Comment on "Rental Property Ownership in the United States"

Oren Ziv

CRIW-NBER

March 2026

1. Introduction

A large amount of academic and popular attention has recently fallen on the housing market, where price increases have led to renewed focus on this sixth of the US economy. The consensus explanation is simple: to understand housing prices, one needs to understand housing supply; variation in the elasticity of the housing supply, in turn, can only be understood by understanding natural and policy constraints. Understand the folly of policy-driven housing constraints, and you understand the housing crisis. However, a recent literature, spurred in part by popular consternation with the post-Global Financial Crisis emergence of large corporate landlords, gentrification, and luxury housing, has begun to interrogate how the peculiarities of the structure of the supply side of the housing market might impact prices (Raymond et al. 2016; Cosman and Quintero 2021; Gurun et al. 2023; Barbieri and Dobbels 2025; Coven 2025; Gorback, Qian, and Zhu 2025). To this end, ownership information has become particularly interesting to housing economists.

Researchers wishing to study ownership in the housing sector quickly come up against an unhappy truth: unlike in many industries, ownership information is scarce and misleading in housing. The authors of "Rental Property Ownership in the United States" have stepped in to rectify this in a comprehensive way. Using a large set of previously unavailable administrative

data, the authors have meaningfully advanced the ability of researchers to draw ownership links between properties.

This comment discusses the contribution of "Rental Property Ownership in the United States." Section 2 discusses the paper's main goal in light of the housing affordability crisis. In particular, I draw out the gaps in our understanding of the heterogeneous nature of the housing affordability crisis, which appears to vary greatly by submarket, and make the case that the heterogeneous evolution of housing costs demands a deeper understanding of the industrial organization of the housing sector. In Section 3, I discuss the method the authors use to construct a comprehensive ownership panel. In Section 4, I present a method for evaluating their results. Section 5 discusses how the financial ownership networks they uncover can be used for further research on the nature of housing investments in the US. Section 6 discusses how some of the authors' findings shed light on puzzles in the housing market. Section 7 discusses the findings and potential contribution for our understanding of real estate portfolios. Section 8 concludes.

2. The Housing Crisis Is a Case Study in the Need for an Industrial Organization of Housing

A popular motivation for much contemporary research in urban economics is the growing housing affordability crisis. According to Albert Saiz, "Housing prices are rising faster than incomes in many areas of the world, reducing well-being and engendering social discontent. Passivity by municipal and national governments is no longer an option" (Saiz 2023). There is a broad consensus that the source of the crisis is the housing supply, and in particular, housing supply constraints which raise the cost or prohibit construction.

This flat view of housing requires no great theory of industrial organization of the housing sector. Implicitly, housing costs are a function of supply curves, which in turn reflect natural and

policy constraints. The former set is immutable while the latter set is a policy choice. Under this view, there is no role for the industrial organization of the housing sector to impact housing costs. On the one hand, this may be a valid long run and first-order accounting of housing costs. On the other hand, if we explore the variation in housing costs across American cities since the Great Recession more deeply, a different and more intricate picture emerges.

Figure 1: Housing Costs to Income Ratios by Region

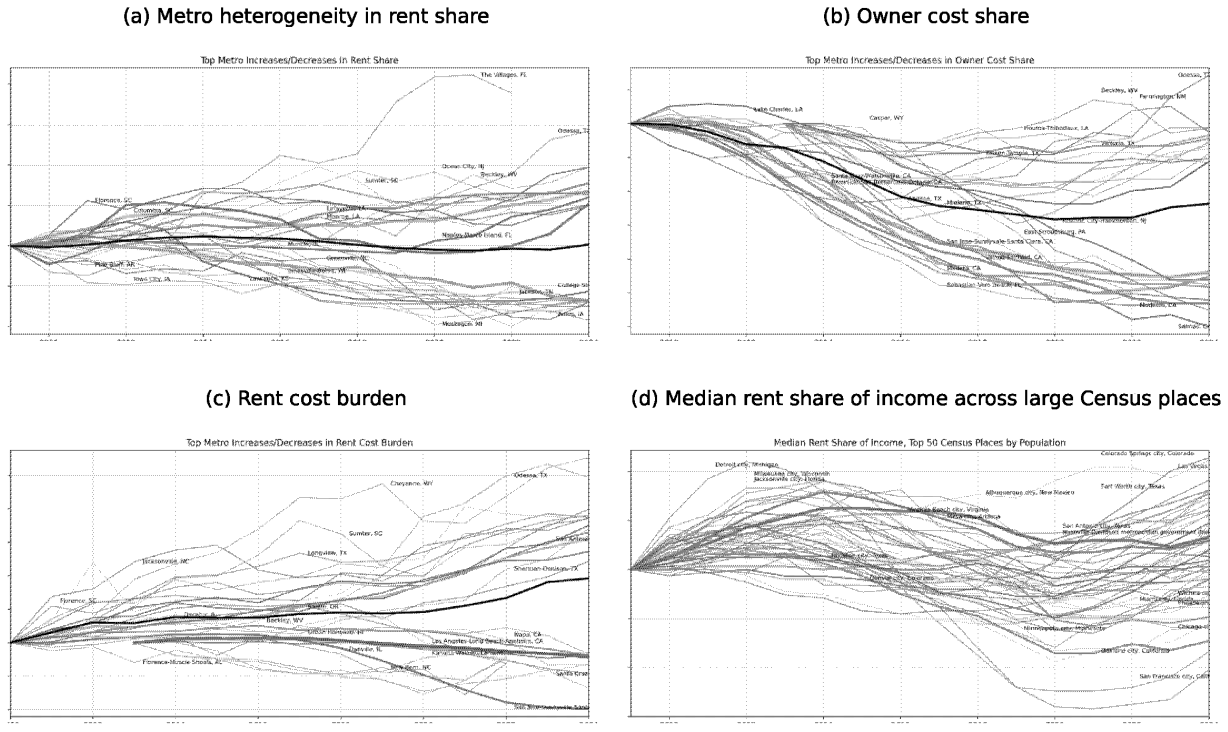

Figure 1 plots the median rent to income ratio for American metropolitan statistical areas from 2009 to 2024, indexed to 2009. The thick black line plots the average across all MSAs, while labeled individual lines plot the top and bottom 15 most extreme MSAs. The first and surprising point to make is that in the average American city, this ratio is extremely stable in this time period. One read of this is that in the aggregate, there simply is no (new) housing affordability crisis.

Indeed, income has risen faster than rents at a large number of MSAs including, among others, Muskegon, MI, Jackson, TN, College Station, TX, and Ames, IA.

But there are also a large number of MSAs for which the affordability crisis is real. In this selected subsample, rent growth has outpaced income growth by a factor of at least 10%. What's more, these cities do not conform to the "simply supply" model of affordability: they are not the larger, older, built-up, and heavily zoning-constrained cities synonymous with unaffordable housing. In fact, they are exclusively "Sunbelt" cities that are synonymous with fast supply growth and relatively lower policy constraints. If there is a housing affordability crisis developing in the last decades, it is not centered in the supply-constrained regions of the country. The notion that colder-older cities experience price increases while the Sunbelt builds might be outdated.

Figure 1b further complicates the simply supply-side story. There, the same graph shows the median of owners' housing cost as a share of income by city and for the cross-MSA average, also indexed to 2009. Incredibly, this share has been dropping since 2009 in the average American city, and increased in only one MSA: Odessa, TX. Why would housing cost to income ratios flatline for renters but drop for owners? In the simplest supply-side story, all housing is policy-constrained, and both renter and owner markets suffer. The clearly diverging patterns in the two figures instead introduce the need to segment these markets. They also suggest that explanations based only on broad housing demand, construction costs, or regional amenities may miss what is specific to the rental market. Rental housing is an intermediated market: renters consume housing services, but landlords own the asset, choose maintenance, set rents, screen tenants, and decide whether to enter or exit. That creates a natural role for industrial organization.

In Figure 1c, I plot the median rent to income ratios by MSA for individuals making under $50,000 a year. Here, a very different picture, one more consistent with the "national housing

crisis" narrative emerges. In the average city, for this sub-population, median share of income devoted to rent is increasing by almost 20%. In outlier MSAs, this number surges past 30%. These MSAs are still in the Sunbelt. So while it appears from median numbers that the narrative in media and academic literature is unsubstantiated by the data, these medians belie clearly different trends in specific submarkets. In these markets catering to below-median income renters, the housing crisis narrative appears to be consistent with the data. But once again, a case must be made that MSA housing markets are segmented to the extent that market conditions vary dramatically for housing catering to median vs below-median income groups, and that the supply structures of these submarkets are fundamentally different.

Finally, Figure 1d plots the median rent share for the top 1% largest places. These are the dense, urban cores or large MSAs. These are the regions where supply is most inelastic and, according to the simple supply-side story, the housing crisis should be most severe. Here again, that story seems to fail us. While Las Vegas, Fort Worth, Albuquerque and other major cities experience 10% increases, others, remarkably including Oakland and San Francisco, experience (mostly modest) declines.

Of course, these figures come with caveats: these are selected samples where market forces are endogenously changing selection over time. In any decomposition there will be outliers. But these figures also illustrate a clear point: if we want to understand the housing crisis, we are unlikely to be able to even trace out its anatomy without accounting for a story capable of capturing the rich heterogeneity. A pure supply-side story does not appear fit for task. To explain the housing crisis, we need to begin to understand why housing prices have skyrocketed in markets where supply has traditionally been viewed as highly elastic, and in submarkets like multifamily rentals, and lower-

income rentals in particular. To do this, we need to build a deeper understanding of the industrial organization of the housing market.

3. To Build an IO of Housing, We Need a Comprehensive Rental Housing Ownership Dataset

This paper delivers a primary element in this larger pursuit: a method for generating data on ownership and ownership dynamics in local housing markets. Ownership structure is a necessary ingredient for understanding firm behavior. Without ownership structure, claims on the nature of firm behavior are therefore difficult to substantiate. Are markups high or mismeasured? Can we reject collusive behavior, or do we just not know how to properly identify colluding and competing products in this market? But ownership structure is exceptionally hard to observe in this market. The solution presented here is therefore critical to a great deal of new work in the industrial organization of housing.

3.1 The Problem with Ownership

Renters reside at addresses and properties can be identified by addresses. Furthermore, tax records with property owners are public record. A large number of datasets aggregate public assessment data and provide the aggregated product to researchers. This can deceptively make ownership information appear to be easy to observe in this market.

The issue at the heart of this paper is that true economic ownership entails an agent capable of optimizing and exercising power over a specified set of variables, in this case price and quantity decisions. The trouble with property ownership data is that the listed owners are not this. In particular, listed owners are often shell companies, LLC's and other paper entities that obscure

another level of ownership: actual decision makers exercising true economic power over a potentially large set of properties. Figure 2 makes these distinctions clear.

Worse yet, the paper entities listed as owners in existing data are often there to deliberately obscure the relationship between properties. In NYC, for example, a building at 53 East 4th street might be owned by the "53 East 4th st LLC." The true owner may own many such buildings nearby and have a large share of the market. But that share would never be observed if each of the firms' properties have their own LLC.

But the error can also go the other way. Many legal entities use the same law firm, registered agent, or mailing address. If we treat that common address as common ownership, we create links that are not really there. So the naive property-record approach can make ownership look too fragmented or too concentrated. We can fail in both directions.

The key data hurdle tackled by this paper is to link ownership entities to actual economic agents tasked with making supply-side market decisions. Until now, very few papers have been able to do this, and only were able to do so in limited markets where this information is available, usually for legal reasons. For example Watson and Ziv (2026, 2025) and Harwood, Ellen, and O'Regan (2025) do this in NYC, where, by law, individual owners with over 20% shares in properties must be listed in publicly available master files. These studies are inherently limited by the laws making this data available. This project offers a method to do this at scale. To do so, the authors must bring substantial new data to bear.

3.2 How to Solve the Obfuscation Problem: Data, Data, and more Data

To link renters to properties, the project currently (and will continue to) uses ACS and the "master address file." This on its own is a large but not uncommon linkage to make. To map properties to owners, the first step, getting us to "ownership entities" uses tax assessor data (e.g. Black Knight).

This is also a large but not uncommon task, and most papers that look into property ownership stop here, inheriting the measurement issues discussed above.

The most difficult step in this link is the last, to move from ownership entities to individual owners. To do this, the authors combine information from the IRS Business Registry, the Longitudinal Business Database, and most crucially, "K1 forms" filed with the IRS.

Figure 2

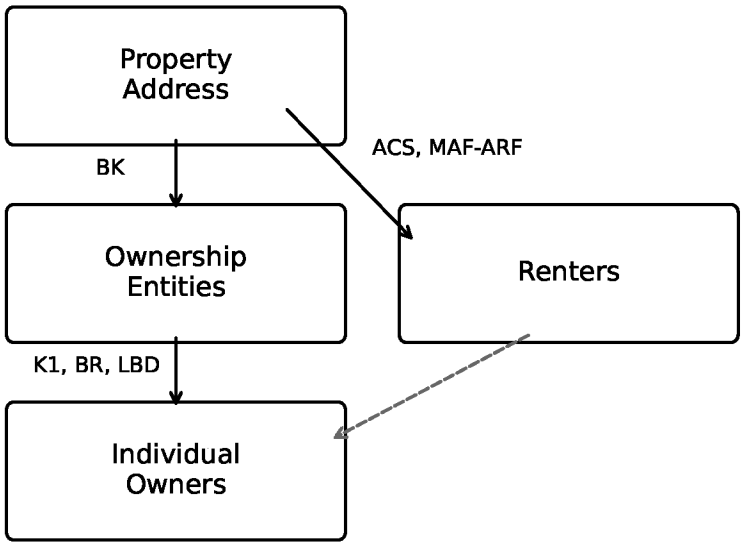

As already demonstrated by the draft at time of writing, this matching process is monumental, and requires a multitude of steps simply to make it computationally feasible. Embedded in each of these steps are decisions which themselves would take paragraphs to explain. The ideal paper gives maximal details on these decisions and their apparent tradeoffs so that future researchers can have as clear as possible an accounting for how this linkage was formed.

3.3 There Is No Conservative String Match

The string matching problem is daunting. There are many potential matches, many judgment calls, and a lot of room for small differences in names and addresses to matter. For example, the authors

choose a cutoff Jaccard similarity score of 0.80 to declare a match. That seems sensible. A hyphen in the wrong place or an extra space can move a match across that boundary. Of course, there is no exact or scientific threshold, nor can there be one: moving the threshold up risks losing true matches, while moving it down risks artificially generating matches. In that sense, there is no unambiguously conservative way to do this, and no way to choose which error, type-1 or type-2 is better to minimize.

When in Watson and Ziv (2025) we undertook a similar, albeit much smaller task of estimating ownership links in New York City, these problems showed up immediately. For example, a significant number of potential overlaps occurred through joint ownerships of individuals with the same last name. But that leaves a larger question open: if family ownership is a relevant margin in a market, how do you deal with people with the same last name? In our New York City data, a common issue we came up against during our string matching process was the existence of common ethnic surnames in ethnic enclaves. Compounding the standard issue in this matching environment of whether two strings are similar is a harder to answer question specific to this undertaking: do two strings refer to the same economic agent, be it a person, family, or ownership network?

On this margin and others, if we miss true links, ownership looks too diffuse. If we create false links, ownership will appear more concentrated than it truly is. Both are problems. We do not have a general prior that one is always worse than the other. With no appropriate null hypothesis to this task, there is therefore no conservative way to make this decision. Nor do we have a principled way of trading off these two errors. The ideal paper will simply tell us what it is doing and, as much as possible, how wrong the resulting ownership network could be.

4. Assessing the Match: A Direct Decomposition

At the time of writing, the biggest missing element of the paper was a direct comparison that helps the reader understand the scope of the mismeasurement in naive assessment-based ownership studies. The authors have indicated that they plan to add this type of exercise to the final draft.

The question we need to answer is: if we did this the naive way, how wrong would we be? Or more generally, how different is one ownership-linking procedure from another? Comparing the properties or ownership types that are matched and unmatched is useful but does not directly answer this question. The key to validation is instead to conceive of the basic unit of observation not as a property but as a (potential) property link. In this way, the pairwise links can be systematically classified and decomposed. This is analogous to work decomposing industry agglomerations Bartelme and Ziv (2023).

In Figure 3, we imagine a set of four properties. The edges between them are potential common-ownership links. The gray links denote all possible pairs. The naive method, in dashed black, gives one network. The authors' preferred method gives another network, in solid black. We can then classify every possible edge into agreements and disagreements.

Figure 3

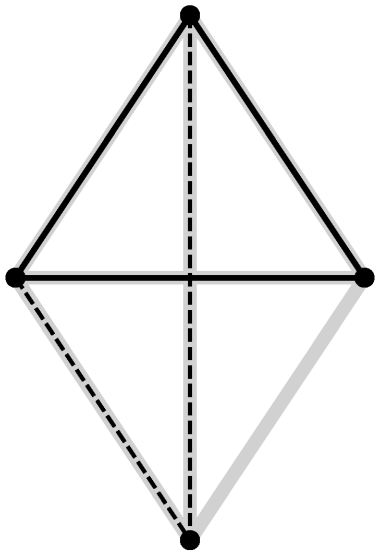

We can characterize the set of all possible links through an accounting decomposition. Let

$$\mathcal{P} = \{\{f, g\}: 1 \leq f < g \leq N\}, \qquad M = |\mathcal{P}| = \frac{N(N-1)}{2}.$$

For each method i, let $P_i \in \{0,1\}^M$ be the vector whose entries index unordered property pairs in $\mathcal{P}$. The m-th entry equals one if method i links the corresponding pair and zero otherwise. Let $\mathbf{1}$ denote an $M \times 1$ vector of ones. Then

$$M = P_i' P_j + (\mathbf{1} - P_i)' P_j + P_i' (\mathbf{1} - P_j) + (\mathbf{1} - P_i)' (\mathbf{1} - P_j).$$

The first term is agreement on a link. The second and third terms are disagreements. The fourth is agreement on a non-link.

Equivalently, let $A_i \in \{0,1\}^{N \times N}$ be the symmetric adjacency matrix implied by method i, with zeros on the diagonal. Let $\mathbf{1}$ denote the $N \times N$ matrix of ones and let I_N denote the identity matrix. Then

$$\frac{N(N-1)}{2} = \frac{1}{2}\left[\langle A_i, A_j\rangle + \langle \mathbf{1} - I_N - A_i, A_j\rangle + \langle A_i, \mathbf{1} - I_N - A_j\rangle + \langle \mathbf{1} - I_N - A_i, \mathbf{1} - I_N - A_j\rangle\right],$$

where

$$\langle B, C\rangle = \sum_{f=1}^{N}\sum_{g=1}^{N} B_{fg}\, C_{fg}.$$

This is the same broad logic as the decomposition in Bartelme and Ziv (2023) into within and across-firm links.

Now, we can further entertain a few intuitive objects as a set of ratios. Let method i denote the benchmark ownership truth and method j denote the method being evaluated. Define

$$TP_{ij} = P_i{}'P_j.$$

$$FN_{ij} = P_i{}'\left(\mathbf{1} - P_j\right).$$

$$FP_{ij} = (\mathbf{1} - P_i)'P_j.$$

$$TN_{ij} = (\mathbf{1} - P_i)'\left(\mathbf{1} - P_j\right).$$

Then

$$\text{FPR}_{j|i} = \frac{FP_{ij}}{FP_{ij} + TN_{ij}} = \frac{(\mathbf{1} - P_i)'P_j}{(\mathbf{1} - P_i)'\mathbf{1}},$$

and

$$\text{FNR}_{j|i} = \frac{FN_{ij}}{FN_{ij} + TP_{ij}} = \frac{P_i{}'\left(\mathbf{1} - P_j\right)}{P_i{}'\mathbf{1}}.$$

Two additional useful statistics are precision and recall:

$$\text{Precision}_{j|i} = \frac{TP_{ij}}{TP_{ij} + FP_{ij}} = \frac{P_i{}'P_j}{\mathbf{1}'P_j},$$

and

$$\text{Recall}_{j|i} = \frac{TP_{ij}}{TP_{ij} + FN_{ij}} = \frac{P_i{}'P_j}{\mathbf{1}'P_i} = 1 - \text{FNR}_{j|i}.$$

There are three ways the authors can implement this decomposition. In a baseline case, this decomposition can be implemented in the authors' work primarily by using their proposed algorithm as the "truth" and the naive assessment-based record as the method to be evaluated (j) to get a quantification of the improvement of the method over the naive approach.

Second, this can be used to test the robustness of the authors' approach to decisions like Jaccard similarity score thresholds and other choices, and finally, using their own method as j and previously published work such as in Watson and Ziv (2026, 2025; Harwood, Ellen, and O'Regan 2025) as the baseline, to implement a validity check on their work.

Third, this decomposition would also be useful by subgroup. I would want to know whether the false-positive and false-negative rates differ by owner type, legal form, neighborhood income, building type, and city. For example, if low-income neighborhoods have more LLC ownership or more complex ownership chains, then matching error could be correlated with the paper's main variables. A national match rate would miss that.

5. Financial Ownership Networks

Once the paper links the individuals behind the entities, it can recover overlapping ownership. This is one of the most interesting parts of the project, and should be foregrounded. These networks are interesting even before we get to pricing.

As an example, one property might be owned in tandem between O1 and O2, while another may be owned by O2 and O3. Maybe this is a family structure. Maybe it is a tax structure. Maybe it is something else. I do not think we should assume we already know why the network looks that way. The fact that it looks that way is itself informative about the nature of real estate investment.

Figure 4

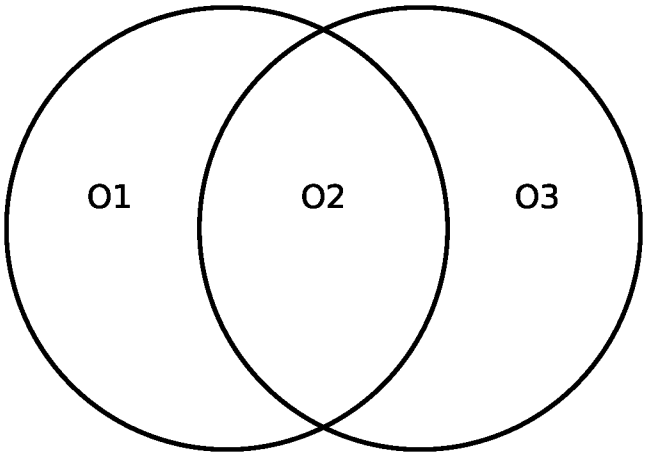

The first descriptive object I would want is the histogram of final ownership network sizes. How many properties are owned by someone who owns no other properties? How many are owned by someone with one other property? This is more informative than starting with the legal ownership breakdown. While the legal form of ownership is useful, for many economic questions, the first-order object is the size and shape of the ownership network.

Now, the authors make a choice to cut FONs when ownership shares are less than 20%. There's a clear logic from the perspective of pricing for this: cross-ownership's impact on pricing decisions attenuates with the dilution of shares.

Why does cross-ownership matter for pricing in the first place? If an owner owns two rental units, and raises the rent p_i at one, that can affect the profitability of the other. This happens through cross-partial elasticities. The profit expression in Equation [eq:prof] captures that intuition:

$$\frac{\partial \Pi_i}{\partial p_i} = \frac{\partial \pi_i}{\partial p_i} + \sigma \sum_{j \in \mathcal{J}_i,\, j \neq i} \frac{\partial \pi_j}{\partial p_i}.$$

where π_i is the profit from rental i, $\mathcal{J}_i$ is the linked portfolio relevant for rental i, and the parameter σ is a stand-in for how much the owner internalizes profits from other rentals. A natural interpretation is common ownership share. As the overlap goes to zero, the properties should behave more independently. As the overlap rises, the properties should be treated more like a single strategic portfolio.

That logic is a good baseline. But it should not be the only way the data speak. For pricing, common ownership and cross-partials are the right starting point. For finance, portfolio exposure may be the right object. For tax or legal questions, the legal entity may be the right object. For management and maintenance, shared management may matter even without full common ownership. The ideal paper will separate these objects as much as possible: legal ownership, economic ownership, strategic ownership, and portfolio exposure.

6. What Do We Learn About Housing Markets?

The ownership heterogeneity in the paper may tie a lot of loose ends together. We already know that large or corporate owners behave differently from smaller owners. We also know that owners in poorer neighborhoods appear to behave differently from owners in richer neighborhoods. In Watson and Ziv (2025), we find that pass-through rates are much higher in poorer areas, and that calculated markups look much higher there as well. That fact dominated the other ways we cut the data.

This paper gives us a possible explanation that ties these facts together. Poorer neighborhoods may not just have different tenants or different demand conditions, they may also have differentially selected owners, who may be larger, more professionalized, more sophisticated, more financially constrained, or simply operating with different strategies. This hypothesis merits further research to verify.

In this context, it is important to be careful about geography, and in particular levels of geography chosen for analysis. In particular, Census tracts are useful, but they are likely too small in an urban context to constitute full markets. While any geography might combine submarkets that renters view as quite different, most likely Census tracts will split a neighborhood that renters view as one market. In particular, these decisions hinge on how local real estate markets really are. A landlord may look large inside a tract but not inside the relevant neighborhood. Conversely, common ownership across adjacent tracts may be missed if the paper studies each tract separately. For questions regarding the competitive environment of a rental market, the market definition should be closer to tenant substitution patterns.

In light of this, the authors should use tract-level facts descriptively, but they should also not stop there. Neighborhood definitions should be above tract where possible. Revealed-preference neighborhood definitions, such as those in Mast and Barca (2025), seem especially useful here. While there is not one perfect geography, the ownership facts the authors present should ideally survive when markets are defined in a way that makes *a priori* the most economic sense.

7. What Do We Learn About Housing Portfolios?

An entirely different contribution made by this paper is its insights into real estate as an investment and as part of a portfolio of assets. Just as we do not know enough about ownership to understand the structure of rental markets, we also do not know enough about ownership to understand how real estate operates in household portfolios.

One fact the authors recover is that around one in ten tax filers own at least one rental unit, or part of a rental unit, as a financial asset. That is a strikingly large number. It means that small-scale real estate ownership is not a niche phenomenon. This in itself creates a range of additional

questions on the nature of real estate investment, especially in smaller portfolios. We still know relatively little about how these assets fit into household balance sheets.

The limitation is that the paper observes only part of the portfolio. We see real estate ownership much better than other assets. But with 1040s and related administrative data, there may be a way to say more about the rest of the household's financial position. Who has W-2 income? Who appears to live primarily on rental income? What share of income comes from rents? What share of asset returns comes from real estate? A household with one inherited rental unit and mostly W-2 income is different from a specialized investor who owns through several entities. A landlord whose wealth is concentrated in one neighborhood is different from one with diversified assets, and landlords who depend on rental income are different from those for whom rental units are simply financial assets.

Exploring this fully will require additional space and effort, and represents a significant branching of the authors' project. Still, its importance to an entirely different literature highlights the foundational nature of this project.

8. Conclusion

Understanding housing markets requires understanding the industrial organization of the housing sector. This paper takes a major step in that direction by recovering an object housing economists have long wanted but rarely had: a map of rental property ownership in the United States. My main comments are meant to sharpen the measurement and make the economics clearer. First, show the linkage from raw records to final ownership networks as transparently as possible. Second, evaluate the method and quantify the total improvement using the proposed decomposition. Third, put financial ownership networks at the center, especially the histogram of network sizes. Fourth, connect the ownership facts to market segmentation, especially the fact that poorer neighborhoods

appear to have different owner types and different pricing behavior. Fifth, make clear that the portfolio facts are a second major contribution. The paper is already a major contribution, and the next iterations should be valuable for both housing economics and household finance.

Bibliography

Barbieri, Felipe, and Gregory Dobbels. 2025. “Market Power and the Welfare Effects of Institutional Landlords.”

Bartelme, Dominick, and Oren Ziv. 2023. “JUE Insight: Firms and Industry Agglomeration.” *Journal of Urban Economics* 133:103372.

Cosman, Jacob, and Luis Quintero. 2021. “Fewer Players, Fewer Homes: Concentration and the New Dynamics of Housing Supply.”

Coven, Joshua. 2025. “The Impact of Institutional Investors on Homeownership and Neighborhood Access.”

Gorback, Caitlin S., Franklin Qian, and Zipei Zhu. 2025. “Impact of Institutional Owners on Housing Markets.”

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

Harwood, Katharine W. H., Ingrid Gould Ellen, and Katherine O’Regan. 2025. “The Rise of Corporate Landlords: An Examination of Behavioral Differences in the Multifamily Market.” *Real Estate Economics*.

Mast, Evan, and Alaina Barca. 2025. “Off the Beaten Tract: Constructing a New Neighborhood Geography Using Revealed Preference.”

Raymond, Elora Lee, Richard Duckworth, Ben Miller, Michael Lucas, and Shiraj Pokharel. 2016. “Corporate Landlords, Institutional Investors, and Displacement: Eviction Rates in Single-Family Rentals.”

Saiz, Albert. 2023. “The Global Housing Affordability Crisis: Policy Options and Strategies.” IZA Policy Paper 203. IZA Institute of Labor Economics.

Watson, C. Luke, and Oren Ziv. 2025. “A Test for Pricing Power in Urban Housing Markets.” *Review of Economics and Statistics*, 1–33.

———. 2026. “Is the Rent Too High? Land Ownership and Monopoly Power.” *Real Estate Economics, Forthcoming*.