Market Power and the Welfare Effects of Institutional Landlords*

Felipe Barbieri† Gregory Dobbels‡

Job Market Paper

March 30, 2026

Click here to download the most recent version

Abstract

In the last decade, large financial institutions in the United States have purchased hundreds of thousands of homes and converted them to rentals. This paper studies the welfare consequences of institutional ownership of single-family housing. We build an equilibrium model of the housing market with two sectors: rental and homeownership. The model captures two key forces from institutional purchases of homes: changes in rental concentration and reallocation of housing stock across sectors. To estimate the model, we construct a novel dataset of individual homes in metropolitan Atlanta, identifying institutional owners of each house and collecting house-level daily prices, rents, vacancies, web page views, and customer contacts from Zillow. Overall, we find that institutional acquisitions decrease rents and increase rental transactions, leading to large welfare gains for renters. This net benefit reflects two opposing forces: while higher concentration raises rents, higher rental supply lowers rents enough to more than offset the effect of concentration, pushing rents down overall. These renter gains come at the expense of homebuyers, whose welfare falls. On the supply side, institutional acquisitions benefit house sellers but harm the average landlord.

JEL classification: L1, L85, R21, R31, R32, R52.

* We would especially like to thank Aviv Nevo, Juan Camilo Castillo and Gilles Duranton for their invaluable advice and support. We are also grateful to Juan Pablo Atal, Sophie Calder-Wang, Ulrich Doraszelski, Fernando Ferreira, Shresth Garg, Joe Gyourko, Benjamin Keys, Stephen Redding, Todd Sinai, and participants from the Wharton Urban and Real Estate Seminar, University of Pennsylvania Industrial Organization Seminar and Empirical Microeconomics Lunch for their helpful comments. This research is supported by the Zell and Lurie Real Estate Center through a grant to the Wharton School of the University of Pennsylvania. Dr. Dobbels was not involved in the data collection for this paper, and no U.S. Government equipment or resources were used in the preparation of this paper. The views expressed are those of the authors and do not purport to reflect the views of the Department of Justice.

† Tuck School of Business at Dartmouth, Email: felipe.barbieri@dartmouth.edu

‡ U.S. Department of Justice, Email: gregory.dobbels@gmail.com

1 Introduction

Until 2011, the U.S. single-family rental housing industry was almost exclusively composed of small landlords.[1] Within a decade, large private equity-backed firms and Real Estate Investment Trusts acquired hundreds of thousands of single-family homes and began renting them out at scale.[2] The rapid growth of these institutional landlords' portfolios raised two concerns. First, their increasing footprint in some rental markets prompted questions about their ability to raise rents. Second, there are worries that the scale of their home acquisitions could reduce the supply of homes for sale, driving up prices and outbidding individual homebuyers. This led to significant press coverage, congressional hearings,[3] and policy proposals such as state[4] and federal[5] legislation targeting institutional landlords.

Despite these concerns about institutional landlords, their economic impact is complex. When institutional investors purchase homes and convert them into rentals, several forces can affect welfare. On the one hand, there is a transfer of houses from the homeownership sector to the rental sector. This increases the rental stock and decreases the homeownership stock, pushing rents down and sales prices up. On the other hand, rental concentration increases, leading to two effects: it raises rents due to increased market power, but it also creates economies of scale, reducing vacancy and occupancy marginal costs. The net effect of these two cost efficiencies on rental prices is itself ambiguous: while lowering occupancy costs reduces rents, lowering vacancy costs increases rents.

In this paper, we measure the welfare effects of institutional ownership of single-family homes. We build an equilibrium model of the housing market with two sectors, rental and homeownership. To estimate the model, we gather new house-level data from Zillow and combine it with data identifying the institutional owner of each house. We then run counterfactuals to quantify the impact of institutional acquisitions on the welfare of renters, homebuyers, home sellers and landlords.

The model is composed of demand, supply, and a matching technology. On the demand side, potential renters and homebuyers make a discrete choice between a set of dif-

[1] In 2011, no owner had more than 1,000 properties in the U.S. Source: U.S. GAO.
[2] As of 2023, U.S. institutional landlords owned 488,000 single-family homes. Source: Urban Institute.
[3] See "How Institutional Landlords are Changing the Housing Market" (Senate Banking Committee).
[4] See California (S.B. 1212), Minnesota (H.F. 685), Nebraska (L.B. 1405), and North Carolina (H.B. 114).
[5] Three bills (H.R.9246, S.5151, S.2224) have been introduced in the U.S. Congress.

ferentiated housing listings advertised for sale or for rent, taking into account the listed price and characteristics of each house. On the supply side, landlords and house sellers compete by setting listing prices in a Nash-Bertrand game, trading off price and *time-on-market*, defined as the time needed to rent or sell a house. Finally, there is a matching technology that bridges demand and supply, serving two purposes. First, it functions as a rationing mechanism: multiple consumers may favor the same house, but each house can only be matched to a single consumer. Second, it translates demand for a listed house into an expected time-on-market, which plays the role of quantity in our model. Indeed, for each house, a higher time-on-market implies a lower occupancy rate, which implies a lower quantity of *house-months* consumed.

To account for the effect of concentration on prices in the rental sector, we make several assumptions regarding landlords' cost structures and competitive behavior. First, landlords have two marginal costs: a marginal cost of occupancy and a marginal cost of vacancy. Reductions in marginal costs have opposing price effects depending on which cost is reduced: lowering occupancy costs pushes landlords to reduce rents, while lowering vacancy costs instead encourages them to *increase* rents. Second, there are two types of landlords: institutional and non-institutional. Although all landlords may own one or more houses, we assume only institutional landlords can jointly price all the houses they own in the market. This ensures we do not overstate the degree of market power of non-institutional landlords.

Multi-product pricing is a key source of institutional landlords' market power: by internalizing how price changes for one listing impact demand for others, institutional landlords are incentivized to raise rents across their portfolio. Observing both rents and ownership at the house level is therefore crucial for accurately measuring the degree of market power of multi-product landlords - an analysis not possible with regional price averages or traditional housing surveys.

To estimate the model, we therefore collect and combine two novel datasets on single-family homes, allowing us to track prices, rents, demand, and ownership at the level of each housing unit. First, we scrape new data from Zillow for single-family houses in metropolitan Atlanta, the urban area with the highest presence of institutional landlords in the United States. For all houses listed on Zillow between May 2023 and September 2024, we observe daily changes in prices, rents, customer web page views, contacts received by landlords and the identity of the manager of the online listing.[6] We also gather data on rent and price histories dating back to 2010 for the universe of all Atlanta parcels. Second, we assemble a dataset identifying the institutional owner of every house in Atlanta. This is a challenging task, since the identity of the institutional owner of a house is hidden through subsidiaries. We link these subsidiary companies listed as owners of houses in the tax data to their parent institutional companies, allowing us to map the entire portfolio of homes owned by institutional landlords. We compare the number of houses we attribute to each institutional landlord in our dataset to S.E.C. disclosures of publicly traded landlords, finding only small discrepancies ranging from 0.1% to 1.5%.

The model has three key primitives, which we estimate using different sources of variation. The first primitive is the price responsiveness of potential renters and homebuyers. This is estimated using high-frequency data on customer contacts and web page views for online listings, which we use as proxies for customer demand. We make the assumption that each individual contacts or views their favorite housing listing.[7] To capture competition between different houses, we construct market shares for every *individual* house using the number of contacts and views received by each house. To address the endogeneity of rents and house prices, we propose a new identification strategy that exploits discontinuities in contact and page view rates around price changes.

The second model primitive is the parameter of our matching technology, which governs how quickly houses with a given contact or viewership rate are matched to renters or buyers. We estimate this parameter using variation in transaction speed across houses with different contact and viewership rates.

The third set of primitives includes the discount factors of house sellers and the marginal costs of landlords, which we retrieve using the supply side of our model. Specifically, we derive first-order conditions from the profit functions of house sellers and landlords, then plug our demand and matching parameter estimates into these conditions to solve for discount factors and marginal costs. For the rental sector, we can identify landlords' *net* marginal cost of occupancy, which is the difference between the marginal cost of occupancy and the marginal cost of vacancy. We estimate large, negative net marginal costs

[6] We collect these records by scraping online listings on a daily basis. See details in Data Appendix C.

[7] If individuals contact or view multiple housing listings, our model will treat each contact or view as a different individual. As long as this weekly propensity to contact or view is randomly distributed across individuals, our market shares will not overestimate or underestimate true demand for each house.

of occupancy, which implies that vacancy costs outweigh occupancy costs and that landlords have significant *aversion* to vacancy.[8] We find that net marginal costs of occupancy are *higher* for institutional landlords than for non-institutional ones. As a consequence, cost efficiencies from institutional concentration puts *upwards* pressure on rental prices, reinforcing the effect of market power.

To understand the effects of institutional entry, we run counterfactual simulations in which we change the market structure of homeownership and rental sectors. Our first counterfactual quantifies the *overall* welfare impact of institutional acquisitions by simulating a world in which institutional landlords never entered the market. In practice, we simulate a market where each house currently owned by institutional landlords returns to its 2009 sector—rental or homeownership—before institutional acquisitions began. We find that rental prices increase by 2.4% and rental transactions decrease by 20%. Higher rents and reduced rental options lead to a welfare loss for the average renter of $2,856 per year. On the other hand, higher revenues and decreased competition benefit the average landlord, whose annual profits increases by $612. These results show that renters *benefited* from institutional acquisitions, whereas landlords were hurt. In the homeownership sector, the influx of additional houses decreases sales prices by 4.8% ($26,000) and increases transactions by 2.5%. Lower prices and higher housing availability raise average homebuyer welfare by $49,950. Conversely, home sellers' average profits decrease by $22,446. In other words, institutional entry made homebuyers worse off but home sellers better off.

Next, we run a second counterfactual to decompose overall welfare changes in the rental sector into two forces: changes in concentration and reallocation of housing stock across sectors. In practice, we simulate a counterfactual in which homes currently owned by institutional landlords are instead owned by smaller landlords. Thus, none of these homes shifts to the homeownership sector, leaving homeownership and rental housing stocks unaffected relative to the status quo, while still decreasing concentration in the rental sector. This isolates the potential effect of lowering rental concentration.[9] We find that eliminating rental concentration decreases rents by 3.8% and increases rental transac-

[8] This can be rationalized by liquidity constraints for small landlords and performance-driven incentives for institutional landlords.

[9] We replace the net marginal costs of houses previously owned by institutional landlords with those of non-institutional landlords.

tions by 3.1%. With lower rental prices and higher quantity of housing consumed, renters are better off and landlords are worse off: average renter welfare increases by $2,232 per year, whereas average landlord annual profit decreases by $924. Overall, this counterfactual shows that the additional concentration from the entry of institutional landlords hurts renters and benefits landlords. However, taken together with our first counterfactual, these results show that the effects of increased concentration are more than offset by the effect of expanded rental supply.

Our findings highlight the welfare tradeoffs associated with the rise of institutional landlords in the single-family rental market. Although these landlords have enough market power to raise rents, the increase in rental supply from their acquisitions more than offsets this effect, ultimately benefiting the average renter. Other winners from this industry shift are homeowners, who can now sell their houses at higher prices. However, this comes at the expense of prospective homebuyers, as single-family homes available for purchase are fewer and more expensive. This points to an important tension in current regulatory debates: while institutional ownership of single-family homes may have made the "American Dream" of homeownership harder to attain, it has simultaneously made living in a single-family house more affordable and accessible for renters.

Related literature

We contribute to a growing literature on the impact of investors in the housing industry. Like Ater, Elster and Hoffmann (2022) and Francke, Hans, Korevaar and van Bekkum (2023), we study how investors affect the allocation of houses between the rental and the homeownership sectors, as well as equilibrium housing prices and rents. Unlike these two papers, however, we focus on a setting with large investors, where concentration and market power are likely to be of first order. In this regard, we add to a literature documenting the rise of single-family institutional investors in the United States (Mills, Molloy and Zarutskie, 2019; Smith and Liu, 2020; Lambie-Hanson, Li and Slonkosky, 2022; Ganduri, Xiao and Xiao, 2023; Hanson, 2023).

Within this strand of literature, our paper is most similar to Coven (2024), who studies how U.S. institutional investors affect homeownership and neighborhood access. Our paper differs in three main ways. First, we focus specifically on quantifying the relative contribution of rental concentration and changes in rental supply on equilibrium housing prices and quantities. Second, we estimate housing demand and landlords' marginal costs using house-level price and demand data, allowing us to identify key demand elasticities and measure market power of multi-product landlords at a granular level. Third, in addition to price and quantity effects, our model also allows us to measure changes in the welfare of renters, homebuyers, home sellers and landlords.

Our work also relates to a recent literature on the effects of market power in rental housing. Watson and Ziv (2023) study the multi-family rental market in New York City. Gurun, Wu, Xiao and Xiao (2023) analyzes the impact of mergers among single-family institutional landlords on neighborhood-level rents, using a difference-in-differences framework. Calder-Wang and Kim (2024) study the effects of algorithmic pricing in the multifamily industry. Our paper adds to this literature by quantifying the distributional effects of single-family institutional landlords.

Our paper also relates to a large literature[10] on demand estimation (Berry, 1994; Berry, Levinsohn and Pakes, 1995), in particular for housing (Bajari and Benkard, 2005; Bayer, Ferreira and McMillan, 2007; Bajari, Chan, Krueger and Miller, 2013). While this literature traditionally models consumer substitution across neighborhoods, we contribute by quantifying substitution across individual housing units. This allows us to estimate house-level and market-level demand elasticities, both of which are of independent interest given the paucity of estimates in the housing literature (Hanushek and Quigley, 1980). Among this literature, Calder-Wang (2021) and Moszkowski and Stackman (2023) most closely relate to our work. Like Calder-Wang (2021), we also study the distributional effects of reallocating housing capital across different sectors. Our notion of vacancy marginal cost is similar to that in Moszkowski and Stackman (2023), though they use a dynamic model, whereas we employ a static approach.

More generally, we contribute to the literature[11] on estimating market power (Nevo, 2001), particularly in search-and-matching markets. Within this strand of literature, our paper is closest to Azar, Berry and Marinescu (2022), who focus on estimating market power in the labor market. However, while they assume a fixed relationship between job applications and successful hires, we explicitly model and estimate the equivalent matching process in the context of housing. Specifically, we parameterize a matching

[10] See Berry and Haile (2021) for an overview.

[11] See Gandhi and Nevo (2021) for an overview.

function that converts customer contacts into successful housing transactions, following a recent literature on the taxi and ridesharing industry (Fréchette, Lizzeri and Salz, 2019; Buchholz, 2021; Castillo, 2023).

Finally, our study relates to a broader literature on rising concentration and markups over time. De Loecker, Eeckhout and Unger (2020) and Hall (2018) document trends at the aggregate level. Several papers focus on specific industries, such as the cement (Miller, Osborne, Sheu and Sileo, 2023), automobile (Grieco, Murry and Yurukoglu, 2023), and retail (Hortaçsu and Syverson, 2015; Smith and Ocampo, 2021) industries. We add to this literature by studying concentration in the housing industry.

2 The Single-Family Housing Market

2.1 Background and setting

74% of the U.S. population lives in a single-family house.[12] The vast majority of single-family rental units have traditionally been owned by small mom-and-pop landlords.[13] After the Global Financial Crisis, a set of national single-family landlords backed by large investors acquired hundreds of thousands of homes and started renting them out at scale. The largest of these investors now operate tens of thousands of single-family rental homes, concentrated in a select number of metro areas.

We focus on the Atlanta metropolitan area as a useful setting to examine the impact of institutional single-family landlords. Atlanta is often described as the "epicenter"[14] of this home-buying spree by institutional landlords. As a consequence, 25% of the stock of single-family rental houses in Atlanta is now owned by landlords with a portfolio of 1,000 or more homes, making it the metropolitan area with the highest share of institutional presence in the United States.[15] Our setting therefore allows us to establish an upper bound on the local impact of institutional landlords, offering a benchmark that is informative for other metropolitan areas where their presence is less pronounced.

[12] Computed from 2020 ACS 5-year averages.

[13] In the 1996 Property Owners and Managers Survey, for instance, less than 3% of detached single-family rentals owned by individuals or partnerships were reported to belong to owners with 50 or more units.

[14] See "Regional Market at Epicenter of Institutional Homebuying" (Parcl Labs).

[15] Source: U.S. Government Accountability Office

2.2 A novel house-level dataset with institutional portfolios

There are no large-scale, readily available data on rental prices and institutional portfolios at the individual house level in the United States. Therefore, we build a new dataset of single-family homes in the Atlanta metropolitan area. We collect and combine two novel datasets. The first contains sales prices, rents, vacancies, high-frequency demand proxies and listing manager identities. The second one identifies the institutional owners of each house, allowing us to map their portfolios.

Advertised Zillow listings First, we gather data on single-family home listings from the largest online real estate marketplace in the United States, Zillow. We construct this dataset by scraping advertised for sale and for rent online property listings on a daily basis in 2023 and 2024.[16]. For listings for rent, we collect advertised rents, property characteristics, and the number of contacts received by the landlord between May 1st, 2023 and May 1st, 2024. For listings for sale, we collect advertised sales prices, property characteristics, and the number of web page views between April 1st, 2024 and September 1st, 2024. Figure 1 shows the granularity of our data by displaying daily aggregate patterns in supply (number of houses advertised) and demand (number of customer contacts or customer views) for homes in our sample.

We observe the identity of the landlord advertising the rental properties on Zillow. Appendix Figure A11 displays the composition of rental listings by splitting them into houses listed by institutional landlords and non-institutional landlords. We further decompose non-institutional landlords into listings for which owners delegated the advertisement to a rental management company, and listings advertised by the owner herself.

Zillow price histories Second, we compile daily rental prices, sales prices, vacancies, and tenancies dating back to 2010 for all single-family parcels in the Atlanta metropolitan area. We collect these by visiting the Zillow web page corresponding to every Atlanta parcel, and gathering the details of each property's price history. Importantly, Zillow has a page for every parcel in Atlanta, including off-market properties, and each page includes a price history with historical sales and rental prices.[17] Figure 2 below displays

[16] Appendix Figure A7 provides examples of real-time Zillow listings.

[17] The Zillow price history dataset also includes records from Multiple Listing Services (MLS), as well as other websites scraped or owned by Zillow, such as Trulia, Hotpads, and several real estate websites.

(a) Rental sector

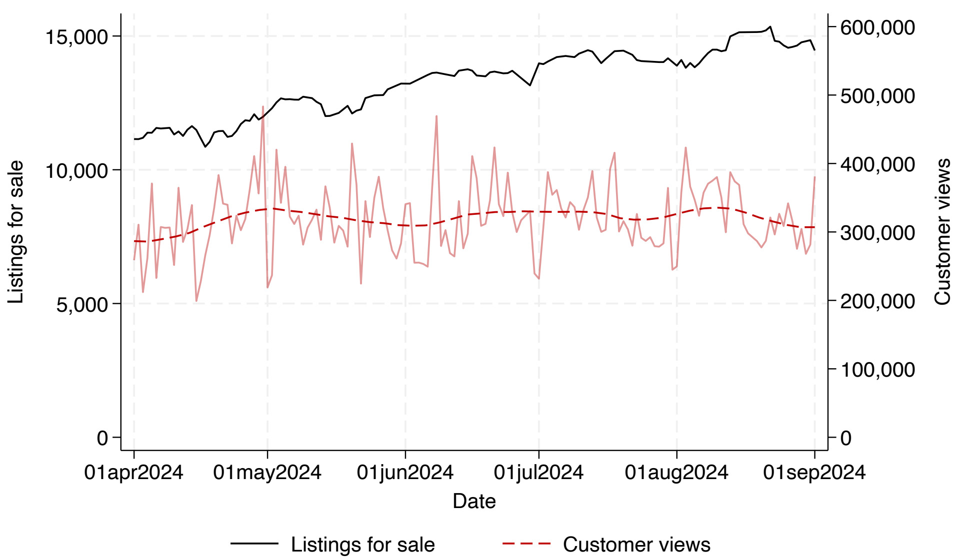

(b) Homeownership sector

Figure 1: Daily patterns in housing supply and demand

Notes: The two figures above plot the daily number of listings for sale and for rent in our sample, as well as the daily number of customer contacts and customer web page views. For rental listings, Zillow displays the number of customer contacts. For sales listings, Zillow does not show the number of contacts; instead, it displays the number of customer views. The black line represents listings. The pale red line shows the high-frequency daily variation in the number of customer contacts and customer views, while the dashed red line represents a smoothed trend of these two time series using Kernel-weighted local polynomial smoothing.

one example of price history in our data for an off-market property. Each listing's start and end dates in the historical data allow us to observe the duration a listing remains on the market (time-on-market). We also observe effective tenancy lengths as the time spent between two successive vacancies. This allows us to compute house-level occupancy and vacancy rates.

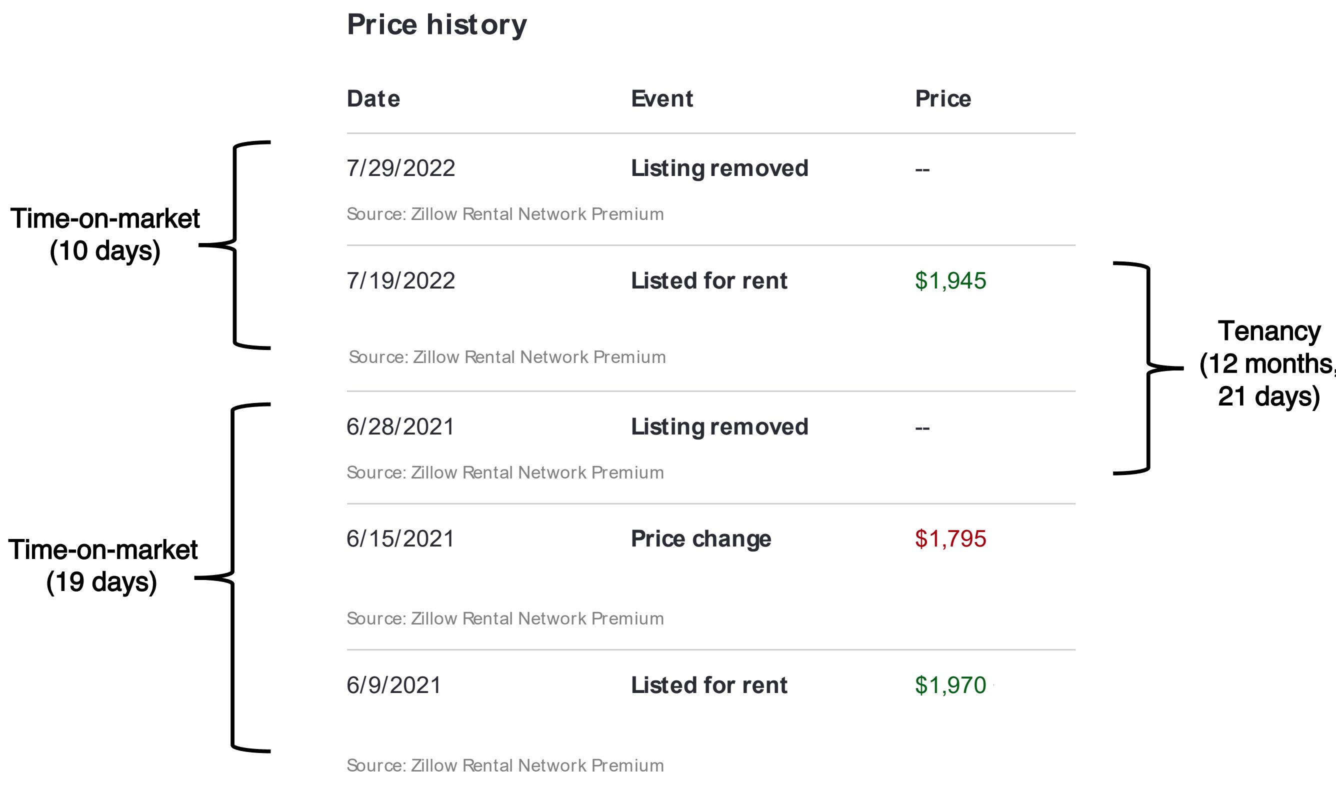
Price history

Date	Event	Price
7/29/2022	**Listing removed**	--
Source: Zillow Rental Network Premium		
7/19/2022	**Listed for rent**	$1,945
Source: Zillow Rental Network Premium		
6/28/2021	**Listing removed**	--
Source: Zillow Rental Network Premium		
6/15/2021	**Price change**	$1,795
Source: Zillow Rental Network Premium		
6/9/2021	**Listed for rent**	$1,970
Source: Zillow Rental Network Premium		

Figure 2: Example of Zillow price history from listing web page

Institutional portfolios Third, we assemble a dataset identifying the institutional owner of every house in Atlanta. This is a challenging task, since the identity of the institutional landlord of each house is hidden through subsidiaries. Figure 3 illustrates this challenge. Tax records only record the name of the *subsidiary* owner, which often corresponds to the name of a Limited Liability Company (LLC) or a Limited Partnership (LP).[18] To identify institutional owners of each house, we map the names of subsidiaries that appear on tax

[18] To shield themselves from liability, institutional landlords silo their holdings into many subsidiary owners instead of holding assets under a single entity.

records to their *parent* company.[19] This allows us to determine the correct portfolios and local market shares of each institutional owner.[20]

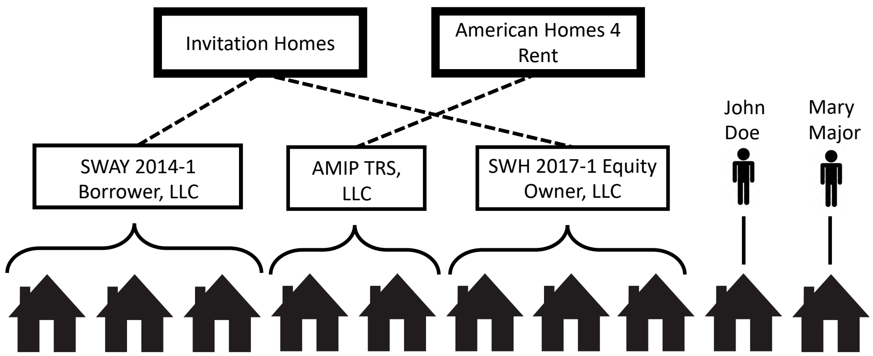

Figure 3: Mapping subsidiaries listed on tax records to institutional owners of each house

We validate our data on institutional portfolios by comparing our home counts with those provided on annual S.E.C. reports of the three publicly-listed investors in our sample for the year 2022. As shown in Figure A1, our numbers align almost exactly with these counts: they represent 98.48%, 100.16% and 99.65% of the portfolios of Invitation Homes, Tricon Residential and American Homes 4 Rent, respectively.[21]

2.3 Key facts on institutional ownership of single-family homes

We show three key facts about institutional ownership of single-family housing in Atlanta. First, we trace the evolution of their portfolio over the last decade, and the rise in rental concentration after 2012. Second, we show that institutional concentration is heterogeneous over space, and large in some market subsegments. Third, we present evidence that institutional landlords expand their portfolios even as the overall number of renter-occupied single-family units falls.

[19] We perform these linkages using a combination of the principal office address of the subsidiary, the name of registered agents and authorizers. Agent names, authorizer names, as well as other information is publicly available on company filings posted on the Georgia's Corporations Division.

[20] We define an "institutional" landlord as a financial institution or corporate entity that owns 1,000 or more rental homes nationally within the United States as of January 1, 2022. While recent legislative bills and industry reports define institutional ownership using thresholds as low as 10, 50, or 100 rental properties, our 1,000-unit cutoff is conservative and isolates the largest national actors.

[21] Slight misalignments could be due to the dates at which these portfolios are measured by different counties. The S.E.C. numbers correspond to the portfolio sizes as of December 31st, 2021 whereas our numbers come from county-level tax records, aggregated by Corelogic.

Fact 1. Institutional ownership of single-family homes drives rental concentration.

We begin by showing trends in institutional ownership of single-family homes in Atlanta, and how the growth of institutional portfolios was the main driver of an increase in rental concentration. Figure 4 shows how concentration has evolved in the single-family rental sector in Atlanta since 2009. We see that before 2012, the largest 10 owners represented a constant 2-3% of all single-family rental homes. However, after institutional landlords began acquiring homes and converting them to rentals, this share started growing fast, reaching about 25% as of 2021.

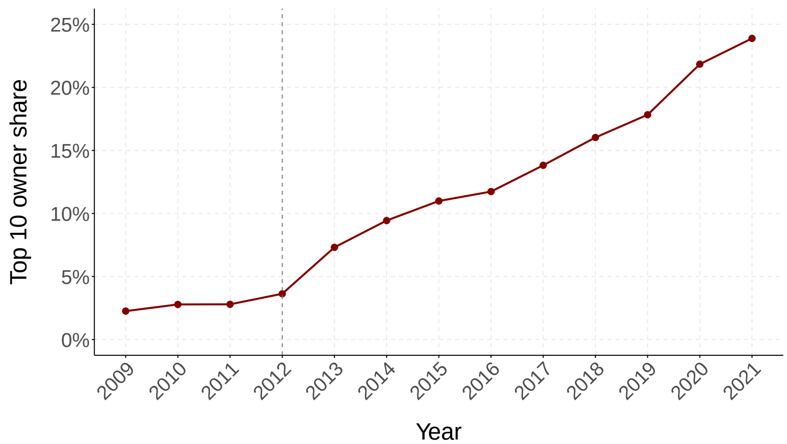

Figure 4: Share of single-family rental homes belonging to top 10 owners, by year

Notes: This figure shows the share of Atlanta's single-family rental homes owned by the 10 largest property owners each year between 2009 and 2021. The numerator of the share is the total number of homes owned by the 10 largest property owners in Atlanta, and is calculated year-by-year using the tax data. The denominator of this share comes from American Community Survey 1-year averages. The dotted vertical line in 2012 represents the first year in which institutional landlords have nonzero holdings in the tax data.

Figure 5 plots the evolution over time of holdings of the 5 largest landlords in Atlanta. Portfolios grew both through consolidation (mergers) and through purchases of individual homes.

Appendix Table Table A1 shows the number of properties owned by the 10 largest institutional landlords in the Atlanta metropolitan area, based on the most recent available tax records from 2022. Of these, 8 have more than 1,000 homes each within Atlanta alone, and the two largest ones have more than 10,000 homes. Together, they represent 25% of the entirety of Atlanta's single-family rental housing stock, which is approximately

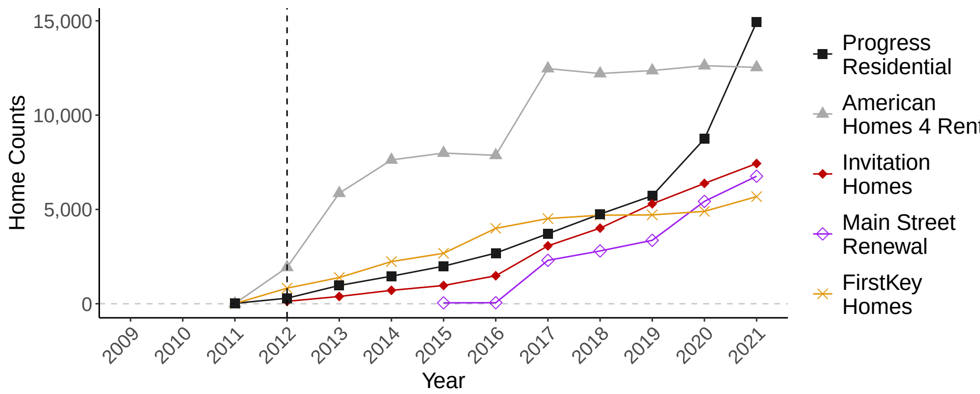

Figure 5: Evolution of Institutional Landlord Portfolios

Notes: This figure plots the number of single-family houses owned by institutional landlords that merged or grew to form the five largest institutional landlords in Atlanta in 2021. In 2015, Starwood-Waypoint acquired Colonial American. In 2016, American Homes 4 Rent (AH4R) acquired American Residential Homes. In 2017, Invitation Homes acquired Starwood-Waypoint. In 2021, Progress Residential's holding company acquired Front Yard Residential. The dotted vertical line in 2012 represents the first year in which institutional landlords have nonzero holdings in the tax data

190,000 as of 2022.[22] However, this share still understates the degree of concentration in rental markets, as it represents overall metropolitan shares and does not distinguish between occupied houses and houses vacant for rent. The following section provides additional insights by exploring heterogeneity in concentration across market subsegments.

Fact 2. Institutional concentration is high in rental market subsegments.

Figure 6 displays the institutional shares of actively listed rental homes at the ZIP Code level. We see substantial spatial heterogeneity in terms of concentration. For a large number of Atlanta ZIP Codes, over half of the advertised listings belong to one of the 10 top institutional landlords.

[22] Source: ACS 1-year average

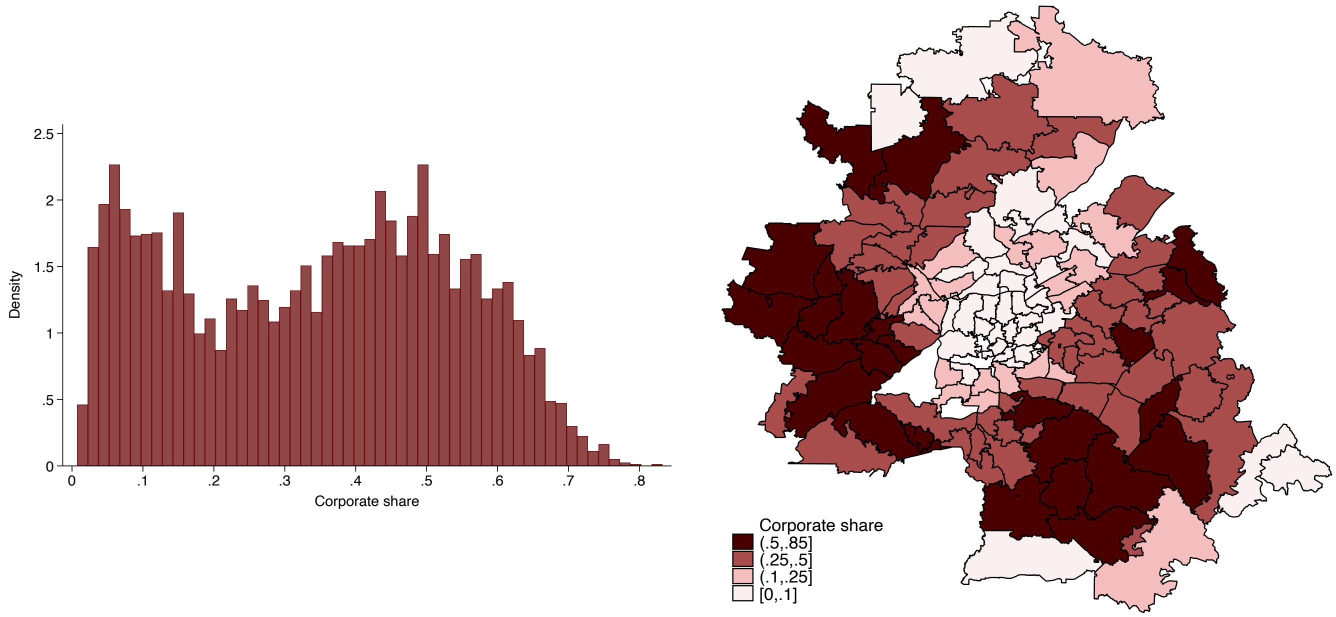

(a) Institutional share by market (ZIP-week)

(b) Institutional share by ZIP Code

Figure 6: Share of institutional rental listings

Notes: These figures show the share of rental listings from institutional landlords. The left panel displays the share across markets (ZIP Code and week couples) for markets with non-zero institutional listings. The right panel plots a map at the ZIP Code level showing the average share across all weeks in our sample for each ZIP Code.

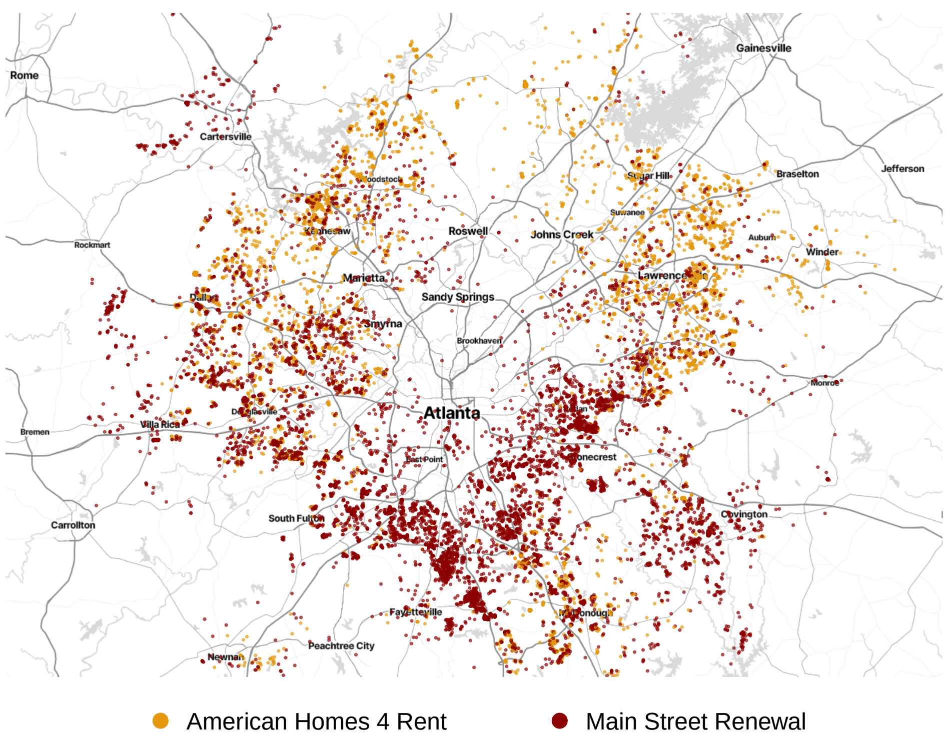

Figure 7: Geographical distribution of houses owned by two institutional landlords

Notes: The map above displays the location of all houses owned by two large landlords in our sample: American Homes 4 Rent (in yellow) and Main Street Renewal (in red).

Next, we show that concentration is high at the market subsegment level because institutional investors focus on certain neighborhoods and certain house types. Our data on institutional ownership provides insight into the spatial distribution of both occupied and vacant homes owned by institutional landlords. We observe the exact latitude and longitude of each individual house. Figure 7 illustrates this by displaying the locations of all homes owned by two institutional landlords: American Homes 4 Rent and Main Street Renewal. The map shows that each of these landlords focuses on different areas of Atlanta's suburbs: while Main Street Renewal concentrates on the Southern part of the metropolitan area, American Homes 4 Rent primarily targets the North East and North West suburbs.

Institutional landlords target "starter" homes at the lower end of the price and quality distribution. Appendix Figure A2a plots the empirical CDFs of the house square footage distributions for institutional and non-institutional owned houses in 2021. Institutionally-owned houses are 369 square feet smaller on average and have a lower variance. Appendix Figure A2b plots the corresponding empirical CDFs of assessed value per square foot, a proxy for house and location quality. Houses owned by institutional landlords are valued at $22 per square foot less than houses not owned by institutional landlords and also have a smaller variance. As suggested in Figure 7, institutional landlords' portfolios are also concentrated geographically. Appendix Figure A4a plots the share of single-family homes owned by institutional landlords in 2021 across Census tracts in the Atlanta MSA. Their holdings form a U-shape pattern, with a larger presence in the Western, Southern, and Eastern suburbs around Atlanta's central business district.

Institutional landlords' portfolios are also differentiated across firms. Among the three largest single-family landlords in our sample, Invitation Homes targets higher-end homes, Amherst Residential (Main Street Renewal) targets lower-end homes, and Progress Residential targets homes in the middle of the quality distribution, relative to other institutional-owned homes. Appendix Figure A3a plots the CDFs of the house size distributions for these three landlords' holdings. Invitation Homes owns larger homes on average, followed by Progress Residential, and then Amherst Residential. Appendix Figure A3b plots the corresponding CDFs of house value per square foot. In addition to owning larger houses, Invitation Homes tends to own higher quality homes, as proxied by value per square foot, than Progress Residential, which in turn owns higher quality houses than Amherst Residential. Figure A4 plots the share of single-family homes owned by each of these owners for Census tract in the Atlanta MSA. In addition to concentrating their holdings in different house size and value per square foot segments of the single-family market, these firms' holdings are also differentiated geographically. The similarity of house characteristics and geographic proximity of houses within a firm suggests that rental units will be closer substitutes with each other than with rental units in the broader housing market. These patterns are also consistent with firms seeking to standardize their holdings to reduce maintenance and tenant screening costs.

Fact 3. Institutional investors transfer homes from the homeownership sector to the rental sector.

Figure 8 displays the evolution of the single-family housing stock in the Atlanta MSA across time, by tenure and by owner type. Panel (a) shows that single-family rentals, in both levels and as a share of single-family homes, grew from the start of the data in 2009 before peaking in 2016, and started decreasing after that. However, even as the rental stock and the rental rate decreased, panel (b) shows that institutional investors continued to expand their portfolios after 2016 while the overall number of renter-occupied single-family units fell.

Figure 8: Single-Family Housing Stock by tenure and owner type

(a) Single-Family Houses by Tenure

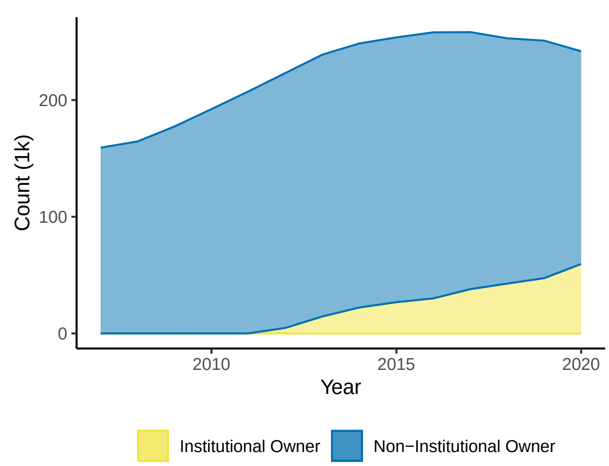

(b) Single-Family Rentals by Owner Type

Notes: Panel (a) plots the total stock of renter-occupied, owner-occupied, and vacant single-family housing units in the Atlanta MSA. Data are taken from 5-year Census ACS averages, centered on the midpoint of the average. Panel (b) plots the total stock of renter-occupied housing, taken from the 5-year ACS, and the total stock of single-family housing units owned by large, institutional investors, derived from property tax records provided by Corelogic, Inc.

3 Model

In this section, we present a model of a single-family housing market. The market is composed of two sectors, rental and homeownership. We begin with the rental sector in section 3.1, followed by the homeownership sector in section 3.2. Finally, section 3.3 describes how stocks and flows of vacant and occupied houses evolve and defines an equilibrium.

3.1 Rental sector

We start with an overview of the components of the rental sector. On the demand side, potential renters make a discrete choice among listed vacant houses. On the supply side, landlords set prices to maximize profits, trading off occupancies and prices. A matching technology determines vacancy lengths.

Rental demand We define a market $t = (Z, w)$ as the collection of individuals considering renting a house in ZIP Code Z during week w. The set of potential houses in market t is given by $\mathcal{J}_t$. Each consumer contacts a single utility-maximizing house, or chooses the outside option, denoted $j = 0$[23]. The utility individual i obtains from renting a house j in market t is given by:

$$u^R_{ijt} = -\alpha^R p_{jt} + \xi_{jt} + \zeta_{ijt} + (1 - \sigma^R)\epsilon_{ijt} \tag{1}$$

where $\delta_{jt} := -\alpha^R p_{jt} + \xi_{jt}$ is the deterministic part of the utility of renting house j, p_{jt} is the rental price, the parameter $\alpha^R > 0$ is the rental price sensitivity, and ξ_{jt} is an unobserved (by the econometrician) demand shock.

The term ϵ_{ijt} is an individual-specific idiosyncratic shock, which we assume is distributed type-1 extreme value. We assume a nested logit structure such that all houses belong to a single, inside goods nest and the outside option $j = 0$ belongs to a separate nest. The parameter $\sigma^R \in [0, 1]$ determines substitution between the outside option nest

[23] Our demand estimation relies on a static model. Fukasawa (2024) identifies three potential sources of bias from applying static demand when demand is dynamic. The most relevant sources in our setting are disregard of state variables and changing consumer expectations, which Fukasawa (2024) shows are both bounded by the conditional choice probability of each product. Given the large number of potential renters and buyers relative to contacts or views received by any individual listing, conditional choice probabilities are small, limiting the scope for these biases.

and the inside goods nest, with larger values of σ^R implies stronger substitution between the inside goods. The term ζ_{ijt} is a demand shock common to all products in the inside nest and has the unique distribution function such that $\zeta_{ijt} + (1-\sigma^R)\epsilon_{ijt}$ is also distributed type-1 extreme value (Cardell, 1997).

Using a nested logit with the outside option in a separate nest has two advantages. First, it relaxes the Independence of Irrelevant Alternatives (IIA) assumption, allowing substitution patterns implied by our model to not be strictly proportional to market shares. Second, it ensures that our definition of market size has little impact on these substitution patterns.

The above assumptions imply that the contact share of house j within the inside option nest $s_{j|g}$ and that the overall contact share s_j for some inside option j are given by:

$$s_{j|g} = \frac{e^{\delta_j/(1-\sigma^R)}}{D}, \qquad s_j = \frac{e^{\delta_j/(1-\sigma^R)}}{D^{\sigma^R}\left[1 + D^{1-\sigma^R}\right]}, \qquad D = \sum_{k\neq 0} e^{\delta_j/(1-\sigma^R)} \tag{2}$$

Rental matching We assume that there is a matching technology that translates the rate at which landlords are contacted by potential tenants into an expected time-on-market τ_{jt}. We parameterize this technology by specifying the following functional form for time-on-market:

$$\tau_{jt}(\boldsymbol{p_t}) = \left(s_{jt}(\boldsymbol{p_t}) \cdot M_t\right)^{\varepsilon^{\tau,\kappa}} \cdot \eta_{jt} \tag{3}$$

where s_j is house j's market share of contacts, M_t is the size of market t, and their product $s_j M_t := \kappa_j$ is house j's daily contacts rate. The term η_{jt} is a time-on-market demand shock, and $\varepsilon^{\tau,\kappa}$ is a structural parameter which corresponds to the elasticity of time-on-market with respect to the contact rate.

Rental supply In each market, landlords simultaneously choose prices for all houses in their portfolio to maximize total portfolio profis. Given a fixed set of advertised houses, landlords set prices while facing a tradeoff between revenue and occupancy: higher-priced units take longer to rent, and therefore have lower occupancy rates.

We first define house j's occupancy rate Ω_{jt} as a function of market prices $\boldsymbol{p_t}$ as

$$\Omega_{jt}(\boldsymbol{p_t}) = \frac{T_j(\boldsymbol{p_t})}{T_j(\boldsymbol{p_t}) + \tau_{jt}(\boldsymbol{p_t})} \tag{4}$$

where T_{kt} is the house-specific tenancy duration, and τ_{jt} is the time-on-market for house j, both of which are endogenously determined by the vector of market prices and the parameters of the matching technology.

Given market rental prices $\boldsymbol{p_t} = [p_{1t}, ..., p_{Jt}]$, steady-state profits of landlord l with portfolio of houses $\mathcal{L}$ are:

$$\begin{aligned} \pi_l(\boldsymbol{p_t}) &= \sum_{j \in \mathcal{L}} (p_{jt} - c_j^o)\Omega_{jt}(\boldsymbol{p_t}) + \sum_{j \in \mathcal{L}} (-c_j^v)(1 - \Omega_{jt}(\boldsymbol{p_t})) \\ &= \sum_{j \in \mathcal{L}} (p_{jt} - \underbrace{[c_j^o - c_j^v]}_{:= c_j})\, \Omega_{jt}(\boldsymbol{p_t}) \quad - \underbrace{c_j^v}_{\text{constant}} \end{aligned} \tag{5}$$

where c_j^o is the marginal cost of one additional time unit of occupancy, c_j^v is the marginal cost of one additional time unit of vacancy, and $c_j := c_j^o - c_j^v$ is the net marginal cost of occupancy.[24]

Since each landlord chooses prices to maximize porftolio profits, this implies $|\mathcal{L}|$ first-order conditions, one for each house j in landlord l's portfolio $\mathcal{L}$. Each of these first-order conditions is given by:

$$\Omega_{jt}(\boldsymbol{p_t}) + \sum_{k \in \mathcal{L}} (p_{jt} - c_j) \frac{\partial \Omega_{jt}(\boldsymbol{p_t})}{\partial p_k} = 0 \tag{6}$$

The key derivatives that matter for pricing are the own- and cross-price derivatives of occupancy $\frac{\partial \Omega_{jt}(\boldsymbol{p_t})}{\partial p_k}$. In section 4.2, we derive analytical expressions for these derivatives as a function of model parameters and detail our estimation strategy.

3.2 Homeownership sector

In this section, we describe the different components of the homeownership sector. On the demand side, potential homebuyers choose between listed houses for sale. On the supply side, house sellers set prices to maximize profits, trading off the time it takes to sell and the sales price of the house. Finally, a matching technology determines time-on-

[24] In equation (5), we omit house j's total fixed costs from the profit function since it does not affect the landlord's pricing decision.

market. The structure of this sector closely mirrors the rental sector, so we omit redundant details

Homeownership demand We assume potential homebuyers gain utility from buying houses. For each market t, an individual views the web page of a single utility-maximizing house. The utility individual i obtains from buying house j in market t is given by:

$$u_{ijt}^{H} = -\alpha^{H} p_{jt} + \xi_{jt} + \zeta_{ijt} + (1 - \sigma^{H})\epsilon_{ijt} \tag{7}$$

where each component of utility follows the same logic as in the rental sector, although here p_{jt} denotes the total *sales* price of the listed house, and parameters have an H superscript for Homeownership.

The above assumptions imply that the page view share of house j within the inside option nest $s_{j|g}$ and that the overall contact share s_j for some inside option j are given by:

$$s_{j|g} = \frac{e^{\delta_j/(1-\sigma^H)}}{D}, \qquad s_j = \frac{e^{\delta_j/(1-\sigma^H)}}{D^{\sigma^H}\left[1 + D^{1-\sigma^H}\right]}, \qquad D = \sum_{k \neq 0} e^{\delta_j/(1-\sigma^H)} \tag{8}$$

Homeownership matching Similar to the rental sector, we assume that there is a matching technology that translates web page views into expected time-on-market $\tau_{jt}(\boldsymbol{p_t})$, which we parameterize as follows:

$$\tau_j(\boldsymbol{p_t}) = (s_{jt}(\boldsymbol{p_t}) \cdot M_t)^{\varepsilon^{\tau,\nu}} \cdot \eta_{jt} \tag{9}$$

where s_{jt} now denotes house j's market share of *web page views*, M_t is the size of market t, and their product $\nu_{jt} = s_{jt}M_t$ is house j's average daily page view rate. The parameter $\varepsilon^{\tau,\nu}$ is the elasticity of time-on-market with respect to views.

Homeownership supply Unlike the rental sector, we assume house sellers choose prices as single-product firms.[25] Therefore, we refer to a "house" or a "seller" interchangeably, indexing both by j. Given a vector of market housing prices $\boldsymbol{p_t} = [p_{1t}, ..., p_{Jt}]$, the net present value of expected profits for house j is given by:

[25] Since we do not see evidence of high levels of concentration among house sellers, we assume all firms in homeownership sector price as single-product firms to keep the model tractable.

$$\Pi_j(\boldsymbol{p_t}) = \beta_j^{\tau_{jt}(\boldsymbol{p_t})} p_{jt} \tag{10}$$

where p_{jt} is house j's advertised sales price, β_j is the seller's discount factor and τ_{jt} is the expected time-on-market given a vector of market prices $\boldsymbol{p_t}$.

Equation 10 implies J first-order conditions for each market, each given by:

$$\beta_j = \exp\left(\frac{-1}{\tau_j \cdot \varepsilon_{jj}^{\tau,p}}\right) \tag{11}$$

The key elasticity determining house sellers' discount factors is the own-price elasticity of time-on-market $\varepsilon_{jj}^{\tau,p}$. In section 4.3, we derive an analytical expression for this elasticity as a function of model parameters, and detail our estimation strategy.

3.3 Equilibrium

This section describes the model equilibrium. We begin by providing an intuitive overview of what our model equilibrium entails and why we adopt this particular notion of equilibrium.

Our equilibrium notion has two key conditions: a Bertrand-Nash equilibrium in prices for vacant houses and a steady-state in terms of housing stocks. The first condition ensures that all sellers strategically set prices to maximize profits. The second condition ensures that housing stocks are constant over time.

The total housing stock of each sector consists of pools of *occupied* and *vacant* homes. Houses flow between occupied and vacant pools. Flows between vacant and occupied pools are determined by the size of the vacant pool and by prices of vacant houses. Flows between occupied and vacant pools are determined by the size of the occupied pool. In such an environment, a steady-state is an intuitive equilibrium notion for the following reason. If most houses are occupied, the occupied-to-vacant flow will be stronger than the vacant-to-occupied flow, leading to a shrinking occupied pool and a growing vacant pool. Conversely, if most houses are vacant, the vacant-to-occupied flow dominates. Thus, the only way for both stocks to remain stable is if these two flows are equalized.

We choose to have a steady-state as a key part of our equilibrium because it is crucial to measure price and welfare effects of large transfers of homes across sectors. To illustrate why, consider a one-time exogenous transfer of homes from the rental sector to the homeownership sector, such as a policy inducing institutional owners to sell their entire portfolio of rental houses to homeowners. If our model only had a notion of partial-equilibrium in prices for house sellers competing in a given week, the answer to our main question of interest – what are the price and welfare effects – would be entirely dependent on the *speed* of these transfers. In one extreme, if all houses were sold in the very first week, a partial-equilibrium model of competition would entail very large price and welfare effects. In the other extreme, if only one house was sold per week, the welfare effects would be essentially zero. However, over time both scenarios will converge to the same steady-state. Our equilibrium concept allows us to abstract from the specific timing of transfers and measure welfare at this steady state, after all price and quantity adjustments have fully played out.

3.3.1 Stocks and flows of houses

We start by describing the evolution of housing stocks and flows in detail. Within-sector flows of homes between the occupied and vacant pools are endogenous, and they are determined by the size and composition of the stock of occupied and vacant houses. Cross-sector flows are exogenous. Figure 9 provides an overview of housing stocks and flows, focusing on a single geographical market.

We use single capital letters to denote housing *stocks*, and double capital letters to denote *flows*. The price vector of vacant houses for each sector is $\boldsymbol{p}_s$, and the vacancy rate of each sector is v_s, where the subscripts $s \in \{h, r\}$ refer to each sector. The total housing stock in the entire market is denoted by N, and is comprised of either homes in the homeownership sector, H, or homes in the homeownership sector, R:

$$N = H + R \tag{12}$$

Within each sector, houses can be either occupied or vacant. For $s \in \{h, r\}$, O_s denotes the stock of occupied houses, and V_s denotes the stock of vacant houses:

$$H = O_h + V_h, \quad R = O_r + V_r \tag{13}$$

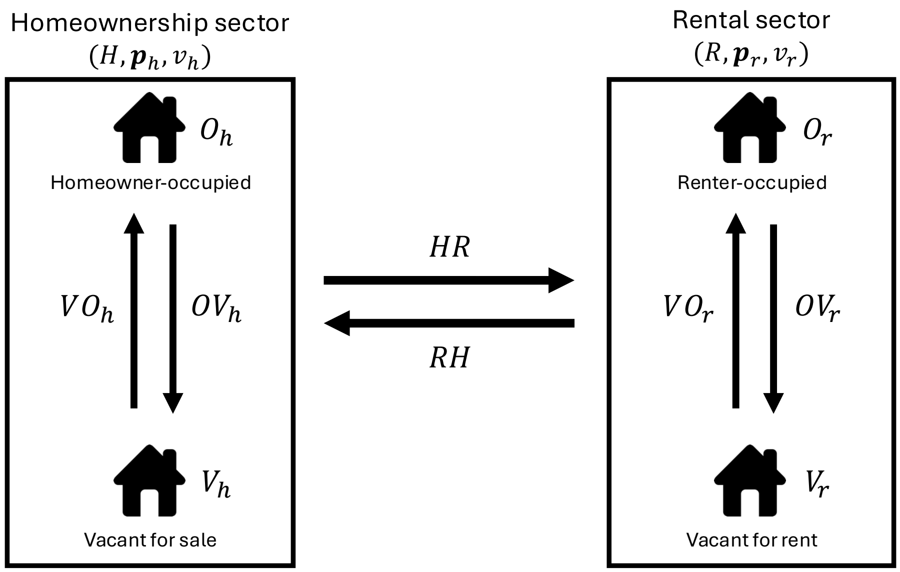

Figure 9: Overview of stocks and flows of houses

Houses flow endogenously between the occupied and the vacant pools within each sector. Occupied-to-vacant flows, or *turnovers*, are denoted by OV_s. Vacant-to-occupied flows, or *transactions*, are denoted by VO_s. We also allow houses to be allocated exogenously across sectors. Exogenous flows from the homeownership to the rental sector are denoted by HR, and rental-to-homeownership flows are denoted by RH.

Sectoral vacancy rates v_s are defined as the share of the sector-specific housing stock that is vacant:

$$v_s = \frac{V_s}{O_s + V_s}, \quad \forall s \in \{h, r\} \tag{14}$$

3.3.2 Linking weekly competition among vacant houses to stocks and flows

We now describe how we link our model of weekly competition among vacant houses specified in sections 3.1 and 3.2 to the evolution of stocks and flows. The key link between the two parts of the model are *transactions*, which are flows between the pools of occupied and vacant homes. To fully describe market behavior, we need to specify how these flows are determined. We do so based on a steady-state equilibrium condition: stocks of houses must be constant over time, which implies that opposing within-sector flows between vacant and occupied pools must equalize.

Weekly competition among vacant houses We begin by specifying the outcomes of weekly competition among vacant houses. The flow of houses between vacant and occupied pools (VO_s) is entirely determined by the part of the model describing competition among vacant houses. In a week, given a specific set of vacant houses $\mathcal{V}_s = \{j_1, ..., j_{V_s}\}$ and a vector of prices $\boldsymbol{p}_s$ for all houses, demand for each vacant house $j \in \mathcal{V}_s$ is a number of customer contacts or customer views given by

$$d_j = D^s_j(\{j_1, ..., j_{V^s}\}, \boldsymbol{p}_s | \theta_D) \tag{15}$$

Expected time-on-market for each house is determined by a matching technology $m(d_j|\theta_m)$ that only depends on each vacant house's demand and the matching parameters θ_m. We assume these matching parameters are common for all houses. Therefore, we can write the vector $\boldsymbol{\tau_s}$ of expected time-on-markets for each sector s as

$$\boldsymbol{\tau_s} = m(\boldsymbol{d_s}|\theta_m) \tag{16}$$

where $\boldsymbol{d_s} = \{d_1, ..., d_{V^s}\}$ is the vector of customer contacts or customer views for all vacant houses.

Transactions (VO_s) The weekly outcomes described above depend on the exact set of houses $\mathcal{V}_s = \{j_1, ..., j_{V_s}\}$ competing in the vacant pool in the short run. One key weekly outcome is the number of *transactions*. For each house, this is a probabilistic outcome. Let $t_j \in \{0, 1\}$ be an indicator for whether a house transacts in a given period. We assume that t_j is a Bernoulli random variable with parameter P_j. If we also assume that house j's weekly probability of transacting is constant over the span of that house's vacancy and is given by P^t_j, this implies that time-on-market τ_j (number of periods before a successful transaction) follows a geometric distribution with parameter P_j and expectation $\frac{1}{P_j}$. We can thus retrieve the weekly probability of house j transacting by inverting the expected time-on-market:

$$E[t_j] = P_j = \frac{1}{E[\tau_j]} \tag{17}$$

To obtain the total expected number of transactions given a set of vacant houses $j \in \mathcal{V}_s$, we can use linearity of expectations and sum across transaction probabilities for all vacant houses. We define the expected flow VO_s of houses between the vacant and occupied pool as the expected number of weekly transactions:

$$VO_s = E\left[\sum_{j\in\mathcal{V}_s} t_j\right] = \sum_{j\in\mathcal{V}_s} E[t_j] = \sum_{j\in\mathcal{V}_s} \frac{1}{E[\tau_j]} \tag{18}$$

Turnovers (OV_s) Next, we define the expected flow OV_s of houses from the occupied pool to the vacant pool, which we call *turnovers*. Let $\mathcal{O}_s = \{k_1, ..., k_{O_s}\}$ be the set of occupied houses in a given period. Let T_k be house $k \in \mathcal{O}_s$ expected tenancy length. For tractability, we do not endogenize occupied houses' tenancy length, but we assume they are exogenously given for each house. At the end of its tenancy, a house turns over and goes back to the vacant pool. Let o_k be an indicator for whether a house turns over in a given period. Then, following the same logic as above, we define the occupied-to-vacant flows OV_s as the expected number of weekly turnovers in sector s:

$$OV_s = E\left[\sum_{k\in\mathcal{O}_s} o_k\right] = \sum_{k\in\mathcal{O}_s} E[o_k] = \sum_{j\in\mathcal{V}_s} \frac{1}{E[T_k]} \tag{19}$$

3.3.3 Equilibrium definition

We can now define our notion of equilibrium. We assume stocks of vacant and occupied homes are in steady state and are constant over time. Within each sector, this implies weekly flows in opposing directions between the pools of vacant and occupied houses are equal.

Definition 1 [Sectoral equilibrium]. *For each sector s, given the total set of vacant and occupied homes $\mathcal{S} = \mathcal{V}_s \cup \mathcal{O}_s$, each endowed with expected tenancy length $\{T_1, ..., T_{\mathcal{S}}\}$, a **sectoral equilibrium** is a vacancy rate v_s and a price vector $p(V_s)$ vector for each $\mathcal{V}_s \subset \mathcal{S}$ such that*

(a) Consumers choose optimally (Equation 15 holds)

(b) Time-on-market is determined by matching according to Equation (16)

(c) Landlords and house sellers choose prices to maximize profits (Equations 6 and 11 hold)

(d) $OV_s = VO_s$ (steady state)

Condition (c) says that landlords – or house sellers, depending on the sector – strategically set prices according to a Bertrand-Nash equilibrium. Condition (d) is the key steady-state condition. It implies that the vacancy rate v_s is high (or low) enough such that the size of the vacant pool implies enough transactions to match the turnover from the occupied pool. This notion of steady-state equilibrium allows us to determine what would happen to each sector if large stocks of houses were to be exogenously transferred across sectors.

4 Estimation

In this section, we describe the identification and the estimation of the model primitives. 4.1 discusses our identification strategy for the price sensitivity, along with descriptive results. 4.2 discusses estimation results for the homeownership sector, and 4.3 discusses estimation results for the homeownership sector. For each sector, we show results for demand parameters, matching technology parameters, and cost parameters.

4.1 Identification of price sensitivity

The key parameter of our model is the price sensitivity α, which determines own and cross-price elasticities. Since landlords are choosing prices as a function of demand, these may be correlated with unobserved demand shocks ξ_{jt}. We tackle this endogeneity challenge by exploiting discontinuities in the rate of customer contacts around price changes as an instrument for the price.

For this identification strategy to be valid, we need our instrument to be both relevant and excluded. Below, we show that our instrument is relevant since both contact rates and page view rates react sharply and discontinuously to price changes, and these discontinuities are proportional to the magnitude of the price change.

One possible threat to identification is that price drops may be more common during periods of low demand, which would induce correlation between demand shocks ξ_{jt} and our discontinuity instrument, violating the exclusion restriction. We argue that our instrument is excluded on the basis that demand shocks are not systematically different *immediately* before and *immediately* after price changes. Hence, variation in the number of customer contacts around rental price changes is only driven by the price elasticity. Having high-frequency demand proxies such as web page views and customer contacts is what allows us to focus on the vicinity of the price change, exploiting discontinuities in these demand proxies.

One additional threat to our identification strategy is that discontinuous price changes happen in conjunction with discontinuous changes in demand, such as on the first day of the month. If that were to be the case, we would likely see spikes in price drops on some specific days of the month. In Appendix Figure A5, we show that there are no spikes in price drops happening at the beginning of the month, or on any specific days.

Data patterns To shed light on the key source of variation identifying price elasticities, Figure 10 presents discontinuities for both small (below-median) and large (above-median) price changes. The top panel shows discontinuities in the daily contact rate of rental listings, while the bottom panel displays the same pattern for the page view rates of listings for sale. We observe that the size of these discontinuities is proportional to the magnitude of the price change: larger price drops are associated with larger discontinuities in the contact and viewership rates.

Event study The patterns in Figure 10 above suggest that price changes imply a substantial and persistent shift in both contact and page view rates. However, the graph also shows a slight downwards time trend in these rates, which is particularly evident in the time series of normalized page view rates for listings for sale. This happens because daily contacts and page views tend to decrease over time for each house over the course of the listing spell. Therefore, to further investigate the persistence of this shift, we perform an event study showing changes in daily page views and daily contact rates around price drops, additionally controlling for this trend by including "days-on-Zillow" fixed effects. Figure 11 below plots the coefficients from this event study.

We see that de-trended contact and viewership rates spike immediately after a price drop, and remain consistently higher even 10 days after the price drop. Thus, we use these discontinuities in our demand estimation as an instrument to identify price elasticities.

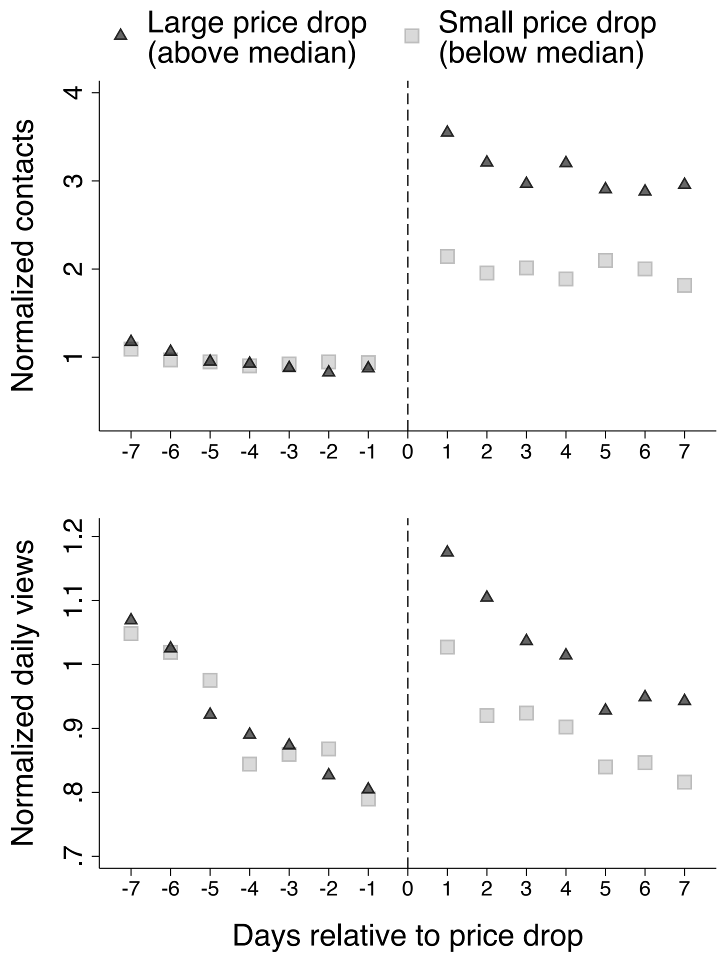

Figure 10: Discontinuities in contacts and page views for small and large price drops

Notes: The top panel plots average contacts before and after a rental price drop. The bottom panel plots average customer views before and after a sales price drop. For each house, contacts and views are normalized by their average value before the price drop. In both panels, houses are split into two groups: those with large (above average) price drops, represented by black triangles, and those with small (below average) price drops, represented by gray squares.

Reduced-form analysis Before estimating the full structural model, we perform a reduced-form estimation of the elasticites of both customer views and customer contacts with respect to listed prices.

We begin with the rental sector. We run a regression of the log-difference between the average number of customer contacts one week before and one week after, $d\log(\kappa_j) := log(\kappa_j^{\text{post}}) - \log(\kappa_j^{\text{pre}})$, on the log-difference between the average price one week before, $d\log(p_j) := log(p_j^{\text{post}}) - \log(p_j^{\text{pre}})$, for the sample of houses with at least one price change:

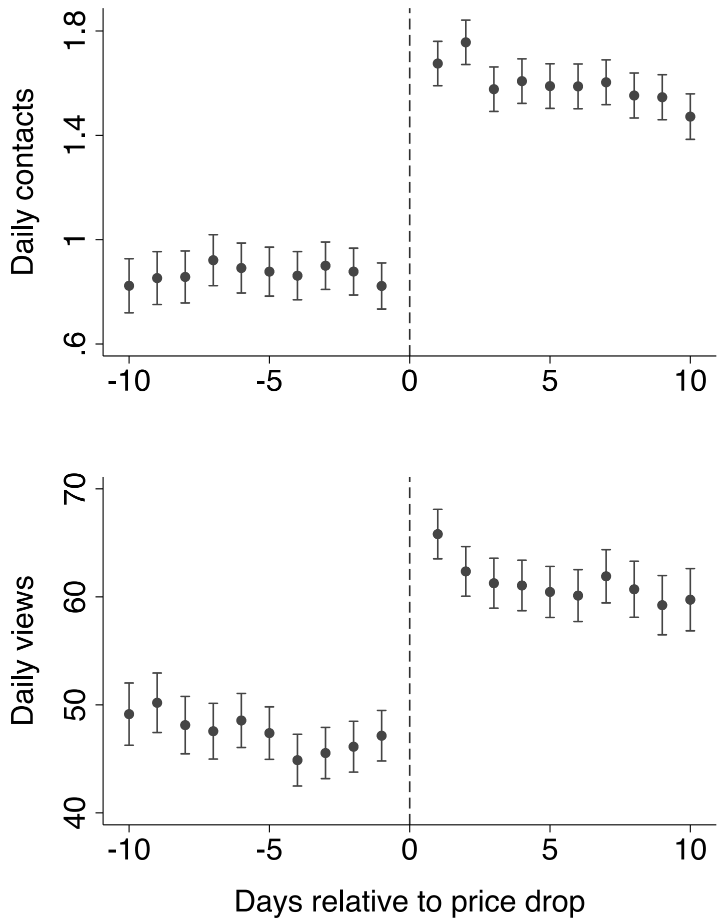

Figure 11: Event studies of discontinuities in contacts and views around price drops

Notes: The figures above display event studies of demand proxies 10 days before and 10 days after a price drop. For the rental sector, we the top panel shows daily contacts. For the homeownership sector, the bottom panel shows daily views. In both panels, we select a sample of houses that experienced one price drop over the course of its listing spell. Both point estimates and 95% confidence intervals are displayed.

$$d\log(\kappa_j) = \beta_0 + \varepsilon^{\kappa,p} \cdot d\log(p_j) + \gamma_{\text{month}} + \gamma_{\text{days-on-zillow}} + \epsilon_j \tag{20}$$

We display the results of the estimation on Table 1, which shows an estimated elasticity of -7.9. We also display the full relationship between changes in prices and changes in contacts around price discontinuities by plotting a binned scatterplot of $d\log(\kappa_j)$ on $d\log(p_j)$ in Figure 12. The plot shows a log-linear relationship between the two, which supports the relevance of our discontinuity instrument in capturing the variation in demand with respect to price changes.

Figure 12: Changes in contacts vs. changes in price

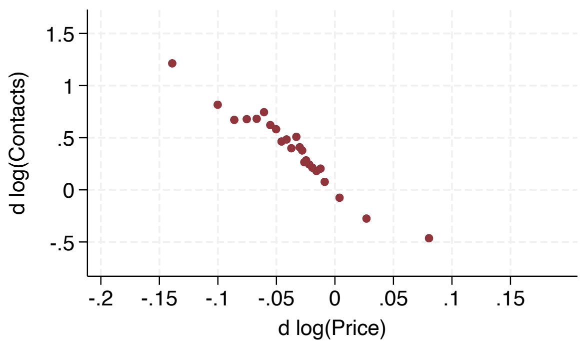

Table 1: Estimation results of Eq. (20)

Price elasticity of contacts	
$\varepsilon^{\kappa,p}$	-7.911***
	(0.099)
# Observations	4,615
Month FE	Yes
Days-on-zillow FE	Yes

Note: Standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Next, we perform the same exercise for the homeownership sector. We run a regression using the same specification as in (20), but instead of contacts, we use the number of customer web page views ν_j as the outcome variable:

$$d\log(\nu_j) = \beta_0 + \varepsilon^{\nu,p} \cdot d\log(p_j) + \gamma_{\text{month}} + \gamma_{\text{days-on-zillow}} + \epsilon_j \tag{21}$$

Table 2 shows an estimated elasticity of -5.6, and Figure 13 displays a similar log-linear relationship as in Figure 12.

Figure 13: Changes in views vs. changes in price

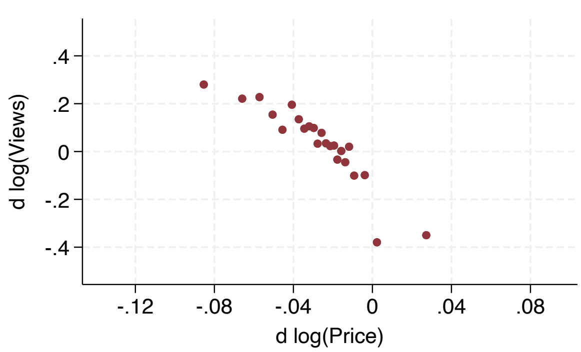

Table 2: Estimation results of Eq. (21)

Price elasticity of views	
$\varepsilon^{\nu,p}$	-5.641***
	(0.266)
# Observations	4,852
Month FE	Yes
Days-on-zillow FE	Yes

Note: Standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

4.2 Rental sector

We now show our main estimation results for the rental sector. We start by explaining how we back out landlords' net marginal costs c_j. Recall from Equation (6) that landlords' first-order-conditions are given by

$$\Omega_{jt}(\boldsymbol{p_t}) + \sum_{k\in\mathcal{L}}(p_{jt} - c_j)\frac{\partial\Omega_{jt}(\boldsymbol{p_t})}{\partial p_k} = 0$$

Since occupancies Ω_{jt} and prices $\boldsymbol{p_t}$ are directly observable in the data, equation (6) implies that estimating the own- and cross-price derivatives of occupancy is enough to infer net marginal costs $c_j = c_j^o - c_j^v$.

To ensure tractability when computing analytical expressions for these derivatives, we make two assumptions. First, we assume that prices only affect time-on-market through contacts:

Assumption 1:

$$\tau_j(\boldsymbol{p}) = \tau_j(\kappa_1(\boldsymbol{p}), ..., \kappa_J(\boldsymbol{p})) \tag{22}$$

Second, we assume that time-on-market for house j only depends on the contacts of house j, and that contacts of any other house k have no effect on time-on-market for house j:

Assumption 2:

$$\boldsymbol{\tau}_j\left(\kappa_j(\boldsymbol{p}), \boldsymbol{\kappa}_{-j}(\boldsymbol{p})\right) = \boldsymbol{\tau}_j\left(\kappa_j(\boldsymbol{p}), \boldsymbol{\kappa}'_{-j}(\boldsymbol{p})\right), \quad \forall \boldsymbol{\kappa}'_{-j}, \boldsymbol{p_t}'_{\boldsymbol{-j}} \tag{23}$$

Equation (2), along with assumptions 1 and 2, imply the following analytical expressions for the price derivatives of occupancy:[26]

$$\frac{\partial\Omega_j}{\partial p_k} = \begin{cases} \Omega_j(1-\Omega_j)\,\alpha^R\,\varepsilon^{\tau,\kappa}\,\left(\frac{1}{1-\sigma^R} - \frac{\sigma^R}{1-\sigma^R}s_{j|g} - s_j\right) & \text{if } j = k \\ -\Omega_j(1-\Omega_j)\,\alpha^R\,\varepsilon^{\tau,\kappa}\,\left(\frac{\sigma^R}{1-\sigma^R}s_{j|g} + s_j\right) & \text{if } j \neq k \end{cases} \tag{24}$$

The analytical expressions in Equation (24) depend on demand parameters α^R and σ^R, matching parameter $\varepsilon^{\tau,\kappa}$, and data. Therefore, once both our demand and matching models are estimated, we can recover implied own and cross-price derivatives of occupancy.

Endowed with occupancy derivatives, we then recover marginal costs by following Nevo (2001). Let $S_{kj} = \frac{\partial\Omega_j(\boldsymbol{p})}{\partial p_k}$, and define the ownership-derivative matrix as:

[26] See Appendix B for the derivation of these analytical derivatives.

$$\boldsymbol{\mathcal{D}} = \begin{cases} S_{kj}, & \text{if } \exists\{j,k\} \in \mathcal{F} \\ 0, & \text{otherwise} \end{cases} \tag{25}$$

Then:

$$\begin{bmatrix} \Omega_1(\boldsymbol{p}) \\ \vdots \\ \Omega_{|\mathcal{L}|}(\boldsymbol{p}) \end{bmatrix} + \boldsymbol{\mathcal{D}} \begin{bmatrix} p_1 - c_1 \\ \vdots \\ p_{|\mathcal{L}|} - c_{|\mathcal{L}|} \end{bmatrix} = \mathbf{0}, \tag{26}$$

Which can be re-written as:

$$\begin{bmatrix} p_1 - c_1 \\ \vdots \\ p_{|\mathcal{L}|} - c_{|\mathcal{L}|} \end{bmatrix} = -\boldsymbol{\mathcal{D}}^{-1} \begin{bmatrix} \Omega_1(\boldsymbol{p}) \\ \vdots \\ \Omega_{|\mathcal{L}|}(\boldsymbol{p}) \end{bmatrix} \tag{27}$$

Equation (27) shows that we can recover net marginal costs c_j by inverting the ownership-derivative matrix $\mathcal{D}$ and plugging in observed prices p_j and occupancies Ω_j.

Rental demand We now detail the estimation of the demand parameters for the rental sector. Recall from Equation (1) that individual i's utility from renting house j is given by

$$u_{ijt} = -\alpha^R p_{jt} + \xi_{jt} + \zeta_{ijt} + (1 - \sigma^R)\epsilon_{ijt}$$

where the outside option $j = 0$ corresponds to not contacting any house listing in Atlanta, or contacting other types of listings such as apartments.

Importantly, we distinguish between two separate concepts. On one hand, there is the number of *contacts* a housing listing receives on a daily basis. On the other hand, there are *successful rentals*. Because we observe both contacts and successful rentals for each house in the data, we exploit both of these outcomes in the estimation of our model parameters.

We assume prospective renters *contact* the house that maximizes their utility. This approach deviates from the traditional housing literature (Bayer et al., 2007), which typically infers preferences from final transactions (matches). Instead, we argue that using contacts offer two advantages for identifying and estimating housing demand. First, contacts provide a more direct measure of consumer preferences. Observed matches are equilibrium outcomes determined by both preferences and supply-side rationing: consumers often fail to secure their favorite home. Therefore, inferring demand from matches without accounting for such choice constraints would bias preference estimates, as the model would conflate supply-side availability with demand-side utility (Agarwal and Somaini, 2025). By modeling the decision to *contact* a listing, we capture the consumer's preferred choice before these supply constraints bind, allowing us to recover preference parameters purged of supply-side rationing. Second, while matches provide only a single outcome per house, daily contact data offers multiple observations for the same listing, enabling the discontinuity identification design described below.

We start by explaining our identification for the price coefficient α^R. Since landlords are choosing rents as a function of demand, these may be correlated with unobserved demand shocks ξ_{jt}. We tackle this endogeneity challenge by exploiting discontinuities in the rate of customer contacts around rental price changes. Figures 10 and 11 illustrate three key points about these discontinuities: (i) customer contacts react sharply and discontinuously at price changes, (ii) the size of the discontinuity is proportional to the magnitude of the price change, and (iii) he contact trend following the rent change remains consistently higher than the trend before the change. We argue that this instrument satisfies the exclusion restriction on the basis that demand shocks are not systematically different right before and right after price changes. Hence, variation in the number of customer contacts around rental price changes is only driven by the price elasticity.

We identify the parameter σ^R, which governs substitution between inside and outside option nests, using local and quadratic differentiation instruments as in Gandhi and Houde (2019). These instruments are based on geodesic distance d_{jkt} between houses j and k within each market t, and are defined as follows:

$$\begin{aligned} z_{jt}^{\text{quad}} &= \sum_{k \in F} d_{jkt}^2 \, , \qquad && \sum_{k \notin F} d_{jkt}^2 \\ z_{jt}^{\text{local}} &= \sum_{k \in F} I\left[d_{jkt} < \text{SD}(d)\right], \qquad && \sum_{k \notin F} I\left[d_{jkt} < \text{SD}(d)\right] \end{aligned} \tag{28}$$

Intuitively, these differentiation instruments capture a localized measure of competition around each individual house—how "crowded" the area is in the vicinity of each property—which we are able to measure using the exact latitudes and longitudes of each house.

We compute the potential market size M_t for each ZIP-week market by scaling the observed aggregate number of contacts, and we allow M_t to vary at the monthly level for each ZIP code. We calibrate this scaling factor by comparing the total volume of online activity for the Atlanta metropolitan area to the number of potential movers estimated from Census migration flows (see Appendix C.3).

We estimate our nested logit model using two-stage least squares. Table (3) shows results for the estimation of our demand specification from Equation (1). Column (1) displays OLS estimates for a logit model without a nest. Column (2) extends this to a nested logit structure. Column (3) instruments for price using discontinuities around price changes. Column (4) is our preferred specification, where we additionally instrument for σ^R using Gandhi-Houde differentiation instruments.

Table 3: Demand estimation results (rental sector)

	(1)	(2)	(3)	(4)
α^R	-0.1767	-0.0230	-0.6108	-1.8434
	(0.0051)	(0.0015)	(0.0461)	(0.0618)
σ^R		0.9553	0.9043	0.4429
		(0.0007)	(0.0019)	(0.0088)
Estimator	OLS	OLS	IV	IV
Nest		✓	✓	✓
House FE			✓	✓
Discontinuity Price Instr.			✓	✓
Gandhi-Houde Instr.				✓
Median Own-Price Elasticity	-0.36	-1.03	-12.72	-6.63
N	182,016	182,016	182,016	182,016

Notes: This table shows demand estimation results for the specifications from equation (1). Standard errors in parentheses.

We see that in columns (1)-(2) with the OLS specifications the price coefficient is low relative to columns (3)-(4). To gauge whether implied own-price elasticities are reasonable, we compare them with the reduced-form estimate from Table 1, which was -7.9. Our results from the OLS specifications imply price elasticities that are too low: median own-price elasticities in columns (1)-(2) are between -0.3 and -1.0. Including one instrument for the price while assuming exogenous inside shares in (3) implies an own-price elasticity of 12.7, about 60% higher than the reduced-form estimate. Finally, when we instrument for both the price and the nest parameter in (4), our estimate for the own-price elasticity is -6.6, reasonably close to the reduced-form estimate of -7.9.

Rental Matching Next, we estimate the parameters of our matching technology from equation (3). The main parameter of interest is the elasticity $\varepsilon^{\tau,\kappa}$ of time-on-market with respect to customer contacts. The key source of variation used to estimate this parameter is the joint distribution of time-on-market and customer contacts. Figure 14 shows that houses with higher average contact rates tend to rent faster than those with lower average contact rates.

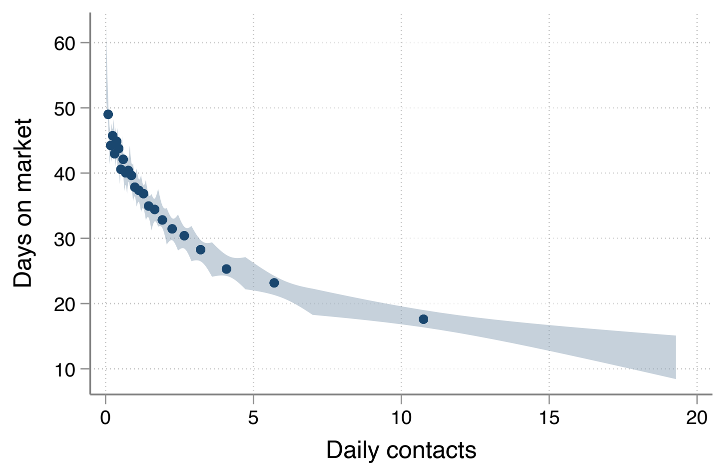

Figure 14: Houses with more contacts rent faster

Notes: This figure plots a binned scatterplot of time on market against average daily customer contacts for our cross-section of listings for sale

Note that, taking logs, equation (3) implies a linear relationship between the logarithm of time-on-market and the logarithm of the contact rate:

$$\log \tau_{jt}(\boldsymbol{p_t}) = \varepsilon^{\tau,\kappa} \log \left(s_{jt}(\boldsymbol{p_t}) \cdot M_t\right) + \log \eta_{jt} \tag{29}$$

We therefore estimate the parameter $\varepsilon^{\tau,\kappa}$ by regressing the logarithm of observed time-on-market on the number of new daily contacts received by landlords.

Table 4 displays the results of our estimation. Columns (1) and (2) present specifications without and with ZIP Code fixed effects, respectively. Both specifications yield an elasticity of time-on-market with respect to customer contacts of approximately -0.3. This implies that a 10% increase in customer contacts corresponds to a 3% decrease in time-on-market. Figure 15 illustrates this relationship through a binned scatterplot of average log

Figure 15: Changes in time-on-market vs. changes in contacts

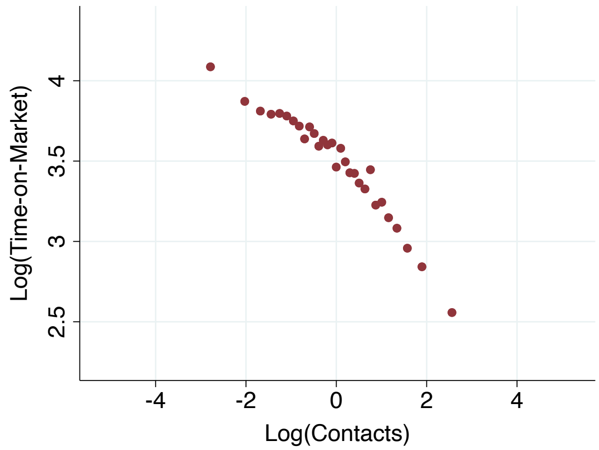

Table 4: Matching estimation results (rental sector)

	(1)	(2)
Log(Contacts)	-0.3026	-0.3095
	(0.0043)	(0.0044)
Constant	3.5333	3.5343
	(0.0044)	(0.0044)
Estimator	OLS	OLS
ZIP Code FE		✓
N	32,214	32,214

Notes: Standard errors in parentheses.

time-on-market against log customer contacts, supporting our log-linear functional form assumption.

Net marginal costs Using our model estimates, we now describe the implied distribution of backed-out net marginal costs of occupancy c_j. Recall from equation (5) that $c_j := c_j^o - c_j^v$ is a *net* marginal cost, representing the difference between the cost of one additional time unit of occupancy and the cost of one additional time unit of vacancy. Although we can only identify the difference between these two marginal costs, there are institutional reasons for why the cost of vacancy may be larger than the cost of occupancy. First, there are direct costs associated with vacancy, such as advertising a property, potentially renovating it, and screening candidate tenants. Second, vacancy incurs a significant opportunity cost because of the foregone rental revenue. Therefore, since the marginal cost of vacancy is likely to be large relative to the marginal cost of occupancy, we expect this net marginal cost to be negative.

Figure 16 shows the overall distribution of recovered net marginal costs at the house-week level, expressed in daily terms. Almost all of the recovered marginal costs are negative, with 70% of the observations ranging between -\$1,000 and \$0. The median daily net marginal cost is -\$638. This means that the median landlord behaves as if increasing vacancy by one day *and* reducing occupancy by one day implies a net cost of \$638.

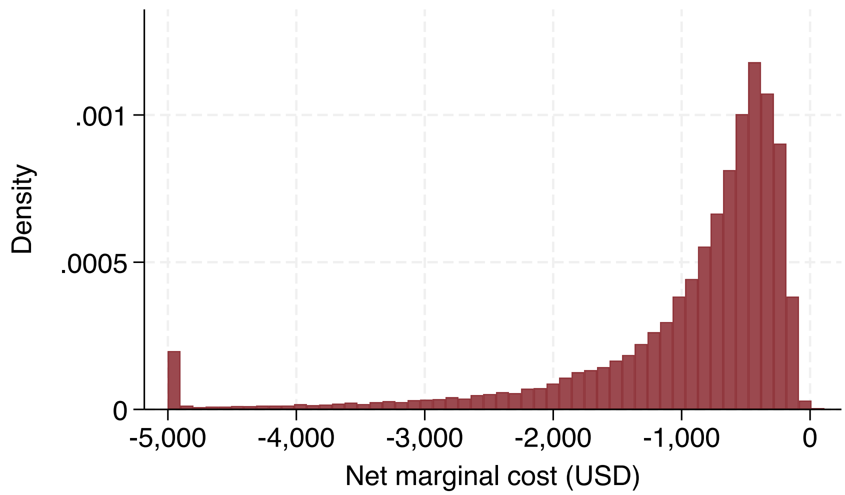

Figure 16: Distribution of daily net marginal costs

While the recovered net marginal costs imply a high daily cost of vacancy, we interpret this not strictly as a monetary cost, but rather as a measure of landlords' aversion to vacancy required to rationalize observed pricing behavior. There are several reasons why this aversion to vacancy may be high. First, having a new vacancy (relative to renewing a lease) involves significant switching costs such as repairs, cleaning, renovating, advertising, tenant screening, as well as the risk of acquiring a new tenant with unknown financial reliability. Second, small "mom-and-pop" landlords may face short-term liquidity constraints that incentivize them to avoid vacancies, even at the expense of substantial rental yield. Third, professional rental managers may also face managerial incentives that reward high occupancy rates in the short run.

We also analyze how the net marginal costs differ depending on whether houses are owned by institutional landlords or not. Table 5 shows the average marginal costs by category of landlord. We see that, on average, institutional landlords have slightly *higher* net marginal costs of occupancy. This result makes sense if the marginal cost of vacancy is large relative to the marginal cost of occupancy: if institutional landlords are cost-efficient in managing vacancies and have a *low* marginal cost of vacancy, this implies they have a *high* net marginal cost of occupancy.[27]

[27] The similarity in net marginal cost distributions between landlord types does not imply identical cost structures. Since we identify only the difference between occupancy and vacancy costs, our estimates are consistent with institutional landlords having offsetting efficiencies in both dimensions—lower occupancy *and* vacancy costs.

Landlord category	Mean	Std. Dev.	N
Non-institutional	-1,022.50	1,215.69	108,414
Institutional	-985.74	1,146.56	73,602

Table 5: Net marginal costs by landlord category (in USD)

4.3 Homeownership sector

We now turn to estimation results for the homeownership sector. We begin by discussing how to recover house sellers' discount factors β_j. Recall from equation 10 that each house seller j sets the listing sales price such that it satisfies the following first-order condition:

$$\beta_j = \exp\left(\frac{-1}{\tau_j \cdot \varepsilon_{jj}^{\tau,p}}\right)$$

Because house sellers act as single-product firms, there are no cross-price elasticities in the expression above. Since time-on-market τ_{jt} is observed, backing out β_j only requires us to estimate the own-price elasticity of time-on-market $\varepsilon_{jj}^{\tau,p}$. For that purpose, note that equation 8 implies that $\varepsilon_{jj}^{\tau,p}$ can be written as:[28]

$$\begin{aligned}\varepsilon_{jj}^{\tau,p}(\boldsymbol{p}|\alpha^H,\sigma^H,\varepsilon^{\tau,\nu}) &= \varepsilon^{\tau,\nu}\ \varepsilon_{jj}^{\nu,p}(\boldsymbol{p}|\alpha^H,\sigma^H) \\ &= -\varepsilon^{\tau,\nu}\ \alpha^H p_j\left(\frac{1}{1-\sigma^H}-\frac{\sigma^H}{1-\sigma^H}s_{j|g}(\boldsymbol{p})-s_j(\boldsymbol{p})\right)\end{aligned} \tag{30}$$

where $\varepsilon^{\tau,\nu}$ is the elasticity of house j's time-on-market with respect to views[29], which is the main structural parameter of the matching technology. We can thus use equation 30 and the first-order condition 11 to express the discount factor β_j as a function of parameters and prices:

$$\beta_j(\boldsymbol{p}|\alpha^H,\sigma^H,\varepsilon^{\tau,\nu}) = \exp\left(\frac{-1}{-\tau_j(\boldsymbol{p})\ \varepsilon_{jj}^{\tau,\nu}\ \alpha^H\ p_j\left(\frac{1}{1-\sigma^H}-\frac{\sigma^H}{1-\sigma^H}s_{j|g}(\boldsymbol{p})-s_j(\boldsymbol{p})\right)}\right) \tag{31}$$

[28] See Appendix B for a derivation of the analytical derivatives and elasticities.

[29] We assume that house j's time-on-market only depends on house j's views. This simplifying assumption allows us to isolate competition between different houses in the demand part of the model only.

Below, we describe how we estimate the demand parameters α^H and σ^H, followed by a discussion on how we estimate the matching technology parameter $\varepsilon^{\tau,\nu}$.

Homeownership demand We now explain how we estimate the demand parameters for the homeownership sector. Recall from equation (7) that individual i's utility from buying house j is given by

$$u_{ijt} = -\alpha^H p_{jt} + \xi_{jt} + \zeta_{ij} + (1 - \sigma^H)\epsilon_{ijt}$$

We define the outside option $j = 0$ as not viewing any house listing in Atlanta, or viewing other types of listings such as apartments.

Our identification strategy for the price sensitivity α^H and the parameter σ^H is identical to the one we use for the rental sector. To identify α^H, we use discontinuities in the rate of web page views around price corrections. To identify σ^H, we build differentiation instruments using the geodesic distances between houses within each market, following Gandhi and Houde (2019). We compute the market size M_t by scaling the aggregate number of page view using Census migration flows data (see Appendix C.3). For a detailed discussion, see Section 4.2.

We implement our estimation of the nested logit model using two-stage least squares. Table (6) shows results for the estimation of our demand specification from Equation (7).

Table 6: Demand estimation results (homeownership sector)

	(1)	(2)	(3)	(4)
α^H	0.0268	0.0006	-0.1728	-0.5120
	(0.0008)	(0.0002)	(0.0201)	(0.0235)
σ^H		0.9651	0.9228	0.7093
		(0.0005)	(0.0011)	(0.0050)
Estimator	OLS	OLS	IV	IV
Nest		✓	✓	✓
House FE			✓	✓
Discontinuity Price Instr.			✓	✓
Gandhi-Houde Instr.				✓
Median Own-Price Elasticity	0.12	0.08	-9.55	-7.54
N	174,059	174,059	174,059	174,059

Notes: This table shows demand estimation results for the specifications from equation (7). Standard errors in parentheses.

Column (1) shows OLS estimates for a standard logit model without a nest. Column (2) shows similar OLS estimates with a nest. We see that the price coefficient α^H is positive on all three specifications, suggesting the presence of omitted variable bias due to potential correlation between price levels and demand shocks. Price elasticities are all positive and unreasonably close to zero.

Column (3) instruments for price using discontinuities around price changes. We see that the price coefficient becomes negative, the median price elasticity also switches sign and increases in magnitude to -9.6. To gauge whether our implied price elasticity is reasonable, we compare it with the reduced-form estimate from Table 2, which was -5.6. The elasticity implied by our estimates from specification (4) is still high in absolute value relative to our reduced-form estimate. In part, we believe this is because our estimate for the nest parameter σ^H is still high and similar to the values obtained in the OLS specification in (2).

Column (4) is our preferred specification, where we additionally instrument for σ^H using Gandhi-Houde differentiation instruments. The estimate for σ^H becomes smaller, and the implied median price elasticity of views is -7.5, closer to the estimate of -5.6 obtained in the reduced-form exercise.

Homeownership matching We now estimate the parameters of our matching technology from equation (32). Similar to the rental sector, the key source of variation used to estimate this parameter is the joint distribution of time-on-market and customer views. Figure 17 shows that houses with higher average web page views tend to rent faster than those with lower average views.

Note that, taking logs, equation 9 implies a linear relationship between the logarithm of time-on-market and the logarithm of the contact rate:

$$\log \tau_{jt}(\boldsymbol{p_t}) = \varepsilon^{\tau,\nu} \log \left(s_{jt}(\boldsymbol{p_t}) \cdot M_t\right) + \log \eta_{jt} \tag{32}$$

We therefore estimate this parameter by regressing the logarithm of observed time-on-market on the number of new daily web page views.

Table 7 displays the results of our estimation. Column (1) shows a specification that regresses, while column (2) adds ZIP Code fixed effects. We obtain an elasticity of time-

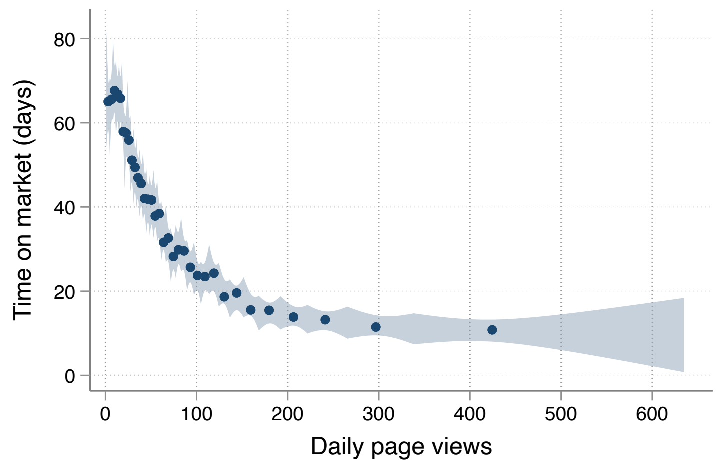

Figure 17: Houses with more views sell faster

Notes: This figure plots a binned scatterplot of time on market against average daily web page views for our cross-section of listings for sale

Figure 18: Changes in time-on-market vs. changes in views

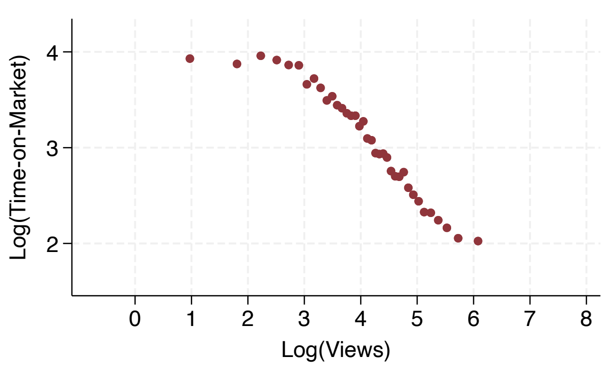

Table 7: Matching estimation results (homeownership sector)

	(1)	(2)
Log(Views)	-0.4999	-0.5403
	(0.0070)	(0.0082)
Constant	5.0953	5.2574
	(0.0292)	(0.0337)
Estimator	OLS	OLS
ZIP Code FE		✓
N	15,177	15,177

Notes: Standard errors in parentheses.

on-market with respect to customer views of about -0.5 for both specifications. The interpretation of this coefficient is that a 10% increase in the number of web page views leads to a 5% decrease in the time-on-market. To visualize the full relationship between these two variables, Figure 18 shows a binned scatterplot of average log time-on-market as a function of log customer views. We see that our log-linear functional form assumption seems in line with what we observe the data.

Discount factors Using our model estimates, we now show the implied distribution of backed-out discount factors β from equation (31). Since the time-on-market we use for the homeownership demand estimation is the time until the first contingency, it underestimates the time before the house is actually sold. We model the relationship between time-to-contingency and expected time-to-sale based on the following assumptions. First, homes listed for sale will always become contingent given enough time: for a given advertised sales price, there always exist a finite expected time-on-market to contingency. Let $\tau_{\text{contingency}}$ be the time to reach contingency. Once contingent, the listing has a 0.85 probability of selling. If the contingent listing sells, it takes an additional 60 days from contingency to sale. If it doesn't sell (0.15 probability), the process starts over from the beginning.[30] Let $E[\tau_{\text{sale}}]$ be the expected time to sale. We can express the expected time to sale as:

$$E[\tau_{\text{sale}}] = \tau_{\text{contingency}} + 0.85 \cdot 60 + 0.15 \cdot E[\tau_{\text{sale}}] \tag{33}$$

Solving for $E[\tau_{\text{sale}}]$ and simplifying,

$$E[\tau_{\text{sale}} | \tau_{\text{contingency}}] = 60 + \frac{1}{0.85} \cdot \tau_{\text{contingency}} \tag{34}$$

We can then plug these expected times-to-sale into equation (31) to back out discount factors. Because one period in our model corresponds to a week, our estimates for β reflect *weekly* discount factors. We can convert them to more commonly used *annual* discount factors by simply exponentiating them $\beta_{\text{annual}} = \left(\beta_{\text{weekly}}\right)^{52}$. We then express these annual discount factors as annual discount rates $r_{\text{annual}} = \frac{1}{\beta_{\text{annual}}} - 1$. Figure 19 below displays the implied densities of annual discount factors and discount rates across house-week couples.

[30] We obtain these numbers from our historical Zillow price histories. See Appendix C.2 for details.

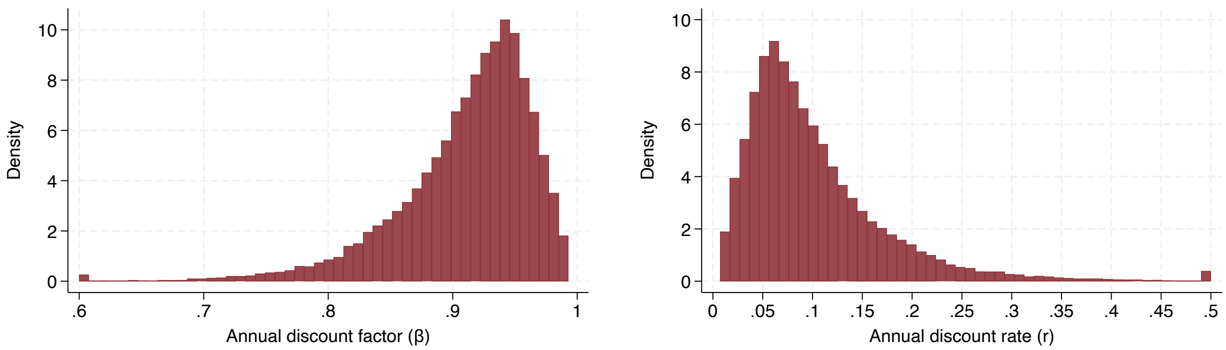

(a) Annual discount factors

(b) Annual discount rates

Figure 19: Implied annual discount factors and discount rates from our estimation results

Notes: The two histograms above plot the entire distribution of implied annual discount factors (β) and annual discount rates (r). One observation is a house-week couple.

Our estimation results imply a median annual discount factor of 92.2% and a median annual discount rate of 8.4%. These numbers are within the typical range found in the literature, particularly for short-term real estate valuations (Giglio, Maggiori and Stroebel, 2014; Bracke, Pinchbeck and Wyatt, 2017), and are consistent with the relatively high interest rates observed during the estimation period (2023–2024).

5 Counterfactuals

Having estimated the parameters of the model, we now turn to counterfactual results.

5.1 Main Results

We first simulate a scenario in which homes owned by institutional investors return to the sector—either rental or homeownership—where they belonged in 2009. This counterfactual, which we call *2009 Ownership*, allows us to quantify the overall welfare effect of purchases by institutional investors. To isolate and quantify the effect of rental concentration, we then consider the case of a policy change inducing institutional landlords to sell *all* their homes to smaller landlords (*All Sold to Landlords*). Finally, to study the potential effects of a large-scale transfer of houses across sectors, we run a counterfactual in which institutional landlords sell their entire portfolio of houses to homeowners (*All Sold to Homeowners*).

In what follows, we describe each counterfactual in detail. Table 8 displays our main counterfactual results.

2009 ownership structure We begin by simulating our main counterfactual, *2009 Ownership*, in which all homes owned by institutional landlords as of 2022 return to their sector of origin as of 2009, before institutional investors started acquiring houses. We use this to quantify the overall welfare effect of purchases by institutional investors. In practice, we implement this counterfactual by allocating each institutionally-owned rental house to the sector it belonged to in 2009. To determine their 2009 sector, we use CoreLogic records. This implies shifting 73% of institutionally-owned homes to the homeownership sector, while keeping 27% of them in the rental sector and assigning them to a small landlord. Column (2) in Table 8 displays the results of this exercise.

We start by looking at the effects in the rental sector. The shift in housing stock across sectors decreases the rental stock by 17.3%. Rental prices increase by 2.4%, or $660 in average yearly rent. Average time-on-market increases only slightly, by 0.9%. The large decrease in housing stock causes transactions to decrease by 20%. As a consequence, average renter surplus decreases by $238 per month, or $2,856 per year. This is a large number, as it represents about 10.4% of the average rental price. For landlords, softer competition increases average annual profits by $612. One interpretation of these results is that if institutional landlords had not entered the market after 2009, renters would be worse off, and landlords would be better off. This is because the effect of increasing rental supply dominates the effect of increasing concentration. Institutional landlords' home acquisition therefore directly benefits renters and hurt landlords.

We now describe the effects in the homeownership sector. Although housing stock only increases by 3.2%, we see large price effects for the average price of homes for sale, which decrease by 4.8%. Time on market decreases by 2.8%, and transactions increase by 2.5%. This implies an increase in homebuyer welfare of around $49,950, which represents 9.1% of average status-quo sales price of $545,000. The additional competition hurts house sellers, whose discounted profits decrease by $22,446. One interpretation of these results is that if institutional landlords had not entered the market after 2010, homebuyers would be better off and home sellers would be worse off.

Table 8: Counterfactual results

		Status Quo (1)	2009 Ownership (2)	All Sold to Landlords (3)	All Sold to Homeowners (4)
Panel A: Housing Stock (occupied + vacant)					
Rental sector	Levels	191,000	158,000	191,000	142,000
	(Δ)	(0%)	(-17.3%)	(0%)	(-25.7%)
Homeownership sector	Levels	1,035,000	1,072,000	1,035,000	1,086,000
	(Δ)	(0%)	(+3.6%)	(0%)	(+4.9%)
Panel B: Prices					
Rental sector (monthly rent)	Levels	\$2,274	\$2,332	\$2,186	\$2,388
	(Δ)	(0%)	(+2.6%)	(-3.9%)	(+5%)
Homeownership sector	Levels	\$544,000	\$519,000	\$544,000	\$512,000
	(Δ)	(0%)	(-4.6%)	(0%)	(-5.9%)
Panel C: Time-on-market (days)					
Rental sector	Levels	67.3	67.2	65.1	67.7
	(Δ)	(0%)	(-0.1%)	(-3.3%)	(+0.6%)
Homeownership sector	Levels	60.0	65.2	60.0	67.6
	(Δ)	(0%)	(+8.7%)	(0%)	(+12.7%)
Panel D: Transactions (per week)					
Rental sector	Levels	226.1	175.1	224.4	153.1
	(Δ)	(0%)	(-22.6%)	(-0.8%)	(-32.3%)
Homeownership sector	Levels	795.4	814.7	795.4	826.4
	(Δ)	(0%)	(+2.4%)	(0%)	(+3.9%)
Panel E: Welfare					
Δ **Consumer Surplus**					
Renters (\$/year)		0	-\$3,027	+\$2,222	-\$5,088
Homebuyers (\$/purchase price)		0	+\$118,894	0	+\$193,057
Δ **Producer Surplus**					
Landlords (\$/year)		0	+\$763	-\$566	+\$1,284
Home sellers (\$/purchase price)		0	-\$17,936	0	-\$25,531

Notes: This table presents counterfactual changes in housing stock, steady-state prices, time-on-market, weekly transactions, and consumer surplus for both the rental and homeownership sectors. Changes are shown relative to the *Status Quo*, shown in Column (1). Counterfactual *2009 Ownership* in column (2) shows a counterfactual where institutionally-owned rental homes return to their historical (2009) sector, with about 27% of them being purchased by landlords, and 73% transferred to homeownership. Column (3) refers to a policy counterfactual where institutionally-owned homes are *All Sold to Landlords* that are not institutional, therefore remaining in the rental sector and reducing rental concentration, without impacting the homeownership sector. Column (4) shows results for a policy counterfactual in which institutionally-owned rental homes are *All Sold to Homeowners*, and are thus all transferred from the rental to the homeownership sector.

All institutional-owned homes sold to landlords To isolate and quantify the effect of rental concentration on renters, counterfactual *All Sold to Landlords* assumes that all institutionally-owned houses are sold to smaller *landlords*. Therefore, all of these houses stay in the rental sector and none shift to the homeownership sector. This approach keeps the homeownership and rental housing stocks unchanged, while reducing concentration in the rental sector. Hence, it isolates the potential effect of decreased rental concentration on renters.

Column (3) in Table 8 shows the results of this simulation. Outcomes for the homeownership sector are identical to outcomes in *Status quo* (Column 1), since this sector is not affected by the policy. In the rental sector, steady-state rents decrease by 3.8%. As rents decrease, average time-on-market also decreases by 3.4%, and rental transactions increase by 3.1%. Overall, the decrease in rental concentration implies an increase in average renter welfare of $186 per renter per month, or $2,232 per year. The reduction of concentration reduces average landlord profits by $924 per year.

All institutional-owned homes sold to homeowners Our final counterfactual, *All Sold to Homeowners*, studies the effect of a policy intervention that induces institutional landlords to sell their entire portfolio of rental houses to *homeowners*. This shifts *all* institutionally-owned homes from the rental to the homeownership sector, and provides an upper bound on the amount of transfers that could occur across sectors.

The results of this counterfactual are shown in column (4) of Table 8. We see that all 50,000 institutionally-owned rental homes are shifted from the rental to the homeownership sector, decreasing the rental stock by 26.6% and increasing homeownership stock by 4.6%. These shifts in sectoral housing stock cause significant price effects on both sectors. In the rental sector, the effect of lower rental supply dominates the effect of lower concentration. Indeed, despite eliminating rental concentration, rents increase by 4.6%. Rental transactions decrease by 30.2%, and average renter welfare decreases by $410 per renter per month, or $4,920 per year. Lower competition among landlords raises mean landlord profits by $1,212 per year. In the homeownership sector, the increase in housing stock lowers sales prices by 6.1% and raises sales transactions by 4.5%. This implies a substantial increase in average homebuyer surplus of $63,680, and a decrease in home seller profits of $35,460.

5.2 Decomposition

Our counterfactual results allow us to quantify the relative importance of the two key forces at play in the rental sector, namely, supply shifts vs. concentration. We illustrate this decomposition in Figures 21 and 22 for changes in the rental price and in renter welfare, respectively.

We begin by decomposing changes in rental prices in Figure 21. Relative to the equilibrium monthly rent in a world in which investors never entered ($2,330, column 2 of Table 8), the monthly rent in the *Status Quo* world with investors ($2,275, column 1 of Table 8) is $55 lower. This 2.4% reduction in rent is the *overall effect* of institutional investors on rents, illustrated by the gray bar in Figure 21. We can decompose this *overall effect* into two forces. To isolate the *concentration effect*, we compare the *Status quo* rents ($2,275, column 1 of Table 8) to the rents in a world without concentration ($2,188, column 3 of Table 8). Concentration increases rents by $87, or 3.7% of the equilibrium monthly rent in a world where investors never entered. This is illustrated by the red bar in Figure 21. Finally, the difference between the concentration effect and the overall effect is the *rental supply effect*, illustrated by the purple bar: higher rental supply lowers monthly rental prices by $140, or 6.1% of the equilibrium monthly rent in a world where investors never entered.

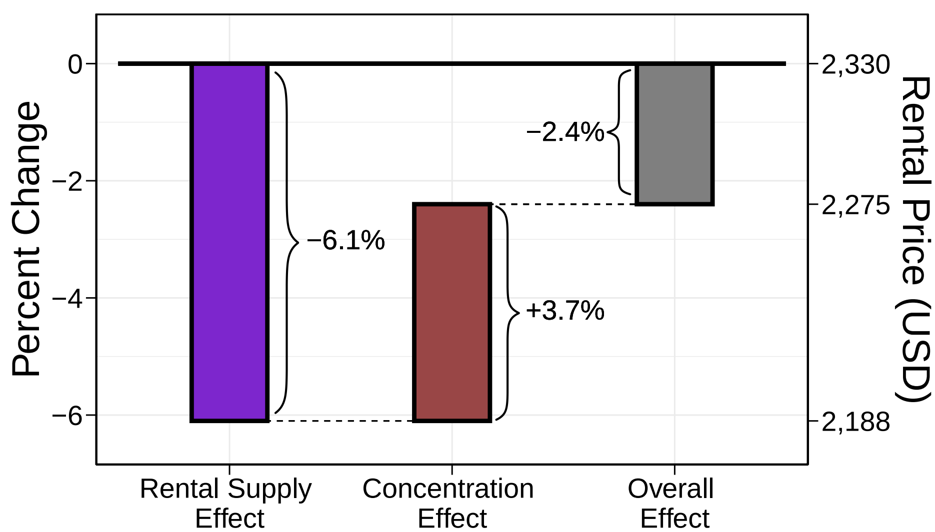

Figure 21: Decomposition of concentration vs. housing supply shift effect (rental price)

Notes: This figure decomposes the rental price effect of institutional acquisitions into two forces. The gray bar shows the overall effect, the red bar shows the effect of concentration and the purple bar shows the effect of increased rental supply.

Similarly, Figure 22 decomposes changes in renter surplus into the two same forces. Panel E of Table 8 shows that, relative to a world in which investors never entered, the average renter gains $2,856 in consumer surplus per year in the *Status Quo*. This *overall effect* is illustrated by the gray bar in Figure 22. To isolate the *concentration effect*, we compute the change in average renter surplus between a world without concentration (*All Sold to Landlords*) and the *Status Quo*: Concentration decreases renter welfare by $2,232 per year, illustrated by the red bar in Figure 22. Again, the difference between the two effects is the *rental supply effect*, illustrated by the purple bar: higher rental supply increases renter welfare by $5,088 per year.

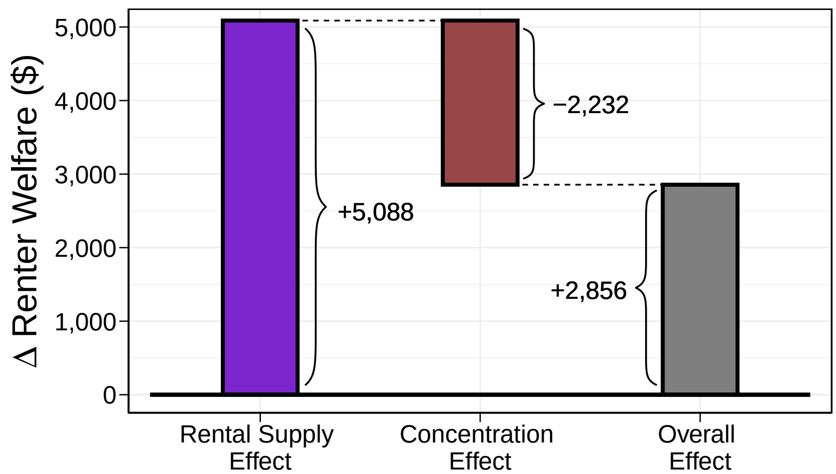

Figure 22: Decomposition of concentration vs. housing supply shift effect (renter surplus)

Notes: This figure decomposes the effect of institutional acquisitions on average renter welfare into two forces. The gray bar shows the overall effect, the red bar shows the effect of concentration and the purple bar shows the effect of increased rental supply.

5.3 Spatial Heterogeneity

The results in Table 8 show average effects across all markets. However, because each individual house has a different estimated price elasticity, and each ZIP code has a different level of institutional concentration, both the counterfactual rental supply effect and the concentration effect will differ by ZIP code. We illustrate this heterogeneity in Figure 23.

Panels (b) and (c) show how the concentration effect and the overall effect for rental prices differ by ZIP code on a map of Atlanta, where we also display the map of the institutional share of single-family rentals in panel (a) for ease of readability. Panels (d) and (e)

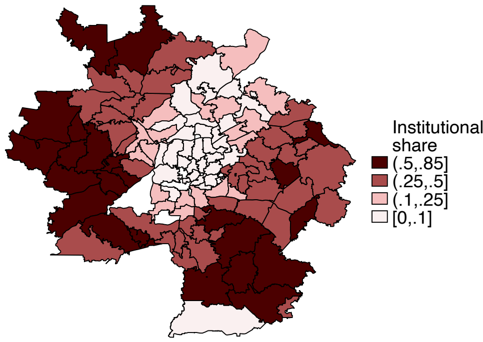

(a) Institutional rental share by ZIP code

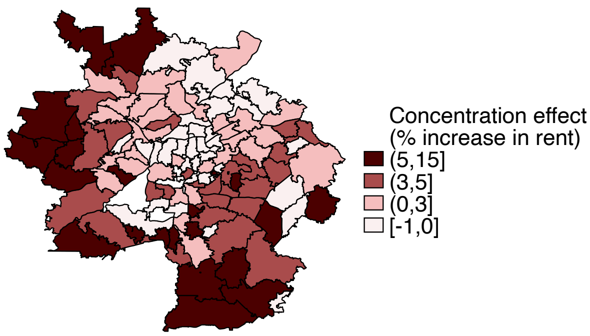

(b) Concentration effect (map)

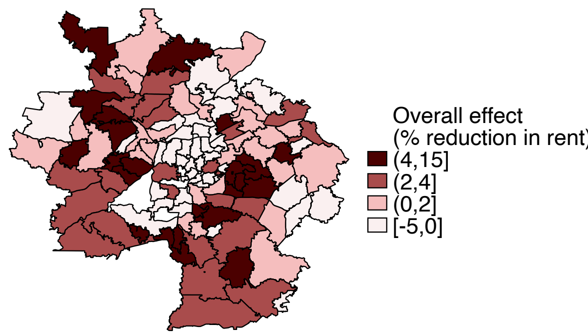

(c) Overall effect (map)

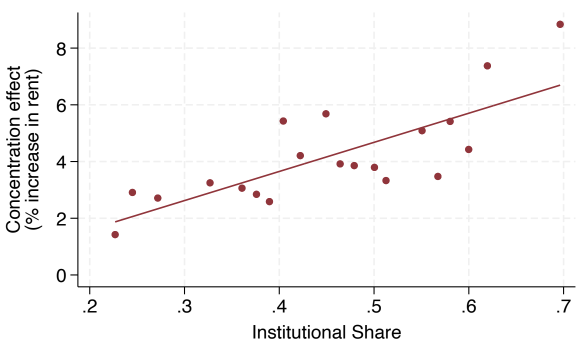

(d) Concentration effect (binscatter)

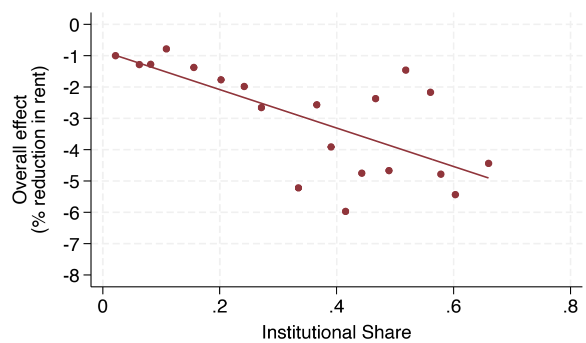

(e) Overall effect (binscatter)

Figure 23: Spatial heterogeneity of price effects from institutional acquisitions

Notes: The figures above illustrate spatial heterogeneity in the price effects of institutional acquisitions. Panel (a) shows the institutional rental share of single-family rental homes by ZIP code. Panels (b) and (c) display maps of the concentration effect and the overall effect on rental prices, respectively. Panels (d) and (e) present binned scatterplots showing how these effects vary with the institutional share of rentals across ZIP codes.

display binned scatterplots showing that these two effects display a monotonic relationship with respect to the share of institutionally-owned listings: although ZIP codes with an institutional share of 70% show a rental price appreciation of more than 8% due to concentration, they show an overall effect of a rental price decrease of around 5%, implying that the rental supply channel lowers rental prices by about 13% in these ZIP codes.

5.4 Robustness analysis

Own-vs-rent tenure choice margin Our model treats consumers as either prospective renters or buyers, without allowing for tenure choice between these options. While this assumption does not affect our estimation results, relaxing it could impact our counterfactuals. For example, in our main counterfactual (*2009 Ownership*), where rental prices increase and sales prices decrease, the key concern is that these price changes could induce renters at the margin between renting and buying to switch towards buying. Specifically, a decrease in the average sales price could lower demand for renting, and an increase in the average rental price could raise demand for buying, thus moderating the effects we find in our *2009 Ownership* counterfactual.

To bound the potential impact of tenure choice on our results, we examine the following extreme-case scenario. First, we assume all renters who switch to the outside option following rent increases in *2009 Ownership* become buyers, which we model as an increase in the homeownership market size. Second, we assume all buyers who switch from the outside option to the inside nest following sales price decreases in *2009 Ownership* were previously renters, and model this as a decrease in the rental market size. This represents an upper bound on cross-sector substitution: In reality, prospective renters who choose not to rent a single-family home might choose options other than buying, such as renting an apartment or staying put. By examining this extreme case, we can bound the extent to which tenure choice might moderate our main findings.

Since demand for rentals is expressed in contacts while demand for house sales is expressed in page views, we need a method to convert between these measures to analyze cross-market substitution. For each market, we proceed as follows. First, we identify the number of rental contacts $\kappa_{I,O}$ that switch from the inside to the outside option following rent increases in the *2009 Ownership* counterfactual. We calculate what share of the overall rental market size this represents: $s^R_{I,O} = \kappa_{I,O}/M_R$. Second, assuming a constant homeownership rate of $2/3$ across markets, we calculate the corresponding increase in the homeownership market size if all these renters became buyers. Since the rental market is half the size of the ownership market, this share in the homeownership sector becomes $s^H_{I,O} = s^R_{I,O}/2$. This allows us to back out the corresponding increase in the number of page views: $\nu^+_{O,I} = s^H_{I,O} \cdot M_S$ (where M_S is the size of the homeownership market). Conversely, we record the number of homeownership page views that move from the outside to the inside option, $\nu_{O,I}$, compute their share of the homeownership market $s^H_{O,I} = \nu_{O,I}/M_S$, calculate the corresponding share that this represents in the rental market: $s^R_{O,I} = 2 \cdot s^H_{O,I}$, and back out the corresponding decrease in the number of rental contacts: $\kappa^+_{O,I} = s^R_{O,I} \cdot M_R$. We then re-compute the steady-state equilibrium in our main "2009 Ownership" counterfactual for both sectors. The results are displayed in column (2b) of Table 9.

We see that even when in the extreme-case scenario where all homeownership-sector individuals who switch to the outside option choose to rent, and vice versa, allowing for an own-vs-rent margin does dampen our results but the overall impact on our key results is small.

In the *rental* sector, the steady-state average rent reaches \$2,321, a 2.0% increase over the status-quo rent of \$2,275. By comparison, in our main "2009 Ownership" counterfactual, average rents rise to \$2,330 — an increase of 2.4% relative to the status quo rent. Average renter's surplus falls by \$2,840 rather than \$2,856. In the *homeownership* sector, the average sales price drops by 4.2% (baseline 4.8% decrease), weekly sales transactions grow by 2.2% (baseline 2.5% increase), and the average homebuyer's surplus improves by \$43,093 instead of \$49,950.

In other words, allowing for an own-vs-rent margin slightly attenuates the magnitudes of our key counterfactual findings, but the overall direction and economic significance of our results remain intact.

5.5 Discussion

Our counterfactual analysis provides four key takeaways. First, institutional purchases benefit renters in overall welfare terms. Second, this overall impact on renters holds because the welfare-improving effect of higher rental supply dominates the welfare-decreasing

Table 9: Counterfactual results: robustness to tenure choice

		Status Quo	2009 Ownership	2009 Ownership **(+ own vs rent)**
		(1)	(2a)	(2b)
Panel A: Housing Stock (occupied + vacant)				
Rental sector	Levels	191,000	158,000	158,000
	(Δ)	(0%)	(–17.3%)	(–17.3%)
Homeownership sector	Levels	1,035,000	1,068,000	1,068,000
	(Δ)	(0%)	(+3.2%)	(+3.2%)
Panel B: Prices				
Rental sector (monthly rent)	Levels	$2,275	$2,330	$2,321
	(Δ)	(0%)	(+2.3%)	(+2.0%)
Homeownership sector	Levels	$545,000	$519,000	$522,000
	(Δ)	(0%)	(–4.8%)	(–4.2%)
Panel C: Transactions (per week)				
Rental sector	Levels	232.8	186.3	186.6
	(Δ)	(0%)	(–20.0%)	(–19.8%)
Homeownership sector	Levels	795.4	815.0	813.0
	(Δ)	(0%)	(+2.5%)	(+2.2%)
Panel D: Consumer Surplus				
Renters ($/year)		0	–$2,856	–$2,840
Homebuyers ($/purchase price)		0	+$49,950	+$43,093

Notes: This table reports counterfactual changes in housing stock, prices, weekly transactions, and consumer surplus for the rental and home-ownership sectors. All percentage changes in parentheses are expressed relative to the *Status Quo* in column (1). Column (2a) reproduces our main *2009 Ownership* counterfactual. Column (2b) applies the same reallocation but adds an extreme-case tenure-choice adjustment: every renter priced out by higher rents buys a home, and every new buyer induced by lower purchase prices would otherwise have rented. Comparing columns (2a) and (2b) therefore bounds the extent to which allowing an own-vs-rent margin can moderate the headline effects.

effect of higher concentration. Third, this increase in renter welfare comes at the expense of homebuyers, for whom the decreased housing supply negatively affects welfare. Fourth, institutional purchases cause an opposite distributional effect for producer surplus: home sellers benefit, while landlords are worse off.

6 Conclusion

This paper studies the welfare effects of institutional ownership in the single-family housing market. To quantify these effects, we construct a novel dataset tracking rental prices, institutional ownership, and high-frequency proxies of demand at the individual house level for the Atlanta metropolitan area. Our dataset provides several new insights on the evolution of rental concentration during the last decade, on the geographical distribution of institutional portfolios, and on the spatial heterogeneity of institutional concentration across different neighborhoods in Atlanta. We build an equilibrium model of the housing market, featuring two sectors: rental and homeownership. Using the granularity of the data, we propose a new identification strategy for the price sensitivity of consumers, exploiting discontinuities in high-frequency demand proxies around price changes. Our model allows us to quantify two channels through which institutional purchases of homes affect rental prices. First, changes in rental concentration, which involve both changes in market power and changes in marginal costs. Second, housing reallocation across the rental and homeownership sectors. Our estimates indicate that institutional acquisitions of single-family homes generate opposing welfare effects on renters and homebuyers, benefiting renters and hurting homebuyers. In the rental sector, institutional purchases increase concentration and hurt renters, but the welfare-improving effect of additional rental supply dominates the effect of concentration, leading to a positive net impact on renter welfare.

References

Agarwal, Nikhil and Paulo Somaini, "Demand Analysis under Latent Choice Constraints," *The Review of Economic Studies*, 10 2025, p. rdaf093.

Ahlfeldt, Gabriel M. and Daniel P. McMillen, "Tall Buildings and Land Values: Height and Construction Cost Elasticities in Chicago, 1870–2010," *The Review of Economics and Statistics*, March 2018, *100* (5), 861–875.

Albouy, David and Gabriel Ehrlich, "Housing Productivity and the Social Cost of Land-Use Restrictions," *Journal of Urban Economics*, September 2018, *107*, 101–120.

Almagro, Milena and Tomas Dominguez-Iino, "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," Working Paper, March 2024.

Ater, Itai, Yael Elster, and Eran B Hoffmann, "Real-Estate Investors, House Prices and Rents: Evidence from Capital-Gains Tax Changes," Working Paper, December 2022.

Austin, Neroli, "Keeping Up with the Blackstones: Institutional Investors and Gentrification," Working Paper, November 2022.

Azar, José, Steven Berry, and Ioana Elena Marinescu, "Estimating Labor Market Power," Working Paper, July 2022.

Bajari, Patrick and C. Lanier Benkard, "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," *Journal of Political Economy*, December 2005, *113* (6), 1239–1276.

__, Phoebe Chan, Dirk Krueger, and Daniel Miller, "A Dynamic Model of Housing Demand: Estimation and Policy Implications," *International Economic Review*, 2013, *54* (2), 409–442.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, August 2007, *115* (4), 588–638.

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

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

__ and Philip A. Haile, "Foundations of Demand Estimation," in Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri, eds., *Handbook of Industrial Organization*, Vol. 4 of *Handbook of Industrial Organization, Volume 4*, Elsevier, January 2021, pp. 1–62.

Billings, Stephen B. and Adam Soliman, "The Erosion of Homeownership and Minority Wealth," Working Paper, November 2023.

Bracke, Philippe, Edward W. Pinchbeck, and James Wyatt, "The Time Value of Housing: Historical Evidence on Discount Rates," *The Economic Journal*, 08 2017, *128* (613), 1820–1843.

Buchholz, Nicholas, "Spatial Equilibrium, Search Frictions, and Dynamic Efficiency in the Taxi Industry," *The Review of Economic Studies*, 09 2021, *89* (2), 556–591.

Calder-Wang, Sophie, "The Distributional Impact of the Sharing Economy on the Housing Market," Working Paper, July 2021.

Calder-Wang, Sophie and Gi Heung Kim, "Algorithmic Pricing in Multifamily Rentals: Efficiency Gains or Price Coordination?," *SSRN*, August 2024.

Cardell, N. Scott, "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity," *Econometric Theory*, 1997, *13* (2), 185–213.

Castillo, Juan Camilo, "Who Benefits from Surge Pricing?," *SSRN*, 2023.

— and Shreya Mathur, "Matching and Network Effects in Ride-Hailing," *AEA Papers and Proceedings*, May 2023, *113*, 244–47.

Christophers, Brett, "How and Why U.S. Single-Family Housing Became an Investor Asset Class," *Journal of Urban History*, 2023, *49* (2), 430–449.

Coven, Joshua, "The Impact of Institutional Investors on Homeownership and Neighborhood Access," *Department of Finance, Stern School of Business, New York University mimeo*, November 2024.

Dachis, Ben, Gilles Duranton, and Matthew A. Turner, "The effects of land transfer taxes on real estate markets: evidence from a natural experiment in Toronto," *Journal of Economic Geography*, 05 2011, *12* (2), 327–354.

Demers, Andrew and Andrea L. Eisfeldt, "Total returns to single-family rentals," *Real Estate Economics*, 2022, *50* (1), 7–32.

Diamond, Rebecca and Tim McQuade, "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development," *Journal of Political Economy*, October 2018, *127* (3), 1063–1117.

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

Duranton, Gilles and Diego Puga, "Urban Growth and Its Aggregate Implications," *NBER Working Paper 26591*, December 2019.

Francke, Marc, Lianne Hans, Matthijs Korevaar, and Sjoerd van Bekkum, "Buy-to-Live vs. Buy-to-Let: The Impact of Real Estate Investors on Housing Costs and Neighborhoods," Working Paper, June 2023.

Fréchette, Guillaume R., Alessandro Lizzeri, and Tobias Salz, "Frictions in a Competitive, Regulated Market: Evidence from Taxis," *American Economic Review*, August 2019, *109* (8), 2954–92.

Fukasawa, Takeshi, "The Biases in Applying Static Demand Models Under Dynamic Demand," *Review of Industrial Organization*, 2024, *65* (2), 561–594.

Gabriel, Stuart A. and Frank E. Nothaft, "Rental Housing Markets, the Incidence and Duration of Vacancy, and the Natural Vacancy Rate," *Journal of Urban Economics*, January 2001, *49* (1), 121–149.

Gandhi, Amit and Aviv Nevo, "Chapter 2 - Empirical models of demand and supply in differentiated products industriesWe thank Pierre Dubois, Phil Haile, Ali Hortaçsu, Felipe Kup Barbieri de Matos, Jing Tao and three referees for helpful comments and insightful suggestions.," in Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri, eds., *Handbook of Industrial Organization, Volume 4*, Vol. 4 of *Handbook of Industrial Organization*, Elsevier, 2021, pp. 63–139.

— **and Jean-François Houde**, "Measuring Substitution Patterns in Differentiated-Products Industries," Working Paper 26375, National Bureau of Economic Research October 2019.

Ganduri, Rohan, Steven Chong Xiao, and Serena Wenjing Xiao, "Tracing the Source of Liquidity for Distressed Housing Markets," *Real Estate Economics*, 2023, *51* (2), 408–440.

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

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel, " Very Long-Run Discount Rates *," *The Quarterly Journal of Economics*, 11 2014, *130* (1), 1–53.

Glaeser, Edward and Joseph Gyourko, "The Economic Implications of Housing Supply," Technical Report w23833, National Bureau of Economic Research September 2017.

Grieco, Paul L E, Charles Murry, and Ali Yurukoglu, "The Evolution of Market Power in the U.S. Automobile Industry*," *The Quarterly Journal of Economics*, 09 2023, *139* (2), 1201–1253.

Gurun, Umit G, Jiabin Wu, Steven Chong Xiao, and Serena Wenjing Xiao, "Do Wall Street Landlords Undermine Renters' Welfare?," *The Review of Financial Studies*, January 2023, *36* (1), 70–121.

Hall, Robert E, "New Evidence on the Markup of Prices over Marginal Costs and the Role of Mega-Firms in the US Economy," Working Paper 24574, National Bureau of Economic Research May 2018.

Hanson, Sebastian, "Institutional investors in the market for single-family housing: Where did they come from, where did they go?," *SSRN*, 2023.

Hanushek, Eric A. and John M. Quigley, "What is the Price Elasticity of Housing Demand?," *The Review of Economics and Statistics*, 1980.

Hortaçsu, Ali and Chad Syverson, "The Ongoing Evolution of US Retail: A Format Tug-of-War," *Journal of Economic Perspectives*, November 2015, *29* (4), 89–112.

Lambie-Hanson, Lauren, Wenli Li, and Michael Slonkosky, "Real Estate Investors and the U.S. Housing Recovery," *Real Estate Economics*, 2022, *50* (6), 1425–1461.

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, "The Rise of Market Power and the Macroeconomic Implications*," *The Quarterly Journal of Economics*, 01 2020, *135* (2), 561–644.

Miller, Nathan, Matthew Osborne, Gloria Sheu, and Gretchen Sileo, "Technology and Market Power: The United States Cement Industry, 1974-2019," *SSRN*, June 29 2023, (4041168).

Mills, James, Raven Molloy, and Rebecca Zarutskie, "Large-Scale Buy-to-Rent Investors in the Single-Family Housing Market: The Emergence of a New Asset Class," *Real Estate Economics*, 2019, *47* (2), 399–430.

Moszkowski, Erica and Daniel Stackman, "Option Value and Storefront Vacancy in New York City," Meyer Fellowship Paper, Joint Center for Housing Studies of Harvard University May 2023.

Nevo, Aviv, "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, March 2001, *69* (2), 307–342.

Ortalo-Magné, François and Andrea Prat, "On the Political Economy of Urban Growth: Homeownership versus Affordability," *American Economic Journal: Microeconomics*, February 2014, *6* (1), 154–181.

Smith, Dominic A. and Sergio Ocampo, "The Evolution of U.S. Retail Concentration," Working Paper 526, U.S. Bureau of Labor Statistics, Office of Prices and Living Conditions January 11 2021. Available from the U.S. Department of Labor, Bureau of Labor Statistics.

Smith, Patrick S. and Crocker H. Liu, "Institutional Investment, Asset Illiquidity and Post-Crash Housing Market Dynamics," *Real Estate Economics*, 2020, *48* (3), 673–709.

Watson, C. Luke and Oren Ziv, "Is the Rent Too High? Land Ownership and Monopoly Power," Working Paper, December 2023.

Appendix

A Tables and Figures

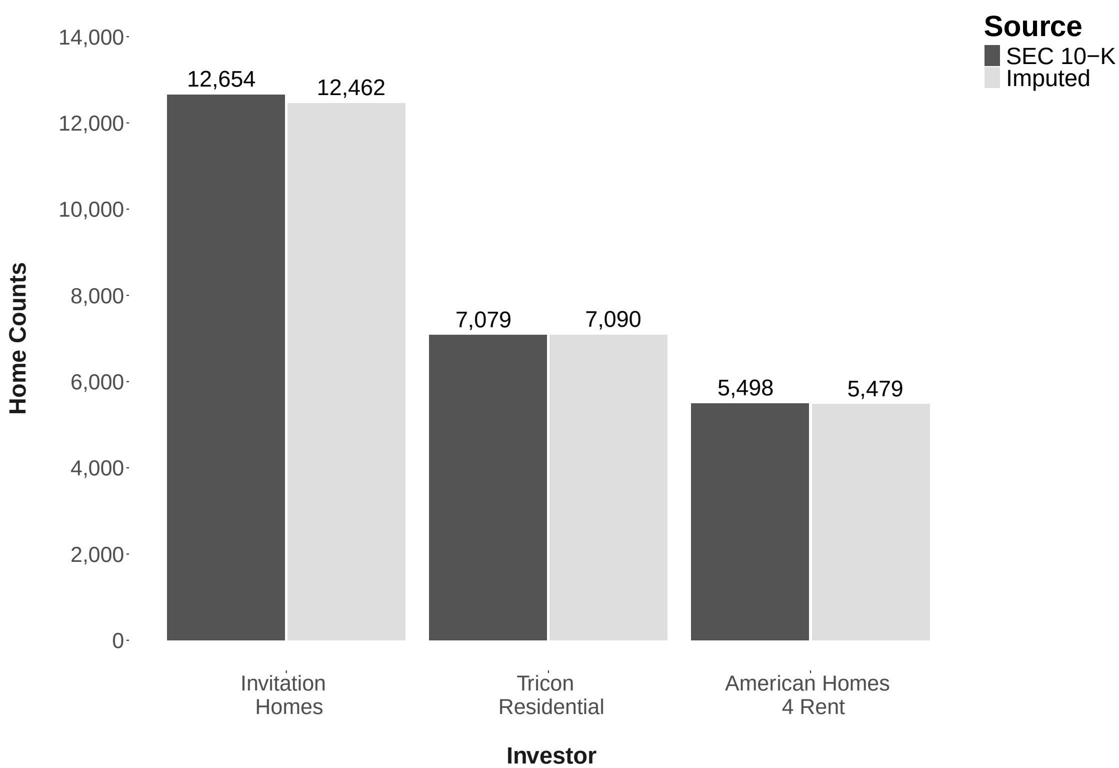

Figure A1: Imputed vs. true institutional portfolio sizes

Notes: This figure shows counts for the number of houses owned by three publicly-listed institutional landlords within Atlanta: Invitation Homes (NYSE:INVH), Tricon Residential (NYSE:TCN), and American Homes 4 Rent (NYSE:AMH). The black bars represent their true Atlanta portfolio sizes, according to their 2022 annual 10-K reports to the S.E.C. The light gray bars represent our counts of homes for each of these investors, based on our mapping of subsidiary companies listed on tax records to their parent institutional owners.

Table A1: Number of rental properties for top 10 largest landlords in Atlanta (2022)

Institutional Landlord	**Number of Properties**	
	Atlanta	**U.S.**
Progress Residential	14,750	70,000
Invitation Homes	12,462	72,245
Main Street Renewal (Amherst)	7,455	45,000
Tricon Residential	7,090	35,908
FirstKey Homes	6,929	32,000
American Homes 4 Rent	5,479	57,878
Home Partners of America	3,824	17,000
Divvy Homes	1,749	7,000
Sylvan Homes	785	4,000
Vinebrook Homes	624	2,838

Notes: This table shows local and national portfolio sizes for the 10 largest institutional landlords in Atlanta. For publicly listed landlords, both Atlanta and U.S. counts are based on 2022 S.E.C. annual 10-K reports. For private landlords, Atlanta counts are based on our own calculations, and U.S. counts are based on numbers from investors' own websites. The total single-family rental housing stock in Atlanta is approximately 190,000 as of 2022.

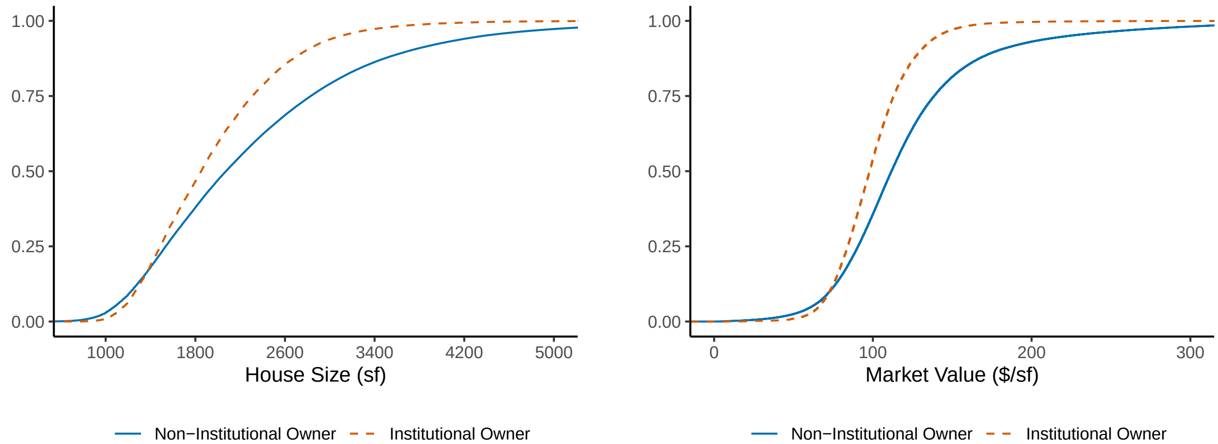

(a) House size by owner type

(b) House value per square foot by owner type

Figure A2: House characteristic distribution by owner type

Notes: Panel (a) plots the empirical CDFs of the distributions of house square footage for institutionally and non-institutionally owned single-family homes in the Atlanta MSA in 2021. Panel (b) plots the corresponding empirical CDFs of assessed market value per square foot.

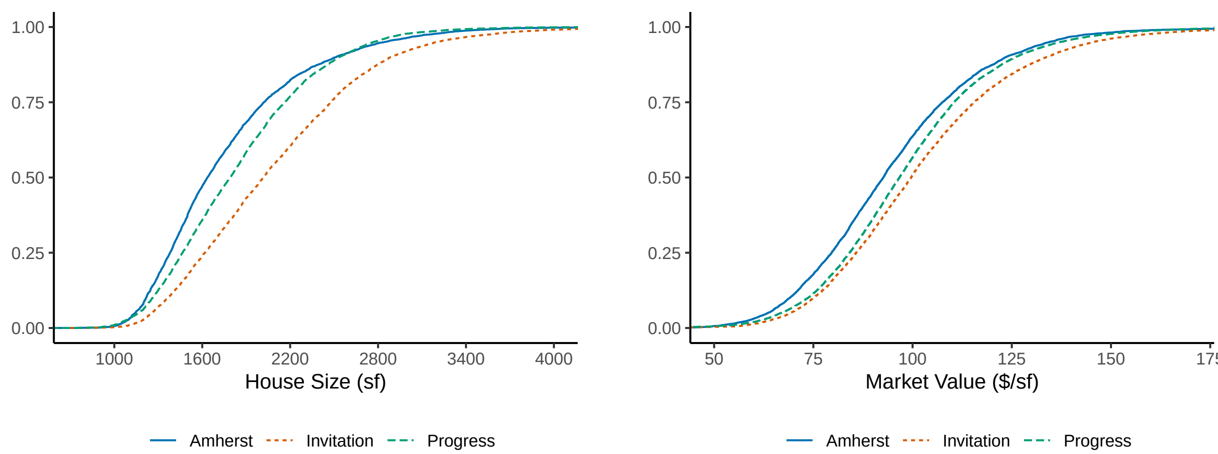

(a) House size by landlord

(b) House value per square foot by landlord

Figure A3: House characteristic distribution by institutional landlord

Notes: Panel (a) and Panel (b) plot the CDFs of the house square footage and house value per square foot distributions, respectively, for houses owned by the three largest single-family landlords in our sample: Invitation Homes, Progress Residential, and Amherst Residential (Main Street Renewal).

Figure A4: Single-Family House Ownership Shares

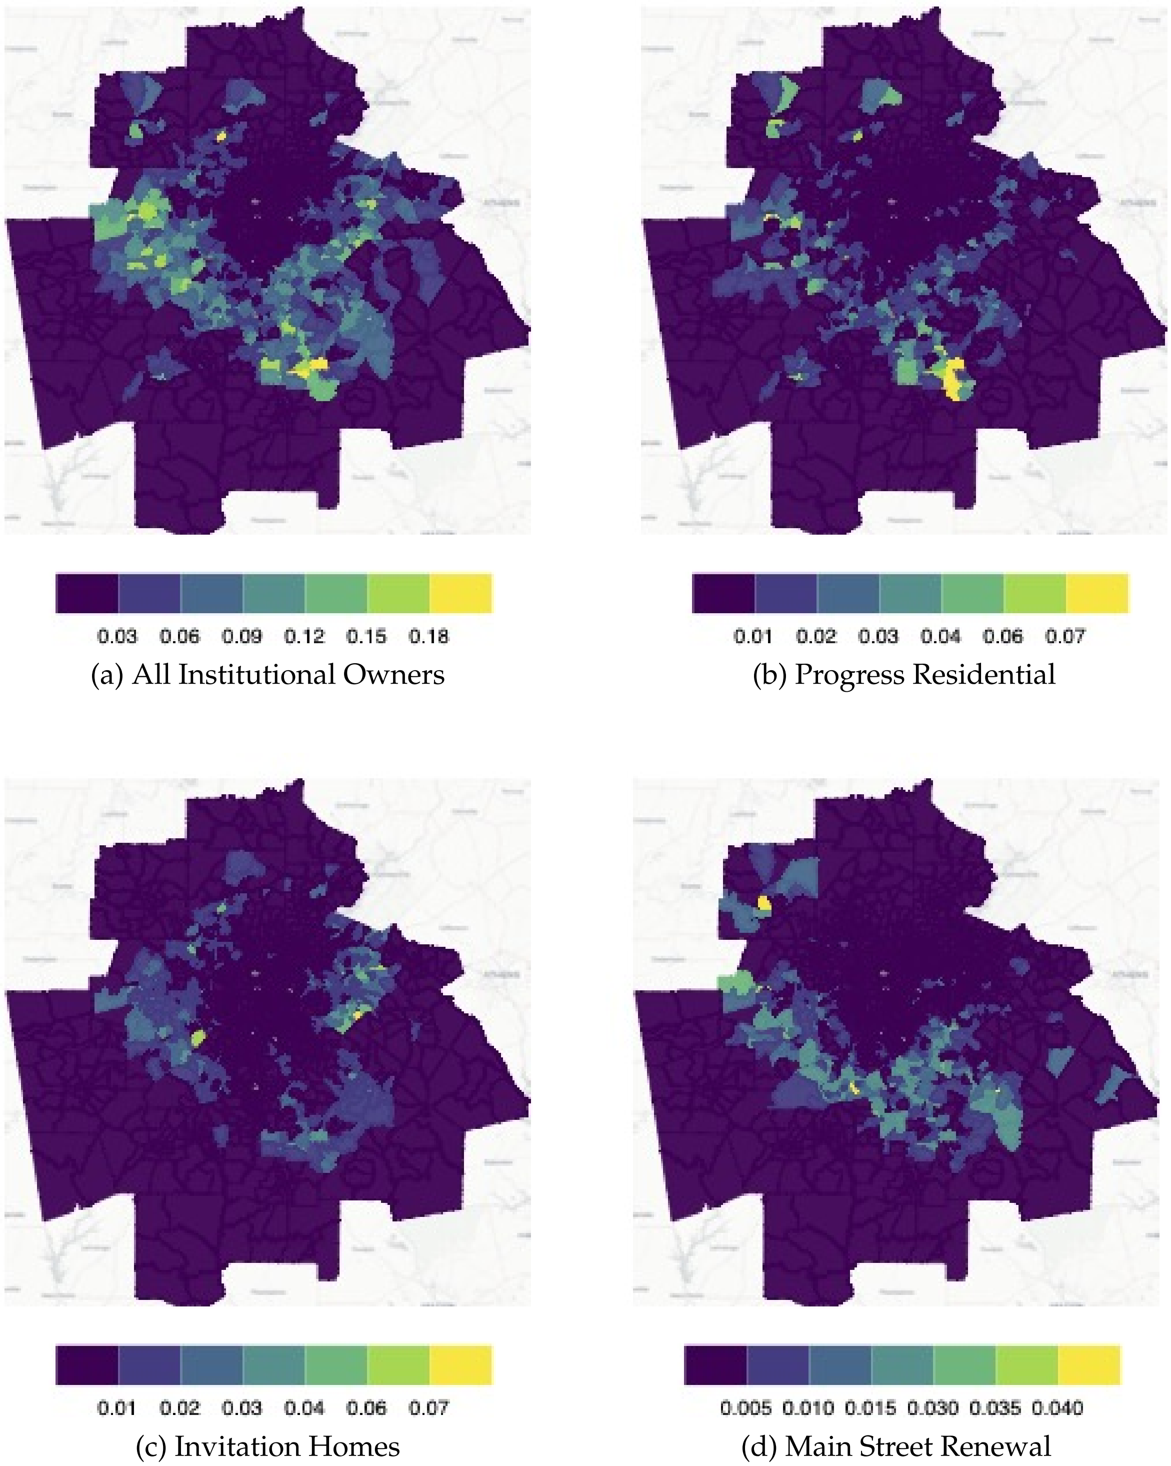

(a) All Institutional Owners

(b) Progress Residential

(c) Invitation Homes

(d) Main Street Renewal

Notes: Panel (a) plots the share of single-family homes owned by any institutional landlord for all Census tracts in the Atlanta MSA. Panel (b) plots the corresponding ownership shares for Progress Residential. Panel (c) plots ownership shares for Invitation Homes. Panel (d) plots ownership share for Main Street Renewal.

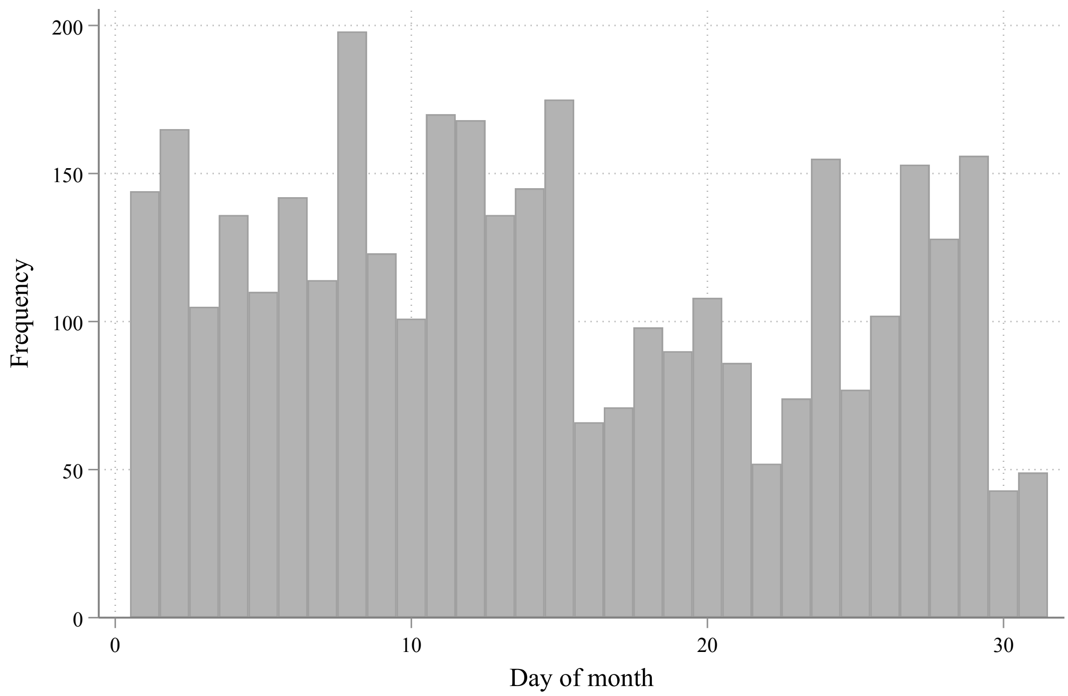

Figure A5: Rental price change frequency by day of the month

Notes: This figure shows the frequency of rental price by the day of month changes for houses in our sample that we observe changing their price at least once during their advertisement spell.

B Analytical derivatives and elasticities

B.1 Rental sector

Below we present the derivation of the occupancy derivatives shown in Equation (24). First, note that the price derivative of time-on-market is given by:

$$\frac{\partial \tau_j(\boldsymbol{p})}{\partial p_k} = \frac{d\tau_j(s_j(\boldsymbol{p}))}{dp_k} = \frac{\partial \tau_j}{\partial s_j}\frac{\partial s_j}{\partial p_k} = \underbrace{\frac{\partial \tau_j}{\partial s_j(\boldsymbol{p})}\frac{s_j}{\tau_j(\boldsymbol{p})}}_{=\varepsilon_{jj}^{\tau,\kappa}} \underbrace{\frac{\partial s_j(\boldsymbol{p})}{\partial p_k}\frac{p_k}{s_j(\boldsymbol{p})}}_{=\varepsilon_{jk}^{\kappa,p}} \frac{\tau(\boldsymbol{p})}{p_k} = \varepsilon_{jj}^{\tau,\kappa}\,\varepsilon_{jk}^{\kappa,p}\,\frac{\tau(\boldsymbol{p})}{p_k} \tag{35}$$

where $\varepsilon_{jj}^{\tau,\kappa}$ is the elasticity of time-on-market τ with respect to contacts κ, a structural parameter which derives from equation 3.

Second, note that the price derivative of the occupancy rate is therefore given by:

$$\frac{\partial \Omega(\boldsymbol{p})}{\partial p_k} = \frac{\partial}{\partial p_k}\left(\frac{T_j}{T_j + \tau_j(\boldsymbol{p})}\right) = \frac{-T_j\left(\frac{\partial \tau_j(\boldsymbol{p})}{\partial p_k}\right)}{(T_j + \tau_j(\boldsymbol{p}))^2} = -\Omega_j(\boldsymbol{p})\ \underbrace{\frac{\tau(\boldsymbol{p})}{T_j + \tau(\boldsymbol{p})}}_{=1-\Omega_j(\boldsymbol{p})} \frac{\varepsilon_{jj}^{\tau,\kappa}\,\varepsilon_{jk}^{\kappa,p}}{p_k} \tag{36}$$

Third, note that equation 2 implies:

$$\frac{\partial s_j}{\partial p_k} = \begin{cases} -\alpha^R s_j\left(\frac{1}{1-\sigma^R} - \frac{\sigma^R}{1-\sigma^R}s_{j|g} - s_j\right) & \text{if } j = k \\ \alpha^R s_k\left(\frac{\sigma^R}{1-\sigma^R}s_{j|g} + s_j\right) & \text{if } j \neq k \end{cases} \tag{37}$$

We can therefore express the elasticity of contacts with respect to the rental price as:

$$\varepsilon_{jk}^{\kappa,p} := \frac{\partial s_j}{\partial p_k}\frac{p_k}{s_j} = \begin{cases} -\alpha^R p_j\left(\frac{1}{1-\sigma^R} - \frac{\sigma^R}{1-\sigma^R}s_{j|g} - s_j\right) & \text{if } j = k \\ \alpha^R p_k\left(\frac{\sigma^R}{1-\sigma^R}s_{j|g} + s_j\right) & \text{if } j \neq k \end{cases} \tag{38}$$

Combining equations 36 and 38, we obtain the following formula for the price derivatives of occupancy:

$$\frac{\partial \Omega_j}{\partial p_k} = \begin{cases} \Omega_j(1-\Omega_j)\ \alpha^R\ \varepsilon_{jj}^{\tau}\ \left(\frac{1}{1-\sigma^R} - \frac{\sigma^R}{1-\sigma^R}s_{j|g} - s_j\right) & \text{if } j = k \\ -\Omega_j(1-\Omega_j)\ \alpha^R\ \varepsilon_{jj}^{\tau}\ \left(\frac{\sigma^R}{1-\sigma^R}s_{j|g} + s_j\right) & \text{if } j \neq k \end{cases} \tag{39}$$

which is Equation (24).

For completeness, we also derive the own- and cross-price elasticities of house j's occupancy with respect to house k's rental price:

$$\varepsilon^{\Omega}_{jk} := \frac{\partial \Omega_j}{\partial p_k}\frac{p_k}{\Omega_j} = \begin{cases} -(1-\Omega_j)\,\varepsilon^{\tau,\kappa}_{jj}\,\varepsilon^{\kappa,p}_{jj} & \text{if } j = k \\ (1-\Omega_j)\,\varepsilon^{\tau,\kappa}_{jj}\,\varepsilon^{\kappa,p}_{jk} & \text{if } j \neq k \end{cases} \tag{40}$$

B.2 Homeownership sector

Below we present the derivation of the analytical expressions for the elasticities shown in Equation (30). Since we assume house sellers price as single-product firms, we only need the expressions for own-price elasticities. Note that equation 8 implies:

$$\frac{\partial s_j}{\partial p_j} = -\alpha^H s_j \left(\frac{1}{1-\sigma^H} - \frac{\sigma^H}{1-\sigma^H} s_{j|g} - s_j\right) \tag{41}$$

Hence, the elasticity of the number of views with respect to price $\varepsilon^{\nu,p}_{jj}$ is given by

$$\varepsilon^{\nu,p}_{jj} = \frac{\partial s_j}{\partial p_j}\frac{p_j}{s_j} = -\alpha^H p_j \left(\frac{1}{1-\sigma^H} - \frac{\sigma^H}{1-\sigma^H} s_{j|g} - s_j\right) \tag{42}$$

Therefore, the elasticity of time-on-market with respect to the sales price $\varepsilon^{\tau,p}_{jj}$ can be written as

$$\varepsilon^{\tau,p}_{jj} = \varepsilon^{\tau,\nu}_{jj}\,\varepsilon^{\nu,p}_{jj} = -\varepsilon^{\tau,\nu}_{jj}\,\alpha^H\,p_j \left(\frac{1}{1-\sigma^H} - \frac{\sigma^H}{1-\sigma^H} s_{j|g} - s_j\right) \tag{43}$$

which is Equation (30).

C Data Appendix

This Appendix describes our data in detail. Subsections C.1 and C.2 describe our "real-time" and "historical" Zillow data, respectively, as well as details about the data collection process we used to create each dataset.

C.1 Real-time Zillow data

We gather data by scraping Zillow in real time between May 2023 and September 2024. We observe daily snapshots of Zillow listings of single-family homes for sale and for rent in metropolitan Atlanta. To narrow down our selection of listings to single-family homes, we filter both rental and sales listings using the "Home Type" filter to include "House" listings only, as shown in Figure Figure A6 below. For logistical reasons, we gather real-time data for all listings in the 11 most central counties of the Atlanta MSA[31] once every day. A detailed map of our coverage area is shown in Figure Figure A10.

We visit the webpage of each house advertised for rent or for sale. Figure Figure A7 shows two examples of such listings. We observe detailed house characteristics, such as the advertised rental and sales price, number of bedrooms and bathrooms, square footage, street address, written description, as well as several other detailed features. Figure Figure A8 displays the distribution of the property area of houses listed for rent in our sample, and Figure Figure A9 shows the distribution of advertised rental prices in U.S. Dollars per month. We also observe the *category* of landlord in charge of advertising and managing the property. Figure Figure A11 splits listing managers into institutional landlords, houses listed by individual owners, and houses whose listing is managed by a real estate rental management company.

[31] These counties are: Cherokee, Clayton, Cobb, DeKalb, Douglas, Fulton, Gwinnett, Henry, Newton, Paulding and Rockdale.

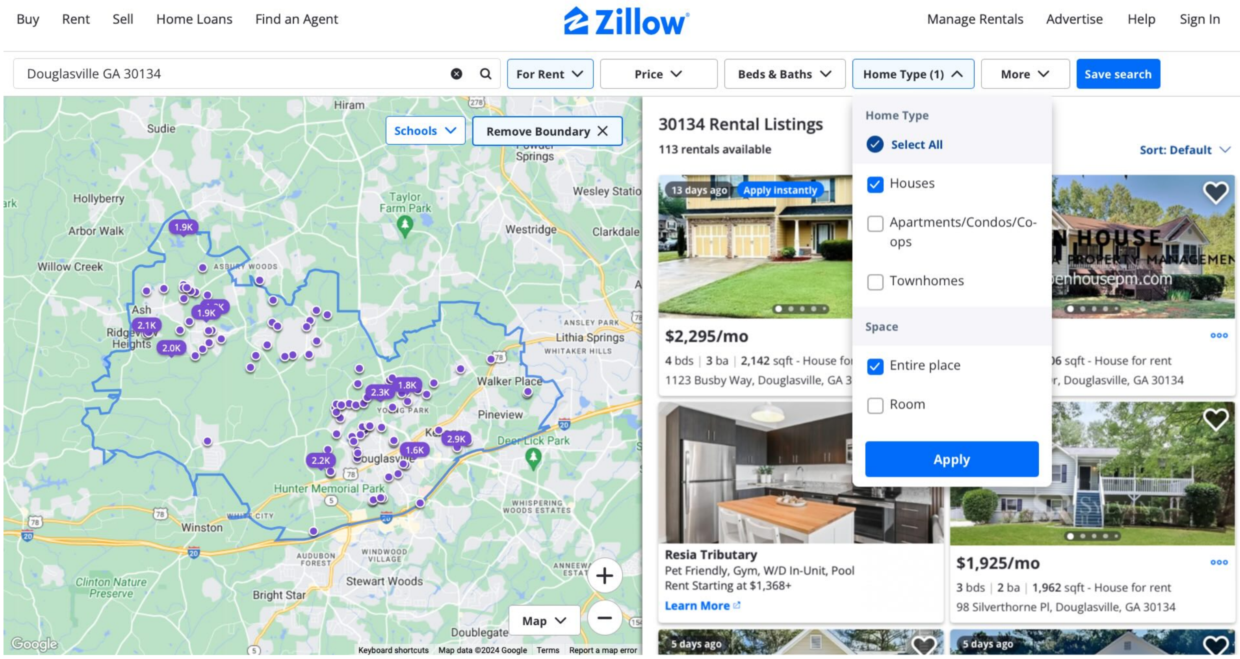

(a) For rent listings

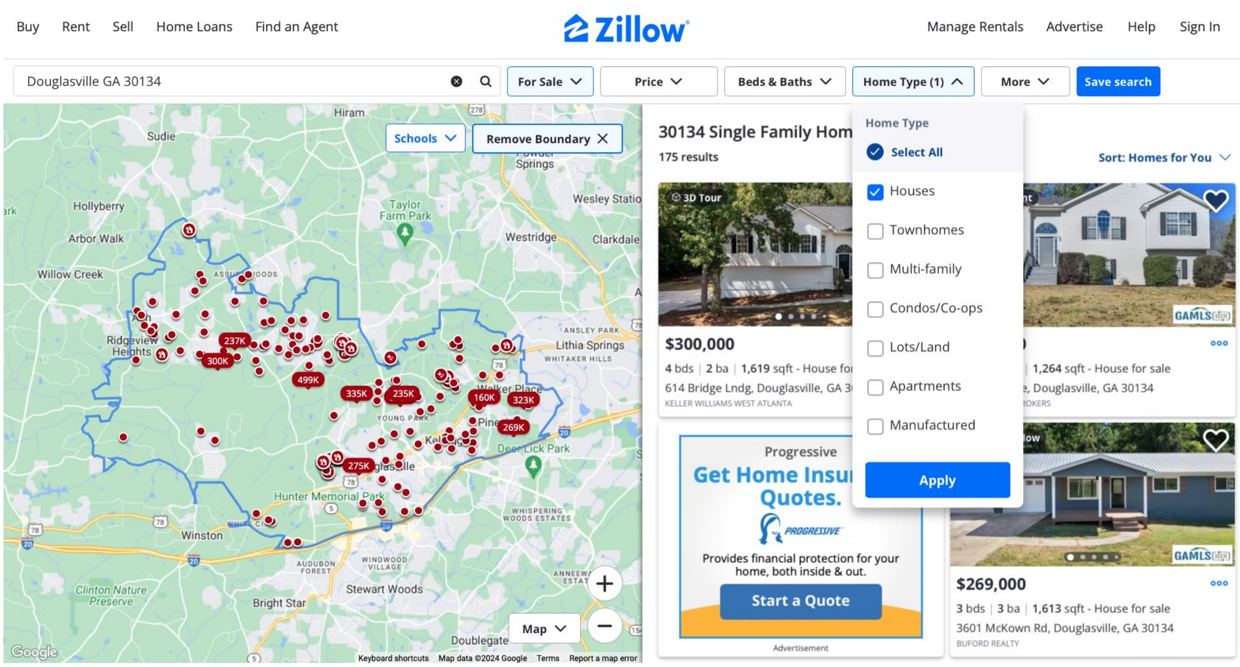

(b) For sale listings

Figure A6: Examples of Zillow search page for single-family home listings for rent and for sale, for one ZIP Code

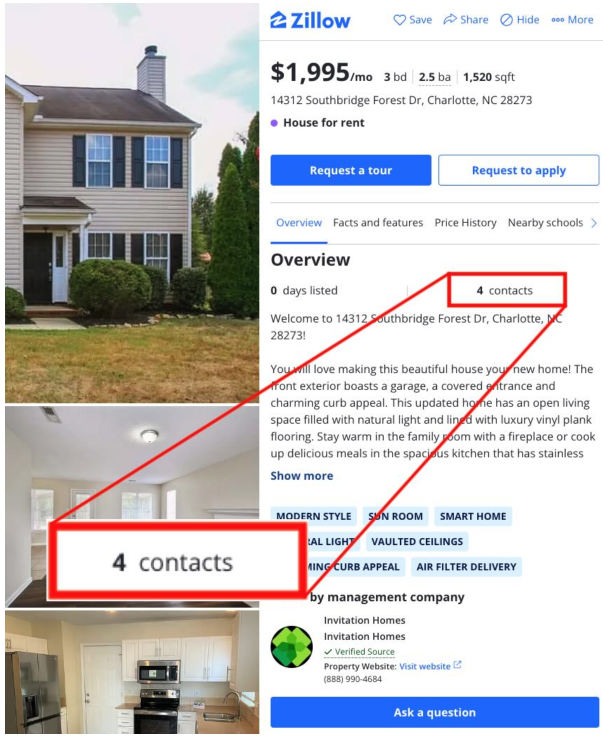

(a) Example of listing for rent (with number of customer contacts)

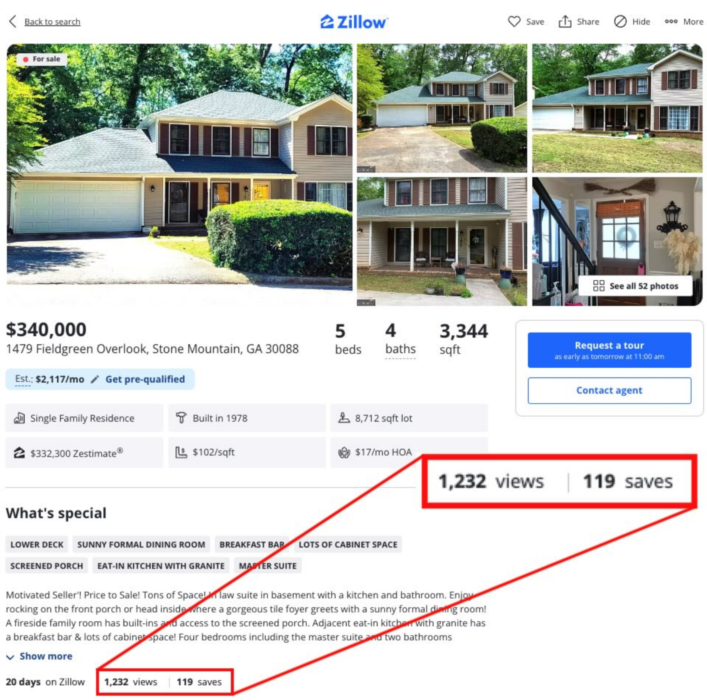

(b) Example of listing for sale (with number of customer views)

Figure A7: Examples of listings scraped in real time between May 2023 and September 2023

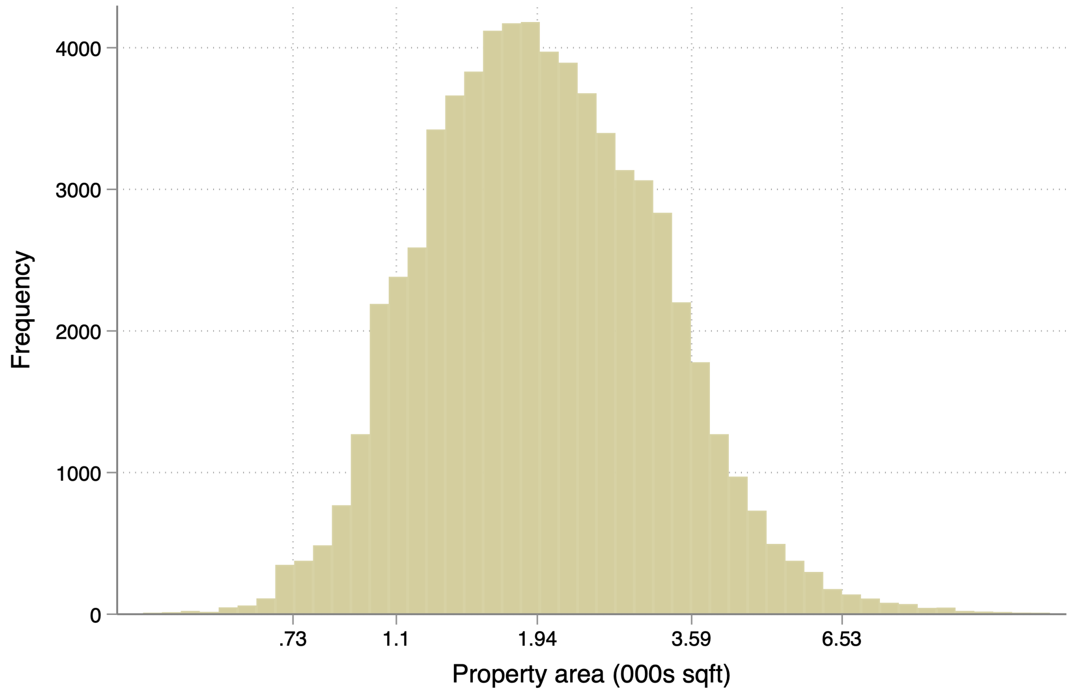

Figure A8: Distribution of Property Area of rental listings in square feet

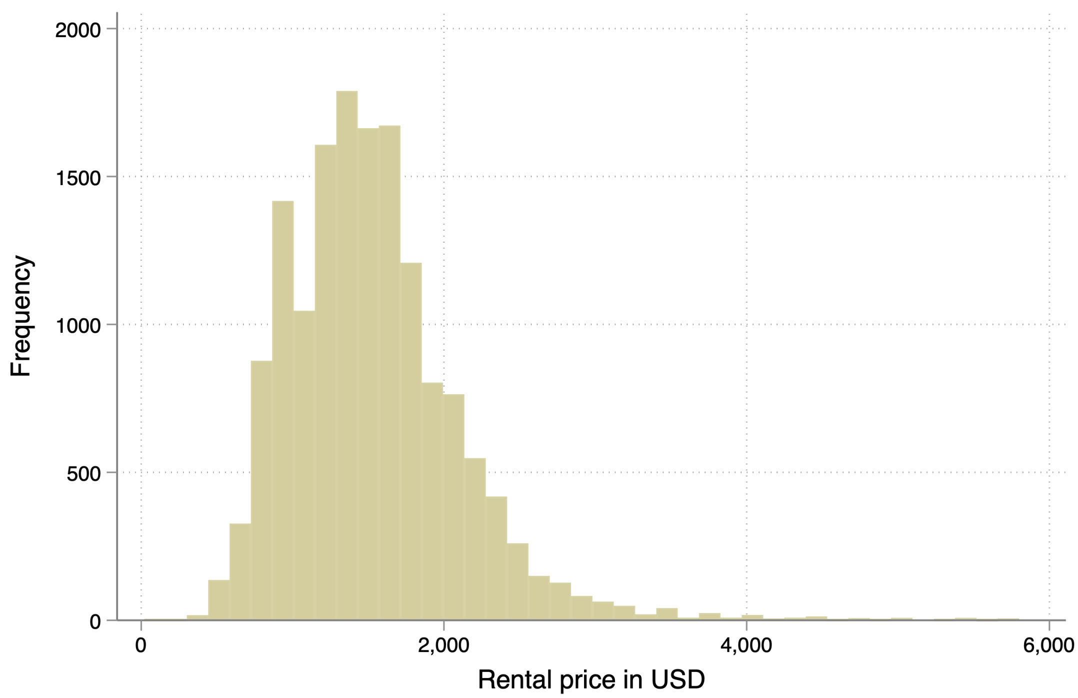

Figure A9: Distribution of rents for advertised rental listings (USD / month)

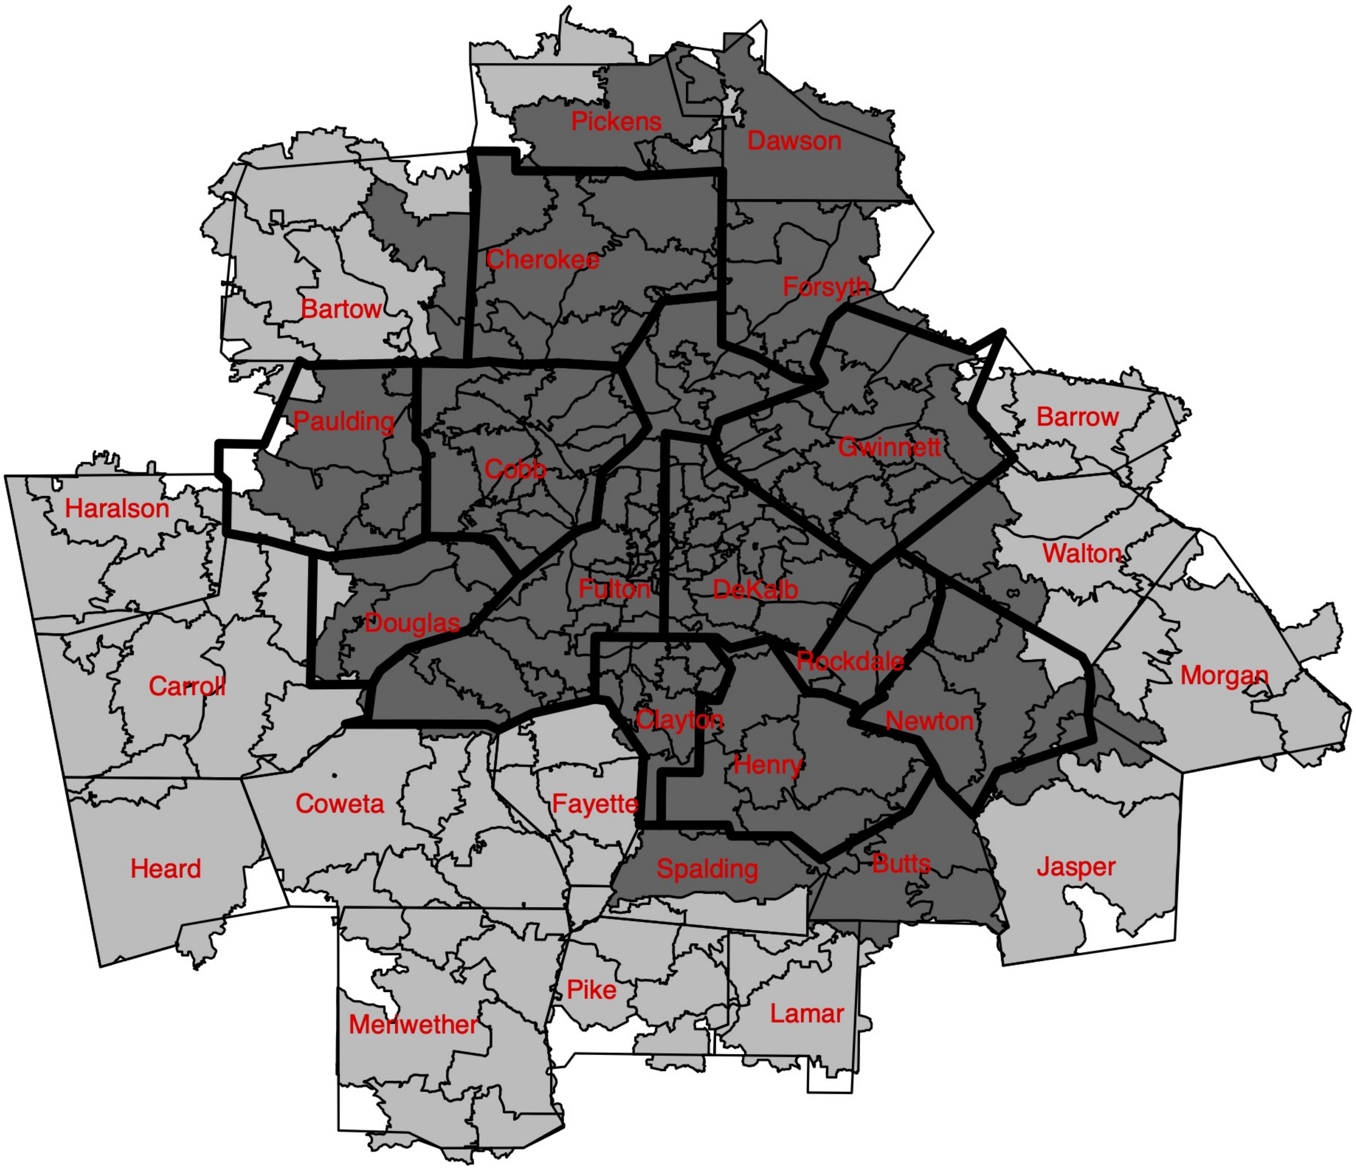

Figure A10: Coverage area of counties and ZIP codes for which we gather real-time daily Zillow data. We focus on ZIP codes located within the 11 most central counties of metropolitan Atlanta, delineated by thick black lines.

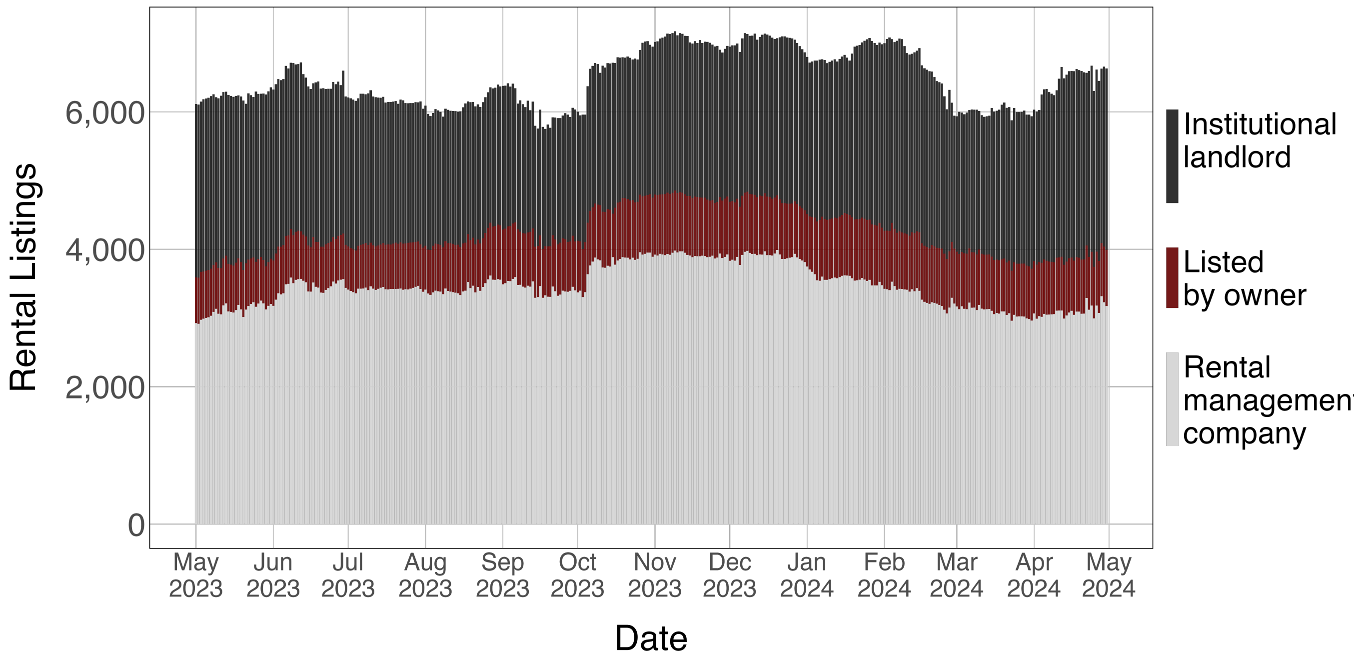

Figure A11: Composition of rental listings by landlord type

Notes: This stacked bar plot decomposes the total number of daily rental listings by the type of landlord advertising the property. Institutional landlords are shown in black. Non-institutional landlords are further decomposed into rental management companies (gray) and owners who list their property themselves (red).

C.2 Historical Zillow data

In addition to real-time data, we also collect "historical" Zillow records dating back to 2010. These records come from price histories of Zillow webpages. As shown in Figure 2, price histories contain information on rental prices, sales prices, begin and end dates of rental and sales advertisement spells, as well as amounts and dates for rents and price changes. Importantly, these price histories exist for all parcels in metropolitan Atlanta. We visit the webpage of every parcel in Atlanta and recover its price history. Figure A12 shows an example of a Zillow map with several on- and off-market parcels, each with their own webpage and price history.

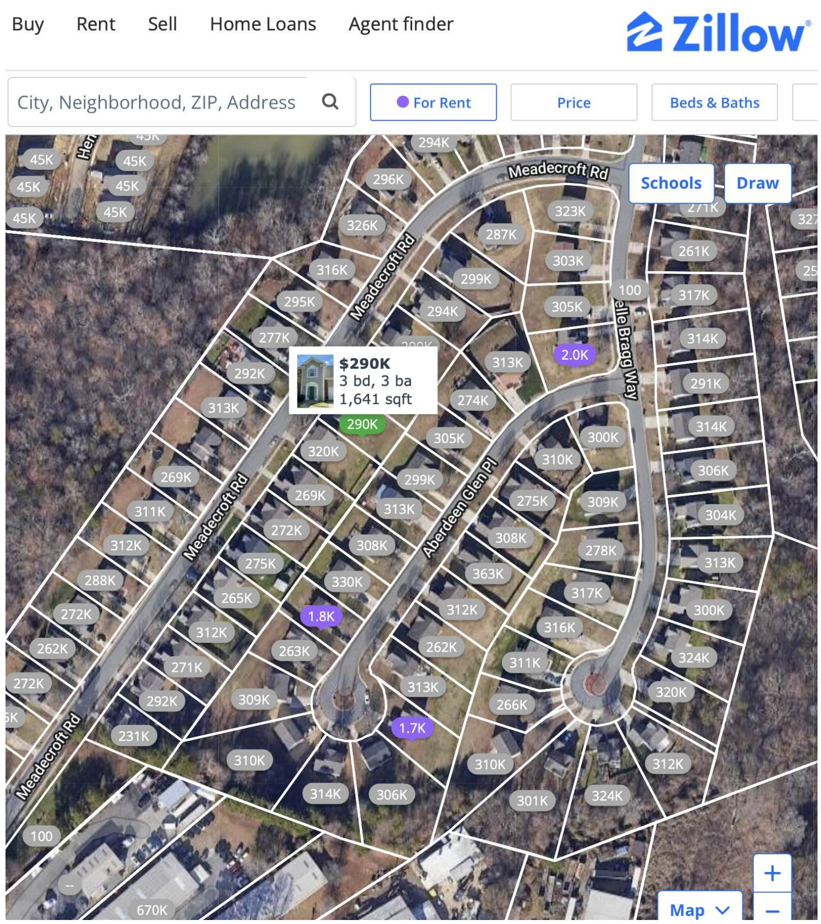

Figure A12: Each off-market parcel in Atlanta has a Zillow webpage with a price history

As shown in Figure 2, we can use gaps between rental advertisement spells to impute rental tenure durations. Figure A13 shows a plot of tenancy lengths. We observe that the majority of tenures have a length of 1-2 years, which aligns with typical lease lengths.

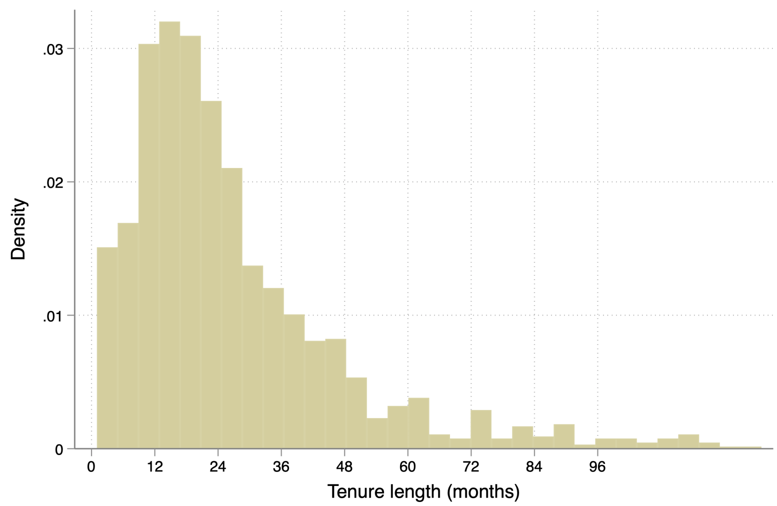

Figure A13: Rental tenure length for historical rental listings

C.3 Market Size Construction

Recall that we define a market $t = (Z, w)$ as the collection of individuals considering a move to ZIP Code Z during week w. We assume that the market size M_t, which corresponds to the potential number of movers to a market, is proportional to the aggregate volume of online activity in that market. In practice, we construct a time-varying M_t for each ZIP code by calculating the maximum weekly number of page views (for the homeownership sector) or contacts (for the rental sector) in a given month, and then multiplying this number by a fixed factor of proportionality λ.

To determine the factor of proportionality $\lambda = \frac{\text{potential number of movers}}{\text{total number of page views}}$, we calibrate the model using the homeownership sector, comparing the volume of web page views to the number of potential movers estimated from Census data. We proceed in three steps.

First, for the numerator, we calculate the *potential number of movers* to the 11 most central counties of the Atlanta MSA[32] using 2016-2020 U.S. Census county-to-county migration flows. To define the number of potential movers, we start with the set of all U.S.

[32] See Figure Figure A10 for a visualization of our geographical scope.

counties with non-zero migration flows to the central 11 counties of Atlanta MSA. For every county of origin o in the United States, we calculate the historical migration probability $\pi_{o\to\text{A}}$ as the share of movers from o who migrated to our sample area A. We then apply these probabilities to the total population of each origin county (Pop_o) to derive the total stock of potential movers, $N^{potential}$:

$$N^{potential} = \sum_{o\in\text{US}} Pop_o \times \pi_{o\to\text{A}}$$

We find that $N^{potential} = 3,210,399$, which represents 54% of the total population of central metropolitan Atlanta (which equals $5,879,040$). By scaling this using the U.S. homeownership rate (0.65), this implies that the number of potential *homeowner* movers to central metropolitan Atlanta represent $0.65 \cdot 0.54 = 35\%$ of the total central metro Atlanta population.

Second, for the denominator, we calculate the aggregate number of page views at the ZIP code level. We do this by taking the maximum number of total weekly page views (across all weeks in our sample) for each ZIP code. To ensure comparability with the Census data, we normalize this measure by the local population: using the 2020 Decennial Census, we find that these aggregate page views represent, on average, 8.3% of a ZIP code's total population.

Finally, to determine the factor of proportionality, we divide the percentage of potential movers by the percentage of observed views. This yields a ratio of $35\%/8.3\% \approx 4.21$. We round this factor to 4, implying that the total market size for the homeownership sector is four times the aggregate number of page views.

To determine the market size for the rental sector, we apply the same proportionality factor calibrated for the homeownership sector. We rely on this calibration because page views are a high-frequency, low-friction proxy for demand that are directly comparable to the population of potential movers: it is reasonable to assume that nearly every potential mover could generate a page view. In contrast, contacts are a noisier measure of demand, and represent a high-friction threshold that only a small fraction of searchers cross, making a direct comparison between the population of movers and the volume of contacts unreliable. We therefore assume that the structural relationship between observed online activity and latent demand is comparable across sectors.