The Impact of Institutional Investors on Homeownership and Neighborhood Access*

Joshua Coven†

October 18, 2024

Abstract

Since 2012, institutional investor entry into the single family rental market in the suburbs of US cities coincided with higher price and rent growth relative to the rest of the country. This paper estimates a structural model where institutional investor landlords benefit from economies of scale and market power, to determine the consequences of their impact on the housing market. While institutional investors' demand shock decreased homeownership and raised prices locally, supply responses dampened these effects resulting in a decrease in homeownership by 0.23 homes for each home purchased and an increase in prices of at most 27% of observed price increase in markets where they entered. Institutional investors decreased rents on net despite the presence of market power because they increased the supply of rentals. Overall, institutional investor entry resulted in a tradeoff: renters benefited through lower rents and more rentals in locations with better schools, but prospective homeowners had a harder time buying homes due to higher prices.

*I would like to thank Arpit Gupta, Robert Richmond, Johannes Stroebel, Cecilia Parlatore, Daniel Greenwald, Sam Chandan, Ingrid Gould Ellen, Sabrina Howell, Holger Mueller, Alexi Savov, Chris Conlon, Milena Almagro, and Marco Giacoletti for their feedback. I would also like to thank Gentry Hoit, John Isakson, Steve Katz, Mitch Rotta, and Alyssa Fantuz for their insight on the single family rental industry. I would also like to thank the seminar participants at NYU Stern Finance, the USC Marshall PhD Conference, and the North American Meeting of the Urban Economics Association for their feedback. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise. All mistakes are my own.

†Department of Finance, Stern School of Business, New York University. Email: joshua.coven@stern.nyu.edu

I Introduction

Historically, large landlords have focused on multifamily housing. However, since 2012, institutional investors entered the single family rental market and purchased up to 8.5% of the housing stock in some zipcodes in the suburbs of 10-20 US cities, including Atlanta, Phoenix, and Tampa. The new institutional investor landlords are different than existing "mom-and-pop" landlords: they have portfolios of up to 85,000 homes rather than 1-3 homes, and are spatially concentrated. Regions where institutional investors bought homes experienced higher price and rent growth than the rest of the country.

These facts call into question whether institutional investors' entry into the single family rental market increased prices, rents, and lowered homeownership. Policy makers, concerned that institutional investors' demand shock lowered homeownership and that their market power raised rents, proposed bans on institutional investors in single family homes[1] and a 5% rent increase cap for corporate landlords.[2] However, the construction sector can respond to a demand shock by building homes and small landlords can sell their homes, which might offset some of the impact on homeownership and prices. Also, the net effect of institutional investors on rents depends on whether market power leads them to decrease the supply of rentals or whether low operating costs leads them to increase the supply of rentals. The implications of institutional investor entry on the housing market depend on the underlying incentives of large landlords and the margins of adjustment through which others in the housing market respond.

This paper examines how institutional investor landlords are different than existing "mom and pop" landlords and the implications of their entry on the housing market. By estimating a structural model of the housing market with heterogeneous landlord types in operating costs and market power, I find that institutional investors raised prices locally by less than the observed price increases in markets where they entered, and decreased the quantity of homes available for homeownership by 0.23 homes for each home purchased. Despite the presence of market power, I show that large landlords increased the supply of rentals by 0.56 homes for each home purchased and decreased rents on net. Renters who moved into institutional investor rentals came from regions with worse schools and historic economic mobility. Policies that would ban institutional investors or impose rent increase caps would increase rents by reducing the rental supply, which is

[1] End Hedge Fund Control of American Homes Act, American Neighborhoods Protection Act

[2] White House calls on corporate landlords to cap rent increases at 5%

the opposite of the intended effect on the rental market. Overall, I find that institutional investors benefited renters by lowering rents and increasing the quantity of rentals in locations with good schools, however prospective homeowners had a harder time buying homes.

I begin by describing the activity of 7 large institutional investors that bought homes since 2012 to supply as rentals. Institutional investors targeted the suburbs of cities with low price to rent ratios that experienced many foreclosures during the Great Recession. While these 7 companies owned only 0.17% of the total housing supply of the US as of February 2021, they purchased 5.7% of the housing stock in Paulding County GA, and up to 19% of the housing stock in some census tracts in only 9 years. Institutional investors purchased homes in areas with few rentals, which led to 7 companies having a combined rental market share of up to 78% in some census tracts. Individual level data show that those who moved into institutional investor homes came from areas with lower median household incomes, middle school math test scores, and historic economic mobility. Institutional investor presence is associated with increases in house prices and rents relative to the rest of the US. The novelty of large landlords in the single family rental market, the concentrated buying, and the association of institutional investor entry with price and rent increases, account for the large amount of public interest in these investors.

Institutional investor landlords operate at lower average costs and scale more efficiently than existing "mom-and-pop" landlords. Two public single family rental REITs, Invitation Homes (INVH) and American Homes for Rent (AMH), pay lower property taxes, operating expenses, and insurance expenses than small landlords as a fraction of rent. This is due to scale, which allows them to efficiently appeal property tax values as documented in Austin (2022), and bargain with insurers and contractors for bulk discounts. The slopes of the cost curves are different as well. Small landlords minimize expenses in ways that do not scale: most do not hire managers and do not have mortgages. Because debt and managers are necessary to scale, small landlords must increase their cash costs to scale and therefore have decreasing returns to scale. Institutional investors on the other hand have vertically integrated management teams, high levels of debt, and a lower cost of debt. INVH's portfolio size increased 64% in a merger and its operating expenses per home did not increase, suggesting that they have constant returns to scale in the portion of the variation that we are able to observe. Lower average costs and constant returns to scale lead institutional investors to accumulate large portfolios where they operate. Large portfolios can lead to market power: the incentive to internalize the impact of quantity choice on rent.

I explore the implications of these cost curve differences by modeling institutional investor behavior in a structural model that includes households, small landlords, and construction. Key features of the model are that small landlords have decreasing returns to scale and large landlords have constant returns to scale. Large landlords can accumulate large enough portfolios to internalize the impact of their quantity choice on rents and therefore behave as Cournot oligopolists.

To answer quantitative questions such as whether institutional investors decrease homeownership, and how many renters would not live in these neighborhoods if institutional investors never entered the market, the model has rich substitution patterns for households. Households choose where to live and whether to own, rent single family, or rent multifamily, by solving a discrete choice problem. I estimate household demand with data from the Census ACS similar to Bayer, Ferreira and McMillan (2007), Calder-Wang (2022), and Diamond (2016). I calibrate geographic substitution patterns with bilateral migration data from Verisk (formerly Infutor). I estimate demand for different household groups based on income levels, which leads to different price and rent elasticities, allowing for institutional investors to have heterogeneous impacts and households to have heterogeneous responses. Elasticities for prices and rents are identified with an instrumental variables strategy.

To allow for supply responses to potentially dampen the effects of demand shocks on prices and homeownership, the model includes an aggregate builder of homes that increases the quantity of homes when prices increase. I use new unit elasticities from Baum-Snow and Han (2024), which allow for a heterogeneous construction response in each region of the model.

The rental supply in the model is determined endogenously by the sum of large and small landlord demands. This allows for the possibility that large landlords crowd out small landlords. I estimate small landlord demand by modeling the decision of an individual small landlord to operate a rental or to sell the property, and then aggregating. Large landlords choose a quantity of homes in each region to purchase and then supply as rentals to maximize profits, taking into account that their quantities will affect equilibrium rents and prices.

I estimate the impact of 3 identical institutional investors who enter the housing market in 2012 and choose locations and quantities of single family homes to buy and supply as rentals. I recover prices, rents, and quantities for Georgia, the epicenter of large landlord entry into single family rentals.

I find that institutional investor entry decreased the number of homes available for owner occupancy by 0.23 for each home they purchased. The demand shock triggered a construction response and crowded out small landlords: builders built 0.33 homes and small landlords sold 0.44 homes. A back of the envelope calculation that does not consider supply responses would overestimate the impact of investors on homeowership by 4x. Institutional investor entry increased the number of rentals by 0.56 for each home they purchased. The model shows that low income renters moved into the rentals supplied by institutional investors, which aligns with data showing that renters from lower income areas moved into these homes and therefore that institutional investors increased neighborhood access for low income renters.

In regions where institutional investors bought the most housing, the estimation implies that institutional investors caused 27% of the observed price increase. In regions where institutional investors bought a smaller share of the housing stock, almost all of the observed price increase would have occurred without institutional investors' demand shock. Institutional investors decreased rents by 0.7pp for every 1pp of the rental stock they own because on net, they increased the supply of single family rentals. This is in the opposite direction of the observed association of institutional investor presence and rent increases. The difference in size between the model implied price increase and the observed price increases, as well as the difference in direction between the model implied rent decrease and the observed rent increase, are consistent with the investors incentives to target regions with expected rent and price growth. I show that areas where institutional investors purchased homes experienced large population increases relative to the rest of the country after investor entry.

To validate that I am correctly measuring the size of the market power channel, I simulate a merger between 2 of 4 large landlords and compare the rent impacts to those in Gurun, Wu, Xiao and Xiao (2023) which uses mergers as quasi-experimental evidence to examine the effect of concentration on rents. I find that this merger increases rents by a median of 0.71% in areas of overlap, which is consistent in magnitude with the rent increase in Gurun et al. (2023).

I then simulate two proposed policies. One levies a tax of $10k per home per year for operators with more than 50 single family units. This more than doubles operating expenses, effectively banning institutional investors from operating single family rentals. I simulate this policy by removing institutional investors from the market in 2019 and find prices decrease, rents increase, and that small landlords gain a large share of the homes that institutional investors are forced to sell. I also simulate a 5% cap on rent increases for large landlords. I find that this lowers the quantity of rentals large landlords supply. The results show that not accounting for the fact that institutional investors increase the quantity of rentals can lead to policies having the opposite of the intended effects on rents.

This paper is the first to develop and estimate a structural model designed to estimate the total market level impacts of institutional investors on the housing market. A number of papers describe institutional investor activity in single family homes and the associations with housing market trends including Mills, Molloy and Zarutskie (2019), Gould Ellen and Goodman (2023), Demers and Eisfeldt (2022), and Giacoletti, Heimer, Li and Yu (2024). Ganduri, Xiao and Xiao (2022) and Lambie-Hanson, Li and Slonkosky (2022) study the role of these investors in increasing prices following the financial crisis. Garriga, Gete and Tsouderou (2023) estimate price impacts of small and medium investors on the purchase and rental markets. Gurun et al. (2023) and Austin (2022) study the effect of increasing concentration on rents, prices, and crime. Billings and Soliman (2023) studies the effect of institutional investors on neighboring housing prices. An (2023) studies associations of institutional investors with decreased black homeownership rates. Gorback, Qian and Zhu (2024) studies the effect of these investors on prices and rents. Francke, Hans, Korevaar and van Bekkum (2023) studies the effect of a ban on investor purchases in the Netherlands on prices of neighborhoods and demographics.

My research differs in five main ways. First, I can compare the size of the market power effect with that of the operational efficiency of the large landlords and find that the market power effect on rents is relatively small, therefore institutional investors caused a net decrease in rents. If policy makers consider only the market power channel, they would not only get the size of the impact on rents wrong, but also the direction. Therefore, policies aimed at lowering rents through banning investors would be counterproductive to lowering rents, as I show in a policy simulation. Second, I use the structural model to quantify the total impact of institutional investors on homeownership. The model accounts for the fact that people can move to become homeowners elsewhere, and then supply can respond in destination locations. The literature has not been able to quantify the total impact of institutional investors on homeowership. A back of the envelope calculation of the impact of investors on homeowership that does not consider supply responses would overestimate the impact by 4x. Third, I do counterfactual analysis to isolate economic channels and simulate policy proposals. I simulate a ban of investors and a 5% rent cap and show that both would reduce the rental supply and increase rents, similar to other research on the effects of rent control like that in Diamond, McQuade and Qian (2019). Fourth, underlying all of these results is a novel analysis of landlord type operating costs. Finally, I use individual location data to show that those who move into institutional investor rentals tend to come from lower income areas with worse schools. This is different from the demographic contribution of Austin (2022) which examines the change in demographics among home buyers.

There are a number of mechanisms through which institutional investors could affect households that I do not study, for example Raymond, Duckworth, Miller, Lucas and Pokharel (2018) finds that corporate landlords in this setting are more likely to evict than other landlords. Additionally, the FTC recently fined Invitation Homes for deceptive fees.[3] Also, it is possible that these investors renovate homes to improve their quality. These are important mechanisms that can affect households, however they are present in multifamily and exist across landlord types. This paper focuses on mechanisms specific to the entry of large landlords into single family rentals.

I also contribute to the literature on market power in housing. Calder-Wang and Kim (2024) studies how algorithmic pricing can lead to market power in multifamily rentals. Watson and Ziv (2024) studies market power in multifamily rentals in New York City. My paper focuses on market power in single family rentals, which are more liquid than multifamily rentals therefore there's a trade-off between rents and quantity rather than rents and vacancy. While market power decreases the quantity of rentals in a region and raises rents, here those rentals then go to either homeowners or other landlords, rather than becoming vacant. Gurun et al. (2023) studies a merger between single family landlords to measure market power. My paper differs in that I can compare the effect of market power on rents to the effect of increased cost efficiency on rents to get a net effect. I find that the cost efficiency channel dominates the market power channel.

I also contribute to the literature on models of asset demand. Koijen and Yogo (2019), Koijen, Richmond and Yogo (2023), Jiang, Richmond and Zhang (2023) use demand systems to examine stock market flows and international asset flows. McFadden (1978), Bayer, McMillan and Rueben (2003), Bayer et al. (2007), Diamond (2016), Calder-Wang (2022), and Almagro and Dominguez-Iino (2024) and use discrete choice models to study housing demand. I use bilateral migration data to estimate heterogeneous moving costs in utility terms for households from different origin locations to each destination region.

[3] https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-takes-action-against-invitation-homes-deceiving-renters-charging-junk-fees-withholding-security

I use the estimated moving costs in the discrete choice model to get realistic spatial substitution patterns in a static model.

II Data

II.A Property Level Data

I use property level data to identify institutional investor homes for the descriptive analysis. Property level data come from the Verisk property files. The dataset consists of 150 million rows at the tax lot level for the US for each cut of the data. I have cross sections for February 2021 and each year from November 2015–2019. The cross sections contain property characteristics, mortgage amounts, and anonymized owner mailing addresses. I include for analysis properties that have an indicator for single family residence, townhouse, apartment, condominium, duplex, triplex, or quadplex. This excludes vacant, commercial, and other types of non-residential properties. I exclude properties with no street information, mobile homes, and remaining properties with a duplicate address indicator. This reduces the dataset to 110 million residential units for analysis. I compare these to zipcode level housing unit counts in the Census ACS 5 year tables for 2020 in Appendix Table A1. The R^2 with the census number of housing units is 94% for all units, 96% for owner-occupied units, and 74% for rental units. The dataset under-counts rental units by around 8%, due to issues with how units for multifamily properties are counted in the data.

I identify institutional owners of properties through the mailing addresses for where property tax forms are sent. Following Ganduri et al. (2022), I look at the mailing addresses with the most properties in the US and then google them to identify companies. In this paper I focus on Invitation Homes, American Homes for Rent, Tricon Residential, Progress Residential, Main Street Renewal, FirstKey Homes, and Home Partners of America. For these 7 companies, I am able to identify 235,057 properties in February 2021. For the publicly traded ones as of 2021 (Invitation Homes, American Homes for Rent, and Tricon Residential), I validate the number of homes I find for them in the data with the numbers listed in their SEC filings. For February 2021 I can identify 86-91% of their properties as shown in Appendix Table A2. Tricon Residential is no longer a public company and was purchased by Blackstone in 2024.

For the descriptive analysis, I aggregate property level information to the owner and census Public Use Microdata Area (PUMA) level. PUMAs are census geographies with around 100,000–200,000 people. I chose PUMAs as the level of aggregation because they are the most granular geography for which the census publicly releases yearly microdata, as opposed to 5 year pooled samples. Any smaller geography would require the use of tables created from 5 year averages, which would be problematic for measuring changes in prices and rents over a 9 year period, especially following the large price swing resulting from the Great Recession. I also aggregate the data to the zipcode level and owner level, as well as the census tract and owner level. I observe the number of foreclosures from 2007-2011 at different geographic levels from Zillow ZTRAX. I also use shapefiles for the geography of the US to get the variable for distance to nearest city. I merge data at the tract level to middle school math test scores from 2013 from Opportunity Insights.

To get housing completions at the PUMA level, I use the Verisk property dataset's year built field and count the number of units at each year built. The resulting sample is accurate up to 2018 when compared with completions in FRED as shown in Appendix Figure A1. To get the amount of new construction at the PUMA level between 2012–2019, I use the annual completion rate for each zipcode between 2012–2018, and then extrapolate that to 2012–2019 to get the total amount of new construction in a PUMA.

The property records contain mortgage data including the original mortgage amount, the origination date, and the term length. I use this mortgage data to study the financing for existing single family rentals for small landlords. Around 10% of the single family rentals with mortgages have mortgage rate information. From this data, I assemble an empirical distribution of mortgage amounts for each PUMA. I then assume linear ammortization for a 30 year mortgage term to get a mortgage balance distribution for each PUMA.

II.B Rental Housing Finance Survey

I study small landlord costs using the Rental Housing Finance Survey (RHFS) from the US Census. The survey samples housing units from the American Housing Survey to get a representative sample of landlords in the US. It asks questions about specific costs and revenues, and distinguishes how many units are in each building as well whether the owner is an individual or a corporation. The survey has been conducted in 2015, 2018, and 2021 so far. Desmond and Wilmers (2019) uses this data to study multifamily costs.

I create a panel of small landlord single family rental costs with two vintages of the data, the 2018 and the 2021 RHS which cover 2017 and 2020 respectively. The survey contains around 500 entries each vintage for 1 unit properties with the ownership category of "Individual Landlords," which excludes corporations, REITs, and LLCs. The dataset does not provide information on how many properties the individual operator owns. From this sub-sample, I select properties with data populated in the columns for rent and market value. I exclude properties with flags for assisted living or rent control. For 2021, I am able to exclude townhouses, however for 2018 this column does not exist. These filters leave me with a sample of 603 small landlords total, covering 2 years.

To study small landlord costs, the ideal dataset would be a panel of landlord cost components or a cross section of average landlord cost components. The RHFS instead is a snapshot of landlord cost components. Some expenses can be large and irregular, like repairs. For categories with large and irregular expenses, I average the category as a fraction of rent across landlords rather than use the landlord level heterogeneity.

II.C Earnings Statements for Public REITs

I use earnings statement supplements to get market level cost components, occupancy, and the number of homes in each market for two public single family rental REITs, Invitation Homes (INVH) and American Homes for Rent (AMH). INVH and AMH became public companies in February 2017 and August 2013 respectively. To examine operating costs, I look at same store costs to exclude homes that were recently purchased or are in the process of getting sold. For the cost component analysis, my goal is to get a measure of market level same store adjusted funds from operations (AFFO) as a fraction of revenues. AFFO is a profitability measure common to REITs that focuses on cash flows, because REITs have large depreciation expenses even though their properties will most likely appreciate in value over time. To do this, I add same store operating expenses, same store recurring capital expenditures, and company level cash expenses including interest expense, management expense, and general and administrative expenses. I divide these by revenue to get AFFO as a fraction of revenue.

II.D Census ACS PUMS Data

For the estimation of housing demand, I create a dataset of housing holdings using data from the census Public Use Microdata Sample (PUMS), which is a survey of around 1%

of the US population each year. For each household, I observe the PUMA of residence, the household income, and whether the household lives in an owner occupied unit, single family rental, or multifamily rental. I aggregate this to the PUMA, year, household income group to get a panel of holdings data for income groups across the US from 2012–2019. The household income groups are 0–25k, 25–50k, 50–75k, 75–100k, and 100k+.

The Census tables at the PUMA level include characteristics that I use for the demand estimation such as the median number of rooms in each PUMA, the fraction of a PUMA that is single family housing, the median age of housing in a PUMA, the fraction of the population that has a commute that takes less than 45 minutes, the fraction of the high school aged population that is enrolled in high school, and the fraction of the high school aged population that is enrolled in private school.

II.E Mover level data

To analyze the differences between who moves into institutional investor homes vs other homes in the same census tract, I use the Verisk location history data. Verisk provides a location history for each individual, including up to the last 10 locations. Each of the locations has an address ID which can be linked to the property dataset. I describe the steps for cleaning this data in the Appendix. I validate the moves to and from a given zipcode with USPS change of address data in Appendix Table A3. I also validate moves from county to county with ACS data in Appendix Table A4. I merge this dataset with the property data on a mover's origin location and destination location to see if the mover moved into an institutional investor home. I can also observe if someone lives in an owner occupied unit, single family rental, or multifamily rental.

II.F Other Data Sources

I include weather data due to the importance of weather as a location amenity as shown in Saiz (2010) and Chodorow-Reich, Guren and McQuade (2023). Weather data at the county level come from the US Department of Agriculture. I include in the demand estimation as characteristics January sunlight, January temperature, July temperature, and July humidity. I create variables for distance to the nearest city and a distance to the nearest top 30 city to use in the household demand estimation. Zipcode-MSA crosswalks come from the US Department of Housing and Urban Development (HUD). Zipcode geographies come from the Census Tigerweb files. To make county level maps, I use a zipcode-county crosswalk from HUD. I aggregate these zipcode characteristics to the PUMA level for the demand estimation.

III Description of Institutional Investor Activity

III.A Location choice and conversion of homes to rentals

I start by describing where institutional investors purchased homes. Appendix Figure B1 shows the fraction of the total housing stock that institutional investors own at the county level for the United States. Institutional investors' largest markets are sunbelt cities that had many foreclosures from 2007–2011, including Atlanta, Phoenix, and Tampa. In Paulding County, GA, these 7 companies that I study own 5.7% of all the housing stock as of 2021. I examine the characteristics of the neighborhoods where institutional investors purchased homes with a descriptive regression where $y_{p,s}$ is a dummy variable if PUMA p in state s has 10 or more institutional investor properties, $\beta_{z,c}$ is a vector of characteristics describing the housing stock, demographics, and location characteristics, and α_c are county fixed effects:

$$y_{p,s} = \beta X_{p,s} + \alpha_s + \varepsilon_{p,s}. \tag{1}$$

Results in Table B2 show that institutional investors are located in PUMAs with low price to rent ratios, which is ideal for the conversion of owner-occupied homes to rentals. The institutional investors I study in this paper have held onto these homes to rent, as opposed to iBuyers like Zillow and Opendoor who purchased homes to flip, as studied in Buchak, Matvos, Piskorski and Seru (2022).

Institutional investors are located in regions that had large population and job growth pre-trends, and a high number of foreclosures per housing unit in the Great Recession. Invitation Homes and American Homes for Rent in IPO filings mentioned that population growth and job growth would lead to rent and price growth. They also targeted regions with low prices relative to prices before the Great Recession. Foreclosures were a critical acquisition channel early on, as documented in Mills et al. (2019) and in SEC filings from INVH and AMH. 37% of the homes Invitation Homes acquired from September 2015 to September 2016 were through distressed sales, showing that foreclosures and short sales played an important role in property acquisition even a number of years after the Great Recession. Additionally, many states have a right of redemption for foreclosures where a homeowner who was foreclosed on has 12 months after the foreclosure sale to pay the purchaser to get the home back. This is a risk for landlords who want to purchase properties to operate as rentals, and this risk can be better absorbed by those with larger diversified portfolios rather than operators looking to buy 1 home to then rent. Spatially, they purchased homes in rings around cities. They bought homes in PUMAs with higher median household incomes, and lower shares of the population white or college educated.

Where they chose to enter the housing market, they amassed concentrated portfolios of housing. In a census tract in Rutherford County, TN, between Nashville and Murfreesboro, these companies own 19% of the entire housing stock: 534 of the 2803 properties in the tract. Figure 1 Panel A shows the fraction of housing stock at the census tract level owned by the 7 institutional investors I study for the Atlanta metro area. The figure shows that they own up to 10% of all the housing in some tracts surrounding the city. I plot the share of housing stock they own against the share of rental stock they own at the tract level for the whole US in Figure 1 Panel B. In some tracts, they make up 78% of the total rental supply. This implies that institutional investors did some combination of expanding the rental supply or replacing existing landlords.

I compare the physical characteristics of the rentals institutional investors supply compared to other rental units in the same zipcode and owner-occupied housing in the same zipcode. I calculate within zipcode differences between the three groups and average the differences weighted by B2Rs zipcode level exposure, and then add back the institutional investor mean. Institutional investors supply rental stock that is newer and larger than the other rental stock in the zipcodes they are present in, as shown in Table B3. On average, the institutional investor rental supply is around 7.48 years newer, has 0.23 more bedrooms and 0.16 more baths, and is 20% more likely to be single family than other rentals in the same zipcodes. The institutional investor rental supply is on average 150 sqft smaller than owner-occupied housing, 7.6% smaller. This is consistent with the idea that they bought starter homes.

III.B Trends in geographies where B2R is most concentrated

I analyze the trends associated with institutional investor entry that highlight why there has been so much attention focused on institutional investors. I first show the association between institutional investors' holdings in 2019 with changes in prices and rents at the PUMA level, relative to the rest of the country. Figure B2 Panel A shows this as a binscatter. Institutional investors are assocaited with price increases of up to 40% more than the rest of the US from 2012-2019. They are associated with rent increases of up to 10% more than the rest of the US. The figure also shows how small of a share institutional investors have of a census PUMA. In their most concentrated PUMA, they own 4.5% of the housing. Panel B shows the association with owner occupancy rates. At the PUMA level, there is no trend in the binscatter of owner occupancy rates with institutional investor presence. The owner occupancy rate in the US decreased from 2012–2016, before rebounding to 2012 levels by 2019. Panel C shows the associations of institutional investor entry with household growth from 2012-2019 and also new single family construction as a share of 2012 population. Institutional investors chose areas that realized large increases in population relative to the rest of the country, and large increases in new building supply. The growth in households exceeds the supply growth in single family homes, which could lead to price increases.

I explore the possibility that the addition of significant rental supply could have had spillovers on broader rental markets by relieving demand for rentals in other regions. The majority of renters new to institutional investor rentals came from areas with at least some institutional investor exposure, or areas adjacent to institutional investor zipcodes. I examine where people who moved into institutional investor owned homes for the first time came from in Figure B4. The entry of institutional investors increased the rental supply for these renters, and therefore rental demand in their origin zipcodes could have decreased. I examine the association between which origin locations had the most people leave to go to institutional investor homes and the rents in the origin locations. I estimate the following descriptive regression of the change in rent for a zipcode Δr_z in county c from 2012–2021 on indicator variables for the number of people I observe who left that zipcode to move into an institutional investor home for the first time $N_{z,c}$ and control for county level fixed effects:

$$\Delta r_{z,c} = \beta N_{z,c} + \alpha_c + \varepsilon_{z,c}. \tag{2}$$

I plot the results from the descriptive regression in Figure B2 Panel D. I find that areas where I am able to observe 100 or more people leaving to go to institutional investor homes are associated with rent decreases of 8.4% . When I restrict the sample to origin locations that do not have any institutional investor homes, the association increases, and 10 to 20 people leaving to go to an institutional investor home is associated with a 7.9% decrease in rents from 2012–2021. The association is negative and increases in magnitude with the number of people. It holds for both zipcodes that contain institutional investor homes and those that do not contain any institutional investor homes. The results show that a model of institutional investors' impact must include geographic spillovers to other markets.

III.C Who moves into institutional investor rentals?

I examine if the rentals that institutional investors supplied increased neighborhood access for the financially constrained. First, I plot the model output of the amount of owner-occupied homes and rentals that each income groups gains or loses when institutional investors enter in Figure 6. Households with incomes between 25–50k lose the most homes for owner-occupancy because homeowners with these incomes are most exposed to the institutional investor shock. By supplying rentals, institutional investors increased the rental supply for the lowest income renters, who are most elastic to changes in rents.

I supplement these results with individual level migration data to examine where those who move into institutional investor homes come from and compare characteristics of the origin locations to the destination locations. I merge tract level data of median incomes, middle school math test scores, and historic economic mobility from Opportunity Insights to both origins and destinations. In Table 1 I show the mean difference in destination tract characteristics relative to origin tract characteristics. Those who moved into institutional investor homes came from areas with 12.2% lower median household incomes, 5.8% worse middle school math test scores, a 6% increase in likelihood to go to jail, and a 3.4% decrease in likelihood to make it to the top income quintile. To the extent that institutional investors increased the rental supply, they increased it in areas with better schools and economic mobility than the origin locations of those who rent from them.

I compare those who move into institutional investor homes with those who move into other homes in the same tract to better isolate whether institutional investors are not just increasing the quantity of rentals in good locations, but also changing the flow of people to these good locations. Among all movers into tracts with institutional investor homes, I create a dummy variable for those who moved into a institutional investor home for the first time. Results are in Table 2. Those who moved into institutional investor homes came from worse areas than those who moved into other homes in the same census tract. This suggests that institutional investors increased access to the tract for those from worse areas by more than if the destination had one additional home. This is due to the fact that institutional investors increase the number of rentals in the new location, and renters come from areas with lower incomes and worse historic economic mobility than homeowners.

III.D Differences between small and large landlords

To figure out if institutional investors affect the housing market, we need to understand how they are different from existing single family landlords and how these differences might matter to the housing market. I start by providing some background information about operator size distributions in the single family rental market in February of 2021 by constructing operator level single family rental portfolios from Verisk property data. Figure 2 Panel A shows the fraction of all single family rental homes owned by operators of different sizes. It also shows how many of the operators within a size bucket operate in only one county as opposed to multiple counties. Institutional investors entered an extremely fragmented market dominated by tiny "mom and pop" landlords. 71% of single family rentals are owned by operators with 1-3 properties. 75% of single family rentals are owned by operators who do not operate in multiple counties. Institutional investors, even in 2021, make up only 1.8% of single family rentals. However, as I showed earlier in Figure 1 Panels A and B, their spatial concentration makes them relevant operators in the regions where they are located. In this fragmented market, institutional investors are much larger, where some own over 80,000 homes. This allows them to have spatial concentration, which is not something that is possible for operators with 1-3 homes.

Next, I examine how economies of scale lower institutional investor landlord operating costs relative to small landlords. Figure 2 Panel B compares cost components for average small landlords from the Rental Housing Finance Survey, to cost components from earnings statements for Invitation Homes and American Homes for Rent which I call the large average. Institutional investors have lower property taxes as a share of rent. This may be due to the fact that institutional investors are more likely to appeal their property tax assessment valuations, as documented in Austin (2022). Institutional investors also have lower operating expenditures. In American Homes for Rent's IPO filing, they mention that they get quantity discounts for materials regularly used including paint, flooring, and blinds. Invitation Homes in its IPO mentions that it is able to get discounts for HVAC systems and discounts for contractors by working directly with vendors, for the work that it does not use in house staff for. Large landlords have much lower insurance costs, 1-2% of rent rather than 5-6%. Large operators bargain with insurance companies for a bulk discount here as well. While in Figure 2 Panel B it appears that large landlords pay more in management costs, that is because 83% of small landlords do not hire any professional management. When they do, they pay an average fee of 10% of rent.

Institutional investors, on the other hand, have vertically integrated management companies and pay 4-7% of towards their internal management operations. While I am not able to break this out in the data, institutional investors also have vertically integrated leasing and acquisition teams. American Homes for Rent, before internalizing its acquisition team, paid a 5% fee on top of all closing costs for each acquisition. Overall, institutional investors benefit from economies of scale that lower operating costs.

Large and small landlords also look very different on the financing side. 63% of small landlords do not have mortgages or similar debt in the RHFS. I verify this number in Verisk and find that 57% of small landlords here do not have a mortgage. Large landlords have higher levels of debt. In 2021, Invitation Homes had a debt to value ratio of 51% and American Homes for Rent had one of 33%. They use many types of debt including asset backed securitizations, bonds, term loans and credit lines. I show differences in interest rates for small landlords, owner occupied homes, and the weighted average interest rates on debt for INVH and AMH in Figure B5. There is a spread between new mortgages for owner occupants and small landlords of around 0.2%. The spread is larger for existing mortgages, around 0.5%. AMH and INVH have significantly lower costs of debt than small landlords and owner occupants. INVH in particular has a lower cost of debt than AMH, possibly due to shorter term lengths.

Small landlords minimize costs in a way that does not scale, suggesting that they have decreasing returns to scale. 83% of small landlords do not use any professional manager and 63% do not have a mortgage. At a certain scale, debt and managers are necessary because people have limited equity and as the number of properties increases, it gets harder for a person to manage them alone. Figure 3 Panel A shows these dynamics in the cross section with a histogram of small landlords, sorted by cash costs as a fraction of rent. The graph also shows the fraction of each histogram bucket that has a mortgage. High cost small landlords are much more likely to have a mortgage than low cost small landlords. Therefore, a move from no mortgage to a mortgage would likely move a small landlord to a higher cost portion of the histogram, suggesting decreasing returns to scale. Additionally, small landlords who purchased homes recently are more likely to have a mortgage than those who purchased decades ago. I show this in Figure B6 Panel A, which buckets landlords in the RHFS 2021 vintage by when they purchased a home, and shows the fraction of that bucket that has a mortgage. Panel B shows the fraction of an original mortgage balance remaining by purchase year for those who still have mortgages. Small landlords appear to get mortgages of around 80% LTV, similar to households, and then pay them off over time. These pattern suggests that small landlords who want to change the number of units they operate are more likely to have a mortgage than those who have operated a property for 20 years, and are more likely to have higher mortgage balances, and therefore have higher interest expenses suggesting decreasing returns to scale within operator.

Large landlords appear to have better returns to scale. While I can't observe company costs before they went public, I can examine how changes in scale affect cash costs for the public REITs. First I examine market level average operating expenses per home for Invitation Homes in Figure B7. The figure shows market average operating expenses per home against the market average number of homes. There is no pattern between the number of homes and the operating expense per home. I examine variation in portfolio size due to a merger to get a different look at how operating expenditures per home might vary with portfolio size. In 2017, Invitation Homes bought Starwood Waypoint Homes and went from 50,000 homes to 82,000 homes. I plot the change in operating expenses per home against the change in number of homes from this merger for each market in Figure 3 Panel B. The black dotted line represents no change in operating expenses per home. While this is not exogenous variation, it is strongly suggestive that these companies can almost double the number of homes they have in a market without raising operating costs. I am able to examine variation in costs only after these companies have already accumulated large portfolios. Most likely, these companies early on experience increasing returns to scale as they are paying fixed costs to build capabilities for management, acquisition, and renovation. Then, they have constant returns to scale, and then decreasing returns to scale. I am only able to observe variation suggestive of the middle part of this curve, the constant returns to scale part. This region where the curve appears constant is the relevant region for when I estimate the impact of these investors with a model because I need the slope to be correct where it intersects demand or marginal revenue curves, and the quantities of large landlords later in the model line up closely in magnitude to those observed around this merger.

III.E STYLIZED MODEL OF SUPPLY AND DEMAND FOR SINGLE FAMILY RENTALS

Now that I have established that small landlords have decreasing returns to scale and large landlords have constant returns to scale over a large range of quantities we are interested in, I use these facts in a stylized model of the single family rental market to illustrate economic channels through which institutional investors can affect the housing market. Figure 4 Panel A shows the supply and demand curves for single family rentals in one region. Households have downward sloping demand, small landlords have upward sloping supply, and large landlords have constant supply, as motivated by the earlier section. Here, if large landlords chooses competitive quantities, the equilibrium moves to the right from A to B, and rent decreases and the number of rentals in the region increases.

To illustrate the market power channel in this stylized example, I assume there is only one large landlord. This large landlord can supply a significant amount of demand that small landlords cannot, and therefore is a monopolist over the residual demand. A profit maximizing large landlord therefore can internalize the impact its quantity choice has on rent, leading to market power. I this show in Figure 4 Panel B. The large landlord chooses a quantity where the residual marginal revenue intersects its cost curve. This shifts the equilibrium left to C, raising rents and lowering quantities. This movement from B to C is the effect of market power, which is derived from the cost advantage of the institutional investors. This cost advantage is derived from scale and spatial concentration.

Supply and demand for single family rentals are modeled as a Cournot (Cournot, 1838) oligopoly where small landlords are price takers and large landlords choose quantities to maximize prices because of two economic features of the single family rental market. First, single family rentals are highly liquid relative to other real estate assets. It's easier to sell a small piece of a single family portfolio than a multifamily portfolio because the discrete units that would be sold are smaller. Additionally, institutional investors don't have to sell to other institutional investors, and could sell to households. This is an additional source of exit liquidity. So rather than raising rents and increasing vacancies, an institutional landlord would sell the vacant units or not have bought them in the first place. There are no barriers to entry: if an institutional investor sells a home, another landlord could operate it if its profitable. If rent rises sufficiently, a small landlord could buy a home from a large landlord or a homeowner and then offer it as a rental. Both the liquidity and the lack of a barrier to entry make single family rentals ideal to model as Cournot. This is in contrast to market power in multifamily housing, where there is a tradeoff between rents and vacancy due to the inability of large landlords to sell units within a building and to sell to households.

IV Model

I estimate the impact of institutional investors entry into the housing market with a structural model, which provides four key benefits. First, it allows me to answer quantitative and global questions, such as, how many households would own homes if the institutional investors did not enter the market? Empirical strategies can isolate a local effect of investors on homeownership, however people can move and become homeowners elsewhere. Second, I can compare economic channels to see what drives net effects. For example, how much of the institutional investor impact on rents is due to market power and how much is due to their ability to operate at low costs? This is important because natural experiments in mergers isolate only the part of institutional investor impact that is due to increasing ownership concentration, and we want to know the net effect of institutional investors on rents. Third, the model allows me to quantify economic channels by shutting off housing market responses like the construction response to price increases. Fourth, the model allows me to simulate policies including what would happen if institutional investors were banned from the single family rental market.

There are a number of trade-offs involved in using a model to estimate the impact of institutional investors. While it provides the benefits listed above, it cannot include forces that are not modeled. This can be a strength, for example, it excludes the population increases that occurred in institutional investor PUMAs which can help isolate the impact of institutional investors on housing markets. However, if does not include forces that may be relevant to investor impact if they are not modeled. For example, if institutional investors upgrade properties, this could increase the demand for properties to increase rents. And I don't model this force. However, if rents rise due to property renovations, it's not clear that that is negative for renters, as opposed to market power where quantities decrease and prices increase.

The model has four agents. Households decide where to live and whether to own or rent by solving a discrete choice problem. The discrete choice demand system for housing is similar to models from Bayer et al. (2007), Calder-Wang (2022), Bayer et al. (2003), McFadden (1978), Diamond (2016), Almagro and Dominguez-Iino (2024), Schubert (2021) for housing and Koijen and Yogo (2019), Koijen and Yogo (2019b), Koijen et al. (2023), and Jiang et al. (2023) for financial markets. Small landlords decide whether to operate or not by comparing expected operating cash flows to the amount of cash they would get from selling their property. In aggregate, they have decreasing returns to scale. Large landlords choose a number of rentals in each region to operate to maximize profits. They are large enough that they internalize the impact of the quantity they choose on the equilibrium rent and are therefore modeled as Cournot oligopolists. An aggregate builder increases the number of homes in a region when prices increase.

The model has three features required to deliver quantitative answers to questions such as how many homeowners would have had homes if institutional investors never bought homes, and how many renters would not live in these neighborhoods if institutional investors never supplied single family rentals. First, the model contains rich substitution patterns for households who have demand for owning, renting from single family, and renting from multifamily in any census PUMA in Georgia. If prices increase, households can stay or move, and switch from owning to renting. I can observe if the household becomes a homeowner elsewhere, leading to the quantifying of a global effect of the investors on homeownership. Second, there is an aggregate builder of homes which increases the supply of homes when prices increase. This may lower the effect of a demand shock on prices and homeownership. And third, rental supply is determined endogenously by the demand of institutional investors and small landlords. This allows the model to take into account that other landlords may have supplied rentals if institutional investors were absent. These margins of adjustment to demand and supply shocks are necessary to recover the impact of institutional investor entry on homeownership and the rental market.

IV.A THE DEMAND FOR HOUSING

The estimation of household demand allows me to measure how many households leave a PUMA if prices increase, where they go to, and whether they become homeowners or renters in the destination location. Households choose where to live and whether to own, rent single family, or rent multifamily in that location. Households are heterogeneous in income, denoted by income group level h, and origin location i. For example, a household with an income of 25–50k owning homes in a suburb of Atlanta decides where to live and whether to own, rent single family, or rent multifamily in that destination. They can also decide to not move. Income heterogeneity allows for groups with different incomes to have different sensitivities to changes in price and rent. Origin location heterogeneity allows the model to incorporate realistic spatial substitution patterns.

A household from origin location i in income group h in asset class l, which could be owning, renting single family, or renting multifamily, solves a discrete choice problem to figure out where to move, j, which could also be not-moving if $j = i$, and whether to own, rent single family, or rent multifamily in that destination, k. The choice is based on characteristics of the destination location and asset class, $X_{j,k}$, and a moving utility cost from location i asset class l to location j asset class k: $\tau_{(i,l)\to(j,k)}$. The asset class in location

j has a log price of $p_{j,k}$, which is the log of the purchase price of owner occupied home if $k = owneroccupied$, or the log of the rent if $k \in \{rent_{sf}, rent_{mf}\}$. Households have a latent demand denoted by $\varepsilon_{h,(i,l)\to(j,k)}$. A household's indirect utility for moving from i, l to j, k is based on these characteristics and their corresponding elasticities β as follows:

$$u_{h,(i,l)\to(j,k)} = \beta_{h,k,0} p_{j,k} + \beta_{h,k,d} X_{j,k} + \tau_{(i,l)\to(j,k)} + \varepsilon_{h,(i,l)\to(j,k)}. \tag{3}$$

The fraction of households of a given income type in a region and in an asset class, $w_{h,j,k}$, is determined by sum of movers to the region and asset class. The share is modeled as a mixed logit:

$$w_{h,j,k} = \sum_{i,l} \frac{\exp(u_{h,(i,l)\to(j,k)})}{1 + \sum_s \exp(u_{h,(i,l)\to(j,s)})} w_{h,i,l}. \tag{4}$$

PUMA and asset class characteristics X include weather, the fraction of the population that has a commute under 45 minutes, the fraction of high school age population that is enrolled in high school, the median number of rooms of housing, the median year built of housing, the year, and other characteristics.

To get accurate geographic substitution, each income group has within group heterogeneity by origin location. Each origin location has a different utility "cost" $\tau_{(i,l)\to(j,k)}$ to each destination location, based on origin destination pair characteristics. These include distance, the number of social connections from Meta's social connectedness index, as well as ownership and renting transitions such as moving from being an owner to a single family rental.

This results in the propagation of the institutional investor impact along a network based on empirical moving patterns that are due to distance and social connections, similar to how in Piazzesi, Schneider and Stroebel (2020) housing market shocks are propagated along a search network. Diamond (2016) models household demand based on birth state heterogeneity. Here, the heterogeneity in moving costs depends on distance and social connections. A number of dynamic housing models like that in Schubert (2021) get accurate spatial substitution for free by estimating migration shares as a function of characteristics. I don't have data that has migration shares broken out by income groups, so here, I combine cross sectional data that has heterogeneity at the income level but not in the migration dimension, with migration data that does not have income level heterogeneity, in order to get a static model with realistic spatial substitution patterns and heterogeneity at the income level.

The outside asset in the model is any PUMA or asset class that has a price below $80k or a monthly rent below $500. The outside asset also includes PUMAs that do not have data for the characteristics of interest. Since I am modeling only Georgia, the outside asset also includes all other PUMAs in the US.

I assume that $\beta_{i,0} < 0$ for all households: that demand is downward sloping. If two PUMAs are identical except for price, an agent would prefer the cheaper one. An estimation that results in an agent favoring higher prices can be interpreted as prices correlating with something desirable to homeowners or investors that is outside of the model. Because of this, $\beta_{i,0}$ is constrained to be less than 1 in the estimation. Portfolio weights for each household group clear:

$$\sum_{(i,l)} w_{h,i,l} = 1. \tag{5}$$

Each income group has N_h households. The quantity of homes demanded by households in each location and asset class is therefore:

$$Q_{d,households,i,l} = \sum_h w_{h,i,l} N_h. \tag{6}$$

IV.B Small landlord demand

Each region j has a stock of small landlords, $N_{small,j}$. A small landlord i in region j operates if cash flows from operating are preferred to selling:

$$E\left[\sum_t \frac{(R_{j,t} - C_{i,j,t})}{(1+r_e)^t}\right] \geq P_{j,t} \times (1 - BrokerFee) - M_{i,j,t}. \tag{7}$$

The left hand side is the expected operating cash flows in each period, $R_{j,t} - C_{i,j,t}$, discounted by the cost of equity for those cash flows r_e. The right hand side is the cash a landlord would get from selling a home. The landlord gets the purchase price $P_{j,t}$ minus a broker fee, and then has to pay back a mortgage $M_{i,j,t}$. Cash flows from operating each period are expected to grow at a rate of g_j, making this a growing perpetuity. g_j is a function of recent population growth in the region. The cost for a landlord i in a region j is:

$$C_{i,j} = P_j \times PropTax_j + InterestExpense_{i,j} + OtherCosts_i + R_j \times ManagerFee_i. \tag{8}$$

For the interest expense, I assume that the typical small landlord has a 30 year mortgage that they pay off over 30 years. A small landlord is aware that this is their likely mortgage payoff behavior and takes this into account when deciding whether to operate or not, so I model their interest expense as the perpetuity equivalent of the present value of their interest payments. This allows the model to capture the fact that some small landlords might barely break even when first purchasing a property, in hopes that after they pay off their mortgage, cash flows will increase. This leads to the decision to operate to be formulated as:

$$R_{j,t} \geq (r_e - g_j) \times (P_{j,t} \times (1 - BrokerFee) - M_{i,j,t}) + C_{i,j,t}. \tag{9}$$

In a region j, I aggregate each small landlord's demand to get the cost curve. At each rent, for a given price, a certain number of these individual landlords will decide to operate or not operate. The number is increasing in rent and decreasing in price. The aggregate small landlord has decreasing returns to scale due to the heterogeneity in costs of the underlying individual landlords. The heterogeneity in costs that leads to this decreasing returns to scale is driven by different interest expenses and operating expenses.

IV.C INSTITUTIONAL INVESTOR DEMAND

A large landlord i chooses a quantity of homes to buy in region j, $Q_{i,j}$, to maximize profits subject to a required cash margin:

$$\max_{Q_i} \{Q_i \times (CashFromOperating - CashToBuy), 0\} \tag{10}$$

$$s.t. \; \frac{Rent - Cost}{Cost} \geq 30\%$$

Both Invitation Homes and American Homes for rent have cash margins of around 30% in each market after subtracting company level cash costs like interest, management, and general administrative costs from their local operating margins.

Large landlords get cash from operating rentals but have to purchase properties as well. They are modeled as Cournot oligpolists who internalize the impact of Q_i on rents and prices given the quantities of others in the market, Q_{other}, which includes small landlords and other large landlords. They pay a renovation cost when purchasing each property and also use a portion of their total debt, D_t, to buy each home:

$$CashFromOperating = E\left[\sum_t \left(\frac{R_{j,t}(Q_{other}+Q_i) - C_{i,j,t}}{(1+r_e)^t}\right)\right] \tag{11}$$

$$CashToBuy = P_{j,t}(Q_{other}+Q_i) \times (1 + renovationCost) - D_t \tag{12}$$

$$C_{i,j} = P_j(Q_{other}+Q_i) \times PropTax_j + IntExp_i + OtherCosts_{i,j} + ManagementCosts_i. \tag{13}$$

Region specific costs depend on the local property tax rate. Unlike small landlords, their interest expense and management costs are not region specific. They have region specific operating costs that are a function of local contractor wages.

Like small landlords, I model the expected operating cash flows as a growing perpetuity with region specific growth, g_j. Both small landlords and large landlords have the same expected growth rate of $R - C$.

This mechanism for market power here is different from that in multifamily housing. For multifamily housing, market power results in a trade-off between raising rents and lowering occupancy, as in Calder-Wang and Kim (2024) and Watson and Ziv (2024). Because single family homes are more liquid, if the institutional investors want to lower quantity, they can sell the homes rather than leaving them vacant. I show the occupancy rates for Invitation Homes' markets in Figure B3. Occupancy rates fluctuate here by up to 4%. They increase sharply after the onset of the COVID-19 pandemic, when many people moved out of cities as documented in Gupta, Mittal, Peeters and Van Nieuwerburgh (2022) and Coven, Gupta and Yao (2023). They do not change substantially around the date of the merger when Invitation Homes bought Starwood Waypoint Homes. High occupancy rates, coupled with the difference in liquidity between single family and multifamily homes, suggest the mechanism for market power here is not raising rents and increasing vacancy. Instead, a company can maximize profits by not buying too many homes to keep rents higher and not cannibalize their own product.

IV.D SUPPLY OF HOUSING

The quantity of single family homes in region j, $Q_{j,own,new}$, is determined by the initial value in 2012, $Q_{j,own,2012}$, plus an amount that varies due to increases in the price of housing. I model this as an aggregate construction sector, with an elasticity of construction with respect to housing prices for each $PUMA_j$ of γ_j:

$$\log\left(\frac{Q_{j,own,new}}{Q_{j,own,2012}}\right) = \max\left\{\gamma_j \times \log\left(\frac{P_{j,own,new}}{P_{j,own,2012}}\right), 1\right\}. \tag{14}$$

The housing supply cannot shrink if prices go down. Elasticities are heterogeneous for each tract, and depend on distance to a city center and on land use features. I aggregate tract level elasticities from Baum-Snow and Han (2024) by taking the weighted average of each tract elasticity in a PUMA, weighted by the number of homes in the tract.

The quantity rentals in each region is determined by the demand of small and large landlords:

$$Q_{s,rent} = Q_{d,smalllandlords} + Q_{d,largelandlords}. \tag{15}$$

There is a multifamily rental asset class in the model that has perfectly inelastic supply.

IV.E Market clearing

Prices are implicitly defined by market clearing, which I rewrite as a function in logarithms and in vectors below:

$$\mathbf{p} = f(\mathbf{p}) = \log\left(\mathbf{P}\cdot\left(\mathbf{Q}_{d,\text{households}} + \mathbf{Q}_{d,\text{smalllandlords}} + \mathbf{Q}_{d,\text{largelandlords}}\right)\right) - \log(\mathbf{Q}_s). \tag{16}$$

For rentals, $Q_{d,largelandlords} = Q_{d,smalllandlords} = 0$ because landlords can own properties and cannot rent them from other landlords in this model. I describe how to compute counterfactual prices in Appendix C.

IV.F Estimation

I estimate household demand by first estimating bilateral migration costs, and then using methods for the estimation of models in the style of Berry, Levinsohn and Pakes (1995) similarly to Conlon and Gortmaker (2020). The household demand estimation uses cross sectional variation from a pool of census PUMAs for the whole US from 2012-2019, and bilateral migration data for the whole US from 2012-2019.

First I estimate $\tau_{(i,l)\to(j,k)}$ by examing the role of distance, social connectedness, and asset class transitions in bilarteral migration data. I use bilateral migration data from Verisk to construct a dataset of migration shares from PUMA i asset class l to PUMA j asset class k, where i can equal j and l can equal k, $w_{(i,l)\to(j,k)}$, which are the fraction of people from (i,l) who move to (j,k). I model migration shares to be based on characteristics of the destination, $X_{j,k}$, log price if k is owner occupied, $p_{j,k}$, log rent if k is a rental, $r_{j,k}$, and origin destination pair characteristics $T_{(i,l)\to(j,k)}$:

$$w_{(i,l)\to(j,k)} = \beta_p p_{j,k} + \beta_r r_{j,k} + \beta_x X_{j,k} + \beta_t T_{(i,l)\to(j,k)} + \varepsilon_{(i,l)\to(j,k)}. \tag{17}$$

Origin destination pair characteristics $T_{(i,l)\to(j,k)}$ include distance, social connectedness, and interactions between all possible asset class transitions and a same PUMA dummy.

I estimate (17) with linear IV. The endogeneity problem here is that the $p_{i,l}$ and $r_{i,l}$ may be correlated with $\varepsilon_{(i,l)\to(j,k)}$. When price and latent demand are positively correlated, this can bias β_p upwards. I instrument for the prices and rents in a PUMA with features of the housing stock and topography of neighboring regions, similar to Calder-Wang (2022), Bayer et al. (2007), and Bayer et al. (2003). The identification assumptions are that for a given PUMA's price and rent, characteristics of neighboring regions matter through a competition channel. But these neighboring PUMAs' characteristics do not affect the utility of one living in that PUMA because they are sufficiently far away, when controlling for own region characteristics.

I use topography describing how hard it is to build 3-10 miles from a region as an instrument, when controling for topography within 3 miles of a region. I construct these measures using the land unavailability measures at the zipcode level from Lutz and Sand (2022), which describe how much water is in each zipcode, how much the zipcode is covered by wetlands, and the overall unavailability of land in a zipcode. I take centroids for each census tract in the US and then create a 3 mile circle around them, and also a 3-10 mile ring around them. I do a geospatial join of this circle and ring to a zipcode level map of land unavailability chatacteristics. This gives me the land unavailability within 3 miles of each census tract, and 3-10 miles from the center of each census tract. I average this up to the PUMA level. This results in a measure for a PUMA that is the average unavailability for each house in the PUMA within 3 miles, and in a 3-10 mile ring away. I use three other instruments that are average values of housing stock characteristics of neighboring PUMAs: the median age of the housing stock, the median number of bedrooms per home, and the fraction of the regions that is single family housing. To support both the relevance and exclusion restriction of the instruments, I regress log prices and log rents on these instruments and other characteristics and show the results in Appendix Table B5. For log rents, all three topography features for within 3 miles have opposite signs to the corresponding features for the 3-10 mile ring. Higher land unavailability in neighboring regions increases rents, possibly due to the inability to build when there's less land available nearby. For log prices, land unavailability due to water for neighboring regions increases price and is of the opposite sign to land unavailability due to water within 3 miles. The statistical significance and the opposite signs suggest that characteristics for neighboring regions are relevant, and affect price and rent through a different channel than the PUMA's characteristics affect price and rent.

I examine the instrument spatially in Figure B10. Panel A shows the mean land unavailability due to water within 3 miles for each PUMA in Georgia, Panel B shows the mean land unavailability due to water 3-10 miles away, Panel C shows the difference between Panel B and Panel A, and Panel D shows log house prices in each PUMA. PUMAs near the city center have more land unavailability due to water in neighboring regions than in their own region. This appears to be correlated with housing prices, which could suggest that land in the city center has high prices because it is harder to build nearby. An alternative hypothesis for this correlation in the regression for the whole US could be that cities occur near ports but on land that is buildable but just near land that is not buildable, and that these ports drive higher house prices and rents. I control for distance to city center in these regressions, so the differential impact of neighboring topography relative to own topography on prices and rents must remain after controlling for proximity to city centers.

Table B4 shows the results of estimating equation (17). People are more likely to move to nearby PUMAs with a large social connectedness index than farther away PUMAs with fewer social connections. People are most likely to stay within the same asset class and within the same PUMA.

I recover estimates for β_t and use those to create the moving costs in utility terms:

$$\hat{\tau}_{(i,l)\to(j,k)} = \beta_t T_{(i,l)\to(j,k)}. \tag{18}$$

Once I have recovered the $\hat{\tau}_{(i,l)\to(j,k)}$, I partition the right hand side of (3) into two terms:

$$u_{h,(i,l)\to(j,k)} = \delta_{h,j,k} + \hat{\tau}_{(i,l)\to(j,k)}. \tag{19}$$

I plug this into (4) to get:

$$w_{h,j,k} = \sum_{i,l} \frac{\exp(\delta_{h,j,k} + \hat{\tau}_{(i,l)\to(j,k)})}{1 + \sum_s \exp(\delta_{h,j,k} + \hat{\tau}_{(i,l)\to(j,s)})} w_{h,i,l}. \tag{20}$$

I then use the contraction map from Conlon and Gortmaker (2020) to recover $\delta_{h,j,k}$. Finally, I estimate the following equation by linear IV

$$\delta_{h,j,k} = \beta_p p_{j,k} + \beta_r r_{j,k} + \beta_x X_{j,k} + \varepsilon_{(i,l)\to(j,k)}. \tag{21}$$

I instrument for prices and rents with the characteristics of neighboring regions as described earlier.

I constrain $\beta_{i,0,l} < 0$ to ensure demand is downward sloping. The demand functions for owning housing and the demand for renting have different coefficients, even for the same group of homeowners. This allows them to have different price and characteristic elasticities for each asset class.

I estimate the demand system for the panel of yearly housing holdings from 2012–2019 in a pooled regression. I do this on the extensive margin by adding one household to each region, resulting in tiny weights in regions where a given household group does not rent or own. The estimates return moving elasticities for price and rent for different groups. These elasticities are the percentage of a group that will leave to go to a different PUMA, not the percentage that will leave a house. The model abstracts away from downgrading within a PUMA. Someone who stays in the same PUMA but downgrades when faced with a price shock would be recorded as inelastic here because these are moving elasticities, and to different PUMAs. Therefore the elasticities will be lower than housing unit elasticities. I expect that as the geography size increases, the elasticity of a group to that geography's housing prices decreases. For example, it's easier to leave a neighborhood if neighborhood prices increase than leave the country if the country's prices increase. Around 80% of counties in the US have a population less than the minimum PUMA population,[4] therefore a person leaving the PUMA would be a larger move than a person leaving their county for the majority of PUMAs in the US. The moving elasticities are sufficient to study the question of how institutional investors increased the prices in an entire PUMA. There is substantial variation in price elasticities by income group, as shown in Figure B8. For both purchasing housing and renting, elasticity decreases with income. Functionally, low moving elasticities for high income groups mean that in the model, when large landlords shock the market, those making 100k+ will not leave to go to a different PUMA.

For the supply side, I use estimates from Baum-Snow and Han (2024). I take the tract level elasticities and aggregate them to the PUMA level through a weighted average by

[4] https://www.census.gov/library/stories/2017/10/big-and-small-counties.html.

the number of homes in each tract. This gives me a supply elasticity that is heterogeneous in each PUMA, which depends on the elasticities in the tracts that make up the PUMAs. These elasticities are heterogeneous and depend on a tract's distance to a city center and on topographic features.

I calibrate small landlord cost curves for each region using cost data from the RHFS, mortgage data from Verisk, and prices and rents from the census. Prices and rents for each PUMA come from the census ACS for 2012. I use a broker fee of 6%, which is standard.[5] For a small landlord i, I get individual costs by sampling N_j times for each PUMA from different cost and mortgage distributions, where N_j is the number of single family rentals in 2012. For the mortgage balance for a landlord, $M_{i,j,t}$, I sample from an empirical distribution of mortgage amounts for PUMA j from Verisk and then assume linear ammortization over 30 years to get a mortgage balance. For $OtherCosts_i$ and $ManagerFee_i$, I sample from a national distribution of operating expenditures and manager fees from the RHFS. I construct a property tax rate with data from Verisk where I take the amount of property tax payments on homes sold and divide them by the sales value, to get a property tax rate on market value rather than assessment value, which gives us $PropTax_j$. Interest expense is calculated from the median interest rate nationally for small landlords, from Verisk in 2012, and multiplying it to the perpetuity equivalent of the present value of expected interest rate payments. Expected interest rate payments are that the small landlord pays off 1/30th of their original mortgage balance each year. For small landlords, I set the discount rate for the equity cash flows to 5%. This is close to the CAPM value of 5.2% if I assume that they have the same asset betas as INVH and AMH from their IPOs to 2024, and if I assume zero leverage. However, there is evidence that small landlords like cash flows from rentals more than would be expected. We can see this in the RHFS in that some small landlords have 0 or negative cash flows, and some have less than the risk-free rate, even if we exclude interest expense. If we excluded all interest expense and set the discount rate to 0, there would still be some landlords who are unprofitable. Including interest expenses and considering the opportunity cost of money tied up in the property make it so that a larger number of landlords in the data should not be operating, but are operating. I calculate g_j in two steps. I first run a regression of rent growth from 2006–2012 at the PUMA level on population growth from 2006–2012, with a state fixed effect. Population growth is a dummy variable of above or below median population growth for the US:

$$g_{j,state} = \beta I_{\Delta pop_{j,state} \geq Median \Delta pop_{j,state}} + \alpha_{state} + \varepsilon_{j,state}. \tag{22}$$

[5] https://listwithclever.com/average-real-estate-commission-rate/#rates-by-state

This gives me a coefficient for the amount of past rent growth that was attributable to state level trends plus an indicator for above median population growth. I set the mean of the resulting distribution of expected rent growth to be the 5 year expected rent growth from the NY Fed's SCE data in 2014, which is the earliest year I have the data for, which is 4%. This process results in an estimation for expected growth that is the sum of a mean expected rent growth, plus a shifter for how much population growth mattered for rent in the past and how much state level trends mattered. The assumption here is that these trends will continue. This is a strong assumption, which is why I use a rough indicator for the relationship between population growth and rent growth rather than use the conditional mean for all levels of past population growth. I use the same expected rent growth for both small and large landlords.

To measure the fit of the small landlord estimation procedure, I plug in 2012 rents and prices into the estimation and compare estimated small landlord quantities with actual small landlord quantities. I compare them in both all of Georgia and in the PUMAs in Georgia where institutional investors combined own over 1000 homes and call these regions "High Investor Activity." Results are in Table B1. Estimated quantities are highly correlated with actual quantities. The estimation underestimates quantities in both Georgia and in the subset of Georgia where institutional investors are most active. The estimation performs better in regions with high institutional investor activity. Overall, the model can't explain why all small landlords operate. This is because there are a number of small landlords who are unprofitable in the RHFS data, which I calibrate landlord costs to. Even if I exclude interest expense and use a discount rate of 0, there are landlords who should not be operating. So it is not surprising that this estimation underestimates the landlords who operate. It does a better job when rent to price ratios are high. There are a few reasons why small landlords might have an additional propensity to be landlords above and beyond the financial returns when compared to other financial assets. First, they have option value to use the home for personal use some day. Second, they have control rights. Third, if a large number of small landlords are retirees, they might value stable cash flows more than the CAPM would suggest. I make up for this underestimate by inserting a residual so that the model matches the 2012 housing market exactly. What matters for the model is not the quantity of the underestimate (because I make up for this with a residual), but the slope of the supply curve where it intersects demand. And differences in elasticities between areas where the model is farther away from actual quantities vs where it is closer are not large.

I calibrate large landlord cost curves for each region using cost data from earnings statement supplements. Prices, rents, property taxes, and growth rates are the same for both landlord types. Large landlords have different operating costs, interest expenses, management costs, renovation costs, and discount rates. I use American Homes for Rent's 2014Q1 average market cost and subtract property taxes. I then apply a cost shifter for regional contractor wages, which are correlated with costs in different markets. I estimate this cost shifter with a regression of market level operating expenditures on state level contractor wages. Given the importance of foreclosures in where large landlords purchased homes, I allow large landlords to purchase homes only in the top 75%-ile of PUMAs in the number of foreclosures per unit from 2007-2011.

I show the large landlord estimation results spatially in Figure B9. Panel A shows the estimated quantities for 3 identical large landlords and Panel B shows the actual quantities for institutional investors in 2019. I choose 3 large landlords to enter because that is the mean number of large landlords in any one PUMA where institutional investors have significant presence. The spatial pattern is similar for both the estimated and the actual. The total number of units in Georgia matches as well, 21k estimated and 22k for the actual. The model over predicts entry in regions with high rent to price ratios but low foreclosures.

V Estimating the Impact of Institutional Investors

V.A Model results

I estimate the equilibrium impact of 3 identical large landlords who enter the housing market in 2012 and choose where to operate and how many units to operate in each PUMA. I choose 3 because in 2021, the mean number of institutional investors in a PUMA where the investors have at least 10 units each is 3. I implement the Newton step algorithm from Koijen and Yogo (2019) described in Appendix C to recover the unique market clearing price vector. I then observe market clearing prices, rents, and quantities of housing.

I begin by analyzing the impact of institutional investors on the number of homes available for homeownership. Institutional investors decreased the housing available for owner occupancy by 23% of the homes they purchased. The impact of the investors on homeownership is significantly less than 1:1 because of two supply responses. I show the impact on homeownership and the supply responses in Figure 5 Panel A. When an institutional investor purchases a home, that puts downward pressure on the number of homes for owner occupancy by 1. However, the institutional investor demand shock triggered a supply response: for each home institutional investors purchased, builders built 0.33 homes and small landlords sold 0.44 homes. On the other hand, institutional investors increased the number of homes available to rent by 0.56 homes for each home they purchased. In Panel B, I show that it is not 1:1 due to the crowding out of small landlords, who sell 0.44 homes. Both impacts are less than 1:1 due to supply responses. A back of the envelope calculation of the impact of institutional investors on homeownership that does not incorporate supply responses overestimates the impact by 4x.

For prices, institutional investors' demand shock caused prices to increase by 2.3pp per 1pp of the total housing stock they bought. In Figure 7 Panel A, I plot the model implied impact of institutional investor entry on home purchase prices, by the share of a PUMAs housing stock that the institutional investors purchased. I also plot the binscatter of the association of these investors and actual price increases from 2012-2019 in excess of the rest of the US. The model implied impact is significantly smaller than the data association, suggesting that the investors targeted regions where prices would have gone up had they not entered. While the model impact is lower than the increases observed in the data, the impact is economically meaningful in the regions where institutional investors purchased the most homes. In the most concentrated PUMAs, prices increased by 10%, which on a $300,000 home is $30,000. The model implied impact is not monotonically increasing in institutional investor share because each PUMA has a different price elasticity due to heterogeneity in those who live there, as well as a different supply elasticity. For the majority of regions where institutional investors entered, almost all of the observed price association is not due to these investors. Additionally, investors only entered a small portion of the country and are not responsible for the broad price increases in the US over this time period.

Institutional investor entry decreased rents on net because they increased the supply of rentals. I show the model implied impact of institutional investors on single family rents in 7 Panel B. The x axis is the share of a PUMAs single family rentals that institutional investors make up after entry. Institutional investors decreased rents by 0.8pp per 1pp of the total housing stock they bought. The model implied rent impact is in the opposite direction of the data association, which could be due to the institutional investors buying homes in areas where rents would have gone up without them. This impact incorporates both the market power of the investors and the operating efficiency. The investors are efficient enough operators that even with market power, they increase the number of rentals and decrease rents. If a policy maker considers only the market power channel, they would get the direction of institutional investor impact on rent wrong.

Institutional investor entry caused a price increase that led to capital gains for those who held homes throughout the period of this price increase. In the model, I calculate the capital gains of groups of households due to the price increase caused by institutional investors. I plot the model implied capital gains for each group in Figure 11. The highest income and middle income households get the most capital gains, due to a combination of who is most exposed to the investors and who is least likely to leave. The middle income groups are the most exposed but are also more likely to sell their homes, and high income homeowners are not as exposed but do not move at all in response to the demand shock.

I examine whether it's likely that the price and rent increases were due to institutional investors targeting regions where prices and rents would have gone up without their entry. In the ACS, I examine the change in number of households from 2012-2019 as a function of institutional investor exposure. I plot the association in Figure B2 Panel C. The areas where institutional investors purchased homes experienced outsized population gains when compared to the rest of the country. The investors in their IPO filings indicated that they targeted areas with expected population growth, to attempt to find areas with expected price and rent appreciation. The shape of the population curve matches the shape of the price and rent associations in Panel A, suggesting that institutional investors did not cause price and rent increases, but population growth caused them. The shape of the price and rent associations is not what one would expect if institutional investors caused the increases. One would expect monotonically increasing prices and rents with institutional investor presence, which is not what we see.

V.B Examining economic channels

To examine how much of the rent impact is due to market power, I simulate a merger between two of the four large landlords who enter the housing market in the model. I simulate a merger with no adjustment costs by having 3 large landlords enter instead of 4. I plot the changes in rents and quantities by PUMA in Figure 8. Panel A shows that in the median area of overlap, single family rents increased by 0.71%. This effect is increasing in the share of rental housing owned by institutional investors. The size of this channel is consistent with that measured in Gurun et al. (2023), which uses quasi-experimental variation due to mergers of large landlords and finds a rent impact of 0.5% in the region of overlap. The market power here is due to these companies operating fewer rental properties. In Panel B, I show the change in the number of single family rentals due to the merger. The quantity of single family rentals decreases by a median of 0.42% in the overlapped regions, and is decreasing in the market share of the institutional investors. As in any Cournot model, a company will decrease its quantity when competitors decrease their quantities. The institutional investors face a larger residual demand over which they are monopolists when there are fewer of them. I show this graphically in Figure 9. A merger here would move the equilibrium from 4 to 3, raising rents and decreasing quantities.

I examine the role of the construction response by estimating the impact of institutional investor entry when the supply of homes is not allowed to adjust. I show the price impact in Figure 10. The price impact would have been 2.5x as large had there been no construction response: 5pp per 1pp of housing stock purchased by institutional investors.

V.C POLICY SIMULATIONS

I simulate two government policies. Two policy proposals are very similar: the End Hedge Fund Control of America Act and the American Neighborhood Homes Protection Act. Both charge a tax of either $10,000 or $50,000 per home per year for each home a landlord owns above either 50 or 75 single family homes. The smaller tax, the $10,000, would more than double the operating costs of AMH and INVH, and therefore effectively ban them from the market. I simulate these policies by removing institutional investors from the market entirely. I estimate the structural model of housing demand and supply to match 2019 exactly, then I remove the 7 large landlords I study from the homes in their exact footprint from 2019. I then observe market clearing prices, rents, and quantities. I show the impact on prices in Figure 10. Prices decrease significantly when institutional investors are forced to leave the market, and rents increase. 58% of the homes vacated go to small landlords.

I also simulate a 5% rent growth cap for corporate landlords. I do this by taking the estimated structural model for the housing market in 2012 and changing the expected rent growth for only large landlords to be capped at 5%. I find that institutional investors enter the market and buy 19% fewer rentals. This leads to higher rents than in the no rent control baseline estimation. Both policies limit the quantity of rentals supplied by large landlords, and therefore lower the rental supply. The policies are counterproductive in the rental market because because they are designed to decrease rent increases from market power, and do not take into account the fact that large landlords on net increased the rental supply and decreased rents.

VI Conclusion

This paper shows that the entry of institutional investors into the single family rental market matters for prices, rents, and homeownership. Institutional investors can scale efficiently, which leads them to accumulate large portfolios and potentially have market power. This matters for the housing market because they can supply rentals cheaper, but also have a higher willingness to pay for housing which can raise prices. The large portfolios also give them market power. I estimate the impact of institutional investors on the housing market by building a structural model that incorporates landlord type heterogeneity.

I find that institutional investors did not cause the majority of the price increases that occurred where they bought homes. A large construction response and the crowding out of small landlords are of first order importance in mitigating the effect of their demand shock on prices and homeownership. I find that market power is not the driving mechanism of investor impact in the rental market, and is outweighed by the operating efficiency of large landlords which on net decreases rents. This estimation shows that concerns about homeownership impacts must consider supply responses, and that concerns about market power must be weighed against the fact that institutional investors increase the number of rentals due to their operating efficiency and on net decrease rents. Landlord operating cost differences matter for the housing market and drive the net effect in the rental market, highlighting that landlord differences are of first order importance to the housing market.

The associations of institutional investor entry with increases in prices and rents appear to be driven by selection: the investors targeted neighborhoods with expected population growth, and the areas they entered ended up experiencing large population increases relative to the rest of the country. Overall, I show that the popular narrative that institutional investors are raising prices is overstated, and that the concern that they increase rents through market power is directionally incorrect. Additionally, large landlords are now starting to build more homes to rent them out. Because this paper provides a flexible framework through which one can study many housing market topics that include landlord type heterogeneity, construction, and household behavior, future work can study the incentives of landlords to build homes and how that can impact the housing market.

References

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

An, Brian Y., 'The Influence of Institutional Single-Family Rental Investors on Homeownership: Who Gets Targeted and Pushed Out of the Local Market?," *Journal of Planning Education and Research*, 2023.

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

Baum-Snow, Nathaniel and Lu Han, 'The Microgeography of Housing Supply," *Journal of Political Economy*, 2024.

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

_ , Robert McMillan, and Kim Rueben, 'An Equilibrium Model of Sorting in an Urban Housing Market: A Study of the Causes and Consequences of Residential Segregation," Working Paper January 2003.

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

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

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru, 'Why is Intermediating Houses so Difficult? Evidence from iBuyers," Working Paper 2022.

Calder-Wang, Sophie, 'The Distributional Impact of the Sharing Economy on the Housing Market," Working paper 2022.

_ and Gi Heung Kim, 'Algorithmic Pricing in Multifamily Rentals: Efficiency Gains or Price Coordination?," Working paper 2024.

Chodorow-Reich, Gabriel, Adam M Guren, and Timothy J McQuade, 'The 2000s Housing Cycle with 2020 Hindsight: A Neo-Kindlebergerian View," *The Review of Economic Studies*, 04 2023, p. rdad045.

Conlon, Christopher and Jeff Gortmaker, 'Best practices for differentiated products demand estimation with PyBLP," *The RAND Journal of Economics*, 2020, *51* (4), 1108–1161.

Cournot, Antoine Augustin, *Recherches sur les principes mathématiques de la théorie des richesses* 1838.

Coven, Joshua, Arpit Gupta, and Iris Yao, 'JUE Insight: Urban flight seeded the COVID-19 pandemic across the United States," *Journal of Urban Economics*, 2023, *133*, 103489. Special Issue: JUE Insight Shorter Papers.

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

Desmond, Matthew and Nathan Wilmers, 'Do the Poor Pay More for Housing? Exploitation, Profit, and Risk in Rental Markets," *American Journal of Sociology*, 2019, *124* (4), 1090–1124.

Diamond, Rebecca, 'The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, March 2016, *106* (3), 479–524.

_ , Tim McQuade, 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), 3365–94.

Ellen, Ingrid Gould and Laurie Goodman, 'Single-family rentals: Trends and policy recommendations," Policy Proposal, Brookings 2023.

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.

Ganduri, Rohan, Steven Chong Xiao, and Serena Wenjing Xiao, 'Tracing the Source of Liquidity for Distressed Housing Markets," *Real Estate Economics*, May 2022.

Garriga, Carlos, Pedro Gete, and Athena Tsouderou, 'The economic effects of real estate investors," *Real Estate Economics*, 2023, *51* (3), 655–685.

Giacoletti, Marco, Rawley Z. Heimer, Wenli Li, and Edison G. Yu, 'Single-Family REITs and Local Housing Markets," Working paper 2024.

Gorback, Caitlin S., Franklin Qian, and Zipei Zhu, 'Impact of Institutional Owners on Housing Markets," Working paper 2024.

Gupta, Arpit, Vrinda Mittal, Jonas Peeters, and Stijn Van Nieuwerburgh, 'Flattening the curve: Pandemic-Induced revaluation of urban real estate," *Journal of Financial Economics*, 2022, *146* (2), 594–636.

Gurun, Umit G., Jiabin Wu, Steven Chong Xiao, and Serena Wenjing Xiao, 'Do Wall Street Landlords Undermine Renters' Welfare?," Review of Financial Studies 2023.

Jiang, Zhengyang, Robert Richmond, and Tony Zhang, 'A Portfolio Approach to Global Imbalances," *Forthcoming, The Journal of Finance*, 2023.

Koijen, R. S. and M. Yogo, 'Exchange rates and asset prices in a global demand system.," Working paper 2019b.

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

_ , Robert Richmond, and Motohiro Yogo, 'Which Investors Matter for Equity Valuations and Expected Returns?," *Forthcoming, Review of Economic Studies*, 2023.

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.

Lutz, Chandler and Ben Sand, 'Highly Disaggregated Land Unavailability," Working paper 2022.

McFadden, Daniel, 'Modeling the Choice of Residential Location," *Spatial Interaction Theory and Planning Models*, 1978.

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.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel, 'Segmented Housing Search," *American Economic Review*, March 2020, *110* (3), 720–59.

Raymond, Elora Lee, Richard Duckworth, Benjamin Miller, Michael Lucas, and Shiraj Pokharel, 'From Foreclosure to Eviction: Housing Insecurity in Corporate-Owned Single-Family Rentals," *Cityscape*, 2018, *20* (3), 159–188.

Saiz, Albert, 'The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Schubert, Gregor, 'House Price Contagion and U.S. City Migration Networks," Working paper 2021.

Watson, C. Luke and Oren Ziv, 'A Test for Pricing Power in Urban Housing Markets," Working paper 2024.

VII Tables and Figures

Figure 1: Tract Concentration of Institutional Investors (2021)

Panel A: Fraction of Housing Units Owned in Atlanta

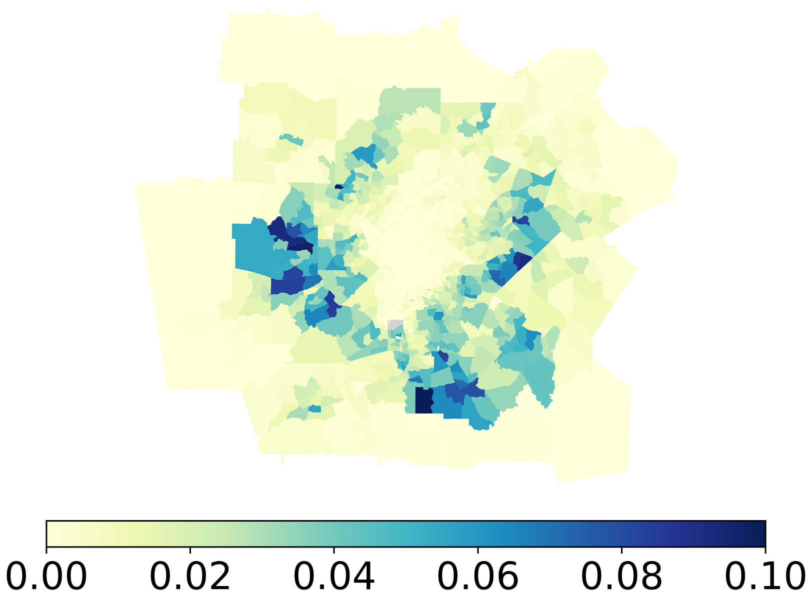

Panel B: Fraction of Rentals Owned vs Housing Units in the US

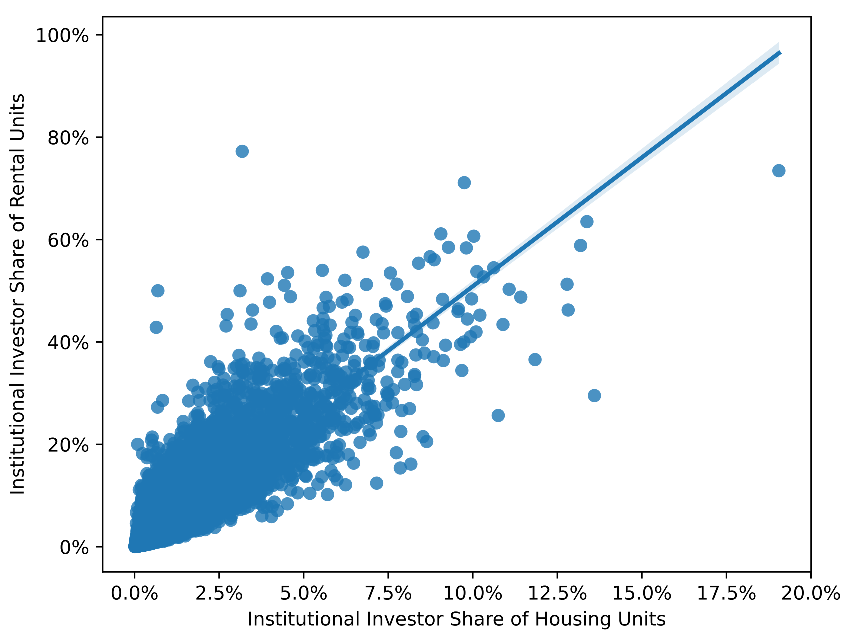

Notes: Panel A shows the fraction of residential housing stock in a census tract owned by the 7 largest institutional investors in February of 2021 in the counties including and surrounding Atlanta, GA. The investors included are Invitation Homes, American Homes for Rent, Tricon Residential, FirstKey Homes, Progress Residential, Main Street Renewal, and Home Partners of America. Panel B plots the fraction of housing stock in a tract owned by one of these investors against the fraction of the rental housing in a tract owned by one of these investors, for all tracts where at least one investor is present.

Figure 2: Scale comparison between small and large landlords

Panel A: Distribution of Single Family Rentals By Operator Size

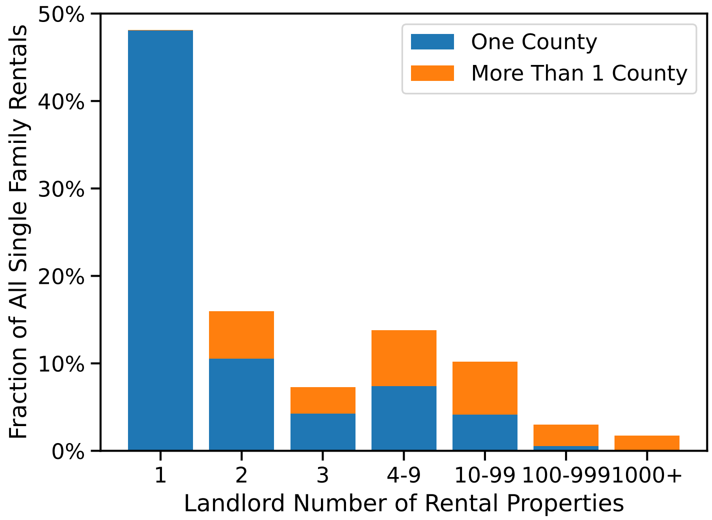

Panel B: Difference between large and small landlord operating costs

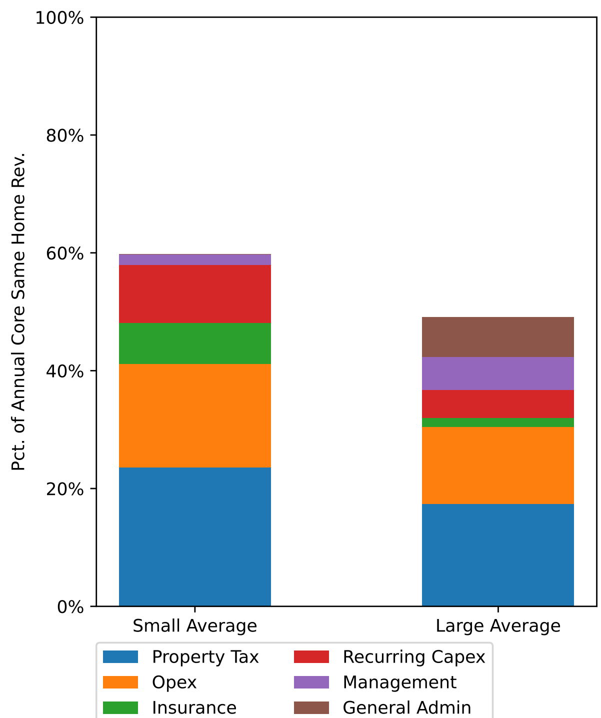

Notes: Panel A shows the fraction of all single family rentals in the US that each operator size bucket owns, and within each operator size bucket, the fraction that operates in only one county or multiple counties. Data come from a Verisk property snapshot from February 2021. Rental status is determined by whether the mailing address is the same as the property address. Panel B shows operating costs for the average 1 unit individual landlord in the Rental Housing Finance Survey, compared to operating costs for the average of Invitation Homes and American Homes for Rent, where data come from their earnings statements supplements. Data are from an average of fiscal year 2017 and 2020. For Invitation Homes and American Homes for Rent, all data except for General Admin, Management, and Insurance are from their "same home portfolios", which excludes recently acquired homes or homes that are preparing to be sold. For the small landlords, recurring capital expenditure is the capital expenditure for categories that include HVAC, roof, and floor.

Figure 3: Scale comparison between small and large landlords

Panel B: Cost Heterogeneity for Small Landlords

Panel A: Change in Opex Per Home From INVH and SWH Merger

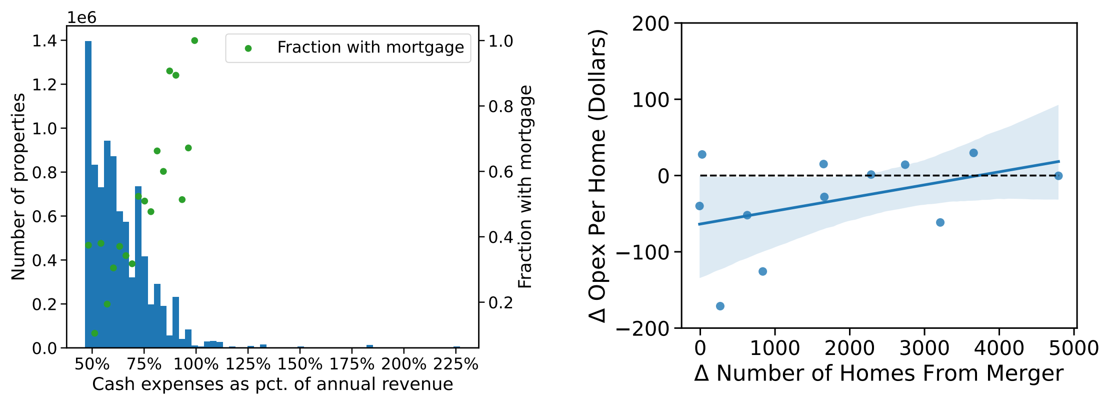

Notes: Panel A shows a histogram of individual 1 unit landlords in the Rental Housing Finance Survey. The entries are distributed by cash expenses as a fraction of rent revenues. The green dots are the fraction of operators in each bucket who have a mortgage. Panel B shows for Invitation Homes, the change in same home operating expenditures per home in each market by the change in number of homes in each market when it merged with Starwood Waypoint Homes. The dotted black line indicates no change in operating expenditures per home.

Figure 4: Stylized Example of Large Landlord Entry

Panel A: Competitive

Panel B: Cournot Oligopoly

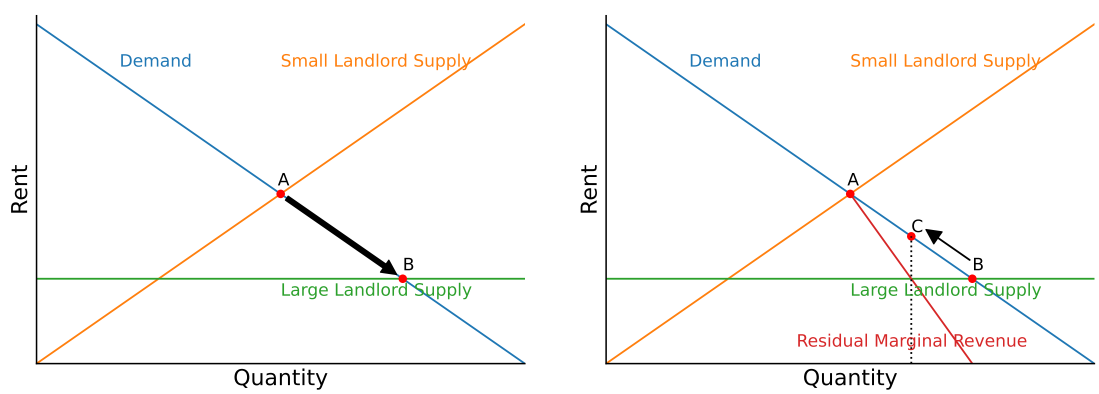

Notes: This figure shows a stylized model of supply and demand for single family rentals. Households have upward sloping demand and small landlords have downward sloping demand. In Panel A, one large landlord with constant returns to scale enters and behaves competitively, shifting the equilibrium from A to B. In Panel B, the large landlord chooses a quantity where residual marginal revenue intersects its cost curve, shifting the equilibrium from B to C.

Figure 5: Quantity changes when institutional investors enter

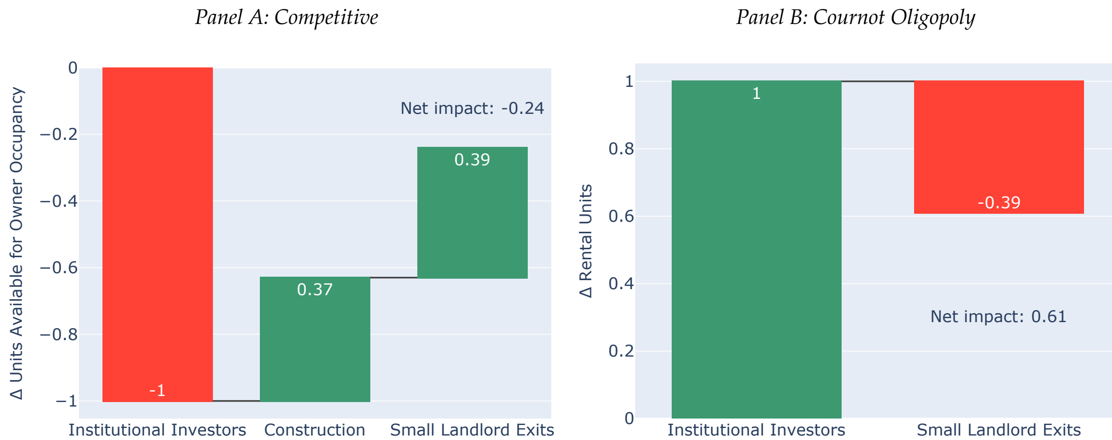

Notes: Panel A shows the change in housing available for owner occupancy due to a purchase of 1 unit by institutional investors. It shows the initial change from the purchase, and then the construction response and the response of small landlords. Panel B shows the change in total rentals available due to the purchase of 1 housing unit by institutional investors. It shows the initial change from the purchase and then the response by small landlords.

Figure 6: How do groups gain homes / lose rentals when B2R Exits?

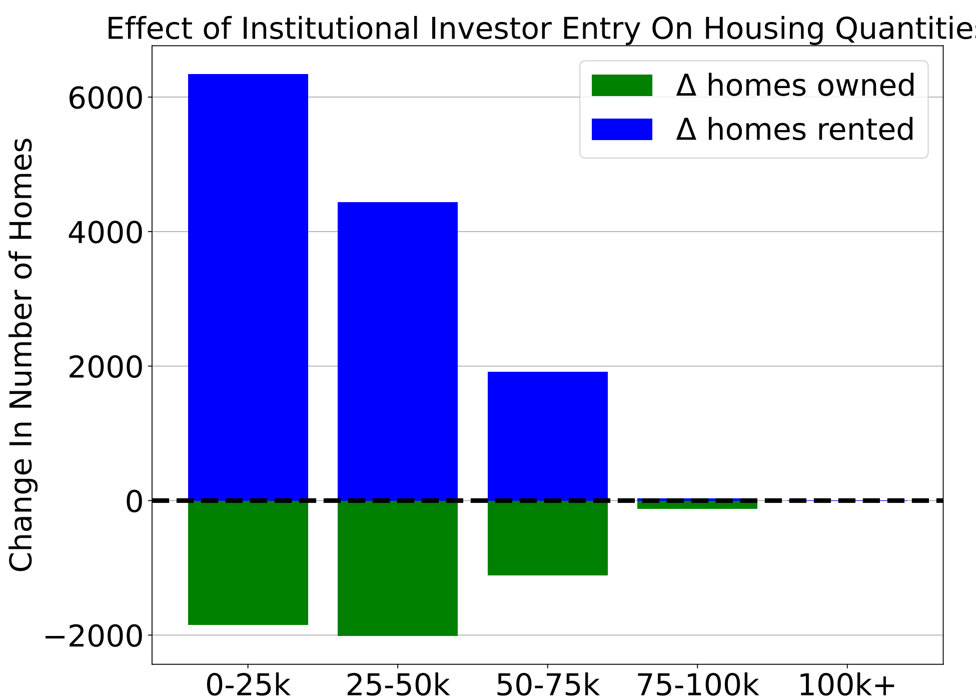

Notes: In the counterfactual where B2R exits the market and rentals that it converted from owner-occupied houses revert back to being owner-occupied, this figures show how each groups' housing allocations change. The green diamonds show the number of homes each group buys when B2R exits the market. The blue triangles show the number of rentals each group loses when B2R exits and the rentals it created revert. They do not sum to the same number because a group can move to the modeled outside asset, which consists of PUMAs with no characteristic data, housing with a median year built of 1939 or older, and housing with low prices and rents, or the unmodeled outside asset which would be equivalent to positive or negative household formation.

Figure 7: Price and rent changes when institutional investors enter

Panel A: Competitive *Panel B: Cournot Oligopoly*

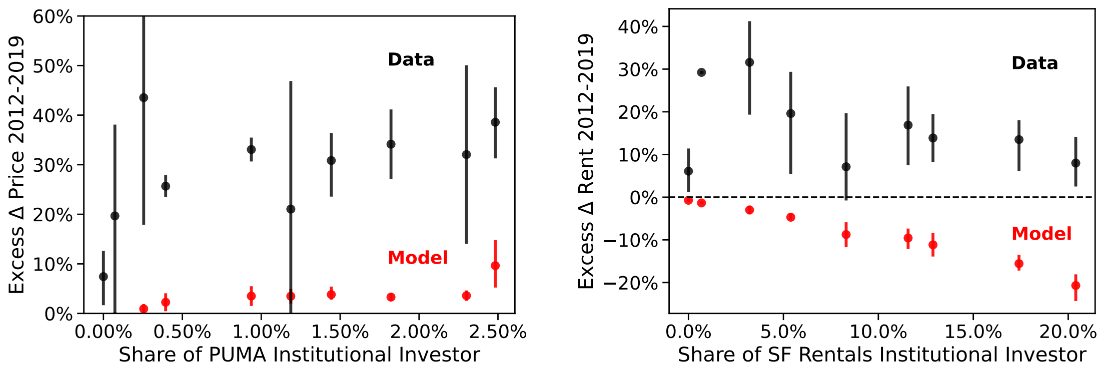

Notes: This figure shows the model implied price and rent impacts of institutional investor entry into Georgia. The x axis shows the share of the entire PUMAs housing stock or rental stock they own, and they y axis shows the excess increase in price or rent compared to the rest of the US. The black binscatter shows the data association from 2012-2019 of these investors with prices and rents. The red binscatter is the model output.

Figure 8: Impact of a large landlord merger

Panel A: Change in Rents *Panel B: Change in Quantities*

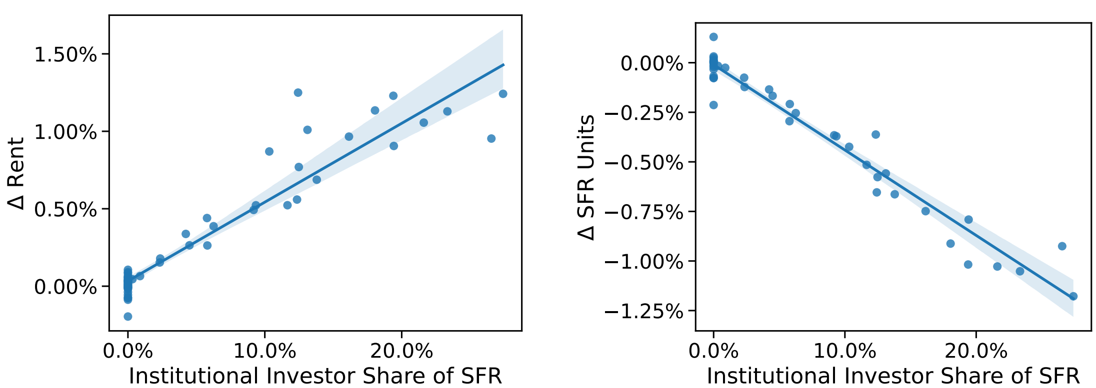

Notes: This figure shows the model implied impact of two large landlords merging, with no adjustment costs. Panel A shows the change in single family rents due to the merger in each PUMA, and Panel B shows the change in the quantity of single family rentals in each PUMA.

Figure 9: Stylized example of merger

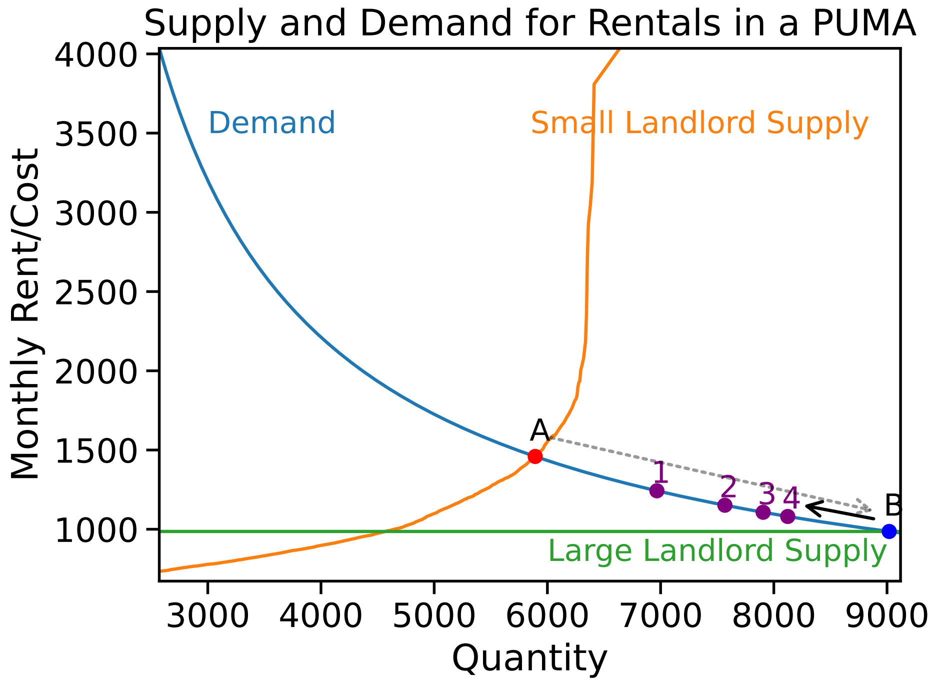

Notes: This figure shows the equilibrium quantities and rents for one PUMAs single family rental market when there are no large landlords (A), when there are either infinite or the ones that exist choose competitive quantities (B), and then when there are 4, 3, 2, or 1 large landlords. A merger between 2 of 4 companies would be a move from (4) to (3).

Figure 10: Price impact in counterfactual experiments

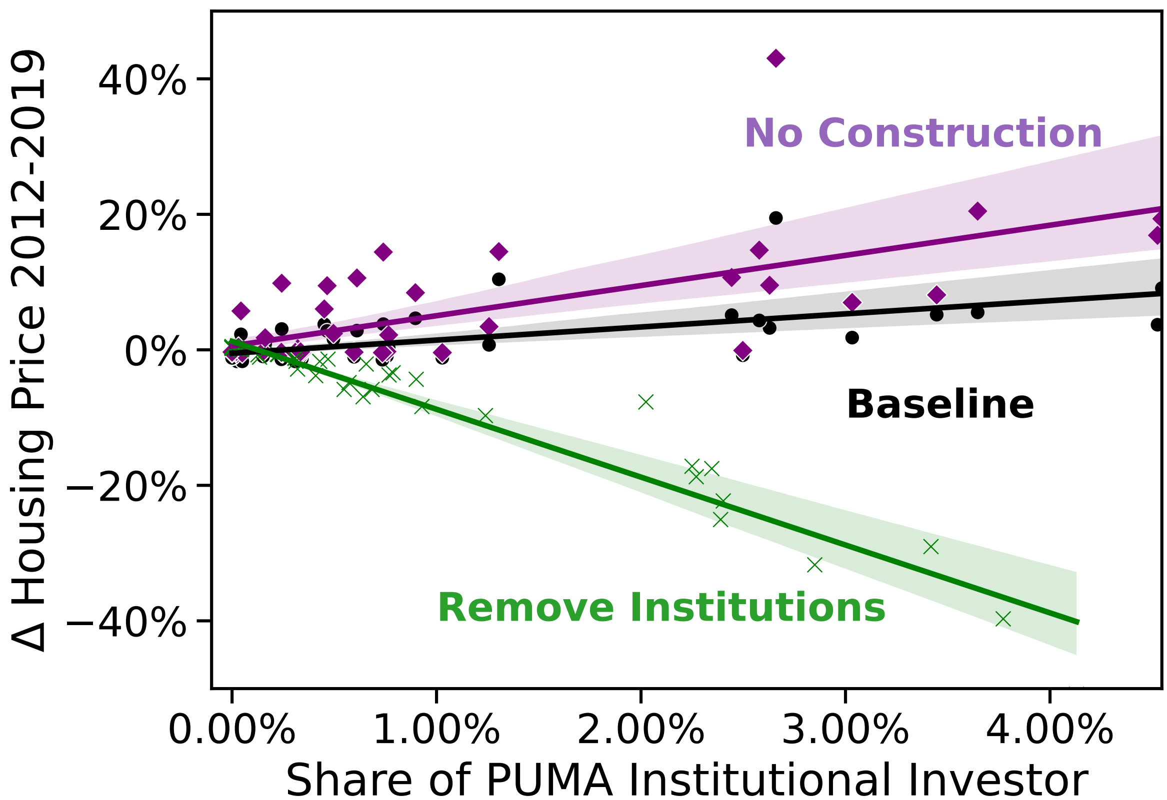

Notes: This figure shows the price impact in each PUMA for the baseline estimation and various counterfactual experiments.

Figure 11: Institutional Investor Impact on Housing Wealth

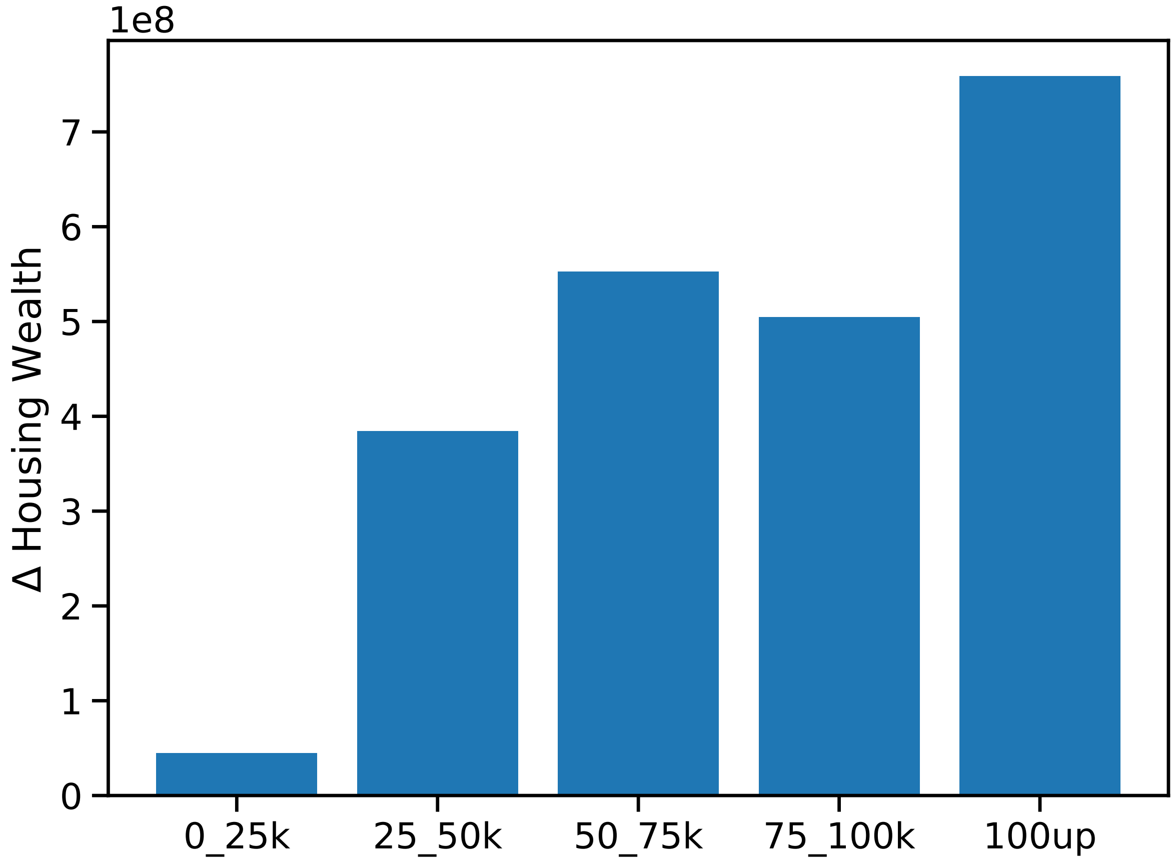

Notes: This figure shows the how many dollars a group gained due to B2R's entry as implied by the counterfactual.

Table 1: Difference in Tract Characteristics For Movers Into Institutional Investor Homes

	Mean Tract Difference
Δ Med. HH. Income	12.2%
Δ Math Scores	5.8%
Δ Jail Rate	-6.0%
Δ Top 20%-ile Income	3.4%

Note: This table shows the mean difference between destination tract characteristics and origin characteristics for those who moved into institutional investor properties for the first time between November 2018 and November 2019. The first row is the difference in median household income from the ACS, the second row is the difference in 2013 math test scores from Opportunity insights, the third row is teh difference in historical likelihood to go to jail from Opportunity insights, and the final row is the historical likelihood to get into the top income quintile from Opportunity insights.

Table 2: Previous Region log Outcome Variables on Moving Into a B2R Home

	log(med. income)	frac college	log(math scores)	log(jail)	log(inc. top quintile)
	(1)	(2)	(3)	(4)	(5)
new to institutional investor home	-0.011***	-0.010***	-0.010***	0.045***	-0.023***
New Tract FE	Y	Y	Y	Y	Y
Observations	591776	591776	591776	591776	591776

Note: *p<0.1; **p<0.05; ***p<0.01. The regression is at the individual mover level for all movers I observe in the US between November 2018 and November 2019, inclusive, who move into a zipcode with B2R presence, who have tract data. Clustering is at the new tract level. It compares within a given new tract, those who move into a B2R home for the first time vs those who move into a non-B2R home or those who move into a B2R home who were in a B2R home previously. The outcome variables are the logs of the measures for the origin geography. The first column is the previous region's median household income from the ACS 5 year tables. Column 2 is the previous region's 3rd grade math test scores from 2013. Columns 3-5 are outcomes of children from a given region including the fraction of children from that region who end up in jail, the fraction who end up in the top income quintile, and the fraction who are married by age 32. These represent historic mobility measures, not future outcomes. The outcome variables in columns 2-5 come from Opportunity Insights.

Appendix

A Data Appendix

A.1 Property Sample

Steps to construct the Verisk property sample:

- Keep properties that have a property indicator for single family residence, townhouse, apartment, condominium, duplex, triplex, or quadplex
- Exclude properties with no street information
- Exclude mobile homes
- Exclude remaining properties with a duplicate address indicator
- owner-occupied: owner-occupied flag of O or S. A, T, and null I call rentals.

Table A1: Verisk Housing Units By Zipcode Compared to Census Units

	Total Units	Owner Occupied	Rental Units
Units	1.064***		
	(0.001)		
Owner Occupied Units		0.991***	
		(0.001)	
Units Not Owner Occupied			1.086***
			(0.004)
Observations	32,456	32,456	32,456
R^2	0.940	0.962	0.744

Note: *p<0.1; **p<0.05; ***p<0.01. This table regresses ACS5 housing units from 2020 (includes all of 2020) on Verisk housing units for February 2021. The first column is all units, the second is owner-occupied, the third is rental.

There are multiple reasons why the sample undercounts rental units. It is possible that units in Verisk that are rented out used to be owned and their owner occupancy codes were not updated. Also, for multi-unit apartment buildings, sometimes Verisk has one row for each apartment in the building and other times it has one row for the entire apartment where the apartment number field details the number of apartments in that row. However, sometimes there are rows that look like they should represent one apartment but have the entire building's number of units as the apartment number. To clean this, I identify any row where the value of the address, either assessed or market, divided by the number of units, is less than 50 thousand. If that's the case, I change the number of units for that row to be 1. I also change to 1 the apartment number of any row that has a living sqft per unit of less than 100.

I validate the number of institutional investor owned homes I am able to identify for each public company in the table below. The identification is described in the data section of the paper.

Table A2: Verisk Number of Units Compared to Public Information

	Data	Target	Percentage
Invitation Homes	73066	80177	91%
American Homes For Rent	47699	53584	89%
Tricon	19630	22766	86%

Note: This table takes the number of units belonging to each company identified with Verisk data and compares it with information available in SEC filings for 2020Q4.

Figure A1: Supply validation

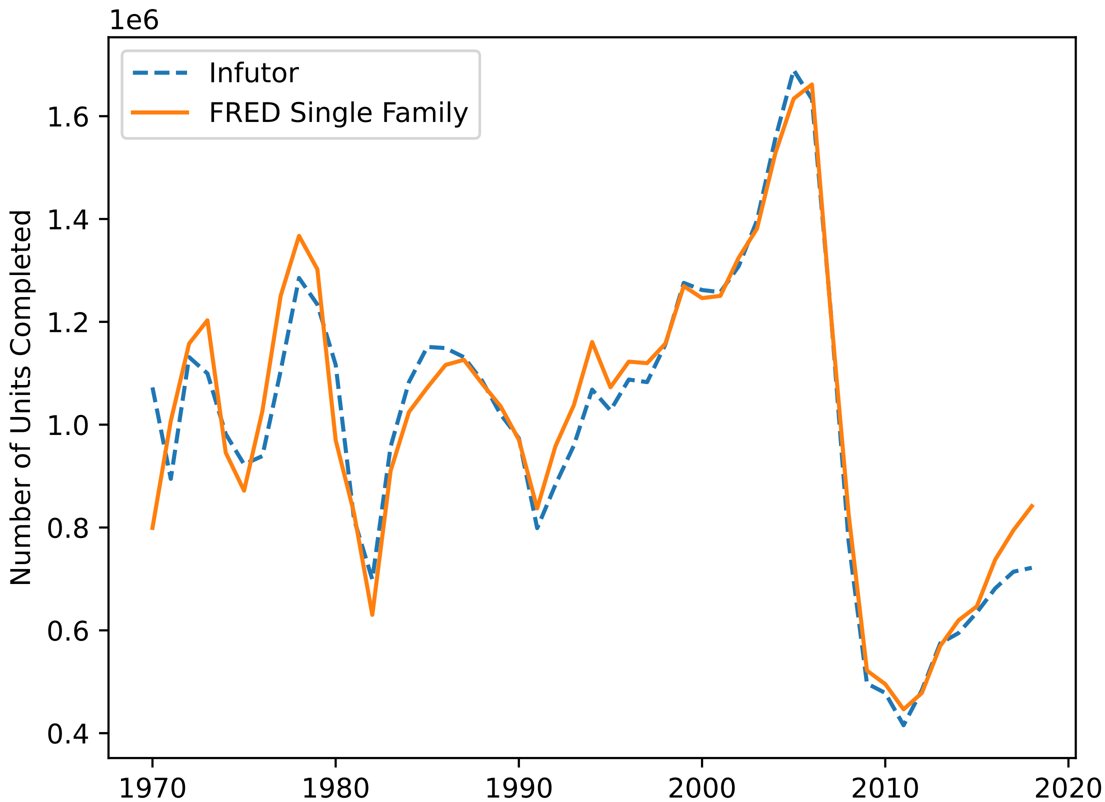

Note: I plot the aggregate number of completions for the US each year from Verisk, where I infer the year a property was built to be its completion date. I compare this with aggregate US completions for single family homes data from FRED.

A.2 Mover Sample

Steps to construct the Verisk moving sample:

- Clean property dataframe for 2019 as described above
- Take the Verisk moving history and exclude anyone who has a deceased flag of Y or null

I use an anonymized dataset so I can't drop duplicate address histories using names, or similar names. Instead, I drop all address histories that are identical. This is possibly an over correction, as different people with the same address history would be dropped.

- Reshape from wide to long to get the dataset at the person ID x address x previous address level, with the date recorded at the previous address and date recorded at the current address
- I drop those with null zipcodes, true duplicates, and duplicates

At this stage, there are many duplicate PID x EFFDATEs. I want to identify one address at a given date for a given person.

- Of the duplicates at the PID x EFFDATE level, I drop those that don't merge to a property identifier from the cleaned property dataframe
- Of the remaining duplicates, I rank them based on their postal delivery designation in the following order: street or residential, rural route, general delivery, high rise or business, PO box, null, firm or company address. I keep the duplicate that has the first rank in that order.
- I drop all entries that have remaining duplicates, and do not select one of the duplicates to keep

For a given year, I want to see who moves. I create a window, where those who moved before the start of the window are considered residents of their most recent location, and those who moved during the window are considered movers from their previous location to their new one. Those who move after the window are not considered movers.

- Create moving window for 201811-201911
- Clean property dataframe for 2018
- For each property, the 2018 owner is the old owner and the 2019 owner is the new owner
- Merge the moving window dataframe to the 2019 property frame
- Use 2019 property frame for the attributes of where people moved from and where they moved to
- Merge to cleaned 2019 demographic file
- Create flags for indicator variables
- Output those who have moved

Table A3: Mover level Validation Moves To and Moves From

	Moved From Zipcode (USPS)	Moved To Zipcode (USPS)
Moved From Zipcode (Infutor)	3.239***	
	(0.002)	
Moved To Zipcode (Infutor)		2.944***
		(0.002)
Observations	291,168	291,168
R^2	0.867	0.872

Note: *p<0.1; **p<0.05; ***p<0.01. This table regresses Verisk moves on USPS zipcode moves. The first column compares moves out of a zipcode in both datasets. The second column compares moves into a zipcode in both datasets.

Table A4: Mover level County To County Validation

	Moved From County 1 to 2 (ACS)
Moved From County 1 to 2 (Infutor)	1.622***
	(0.002)
Observations	426,676
R^2	0.703

Note: *p<0.1; **p<0.05; ***p<0.01. This table regresses Verisk county to county moves on Census county to county moves.

B Additional Figures and Tables

Figure B1: Fraction of Housing Owned by Institutional Investors in February 2021

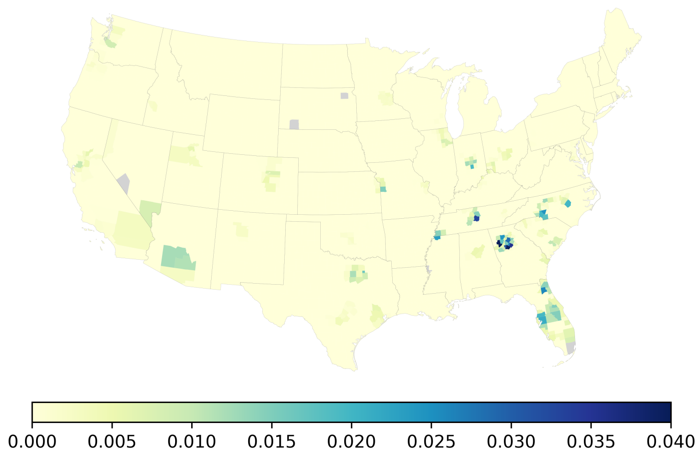

Note: For the US in February 2021, I show the fraction of the total housing stock that 7 institutional investors combined own at the county level. The 7 I include are Invitation Homes, American Homes for Rent, Tricon Residential (now owned by Blackstone), Progress Residential, FirstKey Homes, Main Street Renewal, and Home Partners of America.

Figure B2: Associations with B2R Share of Housing Stock

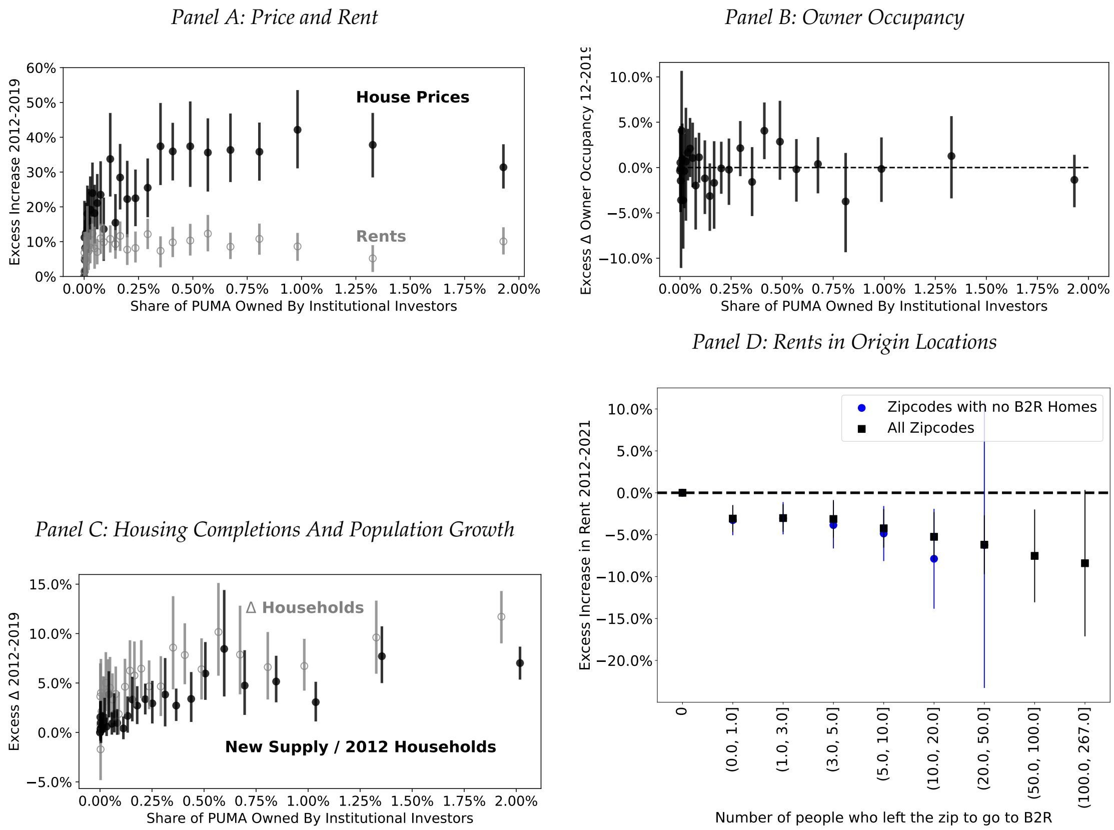

Notes: This figure shows the coefficients of a regression of different zipcode level variables on the fraction of the housing stock owned by B2R. Zipcodes with no B2R presence are the intercept, so each plot shows the excess of the variable relative to the rest of the country. Panel A is the change in the fraction of homes owner-occupied from 2012–2021 from the ACS5. Fraction owner-occupied for a given age group is the number of owner-occupied households in the age group divided by the number of all households. Panel B is the change in fraction renter occupied from 2012–2021 from the ACS5. Panel C is the change in ZHVI from 2012–2021. Panel D is the change in the median contract rent from the ACS5 from 2012–2021. Panel E shows the housing completions per home relative to the rest of the country. Panel F plots the coefficients of a descriptive regression of the change in rent from 2012–2021 on indicators for the number of people who left the zipcode to move to a B2R home. This regression includes county level fixed effects. There was no zipcode where more than 50 people left to go to B2R if there was no B2R home also present in that zipcode.

Figure B3: Market level occupancy rates for Invitation Homes

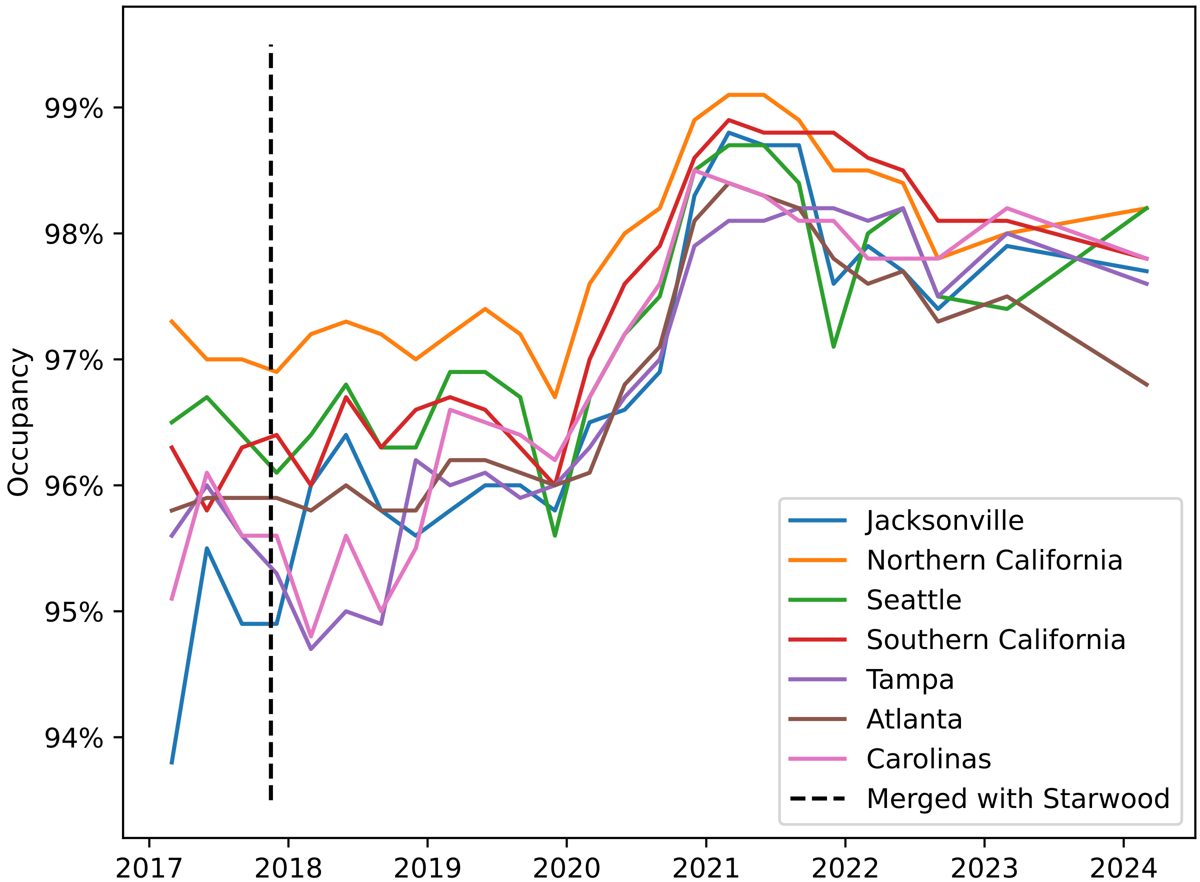

Note: This figure shows the market level occupancy for Invitation Homes for a number of its markets. The dotted black line is the date where Invitation Homes merged with Starwood Waypoint Homes and gained 32,000 homes.

Figure B4: Where did those who rent from institutional investors come from?

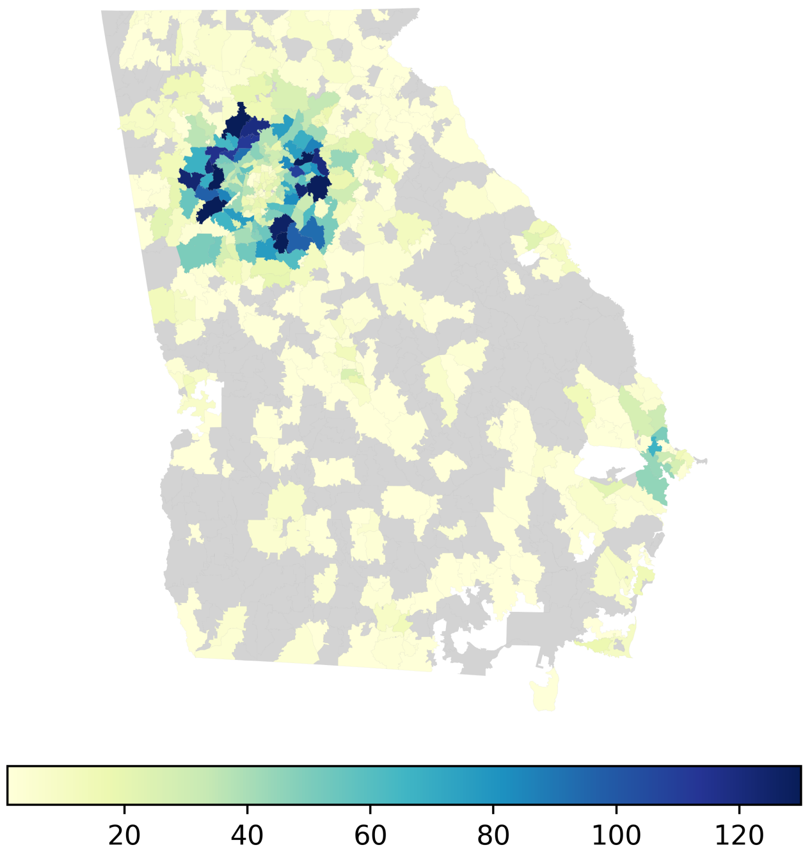

Note: This figure shows the origin locations of those who moved into an institutional investor home between Jan 2012 and Feb 2021. For each home, I include the movers only after the most recent sale date of the institutional investor home. I exclude those who moved from one institutional investor home to another institutional investor home.

Figure B5: Cost of Debt

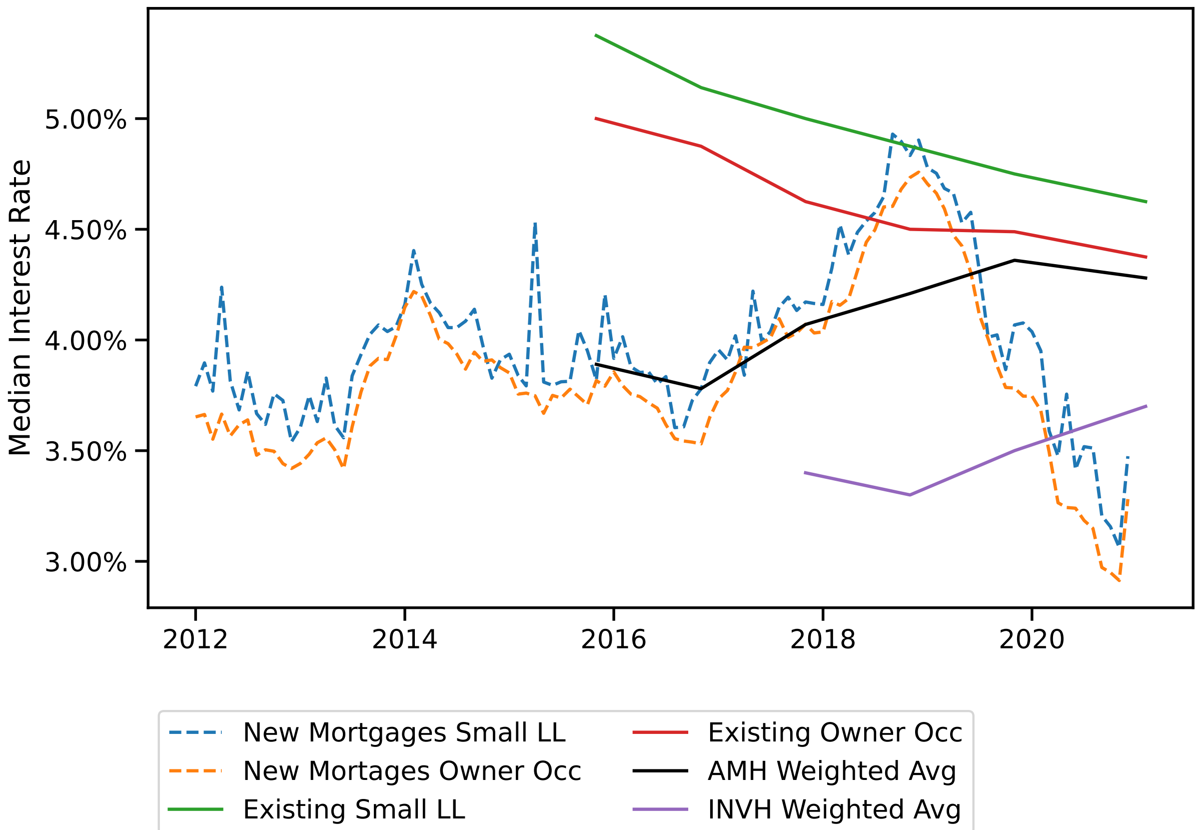

Note: This figure shows median interest rates for new mortgages for small landlords, new mortgages for owner occupied homes, all existing mortgages for small landlords, all existing mortgages for owner occupants, American Homes for Rent's debt, and Invitation Homes' debt. Data for the small landlords and owner occupants come from Verisk. Data for American Homes for Rent and Invitation Homes come from earnings statement supplements. The time series is limited for existing mortgages because I have snapshots of the data starting from November 2015 until February 2021. For Invitation Homes, the timeseries is limited to after its IPO date.

Figure B6: Small Landlord Mortgages

Panel A: Fraction of Small Landlords with Mortgage by Purchase Date

Panel B: Mortgage % Remaining

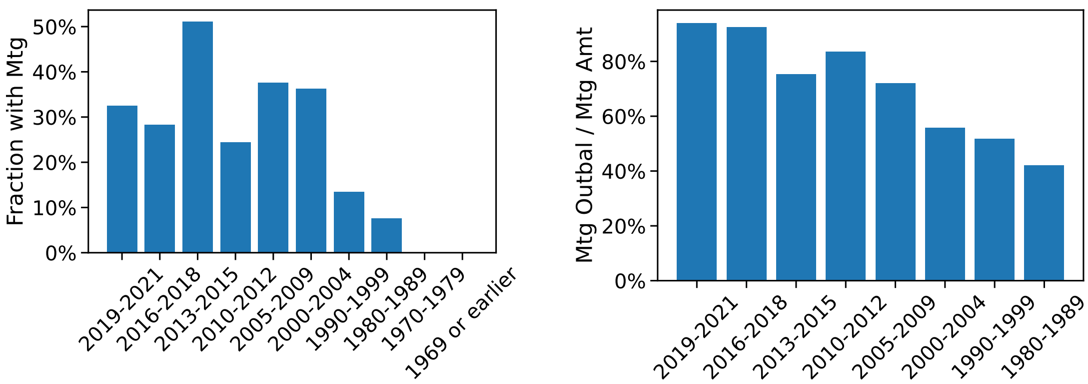

Notes: Panel A shows the fraction of a small landlord bucket that has a mortgage in the 2021 vintage of the RHFS, where each bucket is a purchase date grouping. Panel B shows the fraction of the original mortgage balance remaining for each purchase year bucket for those who still have mortgages.

Figure B7: INVH Market Average Opex Per Home

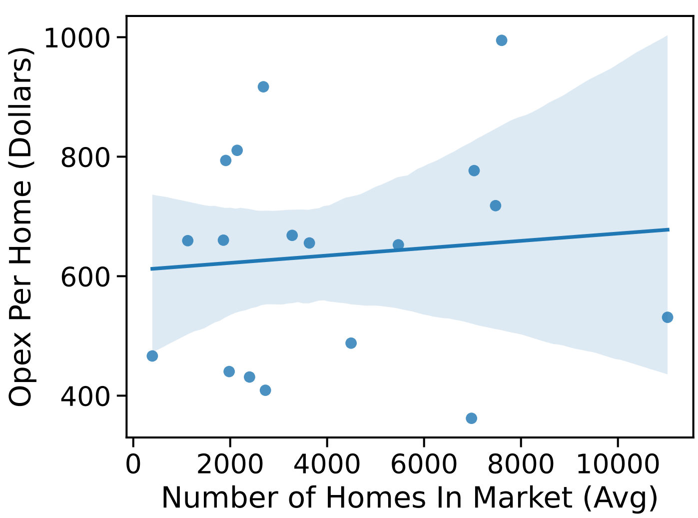

Note: This figure shows Invitation Homes' market level average operating expenditures per home, by the number of homes in each market.

Figure B8: Demand Elasticities

Notes: This figure shows the PUMA level moving elasticities to prices and rents of aggregate groups of households as estimated by the demand system. Elasticities are estimated using ACS1 data from 2012–2019 at the PUMA x year level. The elasticity is the percentage that each group decreases its quantity when price or rent increases.

Table B1: Small landlord estimation

	Georgia	High Investor Activity
Corr($Q_{est,PUMA}$,$Q_{actual,PUMA}$)	98%	99%
Median $Q_{est,PUMA}$ / $Q_{actual,PUMA}$	0.82	0.89
Median Elasticity with Respect to Rent	0.52	0.38
Median Elasticity with Respect to Price	-0.60	-0.48

Note: This table shows estimation results when I plug in 2012 rents and prices into the small landlord estimation. I compare estimated PUMA quantities to actual quantities in 2012, for all of Georgia and for PUMAs where institutional investors have 1000 or more units which I call ”High Investor Activity” regions. I first report the correlations between the estimated quantities and the actual quantities. Then I report the median of the estimated quantity divided by the actual quantity. Then I raise rents by 1% and measure the change in quantity in each region and report the median change as the median elasticity to rent. I then raise prices by 1% and measure the change in quantity in each region and report the median as the median elasticity with respect to price.

Figure B9: Large Landlord Estimation Fit

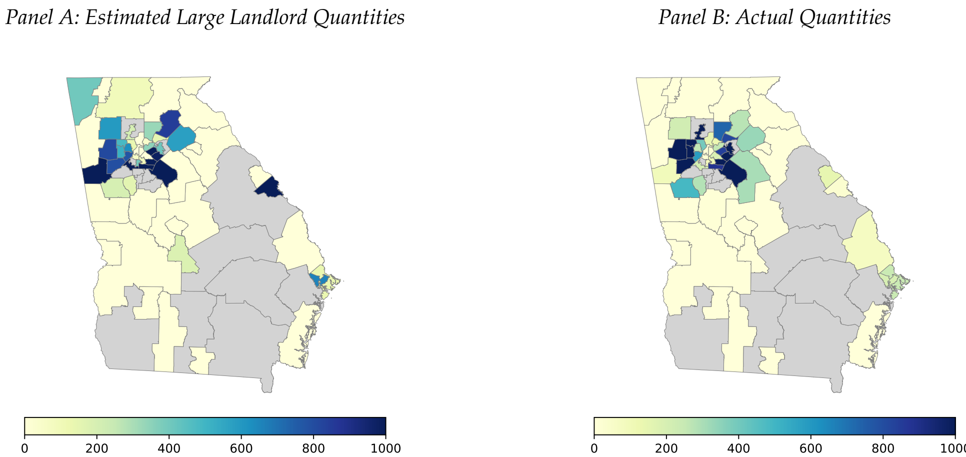

Notes: I estimate the quantities chosen by 3 identical large landlords and plot the sum for Georgia in Panel A. I compare this to the actual institutional investor quantities in Georgia in Panel B. .

Table B3: Within Zipcode Differences in Housing Characteristics

	Rentals Not Institutional	Rentals Institutional	Owner Occupied
avg year built	1985.54	1993.02	1988.93
avg living sqft	1831.05	1834.59	1984.85
avg num. beds	3.13	3.36	3.30
avg num. baths	2.26	2.42	2.43
frac. single family	0.79	0.99	0.91

Note: From Verisk property data I identify institutional investor holdings and compare their physical features in a given zipcode with the features of the rest of the rental stock and of owner-occupied homes. I take these within zipcode differences and compute a weighted average, weighted by the number of units institutional investors have in a given zipcode. This results in a weighted average within zipcode difference in physical characteristics between institutional investor homes and other types of homes.

Table B2: Where are institutional investors located?

	Dependent variable: Institutional Investor Presence	
	(1)	(2)
log(Price)	-0.366***	-0.161**
log(Rent MF)	0.057	0.259***
log(Rent SF)	0.388***	0.308***
ΔPrice 06–12	0.011	-0.135*
ΔPopulation 06–12	0.491***	0.345**
Avg. Annual Job Growth 04–13	1.143**	1.118**
ΔPrice 10–12	-0.014	0.072
ΔRent 10–12	0.009	0.046
Foreclosures per Person	4.636***	4.317***
Dist. To Nearest MSA	0.005**	0.003*
Dist. To Nearest MSA Sq	-0.000**	-0.000*
log(Med. HH Income)	0.077	0.319***
Frac. White	-0.359***	-0.278***
Frac. College Edu	-0.604***	-1.205***
Middle School Math Scores 2013	0.130***	0.024
Housing Stock Controls	Y	Y
Weather Controls	Y	Y
Other Amenity Controls	Y	Y
Fixed Effects		State
Within R-squared	0.349	0.256
Observations	1555	1555

Note: *p<0.1; **p<0.05; ***p<0.01. This is a descriptive regression at the PUMA level. The dependent variable is an indicator variable for a PUMA if institutional investors have 10 or more properties in that PUMA, combined. Column (1) has no fixed effects and Column (2) has state level fixed effects. Prices and rents are median values from the Census ACS 1 year tables for 2012. MF is multifamily, SF is single family. 2006 prices and population counts are from the 2006 Census ACS 1 at the PUMA level. I use a crosswalk from 2000 PUMAs to 2010 PUMAs from the Missouri Census Data Center so I can compare values from 2006 to 2012. Average annual job growth from 2004-2013 and middle school math scores from 2013 come from Opportunity Insights at the tract level, which I aggregate up to the PUMA level. Foreclosures per 2012 population comes from foreclosure data from ZTRAX from Zillow, and population data from the 2012 census. A PUMAs distance to nearest MSA comes from the distance of each zipcode in a PUMA to the center of the nearest MSA, which I then aggregate to get an average at the PUMA level. Housing stock controls are the median year built of the owner occupied housing, the median number of bedrooms of the owner occupied housing, and the fraction of a PUMA that is single family. Weather controls are the January temperature and sunlight, and the July temperature and humidity. Other amenity controls come from the Census ACS 2012 1 year tables and are the fraction of high school age population that is enrolled in high school, the fraction of high school age population that is enrolled in private school, and the fraction of total population that has a commute under 45 minutes long.

Table B4: Migration Share IV

	$log(w_{(i,l)\to(j,k)}/w_{(i,l)\to 0})$
$sqrt(Distance_{i\to j})$	-0.0003***
$log(SCI_{i\to j})$	0.4282***
$ooc \to ooc$, same puma	-0.7401***
$ooc \to ooc$, diff puma	-7.7562***
$ooc \to sf$, same puma	-4.8917***
$ooc \to mf$, same puma	-8.0212***
$ooc \to sf$, diff puma	-8.2135***
$ooc \to mf$, diff puma	-8.8547***
$sf \to ooc$, same puma	-2.5084***
$sf \to ooc$, diff puma	-6.3004***
$sf \to sf$, same puma	-1.1680***
$sf \to mf$, same puma	-6.8053***
$sf \to sf$, diff puma	-6.6063***
$sf \to mf$, diff puma	-7.2718***
$mf \to ooc$, same puma	-2.4026***
$mf \to ooc$, diff puma	-4.3677***
$mf \to sf$, same puma	-3.3508***
$mf \to mf$, same puma	-0.0216
$mf \to sf$, diff puma	-4.6631***
$mf \to mf$, diff puma	-5.3584***
$log(Rent)$	-0.3049***
$log(Price)$	-0.1279***
Topography controls	Y
Weather controls	Y
Housing characteristics controls	Y
Amenity controls	Y
n. obs.	3304840

Note: *p<0.1; **p<0.05; ***p<0.01. This table shows the IV regression of the log share of a PUMA moving from origin i to destination j divided by the share moving to the outside asset. The outside asset is defined as all PUMAs missing a non-price characteristic, or with housing prices below 80k or rent below 500. This is a pooled regression using bilateral migration data from 2012-2019 from Verisk.

Table B5: First Stage For Household Demand

	log(Price)	log(Rent)
log(Land Unavail 0-3mi)	0.0028	-0.0166***
log(Land Unavail 3-10mi)	0.0647***	0.0169***
log(Wetlands 0-3mi)	0.0020	-0.0089
log(Wetlands 3-10mi)	-0.0130	0.0399***
log(Water 0-3mi)	-0.0480***	-0.0520***
log(Water 3-10mi)	0.0811***	0.0811***
med. year built	0.0003**	0.0003***
med. year built neighboring pumas	0.0001**	0.0001***
med. num rooms	-0.2672***	-0.1134***
med. num rooms neighboring pumas	-0.3221***	-0.1898***
frac. SF census	0.5803***	0.2108***
frac. SF census neighboring pumas	-0.2106***	0.1903***
Weather controls	Y	Y
School controls	Y	Y
Amenity controls	Y	Y
Year FEs	Y	Y
Partial F-stat	1000.2	1630.6
n. obs	47033	47033

Note: *p<0.1; **p<0.05; ***p<0.01. This table shows the first stage regression of a pooled instrumental variables regression of indirect utilities on characteristics for the whole US from 2012-2019. Topography characteristics are present both as the amount of that feature within 3 miles of the average tract within a PUMA, and the same feature in a 3-10 mile ring from that tract. These features are the total land unavailability, the amount of water, and the amount of wetlands, all from Lutz and Sand (2022). Other characteristics shown are the median number of rooms, the median year built of housing, and the fraction of single family of a PUMAs housing stock. These three are also included for neighboring PUMAs. Instruments are the characteristics for 3-10 mile rings and the neighboring PUMAs. They are ordered next to their within 3 mile circle for comparison. Partial F-statistics for the instruments are reported below.

Figure B10: Examining The Instrument

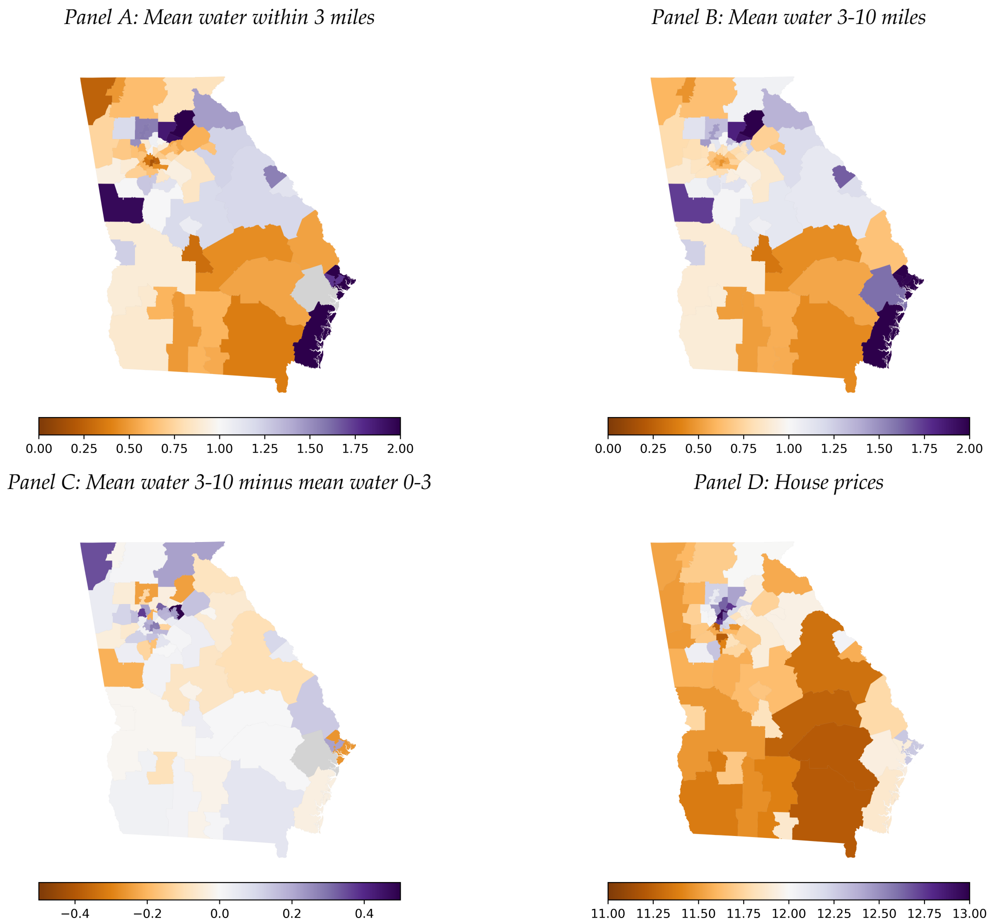

Notes: This figure shows some characteristics of topography and prices for Georgia. Panel A shows the mean amount of water within 3 miles of each census tract in a PUMA. Panel B shows the same but between 3-10 miles away. Panel C shows the average difference between the mean water 3-10 miles away and the mean water within 3 miles of a given tract within a PUMA. Panel D shows the mean of the log house price of a PUMA in 2012.

C CALCULATING COUNTERFACTUAL HOME PRICES

To compute counterfactual market clearing prices I use the Newton's Method algorithm that is used in Koijen and Yogo (2019). The price vector is determined by the market clearing function:

$$\mathbf{p} = f(\mathbf{p}). \tag{23}$$

The price vector is updated based on the slope of the market clearing function:

$$p_{m+1} = p_m + (I - \frac{\delta f(p)}{\delta p})^{-1}(f(p_m) - p_m). \tag{24}$$

Also following Koijen and Yogo (2019), I approximate the Jacobian with its diagonal elements:

$$\frac{\delta f(p_m)}{\delta p_m} \approx \text{diag}\left(\frac{\delta f(p_m)}{\delta p_m}\right). \tag{25}$$

I use a numerical derivative centered around the current price. I iterate until a tolerance level is reached. While I do not prove existence or uniqueness for my setting, demand is downward sloping and supply is upward sloping, and there is no issue with convergence to a stable equilibrium in my simulations.