the number of units allocated to ownership, and similarly for rental.

The supply side is static, in contrast to the dynamic household demand; in counterfactual exercises, the logistic shares adjust fully to the new payoff differentials, so the comparative statics represent the frictionless steady-state allocation and should be interpreted as an upper bound on the magnitude of portfolio reallocation.

The data observe aggregate allocation shares (STR/LTR/ownership counts by product-year), not sequences of individual investor decisions, so there are no individual-level choice probabilities to invert on the supply side; moreover, dynamic estimation would require specifying investor expectations over future regulatory regimes. The static assumption is most restrictive for the outer nest, where selling a property is costly and partially irreversible; for the inner nest, Dutch STR hosts can adjust their Airbnb availability at high frequency, making the annual allocation well-approximated by a static comparison of rental returns.

4.4.1. ***Inner nest: STR vs. LTR allocation***

Taking as given $K_{jt} \equiv S_{jt} + R_{jt}$ rental units in product j (determined in the outer nest), allocated to STR (S_{jt}) and LTR (R_{jt}), per-unit investor's annual gross payoffs from STR and LTR are:

$$v_{jt}^{S} = 365\, p_{jt}^{A}\, \phi_{jt}^{A} + \xi_{jt}^{S}, \tag{23}$$

$$v_{jt}^{L} = r_{jt}^{L}, \tag{24}$$

[40] The logit shares represent equilibrium allocations, not a literal period-by-period reassignment of every unit; unit-level i.i.d. shocks capture heterogeneity in building suitability, existing lease commitments, sitting owner-occupiers, and other frictions that prevent instantaneous arbitrage.

where p_{jt}^A is the per-night STR price, ϕ_{jt}^A is the realized occupancy rate from equation (21), and r_{jt}^L is the annual LTR rent.[41] Only the payoff difference $v_{jt}^S - v_{jt}^L$ determines allocation; we normalize the LTR unobservable to zero. The STR unobservable ξ_{jt}^S captures the net operating cost differential between STR and LTR (platform commissions, cleaning, furnishing, and higher turnover costs, net of any LTR-specific costs such as tenant screening and void periods); it also absorbs location-specific STR suitability and regulatory risk. STR-specific charges such as Amsterdam's tourist tax are levied on accommodation turnover and largely passed through to the guest, so they bear on the visitor price rather than on the host's operating-cost differential; any host-borne residual enters ξ_{jt}^S.

This treatment parallels Almagro and Domínguez-Iino (2025), who model the STR/LTR difference in Amsterdam through an operating-cost wedge absorbed into location and time fixed effects rather than a separately calibrated cost decomposition; we treat ξ_{jt}^S as the analogous sufficient statistic for the net of these components, since they are not separately observed in the AirDNA listing-level data and a calibrated cost spread would add a free parameter without a separate identifying moment. Because ξ_{jt}^S enters the inner-nest estimation as an instrumented structural unobservable, with district and year fixed effects absorbing its location- and time-specific component, absorbing these cost components into it does not bias the estimated reallocation responsiveness. It also absorbs location-specific STR suitability and regulatory risk.

Converting a unit from LTR to STR is more difficult than the reverse due to Dutch tenant protection laws (*huurbescherming*): indefinite-term rental contracts, which are the norm, can only be terminated by the landlord under narrowly defined legal grounds. The model captures the overall slowness of STR/LTR reallocation through the parameter σ_κ: as σ_κ grows, the allocation becomes less responsive to return differentials, consistent with the difficulty of reallocating units when existing tenancies are in place. Direction-specific frictions (LTR-to-STR conversion being harder than the reverse) are not modeled explicitly, so the estimated σ_κ should be interpreted as an average adjustment speed. The unobservable ξ_{jt}^S absorbs cross-sectional variation in conversion costs, including existing lease commitments and expected vacancy durations. At the annual frequency of the model, the allocation represents a stock equilibrium: the share of units in STR reflects the cumulative outcome of past conversion decisions, not a frictionless annual rebalancing. Private-sector tenancies in Amsterdam typically last several years, so at any point in time, only a fraction of LTR units are near the end of a tenancy and available for conversion.

Each unit $\ell \in \{1, \ldots, K_{jt}\}$ receives i.i.d. Type-I extreme value unit-level STR/LTR suitability shocks $(\varepsilon_{jt\ell}^S, \varepsilon_{jt\ell}^L)$ with scale parameter $\sigma_\kappa > 0$, capturing cross-unit hetero-

[41] LTR revenue assumes full occupancy, reflecting Amsterdam's tight rental market where vacancy rates are negligible. STR revenue is explicitly occupancy-adjusted via ϕ_{jt}^A.

geneity in lease status, conversion cost, and location-specific demand. These unit-level shocks do not appear in the payoff equations (23) and (24), which define mean payoffs common to all units in a product; the idiosyncratic shocks enter only at the unit level and aggregate out to produce the logistic share function below. Aggregating unit-level decisions yields the STR share:

$$s_{jt}^{S} = \Lambda\left(\frac{v_{jt}^{S} - v_{jt}^{L}}{\sigma_{\kappa}}\right), \tag{25}$$

where $\Lambda(\cdot)$ is the logistic function and σ_{κ} governs the elasticity of reallocation. The resulting allocation is $S_{jt} = K_{jt}\, s_{jt}^{S}$ units in STR and $R_{jt} = K_{jt}(1 - s_{jt}^{S})$ in LTR, with $K_{jt} \equiv S_{jt} + R_{jt}$ denoting total rental units as in the stock constraint (1).

The *inclusive value* of the rental portfolio is the expected per-unit payoff from optimal STR/LTR allocation:

$$\mathrm{IV}_{jt} = \sigma_{\kappa} \ln\Big(\exp(v_{jt}^{S}/\sigma_{\kappa}) + \exp(v_{jt}^{L}/\sigma_{\kappa}) \Big). \tag{26}$$

4.4.2. *Outer nest: Rental portfolio size*

The exogenous housing stock H_{jt} is split between K_{jt} rental units and $O_{jt} = H_{jt} - K_{jt}$ units in owner-occupation. The rental inclusive value IV_{jt} is weighed against the annualized opportunity cost of the ownership alternative: by selling a unit at price p_{jt}^{O}, an investor frees capital that earns an annual return ρ, so the annual value of the ownership option is:

$$v_{jt}^{O} = \rho \cdot p_{jt}^{O} \cdot (1 - c^{\text{sell}}), \tag{27}$$

where p_{jt}^{O} is the ownership market price, $\rho > 0$ is the investor's annual discount rate, and $c^{\text{sell}} = 0.02$ is the proportional transaction cost of sale (equation (7)).[42]

The annual rental return includes a product-level unobservable:

$$v_{jt}^{K} = \mathrm{IV}_{jt} + \eta_{jt}, \tag{28}$$

where η_{jt}, with the ownership-side unobservable normalized to zero as in the inner nest, captures unobserved factors favoring rental over sale not reflected in IV_{jt} or v_{jt}^{O}. Each unit receives i.i.d. Type-I extreme value shocks with scale parameter $\sigma_{\mathrm{I}} > 0$. Similarly to the inner nest, $\sigma_{\mathrm{I}} > 0$ captures frictions impacting the rental to ownership transition.

[42] Expected capital gains do not affect the rental-versus-sale margin because the sale price p_{jt}^{O} already capitalizes expected future appreciation, so gains cancel in the $v_{jt}^{K} - v_{jt}^{O}$ difference. In contrast, expected capital gains enter the *household* ownership cost rate (13) because they reduce the user cost of owning relative to renting.

Aggregating yields the rental share:

$$s_{jt}^K = \Lambda\left(\frac{v_{jt}^K - v_{jt}^O}{\sigma_{\mathrm{I}}}\right), \tag{29}$$

with $K_{jt} = H_{jt}\, s_{jt}^K$ units in rental and $O_{jt} = H_{jt}(1 - s_{jt}^K)$ allocated to owner-occupation.

4.5. Equilibrium

The model combines dynamic household demand and static visitor demand with static investor supply.

Definition 1 (Equilibrium)**.** *An* equilibrium *is a sequence of prices* $\{p_{jt}^A, r_{jt}^L, p_{jt}^O\}_{j,t}$*, allocations* $\{S_{jt}, R_{jt}, O_{jt}\}_{j,t}$*, and choice probabilities* $\{\Pr(d_{it} = k \mid \Omega_{it})\}_{i,k,t}$ *such that agents are price-takers and the following conditions hold:*

1. ***Household optimality.*** *Choice probabilities over locations, tenure modes, and exit solve the dynamic program* (16) *and* (17) *subject to borrowing constraints* (5) *and* (6)*, given expectations over future prices and state variables.*
2. ***Visitor optimality.*** *Market shares satisfy the logit demand system* (18)*.*
3. ***Investor optimality.*** *Unit allocations satisfy the nested logit* (25) *and* (29)*.*
4. ***Stock constraint.*** $O_{jt} + R_{jt} + S_{jt} = H_{jt}$ *for all* $j \in \mathcal{J}$ *and* t*.*
5. ***Market clearing.*** *For all* $j \in \mathcal{J}$ *and* t*:*
   (a) *STR:* $Q_t^V \cdot s_{jt}^A(p_t^A) = 365\, \phi_{jt}^A \cdot S_{jt}$*;*
   (b) *LTR:* $Q_{jt}^L = R_{jt}$*;*
   (c) *Ownership:* $Q_{jt}^O = O_{jt}$*.*

The equilibrium concept is a competitive equilibrium with adaptive price expectations: agents are price-takers in each market, three prices clear three markets simultaneously (STR revenue enters the inclusive value, which determines portfolio size and hence both LTR and ownership supply), and households forecast future prices and continuation values using an estimated AR(2) process rather than solving a rational-expectations fixed-point (equations (39) to (41)). Existence follows from the logit and nested logit structures: demand in each market is continuous and strictly decreasing in own price, supply in each market is continuous and weakly increasing in own price, and both are bounded between zero and the stock constraint, so a standard fixed-point argument guarantees at least one equilibrium price vector.[43]

The model operates at the product-year level (69 products $\times$ T years). For estimation, we exploit finer spatial resolution where data support it: visitor demand is estimated at the district-month level (districts nest within areas; 5,870 observations) for statistical

[43] Uniqueness is not established analytically. The cross-market linkages through the inclusive value create feedback loops that could, in principle, generate multiple equilibria. In practice, the logit demand functions are log-concave in prices, which limits the scope for multiplicity.

power, and investor supply at the district-year level. Section 5 discusses the mapping between estimation and model aggregation levels.

How STR regulation propagates through the housing market. The equilibrium definition clarifies how regulation propagates through the market: the nested supply structure generates two countervailing forces when regulation tightens (e.g., a binding day-cap reduces ϕ_{jt}^{A}):

(i) **Within-portfolio reallocation (inner nest).** Lower STR returns shift units from STR to LTR, expanding LTR supply R_{jt} and putting downward pressure on rents.

(ii) **Portfolio reallocation (outer nest).** The decline in the inclusive value IV_{jt} reduces the attractiveness of rental portfolios relative to selling, shrinking K_{jt} and expanding ownership supply O_{jt}. This partially offsets the LTR supply gain from (i): some units that leave STR exit the rental sector entirely rather than converting to LTR.

The net effect on LTR supply, and hence on rents, depends on whether (i) dominates (ii). Conversely, when visitor demand grows and STR prices rise, both forces reverse: units shift from LTR to STR (inner nest) while rising rental returns retain units in the rental portfolio that would otherwise have been allocated to owner-occupation (outer nest). Higher rents and ownership prices induce households to sort toward cheaper products, switch tenure, or exit the city. General equilibrium feedbacks (rents adjust to clear the LTR market, which in turn affects the inclusive value and portfolio size) attenuate partial-equilibrium effects.

5. Estimation and Results

We estimate the model in three steps, exploiting the block-recursive structure: each block conditions on prices taken as observed data, so estimation proceeds without solving for equilibrium prices jointly. First, we estimate household residential demand via the CCP method of Hotz and Miller (1993), identifying price sensitivity from cross-type variation in financial moving costs (proportional to origin house value) following Bayer et al. (2016). Second, we estimate STR demand at the district-month level using a logit demand model, instrumenting endogenous Airbnb prices with a regulatory-bite instrument and characteristics of competing listings in adjacent districts. Third, we estimate inner-nest (STR/LTR) and outer-nest (rental/ownership) parameters via Berry (1994)'s inversion of the nested logit, using a Bartik shift-share instrument for global Airbnb demand shifts and ECB-driven mortgage rate changes for ownership prices.

Three moments drive identification of the structural primitives. Cross-owner variation in sale-price-proportional financial moving costs identifies the utility cost of financial moving costs (γ_{fmc}) from the move/stay margin, which then fixes the price coefficient on housing costs in the flow-utility regression. The 2017 day-cap interacted with pre-cap occupancy identifies both visitors' price sensitivity (α^{A}) and the speed

of STR/LTR reallocation (σ_κ), with a Bartik shift-share from global Airbnb demand providing a second instrument. The ECB mortgage-rate shock interacted with pre-period ownership prices identifies investors' return from selling ρ and the speed of rental/ownership reallocation σ_{I} from the rental-vs-ownership share movement.

5.1. Household Demand

Estimation of household residential demand proceeds in four sub-steps. Household types τ are defined by the 300 demographic cells defined in Section 2. Type-specific parameters are parameterized through demographic interactions rather than estimated separately, reducing the parameter space from 300 type-specific vectors to a single set of demographic interaction coefficients. Borrowing constraints generate type-specific feasible sets, and the resulting cross-type variation in choice sets and sorting patterns identifies the demographic interaction coefficients α_Z and β_Z.

5.1.1. Step 1a: Normalized choice-specific value functions

We recover households' choice-specific value functions $\tilde{v}_{kt}^{\bar{\tau}}$ from observed location and tenure choices of movers ($d_{it} \neq d_{i,t-1}$), including within-Amsterdam relocations, exits to the outside option, and new entrants. Entrants (households appearing in the municipal register for the first time with no prior Amsterdam address) are observationally equivalent to movers: they have post-move type $\bar{\tau} = \tau$ (since FMC $= 0$), face no stay option, and choose among $\{0\} \cup \mathcal{K}_t^\tau$.[44] We pool entrants with incumbent movers for estimation of choice probabilities $\widehat{\Pr}_{kt}^{\bar{\tau}}$. Since FMC depends only on the origin property, the post-move type $\bar{\tau}_{it}$ is determined before destination choice, enabling estimation of $\tilde{v}_{kt}^{\bar{\tau}}$ separately for each post-move type.

A mover of post-move type $\bar{\tau}$ therefore chooses k to maximize:

$$k^* = \arg \max_{k \in \{0\} \cup \mathcal{K}_t^{\bar{\tau}}} \left\{ v_{kt}^{\bar{\tau}} + \varepsilon_{ikt} \right\}, \tag{30}$$

where $v_{kt}^{\bar{\tau}} \equiv u_{kt}^{\bar{\tau}} + \beta \bar{V}_{t+1}^{\bar{\tau}}(k)$ is the choice-specific value from equation (15), $\mathcal{K}_t^{\bar{\tau}}$ is the feasible set of inside alternatives, and ε_{ikt} are i.i.d. Type-I extreme value (T1EV) shocks.[45] Under T1EV errors, conditional choice probabilities (conditional on moving) take the standard logit form $\Pr_{kt}^{\bar{\tau}} = \exp(\tilde{v}_{kt}^{\bar{\tau}}) / \sum_{k'} \exp(\tilde{v}_{k't}^{\bar{\tau}})$, where $\tilde{v}_{kt}^{\bar{\tau}} = v_{kt}^{\bar{\tau}} - m_t^{\bar{\tau}}$ is the normalized value function.

[44] Including entrants improves CCP estimation, particularly for young, low-wealth types where incumbent movers are sparse.

[45] This is the key simplification enabled by the origin-only FMC specification. If FMC depended on the destination (e.g., if buyer transaction costs were included in FMC), then movers would face different post-move types depending on where they moved, and the clean separation between Steps 1a and 1b would break down.

We recover $\tilde{v}_{kt}^{\bar{\tau}}$ from observed choice frequencies. Under T1EV errors, the Hotz and Miller (1993) inversion yields a closed-form solution:

$$\widehat{\tilde{v}}_{kt}^{\bar{\tau}} = \ln \widehat{\Pr}_{kt}^{\bar{\tau}} - \frac{1}{|\{0\} \cup \mathcal{K}_t^{\bar{\tau}}|} \sum_{k' \in \{0\} \cup \mathcal{K}_t^{\bar{\tau}}} \ln \widehat{\Pr}_{k't}^{\bar{\tau}}, \tag{31}$$

where $\widehat{\Pr}_{kt}^{\bar{\tau}}$ is the estimated choice probability obtained via kernel-smoothed weighted averages across nearby types in characteristic space (Appendix H.1). The mean-centering removes the type-time constant $m_t^{\bar{\tau}}$, which cannot be identified from choice probabilities alone (Appendix H.2).

5.1.2. Step 1b: Moving cost parameters

We estimate moving cost parameters $(\gamma_{\text{fmc}}, \gamma_{\text{pmc}})$ from stay/move decisions of incumbent households, treating the normalized value functions $\widehat{\tilde{v}}_{kt}^{\bar{\tau}}$ from Step 1a as known. Entrants are excluded from this step, as they have no prior residence and thus no stay option.[46] The identification strategy follows Bayer et al. (2016): $\gamma_{\text{fmc}}^{\tau}$ is identified from variation in FMC across the move/stay margin (equation (7)), while γ_{pmc} is identified from the move/stay margin conditional on demographics. For example, conditional on the inclusive value of available alternatives, lower mobility rates for older households identify higher psychological moving costs for age.

A household of pre-move type τ in alternative $d_{i,t-1}$ stays if:

$$v_{d_{i,t-1},t}^{\tau} + \varepsilon_{i,\text{stay},t} > \max_{k \in \{0\} \cup \mathcal{K}_t^{\bar{\tau}}} \left[v_{kt}^{\bar{\tau}} + \varepsilon_{ikt} \right] - \text{PMC}_{\tau}, \tag{32}$$

where $v_{d_{i,t-1},t}^{\tau}$ and $v_{kt}^{\bar{\tau}}$ are unnormalized lifetime values evaluated at pre- and post-move wealth respectively, PMC_{τ} is the psychological moving cost from equation (10), and $(\varepsilon_{i,\text{stay},t}, \varepsilon_{ikt})$ are i.i.d. T1EV shocks. For stayers, $\bar{\tau} = \tau$ (no FMC), so $\tilde{v}_{d_{i,t-1},t}^{\tau}$ is directly available from Step 1a.

The normalization constant m_t^{τ} represents the baseline utility level for type τ, which depends on wealth. When wealth changes due to moving costs, this baseline shifts by:

$$m_t^{\tau} - m_t^{\bar{\tau}} = \gamma_{\text{fmc}}^{\tau} \cdot \text{FMC}_{it}, \tag{33}$$

where $\gamma_{\text{fmc}}^{\tau}$ is the marginal utility of wealth from equation (9). For renters, $\text{FMC}_{it} = 0$ so normalization constants coincide; for owners, $\text{FMC}_{it} = c^{\text{sell}} p_{j_{i,t-1},t}^{O}$. Substituting normalized values $\tilde{v} = v - m$ into equation (32) and applying equation (33), the T1EV

[46] Including entrants would bias γ_{pmc} downward, since entrants' "moves" reflect initial sorting rather than the decision to leave an established home.

stay probability becomes:

$$\Pr_{\text{stay},it}^{\tau,\bar{\tau}} = \frac{\exp(\tilde{v}^{\tau}_{d_{i,t-1},t})}{\exp(\tilde{v}^{\tau}_{d_{i,t-1},t}) + \exp\left(\ln \sum_{k\in\{0\}\cup\mathcal{K}_t^{\bar{\tau}}} \exp(\tilde{v}^{\bar{\tau}}_{kt}) - \text{PMC}_{\tau} - \gamma^{\tau}_{\text{fmc}} \cdot \text{FMC}_{it}\right)}. \tag{34}$$

The term $\gamma^{\tau}_{\text{fmc}} \cdot \text{FMC}_{it}$ is the utility cost of wealth reduction from moving: variation in FMC across owners (proportional to house prices via c^{sell}) identifies γ_{fmc} from the stay/move margin.

The likelihood contribution of household i's stay/move sequence is:

$$\mathcal{L}_i^{\text{stay}} = \prod_{t=t_0}^{T_i} \left(\Pr_{\text{stay},it}^{\tau,\bar{\tau}} \right)^{\mathbf{1}[d_{it}=d_{i,t-1}]} \times \left(1 - \Pr_{\text{stay},it}^{\tau,\bar{\tau}} \right)^{\mathbf{1}[d_{it}\neq d_{i,t-1}]}, \tag{35}$$

where t_0 and T_i are the first and last periods household i is observed in the sample.

We maximize the log-likelihood treating $\widehat{\tilde{v}}^{\bar{\tau}}_{kt}$ from Step 1a as fixed.[47] Conditional on the Step 1a value functions, the objective is globally concave in $(\gamma_{\text{fmc}}, \gamma_{\text{pmc}})$, corresponding to a standard binary logit with a unique maximum:

$$(\widehat{\gamma}_{\text{fmc}}, \widehat{\gamma}_{\text{pmc}}) = \arg \max_{\gamma_{\text{fmc}},\gamma_{\text{pmc}}} \sum_{i=1}^{N} \ln \mathcal{L}_i^{\text{stay}}(\widehat{\tilde{v}}, \gamma_{\text{fmc}}, \gamma_{\text{pmc}}). \tag{36}$$

We normalize absolute utility to zero for households with zero wealth, giving $m_t^{\tau} = A^{\tau} \cdot \widehat{\gamma}^{\tau}_{\text{fmc}}$, where A^{τ} is the bin-midpoint wealth for type τ and $\widehat{\gamma}^{\tau}_{\text{fmc}} = \gamma_{\text{fmc},0} + \gamma_{\text{fmc},Y} \cdot \ln Y^{\tau}$ is evaluated at the bin-midpoint income Y^{τ}. Unnormalized value functions are then:

$$\widehat{v}^{\tau}_{kt} = \widehat{\tilde{v}}^{\tau}_{kt} + m_t^{\tau} \quad \text{for all } k \in \{0\} \cup \mathcal{K}_t^{\tau}. \tag{37}$$

5.1.3. Step 1c: Flow utilities

We recover flow utilities u^{τ}_{kt} by subtracting continuation values from the unnormalized value functions $\widehat{v}^{\tau}_{kt}$ constructed in Steps 1a and 1b, using equation (15).

The continuation value from choosing $k = (j, h)$ today is the expected log-sum-exp

[47] The sequential approach provides consistent estimates under standard regularity conditions (Hotz and Miller, 1993). Standard errors from this step alone do not account for sampling uncertainty in $\widehat{\tilde{v}}$ from Step 1a; the bootstrap in Step 1d propagates uncertainty from all earlier stages.

of next-period stay and move values:

$$\bar{V}^{\tau}_{t+1}(k) = \mathbb{E}\Bigg[\ln\Bigg(\exp(v^{\tau'}_{k,t+1}) + \exp\Big(\ln \sum_{k\in\{0\}\cup\mathcal{K}^{\bar{\tau}'}_{t+1}} \exp(\tilde{v}^{\bar{\tau}'}_{kt+1}) - \mathrm{PMC}_{\tau'} - \gamma^{\tau'}_{\mathrm{fmc}} \cdot \mathrm{FMC}(k, t+1)\Big)\Bigg) \Bigg| \Omega_t, d_t = k \Bigg], \tag{38}$$

where primes denote next-period types: τ' is the next-period pre-move type (reflecting updated age and wealth from the period boundary transition) and $\bar{\tau}'$ the corresponding post-move type after FMC.

We model the transition of choice-specific value functions and prices as AR(2) processes with deterministic trends:

$$v^{\tau}_{kt} = \psi^{\tau}_{0,k} + \sum_{l=1}^{2} \psi_{1,l}\, v^{\tau}_{k,t-l} + \sum_{l=1}^{2} \psi_{2,l}\, p^{O}_{j(k),t-l} + \psi^{\tau}_{3,k}\, t + \omega^{\tau}_{kt}, \tag{39}$$

$$p^{O}_{jt} = \phi_{0,j} + \sum_{l=1}^{2} \phi_{1,l}\, p^{O}_{j,t-l} + \phi_{2,j}\, t + \nu_{jt}, \tag{40}$$

$$r^{L}_{jt} = \theta_{0,j} + \sum_{l=1}^{2} \theta_{1,l}\, r^{L}_{j,t-l} + \theta_{2,j}\, t + \zeta^{r}_{jt}, \tag{41}$$

where $j(k)$ denotes the product component of alternative k and $(\omega^{\tau}_{kt}, \nu_{jt}, \zeta^{r}_{jt})$ are one-step forecast innovations (mean-zero unforecastable components of the AR processes for the type-alternative value, ownership price, and rent respectively).[48] The one-step-ahead forecast from equation (40) provides $\mathbb{E}_t[\pi_{j,t+1}]$ for the ownership cost rate (equation (13)). Age increases deterministically and income and family size are treated as stable within the estimation window. Wealth evolves through capital gains on owned property ($A_{i,t+1} = A_{it} + \Delta p^{O}_{j,t+1}$ for owners, constant for renters).[49]

We approximate the expectation in equation (38) via forward simulation. For each (Ω_t, k), we draw $N_{\mathrm{sim}} = 1{,}000$ samples of next-period states by sampling errors from the empirical residual distribution, compute the log-sum-exp for each draw, and average.[50] Flow utilities are recovered by rearranging $v = u + \beta \bar{V}$:

$$\widehat{u}^{\tau}_{kt} = \widehat{v}^{\tau}_{kt} - \beta\, \widehat{\bar{V}}^{\tau}_{t+1}(k) \quad \text{for all } k \in \{0\} \cup \mathcal{K}^{\tau}_{t}. \tag{42}$$

[48] We pool AR coefficients $(\psi_{1,l}, \psi_{2,l})$ across types and products to improve precision, while allowing type-and-product-specific intercepts $\psi^{\tau}_{0,k}$ and trends $\psi^{\tau}_{3,k}$.

[49] This approximation treats the household as holding equity equal to one unit of product j. More precisely, wealth changes by equity share $\times\, \Delta p^{O}_{j,t+1}$, but the simplification is standard in the CCP literature.

[50] We draw errors from the empirical distribution of residuals rather than imposing normality. Results are stable when increasing to $N_{\mathrm{sim}} = 5{,}000$.

5.1.4. Step 1d: Preference parameters

We decompose the recovered flow utilities $\widehat{u}_{kt}^{\tau}$ into observed attributes, user costs, and unobserved quality. We parameterize type heterogeneity as:

$$\beta^{\tau} = \beta_0 + \tilde{Z}_{\tau}' \beta_Z, \tag{43}$$

where $\tilde{Z}_{\tau} = (\ln Y_{\tau}, \ln A_{\tau}, \text{age}_{\tau}, \text{fam}_{\tau})$ is defined in equation (11). The marginal utility of wealth $\gamma_{\text{fmc}}^{\tau}$ depends only on income (equation (9)), with coefficients $(\gamma_{\text{fmc},0}, \gamma_{\text{fmc},Y})$. Since those coefficients represent user cost in equation (11) and are estimated in Step 1b, we move $\widehat{\gamma}_{\text{fmc}}^{\tau}$ to the left-hand side of the flow utility, substituting the heterogeneity parameterization (43):

$$\begin{aligned} \widehat{u}_{kt}^{\tau} + \widehat{\gamma}_{\text{fmc}}^{\tau} \cdot \text{UC}_{kt} = \alpha_0 + \alpha_h \cdot \mathbf{1}[h(k) = \text{O}] + \tilde{Z}_{\tau}' \alpha_Z + \alpha_t + \delta_{j(k)} \\ + \beta_0' X_{j(k),t} + (\tilde{Z}_{\tau}' \beta_Z)' X_{j(k),t} + \xi_{j(k),t}^{\tau}, \end{aligned} \tag{44}$$

where UC_{kt} is defined in equation (12) and remaining terms are as in equation (11).[51] The unit of observation is a (type τ, inside alternative $k \in \mathcal{K}_t^{\tau}$, year t) triple, weighted by the number of type-τ households for whom k was feasible. Outside-option parameters $(\alpha_{\text{out},0}, \alpha_{\text{out},Y}, \alpha_{\text{out,age}}, \alpha_{\text{out,fam}}, \alpha_{\text{out},t})$ from equation (14) are identified separately by regressing $\widehat{u}_{0t}^{\tau}$ on type characteristics.

We estimate equation (44) via OLS: the coefficient on user cost is *not estimated* but fixed at $\widehat{\gamma}_{\text{fmc}}^{\tau}$ from Step 1b, so the price coefficient is identified from moving cost variation rather than from potentially endogenous price variation. Location fixed effects $\delta_{j(k)}$ absorb time-invariant product attractiveness. Baseline amenity coefficients β_0 capture average preferences across all types; demographic interactions β_Z are identified from differential sorting (e.g., families disproportionately choosing neighborhoods with high school quality). The homeownership premium α_h is identified from the share of movers choosing ownership versus rental, conditional on affordability. Year fixed effects $\{\alpha_t\}$ capture aggregate time-series variation in housing demand. Standard errors are bootstrapped (500 replications), resampling households with replacement within each year and re-estimating Steps 1a to 1d on each resampled dataset to propagate uncertainty from earlier stages.

5.1.5. Parameter estimates

Table B.7 reports estimated preference and moving cost parameters. The user cost coefficient $\widehat{\gamma}_{\text{fmc}}^{\tau}$ declines with income, generating a two-fold gap in effective price

[51] The structural unobservable ξ_{jt} (equation (11)) is a product-year primitive common to all types. The residual $\xi_{j(k),t}^{\tau}$ here is the OLS residual from this regression, computed separately for each post-move type. It absorbs ξ_{jt} along with any type-specific mismatch not captured by the demographic parameterization.

sensitivity across the income distribution. Product fixed effects reveal strong preferences for central areas, with Centrum and Oud-Zuid ranking highest. Among time-varying attributes, the average household values transit proximity, school quality, and green space. Demographic interactions are systematic: families place additional weight on school quality and older households value green space more.

The outside option baseline is close to the inside baseline, so the net inside-outside gap is governed by demographics and location-specific amenities rather than the intercepts; higher-income households find exit less attractive, consistent with Amsterdam's wage premium being larger for high earners. The psychological and financial moving cost parameters exhibit opposite income gradients: financial costs deter the poor from moving, while psychological costs deter the rich, reflecting accumulated location-specific capital. Moving costs rise sharply with age.

Table 4 translates the raw coefficients into per-period marginal willingness-to-pay (MWTP) figures across income quintiles, following Bayer et al. (2016) and Kvaerner et al. (2023). MWTP is computed as $\widehat{\alpha}_x^\tau / \widehat{\gamma}_{\text{fmc}}^\tau$, where x indexes any preference coefficient of interest and $\widehat{\gamma}_{\text{fmc}}^\tau$ is evaluated at the quintile-median income. Because price sensitivity declines with income (from 0.149 per €1,000 of annual user cost in the bottom quintile to 0.076 in the top), dividing by $\widehat{\gamma}_{\text{fmc}}^\tau$ generates a sharp WTP gradient across quintiles for any given utility coefficient. The homeownership premium translates to €12,000/year for the lowest quintile and €23,700 for the highest (Panel A). Panel B expresses WTP as a share of income; the declining ratios for homeownership (60% to approximately 30%) and school quality (8.7% to 4.3%) confirm that while higher-income households are willing to pay more in absolute terms, housing amenities absorb a smaller fraction of their budgets. The WTP to avoid relocation exceeds annual income for the bottom four quintiles (Panel C), underscoring the strong inertia in low-income housing decisions.

These distributional patterns have direct implications for evaluating STR regulation. Because low-income renters are roughly twice as price-sensitive as high-income renters, a given rent increase imposes a utility cost on the bottom quintile that is nearly double the cost to the top quintile. Combined with prohibitively high moving costs, these households cannot easily escape rising rents, so even modest rent reductions from STR restrictions produce strongly progressive distributional effects.

5.1.6. Model fit

Figure A.4 summarizes fit for the household demand block. The model matches aggregate mobility by construction: the MLE at the move/stay margin (Step 1b) equates model-predicted and observed annual mobility rates (14.9%). Beyond this targeted moment, the model fits non-targeted moments well. Tenure-specific stay rates are within 2 to 3 percentage points of the data (model: 88.6% owners, 80.7% renters; data:

Table 4. Willingness to pay by income quintile

	Q1	Q2	Q3	Q4	Q5
$\widehat{\gamma}^{\tau}_{\text{fmc}}$ (per €1,000)	0.149	0.128	0.112	0.096	0.076
Panel A: MWTP (€1,000/year)					
Homeownership	12.0	14.0	15.9	18.7	23.7
School quality (1 s.d.)	1.7	2.0	2.3	2.7	3.4
Transit proximity	1.5	1.8	2.0	2.4	3.0
Green space	1.0	1.2	1.3	1.6	2.0
Panel B: MWTP / income (%)					
Homeownership	60.0	46.8	39.8	34.1	29.6
School quality (1 s.d.)	8.7	6.8	5.8	5.0	4.3
Panel C: Psychological moving cost (€1,000/move, age 35)					
PMC	37.2	44.1	50.7	60.3	77.2
PMC / income (%)	186	147	127	110	97

Notes: MWTP $= \widehat{\alpha}^{\tau}_{x}/\widehat{\gamma}^{\tau}_{\text{fmc}}$, where $\widehat{\gamma}^{\tau}_{\text{fmc}} = 0.308 - 0.053 \ln Y_{\tau}$ and Y_{τ} is the quintile-median income in €1,000s. School quality, transit proximity, and green space WTP are for a one-standard-deviation increase. PMC evaluated at age 35 and family size 2, using $\text{PMC}_{\tau} = \widehat{\gamma}_{\text{pmc},0} + \widehat{\gamma}_{\text{pmc},Y} \ln Y_{\tau} + \widehat{\gamma}_{\text{pmc,age}} \cdot 35 + \widehat{\gamma}_{\text{pmc,fam}} \cdot 2$.

91.0%, 77.7%). The owner stay rate gap of 2.4 percentage points implies the model overpredicts owner mobility by roughly 25% in relative terms (11.4% vs. 9.0%), a discrepancy consistent with unmodeled frictions such as search costs or attachment to specific dwellings; the renter stay rate is slightly overpredicted (80.7% vs. 77.7%), a smaller relative error consistent with lower switching costs in the rental market. The mover destination-choice model (Step 1a) achieves a McFadden pseudo-R^2 of 0.335 over approximately 115 feasible alternatives per mover. This indicates substantial explanatory power for a discrete-choice model with a large choice set.

Following Bayer et al. (2016), we evaluate the model's ability to predict individual location choices by ranking each mover's feasible alternatives according to predicted mean utility $\widehat{\bar{v}}^{\tau}_{kt}$ (excluding the T1EV idiosyncratic shock). Among movers, 29% chose an alternative ranked in the top 5% of their predicted ranking, and 45% chose from the top 10%.[52] These figures are comparable to those reported by Bayer et al. (2016) (31% and 47% for the San Francisco Bay Area) and confirm that observed location choices cluster around model-predicted high-utility alternatives.

5.2. Visitor Demand

Since the model defines products at the area-size level, we estimate visitor demand at the finer district-month level (districts nest within areas; 5,870 observations) for statistical power. Berry (1994)'s inversion of the logit share equation (18) yields:

$$\ln s^{A}_{wm} - \ln s^{H}_{m} = \delta^{A}_{w} + \beta'_{A} X^{A}_{wm} - \alpha_{A}\, p^{A}_{wm} + \xi^{A}_{wm}, \tag{45}$$

[52] Under the null of a uniform distribution over all feasible alternatives, only 5% and 10% of movers would fall into these bins by chance.

where $s^A_{wm} = Q^A_{wm}/Q^V_m$ is the observed Airbnb market share of district w in month m, constructed from AirDNA listing-level data aggregated to the district-month level, and $s^H_m = 1 - \sum_{w'} s^A_{w'm}$ (sum over all 99 Amsterdam districts) is the hotel outside-option share, common across districts within a month as implied by the logit structure. The regressors are: district fixed effects δ^A_w absorbing time-invariant location characteristics, time-varying attributes X^A_{wm} (average review score, month-of-year dummies capturing seasonality, and indicator variables for major events such as King's Day, Amsterdam Dance Event, and Pride), per-night Airbnb price p^A_{wm}, and unobserved quality ξ^A_{wm}. The estimating equation (45) is linear in price, so α_A is invariant to the level of spatial aggregation and transfers directly to the model's product level.

Instruments. OLS estimates of α_A are attenuated toward zero because high unobserved quality ξ^A_{wm} raises both market shares and equilibrium prices. We instrument p^A_{wm} using two sources of supply-side variation: a regulation-bite instrument exploiting the 2017 day-cap, and three "BLP instruments" constructed from characteristics of competing listings in adjacent districts (Berry et al., 1995).

We estimate equation (45) via 2SLS. The regulation bite instrument is defined as:

$$Z^{\text{Reg}}_{wm} = \mathbf{1}[\text{post-cap}]_m \times \max\left(\phi^{A,*}_{w,\text{pre}} - \tfrac{60}{365},\ 0\right), \tag{46}$$

measuring the gap between pre-cap natural occupancy $\phi^{A,*}_{w,\text{pre}}$ and the regulatory threshold: zero for districts where the cap does not bind, increasing in the occupancy gap for districts above the threshold. The exclusion restriction is that the day-cap's differential impact across districts is uncorrelated with demand quality shocks conditional on district and month fixed effects; this is plausible because the cap was set uniformly citywide and its bite depends on pre-existing supply-side capacity. The remaining three instruments follow Berry et al. (1995): for each time-varying listing characteristic (review scores, number of photos, and ratings), we compute the sum across competing listings in adjacent districts. These shift competitive pressure and hence equilibrium prices without directly affecting own-product demand quality, provided that spatially correlated demand shocks are absorbed by the district and month fixed effects.[53]

Parameter estimates. Table 5 reports visitor demand estimates from equation (45). Column (1) presents OLS with district and year fixed effects; column (2) instruments Airbnb prices with the regulation-bite and BLP instruments described in this section, maintaining the same fixed effect structure; column (3) replaces year fixed effects and seasonal indicators with year-month fixed effects as a robustness check.

The OLS price coefficient is attenuated by the positive correlation between prices

[53]The BLP exclusion restriction requires that characteristics of competitors in nearby districts are uncorrelated with own unobserved quality ξ^A_{wm} conditional on district and month fixed effects. Spatially correlated demand shocks could violate this condition; the stability of $\widehat{\alpha}_A$ across specifications, which absorb aggregate time variation at different frequencies, provides indirect evidence against this concern.

Table 5. Visitor demand estimates

	(1) OLS	(2) IV	(3) IV
Panel A: Primary regression (equation (45))			
α_A (per € per night)	0.003***	0.035***	0.039***
	(0.001)	(0.005)	(0.003)
$\beta_{A,\text{summer}}$ (June to August)	0.318***	0.437***	—
	(0.017)	(0.035)	
$\beta_{A,\text{King'sDay}}$	0.076***	0.370***	—
	(0.021)	(0.059)	
$\beta_{A,\text{ADE}}$	0.061***	0.293***	—
	(0.019)	(0.059)	
$\beta_{A,\text{Pride}}$	0.129***	0.066	—
	(0.015)	(0.039)	
Panel B: District FE decomposition ($\widehat{\delta}_w^A$ from column 2)			
Distance to centrum (km)		−0.515***	
		(0.061)	
Touristiness index		5.365***	
		(1.296)	
Decomposition R^2		0.745	
District FE	✓	✓	✓
Year FE	✓	✓	—
Year-month FE	—	—	✓
First-stage F	—	48.5	17.0
Clusters (districts)	99	99	99
Observations	5,870	5,870	5,870

Notes: Dependent variable: $\ln s_{wm}^A - \ln s_m^H$. Column (1) is OLS. Columns (2) and (3) instrument p_{wm}^A with the regulation-bite instrument Z_{wm}^{Reg} from equation (46) and three BLP instruments (sums of review scores, number of photos, and ratings across competing listings in adjacent districts). Column (2) includes year fixed effects with seasonal and event dummies estimated freely. Column (3) replaces these with year-month fixed effects (60 dummies), absorbing all aggregate time variation; seasonal and event coefficients are subsumed by month fixed effects and hence not separately identified. Panel B decomposes district fixed effects from column (2) by regressing $\widehat{\delta}_w^A$ on district characteristics via OLS with heteroskedasticity-robust standard errors. Hotel outside option normalized to zero. Estimation sample: 2015 to 2019. Standard errors clustered at the district level (Panel A) or HC1-robust (Panel B). First-stage F-statistics reported; the conventional threshold for 10% maximal TSLS relative bias with four excluded instruments is 24.58. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

and unobserved quality. The IV estimate is an order of magnitude larger, consistent with substantial simultaneity bias in OLS. The year-month fixed effect specification in column (3) yields a similar estimate, confirming that the price coefficient is robust to flexible time controls. The first-stage F-statistic of 48.5 in our preferred specification (column 2) comfortably exceeds the Stock and Yogo (2005) critical value.

At the sample mean Airbnb price ($\bar{p}^A \approx €161$/night), the IV estimate implies an own-price elasticity of approximately -5.6 for a typical district product. The estimate is consistent with high substitutability between Airbnb and hotels for short-stay visi-

tors.[54] Seasonality coefficients in column (2) indicate significantly higher demand in summer months: the Airbnb-to-hotel odds ratio increases by approximately 55% in summer. Event indicators are positive and significant, consistent with demand spikes during major city events. Panel B decomposes the district fixed effects $\widehat{\delta}_w^A$ from the IV specification (column 2) by regressing them on district characteristics in a second-stage OLS regression. Distance to the city center has a strong negative effect, consistent with tourist preferences for central locations near major attractions. The touristiness index (a composite measure of tourist-oriented amenities and landmarks) has a large positive coefficient. Together, these two observables account for three-quarters of the cross-district variation in visitor appeal.

5.3. Investor Supply

Three structural parameters govern investor behavior on the supply side: σ_κ (STR/LTR responsiveness) from the inner nest, and σ_1 and ρ (rental/ownership responsiveness and discount rate) from the outer nest. As with visitor demand, we estimate at the finer district-year level to maximize identifying variation. Both nests exploit Berry (1994)'s inversion and identify return-sensitivity parameters from within-district time-series variation in return differentials, after absorbing time-invariant district heterogeneity and common aggregate shocks through fixed effects.

5.3.1. Inner nest: STR vs. LTR allocation

Substituting equations (23) and (24) into the logit inversion and dividing by σ_κ:

$$\ln s_{wt}^S - \ln s_{wt}^L = \frac{1}{\sigma_\kappa} \Delta v_{wt} + \alpha_t + \mu_w + u_{wt}^{\text{inner}}, \tag{47}$$

where $s_{wt}^L = 1 - s_{wt}^S$ is the LTR share, $\Delta v_{wt} \equiv 365\, p_{wt}^A\, \phi_{wt}^A - r_{wt}^L$ is the observed revenue differential, α_t and μ_w are year and district fixed effects, and u_{wt}^{inner} is the regression residual, reflecting the unobserved STR profitability advantage (scaled by $1/\sigma_\kappa$) not captured by fixed effects.

In the inner nest, $1/\sigma_\kappa$ is identified from within-district variation in the STR/LTR revenue differential Δv_{wt}. Two endogeneity concerns motivate instrumental variables. First, simultaneity: high unobserved STR suitability ξ_{wt}^S raises both the STR share and equilibrium STR revenue. Second, measurement error: Airbnb prices and occupancy are scraped from a platform that reports listed rather than realized transaction data, attenuating OLS estimates toward zero.

[54] Estimates in the hospitality demand literature typically range from -2 to -6 depending on market and segment. For comparison, Farronato and Fradkin (2022) estimate accommodation demand elasticities averaging -4.3 across U.S. cities, ranging from approximately -3 to -6 across segments and cities, with lower-end hotel segments toward the more elastic end.

We instrument with two shift-share variables. The first is a Bartik instrument:

$$Z_{wt}^{\text{Bartik}} = \text{GoogleTrends}_t \times s_{w,2015}^{S}, \tag{48}$$

where the shift is the global Airbnb Google Trends index (time-series variation in worldwide platform demand, determined outside Amsterdam) and the share is pre-period STR intensity. The identifying assumption is that pre-period STR intensity is uncorrelated with differential trends in unobserved allocation preferences.[55]

The second is a regulation-bite instrument:

$$Z_{wt}^{\text{Reg}} = \mathbf{1}[\text{post-cap}]_t \times \max\left(\phi_{w,\text{pre}}^{A,*} - \tfrac{60}{365},\ 0\right),$$

measuring the intensity with which the 2017 day-cap binds: zero for districts whose pre-cap natural occupancy was below the threshold, increasing in the gap for districts above. This instrument also appears in the visitor demand equation (Section 5.2), where it instruments Airbnb prices. The two uses are compatible: the visitor demand equation identifies α_A from the effect of instrumented prices on visitor market shares, while the investor supply equation identifies $1/\sigma_\kappa$ from the effect of the instrumented revenue differential on allocation shares. These are distinct structural equations with different dependent variables, so the same source of exogenous variation can serve as a valid instrument in both. The exclusion restriction requires that the cap affects allocation only through STR revenue, not through direct effects on enforcement intensity or regulatory expectations.[56]

5.3.2. Outer nest: Rental portfolio size

Substituting equations (28) and (27) into the logit inversion and dividing by σ_{I}:

$$\ln s_{wt}^{K} - \ln s_{wt}^{O} = \alpha_1 \widehat{\text{IV}}_{wt} - \alpha_2\, p_{wt}^{O} + \alpha_t + \mu_w + u_{wt}^{\text{outer}}, \tag{49}$$

where $s_{wt}^{O} = 1 - s_{wt}^{K}$ is the ownership share, $\alpha_1 = 1/\sigma_{\text{I}}$, $\alpha_2 = \rho(1 - c^{\text{sell}})/\sigma_{\text{I}}$, α_t and μ_w are year and district fixed effects, and u_{wt}^{outer} is the regression residual, reflecting

[55] The Google Trends index is a valid supply-side instrument here because the investor's allocation decision is downstream of tourist demand in the model's recursive structure. In the visitor demand estimation, this same variable would be invalid because it shifts demand directly. A potential concern is that global Airbnb growth correlates with broader global tourism trends that could affect Amsterdam's housing demand through amenity or congestion channels independent of STR revenue. Year fixed effects absorb any aggregate tourism trend common across districts; the exclusion restriction therefore requires only that the *interaction* of global trends with pre-period STR intensity does not predict district-level changes in investor preferences beyond the STR revenue channel.

[56] If the 2017 regulation increased enforcement intensity differentially across districts, it may directly shift ξ_{wt}^{S} beyond the mechanical occupancy-cap channel, violating the exclusion restriction. The Bartik instrument is not vulnerable to this concern, so we report just-identified estimates from each instrument separately: agreement supports both exclusion restrictions; divergence favors the Bartik-only estimate.

unobserved rental attractiveness (scaled by $1/\sigma_{\mathrm{I}}$) not captured by fixed effects.[57] The generated regressor $\widehat{\mathrm{IV}}_{wt}$ is constructed by exploiting the log-sum structure of the inclusive value:

$$\widehat{\mathrm{IV}}_{wt} = r^{L}_{wt} - \hat{\sigma}_{\kappa} \ln(1 - s^{S}_{wt}), \tag{50}$$

which requires only the LTR rent, the observed STR share, and the $\hat{\sigma}_{\kappa}$ estimate.[58]

In the outer nest, σ_{I} and ρ are identified from variation in the rental inclusive value IV_{wt} and ownership price p^{O}_{wt}. Both regressors are endogenous: unobserved rental attractiveness η_{wt} simultaneously raises the rental share, the inclusive value (through correlated ξ^{S}_{wt}), and ownership prices (by reducing ownership supply). Identification requires two instruments, each targeting one endogenous regressor. We instrument the rental inclusive value with the same Bartik shift-share from the inner nest, as it shifts IV_{wt} through global tourist demand growth, which raises STR revenue and hence the rental inclusive value; the exclusion restriction is that global platform trends do not directly affect whether a unit is held as rental or allocated to owner-occupation. Because $\widehat{\mathrm{IV}}_{wt}$ is constructed from inner-nest estimates, the outer-nest exclusion restriction additionally requires that the inner-nest unobservable ξ^{S}_{wt} is uncorrelated with the outer-nest error u^{outer}_{wt} conditional on district and year fixed effects; the district fixed effects absorb the dominant time-invariant component of any such correlation.

We use a mortgage rate instrument for the ownership price, as it shifts p^{O}_{wt} through domestic borrowing costs, which affect ownership demand and prices; the exclusion restriction is that mortgage rate changes do not directly affect STR profitability. The Bartik instrument identifies $\alpha_1 = 1/\sigma_{\mathrm{I}}$ (how the rental share responds to rental returns) and the mortgage instrument identifies $\alpha_2 = \rho(1 - c^{\text{sell}})/\sigma_{\mathrm{I}}$ (how the rental share responds to ownership prices). Since both coefficients share σ_{I} in the denominator, their ratio directly recovers $\rho(1 - c^{\text{sell}})$, providing an internal specification check: an implausible implied discount rate would signal misspecification of either the exclusion restrictions or the nested logit structure.

We define the mortgage rate instrument as follows:

$$Z^{\text{Mortgage}}_{wt} = \Delta r^{m}_{t} \times p^{O}_{w,2015}, \tag{51}$$

[57] The coefficient on p^{O}_{wt} enters with a negative sign in the estimating equation; Table 6 reports $-\hat{\alpha}_2$ directly, so negative values in the IV columns indicate a positive structural $\hat{\alpha}_2$ (higher ownership prices reduce the rental share, as the model predicts).

[58] From the logit share, $1 - s^{S}_{wt} = \exp(v^{L}_{wt}/\sigma_{\kappa})/[\exp(v^{S}_{wt}/\sigma_{\kappa}) + \exp(v^{L}_{wt}/\sigma_{\kappa})]$, so $\mathrm{IV}_{wt} = v^{L}_{wt} - \sigma_{\kappa} \ln(1 - s^{S}_{wt})$. Because $\widehat{\mathrm{IV}}_{wt}$ uses $\hat{\sigma}_{\kappa}$ from the inner-nest first stage, it is a generated regressor; the reported outer-nest standard errors treat $\hat{\sigma}_{\kappa}$ as known and therefore understate uncertainty relative to a joint estimator that accounts for first-stage sampling variability. Point estimates are unaffected by this sequential approach under standard regularity conditions.

where Δr_t^m is the year-on-year change in the Dutch mortgage interest rate (determined by ECB policy, exogenous to Amsterdam). The interaction with pre-period price generates cross-sectional variation: a 50bp rate increase represents €1,000/year on a €200,000 apartment but €2,500/year on a €500,000 apartment. The exclusion restriction requires that mortgage rate changes affect the rental share only through ownership prices, not through direct effects on investor financing costs. The investor payoff specification (27) implicitly assumes equity financing: v_{jt}^O depends on ρ (the opportunity cost of capital) but not on mortgage rates, so mortgage rate changes enter the model only through their effect on ownership demand and hence p_{jt}^O. To the extent that leveraged investors face correlated borrowing costs, the instrument may also capture a direct rental-supply channel.[59] The direction of this contamination is signable: a higher mortgage rate lowers ownership prices and thereby *raises* the rental share through the modeled user-cost channel, whereas a direct financing-cost effect on leveraged landlords would *lower* it, so the two channels are opposite-signed and a residual financing channel attenuates the mortgage instrument's reduced form.

This bears on the ownership-price coefficient the mortgage instrument identifies (and the discount rate ρ recovered from it), not on the outer-nest scale σ_1 that governs the leakage: σ_1 is identified primarily off the Bartik instrument, whose global-tourist-demand shifter is unrelated to mortgage financing costs, so the leakage estimate is largely insulated from this contamination. The marginal-investor framing bounds this concern: we read the contemporaneous mortgage rate r_t^m as the cost of capital faced by marginal buyers and sellers rather than as the financing position of all incumbent investors, and incumbent investors holding below-market fixed financing from the 2015 to 2019 low-rate environment are mechanically less responsive to rate changes through the financing channel. Appendix H.3 provides robustness checks with additional instruments.[60]

5.3.3. Parameter estimates

Table 6 reports OLS and IV estimates of the investor supply parameters across both allocation stages. In Panel A, the OLS coefficient on the revenue differential is roughly one-quarter of the IV estimates, consistent with measurement error attenuation dominating the upward simultaneity bias. The preferred specification is column (4), which features a first-stage F-statistic above conventional benchmarks and combines both instruments: the two just-identified estimates agree closely, supporting the validity of both exclusion restrictions and justifying the more efficient combined estimator.

[59] In this case the IV estimate would converge to the reduced-form effect, combining both the opportunity-cost and financing-cost channels rather than isolating the former.

[60] The mortgage rate instrument has limited time-series variation, with Δr_t^m taking only 4 to 5 distinct values over 2015 to 2019. The first-stage F-statistics inform whether weak-instrument concerns apply.

The estimate of $\hat{\sigma}_\kappa$ in column (4) implies that, at the median STR share of 8%, a €1,000 increase in the annual STR revenue advantage translates to a 0.6 percentage point increase in the STR share. In Panel B, the preferred specification is column (2), which is just-identified with the cleanest instrument pairing: the Bartik shift-share targets $\widehat{\text{IV}}_{wt}$ through tourist demand, while the mortgage rate instrument targets p^O_{wt} through the ownership market. We prefer the just-identified specification for the outer nest because the additional instruments face more stringent exclusion restriction concerns. Columns (3) and (4) serve as robustness checks, and their close agreement with column (2) supports the preferred specification. The estimate of $\hat{\sigma}_\text{I}$ in Column (2) implies that, at the median rental share of 45%, a €1,000 increase in the rental inclusive value translates to a 2.4 percentage point increase in the rental share. At their respective median shares, the rental-ownership margin (2.4 pp per €1,000) is substantially more elastic than the STR/LTR margin (0.6 pp per €1,000).

These findings are consistent with Amsterdam's specific institutional conditions: on the STR side, an increasingly restrictive licensing regime, the binding day-cap, uncertain enforcement risk, and strong tenant protection laws all impede conversion of LTR units to STR; on the ownership side, Amsterdam's liquid transaction market and well-defined transaction costs during the 2015 to 2019 boom facilitate reallocation at the rental-ownership frontier.

The two scale parameters jointly determine the diversion ratios, that is where housing units go when the relative returns to STR and LTR change. Write the stocks in product j across the three tenure options as $S_{jt} = H_{jt}s^K_{jt}s^S_{jt}$, $R_{jt} = H_{jt}s^K_{jt}(1-s^S_{jt})$ and $O_{jt} = H_{jt}(1-s^K_{jt})$, with H_{jt} fixed, and hold rents and ownership prices constant. For a marginal fall in the STR payoff v^S_{jt}, the shares of units leaving STR that are diverted to long-term rental and to owner-occupation are:

$$DR^L_{jt} = \frac{(1-s^S_{jt})\left[\frac{1}{\sigma_\kappa} - \frac{1-s^K_{jt}}{\sigma_I}\right]}{\Psi_{jt}}, \qquad DR^O_{jt} = \frac{\frac{1-s^K_{jt}}{\sigma_I}}{\Psi_{jt}}, \qquad \Psi_{jt} \equiv \frac{1-s^S_{jt}}{\sigma_\kappa} + \frac{(1-s^K_{jt})\,s^S_{jt}}{\sigma_I}, \quad (52)$$

with $DR^L_{jt} + DR^O_{jt} = 1$. Symmetrically, for a marginal fall in the LTR payoff v^L_{jt}, the shares of units leaving LTR that are diverted to STR and to owner-occupation are

$$\widetilde{DR}^S_{jt} = \frac{s^S_{jt}\left[\frac{1}{\sigma_\kappa} - \frac{1-s^K_{jt}}{\sigma_I}\right]}{\widetilde{\Psi}_{jt}}, \qquad \widetilde{DR}^O_{jt} = \frac{\frac{1-s^K_{jt}}{\sigma_I}}{\widetilde{\Psi}_{jt}}, \qquad \widetilde{\Psi}_{jt} \equiv \frac{s^S_{jt}}{\sigma_\kappa} + \frac{(1-s^K_{jt})(1-s^S_{jt})}{\sigma_I}. \quad (53)$$

These ratios depend only on $\hat{\sigma}_\kappa$, $\hat{\sigma}_I$ and observed shares, and are therefore properties of the estimated supply system rather than of any counterfactual (Appendix F.2).

Two features of these diversion ratios deserve special emphasis. First, the term $1/\sigma_\kappa - (1-s^K_{jt})/\sigma_I$ governs whether the STR/LTR margin dominates the sale margin:

Table 6. Investor supply estimates: Berry inversion

	(1) OLS	(2) IV: Spec. A	(3) IV: Spec. B	(4) IV: Spec. C
Panel A: Inner nest (STR/LTR allocation)				
Δv_{wt} $(1/\hat{\sigma}_\kappa)$	0.020^{*}	0.074^{***}	0.080^{***}	0.076^{***}
	(0.010)	(0.017)	(0.027)	(0.020)
Instruments for Δv	—	Bartik	Regulation	Both
First-stage F	—	19.3	13.9	35.4
$\hat{\sigma}_\kappa$ (€000s)	—	13.5	12.5	13.2
Panel B: Outer nest (Rental vs. ownership)				
$\widehat{\text{IV}}_{wt}$ $(\hat{\alpha}_1)$	-0.0379^{***}	0.0977^{***}	0.0975^{***}	0.1002^{**}
	(0.0143)	(0.0376)	(0.0377)	(0.0396)
p^O_{wt} $(-\hat{\alpha}_2)$	0.0008^{***}	-0.0028^{**}	-0.0024^{**}	-0.0023^{**}
	(0.0002)	(0.0012)	(0.0010)	(0.0009)
Instruments for $\widehat{\text{IV}}$	—	Bartik	Bartik, Reg	Bartik, Reg
Instruments for p^O	—	Mortgage	Mortgage	Mortgage, Demo
Kleibergen and Paap rk Wald F	—	11.4	16.2	18.7
$\hat{\sigma}_{\text{I}}$ (€000s)	—	10.2	10.3	10.0
$\hat{\rho}$	—	0.029	0.025	0.023
Year FE	✓	✓	✓	✓
District FE	✓	✓	✓	✓
Observations (Panel A / B)		469 / 690		

Notes: Both panels estimated at the district-year level. Panel A dependent variable: $\ln s^S_{wt} - \ln s^L_{wt}$; district-year cells with $s^S_{wt} < 10^{-3}$ are excluded, yielding 469 observations. Panel B dependent variable: $\ln s^K_{wt} - \ln s^O_{wt}$; full sample of 690 district-year observations. Column (1) is OLS. Columns (2) to (4) are IV specifications with progressively varied instrument sets. In Panel A, columns (2) and (3) are each just-identified using a single instrument; column (4) is overidentified using both. In Panel B, column (2) is just-identified (Bartik for $\widehat{\text{IV}}_{wt}$, mortgage rate for p^O_{wt}); columns (3) and (4) progressively add instruments for overidentification. Standard errors clustered at the area level (23 clusters). Panel A reports the first-stage F-statistic (single endogenous regressor). Panel B reports the Kleibergen and Paap (2006) rk Wald F-statistic, appropriate for weak-instrument testing with two endogenous regressors and clustered standard errors. All monetary variables in €000s. Structural parameters recovered as $\hat{\sigma}_\kappa = 1/\widehat{(1/\sigma_\kappa)}$, $\hat{\sigma}_{\text{I}} = 1/\hat{\alpha}_1$, and $\hat{\rho} = \hat{\alpha}_2/[\hat{\alpha}_1(1 - c^{\text{sell}})]$ with $c^{\text{sell}} = 0.02$. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Preferred specification: column (4) for Panel A (overidentified, both instruments); column (2) for Panel B (just-identified; see text for discussion).

$DR^L_{jt} > 0$ (and $\widetilde{DR}^S_{jt} > 0$) if and only if $s^K_{jt} > 1 - \sigma_I/\sigma_\kappa \approx 0.23$, which holds for every product except Weesp. Second, although this condition holds, diversion to ownership is large. In the median product, $DR^L = 0.27$ and $DR^O = 0.73$: for every 10 units that leave the STR market, roughly 3 are let long-term and 7 are sold to owner-occupiers. A model with only the STR/LTR margin implicitly sets $DR^L = 1$ and would overstate the long-term rental supply gained from an STR restriction by a factor of about $1/DR^L \approx 3.7$. Conversely, when LTR returns fall, $\widetilde{DR}^O = 0.97$: almost all units leaving long-term rental are sold rather than converted to STR. Because these ratios hold prices fixed, they measure the direct investor response.

The high outer-nest responsiveness may also partly reflect the static specification's inability to capture adjustment frictions in the sell/hold decision. A dynamic model with convex adjustment costs would attenuate the outer-nest reallocation response to

any given policy shock, implying that our portfolio contraction estimates should be interpreted as an upper bound on the leakage channel. As a rough bound: if adjustment costs doubled the effective σ_I to 20.4 (€000s), diversion to long-term rental in the median product would rise from 0.27 to approximately 0.63; sales to owner-occupiers would still absorb more than a third of the units leaving STR.

The data-implied discount rate $\hat{\rho} = 0.029$ (column 2) is below conventional benchmarks of $\rho \approx 0.035$ to 0.055 (Dutch government bond yields of 0.5 to 1.5% plus a 3 to 4% real estate risk premium). However, $\hat{\alpha}_2$ is small, so $\hat{\rho}$ is imprecisely estimated. A delta-method 95% confidence interval for $\hat{\rho}$ spans approximately $[0.01, 0.05]$, comfortably including conventional benchmarks of 4 to 5%.

6. Counterfactual Analysis

We use the estimated model to study three policies, all relative to the 2018 baseline, which reproduces the observed allocation of the 69 products. In the first, the city lowers the day cap from 60 to 30 nights, as it did in January 2019. In the second, the 30-night cap is combined with a reduction in the loan-to-income limit from 4 to 3, which tightens the borrowing constraints of households. In the third, instead of the cap's tightening, the city imposes a 10% tax on STR stays, so that visitors pay the gross price and hosts receive the net price. In each case we hold the product-level unobservables ξ^S_{jt} and η_{jt} at their estimated values and solve for the STR prices, free-sector rents and ownership prices that clear the three markets. Table 7 reports the results, and Figure 4 summarizes them.

6.1. Tightening the day cap

Lowering the cap to 30 nights cuts STR nights by 62% and the number of STR units by 26%, from 15,575 to about 11,500 (Panel B). The remaining hosts are constrained by the cap: occupancy falls by about half, and the nightly STR price rises by 17% (Panel A). Of the roughly 4,100 units that leave STR, about 3,100 are sold to owner-occupiers and 1,000 become long-term rentals. Panel C decomposes this change. Within the rental portfolio, the share let long term rises by 4.5% (the inner nest), but the portfolio itself shrinks by 3.2% as investors sell (the outer nest). The net gain in LTR units is therefore only 1.2%, while owner-occupied units rise by 2.4%. The diversion ratios of Section 5.3.3 imply that about 7 in 10 units would be sold at fixed prices; in equilibrium the share is about three in four.

Both markets for residents become cheaper. Free-sector rents fall by 5.3% and ownership prices by 5.1%. Fewer households leave the city, as exits fall by 22%, and purchases by movers and entrants rise by 25% (Panel D). The increase is concentrated among lower-income households: purchases by the two lowest income quintiles rise

by 45%, compared with 13% for the two highest. Lower prices relax the LTV and LTI constraints that keep these households in the rental market.

6.2. The cap and borrowing limits

A loan-to-income limit of 3 leaves the aggregate effects of the cap almost unchanged. STR units still fall by 26% and rents by 5.5%. Ownership prices, however, fall by more (6.9%), because fewer households can borrow to buy. The tighter limit mainly changes who buys: the increase in purchases by the two lowest income quintiles falls from 45% to 28%, and fewer incumbents move. The more binding the borrowing constraints, the smaller the share of the benefit of the cap that reaches constrained renters.

6.3. A tax on STR stays

Taxing STR stays is a policy Amsterdam already uses. From January 2020, the city raised the tourist tax on holiday rentals from 7% to 10% of the overnight price, excluding VAT, while hotels kept the 7% rate and paid a fixed amount of €3 per person per night instead.[61] We study a 10% tax on the nightly STR price, levied in addition to the tourist tax in place in 2018. Hosts receive the net price and visitors pay the gross price. We find that the tax lowers the price received by hosts by 8.1%, while the price paid by visitors rises by only 1.1%. Hosts therefore bear nearly 90% of the tax, because visitors can switch to hotels. The tax removes fewer units than the cap (6.1%), so its effects on rents, prices and purchases are smaller in levels. Per unit removed from STR, however, its effects on residents are similar. Unlike the cap, the tax does not ration the remaining listings: occupancy is unchanged, and the scarcity value that the cap passes on to the remaining hosts through higher nightly prices becomes revenue for the city, about €17 million a year.

6.4. Welfare

Panel E reports the present value of welfare changes. Under the 30-night cap, households gain €1.6 billion: €1.2 billion for owner-occupiers, €0.35 billion for renters and €0.07 billion for entrants. Visitors lose €1.1 billion and investors €3.8 billion, so the total change is a loss of €3.3 billion.

The split between owner-occupiers and investors requires some clarifications. Investor welfare is computed over the whole housing stock, including owner-occupied units, so the fall in the value of owner-occupied homes is recorded in the investor row, while the owner-occupier row reflects only the lower cost of housing. This capital loss is about €3.0 billion under the 30-night cap (a 5.1% fall in the price of 129,302

[61] Gemeente Amsterdam, Verordening Toeristenbelasting 2020, adopted by the city council on 7 November 2019. From January 2024 the tax is a single rate of 12.5% of the overnight price for all accommodation types (Verordening Toeristenbelasting 2024).

owner-occupied units). It is not, however, a welfare loss of the same size for owners. An owner who stays in Amsterdam, or moves within the city, also faces lower prices for the housing she will consume in the future, which largely offsets the lower value of her home (Bajari et al., 2005; Sinai and Souleles, 2005). The capital loss is a welfare loss mainly for owners who leave the city, and for highly leveraged owners, whose equity falls by more than the price. The welfare change of owner-occupiers therefore lies between the €1.2 billion gain in Panel E and a loss of about €1.8 billion if the full capital loss were borne by them. The total is not affected by this attribution.

Per household, the gains are similar in euros across income groups (Panel F). Renters in the lowest income quintile gain €4,000 and those in the highest €5,100 in present value, so relative to income the gains are much larger at the bottom of the distribution. With a loan-to-income limit of 3, renters' gains are unchanged, while owner-occupiers gain more and investors lose more because prices fall further. Because the capital loss on owner-occupied homes also rises, to about €4.0 billion, the larger gain of owner-occupiers under this scenario should be read with the same caution.

The tax yields a much smaller net loss of €0.26 billion, after including €0.34 billion of tax revenue. Per unit removed from STR, households gain about the same under the two policies (€0.39 million under the cap and €0.36 million under the tax), but the aggregate loss is about a third as large under the tax (€0.27 million against €0.79 million).

These totals exclude the amenity and nuisance effects of STR on neighbors, which our model does not include. They can therefore be read as a threshold: the 30-night cap removes about 0.7 million nights a year, so it raises aggregate welfare only if each night removed imposes at least about €230 of external costs on residents. The corresponding threshold for the tax is about €180.

7. Conclusion

Cities restrict short-term rentals to protect residents, but which residents benefit depends on what landlords do with the units that leave the STR market. In this paper we study Amsterdam's 2017 day cap with administrative data on households, housing transactions and mortgages, and with an equilibrium model in which landlords choose between STR, long-term rental and selling to households facing borrowing constraints.

The reduced-form evidence shows that the cap lowered Airbnb activity and house prices in the most exposed districts. The model shows that landlords respond mainly by selling. Tightening the cap from 60 to 30 nights sends about three in four units leaving STR to owner-occupiers, and lowers both rents and house prices by about 5%. The benefit to residents therefore comes mostly through the ownership market, and it reaches lower-income households to the extent that borrowing limits allow them to buy.

Table 7. Counterfactual equilibria: day cap, borrowing limits and a tax on STR stays

	Baseline	(ii) Cap	(iv) Cap 30	(v) Tax
	2018	30 nights	+ LTI 3	10%
Panel A: Prices (baseline level; percent change)				
STR price received by hosts (€ per night)	175.1	+16.7	+16.6	−8.1
STR price paid by visitors (€ per night)	175.1	+16.7	+16.6	+1.1
Free-sector rent (€000s per year)	12.13	−5.3	−5.5	−1.1
Ownership price (€000s)	448.0	−5.1	−6.9	−1.1
Panel B: Allocation of the housing stock (baseline level; percent change)				
STR units S	15,575	−26.3	−25.8	−6.1
Free-sector LTR units R	81,928	+1.2	+2.0	+0.3
Owner-occupied units O	129,302	+2.4	+1.8	+0.5
STR nights (millions)	1.126	−62.2	−61.9	−6.2
STR occupancy rate	0.198	−48.7	−48.6	−0.1
Panel C: Channels of the change in LTR units (baseline level; percent change)				
Rental portfolio $S + R$, outer nest (units)	97,503	−3.2	−2.4	−0.7
LTR share of portfolio $R/(S + R)$, inner nest	0.840	+4.5	+4.6	+1.0
Panel D: Households (baseline level; percent change)				
Households leaving the city (exits)	18,834	−21.7	−21.4	−5.0
Incumbent households moving	29,637	+3.8	+0.8	+0.4
Purchases by movers and entrants	15,427	+24.7	+16.6	+5.2
two lowest income quintiles	4,452	+44.7	+28.2	+9.2
two highest income quintiles	7,969	+12.9	+9.2	+2.8
Panel E: Welfare (million €, present value)				
Households, owner-occupiers		+1,185.3	+1,596.6	+257.9
Households, renters		+352.0	+352.2	+72.1
Households, entrants		+70.2	+72.0	+14.2
Visitors		−1,105.1	−1,100.8	−113.1
Investors		−3,753.4	−4,279.3	−827.1
City (STR tax revenue)		0	0	+340.0
Total		−3,251.0	−3,359.3	−255.9
Panel F: Compensating variation per household (€000s, present value)				
Owner-occupiers, lowest income quintile		+9.05	+11.82	+1.91
Owner-occupiers, highest income quintile		+9.23	+13.19	+2.05
Renters, lowest income quintile		+4.04	+4.18	+0.83
Renters, highest income quintile		+5.10	+5.03	+1.07

Notes: The table presents equilibrium effects relative to the 2018 baseline, which reproduces the observed allocation of the 69 products. The first column gives the baseline level; the other columns give the percent change (Panels A to D) or the welfare change (Panels E and F). Prices are means weighted by baseline units. Shock (ii) lowers the permitted nights from 60 to 30; occupancy follows equation (21) at each product's pre-cap natural rate, scaled to the realized occupancy of 2018. Shock (iv) combines shock (ii) with a loan-to-income limit of 3. Shock (v) imposes a 10 percent tax on the nightly STR price; hosts receive the net price and visitors pay the gross price. In Panel C the two percent changes add up to the percent change of R to first order. Exits are households choosing the outside option. Purchases are the ownership choices of incumbent movers and entrants in the year; income quintiles are those of the household types. Welfare is measured as in Appendix D.1: households by the compensating variation of equation (56), owner-occupiers and renters being the 211,230 incumbents of the 2018 stock by tenure and entrants the 18,834 households entering per year in the steady state of the baseline; visitors by the change in the logit surplus of equation (57) times the visitor flow; investors by the change in the expected per-unit payoff of equation (58) summed over the housing stock, including owner-occupied units. Visitor, investor and city welfare are capitalized at $1/(1-\beta) = 20$; the city's row is the tax revenue and is part of the total. Panel F reports averages over household types weighted by the number of households. Rows may not add to the total because of rounding.

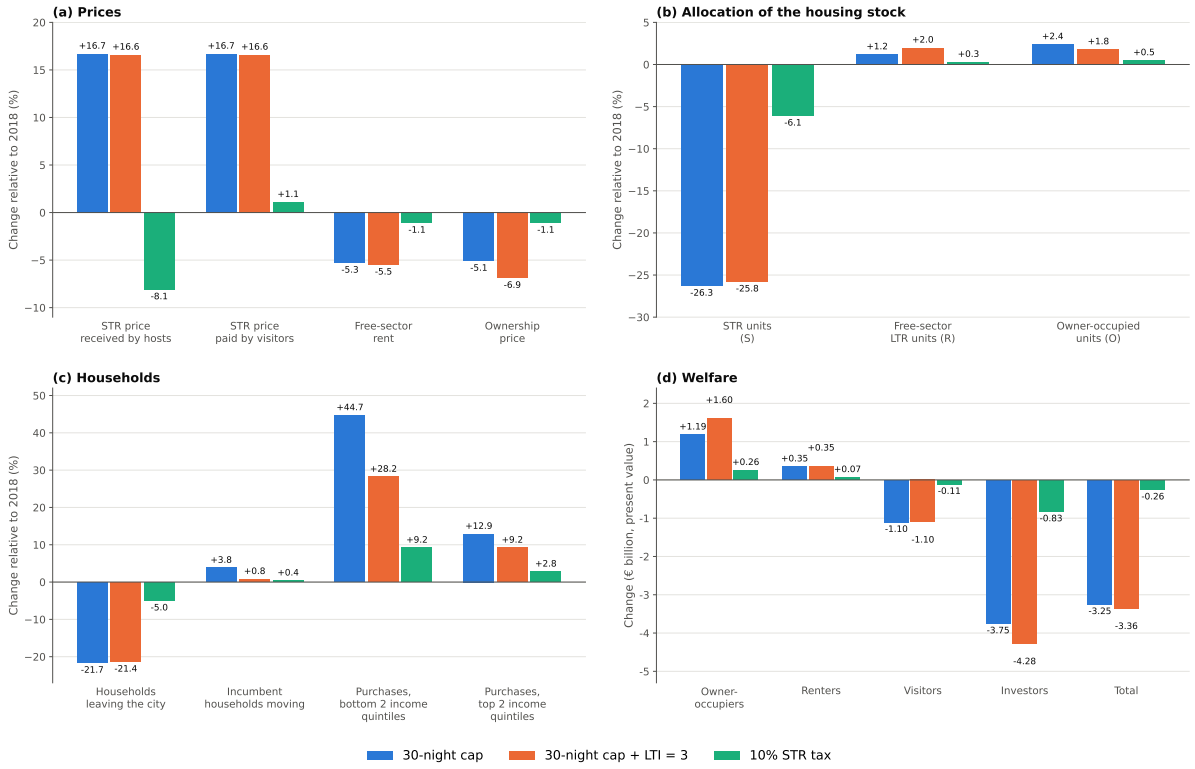


Figure 4. Counterfactual equilibria: 30-night cap, cap with LTI limit of 3, and 10% STR tax
Notes: Changes relative to the 2018 baseline, as in Table 7. Panels (a)–(c) report percent changes; panel (d) reports welfare changes in € billion (present value), with the total including entrants and the city's STR tax revenue.

With a loan-to-income limit of 3, the increase in purchases by low-income households is about a third smaller.

Residents' gains come at a large cost to landlords and visitors. A tax on STR stays, which Amsterdam already uses, delivers similar gains to residents per unit removed from STR at about a third of the aggregate cost, and collects as revenue the scarcity value that a cap passes on to the remaining hosts.

References

Acharya, Viral V., Katharina Bergant, Matteo Crosignani, Tim Eisert, and Fergal McCann, "The Anatomy of the Transmission of Macroprudential Policies," *Journal of Finance*, 2022, *77* (5), 2533–2575.

Almagro, Milena and Tomás Domínguez-Iino, "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," *Econometrica*, 2025, *93* (3), 1031–1071.

Autor, David H., Christopher J. Palmer, and Parag A. Pathak, "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts," *Journal of Political Economy*, 2014, *122* (3), 661–717.

Bajari, Patrick, C. Lanier Benkard, and John Krainer, "House Prices and Consumer Welfare," *Journal of Urban Economics*, 2005, *58* (3), 474–487.

Barron, Kyle, Edward Kung, and Davide Proserpio, "The Effect of Home-Sharing on House Prices and Rents: Evidence from Airbnb," *Marketing Science*, 2021, *40* (1), 23–47.

Bayer, Patrick, Robert McMillan, Alvin Murphy, and Christopher Timmins, "A Dynamic Model of Demand for Houses and Neighborhoods," *Econometrica*, 2016, *84* (3), 893–942.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen, "Inference on Treatment Effects after Selection among High-Dimensional Controls," *The Review of Economic Studies*, 2014, *81* (2), 608–650.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 1994, *25* (2), 242–262.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Calder-Wang, Sophie, "The Distributional Impact of the Sharing Economy on the Housing Market," *SSRN Working Paper*, 2021. Available at SSRN 3908062.

Caloia, Francesco, "Borrower-Based Measures, House Prices and Household Debt," DNB Working Paper 738, De Nederlandsche Bank 2022.

Cavalleri, Maria Chiara, Boris Cournède, and Ezgi Özsöğüt, "How Responsive Are Housing Markets in the OECD? National Level Estimates," OECD Economics Department Working Papers 1589, OECD Publishing, Paris 2019.

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

Eijsink, Gerard and Dorinth Van Dijk, "Borrowing Capacity and House Prices," DNB Analysis, De Nederlandsche Bank 2023.

Farronato, Chiara and Andrey Fradkin, "The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry," *American Economic Review*, 2022, *112* (6), 1782–1817.

Favilukis, Jack, Pierre Mabille, and Stijn Van Nieuwerburgh, "Affordable Housing and City Welfare," *Review of Economic Studies*, 2023, *90* (1), 293–330.

Garcia-López, Miquel-Àngel, Jordi Jofre-Monseny, Rodrigo Martínez-Mazza, and Mariona Segú, "Do Short-Term Rental Platforms Affect Housing Markets? Evidence from Airbnb in Barcelona," *Journal of Urban Economics*, 2020, *119*, 103278.

Garriga, Carlos, Pedro Gete, and Athena Tsouderou, "The Economic Effects of Real Estate Investors," *Real Estate Economics*, 2023, *51* (3), 655–685.

Gemeente Amsterdam, "Toeristische verhuur van woonruimte: Rapportage 1 januari 2017–1 juli 2017," https://openresearch.amsterdam/nl/page/61430/rapportage-toeristische-verhuur-van-woonruimte 2017.

Gete, Pedro and Michael Reher, "Mortgage Supply and Housing Rents," *Review of Financial Studies*, 2018, *31* (12), 4884–4911.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Greenwald, Daniel L. and Adam Guren, "Do Credit Conditions Move House Prices?," *American Economic Review*, 2025, *115* (10), 3559–3596.

Halket, Jonathan, Lars Nesheim, and Florian Oswald, "The Housing Stock, Housing Prices, and User Costs: The Roles of Location, Structure, and Unobserved Quality," *International Economic Review*, 2020, *61* (4), 1777–1814.

Horn, Keren and Mark Merante, "Is Home Sharing Driving up Rents? Evidence from Airbnb in Boston," *Journal of Housing Economics*, 2017, *38*, 14–24.

Hotz, V. Joseph and Robert A. Miller, "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 1993, *60* (3), 497–529.

Jaeger, David A., Joakim Ruist, and Jan Stuhler, "Shift-Share Instruments and the Impact of Immigration," NBER Working Paper 24285, National Bureau of Economic Research 2018.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, *128* (9), 3285–3345.

Kleibergen, Frank and Richard Paap, "Generalized Reduced Rank Tests Using the Singular Value Decomposition," *Journal of Econometrics*, 2006, *133* (1), 97–126.

Koijen, Ralph S. J. and Motohiro Yogo, "A Demand System Approach to Asset Pricing," *Journal of Political Economy*, 2019, *127* (4), 1475–1515.

Koster, Hans R. A., Jos van Ommeren, and Nicolas Volkhausen, "Short-Term Rentals and the Housing Market: Quasi-Experimental Evidence from Airbnb in Los Angeles," *Journal of Urban Economics*, 2021, *124*, 103356.

Kvaerner, Jens Soerlie, Nicola Pavanini, and Yushi Peng, "Leverage Regulation and Housing Inequality," *SSRN Working Paper*, 2023.

Li, Hui, Yijin Kim, and Kannan Srinivasan, "Market Shifts in a Sharing Economy: The Impact of Airbnb on Housing Rentals," *Management Science*, 2022, *68* (11), 8015–8044.

_ , _ , and _ , "Market Shifts in the Sharing Economy: The Impact of Airbnb on Housing Rentals," *Management Science*, 2022, *68* (11), 8015–8044.

Peydrò, J.-L., F. Rodriguez-Tous, J. Tripathy, and A. Uluc, "Macroprudential Policy, Mortgage Cycles and Distributional Effects: Evidence from the UK," *Review of Financial Studies*, 2024, *37* (3), 727–760.

Poterba, James M., "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach," *Quarterly Journal of Economics*, 1984, *99* (4), 729–752.

Rust, John, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

Sinai, Todd and Nicholas S. Souleles, "Owner-Occupied Housing as a Hedge Against Rent Risk," *Quarterly Journal of Economics*, 2005, *120* (2), 763–789.

Sommer, Kamila, Paul Sullivan, and Randal Verbrugge, "The Equilibrium Effect of Fundamentals on House Prices and Rents," *Journal of Monetary Economics*, 2013, *60*, 854–870.

Stock, James H. and Motohiro Yogo, "Testing for Weak Instruments in Linear IV Regression," in Donald W. K. Andrews and James H. Stock, eds., *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge University Press, 2005, pp. 80–108.

Thiel, Jurre H. and Henrik Zaunbrecher, "Mortgage Debt Limits and Buy-to-Let Investors: A Structural Model of Housing with an Endogenous Rental Sector," CPB Discussion Paper 449, CPB Netherlands Bureau for Economic Policy Analysis September 2023.

Tzur-Ilan, N., "Adjusting to Macroprudential Policies: Loan-to-Value Limits and Housing Choice," *Review of Financial Studies*, 2023, *36* (10), 399–4044.

van Bekkum, S., M. Gabarro, R. M. Irani, and J.-L. Peydrò, "The Real Effects of Borrower-Based Macroprudential Policy: Evidence from Administrative Household-Level Data," *Journal of Monetary Economics*, 2024, *147*, 103574.

Zervas, Georgios, Davide Proserpio, and John W. Byers, "The Rise of the Sharing Economy: Estimating the Impact of Airbnb on the Hotel Industry," *Journal of Marketing Research*, 2017, *54* (5), 687–705.

A. Extra Figures

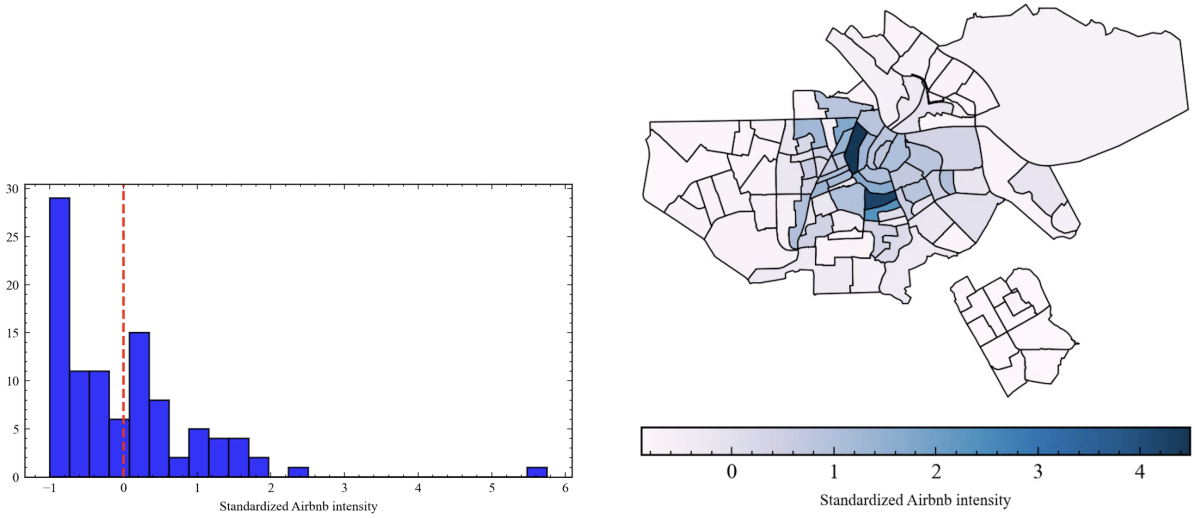


Figure A.1. Distribution of Baseline Airbnb Exposure

Notes: Panel (a): histogram of standardized Airbnb intensity across Amsterdam districts; the pre-regulation average (August 2015 to March 2017) is standardized to mean zero and unit standard deviation, so zero corresponds to the raw mean of 16.8 listings per 1,000 housing units. Panel (b): spatial distribution; darker shading indicates higher exposure. Grey areas indicate insufficient Airbnb data.

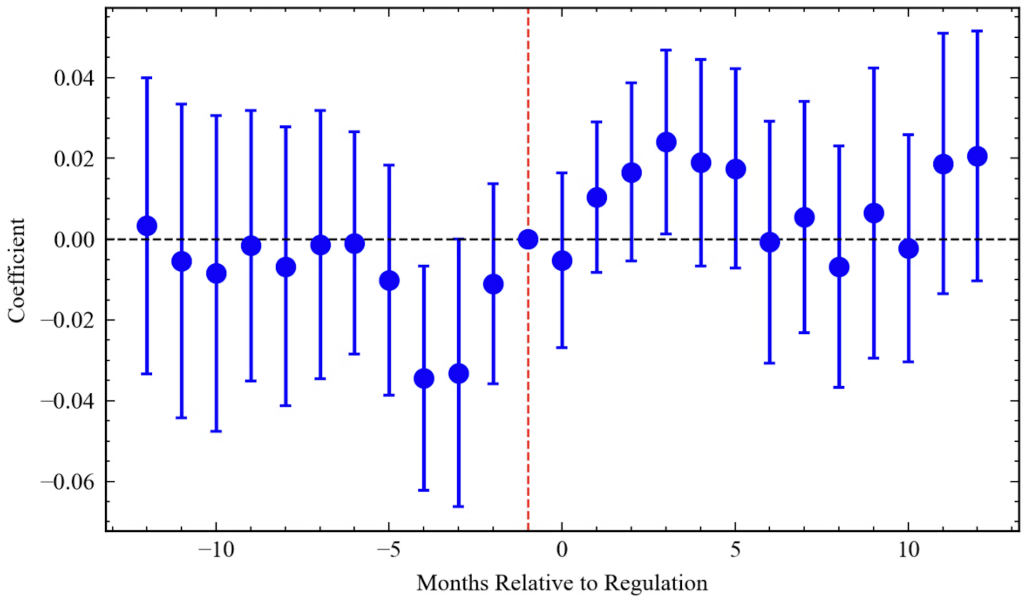


Figure A.2. Event Study: Effect of Regulation on Airbnb Nightly Prices

Notes: Dependent variable is log mean nightly price. Coefficients from equation (3); March 2017 is the omitted reference period and the vertical dashed line marks April 2017, the first post-regulation month. See Figure 2 for full specification details.

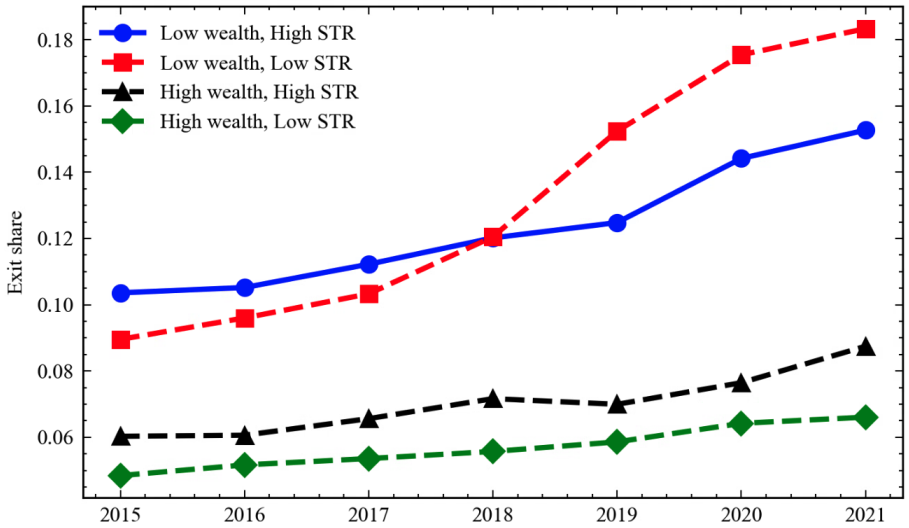


Figure A.3. Annual Exit Rates from Amsterdam's Private Housing Market by Wealth and STR Intensity

Notes: Annual exit rate is the share of households present in year t who have left Amsterdam's liberalized private market by year $t+1$ (including transitions to social housing, other Dutch cities, or abroad). Groups are defined by the interaction of household wealth (above/below median) and area STR intensity (above/below median Airbnb share in the household's area of residence). Sample: liberalized private-market households, 2015 to 2021. Source: CBS microdata (*Gbaadresobjectbus*, *Vehtab*), AirDNA.

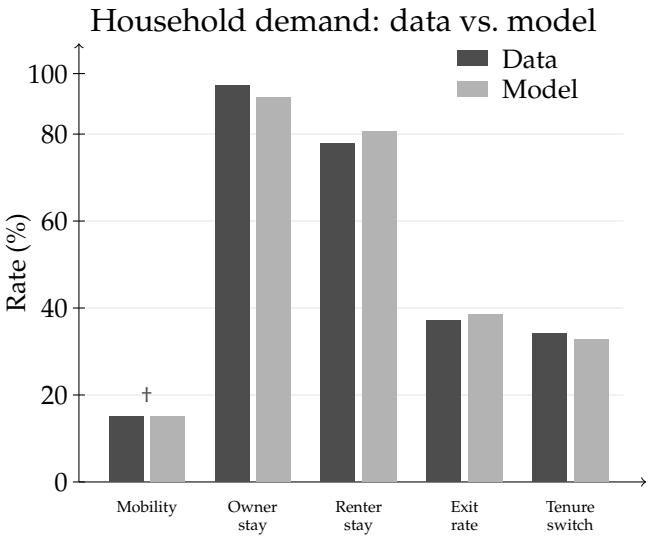


Figure A.4. Household demand: data versus model-predicted rates for five mobility margins. Dark bars are data; light bars are model; the dagger marks the mobility rate matched by construction (Step 1b MLE), while the remaining four moments (owner stay, renter stay, exit rate, tenure switch) are non-targeted.

Notes: Grouped bars compare data (dark) and model-predicted (light) rates for five household mobility margins. † denotes the mobility rate, which is matched by construction (Step 1b MLE); all other moments are non-targeted. McFadden pseudo-$R^2 = 0.335$ over 115 feasible alternatives.

B. Extra Tables

Table B.1. Amsterdam Geography by Borough

Borough	Districts	Pop. (000s)	Housing units	Social (%)	Owner (%)	Free (%)	Airbnb listings	Airbnb/ 1,000
Centrum	10	86.4	53,932	34.0	29.0	37.0	1,744	32.3
West	20	142.8	77,003	43.0	28.0	30.0	2,111	27.4
Zuid	16	141.4	78,238	29.0	30.0	41.0	1,549	19.8
Oost	14	128.7	63,834	52.0	29.0	19.0	977	15.3
Noord	14	82.6	37,869	66.0	27.0	7.0	273	7.2
Nieuw-West	14	146.7	63,802	53.0	30.0	17.0	211	3.3
Zuidoost	11	82.8	37,977	60.0	28.0	12.0	45	1.2
Amsterdam	99	811.4	412,655	46.0	29.0	26.0	6,908	16.8

Notes: Statistics for 2016 (pre-regulation). Social is percentage of housing stock owned by housing corporations. Owner is owner-occupancy rate. Free is share of total housing stock that is free-sector (market-rate) private rental. Airbnb listings is the average number of active entire-home listings per month during the pre-regulation period (August 2015 to March 2017). Airbnb/1,000 is listings per 1,000 housing units. Amsterdam has 110 administrative districts in total; the 99 districts reported here exclude water bodies and industrial areas with minimal residential activity (the housing transaction panel covers 106 districts; the 99-district count reflects the intersected sample with sufficient Airbnb coverage for the main specifications).

Table B.2. Summary Statistics: Airbnb Activity

	Mean	SD	P10	P50	P90
Panel A: Amsterdam (district-month level)					
Active listings	67.3	73.9	5.0	41.0	154.0
Listings per 1,000 housing units	16.8	15.3	1.7	13.2	36.4
Total reservation days	860	1,001	53	471	2,088
Total monthly revenue (€ 000s)	131	164	6	66	334
Mean nightly rate (€)	158	39	115	153	206
Share of listings in top decile			35.3%		
District-months			4,808		
Districts			99		
Panel B: Listing-level characteristics					
			Amsterdam		
Revenue per listing (€/month)	1,993	1,804	0	1,584	4,530
Reservation days	12.8	10.4	0.0	11.0	28.0
Occupancy rate (%)	63	37	0	76	100
Nightly rate (€)	169	84	97	148	264
Bedrooms	1.51	0.84	1.00	1.00	3.00
Listing-months			326,734		
Unique listings			29,159		
			Other Dutch Municipalities		
Revenue per listing (€/month)	1,286	1,194	0	1,053	2,792
Reservation days	13.1	10.1	0.0	12.0	28.0
Occupancy rate (%)	54	36	0	60	100
Nightly rate (€)	103	56	55	90	165
Bedrooms	1.83	1.01	1.00	2.00	3.00
Listing-months			450,124		
Unique listings			30,227		

Notes: Sample restricted to entire-home listings, August 2015 to September 2019. Panel A reports district-month level statistics for Amsterdam districts. The 4,808 district-months reflect 99 districts over 50 months less 142 district-months with no active listings, which are excluded from the panel. Panel B reports listing-month level statistics; a listing-month with zero reservations contributes zero revenue and a zero occupancy rate. Other Dutch Municipalities includes Rotterdam, The Hague, Utrecht, and other municipalities. Mean nightly rate in Panel A differs from Panel B because Panel A averages across district-months while Panel B is at the listing-month level.

Table B.3. Summary Statistics: Housing Transactions by Airbnb Exposure

	Mean	SD	P10	P50	P90
Panel A: High Airbnb exposure (transaction level)					
Price (€ 000s)	405.1	190.8	205.0	360.0	670.0
Price per m^2 (€)	5,114	1,303	3,491	5,000	6,908
Living area (m^2)	79.4	30.9	47.0	71.0	125.0
Building age (years)	92.8	70.0	20.0	97.0	118.0
Rooms	3.1	1.0	2.0	3.0	4.0
Apartment (%)	98.2				
Social housing (%)	11.2				
Transactions			22,062		
Districts			33		
Panel B: Low Airbnb exposure (transaction level)					
Price (€ 000s)	268.9	128.8	141.5	240.0	435.0
Price per m^2 (€)	2,867	953	1,860	2,688	4,066
Living area (m^2)	93.3	27.1	63.0	90.0	131.0
Building age (years)	36.7	22.1	9.0	35.0	62.0
Rooms	3.7	1.1	2.0	4.0	5.0
Apartment (%)	67.4				
Social housing (%)	21.5				
Transactions			6,831		
Districts			33		
Panel C: District-month level (full Amsterdam sample)					
Transactions per month	6.7	5.4	1.0	5.0	14.0
Mean price (€ 000s)	373.8	159.3	192.4	347.6	585.9
Mean price per m^2 (€)	4,257	1,466	2,289	4,174	6,275
Sparse cells (%)	12.7	33.3	0.0	0.0	100.0
District-months			6,654		
Districts			106		

Notes: Sample period January 2014 to December 2019. High/low exposure defined as top/bottom tercile of Airbnb listings per 1,000 housing units in the pre-regulation period; middle tercile (33 districts) omitted from Panels A and B. Apartment and Social housing rows report means only (no distributional statistics). Panel C includes all Amsterdam districts. Social housing indicates transactions involving a housing corporation. Sparse cells is a binary indicator equal to 100 if the district-month has fewer than three transactions and 0 otherwise; the mean of 12.7 reflects the fraction of such cells.

Table B.4. Summary Statistics: Rental Market by Airbnb Exposure (Survey Data)

	High Exposure				Low Exposure			
	Mean	SD	P50	P10/P90	Mean	SD	P50	P10/P90
Panel A: Regulated sector								
Monthly rent (€)	526	169	540	317/695	599	146	598	425/728
Living area (m^2)	60	21	57	36/85	68	20	70	42/91
WOZ value (€ 000s)	359	134	343	209/528	223	107	205	107/362
Construction year	1938	50	1929	1891/1998	1978	21	1978	1956/2008
Observations			8,025				4,701	
Panel B: Free sector								
Monthly rent (€)	1,467	482	1,399	916/2,045	1,318	565	1,183	894/1,889
Living area (m^2)	93	29	91	58/132	90	24	91	63/117
WOZ value (€ 000s)	543	179	523	345/771	443	202	401	233/729
Construction year	1979	72	1998	1912/2014	2007	18	2016	1983/2018
Observations			18,819				6,837	
Panel C: Market structure								
Free-sector share (%)			70.1				58.5	
Regulated share (%)			29.9				41.5	

Notes: Survey observations only, from the CBS Huurenquête, 2015 to 2023 (summary statistics use the full survey window for descriptive context; the rent regression analysis uses imputed population-level rents aggregated to district-year means, as described in Section 2.3). Amsterdam observations only. High/low exposure defined as top/bottom tercile of Airbnb listings per 1,000 housing units in the pre-regulation period; middle tercile omitted. Regulated sector includes social housing and regulated private rentals. Free sector comprises private rentals above the WWS threshold (∼142 points). WOZ = assessed property value. Free-sector shares reflect the survey sample composition and are not representative of the full housing stock; imputed population shares are reported in the regression analysis.

Table B.5. Household panel

	Statistic	Count
Panel A: Dimensions		
Household-year observations		961,602
Unique households		209,921
Mean years observed	4.58	
Panel B: Annual mobility		
Move rate (overall)	14.9%	126,636
Within liberal market	36.9%	46,675
Exit liberal market	63.1%	79,961
Stay rate: owners	91.0%	
Stay rate: renters	77.7%	
Tenure switch (within-Amsterdam movers)	46.8%	
Entrants (annual)		12,406
Panel C: Demographics (household-year means)		
Gross income (€ 000s/yr)	89.5	
(SD)	(72.4)	
Wealth (€ 000s)	260.7	
(SD)	(406.7)	
Negative wealth share	12.8%	
Owner-occupier share	58.8%	
Panel D: By mover status		
	Movers	Stayers
Mean age	41.6	50.7
Median income (€ 000s)	87.5	89.8
Median wealth (€ 000s)	170.8	276.2
Choose ownership (%)	53.4	

Notes: Liberalized private-market residents aged 18+, 2015 to 2019. Source: CBS microdata (*Gbaadresobjectbus*, *Gbapersoontab*, *Inhatab*, *Vehtab*, *Eigendomtab*). A "move" is a change in product-tenure bundle between consecutive January 1 snapshots. "Exit liberal market" includes transitions to social housing, other Dutch cities, or abroad. Within/Exit percentages are shares of movers. Wealth A_τ equals liquid wealth for renters and net wealth (liquid plus housing equity) for homeowners. "Choose ownership" is defined for within-Amsterdam movers only; stayers retain their current tenure by definition.

Table B.6. Housing stock allocation before and after the 2017 day-cap

	Centrum		Rest of Amsterdam	
	2016	2018	2016	2018
Stock allocation (%)				
Owner-occupied	39.9	39.4	55.0	52.5
LTR	51.0	52.8	39.7	42.2
STR	9.1	7.8	5.3	5.3
STR profitability				
Annual STR revenue / LTR rent	1.24	1.55	1.19	0.98
Mean Airbnb price (€/night)	184	219	193	153
Mean LTR rent (€/month)	891	1,003	847	935

Notes: Annual STR revenue = mean nightly price $\times$ actual observed occupied nights per year (not the regulatory cap); the revenue premium annualizes LTR rent ($\times 12$) for comparability. Monthly rents are reported in the table for readability. 2016 is the last pre-regulation year; 2018 is the first full year under the 60-night-cap. Shares are percentages of liberalized private stock H_{jt} and sum to 100% by the stock constraint.

Table B.7. Household preference and moving cost estimates

Parameter	Estimate	Std. Error
Price sensitivity (marginal utility of wealth)		
$\gamma_{\text{fmc},0}$ (baseline, per €1,000)	0.308	(0.002)
$\gamma_{\text{fmc},Y}$ (log income)	−0.053	(0.0004)
Homeownership premium		
α_h (non-pecuniary value of owning)	1.792	(0.074)
Baseline preferences		
α_0 (constant)	38.230	(0.389)
$\alpha_{Z,\ln Y}$ (log income)	−5.445	(0.049)
$\alpha_{Z,\ln A}$ (log wealth)	1.443	(0.011)
$\alpha_{Z,\text{age}}$ (age, years)	0.004	(0.002)
$\alpha_{Z,\text{fam}}$ (family size)	0.017	(0.069)
Amenity preferences (baseline)		
$\beta_{0,\text{transit}}$ (metro proximity)	0.230	(0.060)
$\beta_{0,\text{school}}$ (school quality)	0.260	(0.050)
$\beta_{0,\text{green}}$ (green space)	0.150	(0.040)
Amenity-demographic interactions		
$\beta_{Z,\text{fam}\times\text{school}}$	0.080	(0.020)
$\beta_{Z,\text{age}\times\text{green}}$	0.004	(0.001)
Outside option		
$\alpha_{\text{out},0}$ (baseline)	38.450	(2.777)
$\alpha_{\text{out},Y}$ (log income)	−3.405	(0.521)
$\alpha_{\text{out,age}}$ (age, years)	0.004	(0.027)
$\alpha_{\text{out,fam}}$ (family size)	0.150	(0.090)
Psychological moving costs (utility units)		
$\gamma_{\text{pmc},0}$ (baseline)	3.012	(0.022)
$\gamma_{\text{pmc},Y}$ (log income)	0.222	(0.004)
$\gamma_{\text{pmc,age}}$ (age, years)	0.053	(0.0003)
$\gamma_{\text{pmc,fam}}$ (family size)	0.006	(0.008)

Notes: Bootstrapped standard errors (500 replications of Steps 1a to 1d) in parentheses. Price sensitivity (γ_{fmc}) estimated via MLE at the move/stay margin (Step 1b). Preference parameters (α, β) estimated via OLS decomposition of recovered flow utilities (Step 1d). Model fit statistics reported separately in Figure A.4.

Table B.8. Borrowing constraint anatomy: renter and owner households by wealth quartile

		Constraint status (%)				Median	Can afford
	N	Only LTV	Only LTI	Both	Neither	affordable	median (%)
Panel A: Renter households							
Q1 (€ -549k to € 0k)	14,391	17.3	0.0	82.7	0.0	0	0.0
Q2 (€ 0k to € 6k)	17,976	14.3	0.0	85.7	0.0	0	0.0
Q3 (€ 6k to € 37k)	16,183	18.6	22.3	44.3	14.8	1	14.8
Q4 (€ 37k to € 4,575k)	16,184	0.0	30.8	0.0	69.2	57	69.2
Panel B: Owner households (memo)							
Q1 (€ -508k to € 188k)	18,086	13.7	15.8	7.5	63.0	45	63.0
Q2 (€ 188k to € 340k)	18,086	0.0	5.1	0.0	94.9	63	94.9
Q3 (€ 340k to € 577k)	18,086	0.0	0.1	0.0	99.9	67	99.9
Q4 (€ 577k to € 4,466k)	18,086	0.0	0.0	0.0	100.0	69	100.0

Notes: Constraint status evaluated at the median-priced ownership product (p^O_{med}). LTV binds if $A_\tau < (1 - \overline{\text{LTV}}) \cdot p^O_{\text{med}}$; LTI binds if $p^O_{\text{med}} - A_\tau > \overline{\text{LTI}} \cdot Y_\tau$. "Median affordable" is the median number of ownership products (out of 69) for which both constraints are satisfied. "Can afford median" is the share of households that can afford the median-priced product. Parameters: $\overline{\text{LTV}} = 0.95$, $\overline{\text{LTI}} = 4$ (see equation (5)).

C. Contemporaneous Housing Policies

Three contemporaneous Dutch housing-policy reforms overlap with our sample period, but each operates on a margin distinct from the rental-to-short-term-rental stock withdrawal this paper studies. First, the assessed property value (WOZ) entered the WWS point system in October 2015, mechanically pushing some high-WOZ central-district units above the 142-point cutoff into the free sector. Second, a new legislation (*Wet doorstroming huurmarkt*) introduced temporary free-sector contracts in July 2016. Third, investor purchases of low-priced peripheral-district properties intensified. The first two reforms reclassify or re-contract units within the rental market without withdrawing them, while investor purchases concentrated in peripheral districts anti-correlated with Airbnb's central tourist core. The supply-withdrawal mechanism is therefore unlikely to be confounded, and the free-sector rent prediction in particular is insulated because sector classification is dominated by ownership type rather than the WWS-point cutoff through which the reform operates (Appendix I.3). We do not, however, separately identify reform-driven and Airbnb-driven contributions to central-Amsterdam house-price capitalization, where WOZ-WWS reclassification and Airbnb both load on central geography. We start the sample in 2015 to align with the WOZ-WWS reform and with AirDNA's earliest coverage of all four major Dutch metros.

D. Platform Cross-Listings

Gemeente Amsterdam (2017) documents that advertised listings across all six monitored platforms (Airbnb, Booking, Wimdu, HomeAway, Flipkey, 9flats) moved in parallel through the first half of 2017 and that the agreement with Airbnb did not produce migration of activity to other platforms, the opposite of what platform-substitution would predict. This pattern is consistent with the regulation's institutional design: the 60-night cap is a legal restriction on entire-home short-term rentals that applies to all platforms equally, so a host bound by the cap on Airbnb cannot lawfully continue the activity on another platform, and platform-shifting can evade the cap only through under-enforcement of non-Airbnb listings. No listing-level cross-platform panel is available to test proportional co-movement directly; we require only the weaker condition that other-platform activity not grow at a higher proportional rate than Airbnb in response to the identifying shocks, under which our instrumental-variables estimates are lower bounds on the effect of total short-term rental activity. Because the report covers only the first half of 2017, we apply this evidence to the full sample under the assumption that the platform-neutral enforcement environment and platform shares did not shift materially, and we treat Airbnb activity as a proxy for total short-term rental exposure throughout.

E. Model Discussion

First, the two-level nested logit on the investor side is new to the STR literature. The nested structure generates the paper's central identification: the separation of the STR/LTR margin (inner nest) from the rental-versus-ownership margin (outer nest) allows the portfolio-reallocation channel to emerge as a distinct force from within-portfolio reallocation. Single-margin STR models (Calder-Wang, 2021) correspond to the inner nest alone; demand-based asset-pricing models following Koijen and Yogo

(2019) correspond to the outer nest alone. Neither captures the effect of a binding day-cap that reallocates units from STR to LTR through the inner nest, but simultaneously expands ownership supply through the outer nest, attenuating the gross LTR supply gain substantially at median shares.

Second, the supply side is static while the household demand side is dynamic. This asymmetry reflects the data structure: aggregate allocation shares are observed but not sequences of individual investor decisions, so dynamic estimation would require specifying investor expectations over future regulatory regimes that the data cannot pin down. The static assumption is most restrictive for the outer nest, where selling a property is costly and partially irreversible; estimated leakage fractions should therefore be interpreted as upper bounds on the steady-state portfolio-reallocation response. The static-vs-dynamic direction of bias is, however, time-horizon dependent: forward-looking investors anticipating further regulatory tightening would preemptively reallocate, so the cumulative portfolio-reallocation response over a multi-year horizon could exceed the static estimate, in which case our magnitudes are conservative on the headline reallocation finding rather than on the leakage fraction. A dynamic extension with convex adjustment costs is the natural next step; empirical resolution of these competing forces would require either panel data on individual investor portfolios or a structural dynamic model with explicit regulatory-anticipation expectations. A related limitation concerns partial-listing Airbnb activity: approximately 15% of Amsterdam listings are private or shared rooms (Section 2.3), and the 60-night cap as enforced applies only to entire-home listings, so the cap raises the option value of owner-occupation differentially, since an owner retains the uncapped option to monetize spare capacity via room rental while a long-term renter does not. The model treats Airbnb activity as homogeneous and the cap as binding on all listings, so this option-value channel is unmodeled; to the extent the cap raises owner-occupation's option value relative to long-term rental, the model understates the share of freed STR units that ultimately flow to ownership, the same direction as the static-supply bias, leaving the portfolio-reallocation magnitude conservative. The exogenous-stock assumption is instead a medium-run scope condition: the long-run construction response it omits could attenuate the leakage rather than amplify it (Section 2.4).

Third, the model abstracts from the mortgage market: borrowing constraints enter through LTV and LTI caps evaluated at current prices, but there is no explicit balance sheet for mortgage lenders and no pass-through from monetary policy to housing demand beyond the mortgage-rate instrument used in Section 5.3. Extending the model to incorporate the mortgage market, with leveraged investors facing borrowing costs and households facing endogenous credit supply, is the planned next-stage contribution of this project and would deliver a quantitative counterpart to the evidence in Section 3 on the channel from STR regulation to house prices.

Beyond these design choices, the equilibrium abstracts from several features of the Dutch housing market and the modeled agent population. Our framing is that the model captures a medium-run steady state at the cross-sectional product level rather than year-by-year market dynamics, and the following abstractions deserve explicit acknowledgment.

Contract frictions and steady-state framing. Dutch tenant protection (*huurbescherming*) and the indefinite default rental contract make rental tenure persistent: in the renter panel, annual stay rates are 77.7% (Section 2), implying a roughly 22% annual turnover rate at which contracts can reprice at market terms. The free sector also operated

without a statutory rent ceiling before 1 May 2021, and government-set annual maxima cap regulated-sector rent increases. The model's equilibrium prices clear the LTR market in each period without explicitly modeling these contract frictions, so the equilibrium is best read as the medium-run cross-sectional configuration toward which year-by-year repricing would converge under the observed cross-product pattern of STR, LTR, and ownership exposure. Identification of the household price-sensitivity parameter $\gamma^{\tau}_{\text{fmc}}$ comes from cross-sectional variation in moving and staying decisions across product-types rather than from regulation-driven repricing within a single product, so the steady-state read does not threaten the demand-side estimates. On the supply side, identification does draw in part on the 2017 day-cap: the regulation-bite instrument is the post-2017 indicator interacted with cross-sectional pre-cap occupancy gaps, so a change in frictions is used to identify a model without frictions. The robustness to this is that the Bartik instrument, which shifts STR returns through global tourist-demand growth rather than through the regulation, yields a closely agreeing just-identified inner-nest estimate (revenue-differential coefficient 0.074 versus 0.080 for the regulation-bite instrument; Section 5.3); a non-regulation shifter and the regulation-change shifter delivering the same supply elasticity means identification does not hinge on the regulation-change variation, which is what makes the steady-state reading safe.

Interest-rate framing. The model's user-cost specification depends on the contemporaneous post-tax mortgage rate r^m_t. In reality, Dutch household mortgages are typically fixed for ten or more years at origination and many incumbents hold mortgages dating from the low-rate environment of 2015 to 2019; investor portfolios assembled in the same period likewise carry below-market financing. We read the contemporaneous r^m_t as a summary of the steady-state cost of capital faced by marginal buyers and sellers, not as the financing position of all incumbent owners or all incumbent investors. The static specification therefore overstates household sales responses to rate-change counterfactuals because lock-in is absent, and conversely understates investor leverage responses because mortgage-financed portfolio expansion is not modeled. We do not run rate-change counterfactuals in this paper; the counterfactuals hold r^m_t fixed and so are not affected by the lock-in/leverage abstraction.

Household-vs-investor classification of multi-property dwellings. The model treats households (decision-makers over location and tenure) and investors (allocators of the dwelling stock across STR, LTR, and owner-occupation) as distinct agent types; in reality, some households own multiple dwellings and act simultaneously on both margins. In the CBS administrative data, the household block uses individual-level population-register records to characterize where each household lives, while the investor block characterizes dwelling-level allocation at the product-year level. A household that owns more than one Amsterdam dwelling enters the household block through their primary-residence choice and contributes to the investor stock through any non-primary dwellings they hold. The model does not separately label multi-property households as a distinct agent type and so cannot trace within-household coordination between primary-residence decisions and second-home allocation choices; a richer household-investor-coordination extension is left to future work. Because the outer- and inner-nest reallocation elasticities are identified from market-level stock reallocation under the as-if-representative-marginal-investor assumption (Section 4.4), the composition of the actors behind that reallocation, multi-property households included, does not bias the portfolio-reallocation magnitude itself; what the sharp split omits is the within-household feedback between a multi-property household's

primary-residence and second-home decisions, which would affect counterfactuals on the joint tenure margin rather than the market-level reallocation elasticity the paper identifies.

Synthesis. The model abstracts from contract rigidity, mortgage lock-in, dynamic investor leverage, and the sharp household/investor split for multi-property owners. Each abstraction leaves the paper's headline portfolio-reallocation finding intact, making the frictionless comparative statics an upper bound on its magnitude, not bearing on the counterfactuals we run, or being too small to matter: contract frictions slow rather than reverse the price-quantity adjustments the model identifies, so the static steady state bounds the realized reallocation from above rather than arising as an artifact of the frictionless assumption; lock-in and leverage are absent from rate-change counterfactuals we do not perform; the household/investor split affects only within-household coordination at the joint tenure margin, leaving the market-level reallocation elasticity unchanged; and the LTI institutional drift shifts the household borrowing set by at most roughly 3%, small relative to the cross-sectional income and wealth variation that drives the household-side response to STR counterfactuals.

The block-recursive structure of the model, in which household and visitor demand condition on observed prices, makes sequential estimation feasible: household preferences are recovered from choice probabilities via the Hotz and Miller inversion, visitor demand from Berry inversion of the logit share equation, and investor supply from the nested-logit inversion conditional on the demand-side estimates. Section 5 develops this procedure and reports the estimated parameters.

F. Theory

F.1. Welfare Measures

The model implies welfare measures for each group of agents that can be computed at any set of equilibrium prices and allocations. We define these measures here; their evaluation at estimated parameter values under counterfactual policies is left for future work.

Household welfare. For a household of type τ at state Ω_{it}, the ex-ante expected utility follows from the logit structure of the dynamic program. Taking the expectation over i.i.d. T1EV shocks in equation (16) yields the closed-form ex-ante value function:

$$\mathbb{V}_{it}(\Omega_{it}) = \ln\left(\exp\left(v^{\tau}_{d_{i,t-1},t}\right) + \exp\Big(\ln \sum_{k\in\{0\}\cup\mathcal{K}^{\bar{\tau}}_t} \exp\left(v^{\bar{\tau}}_{kt}\right) - \mathrm{PMC}_{\tau}\Big)\right), \tag{54}$$

where we omit the additive Euler-Mascheroni constant ($\gamma_E \approx 0.5772$), which cancels in welfare differences. The financial moving cost enters through the type transition $\tau \to \bar{\tau}$: the mover's value $v^{\bar{\tau}}_{kt}$ is evaluated at post-move wealth (equation (8)), so no separate FMC term appears. Changes in $\mathbb{V}_{it}$ across policy regimes measure the welfare change in utils. To convert to monetary units, we divide by the marginal utility of wealth $\gamma^{\tau}_{\text{fmc}}$ from equation (9):

$$\mathrm{CV}^{\tau}_{t} = \frac{\mathbb{V}_{it}(\text{counterfactual}) - \mathbb{V}_{it}(\text{baseline})}{\gamma^{\tau}_{\text{fmc}}}. \tag{55}$$

Since $\mathbb{V}_{it}$ depends on the household's current location $d_{i,t-1}$, which varies within type, the type-level CV is an average over the within-type state distribution: $\mathrm{CV}_t^\tau \equiv \mathbb{E}_{d_{i,t-1}|\tau}\left[(\mathbb{V}_{it}(\text{counterfactual}) - \mathbb{V}_{it}(\text{baseline}))/\gamma_{\text{fmc}}^\tau\right]$, where the expectation is taken over the observed distribution of locations among type-τ households at t. Aggregate household welfare is then $\Delta W_t^{HH} = \sum_\tau N_t^\tau \cdot \mathrm{CV}_t^\tau$, where N_t^τ is the number of type-τ households. This measure captures both the direct effect of price changes on housing costs and the indirect effects on mobility, sorting, and feasible set changes through borrowing constraints.

Visitor welfare. The logit consumer surplus for the representative visitor is:

$$\mathrm{CS}_t^V = \frac{1}{\alpha_A} \ln\left(1 + \sum_{j \in \mathcal{J}} \exp\left(\bar{\delta}_{jt}^A\right)\right), \tag{56}$$

where $\bar{\delta}_{jt}^A$ is the mean visitor utility defined in equation (18). Since CS_t^V is a per-period (annual) flow, we capitalize to make it commensurable with the household present value: $\Delta W_t^V = \frac{Q_t^V}{1-\beta} \cdot [\mathrm{CS}_t^V(\text{counterfactual}) - \mathrm{CS}_t^V(\text{baseline})]$, where the factor $1/(1-\beta)$ converts the annual surplus change to a present-discounted value under the assumption that the policy change is permanent. Since the hotel outside option is normalized to zero, this measures welfare relative to a world where all visitors stay in hotels.

Investor welfare. For investors, the expected per-unit payoff from the outer nest is:

$$\Pi_{jt} = \sigma_{\text{I}} \ln\Big(\exp\left(v_{jt}^K/\sigma_{\text{I}}\right) + \exp\left(v_{jt}^O/\sigma_{\text{I}}\right)\Big), \tag{57}$$

where v_{jt}^K and v_{jt}^O are the rental and ownership payoffs defined in equations (28) and (27) respectively, and total investor surplus is $W_t^I = \frac{1}{1-\beta} \sum_{j \in \mathcal{J}} H_{jt} \cdot \Pi_{jt}$, where the capitalization factor $1/(1-\beta)$ converts the annual per-unit payoff to a present-discounted value, consistent with the household welfare measure. Changes in W_t^I across regimes capture the effect of regulation on property owners' returns.[62]

Total welfare. Total welfare is the sum $\Delta W_t = \Delta W_t^{HH} + \Delta W_t^V + \Delta W_t^I$, measured in present-discounted EUR. Distributional analysis decomposes ΔW_t^{HH} by household type (income, wealth, tenure, location) to assess how regulation redistributes across the wealth distribution.

F.2. Comparative Statics Derivations

This subsection derives the supply-side diversion ratios reported in Section 5.3.3.

Derivation of the diversion ratios. Throughout, let

$$D_{jt} \equiv v_{jt}^S - v_{jt}^L = 365\, p_{jt}^A \phi_{jt}^A + \xi_{jt}^S - r_{jt}^L \tag{58}$$

[62]The level expression omits the additive constant $\sigma_{\text{I}}\gamma_E$, which cancels in welfare differences ΔW_t^I. The analogous inner-nest constant $\sigma_\kappa \gamma_E$ is absorbed into the estimated unobservable η_{jt} during Berry inversion and is therefore already present in the estimated v_{jt}^K.

denote the STR/LTR payoff differential, using equations (23) and (24). Factoring $\exp(v^L_{jt}/\sigma_\kappa)$ out of equation (26) gives $IV_{jt} = v^L_{jt} + \sigma_\kappa \ln\left(1+\exp(D_{jt}/\sigma_\kappa)\right)$. Substituting $v^L_{jt} = r^L_{jt}$ and $1 - s^S_{jt} = 1/\left(1+\exp(D_{jt}/\sigma_\kappa)\right)$ yields the equivalent expression $IV_{jt} = r^L_{jt} - \sigma_\kappa \ln(1-s^S_{jt})$, used in the derivations below.

Derivation of the diversion ratios. Hold rents, ownership prices and ξ^S_{jt} fixed. From the logit shares (25) and (29), and using $\partial \mathrm{IV}_{jt}/\partial v^S_{jt} = s^S_{jt}$ and $\partial \mathrm{IV}_{jt}/\partial v^L_{jt} = 1 - s^S_{jt}$,

$$\frac{\partial s^S_{jt}}{\partial v^S_{jt}} = \frac{s^S_{jt}(1-s^S_{jt})}{\sigma_\kappa}, \qquad \frac{\partial s^K_{jt}}{\partial v^S_{jt}} = \frac{s^K_{jt}(1-s^K_{jt})}{\sigma_I}\, s^S_{jt},$$
$$\frac{\partial s^S_{jt}}{\partial v^L_{jt}} = -\frac{s^S_{jt}(1-s^S_{jt})}{\sigma_\kappa}, \qquad \frac{\partial s^K_{jt}}{\partial v^L_{jt}} = \frac{s^K_{jt}(1-s^K_{jt})}{\sigma_I}\,(1-s^S_{jt}).$$

Differentiating $S_{jt} = H_{jt}s^K_{jt}s^S_{jt}$, $R_{jt} = H_{jt}s^K_{jt}(1-s^S_{jt})$ and $O_{jt} = H_{jt}(1-s^K_{jt})$ with respect to v^S_{jt} gives

$$\frac{\partial S_{jt}}{\partial v^S_{jt}} = H_{jt}s^K_{jt}s^S_{jt}\,\Psi_{jt},$$
$$\frac{\partial R_{jt}}{\partial v^S_{jt}} = -H_{jt}s^K_{jt}s^S_{jt}(1-s^S_{jt})\left[\tfrac{1}{\sigma_\kappa} - \tfrac{1-s^K_{jt}}{\sigma_I}\right],$$
$$\frac{\partial O_{jt}}{\partial v^S_{jt}} = -H_{jt}s^K_{jt}s^S_{jt}\,\tfrac{1-s^K_{jt}}{\sigma_I},$$

with Ψ_{jt} as defined in equation (52). The three derivatives sum to zero because H_{jt} is fixed. Dividing the second and third by minus the first yields DR^L_{jt} and DR^O_{jt} in equation (52). The common positive factor $H_{jt}s^K_{jt}s^S_{jt}$ cancels, so the sign of DR^L_{jt} is that of $1/\sigma_\kappa - (1-s^K_{jt})/\sigma_I$, which is positive if and only if $s^K_{jt} > 1 - \sigma_I/\sigma_\kappa$. At estimated parameters ($\hat\sigma_I = 10.2$, $\hat\sigma_\kappa = 13.2$) the threshold is approximately 0.23, and the condition holds for all products except Weesp (rental share 12.6%, negligible market weight).

Analogously, differentiating with respect to v^L_{jt},

$$\frac{\partial R_{jt}}{\partial v^L_{jt}} = H_{jt}s^K_{jt}(1-s^S_{jt})\,\widetilde{\Psi}_{jt} > 0,$$
$$\frac{\partial S_{jt}}{\partial v^L_{jt}} = -H_{jt}s^K_{jt}s^S_{jt}(1-s^S_{jt})\left[\tfrac{1}{\sigma_\kappa} - \tfrac{1-s^K_{jt}}{\sigma_I}\right],$$
$$\frac{\partial O_{jt}}{\partial v^L_{jt}} = -H_{jt}s^K_{jt}(1-s^S_{jt})\,\tfrac{1-s^K_{jt}}{\sigma_I},$$

and dividing the second and third by minus the first yields $\widetilde{DR}^S_{jt}$ and $\widetilde{DR}^O_{jt}$ in equation (53). LTR supply falls with v^L_{jt} for all parameter values, whereas the number of STR units rises only if $s^K_{jt} > 1 - \sigma_I/\sigma_\kappa$.

G. Empirical Facts: Additional Results and Robustness

This appendix collects the deferred material referenced in Section 3: the full regression tables behind the body's summary figures, robustness checks across alternative specifications, borough-trends and binary-treatment alternatives, the Bartik first-stage table, the heterogeneity analysis that pins down the supply-withdrawal mechanism, the rent-IV robustness across alternative shift variables, and the unselected event-study panels. The appendix is organized into four subsections mirroring the body structure: G.1 Effect of Regulation on Airbnb Activity, G.2 Effect of Airbnb on House Prices, G.3 Effect of Airbnb on Rents, and G.4 Effect of Airbnb on Mortgages.

G.1. Effect of Regulation on Airbnb Activity

Table G.1 reports the baseline difference-in-differences specification from equation (2) applied to five Airbnb activity outcomes, along with the borough-trends specification that augments the baseline with borough-specific linear time trends. The extensive-margin coefficients fall by roughly 6 to 7% per standard deviation of baseline exposure across both specifications; the intensive-margin total reservation-days coefficient is unchanged at -0.154. The nightly-price coefficient rises from 0.019 to 0.024 under borough trends, consistent with some of the supply-side price pass-through being absorbed by borough-level appreciation trends in the baseline.

Table G.2 summarizes alternative specifications: borough trends (same as column 2 of Table G.1, repeated here for comparison), a binary treatment indicator for the top tercile of baseline exposure, and a sample excluding the Centrum borough. The main quantity findings survive across all alternatives; the price finding loses significance under the binary specification (consistent with reduced power from discretization) but is robust to the Centrum exclusion.

Table G.3 tests whether the decline in Airbnb activity is concentrated in districts where the supply-withdrawal mechanism operates. The triple-interaction coefficient on Exposure $\times$ Post $\times$ High Private Rental is negative and marginally significant for the extensive margin (-0.078 and -0.073 for log listings and log listings per 1,000, $p < 0.10$), while the main effect is near zero and statistically insignificant. This pattern supports the mechanism precondition: regulation predominantly displaces private-sector hosts from the short-term market, consistent with the interpretation that the units withdrawn from Airbnb are the same units competing in the long-term rental market.

The unselected event-study panels are collected in Figures G.1 and G.2. The raw log listings panel in Figure G.1 complements the per-1,000-housing-units panel in the body by using unnormalized counts; the two deliver the same qualitative pattern. The mean reservation days per listing panel in Figure G.2 shows the intensive margin on a per-listing basis rather than the total; surviving listings average more days per listing once marginal hosts exit. This is consistent either with a composition effect, in which exiting hosts are marginal operators with low baseline utilization while a residual pool of dedicated operators uses the full 60-night allowance, or with surviving hosts strategically concentrating bookings once marginal competitors exit; the data do not allow separate identification of these two mechanisms.

Table G.1. Difference-in-Differences: Effect of April 2017 Regulation on Airbnb Activity

	(1) Baseline	(2) Borough Trends
Extensive margin		
ln(Listings)	-0.067^{***}	-0.049^{**}
	(0.022)	(0.022)
ln(Listings per 1,000)	-0.065^{***}	-0.046^{**}
	(0.021)	(0.021)
Intensive margin, totals		
ln(Total reservation days)	-0.154^{***}	-0.154^{***}
	(0.023)	(0.025)
Intensive margin, per-listing means		
ln(Mean reservation days per listing)	-0.078^{***}	-0.094^{***}
	(0.011)	(0.010)
Price		
ln(Mean nightly price)	0.019^{**}	0.024^{***}
	(0.009)	(0.009)
District FE	Yes	Yes
Year-Month FE	Yes	Yes
Borough trends	No	Yes
R^2 (reservation days)	0.971	0.973
Observations	2,177	
Districts	99	

Notes: Unit of observation is district-month. Sample period April 2016 to March 2018 (12 months before and after the April 2017 regulation). Each row reports a separate regression of the outcome on Exposure $\times$ Post, with district and year-month fixed effects. Exposure is standardized baseline Airbnb density (listings per 1,000 housing units, pre-regulation mean; standardized to mean zero and unit standard deviation). Borough trends are borough-specific linear time trends ($\delta_s \cdot t$). All outcomes in logs. Observations weighted by baseline listing counts. Standard errors clustered at the district level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table G.2. Robustness: Effect of April 2017 Regulation on Airbnb Activity (Alternative Specifications)

	Extensive Margin		Intensive Margin		Price
Specification	$\ln(\text{List.})$	$\ln(\text{List.}/1000)$	$\ln(\text{Days})$	$\ln(\bar{\text{Days}})$	$\ln(\bar{P})$
Baseline	-0.067^{***}	-0.065^{***}	-0.154^{***}	-0.078^{***}	0.019^{**}
Alternative specifications					
Borough trends	-0.049^{**}	-0.046^{**}	-0.154^{***}	-0.094^{***}	0.024^{***}
Alternative treatment definitions					
Binary (top tercile)	-0.044^{*}	-0.042^{*}	-0.117^{***}	-0.066^{***}	0.020
Sample restrictions					
Exclude Centrum	-0.060^{*}	-0.056^{*}	-0.188^{***}	-0.115^{***}	0.030^{**}

Notes: Each cell reports the coefficient on Exposure × Post from a separate regression; standard errors omitted for brevity, full results available on request. Baseline specification as in Table G.1. Borough trends adds borough-specific linear time trends. Binary treatment is an indicator for top tercile of baseline Airbnb density (bottom two terciles omitted). Exclude Centrum drops the ten districts in Amsterdam's historic center ($N = 2{,}164$). Variables with bars denote per-listing means; $\bar{P}$ denotes mean nightly price. All outcomes in logs. All baseline coefficients significant at 1% or better except $\ln(\bar{P})$ (significant at 5%). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table G.3. Heterogeneous Effects of Regulation on Airbnb Activity by Private Rental Share

	Extensive Margin		Intensive Margin		Price
	$\ln(\text{List.})$	$\ln(\text{List.}/1000)$	$\ln(\text{Days})$	$\ln(\bar{\text{Days}})$	$\ln(\bar{P})$
Exp. × Post	0.004	0.004	-0.153^{***}	-0.144^{***}	0.031^{**}
	(0.038)	(0.038)	(0.039)	(0.014)	(0.016)
Exp. × Post × High Priv. Rental	-0.078^{*}	-0.073^{*}	-0.032	0.045^{**}	-0.004
	(0.045)	(0.044)	(0.052)	(0.019)	(0.021)
Total effect (high priv. rental)	-0.074	-0.069	-0.185	-0.099	0.027
District FE	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes
Borough trends	No	No	No	No	No
Observations			2,177		

Notes: High Private Rental is an indicator for above-median share of private rental housing in the district (baseline period). Specification includes all lower-order interactions. Total effect row reports the sum of main and interaction coefficients; standard errors are not reported for this row. Variables with bars denote per-listing means; $\bar{P}$ denotes mean nightly price. All outcomes in logs. Standard errors clustered at district level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

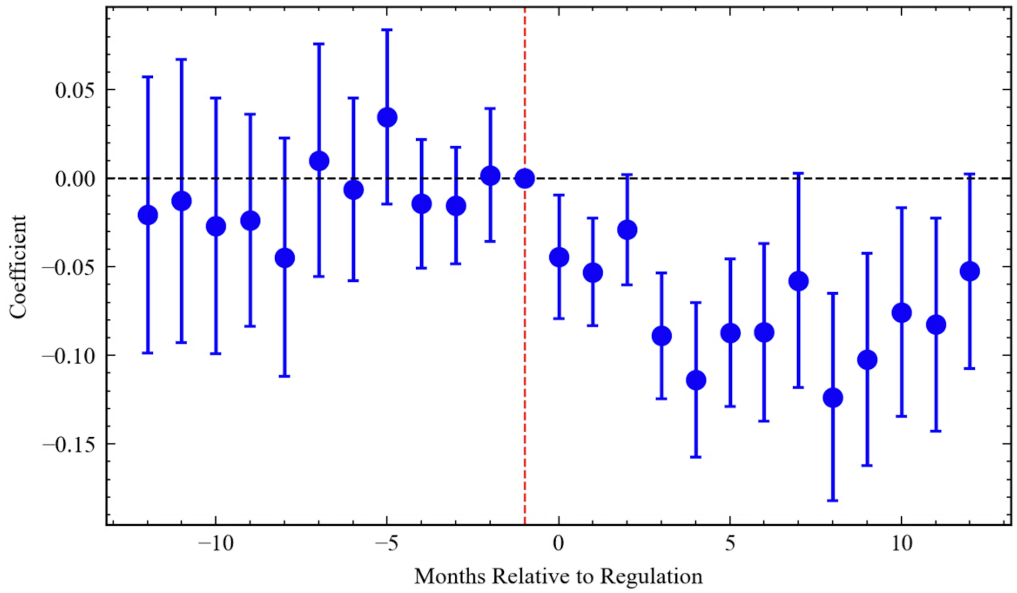


Figure G.1. Event Study: Effect of Regulation on Airbnb Listings (unnormalized)

Notes: Dependent variable is log total listings (raw, not normalized by housing stock). Coefficients from equation (3). Complements Figure 2 in the body, which normalizes by housing units.

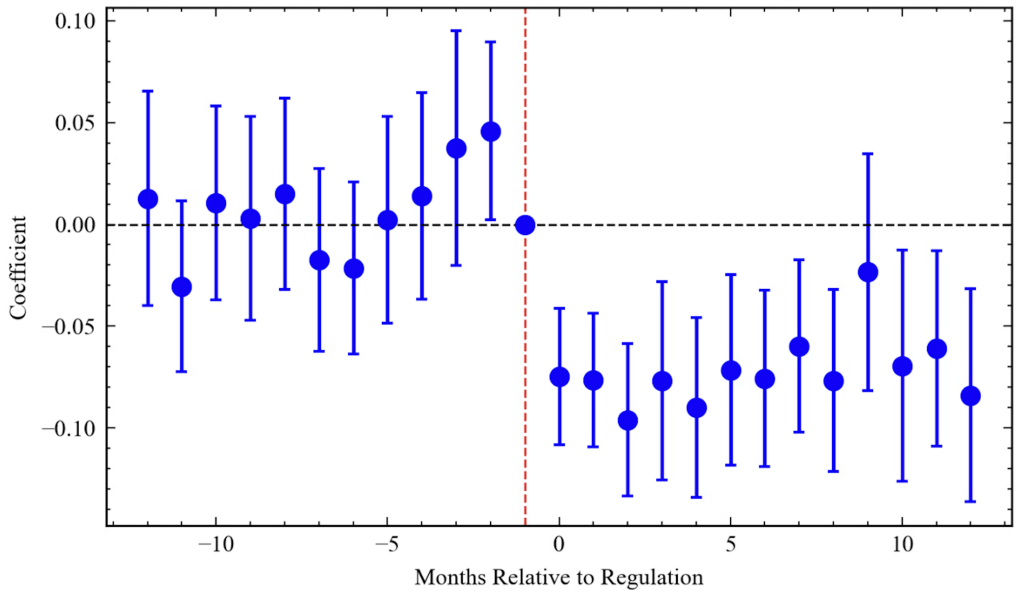


Figure G.2. Event Study: Effect of Regulation on Airbnb Utilization (per listing)

Notes: Dependent variable is log mean reservation days per listing. Coefficients from equation (3). Complements Figure 2 in the body, which uses total reservation days. The larger magnitude of the per-listing decline relative to the total indicates that surviving listings concentrate bookings once marginal hosts exit.

G.2. Effect of Airbnb on House Prices

Table G.4 reports the DiD specification applied to three house-price outcomes: the hedonic residual from the post-double-selection LASSO, the hedonic residual per square meter, and log mean price per square meter without hedonic adjustment. All three estimates are negative and statistically significant; the hedonic-residual specification is the preferred outcome in the body.

Table G.4. Difference-in-Differences: Effect of April 2017 Regulation on House Prices

	Hedonic (1)	Hedonic/m^2 (2)	ln(Price/m^2) (3)
Exposure × Post	-0.011^{*}	-0.013^{**}	-0.013^{**}
	(0.006)	(0.005)	(0.006)
R^2	0.854	0.832	0.959
District FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Borough trends	No	No	No
Observations	2,177	2,177	2,177

Notes: Unit of observation is district-month. Sample period April 2016 to March 2018 (12 months before and after the April 2017 regulation). Exposure is standardized baseline Airbnb density (listings per 1,000 housing units, pre-regulation mean; standardized to mean zero and unit standard deviation). Dependent variables: column (1) is the district-month mean hedonic residual from a post-double-selection LASSO regression of log transaction price on property characteristics (see Appendix I.2); column (2) is the hedonic residual per square meter; column (3) is log mean transaction price per square meter, a simpler measure that does not rely on the hedonic model. All specifications weighted by number of transactions. Robustness checks including borough-specific linear trends are reported in Table G.7. Standard errors clustered at district level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

We also employ an alternative identification strategy relying on a Bartik instrument, which interacts a pre-regulation district touristiness index (share) with total Airbnb listings in other Dutch cities (shift). The exclusion restriction is that platform expansion outside Amsterdam affects Amsterdam house prices only through Amsterdam's own Airbnb penetration, conditional on district and year-month fixed effects. The shift-share design admits two identifying decompositions: Goldsmith-Pinkham et al. (2020) rest identification on the conditional exogeneity of the exposure shares, while Borusyak et al. (2022) rest it on the orthogonality of many independent aggregate shifts. Because our identifying variation comes from a single aggregate expansion shock captured by two correlated shift measures, the design rests primarily on share exogeneity: the touristiness components were measured in 2015, prior to the platform's rapid expansion, and the income balance in Table 1 rules out the simplest gentrification violation. The complementary requirement is shift orthogonality to Amsterdam-specific shocks; the leading threat is inter-city substitution (Jaeger et al., 2018), if Amsterdam's regulation displaced tourists or hosts toward other Dutch cities. Amsterdam's regulatory pressure is small relative to other-Dutch-city growth (other cities continued expanding while

Amsterdam contracted, Figure 1), so even full diversion of Amsterdam's reduction would shift the aggregate growth path by a small fraction.

Table G.5 reports the first stage: the two instruments have individual Kleibergen and Paap F-statistics of 25.14 and 18.85, both well above conventional weak-instrument thresholds, while the joint F-statistic of 13.39 is lower because the two shift measures are correlated; single-instrument specifications confirm the results are stable using either instrument alone. Figure G.3 visualizes the relationship as a binned scatterplot of residualized log Airbnb density against the national growth instrument. Table G.6 reports OLS and 2SLS estimates. OLS is positive but modest: a 10% increase in Airbnb density is associated with a 0.4% price increase. The 2SLS estimates are roughly eight times larger: a 10% increase in Airbnb density raises house prices by 3.1 to 3.4% across the three outcome measures (Table G.6, Panel B). This gap runs counter to the standard prior that OLS is biased upward by Airbnb's concentration in appreciating high-amenity districts. Either net downward endogeneity (tourists and locals sorting across districts in opposing directions) or LATE heterogeneity (complier districts with structurally stronger supply-withdrawal dynamics) could open the gap; the heterogeneity analysis in Appendix G does not support structurally larger elasticities in touristy districts, so we attribute the gap primarily to net downward endogeneity purged by the instrument. Because Airbnb accounts for approximately 80% of advertised STR listings (Section 2.3), the IV coefficient on log Airbnb density is a lower bound on the per-unit elasticity with respect to total short-term rental supply; rescaling by $1/0.80$ gives 3.9 to 4.3% as an illustrative upper bound.[63]

Table G.6 presents OLS and 2SLS estimates of the elasticity of house prices with respect to Airbnb density. OLS estimates are positive and statistically significant but modest in magnitude (0.040); 2SLS estimates are roughly 7 to 8 times larger (0.309 to 0.335), consistent with attenuation bias in OLS dominating the upward selection bias expected from the concentration of Airbnb in high-amenity districts.

Table G.7 reports the full set of robustness checks for both the DiD and IV designs. The DiD borough-trends specification yields an estimate of -0.051 compared to the baseline -0.011; the borough-trends amplification has a natural interpretation in terms of within-borough appreciation trajectories and is discussed in the body. The IV estimates are stable across instrument choices except for the Netherlands-specific Google Trends shift, which is dominated by Amsterdam itself and therefore partially reflects Amsterdam's own Airbnb dynamics rather than external platform growth.

The unselected event-study panels for the house-price outcomes are collected in Figures G.4 and G.5. The level hedonic-residual panel in Figure G.4 omits the per-m^2 normalization; the no-hedonic panel in Figure G.5 omits the LASSO adjustment. Both panels deliver the same qualitative pattern as the body figure.

G.3. Effect of Airbnb on Rents

Table G.8 tests the sensitivity of the segment-specific rent IV estimates to the choice of shift variable. Six shift measures are used, each interacted with the same touristiness

[63] Precisely, the lower bound requires other-platform activity not to grow at a higher proportional rate than Airbnb in response to the identifying shocks; the rescaling upper bound additionally requires other-platform activity not to decline as Airbnb grows. The cross-platform evidence in Section 2.3 is consistent with both conditions at the aggregate level but does not directly verify them at the within-district margin. The proxy reasoning bears on the level of the elasticity, not the OLS to 2SLS gap: it applies to OLS as well, so it cannot explain why the 2SLS estimate exceeds OLS.

Table G.5. First Stage: Touristiness and Airbnb Penetration

Dep. var.:	(1)	(2)	(3)
	ln(Airbnb listings per 1,000)		
Touristiness × National Growth	0.174***		0.127***
	(0.035)		(0.029)
Touristiness × Google Trends		0.239***	0.095*
		(0.055)	(0.052)
Kleibergen and Paap F-statistic	25.14	18.85	13.39
p-value	0.000	0.000	0.000
District FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Borough trends	No	No	No
Observations	4,808	4,808	4,808

Notes: Unit of observation is district-month. Sample period August 2015 to September 2019. Touristiness is the standardized first principal component of tourist infrastructure measured in 2015, prior to significant Airbnb penetration. National Growth is total Airbnb listings in other Dutch cities excluding Amsterdam. Google Trends is worldwide search intensity for "Airbnb." Regressions weighted by number of transactions (for consistency with the second stage). Standard errors clustered at district level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

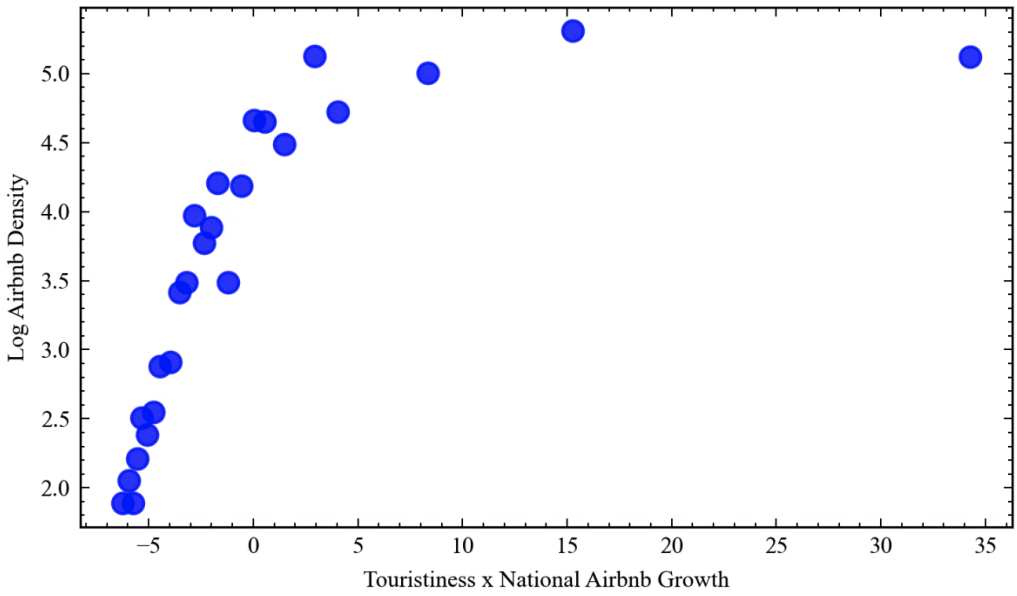


Figure G.3. First Stage: Touristiness and Airbnb Growth

Notes: Binned scatterplot of log Airbnb density against the touristiness × national growth instrument, after partialling out district and year-month fixed effects. Each dot represents a bin of districts. The positive slope identifies the first stage: districts with higher pre-determined touristiness saw proportionally larger Airbnb growth as the platform expanded across the Netherlands.

Table G.6. Effect of Airbnb Density on House Prices: OLS and 2SLS Estimates

	Hedonic (1)	Hedonic/m^2 (2)	ln(Price/m^2) (3)
Panel A: OLS			
ln(Airbnb/1000)	0.040***	0.041***	0.039***
	(0.015)	(0.015)	(0.014)
Panel B: 2SLS			
ln(Airbnb/1000)	0.313***	0.335***	0.309***
	(0.091)	(0.096)	(0.080)
First-stage F	13.39	13.39	13.39
District FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Borough trends	No	No	No
Observations	4,808	4,808	4,808

Notes: Unit of observation is district-month. Sample period August 2015 to September 2019. Endogenous variable: log active entire-home Airbnb listings per 1,000 housing units. Dependent variables: column (1) is the district-month mean hedonic residual; column (2) is the hedonic residual per square meter; column (3) is log mean transaction price per square meter. Panel A reports OLS estimates. Panel B reports 2SLS estimates using both instruments (Touristiness $\times$ National Growth and Touristiness $\times$ Google Trends); coefficients are elasticities, so 0.313 means a 10% increase in Airbnb density raises prices by 3.1%. Regressions weighted by number of transactions. Standard errors clustered at district level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

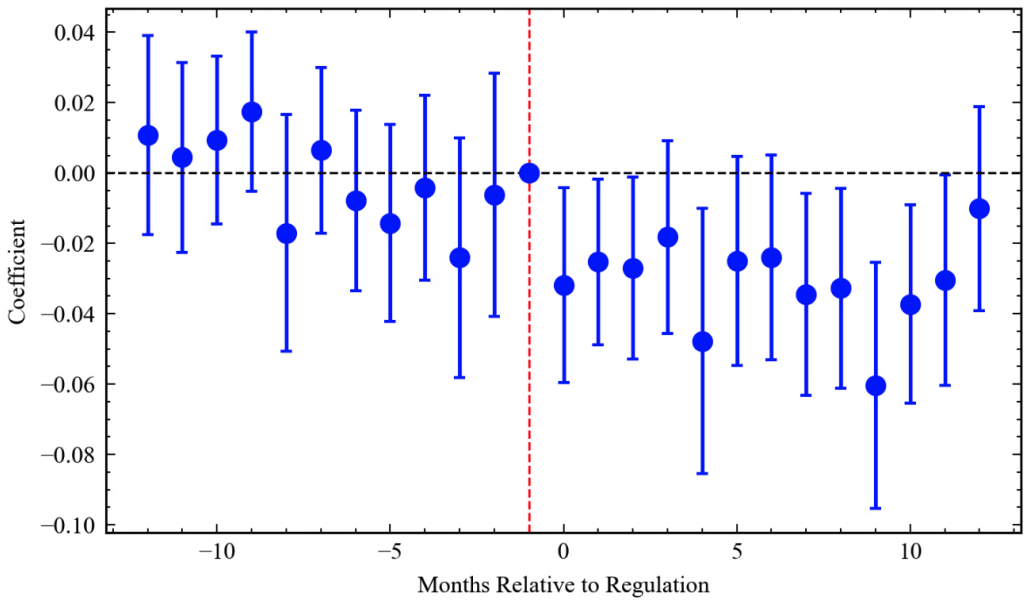


Figure G.4. Event Study: Effect of Regulation on House Prices (hedonic residual)

Notes: Dependent variable is the hedonic residual from a post-double-selection LASSO regression of log transaction price on property characteristics (without per-m^2 normalization). Coefficients from equation (3). Complements Figure 3 in the body, which reports the per-m^2 hedonic specification.

Table G.7. Robustness: Effect of Airbnb on House Prices (Alternative DiD and IV Specifications)

Specification	Hedonic residual	
	Estimate	SE
Panel A: Difference-in-Differences		
Baseline	-0.011^{*}	(0.006)
Borough trends	-0.051^{***}	(0.008)
Binary treatment (top tercile)	-0.004	(0.007)
Exclude Centrum	-0.011^{*}	(0.006)
Panel B: Instrumental Variables		
Baseline (both instruments)	0.313^{***}	(0.091)
National growth instrument only	0.379^{***}	(0.121)
Google Trends (world) only	0.266^{***}	(0.080)
Google Trends (NL) only	0.126	(0.106)
Exclude Centrum	0.326^{***}	(0.125)

Notes: Dependent variable is hedonic residual (weighted by transaction counts) in all specifications. Panel A reports the baseline specification (district and year-month fixed effects) and robustness checks; coefficient on Exposure × Post with standard error in parentheses. Borough trends adds borough-specific linear time trends ($\delta_s \cdot t$). Binary treatment uses an indicator for the top tercile of baseline Airbnb density. Panel B: IV coefficient on log Airbnb density with standard error in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

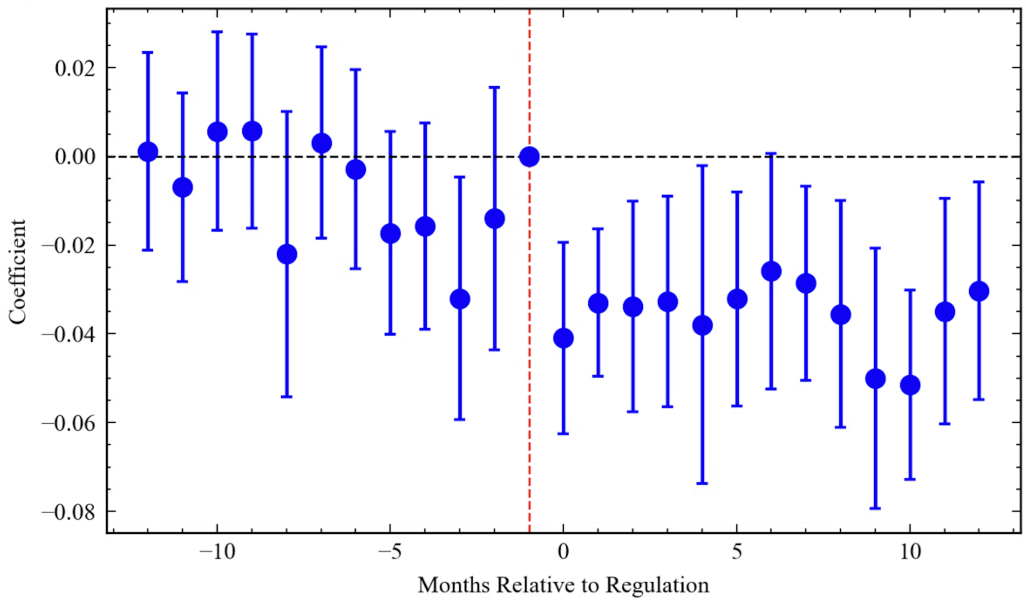


Figure G.5. Event Study: Effect of Regulation on House Prices (log price per m^2, no hedonic adjustment)

Notes: Dependent variable is log mean transaction price per square meter, aggregated to district-month without the post-LASSO hedonic adjustment. Coefficients from equation (3). Provided as a specification check that the headline result does not depend on the LASSO hedonic model.

share: Airbnb listings in other Dutch cities (the baseline), Google Trends (Netherlands), Google Trends (worldwide), foreign tourist stays, foreign tourist guests, and total tourist stays in Amsterdam. Free-sector estimates range from 0.067 to 0.075 across all six instruments, a range narrow relative to the baseline standard error of 0.032. The regulated-sector estimates cluster tightly around zero, ranging from -0.016 to -0.024 with p-values between 0.391 and 0.678. The stability of both the point estimates and the segment contrast across instruments provides strong support for the robustness of the free-sector result.

Table G.8. Robustness: Segment-Specific Rent Effects Across Shift Variables

	IV Coefficient (SE)		p-value	
Shift Variable	Regulated	Free	Reg.	Free
Airbnb listings (other cities)	-0.018	0.073^{**}	0.620	0.023
	(0.036)	(0.032)		
Google Trends (NL)	-0.019	0.072^{**}	0.585	0.019
	(0.034)	(0.030)		
Google Trends (World)	-0.016	0.075^{**}	0.673	0.030
	(0.038)	(0.034)		
Tourist stays (foreign)	-0.016	0.074^{**}	0.678	0.029
	(0.038)	(0.034)		
Tourist guests (foreign)	-0.019	0.071^{**}	0.570	0.017
	(0.034)	(0.030)		
Tourist stays (total)	-0.024	0.067^{***}	0.391	0.005
	(0.027)	(0.024)		

Notes: Each row reports IV estimates from a separate regression using a different shift variable. All specifications include district, year, and segment fixed effects with 586 observations and 98 clusters. Standard errors clustered at district level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The rent estimates face three channels of attenuation, which imply they should be interpreted as conservative lower bounds on the true effect. First, imputation error in the ML-constructed district-year rent panel introduces classical measurement error in the dependent variable, inflating residual variance and widening confidence intervals without biasing point estimates. Second, the stock-based rent measure responds slowly to market shocks, so effects may accumulate beyond our three-year window. Third, because the dependent variable is itself estimated, the reported standard errors do not account for first-stage imputation uncertainty and are therefore likely understated; this generated-regressor problem is mitigated by the large size of the imputation training sample (over 58,000 survey observations) relative to the 586-observation regression panel, but the confidence intervals should be read as conservative bounds. Re-estimating the rent regressions on the survey-only subset (without imputation) is structurally infeasible: stratified *Huurenquête* sampling yields too few observations per district-year-segment cell to construct an Amsterdam-wide panel (Section 2.3), which is precisely why the imputation strategy is necessary.

G.4. Effect of Airbnb on Mortgages

We use mortgage originations for home purchases in the four largest Dutch cities over 2013–2019, deduplicated to one origination per borrower and aggregated across loan parts, trimmed at the 1st and 99th percentiles of loan size, LTV, and LTI. The within-Amsterdam estimation sample contains 19,605 originations (2015–2018). Table G.9 reports summary statistics.

Table G.9. Summary Statistics: Mortgage Origination Sample

	Mean	SD	P10	P50	P90
Panel A: Origination characteristics (all four cities)					
Loan amount (€)	283,227	171,896	118,656	237,930	509,150
LTV (%)	85.4	19.5	56.0	91.5	103.0
LTI (ratio)	3.94	1.35	2.18	4.01	5.39
Interest rate (%)	2.46	0.71	1.70	2.33	3.45
Panel B: Originations by city					
Amsterdam (363)	22,988				
The Hague (518)	11,538				
Rotterdam (599)	9,772				
Utrecht (344)	6,766				
Originations (total)	51,064				

Notes: The table presents summary statistics for the mortgage origination sample. The sample consists of home-purchase mortgage originations from the DNB Loan Level Data, 2013–2019, deduplicated to the borrower level and trimmed at the 1st/99th percentiles of loan amount, LTV, and LTI. LTV is measured relative to appraised value, and LTI is total mortgage debt over gross household income.

The price and rent sections are estimated at the district (*wijk*) level. The Airbnb exposure measure *can* be built at the district level from the geocoded listing data. The mortgage data, however, geocode each loan only to its four-digit postal-code (PC4) area, and PC4 areas do not nest within districts (a PC4 spans 1.4 districts on average). Loans therefore cannot be assigned to districts directly, so we retain PC4 as the spatial unit, the postal-code counterpart of the districts used elsewhere in the chapter. We index PC4 areas by j.

Exposure is defined exactly as in the price and rent sections: Exposure_j is pre-regulation Airbnb intensity, defined as active entire-home listings per 1,000 housing units, averaged over the pre-regulation months and standardized to mean zero and unit standard deviation, so the coefficient is a per-standard-deviation effect.

We also define the difference-in-differences regression as in the previous sections, where for mortgage outcome M of origination i in PC4 area j and month t:

$$M_{ijt} = \alpha_j + \gamma_{\ell t} + \beta \left(\text{Exposure}_j \times \text{Post}_t\right) + \varepsilon_{ijt}, \tag{59}$$

where α_j are area fixed effects, $\gamma_{\ell t}$ are lender-by-month fixed effects, Post_t indicates months from April 2017, and standard errors are clustered at the PC4 level. For the interest rate we additionally absorb loan-product fixed effects. The event study replaces Post_t with quarter indicators (omitting 2017q1), and we report joint and linear pre-trend tests.

The second strategy instruments log Airbnb density with the shift-share $Z_{jt} = \text{Touristiness}_j \times \text{AirbnbGrowth}_t$, where Touristiness_j is the first principal component of 2015 CBS tourist-infrastructure proximity (hotels, museums, restaurants, cafés and bars, attractions, and centrality), which is the PC4 analog of the district touristiness index of Section 2.4.3, and AirbnbGrowth_t is total entire-home listings in other Dutch cities, indexed to August 2015. We estimate the first stage and two-stage least squares at the PC4-by-month level (Amsterdam), weighting by the number of originations and clustering at the PC4 level.

G.4.1. Results

Table G.10 reports the difference-in-differences estimates. A one-standard-deviation increase in Exposure is associated with a 1.7% smaller loan after the regulation. This is the mortgage-market counterpart of the price decline in Section 3.2. LTV, LTI, and the interest rate do not respond (columns 2–4): the regulation moved loan *amounts*, tracking prices, rather than the *terms* of credit.

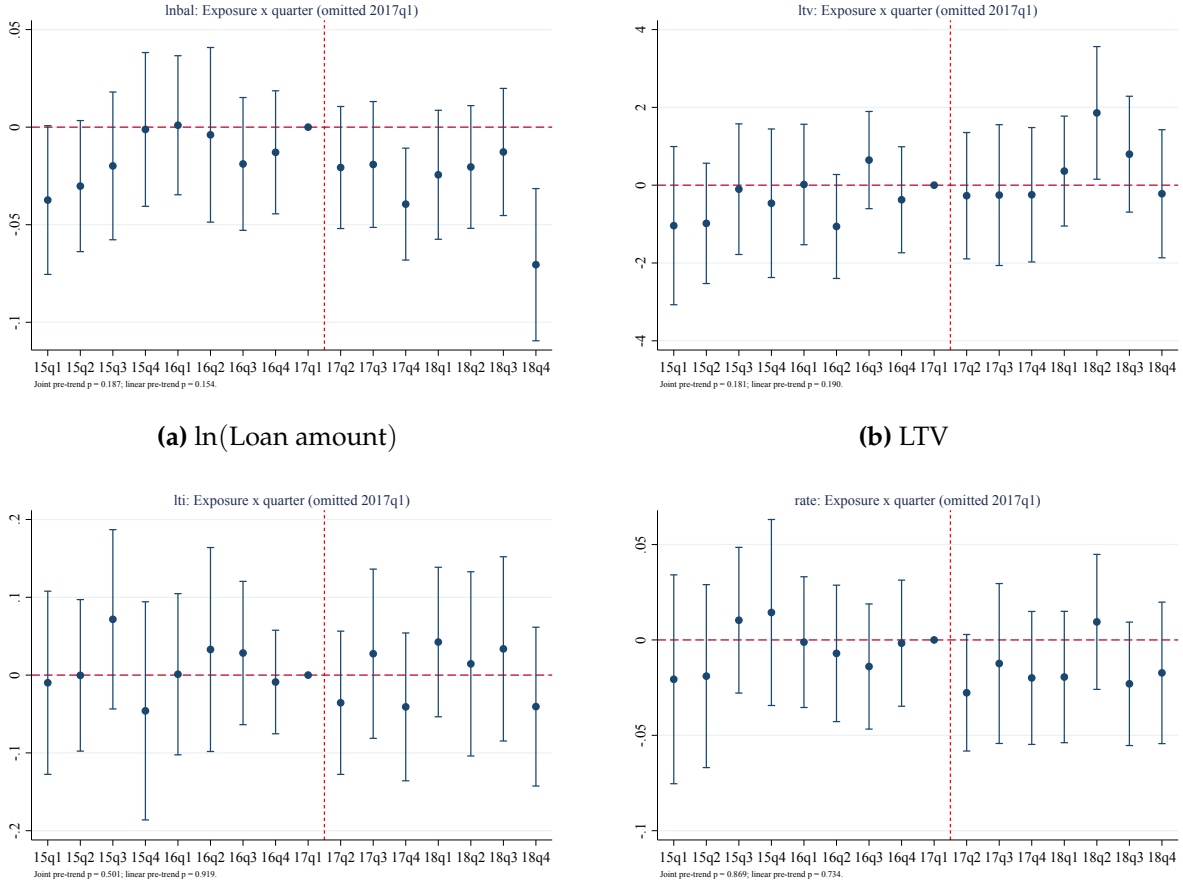
Table G.10. Effect of the Airbnb Regulation on Mortgage Outcomes (Difference-in-Differences)

	(1) ln(Loan)	(2) LTV	(3) LTI	(4) Rate
$\text{Exposure}_j \times \text{Post}_t$	-0.017^{**}	0.587	-0.008	-0.008
	(0.007)	(0.369)	(0.020)	(0.006)
Area (PC4) FE	Yes	Yes	Yes	Yes
Lender × Month FE	Yes	Yes	Yes	Yes
Product × Lender FE	–	–	–	Yes
Observations	19,605	19,605	19,605	19,575

Notes: The table presents difference-in-differences estimates of the effect of the regulation on mortgage outcomes. Each column reports the coefficient on $\text{Exposure}_j \times \text{Post}_t$ from equation (59), estimated within Amsterdam over 2015–2018. ln(Loan) is the log origination amount, LTV and LTI are in level units, and Rate is in percentage points. Exposure_j is standardized pre-regulation Airbnb intensity, so coefficients are per standard deviation. Product×Lender FE denotes loan-product-by-lender fixed effects, where a loan product is a combination of the fixed-rate period, the repayment type, and related contract characteristics. Standard errors are clustered at the PC4 level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The event study (Figure G.6) shows flat pre-trends for all four outcomes; neither the joint nor the linear pre-trend test rejects. One caveat applies to the loan-size result: the post-period coefficient turns clearly negative only in the final quarter of the window (2018q4), so the average effect is partly driven by the end of the sample. Dropping that quarter reduces the estimate to -0.011 and renders it insignificant. The loan-size effect should therefore be read as suggestive; the leverage and rate nulls, by contrast, are flat throughout.

Table G.11 reports the Bartik IV. The instrument is strongly relevant: the first-stage F-statistic is 53.9. The OLS association between log Airbnb density and loan size is positive, and the 2SLS estimate is larger, implying that a 10% increase in Airbnb density raises loan size by about 1.9%. The magnitude and the OLS<IV pattern match the



(a) $\ln(\text{Loan amount})$ **(b)** LTV

(c) LTI **(d)** Interest rate

Figure G.6. Event Study: Effect of the Airbnb Regulation on Mortgage Outcomes

Notes: The figure presents event-study estimates of the effect of the regulation on mortgage outcomes. Each panel plots the coefficients on $\text{Exposure}_j \times \mathbf{1}[q(t) = k]$ from the event-study analog of equation (59), estimated within Amsterdam by calendar quarter relative to the omitted reference quarter 2017q1; the vertical line marks the April 2017 regulation. In panel (a), the loan-amount coefficient turns clearly negative only in 2018q4. All regressions include area (PC4), lender×month, and loan-count fixed effects, and the bars are 95% confidence intervals with standard errors clustered at the PC4 level.

house-price elasticity in Section 3.2. As in the difference-in-differences, LTV, LTI, and the rate show no significant IV effect. One caveat is that the IV is estimated at the PC4 level rather than the district level used for prices and rents.

Table G.11. Effect of Airbnb Density on Mortgage Outcomes: OLS and 2SLS

	(1) ln(Loan)	(2) LTV	(3) LTI	(4) Rate
Panel A: OLS				
ln(Airbnb density)	0.043**	−1.143	−0.065	0.059***
	(0.021)	(0.960)	(0.061)	(0.018)
Panel B: 2SLS				
ln(Airbnb density)	0.191***	−3.673	−0.081	0.052
	(0.045)	(2.421)	(0.134)	(0.041)
First-stage F		53.9		
Area (PC4) FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Observations (PC4-month)	2,447	2,447	2,447	2,447

Notes: The table presents OLS (Panel A) and two-stage least squares (Panel B) estimates of the effect of Airbnb density on mortgage outcomes. The unit of observation is the PC4-month (Amsterdam, 2015m8–2019m9); regressions are weighted by the number of originations, and standard errors are clustered at the PC4 level. The endogenous variable is log Airbnb density (active entire-home listings per 1,000 housing units), the excluded instrument is $\text{Touristiness}_j \times \text{AirbnbGrowth}_t$, and coefficients are elasticities. The 2SLS is implemented with PC4 and year-month dummies. The instrument's first-stage coefficient is +0.048 (SE 0.007); a shift-share instrument is identified only up to sign and is oriented here so that the first stage is positive, which leaves the 2SLS estimates unchanged. The number of observations in the first stage is the number of PC4-month cells with at least one origination and a positive number of listings; the dependent variables are PC4-month averages of the origination-level outcomes. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G.4.2. Prepayment and Arrears

The results so far concern new mortgage originations. As a final exercise, we examine whether the regulation affected the repayment behavior of existing borrowers, using the monthly loan performance panel of the four largest lenders. We construct prepayment by differencing the outstanding balance of each loan: the smallest regular monthly decline is the scheduled payment, and any excess decline is classified as a prepayment; a borrower's annual prepayment is voluntary when it totals at least €10,000. Arrears are observed directly. We collapse the data to a borrower-by-year panel and estimate equation (59) on the Amsterdam sample (2015–2018), with Post equal to one from 2017 onward, controlling for the log current home value and the begin-of-year LTV, with area and lender-by-year fixed effects.

We expect little effect on this margin: the regulation operated through the prices facing new buyers, while the repayment incentives of existing borrowers, such as interest rates and income, were untouched. Table G.12 confirms this. In the estimation sample, 10.4% of borrower-years record a voluntary prepayment and 4.1% record arrears, and the $\text{Exposure}_j \times \text{Post}_t$ coefficients are economically negligible and statistically

indistinguishable from zero for the prepayment probability, the prepayment amount, the arrears probability, and the arrears amount (Panel A). The instrumental-variables design leads to the same conclusion: aggregating the panel to PC4-year cells and instrumenting log annual Airbnb density with the Bartik instrument, the first stage remains strong at the annual frequency ($F = 74.0$), and the OLS and 2SLS coefficients are small and insignificant for all four outcomes (Panels B and C). The absence of a response on the repayment margin is consistent with the narrow footprint documented above: the regulation moved the size of new loans, not the behavior of existing borrowers.

Table G.12. Prepayment and Arrears of Existing Borrowers

	(1) Pr(Prepay)	(2) ln(Prepayment)	(3) Pr(Arrears)	(4) ln(Arrears)
Panel A: Difference-in-differences (borrower-year level)				
$\text{Exposure}_j \times \text{Post}_t$	−0.0020	−0.0198	0.0009	0.0028
	(0.0036)	(0.0367)	(0.0025)	(0.0159)
Observations (borrower-year)	45,280	45,280	45,280	45,280
Panel B: OLS (PC4-year level)				
ln(Airbnb density)	0.0066	0.0700	0.0139	0.1104*
	(0.0112)	(0.1112)	(0.0089)	(0.0557)
Panel C: 2SLS (PC4-year level)				
ln(Airbnb density)	0.0133	0.1249	0.0063	0.0412
	(0.0185)	(0.1878)	(0.0135)	(0.0862)
First-stage F		74.0		
Observations (PC4-year)	251	251	251	251

Notes: The table presents estimates of the effect of the regulation on the repayment behavior of existing borrowers. The sample is a borrower-by-year panel constructed from the monthly loan performance data of the four largest lenders (Amsterdam, 2015–2018, loans originated after 2013, no NHG). Prepayment is constructed by differencing the loan's outstanding balance: the smallest regular monthly decline is the scheduled payment, the excess decline is a prepayment, and a year's prepayment is voluntary when it totals at least €10,000; arrears are amounts in arrears. In Panel A each column reports the coefficient on $\text{Exposure}_j \times \text{Post}_t$, where Post equals one from 2017 onward because the panel is annual; the regressions control for the log current home value and the begin-of-year LTV and include area (PC4) and lender×year fixed effects. In Panels B and C the panel is aggregated to PC4-year cells, weighted by the number of borrowers; the dependent variables are cell means, the regressor is log annual Airbnb density (instrumented in Panel C by $\text{Touristiness}_j \times \text{AirbnbGrowth}_t$, as in Table G.11), and the regressions include area (PC4) and year fixed effects. Standard errors are clustered at the PC4 level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G.4.3. Summary

Applying our empirical design to the mortgage market, we find that the regulation left a measurable but narrow footprint: loan amounts fell where Airbnb exposure was highest, consistent with the house price declines and with the IV elasticity of loan size with respect to Airbnb density, while leverage ratios and interest rates were essentially unchanged. Consistent with this interpretation, the prepayment and arrears behavior of existing borrowers shows no response.

H. Estimation and Computational Details

H.1. Kernel Smoothing

With 300 household types, approximately 115 feasible alternatives per type, and 5 years of data, the raw cell counts for some type-alternative-year combinations can be small, generating noisy choice probability estimates. Following Bayer et al. (2016), we address this by computing kernel-smoothed choice probabilities that borrow information from households with similar post-move characteristics.

For each target type $\bar{\tau}$ with characteristics $\bar{Z}_{\bar{\tau}}$, the smoothed unconditional probability of choosing alternative k in year t over the full mover population (including those who exit to the outside option) is:

$$\widehat{\Pr}_{kt}^{\bar{\tau}} = \frac{\sum_{i=1}^{N_t} \mathbf{1}[d_{it} = k] \cdot W^{\bar{\tau}}(\bar{Z}_{it})}{\sum_{i=1}^{N_t} W^{\bar{\tau}}(\bar{Z}_{it})}, \tag{60}$$

where the sums run over all N_t movers (including entrants and exiters) in year t, and k ranges over the full choice set $\{0\} \cup \mathcal{K}_t^{\bar{\tau}}$ including the outside option. The kernel weight assigned to household i with post-move characteristics $\bar{Z}_{it} = (Y_\tau, \bar{A}_{it}, \text{age}_\tau, \text{fam}_\tau)$ (as defined in equation (8)) is the product of $L = 4$ Gaussian kernels (Bayer et al., 2016, eq. (22)):

$$W^{\bar{\tau}}(\bar{Z}_{it}) = \prod_{l=1}^{L} \frac{1}{b_l^{\bar{\tau}}} \varphi\left(\frac{\bar{Z}_{it}(l) - \bar{Z}_{\bar{\tau}}(l)}{b_l^{\bar{\tau}}}\right), \tag{61}$$

where $\varphi(\cdot)$ is the standard normal density, $\bar{Z}(l)$ is the lth element of the household characteristic vector (log income, log wealth, age, family size), and $b_l^{\bar{\tau}}$ is the bandwidth for type $\bar{\tau}$ in dimension l. Each characteristic is standardized to unit variance before applying the kernel, so bandwidths are measured in standard-deviation units.

Our implementation computes unconditional CCPs over the full choice set $\{0\} \cup \mathcal{K}_t^{\bar{\tau}}$, including the outside option, and feeds these directly into the mean-centered inversion (31). Appendix H.2 shows that this choice is without loss of generality: all three normalization approaches in the literature recover identical normalized value functions for inside alternatives. We include the outside option in the smoothing because it avoids a separate exit-probability estimation step and because the mean-centering distributes the influence of exit rates across all alternatives, reducing sensitivity to noisy exit probabilities for types with very high or very low exit rates.

The Gaussian kernel ensures strictly positive weight from every household in the sample, eliminating zero-cell problems that would cause the log-inversion to fail. Households with characteristics close to the target type receive high weight; those far away receive negligible but nonzero weight.

Bandwidths $b_l^{\bar{\tau}}$ are chosen by leave-one-type-out cross-validation, maximizing the log-likelihood of mover destination choice. We allow dimension- and type-specific bandwidths. As a robustness check, we verify that preference parameter estimates are stable under Epanechnikov kernels (which have compact support) and under bandwidths scaled by factors of 0.75 and 1.25.

H.2. Hotz and Miller Inversion

This subsection derives the mean-centered inversion in equation (31). Under the T1EV distributional assumption, the conditional choice probability for a type-$\bar{\tau}$ mover choosing alternative k from feasible set $\{0\} \cup \mathcal{K}_t^{\bar{\tau}}$ is:

$$\Pr_{kt}^{\bar{\tau}} = \frac{\exp(v_{kt}^{\bar{\tau}})}{\sum_{k' \in \{0\} \cup \mathcal{K}_t^{\bar{\tau}}} \exp(v_{k't}^{\bar{\tau}})}, \tag{62}$$

where $v_{kt}^{\bar{\tau}} \equiv u_{kt}^{\bar{\tau}} + \beta \bar{V}_{t+1}^{\bar{\tau}}(k)$ is the full choice-specific value from equation (15). Note that $v_{kt}^{\bar{\tau}}$ and the normalized $\tilde{v}_{kt}^{\bar{\tau}} = v_{kt}^{\bar{\tau}} - m_t^{\bar{\tau}}$ generate identical choice probabilities, since $m_t^{\bar{\tau}}$ cancels in numerator and denominator. The derivation below shows how this cancellation makes $m_t^{\bar{\tau}}$ recoverable from the stay/move margin (Step 1b) rather than from destination choices alone. Taking logs:

$$\ln \Pr_{kt}^{\bar{\tau}} = v_{kt}^{\bar{\tau}} - \ln \sum_{k'} \exp(v_{k't}^{\bar{\tau}}). \tag{63}$$

The second term on the right-hand side is the log-sum-exp, which is common across all alternatives k for a given $(\bar{\tau}, t)$. Averaging equation (63) over all $K_t^{\bar{\tau}} \equiv |\{0\} \cup \mathcal{K}_t^{\bar{\tau}}|$ alternatives:

$$\frac{1}{K_t^{\bar{\tau}}} \sum_{k'} \ln \Pr_{k't}^{\bar{\tau}} = \frac{1}{K_t^{\bar{\tau}}} \sum_{k'} v_{k't}^{\bar{\tau}} - \ln \sum_{k'} \exp(v_{k't}^{\bar{\tau}}). \tag{64}$$

Subtracting from equation (63):

$$\ln \Pr_{kt}^{\bar{\tau}} - \frac{1}{K_t^{\bar{\tau}}} \sum_{k'} \ln \Pr_{k't}^{\bar{\tau}} = v_{kt}^{\bar{\tau}} - \frac{1}{K_t^{\bar{\tau}}} \sum_{k'} v_{k't}^{\bar{\tau}}. \tag{65}$$

Defining the normalized value function $\tilde{v}_{kt}^{\bar{\tau}} \equiv v_{kt}^{\bar{\tau}} - m_t^{\bar{\tau}}$ where $m_t^{\bar{\tau}} \equiv (1/K_t^{\bar{\tau}}) \sum_{k'} v_{k't}^{\bar{\tau}}$:

$$\widehat{\tilde{v}}_{kt}^{\bar{\tau}} = \ln \widehat{\Pr}_{kt}^{\bar{\tau}} - \frac{1}{K_t^{\bar{\tau}}} \sum_{k'} \ln \widehat{\Pr}_{k't}^{\bar{\tau}}, \tag{66}$$

which is equation (31). The key property is that the normalization constant $m_t^{\bar{\tau}}$ cancels in the differencing, so the inversion recovers value functions up to an additive type-time-specific constant. This constant is subsequently identified in Step 1b from the stay/move decision.

The mean-centered inversion used here computes CCPs over the full choice set $\{0\} \cup \mathcal{K}_t^{\bar{\tau}}$ including the outside option, yielding value functions identified up to the type-time constant $m_t^{\bar{\tau}}$. Two alternatives appear in the literature. The standard Hotz and Miller representation normalizes relative to the outside option directly: $\ln \Pr_{kt}^{\bar{\tau}} - \ln \Pr_{0t}^{\bar{\tau}} = v_{kt}^{\bar{\tau}} - v_{0t}^{\bar{\tau}}$. Bayer et al. (2016) condition on inside-option movers when estimating destination choice probabilities, handling the exit probability through a separate likelihood component. All three approaches recover normalized value functions for inside alternatives that are identical up to an additive constant; this constant differs across normalizations but is absorbed by the type-time constant $m_t^{\bar{\tau}}$ identified in Step 1b from the stay/move margin.

H.3. Extra Instruments Investor Supply

For overidentified specifications of the investor supply econometric model, we supplement with Z_{wt}^{Reg} as a second instrument for $\widehat{\text{IV}}_{wt}$ and a demographic instrument for p_{wt}^{O}:

$$Z_{wt}^{\text{Demo}} = \Delta\text{Pop}_{n,t} \times \text{OwnerShare}_{w,2015}, \tag{67}$$

where $\Delta\text{Pop}_{n,t}$ is area-level population growth and $\text{OwnerShare}_{w,2015}$ is the pre-period ownership share. The exclusion restriction requires that population growth does not directly shift investor rental preferences conditional on prices and fixed effects. This is the most vulnerable instrument: population growth may reflect neighborhood attractiveness trends correlated with η_{wt}, so we treat the just-identified specification as preferred and use overidentified specifications as robustness checks.

I. Data Appendix

This appendix documents the construction of three data objects used in the reduced-form analysis: the touristiness index that enters the Bartik instrument, the post-LASSO hedonic price index that defines the main house-price outcome, and the two-stage machine-learning imputation that produces the district-year rent panel.

I.1. Touristiness Index Construction

Table I.1 reports the PCA loadings and component statistics for the touristiness index constructed in Section 2.3. The first principal component explains 63.7% of total variance, with all four inputs loading positively: horeca establishments (restaurants, cafés, and bars) most heavily (0.59), followed by hotel beds (0.53), centrality (0.51), and cultural facilities (0.33). The moderate correlation between the index and log Airbnb density ($r = 0.53$, $R^2 = 0.28$) reflects that Airbnb location depends on factors beyond traditional tourism infrastructure, particularly housing stock characteristics. This partial overlap satisfies both conditions for a valid instrument: the index is sufficiently correlated with Airbnb penetration to yield a strong first stage, and the correlation operates through predetermined 2015 tourism infrastructure rather than contemporaneous Amsterdam-specific housing demand shocks.

The component-level correlations with Airbnb density are heterogeneous. Cultural facilities carry the highest correlation with Airbnb (0.80) despite loading least heavily on the first principal component. This reflects that cultural facilities are more tightly spatially concentrated than hotels or horeca establishments, making them strong predictors of the specific canal-district areas where Airbnb penetrated most deeply. All components are measured in 2015, prior to significant Airbnb-driven reallocation, making reverse causality from platform activity to tourism infrastructure implausible within the sample window. Hotel beds, by contrast, correlate only 0.24 with Airbnb, consistent with hotels and short-term rentals serving partially overlapping but distinct market segments.

I.2. Hedonic Price Index Construction

We construct quality-adjusted house price residuals using a two-step procedure. In the first step, we select hedonic controls from a large candidate set of property-level characteristics using the post-double-selection LASSO estimator of Belloni et al. (2014),

Table I.1. Touristiness Index: Components and Construction

	Mean	SD	PC1 Loading	Corr. w/ Index	Corr. w/ Airbnb
Hotel beds	624	1,265	0.53	0.85	0.24
Cultural facilities	117	108	0.33	0.53	0.80
Horeca establishments	51	70	0.59	0.94	0.55
Centrality (inv. distance)	—	—	0.51	0.81	0.26
Touristiness index	0.00	1.00	—	1.00	0.53
PCA diagnostics					
First PC eigenvalue	2.55 (63.7% of variance)				
Second PC eigenvalue	0.94 (23.4% of variance)				
Predictive power					
Coefficient (log Airbnb $\sim$ Index)	0.46				
R^2	0.28				

Notes: Components measured in 2015 across Amsterdam districts. PC1 loading is the weight on each standardized component in the first principal component. Corr. w/ Index is correlation between each component and the touristiness index. Corr. w/ Airbnb is correlation with log listings per 1,000 housing units in the pre-regulation period. Centrality computed as inverse distance to Dam Square; raw values not reported as units are not interpretable.

which guards against omitted variable bias in the hedonic regression by selecting controls associated with both the outcome (log price) and the treatment (Airbnb exposure). The candidate covariate set includes continuous variables (log living area, rooms, building age, and their squares and pairwise interactions), categorical indicators (property type, construction period, energy label, housing corporation ownership), and dwelling attributes (bathrooms, garden, garage, floor level for apartments). No district or time fixed effects enter the hedonic; location and time variation are absorbed entirely by the fixed effects in the main regression specifications. The LASSO selects 34 covariates from this candidate set; all selected covariates enter the hedonic regression without penalization.

In the second step, we estimate the hedonic regression by OLS on the full transaction sample (January 2014 to December 2019) and compute residuals as:

$$\hat{u}_{it} = \ln P_{it} - \hat{c} - \hat{X}'_{it}\hat{\gamma},$$

where P_{it} is the transaction price, $\hat{c}$ is the estimated intercept (no district or time fixed effects are included), X_{it} is the vector of selected hedonic controls, and $\hat{\gamma}$ are the OLS coefficients. Because no location fixed effects enter the hedonic, the residuals retain district-level location premia; these are absorbed by the district fixed effects α_j in the main specifications. The residuals are aggregated to the district-month level as transaction-count-weighted means and represent price variation orthogonal to observable dwelling quality. The hedonic R^2 is 0.74, indicating that observable property characteristics explain approximately three-quarters of the variance in log prices. The identifying variation in the main specifications is the within-district, within-time-period component of these residuals, which is the standard difference-in-differences variation.

I.3. Rent Imputation Details

The two-stage imputation described in Section 2.3 achieves near-perfect sector classification (AUC-ROC = 0.9995) and strong rent prediction performance ($R^2 = 0.83$

for the free sector, $R^2 = 0.65$ for the regulated sector), with the 2015 survey year excluded due to limited geographic coverage. This appendix documents the full model performance statistics, calibration diagnostics, and validation figures.

Stage 1: Sector classification. Table I.2 reports classification performance and feature importance for the gradient boosting model that predicts free-sector status. The near-perfect discrimination (AUC-ROC = 0.9995) reflects the near-deterministic role of ownership type: housing corporation properties are regulated by statute regardless of their WWS score, so ownership alone accounts for 81.5% of predictive gain. Figure I.1 displays calibration and discrimination diagnostics. The calibration plot reveals slight underprediction in the intermediate probability range (0.2 to 0.8), but as fewer than 1% of predictions fall in this region, the impact on district-level aggregates is negligible. The discrimination plot confirms that predicted probabilities are strongly bimodal, with mass concentrated near zero for regulated dwellings and near one for free-sector dwellings, consistent with the near-deterministic classification.

Table I.2. Rent Imputation: Model Performance

	Classification	Free-sector Rent	Regulated Rent
Model performance			
AUC-ROC / R^2	0.9995	0.83	0.65
Accuracy / MAPE (%)	99.0	9.4	12.0
F1 score (free-sector) / Brier	0.93 / 0.003	—	—
Top features (importance share %)			
Ownership type	81.5	—	—
Construction year	8.1	—	—
District	4.5	—	—
WOZ value	3.4	—	—
Living area	2.5	—	—
Training sample			
Observations	195,817	40,523	18,176
Free-sector share (housing stock) / Mean rent	3.95%	€1,467	€526

Notes: Performance from 5-fold cross-validation. Classification model predicts free-sector status for full housing stock; the 3.95% free-sector share reflects population prevalence (dominated by social housing), not survey composition. Rent models trained separately on CBS Huurenquête observations by sector. Feature importance measured by total gain in gradient boosting models. These five features account for all non-zero importance; all other features in the candidate set received zero gain.

Stage 2: Rent prediction and validation. Figure I.2 validates the Stage 2 rent predictions and documents their time-series properties. The top-left panel confirms that predicted rents track actual rents closely along the 45-degree line, with dispersion widening at higher values consistent with greater heterogeneity among high-value free-sector units. The top-right panel shows mean imputed rents by sector from 2015 to 2022 (the analysis sample runs through 2019; 2020 to 2022 are shown for out-of-sample context): free-sector rents grew from approximately €1,050 in 2016 to €1,400 by 2022, consistent with Amsterdam's housing boom, while regulated rents tracked administrative ceiling adjustments from approximately €500 to €650 over the same period.

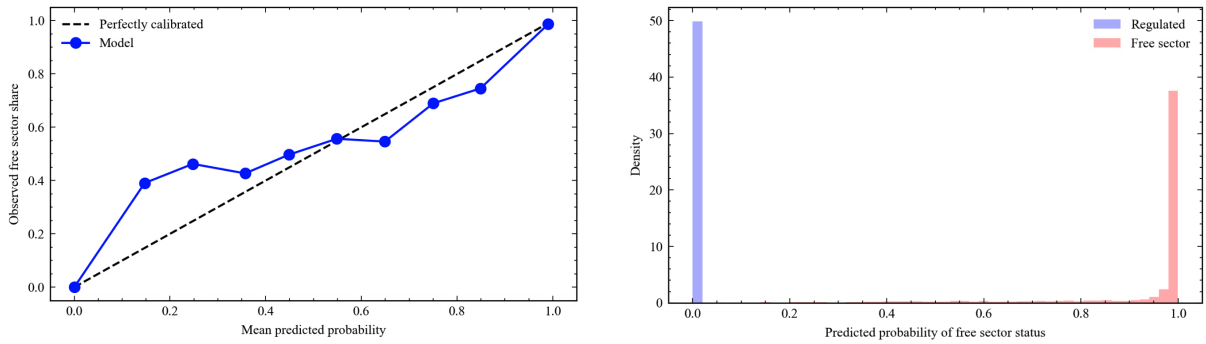


Figure I.1. Free-sector Status Classification: Calibration and Discrimination

Notes: Left panel: calibration plot showing mean predicted probability (x-axis) against observed free-sector share (y-axis) by probability decile; dashed line indicates perfect calibration. The model tracks below the perfect calibration line at intermediate probabilities (0.2 to 0.8), reflecting slight underprediction in that range; in practice over 99% of predictions fall outside this interval. Right panel: density of predicted probabilities separately for regulated (blue) and free-sector (pink) dwellings, showing strong separation with mass concentrated near 0 and 1 respectively. AUC-ROC = 0.9995; Brier score = 0.003.

The anomalously low 2015 regulated rent figure reflects weak model performance in the first survey year ($R^2 = 0.34$) due to limited geographic coverage; we therefore use 2016 as the effective start of the rent analysis, yielding a 2016 to 2018 panel (three years $\times$ 98 districts $\times$ 2 segments $=$ 588 observations, reduced to 586 after dropping two sparse cells). The bottom panel shows the distribution of the free-sector rent premium across comparable unit pairs; the median premium is approximately 150%, with a right-skewed distribution reflecting heterogeneity in location and dwelling quality.

I.4. Post-2019 Ownership Data: Source Change and Coverage

Two coverage details of the BKW transaction register affect the housing-price analysis in Section 3.2. We document them here for completeness; neither materially changes the headline estimates.

Source change in CBS ownership data (from 2019). From the 2019 reference year, CBS switched the data sources behind its annual ownership-stock register, deriving ownership from the Land Registry's cadastral ownership records (*Basisregistratie Kadaster*) and slightly adjusting the underlying methodology. As CBS documents, part of the apparent change in ownership between 2018 and 2019 reflects this source switch rather than real transactions, with the documented direction being fewer dwellings of unknown type, more owner-occupied dwellings, and a slight increase in dwellings classified as housing-corporation. Sale prices and Airbnb density are unaffected: sale prices come from the separate BKW register of recorded sales (actual market transactions with documented prices), not from the cadastral ownership records that feed the annual stock register. Ownership type is, however, a control in our transaction regressions and the dominant feature in the rent-imputation classifier; the 2019 source change could mechanically flip the housing-corporation indicator for a small share of 2019 transactions. Three design properties absorb this. First, the symmetric difference-in-differences window (April 2016 to March 2018) entirely pre-dates the methodology change. Second, year-month fixed effects absorb uniform 2019 shifts in classification, assuming the

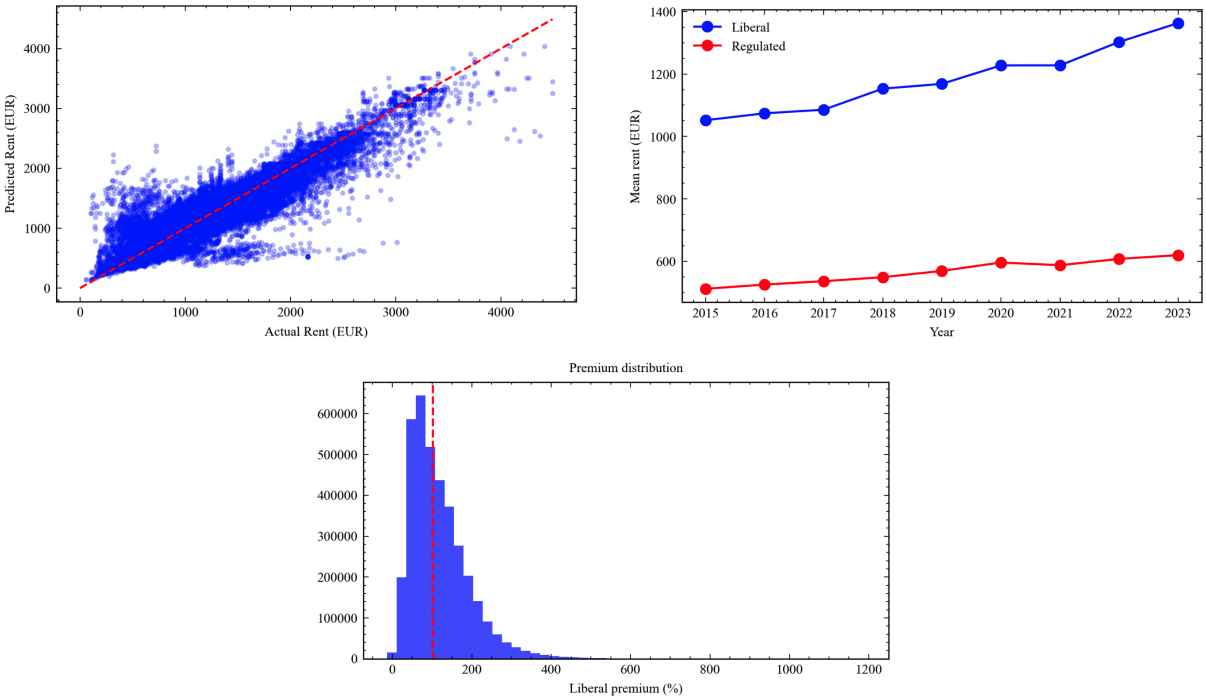


Figure I.2. Rent Imputation: Validation and Trends

Notes: Top left: actual versus predicted rent in the hold-out validation sample; dashed line indicates perfect prediction. Top right: mean imputed rent by sector and year for the full Amsterdam housing stock; the anomalously low 2015 regulated rent reflects weak model performance in the first survey year ($R^2 = 0.34$) due to limited geographic coverage. Data from 2020 onward are shown for out-of-sample context only; the analysis sample runs through 2019. Bottom: distribution of the free-sector rent premium (ratio of free-sector to regulated imputed rent for comparable units), with dashed line at the median of approximately 150%.

methodology change applies symmetrically across Amsterdam districts (CBS documents the change at the national level without specifying its within-municipality spatial pattern). Third, district fixed effects absorb time-invariant district-level classification baselines. The break would confound the elasticity only if the methodology change produced district-specific differential reclassifications correlated with Airbnb exposure in 2019, a pattern the documented source switch does not suggest.

Owner-occupier-only coverage. The CBS BKW register covers transactions of owner-occupied dwellings; sales of dwellings that were investor-held at the time of sale (whether transferred to another investor or to an owner-occupier) are out of coverage. We expect this to bias the transaction count downward in peripheral districts where investor-held units are concentrated (Section 2), but not the price estimates, since the price effects we identify are based on transactions of owner-occupied units. We therefore treat the coverage gap as a sample limitation rather than as a bias on the headline estimates.

J. Additional Tables and Figures

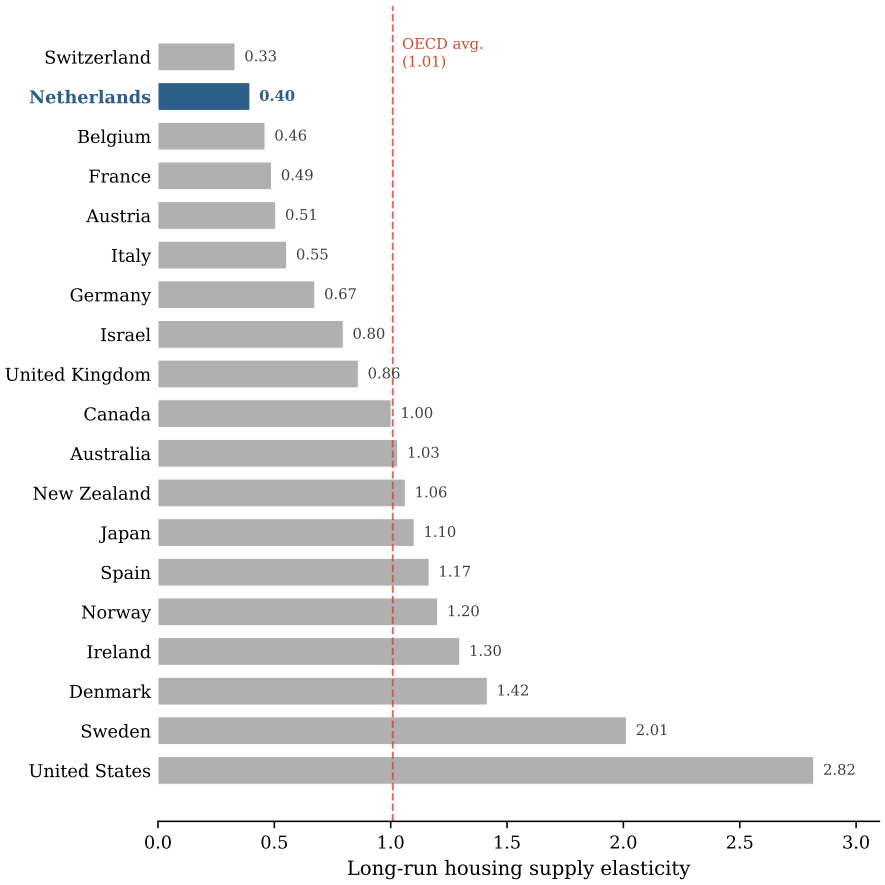


Figure J.1. Long-Run Housing Supply Elasticities: OECD Countries

Notes: Long-run housing supply elasticities for 19 OECD countries. The Netherlands (0.40) is among the lowest, well below the OECD average of 1.01. Source: Cavalleri et al. (2019).

Table J.1. Data Sources Overview

Data	Source	Coverage	Frequency
Airbnb listings	AirDNA	2015 to 2019	Monthly
Housing sales (universe, price & area)	BKW (CBS)	2014 to 2019	Transaction-level
Housing sales (characteristics)	NVM	2014 to 2019	Transaction-level
Rental prices	Huurenquête (CBS)	2015 to 2019	Annual
District characteristics	CBS / Gemeente Amsterdam	2015 to 2016	Cross-section
Household address spells	Gbaadresobjectbus (CBS)	2015 to 2019	Spell (Jan. 1 snapshots)
Household demographics, income, wealth	Gbapersoontab, Inhatab, Vehtab (CBS)	2015 to 2019	Annual
Dwelling stock and characteristics	Levcyclwoonnietwoonbus, BAG (CBS)	2015 to 2019	Annual
Tenure and ownership category	Eigendomtab, Eigenaartab (CBS)	2015 to 2019	Annual
Geographic crosswalk	Vslgwbtab (CBS)	2015 to 2019	Annual
Tourism statistics (visitor nights)	OIS	2015 to 2019	Monthly

Notes: AirDNA is a commercial provider of short-term rental market data. BKW (*Bestaande Koopwoningen*) is a CBS administrative register covering the universe of ownership transactions; it is used alongside NVM for its universe coverage, though unit-level characteristics are not available (see Section 2.3 for how BKW and NVM are combined). NVM covers approximately 85% of transactions with full unit-level characteristics (construction year, property type, floor area, number of rooms). Huurenquête is the CBS annual rental dwelling survey. District characteristics include demographics, income, housing stock composition, tenure shares, residential mobility, and tourism infrastructure (hotel beds, cultural facilities, hospitality density, inverse distance to Dam Square); all variables are measured at the district level from CBS Kerncijfers Wijken en Buurten and the municipality of Amsterdam, with tourism infrastructure components measured in 2015 to ensure predetermination for the Bartik instrument. The household and housing-stock CBS registers are accessed via the CBS remote access environment under standard microdata protocols; registers are linked at the dwelling level via BAG object identifiers (*Rinobjectnummer*) and at the person level via pseudonymized identifiers (*Rinpersoon*). OIS is the Amsterdam Office of Research and Statistics (*Onderzoek & Statistiek*); its tourism statistics and area characteristics are publicly available.

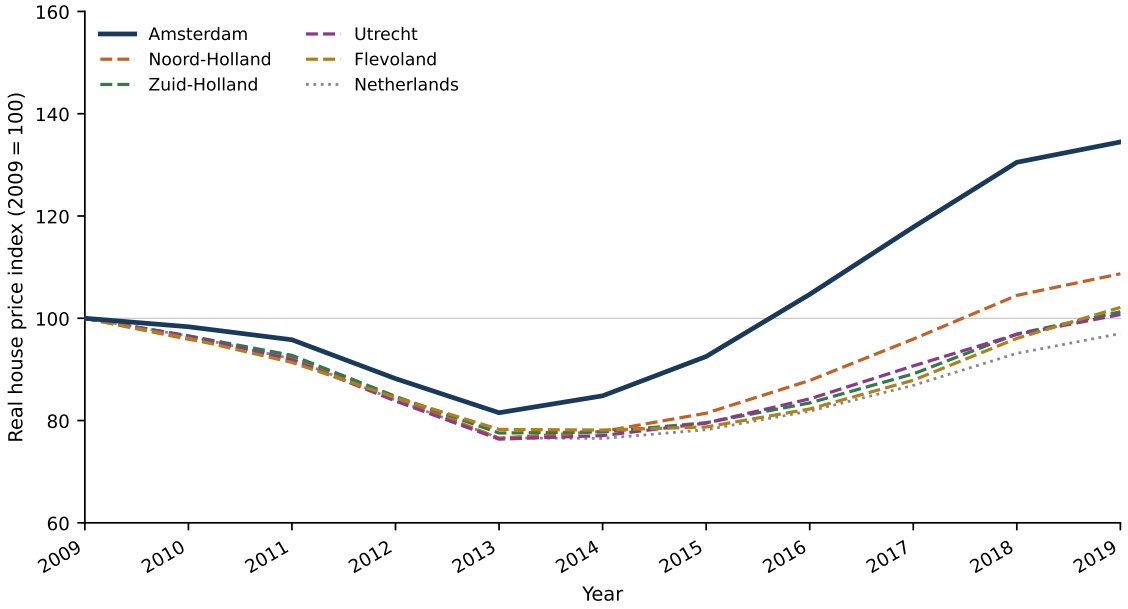


Figure J.2. Real House Price Index by Region, 2009 to 2019

Notes: Real house price index (2009 = 100) for Amsterdam and five comparison regions. Amsterdam appreciates faster than all comparisons from 2015 onward, consistent with supply constraints amplifying demand-driven price growth. Source: NVM transaction registry.

Table J.2. Product aggregation: within-cell price dispersion

	Ownership prices	LTR rents
R^2 (area × size × year FE)	0.755	0.507
Median within-cell CV	0.261	0.246
90th pctl within-cell CV	0.383	0.324
Number of cells	345	345
Number of observations	47,002	1,201,262

Notes: R^2 from OLS regression of individual transaction prices (BKW) or dwelling-level imputed rents on area × size × year fixed effects. Ownership observations are arm's-length transactions; rent observations are imputed liberal rents for the full non-social dwelling stock. CV = coefficient of variation (standard deviation divided by mean) within each cell.

Table J.3. Summary of Main Estimates

Outcome	Method	Estimate	SE
Effect of Regulation on House Prices (DiD)			
Hedonic residual	Exposure × Post	-0.011^{*}	(0.006)
Hedonic residual per m^2	Exposure × Post	-0.013^{**}	(0.005)
$\ln(\text{Price}/m^2)$	Exposure × Post	-0.013^{**}	(0.006)
Effect of Airbnb on House Prices			
Panel A: OLS			
Hedonic residual	OLS	0.040^{***}	(0.015)
Panel B: 2SLS			
Hedonic residual	2SLS	0.313^{***}	(0.091)
Hedonic residual per m^2	2SLS	0.335^{***}	(0.096)
$\ln(\text{Price}/m^2)$	2SLS	0.309^{***}	(0.080)
Effect of Airbnb on Rents (IV)			
Free sector	2SLS	0.073^{**}	(0.032)
Regulated sector	2SLS	-0.018	(0.036)

Notes: DiD estimates are coefficients on Exposure × Post from equation (2). House price IV uses log Airbnb per 1,000 housing units; rent IV uses log Airbnb listings. Full specifications and robustness checks in Appendix G. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.