Three Essays on Rental Housing Markets in the United States

by

Benjamin J Preis

Bachelor of Science in Physics and Peace & Justice Studies, Tufts University (2015)
Master in City Planning, Massachusetts Institute of Technology (2019)

Submitted to the Department of Urban Studies and Planning
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Urban Science and Policy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2023

Authored by: Benjamin J Preis
Department of Urban Studies & Planning
August 4, 2023

Certified by: Justin Steil
Associate Professor of Law and Urban Planning
Thesis Supervisor

Accepted by: Jinhua Zhao
Edward and Joyce Linde Associate Professor of City and Transportation Planning
Co-Chair, PhD Committee
Department of Ubran Studies and Planning

Three Essays on Rental Housing Markets in the United States

by
Benjamin J Preis
Submitted to the Department of Urban Studies and Planning
on August 4, 2023
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Urban Science and Policy

Abstract

What do rental markets look like in the United States today? Over the course of three papers, I investigate rental housing markets in ten of America's most populous cities: Boston, MA; Columbus, OH; Dallas, TX; Kansas City, MO; Minneapolis, MN; Nashville, TN; Omaha, NE; Philadelphia, PA; Seattle, WA; and Washington, DC. I compare the locations of landlords and rentals, examine the extent to which rental markets have concentrated ownership, and consider the difficult identification problem facing researchers who try to use administrative data to identify rental properties in the United States.

In the first paper, I geolocate rental properties and landlords in eight cities. I find that the median landlord has a mailing address within 10 miles of their rental property, and that a majority of landlords with a residential mailing address are located within the same region as their rental properties. Landlords with residential mailing addresses are located in neighborhoods that are whiter, richer, and have more college graduates than the neighborhoods in which they own properties. I also find that many landlords are located far away from their rental properties, in superstar cities and throughout the country. I use a network-science approach to identify the core locations of landlords, which I call the "landlord market area."

In the second paper, I identify landlords who have significant market shares in a given city or neighborhood. I use machine learning to deduplicate different landlord records, a required step in the age of corporate landlords. I find that many neighborhoods have moderate and high levels of ownership concentration. Higher levels of concentration are correlated with higher rent levels among the cities I study. I use an instrumental variable approach to investigate the interaction between city-wide wage increases and ownership concentration, finding that neighborhoods with higher levels of concentration see larger rent increases.

In the third paper, I investigate the most common methods to identify rental properties in the United States. Most studies heretofore have relied on tax assessment databases to identify rental properties, yet I find that the most common approaches to do so are overinclusive of some types of units, and underinclusive of others. I compare these methods to the rental properties identified by rental registries and to American Community Survey estimates. I identify best practices with regard to rental registries, and interrogate when different approaches are more suitable than others.

Dissertation Supervisor: Associate Professor Justin Steil

Title: Associate Professor of Law and Urban Planning

Members: Ingrid Gould Ellen (New York University), Clio Andris (Georgia Institute of Technology)

Acknowledgements

I received funding from the Sagalyn-Hack Dissertation Award, the Harold Horowitz Research Fund, and the NSF Graduate Research Fellowship Program (Grant No. DGE-1745302.) Any opinions, findings, conclusions, or recommendations expressed in this material are my own and do not necessarily reflect the views of the National Science Foundation. Property data was provided by Regrid. More information regarding this data source can be found at https://regrid.com/. The results and opinions are those of the author and do not reflect the positions Regrid.

The research in this dissertation is based on data collected from local governments. In order to acquire this data, I had to send numerous right-to-know requests. The team at the BU Technology Law Clinic prepared many of these right-to-know requests, kept me on track, and provided valuable information about what I can do with this data. I would like to thank Christopher Conley, Yifu Dong, Jared Lockhart, Evelyn Pacheco, and Janelle Robins for their help. Thank you to Matthew Zook at the University of Kentucky and Margaret Haltom at MIT who both submitted public records requests on my behalf in states that required a local resident.

There are many people who made this dissertation a reality. I would like to begin by thanking the many faculty and staff at MIT and other universities who helped me get to the finish line. This includes Larry Susskind, who mentored me as a master's student, guided me in my PhD decision making process, and advised me during my general exams. I would like to thank Sarah Williams for advising me as a first year PhD student. At MIT, devin bunten, Eric Huntley, Albert Saiz, and Kathy Thelen, all provided valuable feedback on my dissertation, and made it a better document. Zak Neal at Michigan State University graciously agreed to participate on my general exam committee. Lastly, I want to thank Sandy Wellford, Sandra Elliott, and Ellen Rushman for helping me navigate the rules and regulations of MIT. I would not have finished the PhD without all of these generous and kind individuals.

I would also like to thank some of my classmates who made these six years of graduate school bearable, and even enjoyable. Enjoli Hall and Nick Allen both provided friendship and feedback throughout the PhD process. I am so glad to have met MacKenzie Scott and Brian Highsmith from outside of DUSP, from whom I learned so much and who brought light to the early, dark days of the pandemic. I am also indebted to Madeleine Daepp and Arianna Salazar Miranda for their guidance and recommendations throughout the PhD. My classmates from my Master in City Planning are among my favorite people on earth, and I cannot imagine my time at MIT without them.

Thank you to my committee members, Justin Steil, Ingrid Gould Ellen, and Clio Andris for their guidance, support, advice, feedback, and time over the last 18 months. This dissertation is better because of you, and I am a better researcher because of you. Thank you especially to Ingrid and Clio, who agreed to advise a student they did not know from a university with which they are not affiliated.

I would like to thank two Boston friends, Jonah Harris and Noah Adler, in particular, for their support during the latter part of the PhD. Thank you for hosting me, even when inconvenient, in your apartments, when I was commuting between DC and Boston. I so enjoyed your company, and you made the trips to Boston better. I don't know what I would have done if not for your graciousness in opening your homes to me. Thank you to Yael Wollstein for reviewing and editing parts of this dissertation. Though there are too many to name here, I would also like to thank all of my friends who supported me throughout the PhD. You listened to my problems and asked important and helpful questions, but you also provided a respite from the drudgery of a dissertation. I am so thankful to have you all in my life.

Lastly, thank you to my Godmother, Kass, and her wife, Kathleen, for their unwavering enthusiasm and interest, their support of me, and their love over all these years. I have learned so much from the both of them. Thank you to my mom, Annie, for all that she has always done for me. Her neighborhood tours of Pittsburgh growing up, her encouragement of me, and her belief that I was making the right decision, every step of the way, is what got me to this point. Her support of my pursuits and pride in what I have achieved has propelled me forward. Finally, thank you to my fiancée, Hannah Seigel, for her love and support during these past six years, and especially during these past six months. I am so grateful to you. She kept me fed, grounded, and believed in me even when I could not see a path forward in the research. I do not know how I would have finished without her, and I cannot wait for our next adventure together.

Contents

List of Figures

List of Tables

CHAPTER 1 | INTRODUCTION

§1.1 — Overview

Rental housing is of profound importance to most center cities in the US. Most households in large center cities are renters. Most low income households, and most Black and Latino households are renter households. Additionally, Americans spend more on shelter than on healthcare, transportation, or food, making housing costs uniquely important at the individual and household levels (Bureau of Labor Statistics 2021). According to the most recent American Community Survey data, the median renter household pays approximately 30 percent of its income in rent, while more than one in five renter households are severely cost burdened, paying more than 50 percent of their monthly income in rent. In practice, this means that every month, renter households transfer approximately a third of their earnings, sometimes more than half of their earnings, to their landlord. In return, their landlord provides them with shelter, and is also expected to provide upkeep of the property. For many renters, the alternative — ownership — is not financially feasible, and may not even be a sound investment strategy (Herbert, McCue, and Sanchez-Moyano 2014; Reid 2013). Thus, to study rental markets is to study one facet of wealth and inequality in America. In order to understand local housing dynamics of central cities in the United States, it is necessary to understand the dynamics of the rental market.

The ownership of rental housing in the United States is composed of individuals, corporations large and small, and investment vehicles like real estate investment trusts (REITs). Most rental properties are owned by "sole proprietors" who own the property in their own name.

This is especially true among smaller rental properties. According to the 2021 Rental Housing Finance Survey, 71 percent of single-family rentals are owned by individual investors, and 63 percent of rental buildings with two to four units are owned by individual investors. However, individual investors only own a minority, 48 percent, of rental units (US Census Bureau 2021).

Beyond raw numbers, our understanding about landlords is fairly limited. Ethnographic research provides important insight into landlords and rental markets, particularly in terms of low-income renters' ability to rent and remain in rental housing (Cossyleon, Garboden, and DeLuca 2020; Garboden and Rosen 2018; Gilderbloom 1985; Rosen 2014). Several quantitative housing scholars in sociology and economics now study corporate ownership in the rental market, and its impact on tenants (Ellen, Harwood, and O'Regan 2022; Gomory 2021; Immergluck et al. 2020; Mills, Molloy, and Zarutskie 2019; Raymond et al. 2018). But more fundamental questions elude us, such as how many landlords there are in the United States. Richardson (2018) tries to answer this question using a combination of tax and survey data, but the number is imprecise: between 10 and 12 million, as of 2015. Whether any landlords have significant market concentration is also an as-of-yet unanswered question, stymied by data shortfalls and open epistemological questions.

In three papers, this dissertation explores US rental housing markets. In doing so it engages in — and begins to answer — these fundamental questions. Paper one, *Where the Landlords Are: A Network Approach to Landlord-Rental Locations*, examines the location of landlords relative to their rental properties, with a particular focus on landlords with residential mailing addresses. Paper two, *The Relationship Between Local Rental Market Ownership Concentration and Rent*, looks at levels of ownership concentration, and their relationship to rent levels, and the interaction between concentration, wage growth, and rents. Finally, paper three, *Is This a Rental? Comparing Methods for Identifying Rental Units,* examines the tricky problem of using administrative data to identify whether residential properties in the United States are owner- or renter-occupied, and compares three different methods for identifying residential properties as rentals.

All three papers examine a sample of America's most populous cities. As described in §1.2, the main data source for identifying rental properties, and all subsequent analysis, comes from data collected by local governments through rental registration ordinances. Owing to the differences in data availability and quality, Paper One examines rental markets in: Columbus, OH; Dallas, TX;

Minneapolis, MN; Nashville, TN; Omaha, NE; Philadelphia; PA; Seattle, WA; and Washington; DC. Paper Two examines rental markets in: Boston, MA; Columbus, OH; Dallas, TX; Kansas City, MO; Minneapolis, MN; Seattle, WA; and Washington, DC. Paper Three examines rental markets in Columbus, OH; Minneapolis, MN; Philadelphia, PA; Nashville, TN; and Washington, DC. These overlapping sets of cities are all among the 50 most populous local governments in the United States, and they all have rental registration ordinances from which I was able to acquire data. Yet they are all very different cities: located within differing state contexts, with different powers, limitations, and housing markets. Thus, while these cities do not represent a random sample of large American cities, they are meant to represent a sample of large American cities that may or may not have similarities among their rental housing markets.

§1.1.1 — Landlords' Locations

The first paper takes a network approach to understanding rental housing. It begins by identifying where landlords' mailing addresses are relative to their rental properties. Among rental properties with landlords who have a residential mailing address, a slim majority seem to live in the city where they own properties. In some cities, weighting by units, rather than properties, results in a higher share of in-city landlords, while in others, it results in a lower share of in-city landlords. Across all cities, the majority — between 80 and 90 percent — of rental properties with landlords who have a residential mailing address have a landlord within the same metropolitan statistical area (MSA). However, these numbers become less uniform among all rental properties, regardless of landlords' mailing address type: in Dallas, Nashville, Minneapolis, and Columbus, a minority of landlord mailing addresses are in the same city, while in Philadelphia, Seattle, Omaha, and DC, more than 60 percent of landlord mailing addresses are in the same city. These numbers show the differing patterns of local and out-of-town ownership across large cities in the United States.

Considering the characteristics of rental property and landlord locations, I find, perhaps not surprisingly, that landlords' residential mailing addresses are typically located in neighborhoods that are higher-income, whiter, and have more college graduates than the neighborhoods in which they own property. It seems that most landlords live in neighborhoods that are quite different from where they own properties, even if they live in the same city or MSA. This raises questions about inequality, social networks, and power in the private US rental housing market. Landlords who live in different neighborhoods than the types of neighborhoods in which they own property raises concerns about peer effects and segregated social networks. At the same time, regional inequality may be heightened by transfers of rent from tenants to landlords who live in different neighborhoods and municipalities.

Finally, this paper uses methods from network science to define a "landlord market area," based on dense clusters of landlords' locations throughout the United States. Researchers have typically defined "housing market areas" based on where individuals search for housing in which to live (Jones and Coombes 2013; Royuela and Vargas 2009). In contrast, "landlord market areas," are defined by where landlords are located relative to where they search for property in order to invest capital. I use a network-based approach to argue that one can define a landlord market area by extracting a network "backbone" (Z. P. Neal, Domagalski, and Sagan 2022) of the largest clusters of landlord mailing addresses. Unlike commuters who are mostly found within the same metropolitan statistical area as where they work, landlords are found in large numbers throughout the US. Rental housing is a national (indeed, international) market, and we cannot fully understand how they function by focusing solely on local actors.

§1.1.2 — Landlords' Holdings

The second paper aims to understand the extent to which, in some neighborhoods and some markets, landlords have pricing power due to high levels of ownership concentration of rental units. Over the last thirty years, the composition of landlords at the national level has changed substantially: large and corporate landlords have a substantially greater presence today than ever before. The National Multifamily Council, which has tracked the largest multifamily owners for the past 30 years, reports that the largest fifty multifamily landlords in the US own fifty percent more units in 2022 than were owned by the largest fifty multifamily landlords in 1992 (National Multifamily Housing Council 2022). In 1991, approximately 20 percent of large[1] apartment buildings were owned by individual investors (US Census Bureau 1996); by 2018, that number had shrunk to only 9 percent (US Census Bureau 2021). The increase in institutional ownership of single-family rentals has concerned the general public and academics alike. Anecdotal reporting

[1] More than 50 units

points to a phenomenon wherein some subdivisions and neighborhoods now have one corporate owner (Dezember and Kusisto 2017). While economic theory typically assumes that the housing market exists in a state of perfect competition, recent scholarship has shown how some individual landlords may have highly concentrated holdings (Linger, Singer, and Tatos 2022; Tapp and Peiser 2022; Watson and Ziv 2022), and how increased consolidation in the construction market has also been associated with higher prices (Cosman and Quintero 2021). Does this hold true for landlords and renters in America's largest cities?

I find that, at the city level, America's rental markets do not show signs of concentration. Few landlords own more than one percent of the city-wide rental market, and the average landlord owns very few units. However, the city level may not be the right level of geography at which to look at market concentration. When households search for a new place to live, they often select a small set of neighborhoods to search, biased towards short-distance moves and shaped by racial segregation and differences in neighborhood housing costs (Bruch and Swait 2019; Carrillo et al. 2016; Krysan 2008; Rae and Sener 2016). Thus, one may wish to look at concentration at the neighborhood scale. At smaller geographies, there are significantly higher levels of concentration in the cities I examine. At the ZIP code level, higher levels of concentration are correlated with higher asking rents. Using an instrumental variable approach, I also ask whether landlords in concentrated neighborhoods raise rent more than landlords in less concentrated neighborhoods when there is a city-wide wage increase. I find that the answer is yes: neighborhoods with higher levels of concentration have larger rent increases in cities with wage growth between 2014 and 2022.

§1.1.3 — Identifying Rentals

Both papers one and two use data from rental registries, which are created by local governments specifically for the purpose of identifying rental properties. Yet this is not a common source of data among academic and applied researchers (Coulton et al. 2020; Ellen, Harwood, and O'Regan 2022; Haider 2021; Lynch 2022; Watson and Ziv 2022 are among the few who also use this data source). Instead, most researchers trying to understand the rental market use information provided by tax assessment databases to identify residential rental properties, as compared to owner-occupied or commercial properties (see, e.g., Coulton et al. 2020; Ferrer 2021; Freemark, Noble, and Su 2021;

Gomory 2021; Hangen and O'Brien 2022; Linger, Singer, and Tatos 2022; Raymond et al. 2018; Rose and Harris 2021; Travis 2019). However, tax assessment databases were not designed to identify rental properties, and there are drawbacks to their use. In line with other studies that examine the accuracy of and challenges with local administrative data (Folch, Spielman, and Manduca 2018; Molfino et al. 2017), my third paper compares the accuracy of rental registration ordinances to tax assessment databases, in trying to identify rental properties.

The third paper finds that the two methods for identifying rental properties via tax assessment databases — identifying residential properties that lack a homeowner exemption (the "homeowner exemption" approach) or comparing mailing addresses to property addresses (the "mailing address" approach) — are successful in identifying many of the same properties as rentals. However, neither method is universally applicable. In some cities, the "mailing address" approach seems to work better, while in others, the "homeowner exemption" approach would be preferred. However, neither method captures all rental properties within a city: between 5 and 20 percent of rental properties that are identified via rental registries approach are not identified as rentals via either the mailing address or homeowner exemption approach.

All three approaches present their own challenges: some owner-occupiers mail their property tax bill to a different address, while some landlords mail their property tax bill to their rental property. It appears that some landlords continue to claim a homeowner exemption, even when they are ineligible for it. Additionally, not all cities have homeowner exemptions that apply to all homeowners, and not all eligible homeowners claim the exemption. On the other hand, rental registration ordinances suffer from coverage problems. Relatively few local governments have rental registration ordinances, not all ordinances cover all rental properties, and not all cities enforce their rental registries well; only 10 percent of single-family rentals in Nashville appear to be registered with the local government, compared to 60 percent in Philadelphia.

It is critically important to accurately identify rental properties in the study of rental markets. This is especially true among studies (such as the second paper of this dissertation) that aim to identify market concentration and its effects. This paper calls into question the accuracy of the tax assessment approaches. While rental registries are far from a perfect substitute, they represent an accurate subsample of the rental market: legal (registered) rental properties.

§1.2 — A Note on the Data

This dissertation primarily uses data from rental registration ordinances in order to identify rental properties in the United States. While Chapter 4 discusses, in detail, the peculiarities and challenges of rental registration ordinances, I briefly sketch here the procedure for acquiring rental registration ordinances, and the reason for different cities' inclusion in different papers.

Of America's 50 most populous cities, approximately 30 require some landlords to register with the city. Of these 30 that require landlord registration, I was able to obtain data from 19. However, I only use data from 10 cities across the three papers, and no paper uses the same set of cities. The reasons for this small sample are legion. Of the 30 cities that require landlords to register their properties, five exclude small properties[2] from registering. For two cities, I would have had to have paid $450 to acquire rental registry. Six others denied or failed to respond to a public records request for the data. Of those 19 cities for which I was able to acquire the rental licensure data, not all of it was usable. In three cities, officials would only provide me with the addresses of rentals — they would not provide parcel numbers, landlord names or mailing addresses, or any other data collected. Some cities do not collect unit counts, making the data significantly less useful. Others have such poor coverage that their usage made little sense.

These limitations necessarily restricted the scope of this study to the cities with usable rental registry data that I was able to acquire. While rental registration ordinances are far from widespread, they exist in a variety of local government contexts. For instance, beyond the most populous cities, many college towns in the US have rental registries; a cursory search shows that Ithaca, NY (home to Cornell University); State College, PA (home to Pennsylvania State University); and East Lansing, MI (home to Michigan State University) all require landlords to register rental properties. Smaller US cities, such as Buffalo, NY; Scranton, PA; and Portland, ME require landlords to register rental properties with the city. These are an underutilized resource for academic researchers. Additionally, their widespread implementation leaves significant heterogeneity to understand their differing impact, effectiveness, and enforcement (see Haider 2021

[2] Small is variously defined as properties with fewer than 5, 4, or 3 units.

on differences between major cities with regards to rental code enforcement; and Samuel, Schwartz, and Tan 2021 on the impact of rental licensure).

§1.3 — Research Questions and Aims

Each paper in this dissertation aims to answer a set of distinct research questions, though they are all related to rental housing in the United States. The questions for the three papers are as follows:

1. Where are landlords' mailing addresses relative to their rental properties: in the same city, same MSA, same state, or far away? How do the neighborhoods of landlords differ from the neighborhoods of rental properties? Can we define a landlord market area by examining dense clusters in the location of landlords and renters? If so, how does that definition of a landlord market map onto existing definitions, or does it demonstrate a distinct landlord geography?
2. How many owners control what share of rental units in a given housing market? Are higher levels of ownership concentration at the neighborhood scale associated with higher levels of asking rents? Among cities that experience a wage shock, do rents rise more in neighborhoods with higher levels of concentration than in neighborhoods with lower levels of concentration?
3. What are the different ways to identify rental units at the local level in the United States? Do these different methods identify the same properties as rental units? What are the challenges in identifying rental units from local administrative data, and are there some types of units or neighborhoods that are systematically under-counted by certain methods?

These questions were developed in order to broaden our understanding of the rental market in the United States today. Scholars from a range of disciplines have answered related questions on the experience of tenants in rental housing, the practices of different types of landlords, and the levels of concentration in some American housing markets. Yet there is still much more to learn. The quality of data is a severe limitation in all studies, and the in this dissertation suffers from that shortcoming as well, leading to the need for the third dissertation paper. These studies are all linked to important policy-related challenges in the American context today: the provision of rental housing, its cost, and appropriate safeguards (or lack thereof) for renters. Only by accurately identifying rental properties, accurately identifying landlords, quantifying landlords' holdings and identifying landlords' locations can we begin to accurately understand the contours of the US rental housing market. By understanding these contours, in turn, policymakers can more appropriately intervene in the housing market, be it through regulation, subsidy, or other policy tools, to make rental housing a better, safer, and more equitable experience for one third of the US population.

CHAPTER 2 | WHERE THE LANDLORDS ARE: A NETWORK APPROACH TO LANDLORD-RENTAL LOCATIONS

Abstract

The US is home to more than 100 million renters, and approximately 11 million landlords, yet these two sides to the rental market are rarely studied in tandem. This study uses a multiscalar network-based approach to identify landlord market areas. Building on administrative data of rental properties' and landlords' locations, I define landlord-property networks as a spatial bipartite network, where landlord addresses are connected to their properties' addresses, and vice versa. I first examine the location of landlords with residential mailing addresses. I then compare the differences in socioeconomic characteristics in landlord and rental tracts. I simplify this network by extracting its backbone, defining a core component of a landlord market. I compare these networks to Metropolitan Statistical Areas and commuting networks, in order to evaluate the performance of the backbone extraction method. I find that most landlords are local, and, perhaps unsurprisingly, that landlord neighborhoods are richer, whiter, and more expensive than where their properties are located. Extracting the backbone of the commuting network results in a network that mirrors a regional definition, while the landlord market area is much more national in scope. These two networks differ geographically, and also with regards to their network statistics. While renters and homeowners search within a region for new housing, landlords and capital can search nationally for locations in which to invest. This paper provides a new, robust foundation to understand rental market investor dynamics and the relationship between owner, renter, and property.

§2.1 — Introduction

Rental housing has become a flashpoint in American politics. On the left, there are calls to "abolish landlords" (Medwell 2021). States and localities nationwide have enacted new laws that enhance eviction protection, rent control, and right-to-counsel for tenants (see Appendix A of Collinson et al. 2022). In part, these tensions have been brought forth over changes to the rental market: the increased proportion of rental housing that is owned by corporations (Lee 2017), the financialization of single-family rentals (Fields 2018), and the new geography of single-family rentals (Charles 2020a).

Yet our understanding of rental housing ownership remains rudimentary. It is hard to pin down how many landlords there are in America — estimates range from 10 to 12 million (Richardson 2018). Recent qualitative work examines the relationship between landlords and their tenants, including the impact of professional management (Garboden and Rosen 2022; Shiffer–Sebba 2020). Quantitative housing scholars have also examined the impact of corporate landlords (Ellen, Harwood, and O'Regan 2022; Travis 2019), absentee landlords (Rose and Harris 2021), and landlords who own in low-income neighborhoods (Desmond and Wilmers 2019; Shelton 2018). Tax and social registry data abroad paints a more detailed picture of landlords as high-income and high-wealth individuals (Hochstenbach 2022; Statistics Canada 2022).

This paper aims to contribute to the burgeoning literature on landlords by considering where landlords are located relative to the properties that they own — the geography of landlords. It systematically collects rental registry data from eight large cities in the United States: Philadelphia, PA; Minneapolis, MN; Omaha, NE; Seattle, WA; Dallas, TX; Nashville, TN; Columbus, OH; and Washington, DC. Landlord and rental property addresses are geocoded, and analyzed as a network, where *landlords* are connected to their *properties*, both of which exist in geographic space. It examines the fraction of rental properties that have landlords with mailing addresses in the same city, metropolitan area, or state. It calculates the socioeconomic differences between rental neighborhoods and landlord neighborhoods. It uses network-science approaches to identify the core component of the rental property network, which is termed the landlord market area. It compares this network to the core commuting network and the Metropolitan Statistical Area (MSA) boundaries.

Why care about landlords' locations with respect to their properties? First, rich landlords investing in poor neighborhoods raises concerns about exploitation and inequality. Shelton (2018) argues that areas of racially and ethnically concentrated poverty are *directly linked* to areas of concentrated affluence via landlords' property ownership ties in Louisville, KY. Similarly, Hochstenbach (2023) finds that, in the Dutch case, landlords are more likely to live in less-dense, higher-income, and higher-value neighborhoods than the neighborhoods in which they own property. Harvey (2009, 100) argues that "our existing technology should be used to extend our understanding of interregional income transfers … insofar as these have actual or potential consequences for the distribution of income in society." Insofar as most rental properties remain in the hands of private individuals (US Census Bureau 2021), and the median renter is cost burdened, the monthly payment of rent represents an income transfer from renters to landlords; this paper demonstrates the intra- and inter-regional flows of that rent.

Second, spatial and economic inequalities can compound, meaning that it may be preferable that rich landlords live in the neighborhoods where they own property. This could be true for the purposes of peer effects and socioeconomic integration (Chetty et al. 2022; Small 2010). If landlords were to live in the same neighborhood — or similar neighborhoods — as where they own property, we might expect less "landlord paternalism" (Rosen and Garboden 2022). From a property maintenance point of view, nearby landlords may be less likely to "milk" properties (Mallach 2010), because rental properties and their tenants would directly impact a resident landlord, while absentee landlords can be shielded from such negative outcomes (Rose and Harris 2021). Additionally, out-of-town landlords may be attractive targets for revenue raising (Vigdor 2004), while resource hoarding and municipal fragmentation in the United States (Freemark, Steil, and Thelen 2020) means that landlords and renters living in different municipalities may contribute to regional inequalities. More activity from small- and medium-sized real estate investors have also been found to result in increased home prices and rents (Garriga, Gete, and Tsouderou 2023). Finally, technological changes means that landlords may no longer need to be local to acquire and manage properties, but can instead find, acquire, and manage properties from afar (Fields 2022). Taken together, these concerns about spatial and economic inequality compel us to understand where landlords operate and where they own property.

§2.2 — The Relationship Between Regionalization, Housing Market Areas, and Landlords

This paper contributes to several conversations within the geography and housing literatures. First, I contribute to the literature on the differing types and scales of real estate investment and landlords. Özogul and Tasan-Kok (2020) find that researchers typically categorize real estate investors by: their spatial scale of operations; size and social composition; investment objective; or social behavior. Relating to spatial scale, Rose and Harris (2021) find that absentee landlords receive more code violations than owner-occupants; Immergluck and Law (2014) find that most investors in foreclosed properties in Atlanta operated in Georgia; Crook, Ferrari, and Kemp (2012) find that most landlords in Scotland intentionally buy properties near to where they live; and D'Lima and Schultz (2021) find that real estate investors who are local typically gain higher returns compared to non-local investors. Relating to size, scholars have found that larger landlords may be more likely to evict tenants (Immergluck et al. 2020; Raymond et al. 2018), and that deliberate landlords (as compared to "circumstantial" landlords who fall into landlordship) have larger portfolios and different logics relating to rent raising, management, and neighborhood investment (Shiffer–Sebba 2020). I intervene in these conversations by examining the spatial location of landlords, including how the spatial locations differ based on size and corporate structure.

Second, I call the geographical area bounded by landlord and rental locations a "landlord market area." The concept of a landlord market area is derived from a long-established literature on housing market areas generally (Brown and Hincks 2008; Jones 2002; Jones and Coombes 2013; Royuela and Vargas 2009). These housing markets areas are designed to support regional housing development and are based on subnational migration or travel-to-work patterns. While most papers define a housing market area specific to homeowners, Jones and Coombes (2013) differentiate between the migration patterns of homeowners, renters, and social housing tenants. They find that these tenure-specific housing market area definitions differ from each other, in some cases substantially. Housing market areas are meant to capture the geographic areas where most people would search for housing in which to live, given that they already live nearby. Thus, housing market areas subset larger geographies into spatial regions where housing search is generally self-contained; that is, where most people looking for housing only look within the region, and not outside of it.

While labor — and therefore most home-searchers — is often bound by geography, capital is not (Sassen 1990). Where do landlords look when they wish to purchase property? One answer to this question is that they seek places with a rent-gap, such that they may capitalize on asset appreciation (Smith 1979). Recent scholarship has found that institutional-investor-owned suburban single-family rentals are often concentrated in middle-income white or low-income Black and Hispanic neighborhoods (Charles 2020a). Yet the geographies of institutional investors may follow distinct logics from the geographies of mom-and-pop landlords. Researchers have found that mom-and-pop landlords prefer to live near the properties that they own as it provides them with more information and reduces risk (Crook, Ferrari, and Kemp 2012), resulting in higher returns (D'Lima and Schultz 2021). In the Netherlands, Hochstenbach (2023) found that 70% of landlords who own rental properties in Amsterdam live within 20 kilometers of their property. Yet platform technologies have enabled real estate investment to occur outside the bounds of where a landlord lives or works, identifying, acquiring, and managing properties from afar, without ever leaving their house (Parker and Friedman 2022).

I build off of these two strands of the literature to define a "landlord market area" for a given city as the area in which a substantial share of the city's landlords can be found. This definition mirrors the "housing market area" literature by identifying a reasonable subset of space where there is currently landlord activity, as compared to a reasonable subset of space in which most homeowners are looking to buy. However, absent capital controls, the true market area for a given parcel is global, while housing market areas are predominantly small regions. Technology has enabled real estate investors to identify, acquire, and manage properties from afar, at great scale (Fields 2022). Yet there must still be some logic to the location of landlords in a given city. For instance, Sakong (2021) examines the geography of Chinese investment in the US housing market, finding that US housing markets that have a shorter flight time to China have a higher proportion of Chinese owners.

What would we expect the landlord market area to look like? If being a landlord were a job that involved property maintenance, we might expect the landlord market area to resemble a commuting network. Given the work of Crook, Ferrari, and Kemp (2012), D'Lima and Schultz (2021), and Hochstenbach (2023), we might expect most of the landlords to be local. However, if being a landlord is a class position based on capital ownership (Hochstenbach 2022), enabled by property management companies and technology (Fields 2018; Fields 2022), then perhaps we would expect the landlord market area to reflect the global city network (Sassen 2005), as capital from superstar cities chases rent-gaps and returns in other locales.

§2.3 — Methods and Data

§2.3.1 — Data

I use rental registry data collected by eight cities and acquired via right-to-know requests and open data portals. After reviewing which of the 50 most populous US cities had rental registries, I selected these eight cities based on data availability, data quality, and because these eight cities have very few exemptions for registration. While this is certainly a convenience sample, Columbus, Dallas, Minneapolis, Nashville, Omaha, Philadelphia, Seattle, and Washington, DC represent very different parts of the United States and have substantially different housing markets. Some are predominantly owner-occupied (Omaha, Philadelphia, and Nashville); some are in the Sun Belt or Midwest, sites of institutional investor interest (Mills, Molloy, and Zarutskie 2019); some have high proportions of single-family rentals (Philadelphia, Columbus, Nashville), while others have a broader mix including large multifamily buildings. Their housing stock, industrial composition, and regulatory powers also differ.

These data include the location and number of rental units, and names and mailing addresses of landlords. Rental units are either represented as coordinates or parcels, depending on the data source. Landlords are located using the mailing address provided by the rental registry. I use Dedupe, a machine learning software, to match landlord addresses to a complete list of parcels in the United States, provided by Regrid. Unmatched landlord addresses are geocoded using the Google Maps API. Between exact matching, matching with Dedupe, and geocoding with Google, I estimate that I am able to geocode 98% of all addresses with 90% accuracy, far exceeding the commonly accepted threshold of an 85% geocode rate (Briz-Redón, Martinez-Ruiz, and Montes 2020).

Landlord mailing addresses from rental registries present two challenges. First, some large-scale landlords who operate in many cities use local mailing addresses, while others provide mailing addresses for their national headquarters. Second, some landlords provide the mailing address of a management company, rather than their own mailing address. While I am unable to directly address these challenges, Regrid provides information from the United States Postal Service as to whether a mailing address is residential, known as the Residential Delivery Indicator (RDI).[3] Because the majority of rental properties in the United States are owned by individual investors, the vast majority of whom manage their own properties (US Census Bureau 2021),[4] I expect that residential mailing addresses for landlords would be their actual residences, but I cannot conclude that with certainty. Nonetheless, for some research questions, I restrict the universe of rentals to only those with landlords who list a residential mailing address. The vast majority of smaller landlords (those who own five or fewer units) report residential mailing addresses.[5] When I use all landlord mailing addresses, I qualify the findings that I am only able to identify the first-order flows of rent, recognizing that landlords that operate nationally may provide local mailing addresses.

§2.3.2 — Methods

First, I estimate what percentage of rental properties have landlords who live in the same municipality, metropolitan statistical area (MSA), or state as the location of their rental property. Based on owner name information, I differentiate among corporate landlords and individual investors.[6] Based on my work for Chapter 3 in this dissertation, wherein I deduplicate landlords to

[3] Nationally, Regrid only provides RDI information on 68% of parcels. However, among the owner addresses identified in this study, 92% of mailing addresses were identified. Most cities had upwards of 95% of owner addresses identified as either residential or not, with the exceptions of Dallas (77%), and Columbus (84%).

[4] According to the 2021 Rental Housing Finance Survey, a full 78% of rental properties are managed by their owners; 84% of rental properties owned by individual investors are managed by their owners. Even among rental properties owned by LLCs, LPs, or LLPs, 56% are managed by the "property owner or unpaid agent of the owner." Based on the questionnaire language, this seems to be the actual human owner, as compared to a "management agent directly employed by the owner," or a management company.

[5] For instance, in Philadelphia and Washington, DC, between 87 and 89% of small landlords, respectively, provide a mailing address that Regrid identifies as residential.

[6] Specifically, I use Regular Expressions to search for words such as "LLC" "Corp" "Inc," etc.

identify landlords who own multiple properties, I also distinguish between large, medium, and small landlords. I define large landlords as those who own more than 50 rental units, medium landlords as those who own between 6 and 50 units, and small landlords as those who own between 1 and 5 units.

Second, I construct a spatial bipartite network, where landlords are connected to their properties, both of which exist in geographic space. This landlord-property network is a weighted graph, where the weights of the connection are the number of units in a property. From this spatial bipartite network, I assign owners and renters to census tracts, connecting these geographies with weights based on the cumulative number of owners and units. Thus, the main network of interest is a weighted unipartite projection of a bipartite network, where census tracts are connected if there are landlords and properties that connect them, and the weights between nodes are the cumulative number of units.[7]

Among rental properties with a landlord who lists a residential mailing address, I compare the differences between the tract where the landlord's mailing address is to the tract where the rental property is. In order to generate a null distribution against which I compare the empirical results, the average differences in socioeconomic variables between tracts are compared to a simulated random network (Andris et al. 2021). I run 1,000 simulations for each city, where one end of an edge is randomly reconnected to any other node in the network, creating a random network with the same number of nodes and edges.

Third, I construct a reduced network based on the highest density of landlord locations, which I term the "landlord market area." Among the cities I consider, landlord-property networks are quite large and unequal: in Philadelphia, for instance, 49% of edges have a weight of one. To include all edges in such a landlord market area would make this term meaningless. To detect the core structure of the network, I extract the network backbone of the landlord network through the "backbone" package in R (Z. P. Neal 2022). Network backbones can be thought of as the "core" part of the network, and there are many different ways to extract a network backbone. I use the

[7] Not all rental registries provide unit counts. In these cases, I weight by the number of properties.

“disparity” filter, from Serrano, Boguñá, and Vespignani (2009). For every edge in a network, the disparity filter compares the actual edge weight to a null distribution, where each node has equally weighted edges. Edges are retained if its weight is statistically significant, relative to this null model. It thus reduces the number of edges in the model by eliminating edges that are *locally* unimportant. I use the disparity filter, rather than an absolute measure (such as including all edges with a weight greater than 10) to account for city-by-city differences in the distribution of edge weights.[8]

Because the geography of landlords is an as-of-yet unexamined question, there is no clear baseline for what we would expect the landlord market area to look like. Thus, I compare the landlord market area to the commuting network in the same city. I do this both to ensure that the disparity filter produces reasonable results — I expect the commuting backbone to mirror MSA boundaries — and to compare the empirical results to the landlord market area. To create the commuting network, I use American Community Survey (ACS) commuting data, created by a special tabulation via the Census Transportation Planning Products, provided by Dash Nelson and Rae (2016), and extract the backbone via the disparity filter. I make the initial network of commuters those who commute to census tracts in the central cities. I map the resulting commuting and landlord-property backbones, and overlay the local MSA on these maps, showing the stark divergence between MSA definitions and landlord market areas. I compare the basic statistics associated with their network structure.

§2.4 — Results

§2.4.1 — Coverage of the Data and Geocoding Accuracy

Table 2.1 depicts the coverage of the rental registries in the eight cities analyzed. There are significant differences in coverage among these cities, in terms of the number of rental units registered as compared to the number of rental units believed to be in the city per the ACS. Cities such as Philadelphia and Seattle have nearly 100% coverage, indicating that most of their rental

[8] While the edge-weight distributions are similar among cities, there are quite large differences at the lower end. 94% of edges in Philadelphia have a weight 10 or smaller, compared to 82% in Dallas.

properties are registered[9] with the city. Surprisingly, Minneapolis, has more registered properties than the ACS reports. This is possible either because of lagged Census data collection (wherein new rental construction isn't yet reflected in the Census) or because of error on the city side. Likely, it is a combination of both. For a greater discussion about the coverage and accuracy of rental license data, see Chapter 4.

It is worth noting that PO Boxes are a major challenge in this work, as PO Boxes can only be geocoded to a post office or brick-and-mortar shipment business associated with the PO Box ZIP code. The use of PO Boxes ranges significantly among cities, with a low of 0.3% of mailing addresses in Minneapolis, to a high of 16% in Omaha. This likely reflects differences in rental registration ordinances: the Minneapolis rental license application explicitly states that the owner address cannot be a PO Box. It also may indicate differences in landlord types and logics, as we might expect more professionalized landlords to use devices such as incorporation or PO Boxes to separate their personal and professional activities (Shiffer–Sebba 2020).

Reflecting on the type of owner mailing addresses found in the data, we find that the majority of landlord mailing addresses are residential, as indicated by the USPS Residential Delivery Indicator collected by Regrid, yet that breakdown is not equal across cities. Indeed, in Columbus, only 52% of all owner mailing addresses were found to be residential. These numbers are not necessarily a count of sole proprietors, however, since many corporate landlords have residential mailing addresses. Again, these differences may account for differences in the professionalization of the landlord industry in these cities (Shiffer–Sebba 2020).

[9] Or have been registered recently. I retained "Inactive" or "Expired" licenses that expired in or after 2019, to account for frequent lapses and lack of timely renewals.

Table 2.1: Descriptive Statistics

City	City Population, ACS 2017-2021	Number of Licenses	Number of Units	Number of Rental Units, ACS 2017-2021	Number of PO Box Mailing Addresses	Number of Exact License Matches with Regrid	Number of Deduped License Matches with Regrid
Columbus	898,143	57,368	109,007	208,644	6,082	23,283	14,986
Dallas	1,300,239	10,891	251,560	301,691	807	4,339	3,494
Minneapolis	425,091	22,634	103,300	94,741	62	10,296	8,702
Nashville	682,646	13,839[10]	-	129,737	831	648	6,485
Omaha	488,059	16,925	80,495	81,394	2,745	6,511	5,000
Philadelphia	1,596,865	110,983	280,918	307,740	946	61,488	34,382
Seattle	726,054	25,992	149,497	184,866	2,092	13,451	6,845
Washington, DC	683,154	33,966	175,429	181,384	2,122	11,068	12,236

	Licenses Analyzed	Units Analyzed	Number of Unmatched Licenses	Number of Unmatched Units	Number of Residential Mailing Addresses	Number of Non-Residential Mailing Addresses	Number of Unknown Type Mailing Addresses
Columbus	50,444	95,629	893 (1.6%)	1,707 (1.6%)	29,436 (51.3%)	12,108 (21.1%)	8,900 (15.5%)
Dallas	9,998	242,810	2,251 (20.7%)	57,581 (22.9%)	5,816 (53.4%)	1,929 (17.7%)	2,253 (20.7%)
Minneapolis	22,421	100,670	152 (0.7%)	2,398 (2.3%)	18,000 (79.5%)	3,215 (14.2%)	1,206 (5.3%)
Nashville	11,647	-	678 (4.9%)	NA (NA%)	8,909 (64.4%)	2,440 (17.6%)	298 (2.2%)
Omaha	14,103	71,478	77 (0.5%)	557 (0.7%)	9,615 (56.8%)	3,547 (21%)	941 (5.6%)
Philadelphia	109,013	270,845	1024 (0.9%)	6,977 (2.5%)	82,252 (74.1%)	21,062 (19%)	5,699 (5.1%)
Seattle	24,381	150,486	207 (0.8%)	2,423 (1.6%)	19,545 (75.2%)	4,262 (16.4%)	574 (2.2%)
Washington, DC	30,793	162,129	486 (1.4%)	2,789 (1.6%)	23,331 (68.7%)	6,835 (20.1%)	627 (1.8%)

Note: Percentages in parentheses refer to the fraction of licenses or units, as appropriate, from the first half of the table.

[10] Nashville only provided property counts, not unit counts.

§2.4.2 — Comparing Landlord and Rental Locations

Table 2.2 shows the fraction of landlords with residential mailing addresses who have a mailing address in the same municipality, same MSA, or same state as their rental properties. This table reports a smaller number of landlords than in Table 2.1, because I deduplicated the landlords so as to not count landlords who own multiple properties multiple times (see Chapter 3). Tables 2.3 and 2.4 show the differing locations of landlords based on whether they appear to be a "corporate landlord."[11] Table 2.5 shows the differing locations of landlords based on the number of units they own, again, using groupings of landlords identified in Chapter 3. In line with my interest in the inequality generated by landlords, I restrict all three tables to landlords who have a residential mailing address, as I expect these mailing addresses to correspond to a landlords' residence.

Table 2.2: Fraction of Landlords with a Residential Mailing Address, in Same City, MSA, or State

City	Number (Fraction of All Landlords)	Same City	Same MSA	Same State
Columbus	14,451 (72.1%)	48%	86%	89%
Dallas	2,811 (69.8%)	54%	85%	88%
Minneapolis	12,123 (89.4%)	50%	91%	92%
Nashville	4,053 (92%)	51%	80%	82%
Omaha	3,764 (80%)	68%	91%	90%
Philadelphia	36,686 (81.2%)	59%	88%	84%
Seattle	15,144 (85.3%)	61%	87%	90%
Washington, DC	19,457 (84.6%)	58%	84%	58%

Among landlords with a residential mailing address, this remains a local business. A slim majority of landlords with residential mailing addresses live in the same city in which they own property, and a substantial majority live in the same MSA. Given that I would expect landlords who have a residential mailing address to be a mom-and-pop or "circumstantial" landlord, it is no surprise that they live close — of the landlords interviewed by Shiffer–Sebba (2020), the majority of circumstantial landlords previously lived in the properties they now rent out (see also Crook, Ferrari, and Kemp 2012).

[11] For these two tables, I do not deduplicate landlords.

Table 2.3: Fraction of Properties with a Residential Mailing Address with a Corporate Landlord, in Same City, MSA, or State

City	Number (Fraction of All Properties with a residential mailing address)	Same City	Same MSA	Same State
Columbus	11,589 (39.4%)	37%	81%	84%
Dallas	2,298 (39.5%)	54%	83%	85%
Minneapolis	961 (5.3%)	46%	92%	92%
Nashville	1,770 (19.9%)	37%	60%	62%
Omaha	5,312 (46.3%)	65%	89%	88%
Philadelphia	21,682 (26.4%)	61%	89%	86%
Seattle	1,906 (9.8%)	67%	91%	94%
Washington, DC	5,387 (23.1%)	66%	91%	66%

Across these eight cities, the spatial spread of corporate landlords with a residential mailing address is different from all landlords, but no clear patterns emerge. In Columbus and Nashville, the fraction of corporate landlords with a residential mailing address in the same city or MSA are substantially smaller than the overall population of landlords. Yet in Washington, DC, and Seattle, there are more corporate landlords with a residential mailing address in the same city and MSA.

Table 2.4: Fraction of Properties without a Residential Mailing Address with a Corporate Landlord, in the Same City, MSA, or State.

City	Number (Fraction of All Properties without a Residential Mailing Address)	Same City	Same MSA	Same State
Columbus	17,086 (81.3%)	38%	59%	67%
Dallas	2,905 (69.5%)	37%	60%	70%
Minneapolis	766 (17.3%)	33%	64%	65%
Nashville	1,785 (65.2%)	23%	33%	33%
Omaha	2,569 (57.2%)	60%	81%	80%
Philadelphia	13,935 (52.1%)	68%	89%	87%
Seattle	1,801 (37.2%)	66%	91%	92%
Washington, DC	4,593 (61.6%)	67%	94%	67%

Insofar as it is simple and recommended for landlords to incorporate LLCs (Travis 2019), the lack of a pattern in the spatial distribution of corporate owners reflects the fact that being a "corporate landlord" does not necessarily mean a professionalized landlord, even if corporate landlords differ in other ways from non-corporate landlords (Ellen, Harwood, and O'Regan 2022). Table 2.4 shows some differences from Table 2.3, however. For instance, in Dallas, Minneapolis, and Nashville, corporate landlords without a residential mailing address are less likely to have a mailing address is the same city than corporate landlords with a residential mailing address. Comparing these two tables also shows that the likelihood of a corporate landlord having a residential mailing address differs across cities: in Minneapolis, Omaha, Philadelphia, Seattle, and Washington, more corporate landlords have residential mailing addresses than non-residential mailing addresses.

Table 2.5: Fraction of Properties with Landlords with a Residential Mailing Address, in Same City, MSA, or State, by landlord size

City	Landlord Size	Number (Fraction of Properties with a Residential Mailing Address)	Same City	Same MSA	Same State
Columbus	Large	80 (0.6%)	29%	67%	69%
	Medium	946 (6.5%)	36%	89%	91%
	Small	13,497 (92.9%)	56%	88%	91%
Dallas	Large	284 (10.1%)	72%	79%	80%
	Medium	322 (11.5%)	56%	90%	91%
	Small	2,205 (78.4%)	48%	83%	88%
Minneapolis	Large	125 (1%)	61%	97%	97%
	Medium	976 (8.1%)	42%	96%	96%
	Small	11,022 (90.9%)	52%	89%	90%
Nashville	Large	5 (0.1%)	29%	38%	38%
	Medium	212 (5.2%)	52%	92%	93%
	Small	3,836 (94.6%)	52%	79%	81%
Omaha	Large	80 (1.8%)	72%	97%	93%
	Medium	693 (15.9%)	66%	92%	92%
	Small	3,577 (82.2%)	68%	88%	88%
Philadelphia	Large	358 (1%)	78%	97%	96%
	Medium	4,534 (12.4%)	55%	92%	86%
	Small	31,793 (86.7%)	59%	84%	80%
Seattle	Large	217 (1.4%)	71%	97%	97%
	Medium	1,207 (8%)	65%	94%	97%
	Small	13,721 (90.6%)	59%	85%	88%
Washington, DC	Large	552 (2.9%)	79%	97%	79%
	Medium	671 (3.6%)	69%	93%	69%
	Small	17,630 (93.5%)	57%	82%	57%

Note: Large landlords hold more than 50 units, medium landlords hold between 6 and 50 units, and small landlords hold 5 or fewer units. In the case of Nashville, I use property, rather than unit counts.

A number of patterns emerge. First, there are relatively few large landlords who list residential mailing addresses. This makes sense: larger holdings are associated with professional landlords (Immergluck et al. 2020; Shiffer–Sebba 2020), who would be unlikely to continue to operate out of a home. In Dallas, Minneapolis, Omaha, Philadelphia, Seattle, and Washington, "large" landlords with residential mailing addresses are substantially more likely to be located in the same city or MSA as their property holdings. While it may be surprising that smaller landlords are less likely to be in the same city, it is perhaps a result of real estate investing platforms allowing for easy out-of-town investing (Parker and Friedman 2022), or circumstantial landlords who move away from their previous home locations. This presents an unusual dynamic, because while there are *many more* small landlords, they own many fewer properties, compared to large landlords, based on the breakdown that I have chosen. Table 2.5, however, is somewhat misleading, as the majority of large landlords have non-residential addresses — see Table 2.6.

Table 2.6: Fraction of Large Landlords with a Residential Mailing Address

City	**Number of Large Landlords (Fraction of All Large Landlords)**
Columbus	80 (31.6%)
Dallas	284 (42.8%)
Minneapolis	125 (41.5%)
Nashville	5 (50%)
Omaha	66 (31.6%)
Philadelphia	358 (48.8%)
Seattle	217 (48.4%)
Washington, DC	148 (40.3%)

Finally, Table 2.7 shows the mean and median distances between building and mailing addresses, subset by whether or not the mailing address is residential. As we have seen in the other tables presented in this subsection, we find that most landlords are local. In Columbus, Dallas, Minneapolis, Nashville, Seattle, and Washington, DC, the median residential mailing address is closer than the median non-residential mailing address, furthering the belief that "mom and pop" landlords prefer to be closer to their properties (D'Lima and Schultz 2021; Shiffer–Sebba 2020). One number stands out in this table: the median landlord without a residential mailing address in Nashville is located 202 miles from their rental property. This shocking number stems from the large number of properties owned by institutional investors in Nashville. Looking at the registry data, there are hundreds of homes registered to American Homes 4 Rent, "SFR Borrower III, LLC," and other institutional landlords. However, this may also reflect the likelihood of different types of landlords to register, given the relatively low compliance rate with Nashville's rental registries.

Table 2.7: Distance Between Landlord and Rental Properties, by Residential Mailing Address

City	Residential Mailing Address?	Mean (mi)	Median (mi)
Columbus	Overall	172.6	6.5
	Yes	119.4	5.3
	No	367.4	12.5
Dallas	Overall	197.3	8.6
	Yes	151.6	6.0
	No	123.3	9.0
Minneapolis	Overall	97.0	4.0
	Yes	85.6	3.9
	No	148.9	4.6
Nashville	Overall	309.5	9.3
	Yes	176.4	8.2
	No	817.0	202.1
Omaha	Overall	81.8	5.3
	Yes	73.4	5.4
	No	65.5	4.3
Philadelphia	Overall	47.0	4.2
	Yes	47.7	4.3
	No	40.1	3.7
Seattle	Overall	132.1	3.6
	Yes	131.6	3.5
	No	88.7	3.6
Washington, DC	Overall	128.3	3.3
	Yes	146.2	3.2
	No	50.2	3.3

§2.4.3 — Comparing Landlord and Rental Census Tracts

Table 2.8 shows the differences in demographic, housing, and income characteristics between landlords' mailing address tracts and their rental property's tracts. The table highlights just how different landlords' neighborhoods are from the neighborhoods where they own property. Because many of these socioeconomic variables would be skewed by the inclusion of downtown or commercial tracts, the tables presented here restricted the comparison to only those rental properties that had a landlord with a residential mailing address. Most of the simulated means differences are near zero, reflecting that, if landlords and properties were randomly distributed, the differences between landlord tracts and rental tracts would be minimal. I use this simulated mean to calculate a t-value based on the difference-in-means between the actual mean and simulated mean, most of which are statistically significant at the 0.05 level or smaller.

Table 2.8: Differences in Demographic Characteristics, Rental and Ownership Tracts.

Variable	Statistic	Columbus	Dallas	Minneapolis	Nashville	Omaha	Philadelphia	Seattle	Washington, DC
Difference in Percentage Non-Hispanic White	t-value	-96.87	-46.70	-66.80	-55.99	-74.76	-192.31	-5.66	-34.77
	Actual Mean	-14 PP	-18 PP	-12 PP	-16 PP	-21 PP	-22 PP	-1 PP	-6 PP
	Simulated Mean	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP
Difference in Percentage Black	t-value	100.68	28.85	66.59	56.72	59.83	169.37	20.19	49.90
	Actual Mean	13 PP	9 PP	8 PP	14 PP	13 PP	19 PP	1 PP	9 PP
	Simulated Mean	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP
Difference in Percentage Asian	t-value	-29.75	-32.56	3.21	-16.38	5.66	-54.10	-7.70	-44.30
	Actual Mean	-1 PP	-4 PP	0 PP	-1 PP	0 PP	-2 PP	-1 PP	-2 PP
	Simulated Mean	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP
Difference in Percentage Hispanic	t-value	25.69	39.95	36.61	22.25	41.38	78.78	-10.32	-7.81
	Actual Mean	1 PP	13 PP	3 PP	3 PP	7 PP	6 PP	-1 PP	-1 PP
	Simulated Mean	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP
Difference in Percentage College Graduates	t-value	-96.10	-43.04	-29.85	-44.35	-79.51	-180.08	27.78	3.97
	Actual Mean	-14 PP	-16 PP	-4 PP	-11 PP	-16 PP	-16 PP	3 PP	1 PP
	Simulated Mean	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP
Difference in Percentage Owner-Occupied	t-value	-115.33	-34.79	-74.44	-55.45	-63.71	-114.81	-62.38	-84.20
	Actual Mean	-19 PP	-13 PP	-16 PP	-16 PP	-16 PP	-10 PP	-11 PP	-16 PP
	Simulated Mean	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP

Variable	Statistic	Columbus	Dallas	Minneapolis	Nashville	Omaha	Philadelphia	Seattle	Washington, DC
Difference in Median Income	t-value	-133.96	-50.27	-76.42	-69.09	-88.01	-207.08	-37.98	-54.28
	Actual Mean	-$32,622	-$33,482	-$21,406	-$33,452	-$31,632	-$28,949	-$11,395	-$19,087
	Simulated Mean	$5	$0	$3	-$3	-$14	-$5	$3	$15
Difference in Median Rent	t-value	-73.32	-39.55	-50.08	-45.27	-39.97	-144.56	-28.12	-28.84
	Actual Mean	-$182	-$290	-$152	-$294	-$171	-$220	-$103	-$126
	Simulated Mean	$0	$0	$0	$0	$0	$0	$0	$0
Difference in Median House Value	t-value	-96.08	-37.78	-50.37	-51.58	-70.87	-180.69	-14.01	-20.68
	Actual Mean	-$124,168	-$159,336	-$63,260	-$163,901	-$109,564	-$131,542	-$29,334	-$39,648
	Simulated Mean	$9	-$88	$21	-$15	-$4	$5	$22	$125
Difference in Percentage with Income from Rent, Dividends, Interest	t-value	-135.23	-49.78	-70.77	-68.95	-86.46	-213.76	-26.78	-55.69
	Actual Mean	-10 PP	-11 PP	-6 PP	-11 PP	-11 PP	-10 PP	-2 PP	-6 PP
	Simulated Mean	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP	0 PP
Difference in Total Income from Rent, Dividends, Interest	t-value	-85.62	-28.28	-44.08	-45.52	-45.38	-107.73	-16.39	-20.31
	Actual Mean	- $7,395,183	-$9,894,803	-$6,226,562	-$12,210,949	-$6,263,388	-$5,516,849	-$3,542,248	-$3,737,347
	Simulated Mean	$389	-$647	$5,506	$8,363	-$271	-$1,062	$4,959	$5,885

Note: Restricted to rental properties where the owner had a residential mailing address. PP is "percentage points." Population with a College Degree is limited to those ages 25+. Estimates are from the 2015-2019 ACS, from Manson et al. (2021)

Overall, the results are striking. Rental tracts have significantly smaller proportions of white and Asian residents, and higher proportions of Black and Hispanic residents. As this is comparing differences across tracts, we can say that, on average, a Philadelphia landlord with a residential mailing address has an address in a census tract that has a 22 percentage point higher white population than the tract in which the landlord owns property. These patterns hold with three exceptions: rental tracts in Omaha and Minneapolis have slightly larger Asian populations than their respective landlord tracts, while rental tracts in Seattle and Washington, DC have slightly smaller Hispanic populations, and higher proportions of college-educated residents.

Not surprisingly, rental tracts have significantly lower homeownership rates than the tracts where their landlords are found, ranging from 10 percentage points lower in Philadelphia to 19 percentage points lower in Columbus. Rental tracts have lower household incomes, lower rents, and lower house values than their respective landlords' tracts. Rental tracts also have smaller populations reporting income from rentals, interest, dividends, royalties, estates, or trusts to the ACS. Throughout these cities, rental tracts, have, on average, 17 percent of households reporting rental or other non-wage income, compared to 26 percent of households in landlord tracts. These differences add up to serious sums of money: in Nashville, for example, landlord tracts report a cumulative $12.2 million more in this type of income, every year, than rental tracts.

§2.4.4 — Comparing Commuting and Landlord Market Networks

Figures 2.1 and 2.2 bring visual clarity to the stark differences between the landlord market areas and commuting backbones. The landlord backbones are mapped first (Panel A), and comprise the "landlord market area," while the commuting backbones are mapped below (Panel B). The borders of the center city are highlighted in black, while the MSA border is in red. It is clear that, while commuting backbones mostly remain within the MSA boundary, the landlord market area extends far beyond the MSA border.

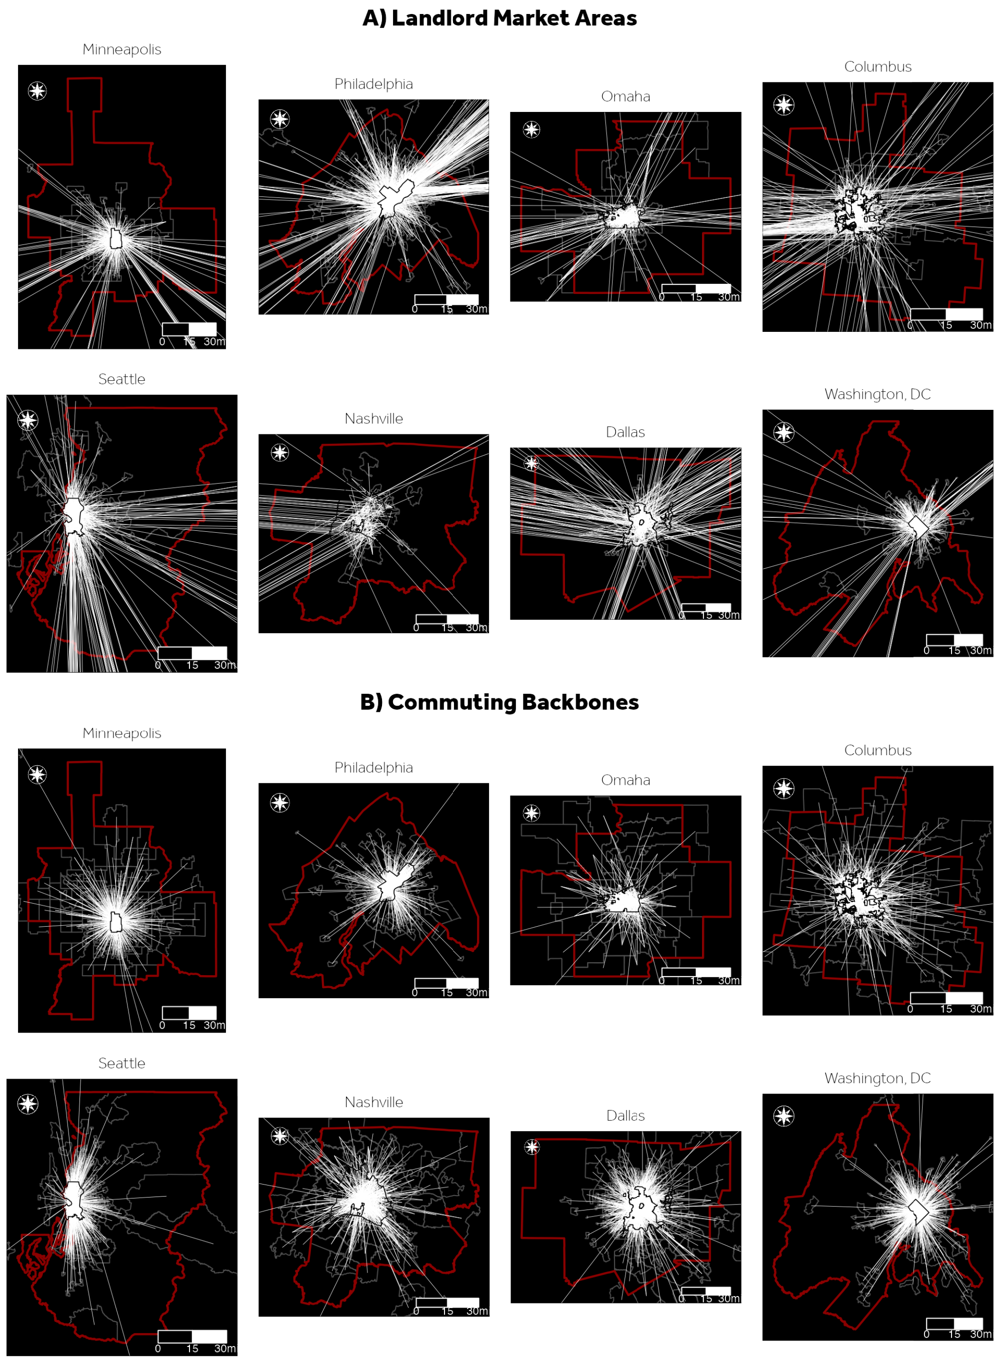

Figure 2.1: Landlord-property backbones and commuting backbones.

Note: Network edges are in white, central city borders are in black, MSA outlines are in red, network nodes (tracts) are outlined in grey.

Using the disparity filter to generate the backbones results in a substantially more manageable network size. For instance, the Washington, DC landlord-property network originally contained 4,086 census tracts with 15,846 edges among them. The landlord market area, created via the disparity filter, reduced the number of edges by 92.5%, and reduced the number of connected nodes by 90.6%, resulting in a landlord market area with 384 nodes and 959 edges. The average weight of a retained edge is 4.84, compared to a mean edge-weight of 1.94 in the landlord-property network. Across all eight cities, the disparity filter reduces the number of nodes and edges by approximately 70-95%.

Using the disparity filter to create the commuting backbone provides us with a baseline result for the accuracy of this approach for commuting and landlord networks. We would expect the commuting backbone to be mostly contained within the MSA boundary, and we would expect much of the MSA to be represented in the commuting backbone. Both of these expectations hold true. In turn, we can thus believe that the landlord market area provides us with an accurate picture of where there are large concentrations of landlords, just as the commuting backbone provides us with a picture of where there are large concentrations of commuters.

The landlord market areas in Figure 2.1 tells two stories. On the one hand, landlord market areas remain predominantly local. As we observed Section 2.4.2, most landlords have mailing addresses in the same city or MSA as their rental properties. In Minneapolis, there are 81 rental tracts in the landlord market area found outside the MSA — approximately 10 percent of the tracts found in the backbone itself. The Philadelphia landlord market has a number of nodes in cities of global capital, as well as other regionally important cities nearby: four out of five of New York City's boroughs are included; as are Rochester, Albany, Cleveland, Miami, Houston, Los Angeles, and many other cities. Rental ownership is much more national[12] than commuting.

Figure 2.2 shows the landlord market areas laid upon each other at the national scale, while Figure 2.3 includes the commuting backbones as well. These two national maps show that the landlord market area is predominantly local. Indeed, the white circles show that between 60% of landlord nodes are within the MSA in Dallas, and 92% of landlord nodes are within the MSA in Washington, DC. Yet the landlord market area is decidedly not exclusively local. At the same time, there are sites of national importance that appear across many of the cities as well. California, a

[12] Landlord-property networks are also more international than commuting networks, though the international addresses were dropped from this analysis.

powerhouse of the information technology industry; Texas, headquarters of Invitation Homes, among the largest institutional investor for single-family rentals; and Georgia, a site of significant single-family rental investment (Charles 2020a) are represented across all eight networks. While a few of the network nodes in the commuting networks are far from the center city, they appear to be fragments of the data process, rather than a pattern of long-distance connections we see in the landlord market area.

Landlord Market Areas Across the US

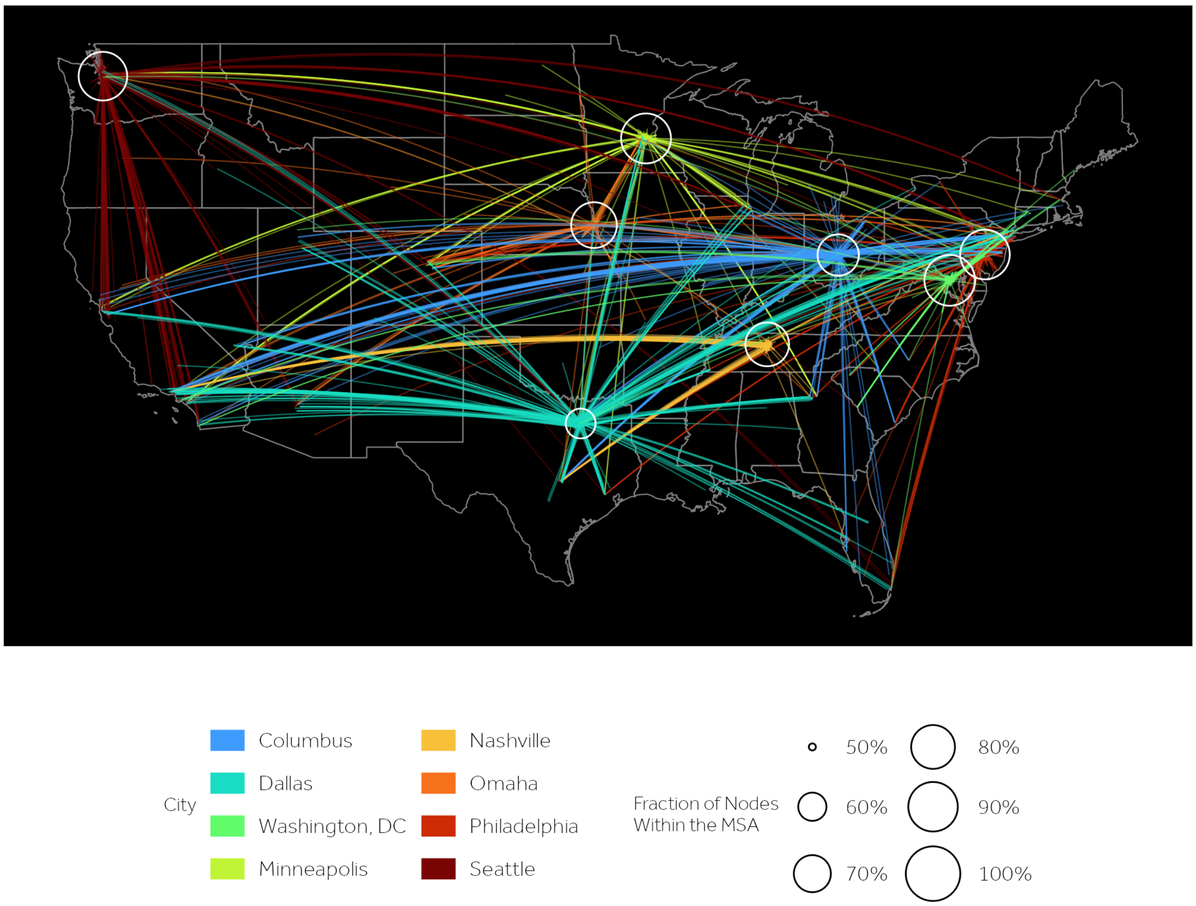

Figure 2.2: Landlord Market Areas for All Cities Analyzed

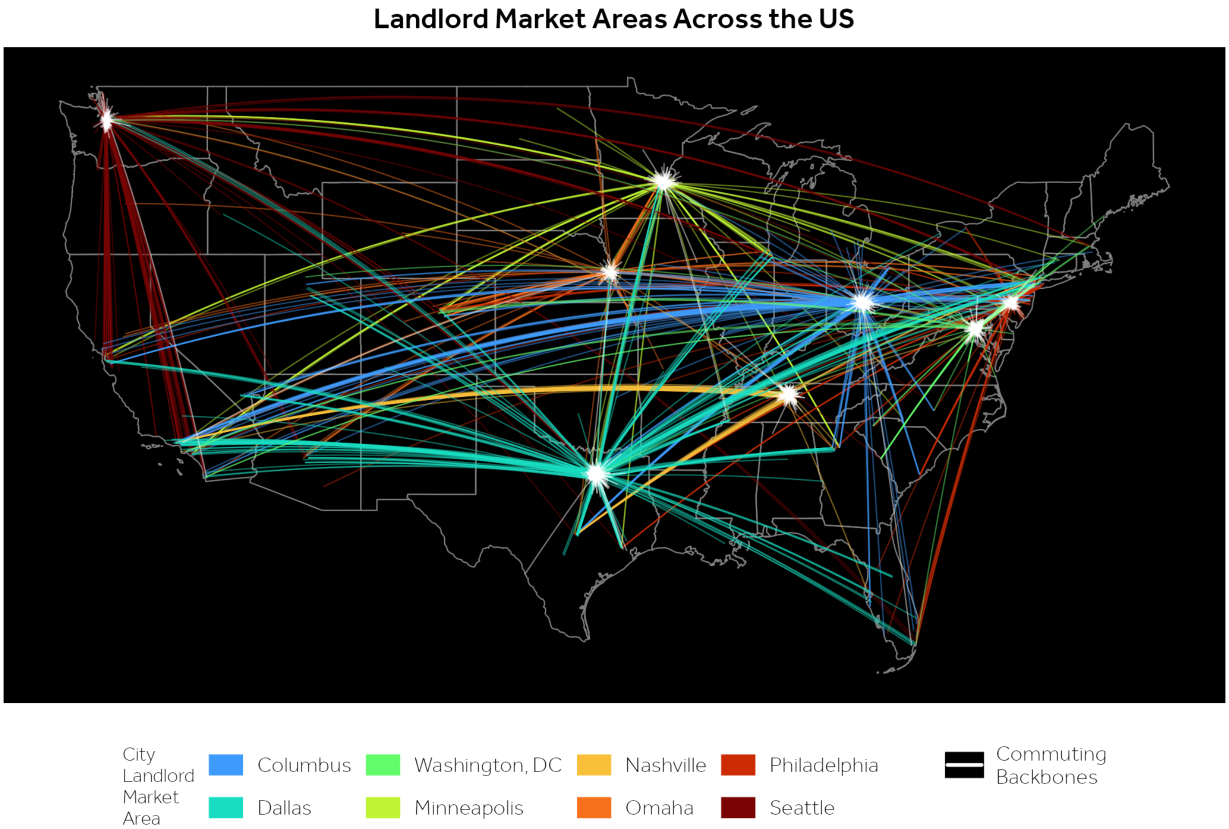

Figure 2.3: Landlord Market Areas for All Cities Analyzed, with Commuting Networks

§2.4.5 — Comparing Network Statistics

Table 2.9 shows the differences in network characteristics between the landlord market area and the commuting backbone. From these statistics, we see that the number of nodes and edges in the landlord and commuting backbones of a given city are often remarkably similar. Both the commuting networks and the landlord networks are extremely sparse, with low edge density, which is expected given the way that the backbones were constructed.[13] This is a product of using of the disparity filter, showing the utility of this approach in generating core spatial networks from larger and messier networks. Additionally, in half of the cities, the average degree is higher among the commuting networks, in the other half, among the landlord networks. This may reflect differing

[13] Edge density is a measure of the number of actual edges divided by the number of possible connections

patterns of commuting or rental property distribution within a city, as a higher average degree means that, on average, a single census tract is more connected to other census tracts within the network.

Table 2.9: Network Statistics Across Landlord and Commuting Backbones

Network	N Nodes	N Edges	Average Degree	Max Degree	Edge Density	Average Clustering Coefficient
Columbus Landlord	421	931	4.42	42	0.011	0.070
Columbus Commuting	408	1335	6.54	260	0.016	0.450
Dallas Landlord	487	512	2.10	15	0.004	0.008
Dallas Commuting	814	1705	4.19	444	0.005	0.254
Minneapolis Landlord	411	879	4.28	37	0.010	0.103
Minneapolis Commuting	541	927	3.43	323	0.006	0.534
Nashville Landlord	187	254	2.72	25	0.015	0.028
Nashville Commuting	322	875	5.43	243	0.017	0.478
Omaha Landlord	239	471	3.94	41	0.017	0.083
Omaha Commuting	252	778	6.17	151	0.025	0.426
Philadelphia Landlord	942	3103	6.59	48	0.007	0.137
Philadelphia Commuting	664	1435	4.32	286	0.007	0.184
Seattle Landlord	400	1079	5.40	54	0.014	0.137
Seattle Commuting	426	984	4.62	275	0.011	0.453
Washington, DC Landlord	384	959	4.99	45	0.013	0.077
Washington, DC Commuting	755	1382	3.66	340	0.005	0.474

Yet two clear patterns do stand out from these networks: first, commuting networks all have significantly larger maximum degree: for example, Minneapolis's most connected commuting node (tract) has 323 edges, while its most connected landlord node has only 37 edges. Second, clustering in the commuting networks is often higher than clustering in the landlord network. Both of these reflect real differences between commuting towards a downtown, as compared to a diffusion of rental properties throughout a city. These two patterns again reinforce that the backbone extraction produces networks that follow intuition. Backbone extraction produces reasonable core networks, and thus, we can infer accuracy about the geographical extent of the landlord and commuting networks.

§2.5 — Discussion and Conclusion

Creating and analyzing the landlord-property networks in eight large US cities shows us the complicated contours of the rental property market today. The maps of the landlord-property backbones — the landlord market areas — show us that these landlord networks are clearly national in scope. Many landlords exist beyond the region. Indeed, these maps likely understate the extent of this phenomenon, given that some landlords provide the addresses of property managers when complying with rental registration ordinances. Similarly, many institutional investors provide local mailing addresses rather than their national headquarters. At the same time, even if the vast majority of landlords are local, a sizable proportion of landlords have residential mailing addresses outside of the MSA. Many landlords are local, but the flow of capital to the landlord market area is national.

Landlord market areas were created using the disparity filter to extract the network backbone from the overall landlord-property network. Given that we had no baseline for what the landlord backbone should be, we compared that backbone to a commuting backbone for the same eight cities. The disparity filter produced reasonable approximations for the metropolitan statistical area, based on commuting data, demonstrating its utility in producing a core network from a larger and messier data set.

These landlord-property networks are substantially different from commuting networks and MSAs. MSAs, which are largely based on commuting flows between counties, cleanly contain the vast majority of commuting backbone nodes and edges. Most commuters — at least at the time that these data were collected — lived close to where they worked, and defining the contours of a region based on those commuting trips was possible. Landlord market areas have no such clean boundaries. Instead, while landlord market areas are primarily regional, most cities have landlord tracts in far-away cities. Indeed, in all cities studied, California, Georgia, and Texas tracts were part of the landlord market area. New York City, Dallas, San Francisco and Chicago tracts were also all heavily represented in these eight landlord market areas. These, in turn, are global cities and sites of global capital, technological innovation, and headquarters for large-scale landlords.

Network methods are what provide us with this avenue for understanding. Only by recognizing that rental properties and landlords are but dyads within a larger landlord-property network structure can we begin to piece together the way that modern rental property ownership differs from other types of relational flows in our cities. Extracting the landlord-property backbone allows us to see more clearly the boundaries of the landlord-property network, which would otherwise by obfuscated by the overwhelmingly large number of landlords and rentals. Network statistics reveal the differing structures between landlord-property networks and commuting networks.

Within the landlord markets themselves, areas where landlords live and areas where they own property are substantially different. Landlords have residential mailing addresses in areas with richer, whiter, and more educated neighbors than the areas where they own property. By comparing the socioeconomic differences between landlord nodes and renter nodes, this is not merely a tautology restating the fact that, in America, homeownership is correlated with many of these socioeconomic characteristics. Rather, it is clearly demonstrating that landlords own properties in neighborhoods that are different from their own. While not surprising, it is not immediately obvious why this need to be true. Owner-occupied rental properties, or landlords who own properties down the street, would have no differences in the socioeconomic variables calculated in this paper. Indeed, it is a peculiarity of the *distance* between modern-day landlords and their rental properties that produces this difference.

CHAPTER 3 | THE RELATIONSHIP BETWEEN LOCAL RENTAL MARKET OWNERSHIP CONCENTRATION AND RENT

Abstract

Over the last 30 years, the rental housing industry in the United States has seen a shift towards larger landlords who more often take a corporate form. At the same time, rental prices have continued to increase at a rapid pace. This study examines the relationship between concentration of ownership of rental units and rent levels and rent changes. It uses rental registry data from seven large American cities — Boston; Columbus; Dallas; Kansas City; Minneapolis; Seattle; and Washington, DC. It first links disparate owners using machine learning software trained to group landlords based on similar names, related addresses, and additional information. It then calculates the level of concentration at a variety of geographical scales. At the city scale, rental markets are not particularly concentrated. However, all cities studied have ZIP codes, census tracts, and elementary school boundaries that have high levels of concentration. Using rent data from a variety of sources, it finds that higher levels of ownership concentration at the ZIP code scale are associated with higher levels of asking rents. Additionally, it uses a shift-share instrument to identify the effect of higher levels of concentration on changes in rent, when accounting for city-wide income shocks. It finds that higher levels of concentration likely lead to higher gross paid rents. This research furthers our understanding of rental markets in the United States, pointing towards concerning levels of concentration in some urban neighborhoods, and their impacts on asking and actual rents.

§3.1 — Introduction

In the 21st century, rental markets in the United States have changed in numerous ways. First, the increased ease of incorporating a Limited Liability Corporation (LLC) since the 1990s has meant that landlords have increasingly taken on corporate forms (Travis 2019). Second, following the 2007 financial crisis, institutional investors — such as private equity firms, Real Estate Investment Trusts (REITs), and pension funds — expanded their presence in the single-family rental market (Mills, Molloy, and Zarutskie 2015). At the same time, rents have sharply increased in many major cities, with median rental price growth having exceeded renter income growth since 2000. The increased frequency of landlords using corporations as a means to own real estate, the increased presence of institutional investors in the housing market, and increasing prices for consumers has left the public, policymakers, and academics alike to wonder to what extent these changes are interrelated.

This paper aims to answer three questions: first, how concentrated is the ownership of rental properties in different housing markets in the United States? Second, are higher levels of concentration associated with higher rent levels? Third, do higher levels of concentration cause larger rent increases? It answers these questions using data from seven large American cities: Boston, MA; Columbus, OH; Dallas, TX; Kansas City; MO; Minneapolis, MN; Seattle, WA; and Washington, DC. It uses these seven cities because of unique data availability, as explained below. Nonetheless, these cities are very different — in size, levels of concentration, and industry, meaning that these results may be applicable to a broader set of US cities.

Across these seven cities, it finds low levels of concentration at the city level, but that there are numerous subgeographies with high levels of concentration. Higher levels of concentration are associated with higher asking rents. Using a Bartik-like instrument to measure expected income growth, it finds weak indications that neighborhoods with higher levels of ownership concentration have larger increases in rent than neighborhoods with lower levels of concentration. This paper looks at concentration across all landlords in the formal rental market; it does not address the question of the differences between institutional landlords and corporate landlords, a topic that is covered elsewhere (Desmond and Wilmers 2019; Ellen, Harwood, and O'Regan 2022; Fields 2018; Robinson and Steil 2021; Rose and Harris 2021; Rosen and Garboden 2022; Travis 2019).

The first question aims to tackle the problem of scale. While institutional investors, such as private equity firms, own only approximately one- to two-percent of the nation's single-family rental housing stock, many of these large landlords have concentrated their purchases in a select set of cities (Fields, Kohli, and Schafran 2016), namely those in the Sunbelt and West. The largest multifamily owners have grown substantially in size over the last 30 years, now controlling 50 percent more than they did in 1992 (National Multifamily Housing Council 2022). Additionally, while institutional investors may be a relatively rare phenomenon in the national rental market, there may be locally "large" investors that are important regional players that aren't nationally known. I aim to examine variation in the concentration of ownership across and within cities. I use multiple differing definitions of neighborhood scale, because housing submarkets are not clearly defined areas that can be cleanly captured by administrative areas.

The second question aims to elucidate whether concentration at the neighborhood scale is correlated with rent levels. While a handful of studies have begun to examine the relationship between rent levels and concentration (Linger, Singer, and Tatos 2022; Watson and Ziv 2022), this line of questioning is far from settled. It is not yet known whether there are submarkets where the level of concentration in the housing market has reached such a point where landlords have potentially oligopolistic or even monopolistic pricing power. Ownership over a parcel of land obviously gives the owner a monopoly over the parcel itself — the question is whether a landlord, owning a significant number of properties in a neighborhood, is able to translate their monopolistic control of numerous parcels into higher prices than would otherwise be expected. While this cannot be directly answered with the data at hand, it is possible to take a first look at this question by analyzing the relationship between neighborhood-level concentration and average asking rents.

The third question is related to the second, and asks whether higher levels of concentration result in higher rent increases when a city experiences an income shock. This causal question thus differentiates between different sources of rent increases; namely, the city-wide income changes which could lead to higher willingness to pay, and neighborhood-level concentration in the ownership of the rental market. In order to analyze these two questions, I use a Bartik-like shift-share instrument (Bartik 1991) to estimate changes in income.

Identifying concentrated ownership of rental properties in the US context presents a number of challenges. First, it is difficult to distinguish between which housing units are owned solely for the purpose of owner occupation, which for rentals, and which for another purpose (such as second homes, short term rentals, or speculative owners who leave units vacant). More detail is discussed in Chapter 4 of this dissertation. The typical approach is to call all properties "rentals" where the owner's mailing address is different than the property's address (Mills, Molloy, and Zarutskie 2019; Shelton 2018). Unfortunately, these approaches are *under*inclusive of those owners who send mail to their rental properties (Sakong 2021, 15), condominiums, cooperatives, and mixed use buildings. Indeed, many prior studies explicitly exclude mixed-use buildings and condominiums, and by design are unable to distinguish among cooperative ownership (Gomory 2021; Hangen and O'Brien 2022). These approaches may also be *over*inclusive with respect to properties that have a tax mailing address that is different than their physical address, which might include second homes (Shelton 2021), vacant properties, properties under redevelopment, bank-owned properties, short term rentals, and properties where the owner chooses to mail their property tax bill to an accountant or bank. To address these concerns, I used city-administered rental registries in seven cities. These registries suffer from enforcement problems, but nonetheless represent the legal, or formal, rental housing market in these cities.

Second, the rise of the corporate form of landownership means that identifying rental properties which have owners with distinct corporate names yet are actually controlled by the same overarching entity is quite challenging. Owners often use unique LLCs for each rental property. Scholars have taken multiple different approaches in recent years, often relying on natural language processing or machine learning algorithms trained on company names and mailing addresses, and sometimes combined with corporate registry data (Gomory 2021; Hangen and O'Brien 2022; Immergluck et al. 2020; Robsky Huntley et al. 2022; Watson and Ziv 2022). Rental registries again address some of these concerns, as they often ask landlords to identify the actual individual owners or other contact names that point to common ownership, in addition to other identifying information, such as phone numbers or email addresses. Relying on rental registry data allows for the capture of the formal rental housing market in a city, including units in condominium, cooperative, and mixed-use buildings.

§3.2 — Literature Review

This paper contributes to several long-running strands of regional studies and urban economics literature, including the extent to which housing markets are competitive and the reasons as to why landlords may be at an advantage in the context of the rental housing market. Recent scholarship has examined concentration of the rental housing market; this paper directly contributes to this growing body of work. The geography of search is an important component to the validity of this study, and thus I draw on the literature relating to housing search to ground my methodological decisions. Finally, this paper relies on a shift-share instrument, a longstanding technique within the urban economics literature used to estimate the causal impact of shocks at the local level.

Scholars have long questioned the extent to which rental markets are competitive. This question attracted significant interest in the 1970s and 1980s, with a number of papers questioning the economic orthodoxy that housing markets generally, and rental markets specifically, represented a competitive market (Arnott 1995; Cherry and Ford Jr. 1975; Gilderbloom 1985; Gilderbloom 1989; Gilderbloom and Appelbaum 1987). In particular, Gilderbloom and Appelbaum (1987) argued that many of the characteristics of rental markets mean that they are not perfectly competitive: rental markets have relatively few owners, who may collude (implicitly or otherwise); there are high entry and exit costs for both landlords and tenants; knowledge of the market is imperfect, as is knowledge of the good itself; there are often substantial constraints on the supply of new housing; and housing is a heterogenous product with substantial differentiation among units. Using census data across 140 urban areas, they found that median income and professionalization of the landlord industry were primary drivers of rental prices.

There are a number of reasons one might expect landlords to have the upper hand within the rental market today. For one, information — about comparable lease terms (as compared to asking rents), long-term quality, rental market history — is not freely available and landlords may have better information than tenants. Tenants only enter the search process irregularly, necessitating search and information acquisition activities (Maclennan and O'Sullivan 2012). On the other hand, large landlords can gather information based on their current rental portfolio (Cherry and Ford Jr. 1975). Indeed, some institutional investors today use advanced statistical models to set and raise rent (Abood 2018), while smaller landlords are often reluctant to raise rent on an existing tenant (Shiffer–Sebba 2020). These new models and sources of proprietary information allow landlords to set rents closer to the market, while tenants do not have access to similar information. Additionally, landlords might implicitly or explicitly collude with one another, sharing contract rental prices to inform their own understanding of what rent to charge (Vogell 2022). Tenants, on the other hand, typically can only access information on asking rents, not contract rents, and may not be able to access historical asking prices. Additionally, segmented housing markets mean that, if people are only searching in a particular geography or for a particular type of rental (like a single-family home in a suburban neighborhood), landlords may be able to exert pricing power within that submarket.

In recent years, there has been a number of papers examining concentration in rental markets (Hangen and O'Brien 2022; Lennartz 2014; Linger, Singer, and Tatos 2022; Tapp and Peiser 2022; Watson and Ziv 2022). Of these, two examine the relationship between rent and concentration, with Watson and Ziv offering evidence that higher levels of concentration are associated with higher rents in New York City, while Linger et al. use an instrumental variable approach to identify a causal relationship between higher levels of concentration and higher rents and rental inflation in Florida. In a related literature, Cosman and Quintero (2021) found that lack of competition in the construction industry leads to higher prices in local housing markets, demonstrating that housing markets are subject to concentration effects.

One challenge with identifying concentration in rental markets — and then evaluating the relationship between concentration and other outcomes — is that the ownership of rental housing is often obscured in the United States. The first step in identifying rental market concentration is to identify which properties in a city are renter-occupied, and how many units there are in a property. Surprisingly, this is no easy task, as I explore further in Chapter 4. It is quite difficult to differentiate between owner-occupied homes, rental properties, or other residential property tenures. The most common approach taken is to rely on property tax assessment databases (Gomory 2021; Hangen and O'Brien 2022; Immergluck et al. 2020; Robsky Huntley et al. 2022; Rose and Harris 2021; Tapp and Peiser 2022). Yet this approach obscures issues with second homes (Shelton 2021), foreign investors (Sakong 2021), and on-site management of multifamily properties (Carswell 2018).

Another challenge in identifying concentration is that it requires a count of the number of rental units, not just a count of the number of rental properties. Few cities collect this information within the tax assessment process, with many recording ranges of apartment building sizes, though there are some examples where property tax information includes unit counts.[14] Prior scholarship has addressed this challenge by relying on data from private sources which provide unit counts on multi-family rentals (Tapp and Peiser 2022), using multi-family rental registry data (Watson and Ziv 2022), or by restricting the study to only single-family rentals (Fields and Vergerio 2022; Gurun et al. 2022).

To overcome some of these barriers, I use data on property ownership gathered from rental licenses as required by city rental registration ordinances, and augmented with information from tax assessment databases. These ordinances are common in a number of large American cities: among the 50 most populous US cities, 30 have rental registration ordinances of some kind. Of these 30, 24 have no or few exclusions: all rentals that are not occupied by their owners or an owner's family member are typically required to register. Compliance and enforcement ranges significantly across cities. One report estimates that only 15% of Detroit's rental properties are properly licensed, with significant heterogeneity between small "mom and pop" landlords who own only one or two properties (only 2% of which are estimated to be registered), and larger, more professional landlords (of which 51% are estimated to be registered) (Lynch 2022). On the other hand, it appears that nearly all units are registered in Washington, DC with 175,000 registered rental units compared to 195,000 rental units estimated by the ACS. Rental registration ordinances are often building-level, requiring landlords and management companies to disclose the number of units at a particular address that they rent. Some cities also require LLCs that own rental property to disclose information about beneficial owners, allowing for clear linkages of disparate

[14] For instance, while Minneapolis and Boston provide unit counts for multifamily apartment buildings, Philadelphia, Seattle, Kansas City, DC, and Columbus provide unit ranges.

LLCs to actual ownership.[15] Almost all registration ordinances require a building owner to identify a manager if they are not themselves managing the building.[16]

This study builds on recent scholarship that has sought methodological advances in the linking of corporate landlords. Compounding the challenges of identifying concentrated ownership in the rental market is the fact that multiple, distinct corporations can be owned in common and act as if they were one common entity. Immergluck et al. (2020) and Robsky Huntley et al. (2022) both use "Dedupe," a machine learning package in Python to try and identify common owners through similar names and mailing addresses. Hangen and O'Brien (2022), who use an approach similar to Gomory (2021), criticize this approach, charging that Dedupe can be both too inclusive with regards to certain names — i.e. it generates false positives — and too exclusive — i.e. it generates false negatives. Instead, they propose a natural-language processing approach combined with corporate agent data; both papers focus on Boston, Massachusetts.

One final challenge in the literature is worthy of mention: the geography of a rental market is not obvious. There has been longstanding interest in defining housing market areas, including housing market areas that differentiates between owners and renters (Jones 2002; Jones and Coombes 2013; Royuela and Vargas 2009). These studies question the typical approach of using the metropolitan statistical area as coterminous with a housing market. Yet, for the purposes of search, it is not clear that renters would search an entire city or region; indeed, we can imagine that some neighborhoods or municipalities are completely off-limits to renters, given restrictions on multifamily construction, on renters,[17] or on unrelated members in a single household (Airgood-Obrycki and Wedeen 2022). The literature on housing search is particularly illustrative here: when households search for a new place to live, they often select a small set of neighborhoods to search, biased towards short-distance moves and impacted by racial segregation and differences in neighborhood housing costs (Bruch and Swait 2019; Carrillo et al. 2016; Krysan 2008). Rae and

[15] Among the cities included in this paper, Minneapolis, Washington, DC, and Philadelphia all require beneficial owner information.

[16] Among the cities included in this paper, only Kansas City does not explicitly collect management information

[17] Here, I am imagining homes in communities or condominiums with homeowner associates (HOAs) that explicitly limit the number of renters, or prohibit them altogether.

Sener (2016) use data from an online housing search portal that allowed users to draw their own boundaries. They find that a majority of users draw a search geography that is less than five square kilometers (Rae and Sener 2016, 142). For comparison, the average ZIP code in Washington, DC is approximately seven square kilometers, while in Dallas it is 20 square kilometers.

Census tracts have typically been the geography of choice for those studying rental concentration (Linger, Singer, and Tatos 2022; Tapp and Peiser 2022; Watson and Ziv 2022). However, census tracts are not likely a geography that would be particularly meaningful to renters seeking a home. While census tracts are a typical avenue for operationalizing neighborhood definitions in the United States, it is not clear that an area as small as a tract represents a housing submarket that would be relevant to either renters or policymakers. Indeed, scholarship looking specifically at the question of housing submarkets often aggregate census tracts into geographically contiguous areas based on a number of factors (Goodman and Thibodeau 2003; Hwang 2015). Being precise in defining geographic housing submarkets is crucial, given how geography is of particular interest to households when engaging in search for a new housing unit, with households looking in specific areas (Bruch and Swait 2019). Though census tracts may at times represent neighborhoods as understood on the ground, they also may bisect communities. Thus, this paper first examines concentration at a variety of different levels of aggregation, including tract, ZIP code, elementary school attendance zone, and city-defined neighborhood boundaries. I primarily rely on the ZIP code level for my regression analysis, given other data availability. However, I include the census tracts to compare to Linger, Singer, and Tatos (2022) and Watson and Ziv (2022); I include school attendance zones given the important interconnections between schools and neighborhood choice in the US context (Owens 2020).

Additionally, the statistical analysis in this paper asks: when there is a city-wide income shock, do neighborhoods with higher levels of concentration experience larger price increases? This contributes to a longstanding literature examining the impact of shocks on rents and housing prices, such as immigration (Gonzalez and Ortega 2013; Saiz 2007; Saiz and Wachter 2011). It directly builds on the work of Guerrieri, Hartley, and Hurst (2013) who use the same instrumental variable. In their work, they sought to understand the relationship between house price appreciation and gentrification, which they believed was mediated by distance between poor and rich neighborhoods; they used city-wide income shocks to understand the interrelated nature of these distinct phenomena. Their work with a Bartik-style instrument (Bartik 1991) is part of a long line of urban economics research that employs this technique (Baum-Snow and Ferreira 2015).

§3.3 — Data and Methods

§3.3.1 — Data

This paper primarily uses data on rental properties as collected by cities through a rental licensing or rental registration process. I acquired these data through open data portals, right-to-know requests, or through web scraping of public websites. Among public record requests sent to 20 cities, and another set collected through open data portals, seven had sufficient information to be included in the present paper. In order to be used in this paper, rental registries (and responses to the public records request) had to provide unit counts, property addresses, names and mailing addresses of owners. Some registries provide additional information, including the phone number or email address of the property owner; some also require that, in the case that the property owner is a corporation, the name(s) of controlling owners for the corporation are also disclosed. I used all provided information about the properties and their owners in order to identify common owners. In some cases, I combine registry data with property tax assessment data from Regrid, a national data provider of standardized parcel data. See Appendix A for more information. A rental registry, when I received it from the local government, may look as follows:

Table 3.1: Example Registry

Row Number	License Number	Property Address	Number of Units	Owner Name	Owner Other Name	Owner Address
1	50000040	29 Galveston St	12	Gresham Shawn		1636 Kenyon Str NW Washington DC 20010
2	565901	1115 Monroe St	27	Shawn Gresham		1636 Kenyon St NW Washington DC 20010
3	709378	3021 Chicago Ave #103	1	Blueberry 107 LLC	John Bilikam	70 Ardmore St Hamden CT 06517
4	536711	401 Ridgewood Ave	3	John Bilikam	Victoria Husband	5736 N Maryland Ave Portland OR 97217
5	889633	3537 Park Ave	12	Blueberry 107 LLC	John Bilikam	70 Ardmore St Hamden CT 06517

Note: This table was constructed based on actual information from the rental registries used in this paper, but is not a subset from any one registry. It was created in order to most easily illustrate the approach to deduplicating landlords.

In the cases where the rental registries do not require additional information from corporate landlords, I try and match corporate names to the OpenCorporates database. OpenCorporates is a website dedicated to providing information about corporations: their owners, addresses, and membership structure. They provide a free API for public use. In some US states, OpenCorporates is able to provide the names of agents or officers for corporations registered with that state, thus allowing matching of corporate owners based on human names, in addition to corporate names, as in Hangen and O'Brien (2022). I first match corporate names to the OpenCorporates database using the OpenCorporates reconciliation API via the software OpenRefine, wherein I provide OpenCorporates the corporation's name and state; in turn, OpenCorporates provides the most likely corporate match. I then use the OpenCorporates API to gather the names of agents or corporate members of each matched corporation. I exclude agents' names that seem to be from a general incorporation service.[18]

Data for covariates comes from the American Community Survey (ACS) 5-year estimates (2010-2014 or 2017-2021), provided by the National Historical Geographic Information System (Manson et al. 2021). To calculate income shocks, I rely on ACS 1-year (2014 and 2021) and 5-year (2010-2014, 2017-2021) microdata provided by IPUMS USA at the metropolitan statistical area level, reconstructed via Public Use Microdata Areas (PUMAs) (Ruggles et al. 2023). I rely on the location of central business districts, as calculated by Manduca (2020). Elementary School Attendance Boundaries are provided by the School Attendance Boundary Survey (US Department of Education National Center for Education Statistics 2016).

I primarily rely on Zillow's Observed Rent Index (ZORI) as the measure of asking rents at the ZIP code level. In general, data on actual and asking rents in the US are difficult to come by. ZORI is a repeat-rent index of a "typical" rental, weighted to represent the entire market. Unfortunately, ZORI has a problem with missing data. Only 20 percent of US ZIP codes are provided by ZORI, a number that shrinks further for their historical data. These numbers overstate the problem somewhat, as ZORI availability is significantly higher in central cities. In the context

[18] These are agent names such as Registered Agents, Inc.; Incorporation Services, LLC; Comercial Business Services; etc. I also exclude names of some human agents, namely those that seem to be part of an incorporation service.

of this study, ZIPs with ZORI coverage have lower levels of HHI than tracts missing ZORI coverage, but larger number of rental units.[19] See Appendix B1 for balance distributions among the covariates. In general, ZORI provides data on 2022 rents for 82 percent of ZIP codes in my study. However, ZORI provides data on 2014 rents for only 43 percent of ZIP codes in my study. This seems primarily driven by the fact that Zillow only releases data on rents and house prices at smaller geographies when they can be assured of data quality.[20] However, because of the missing data associated with ZORI, I also use asking rent data in 2022 from Rentometer, a Zillow competitor with greater geographic coverage,[21] and with median gross rent from the 2017-2021 American Community Survey.

§3.3.2 — Methods

§3.3.2.1 — Deduplication

I first identify owners who own multiple properties, a process of deduplication and clustering. In general, rental registry data is inconsistently recorded. Thus I first clean owners' name(s) and mailing address(es). I standardize street names and addresses, phone numbers, common abbreviations, and identify which owners are likely to be corporations based on keyword matching across a host of words likely to appear in a corporate owner's name.[22] I identify corporate landlords in order to match information from the registry with information from OpenCorporates. I first deterministically group different rows based on the owner's name, a secondary owner name (if

[19] This is because, across all seven cities, HHI is negatively correlated with rental unit counts. See Figure 3.1

[20] Gupta et al. (2022) analyzed the missingness of ZORI data, finding that Zillow has greater coverage closer to the city center, as well as in neighborhoods with higher proportion of renters. In the present study, the average ZIP code without Zillow data has only 1,200 rental units, compared to 3,800 rental units for ZIP codes with Zillow data in 2022 only, and 6,600 rental units for ZIP codes with Zillow data in both 2014 and 2022.

[21] Rentometer and ZORI have a correlation coefficient of 0.90 for the ZIP codes in my study area, yet Rentometer, in general, shows higher rents than ZORI across these ZIP codes.

ZORI does not release its methodology, so it isn't possible to say why Rentometer has greater coverage than ZORI. However, Rentometer does release the number of observations per ZIP code, and will release an estimate with as few as 2 observations in a given ZIP. I only retain Rentometer data for ZIPs with 10 or more observations.

[22] This includes looking for words like LLC, LP, Limited, Development, Associates, Real Estate, etc., in addition to names that start with numbers. This latter keyword search is because it is common for landlords to use the property address as its own LLC name, such as "123 Main Street LLC."

provided), and the owner address.[23] If all three are identical (after data standardization), I say that the two rental properties are owned by the same landlord — above, in Table 3.1, rows 3 and 5 would be classified as being in the same cluster, i.e. have the same owner.

I then use "Dedupe," a machine learning package in python to probabilistically cluster owners with similar names, mailing addresses, and other identifying information. Despite the concerns raised by Hangen and O'Brien (2022), I find Dedupe is the most appropriate tool to handle the multitudinous data I use to identify common owners. Dedupe allows one to compare multiple different fields as different types of data. For instance, I can program Dedupe to compare entries within the "Owner Name" column as both strings of characters (which is helpful for catching typos), and as sets of words (which is helpful for identifying matches when names are presented in different orders). Dedupe can also compare across columns, looking for similarities in both, say, the "Owner Name" and "Owner Other Name" columns.

Dedupe identifies rules to probabilistically group rows based on a training session in which the user is given pairs to match, and the user tells the program whether that pair is or is not a true match (or if the user is unsure.) I train Dedupe separately for each city, given the differences in information provided by the cities, and the unique characteristics of city address conventions. During this training session, I manually identify the names and mailing addresses of management companies,[24] and remove those names and mailing addresses from the dataset so as to match only on owner information, and not management company information.[25] In general, I require two pieces of similar information in order to tell Dedupe that a pair is a match. For example, in the table above, I would tell Dedupe that rows 1 and 2 are a match, but I would not tell Dedupe that row 4 matches with any other row provided. However, if I also had owner phone numbers, and the

[23] I borrow Hangen and O'Brien's (2022) distinction between "deterministic" matching and "probabilistic" matching.

[24] During early stages of the deduplication process, I note companies with names that may indicate that they are property management companies. I also look for addresses that appear across multiple entries, but with owner names that seem distinct. I then searched for these companies online to see if I could identify them as a property management company that works with multiple owners.

[25] Despite the requirement that owner names and mailing addresses be used, it was not uncommon to identify management companies in this work. Because I cannot disentangle the effects of management concentration from the effects of ownership concentration, I remove management information.

phone numbers for rows 3, 4, and 5 were all the same, then I would tell Dedupe that the pair of 4/5 or 3/4 is a match, since they would share a common name and phone number. I provide Dedupe with anywhere between 25 and 100 positive and negative matches per training session.

This is an inherently subjective process. While I would tell Dedupe that John Smith and Jane Smith with the same mailing address should be part of the same pair, I may not tell Dedupe that Jane Doe and Barbara Smith with the same mailing address are a match, barring additional information. Corporate matching is no easier without additional information. While many companies may have similar names that, at first glance, might lead to the identification of a cluster, companies are only deemed to be a match if they shared an owner name, phone number, or mailing address, regardless of naming similarity. Additionally, while care is taken to remove property manager names and addresses during the matching process, some management companies both own and manage their properties in addition to properties owned by others, which, from the data available, cannot be easily differentiated.

In order to ensure that the groupings of landlords are not overinclusive, I manually check a subset of the clusters after Dedupe has been trained and has returned a table of likely matches. Because of my concern with concentration, I check all groupings of landlords where more than 10 individual landlords are linked. I also check ten percent of all clusters that have between two and nine distinct entries. If both sets of ('large' and 'small') clusters have fewer than five percent false-positives, I accept the clusters suggested by Dedupe. Here, I define a false positive as an entry being part of a cluster when it should not be — each entry that is incorrectly clustered is one false positive. For example, Dedupe might cluster "Daniel J Smith," "Daniel C Smith," and "Daniel R Smith," which I would call three false positives. If Dedupe has more than five percent false positives, I run a new training session until there are fewer than five percent false positives. Generally, I do not check for false negatives; that is, rows that should be clustered but are not. However, in some cases, I will join two large clusters together that ought to be combined.[26] For ease of reading, I refer to

---

[26] For example, there may be 10 entries that are clustered all with the name "SFR3-030" and another 15 that are clustered with the name "SFR3-030 LLC." In that case, I may either choose to retrain Dedupe so as that cluster is merged probabilistically, or I will manually merge it.

clustered landlords (where multiple landlords with similar or identical names are believed to be the same entity), simply as "landlords." I report results from Deduplication in Appendix B2.

§3.3.2.2 — Concentration

In order to identify levels of concentration, I use the Herfindahl-Hirschman Index (HHI). The HHI is a common measure of market concentration, and is calculated by taking the sum of the squares of market share. In this case, market share is measured by the percentage of all housing units in the market owned by a single landlord. The US Department of Justice considers a market with an HHI between 1,500 and 2,500 to be "moderately concentrated" and a market with an HHI above 2,500 to be "highly concentrated,"[27] thresholds that have been adopted by housing scholars elsewhere (Tapp and Peiser 2022; Watson and Ziv 2022).

Calculating the HHI requires the definition of a market. To examine the simple question — is the rental market concentrated — I take no *a priori* definition of a rental market. Instead, I calculate the HHI over the following definitions of a market: the city, the five-digit ZIP code, the census tract,[28] and the elementary school attendance zone. For further statistical analyses, however, I only rely on the ZIP code, given their use in defining geographic submarkets, and the availability of asking rent data at this geography. While an imperfect measure of a rental market, prior literature on housing submarkets have found that ZIP codes characterize housing submarkets well despite their relative arbitrariness (Goodman and Thibodeau 2003). HHI_a, the level of concentration in a market area *a* is thus defined:

$$HHI_a = \sum_{i}^{n} s_i^2$$

Where the market share of a landlord, *i*, in area *a*, is defined

$$s_{i,a} = \frac{\sum_{i,a} units}{\sum_{a} units} \times 100$$

[27] https://www.justice.gov/atr/herfindahl-hirschman-index

[28] Census tracts and ZIP codes may cross city boundaries. I calculate the HHI based solely on data from rental registries. Thus, when a tract crosses city boundaries, I report the level of concentration of within city limits.

§3.3.2.3 — Cross-Sectional Analysis

To conduct a statistical analysis of the relationship between rent and income, I use OLS regression to examine the relationship between rental prices and concentration levels at the ZIP code level. My estimating equation is as follows:

$$y_i = \beta_0 + \beta_1 HHI_i + \boldsymbol{\beta}\mathbf{X_i} + \gamma_c + \epsilon$$

Where y_i is the log of the rent prices in ZIP code *i*, from ZORI, Rentometer, or the American Community Survey, HHI_i is the HHI in ZIP code *i*, X_i is a set of ZIP-code level covariates, including percentage of the population that is white, percentage of the population that has a college degree, median house value, median household income, number of renter-occupied housing units, fraction of housing units that are owner-occupied, fraction of the housing units that are vacant, and distance to the central business district.[29] I include city-level fixed effects, γ_c; ϵ is the error term. Standard errors are clustered at the city level.

§3.3.2.4 — Income Shock Analysis

I ask one final question: when there is a city-wide income shock, how does neighborhood-level concentration affect rent changes? Following Guerrieri, Hartley, and Hurst (2013), I calculate *expected* income growth using a shift-share instrument based on wage growth and industry composition between 2014 and 2021, at the MSA-level. That is, given industry composition in 2014 and national industry wage growth between 2014 and 2021, I predict the city's wage growth in 2021 based on the wages and industry shares in that city in 2014 and national wage growth to generate $\widehat{incomegrowth}_c$. I use the 3-digit industry classification level. To ensure reliability, a regression of expected income growth on actual income growth yields a coefficient of 2.27 with an F-statistic of 9.57. Because I am looking at an income shock over time, I also opt to examine rent price *changes*, rather than rents themselves.

[29] I calculate the centroid-to-centroid distance, recognizing that ZIP codes are often large geographies where there are some areas that may be substantially closer to the CBD.

My estimating equation here is as follows:

$$\frac{\Delta y_i}{y_{i,2014}} = \beta_0 + \beta_1 HHI_i + \beta_2 HHI_i \times \widehat{incomegrowth_c} + \beta_3 \times \widehat{incomegrowth_c} + \boldsymbol{\beta}\mathbf{X_i} + \epsilon$$

Where $\frac{\Delta y_i}{y_{i,2014}}$ is now the percentage difference between rents in 2022 and 2014 (the earliest year for which ZORI rents are available). My main coefficient of interest is β_2, which represents the interaction between the city-wide income growth and rental ownership concentration. The variables X_i, HHI_i, and ϵ, are as above. In these regressions, I do not include city-level fixed effects because including both city-level fixed effects and the income shock (which is unique to each city) introduces collinearity problems. Because of the sparse data availability from ZORI, this analysis is re-run using changes in the median gross rent between the 2010-2014 ACS and the 2017-2021 ACS, again at the ZIP code level.

§3.4 — Results

§3.4.1 — Concentration by geography

Table 3.2 shows the concentration of the rental market at a variety of geographic levels in the seven cities studied here. At the city level, the level of concentration is generally quite low, with HHIs indicating that, city-wide, few owners control more than a tiny portion of the overall rental stock. Nonetheless, the citywide HHI does hide some aspects of the local rental markets. In Washington, DC, the largest landlord controls approximately 6 percent of the rental stock, while only 102 landlords, or 0.4 percent of all registered landlords, control 50 percent of the total rental stock. In Seattle, the city with the lowest city-wide HHI, the largest landlord only controls 1.3 percent of the housing stock, and 234 landlords — or 1.2 percent of all landlords registered with the city — control 50% of the rental units.

Table 3.2: Geographic Concentration

		% of Rental Units Owned by Corporate LLs	City-wide HHI	Rental ZIP Code HHI	Rental Elementary School Attendance Boundary HHI	Rental Tract HHI
Boston		36%	26			
	Mean			1,364	2,799[30]	1,392
	Median			391	2,799	683
	Max			10,000	2,856	10,000
Columbus		65%	23			
	Mean			726	2,925	1,551
	Median			370	2,031	804
	Max			4,390	10,000	10,000
Dallas		70%	28			
	Mean			1,828	2,809	3,401
	Median			948	2,086	2,639
	Max			10,000	10,000	10,000
Kansas City		84%	34			
	Mean			2,049	2,529	2,194
	Median			872	1,933	1,548
	Max			10,000	9,903	9,884
Minneapolis		17%	32			
	Mean			497	212	626
	Median			158	151	434
	Max			3,348	596	4,252
Seattle		49%	19			
	Mean			473	580	656
	Median			192	312	420
	Max			5,695	2,853	3,434
Washington, DC		58%	114			
	Mean			803	1,124	1,533
	Median			499	613	873
	Max			3,382	8,175	9,141

At the ZIP code level, the average level of HHI across all seven cities is 1,252. However, weighting by unit counts results in a weighted mean HHI of 570. Areas with fewer housing units tend to have higher levels of concentration, see Figure 3.1. This is not to say that there are no large ZIP codes with high levels of concentration; the average ZIP code across these seven cities has approximately 4,200 rental units; the average ZIP code with an HHI above 1,500 has 1,600 rental

[30] Because of how elementary school attendance assignments work in Boston, there are only two schools for which data is available.

units. Nonetheless, three cities have average HHI at the ZIP code level that are at or near levels of moderate concentration: Boston, Dallas, and Kansas City.

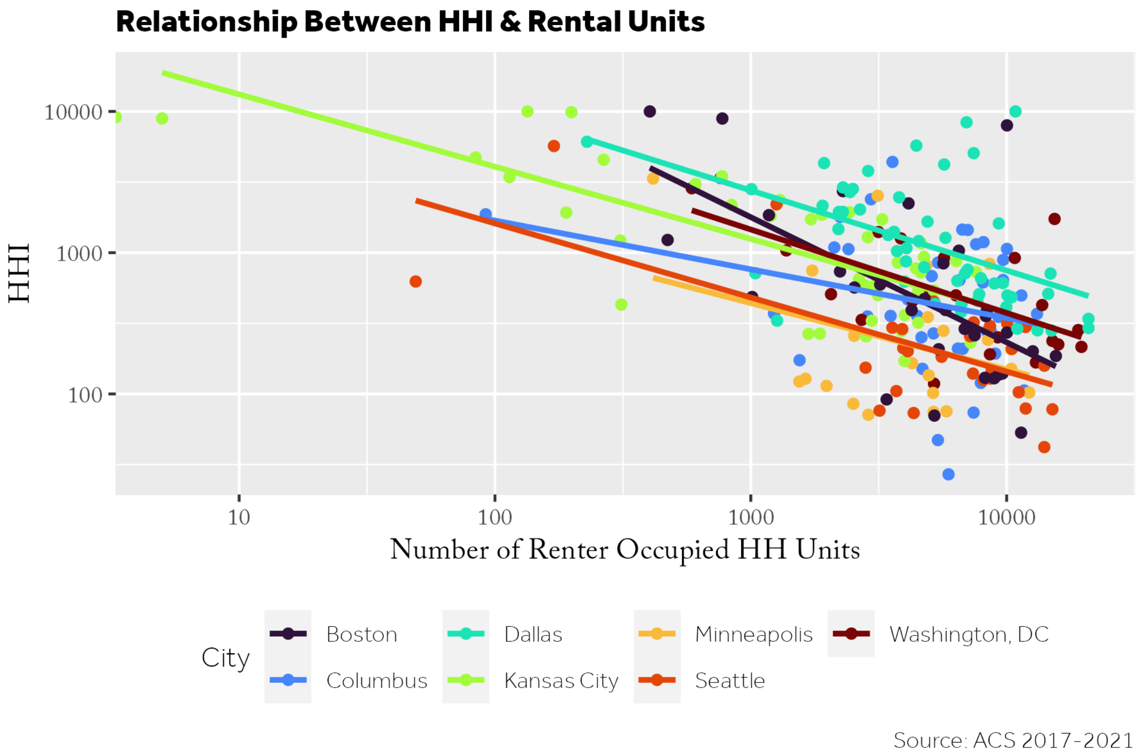

Figure 3.1: Relationship Between Housing Unit Counts and HHI at the ZIP Code Level

At smaller levels of geography, there are higher levels of both average and maximum concentration. Taking a weighted mean of the HHI at the tract level across my seven cities,[31] I find that the average HHI at the census tract level is approximately 1,800. Columbus, Dallas, Kansas City, and Washington, DC all have average tract-level HHIs above 1,500, a "moderately high" level of concentration. Seattle and Minneapolis both have low average HHIs at the Census Tract level, approximately 626 and 656, respectively. Watson and Ziv (2022) found the average HHI at the census tract level among multifamily buildings in New York City to be 2,000, while Linger, Singer, and Tatos (2022) found the average HHI at the census tract level to be approximately 600 across all of Florida, so my estimates seem to be in line with the existing literature. It is no surprise that smaller geographies have higher levels of concentration, as there are fewer rental units, so large owners can control a larger market share. Figure 3.2 displays the median HHI across the three sub-

[31] I weight each tract by the number of units found in that tract, per the rental registry.

city geographies, repeating information from Table 3.2 in a visual format. Here, it is easier to see that median HHIs are generally correlated across geographies at the city level, with Boston's and Columbus's school zones being obvious exceptions.

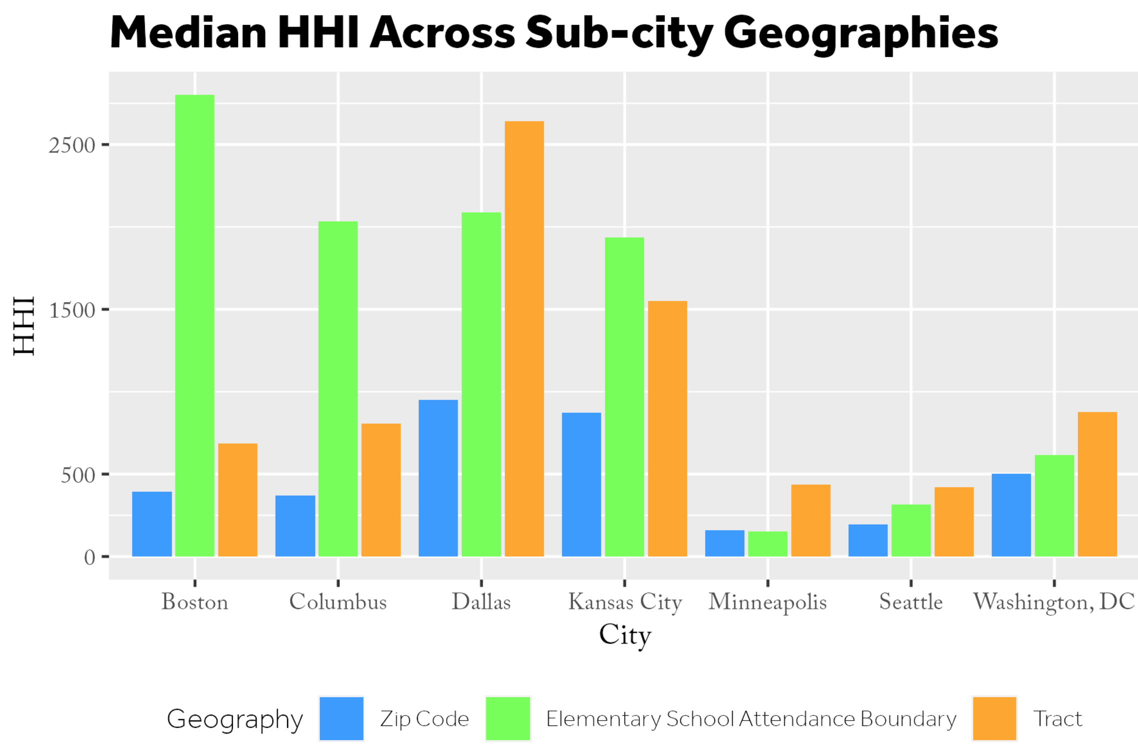

Figure 3.2: Median HHI Across Sub-city Geographies

There are some notable outliers related to geographic scale. In Columbus, Boston, and Dallas, there is at least one census tract with an HHI of 10,000 — the maximum possible HHI indicating that a singular landlord controls the entire housing stock. The singular ZIP code in Kansas City that has an HHI of 10,000 is a ZIP code that is predominantly industrial, but has an apartment complex that sits on its edge. The ZIP code in Boston that has an HHI of 10,000 is exclusively the campus of Harvard Business School, where two apartment buildings serving students are the only housing options.

The level of concentration among elementary school attendance zones is somewhat concerning.[32] In Dallas, Columbus, and Kansas City, the average elementary school attendance zone has an HHI that would be considered "moderately" or "highly" concentrated by the Department of Justice. In practice, that means that near Rhoads Learning Center in Dallas — whose attendance boundary has an HHI of 5,229 — the largest landlord controls 72 percent of the 527 rental housing units, while the ten largest landlords control 85 percent of the housing stock.

Both Table 3.2, above, and Table 3.3, below, show us that there is not a strong relationship between the percentage of the housing stock that is owned by corporate landlords and the level of concentration. The three cities with the highest average concentration at the ZIP code level — Boston, Dallas, and Kansas City — range significantly in the percentage of rental units that are owned by corporate landlords. At the ZIP code level, Columbus, Kansas City, Minneapolis, and Washington, DC have a positive correlation between the fraction of rental units that are owned by corporate landlords and the HHI at the ZIP code level, but the correlation is not particularly strong. In Boston, Dallas, and Seattle, there is a negative correlation. See Table 3.3.

Table 3.3: Correlation Between Corporate Ownership and HHI, ZIP Code Level

City	Correlation
Boston	-0.43466
Columbus	0.168717
Dallas	-0.17847
Kansas City	0.542489
Minneapolis	0.236708
Seattle	-0.33279
Washington, DC	0.159755

§3.4.2 — Cross-Sectional Regression Results

The results from the first set of regressions can be found in Table 3.4. The coefficient of interest is the top one displayed in the table, relating the log of the HHI to the log of the average rent. Model I uses ZORI data. Model II uses data from Rentometer, while Model III uses data on median rent

[32] Excluding Boston, the average elementary school attendance zone has 2,500 rental housing units across the remaining six cities. This is larger than a census tract — which has an average of 850 rental housing units, but smaller than a ZIP code — which has an average of 4,200 rental housing units.

from the American Community Survey. All three models are restricted to the 194 ZIP codes for which ZORI has data. All three models include city-level fixed effects.

Table 3.4: Cross Sectional Regression, ZORI ZIP Codes

	Model I: log(Asking Rent, ZORI)	Model II: log(Asking Rent, Rentometer)	Model III: log(Median Rent, ACS)
Log(HHI)	0.016*	0.014*	0.003
	(0.009)	(0.008)	(0.016)
Frac. White	-0.025	-0.075***	-0.030
	(0.040)	(0.025)	(0.079)
Frac. 25+ w/ Bachelor Degree	0.089	0.019	0.346***
	(0.058)	(0.079)	(0.052)
Log(Median Household Income)	0.120	0.252***	0.503***
	(0.104)	(0.071)	(0.027)
Log(Number of Renter-Occupied Housing Units)	-0.050*	-0.030	-0.007
	(0.027)	(0.020)	(0.021)
Log(Distance to CBD)	-0.045**	-0.053*	-0.012
	(0.020)	(0.031)	(0.014)
Frac. Vacant Housing Units	-0.520*	-0.073	0.158
	(0.271)	(0.197)	(0.379)
Frac. Owner-Occupied Housing Units	-0.066	0.041	-0.332***
	(0.091)	(0.185)	(0.035)
Log(Median House Value)	0.048	0.136***	-0.143***
	(0.062)	(0.038)	(0.028)
Constant	6.422***	3.689***	3.694***
	(0.749)	(0.680)	(0.301)
N	194	194	194
R^2	0.850	0.916	0.923
Adjusted R^2	0.837	0.909	0.916
Residual Std. Error (df = 178)	0.132	0.108	0.088
F Statistic (df = 15; 178)	67.161***	130.016***	142.180***

Note: ***p < .01; **p < .05; *p < .1
All Models Include City-Level Fixed Effects
Standard Errors Clustered At the City Level

Among the ZIP codes for which ZORI has rent levels, Models I and II agree remarkably well on the relationship between rental concentration and asking rents, with a 1 percent increase in HHI associated with a 0.016-0.019 percent increase in asking rents, statistically significant at the p < 0.1 level.[33] While this number sounds quite small, it translates into a meaningful economic difference in real terms. As an example, consider Minneapolis, where the typical asking rent in 2022 was approximately $1,600[34] and the average ZIP code level concentration is 500. If that ZIP code were instead to have an HHI of 2,500 — the level that is considered "highly concentrated," yet far below the most concentrated ZIP code in the city — we would expect the asking rent to be between $90 and $104 higher per month, or between $1,080 and $1,248 per year.

As is expected, the relationship between the median gross paid rent and the level of concentration is near zero, as is shown in Model III. We would expect this number to be smaller than the asking rent for two reasons. First, because the ACS only provides data at this level of geography for the five-year samples, we are looking at the relationship between concentration in 2022 and the actual paid rents in 2017 - 2021. If concentration is higher now than it was in the past, then the effect of concentration wouldn't yet be present in the data. Second, gross paid rent and asking rents differ in substantive ways (Boeing, Wegmann, and Jiao 2020): contract rent is the result of a negotiation between landlord and tenant, and may be different than the advertised asking rent; gross rent includes utilities while asking rents often do not; and landlords may raise rent differently for new as compared to existing tenants (Shiffer–Sebba 2020).

Because Zillow only releases information on a restricted sample of ZIP codes, I rerun Models II and III for all ZIP codes with data available from Rentometer (Model IV) and the ACS (Model V), shown in Table 3.5. Model IV shows a larger effect in all ZIP codes than in the ZIP codes that ZORI identifies. This may be because of the fact that Zillow is missing many ZIP codes with few rentals (see Appendix B1), and HHI and the number of rentals are negatively correlated (See Figure 3.1). Thus, across all ZIP codes, we see that, for every 1 percent increase in HHI, we would expect a 0.022 increase in asking rents, based on Rentometer data. Again, in economic terms, this means that in Washington, DC, moving from the typical asking rent ($2,450) in a ZIP code with an average HHI (800) to a ZIP code with an HHI of 2,500, we would expect monthly asking

[33] This estimate is more conservative than that of Watson and Ziv (2022), who find, at the census tract level, that a 10% increase in concentration results in a 0.5% increase in rent.

[34] Typical asking rent is also provided by Zillow at the city level, the source of this estimate.

rents to be $112 higher, or $1,345 per year. As in Model III in Table 3.4, Model V in Table 3.5, using median gross rent from the ACS, produces an estimate near zero without statistical significance.

Table 3.5: Cross Sectional Regression, All ZIP Codes

	Model IV: log(Asking Rent, Rentometer)	Model V: log(Median Rent, ACS)
Log(HHI)	0.022**	0.009
	(0.011)	(0.012)
Frac. White	-0.023	-0.067
	(0.066)	(0.074)
Frac. 25+ w/ Bachelor Degree	0.094	0.381***
	(0.107)	(0.063)
Log(Median Household Income)	0.214***	0.469***
	(0.073)	(0.030)
Log(Number of Renter-Occupied Housing Units)	-0.025	-0.006
	(0.020)	(0.012)
Log(Distance to CBD)	-0.061*	-0.021
	(0.031)	(0.013)
Frac. Vacant Housing Units	0.087	-0.006
	(0.102)	(0.521)
Frac. Owner-Occupied Housing Units	0.127	-0.273***
	(0.185)	(0.047)
Log(Median House Value)	0.099**	-0.113***
	(0.046)	(0.029)
Constant	4.404***	3.672***
	(0.620)	(0.371)
N	226	227
R^2	0.905	0.907
Adjusted R^2	0.898	0.900
Residual Std. Error (df = 178)	0.115 (df = 210)	0.100 (df = 211)
F Statistic (df = 15; 178)	133.327*** (df = 15; 210)	136.403*** (df = 15; 211)
Note:	***p < .01; **p < .05; *p < .1 All Models Include City-Level Fixed Effects Standard Errors Clustered At the City Level	

§3.4.3 — Income Shock Results

Table 3.6 shows the results from the third set of regressions. Models VI and VII both use change in ZORI asking rent. Model VI shows the simplest regression, where the only regressors are the expected income shock and the HHI. Model VII includes a suite of covariates, while Model VIII uses change in ACS median gross paid rent, as compared to gross asking rent. Models VI, VII, and VIII are all restricted to the 104 ZIPs for which ZORI had data in both 2014 and 2022, while Model IX uses all ZIPs for which the ACS has data. Because Model VIII and Model IX use five-year ACS estimates for the median gross rent, the income estimates were also generated using the five-year ACS estimates.

In order to generate the expected income growth, I take the share of workers[35] working in each industry[36] in each MSA in 2014, as well as the average income of those workers. For each industry-MSA pair, I calculate the expected income growth based on national income growth in each industry, *excluding* the income growth in a given MSA. I then calculate the MSA-wide expected income growth, $\widehat{incomegrowth}_c$, by multiplying the population share working in each industry in 2014 by the national income growth for each industry. This produces an estimate for the income growth in each MSA. Because I calculate expected industry growth for each industry-MSA pair by excluding the MSA itself, I am able to exclude possible local shocks that may affect both wages and rent levels. This shift-share instrument is a common approach in the urban economics literature for identifying causal impacts of local shocks (Bartik 1991; Baum-Snow and Ferreira 2015).

[35] Workers here are defined as individuals between the ages of 25 and 55 working at least 30 hours per week and at least 48 weeks per year.

[36] Industry defined at the 3-digit SIC code level.

Table 3.6: Income Shock Regressions

	Rent Change			
	Model VI: Delta ZORI	Model VII: Delta ZORI	Model VIII: Delta ACS	Model IX: Delta ACS
Log(HHI)	-2.166	-0.784	0.287	-1.794***
	(1.387)	(0.889)	(1.517)	(0.559)
Log(HHI)*Expected Income Growth	9.361	3.356	-1.213	8.230***
	(6.006)	(3.871)	(6.894)	(2.550)
Expected Income Growth	-55.280	-16.361	-27.814	-77.639***
	(46.085)	(29.048)	(31.996)	(9.974)
Frac. White		0.210	0.072	-0.046
		(0.157)	(0.130)	(0.104)
Frac. 25+ w/ Bachelor Degree		-0.636***	-0.440***	-0.507***
		(0.168)	(0.134)	(0.190)
Log(Median Household Income)		-0.048	0.155*	0.146
		(0.093)	(0.089)	(0.098)
Log(Number of Renter-Occupied Housing Units)		-0.028	0.029	0.002
		(0.038)	(0.032)	(0.011)
Log(Distance to CBD)		0.118***	-0.001	-0.029
		(0.032)	(0.021)	(0.022)
Frac. Vacant Housing Units		-0.046	0.027	-0.211
		(0.459)	(0.362)	(0.240)
Frac. Owner-Occupied Housing Units		-0.042	-0.276	-0.300**
		(0.167)	(0.184)	(0.125)
Log(Median House Value)		0.022	0.123***	0.160***
		(0.061)	(0.033)	(0.060)
Constant	13.195	4.837	3.166	14.173***
	(10.531)	(6.395)	(7.292)	(1.865)
N	104	104	104	228
R^2	0.037	0.607	0.443	0.353
Adjusted R^2	0.008	0.560	0.376	0.320
Residual Std. Error	0.212 (df = 100)	0.141 (df = 92)	0.123 (df = 92)	0.151 (df = 216)
F Statistic	1.280 (df = 3; 100)	12.920*** (df = 11; 92)	6.652*** (df = 11; 92)	10.694*** (df = 11; 216)
Note:	***p < .01; **p < .05; *p < .1 Standard Errors Clustered At the City Level			

The coefficient of interest is the interaction term Log(HHI)*Expected Income Growth. Models VI and VII both show coefficients in the expected direction, but are lacking in statistical significance. ZORI rent data is sparse — and is more sparse the further back in time one tries to observe: there are only 104 observations in these two regressions, compared to 230 ZIP codes overall. This likely contributes to the lack of statistical significance. Model VIII uses the ACS data, restricted only to the ZIP codes with ZORI data, and shows a negative coefficient that is not statistically significant. Model IX, which has more complete observations, has a similar magnitude to models VI and VII, but is statistically significant at the 0.01 level.

To test the robustness of these regressions, I rerun models VII, VIII, and IX three times, interacting the estimated income shock with the three covariates that are statistically significant in Model IX: fraction of college graduates, owner occupancy rate, and the house value. The full results of these regressions can be found in Appendix B3. In general, the magnitude, direction, and statistical significance of the main coefficient of interest remains similar, even with the inclusion of the additional interaction term. However, when I rerun Model IX with an interaction for the median house value, the coefficient of interest drops to near zero and loses statistical significance.

As in the previous section, I will compare Models VII and IX to illustrate the substantive meaning of the coefficients. A one standard deviation change in nominal expected income growth is between 0.56 percent (using the 1-year ACS estimates) and 0.86 percent (using the 5-year ACS estimates). Consider two ZIP codes in a city that experiences a one standard deviation expected income shock — one at a median concentration level of 500 HHI and one at a concentration level of 1,500 HHI, corresponding to approximately the 75th percentile among the ZIPs in this paper. Based on Model VII and the 1-year ACS estimates, we would expect the ZIP code with a moderate concentration level to have an increase in asking rents of approximately 2 percent larger than the neighborhood with a low concentration level according to Model VII.[37] Based on Model IX and the 5-year ACS estimates, we would similarly expect the ZIP code with a moderate concentration

[37] This is obtained by: $log(1500/500) \times 3.356 \times 0.0056$

level to have median gross that increased by 7.7 percent more than a similarly situated ZIP code at a low level of concentration.[38]

It is somewhat surprising that Model IX produces a larger estimate for the coefficient of interest than Model VII, since we would expect asking rents to be more sensitive to wage shocks than gross paid rents. The difference between models VIII and IX may indicate that the interaction between income shocks and concentration may be more pronounced in the areas where Zillow is lacking data: those with fewer rental units, farther from the city center. Zillow also has better coverage in some cities (Boston, Seattle, and Washington, DC) and worse coverage in others (Columbus, Kansas City, and Minneapolis), which may impact the findings (see Appendix B1). Further research is necessary, with more complete data sources, to disentangle the differences between geography and types of reported rents.

§3.5 — Discussion and Conclusion

This study aims to examine the relationship between concentration of rental property ownership at the city level and rents. By taking a novel source of rental ownership information — namely, rental registries — and identifying groups of rental property owners who appear to be distinct in the data, I am able to characterize the ownership concentration of rental properties at the local level within the formal rental market. The use of a probabilistic matching process that incorporates information about the officers and owners of corporate landlords, in addition to registrant information like email addresses and cell phones, results in a reduction in the total number of "unique" landlords of between 7 percent and 25 percent, compared to only grouping landlords based on exact matches of names and mailing addresses.

At the present time, the cities in this study do not have concentrated rental markets. While some cities do have individual landlords that control a sizable portion of the overall rental market, the low levels of HHI indicate that few owners control more than 1 percent of the rental market

[38] This is obtained by: $log(1500/500) \times 8.230 \times 0.0086$

Though the coefficient for the interaction term in Model IX is approximately double that of the same coefficient in Model VII, the differences are compounded by a larger standard deviation in the expected income growth in the 5-year ACS data.

city-wide. Nonetheless, there are higher levels of concentration at the neighborhood scale that raise questions about the extent to which rental submarkets are competitive. In all cities examined, there is at least one ZIP code, one elementary school attendance boundary, and one census tract with HHI levels consistent with "highly concentrated" markets.[39] While these are administrative boundaries, academic research on the importance of schools and neighborhoods (Owens 2020) and housing submarkets (Goodman and Thibodeau 2003), demonstrate their importance within the housing search process. In general, Minneapolis and Seattle appear to have the lowest level of overall concentration across all geographies.

This study can only speculate on the causes of the differences in levels of concentration among these seven cities. At sub-city levels of geography, some cities show a correlation between average HHI and the fraction of rental units that are owned by corporate entities. As prior studies have noted (Cherry and Ford Jr. 1975; Gilderbloom and Appelbaum 1987), one would expect that the professionalization of the landlord industry — indicated here by fewer sole proprietorships and more corporate landlords — would decrease competition in the housing market. It is also conceivable that more professionalized landlords could use software tools to help them set rent (Vogell, Coryne, and Little 2022), relating higher levels of professionalization to higher levels of rent.

Indeed, the cross-sectional regression results indicate that higher levels of concentration at the ZIP code level are correlated with higher levels of rent, compared to similar neighborhoods within the same city. The coefficients are in line with other studies, including those that have used instrumental variables to attempt a causal explanation for the relationship between concentration and rent levels (Linger, Singer, and Tatos 2022). While a 1 percent increase in HHI only causes a 0.016 - 0.022 percent increase in asking rent levels, the large range in ZIP code HHIs means that this can translate to substantial real differences in rent. Across the seven cities studied, the maximum HHI at the ZIP code level is, on average, 6,000 percent higher than the minimum HHI.

[39] With one exception: Minneapolis does not have a highly concentrated elementary school attendance zone.

Thus, going from the average minimum HHI to the average maximum HHI would result in an expected difference in asking rent of between 98 percent and 130 percent.

The results of the income shock regressions also indicate that higher levels of concentration may result in larger increases in rents when a city experiences a city-wide increase in wages. While the regression models that use asking rent data from Zillow did not produce statistically significant estimates, the model using median gross rent from the American Community Survey resulted in statistically significant results with fairly similar coefficients. It appears that, when landlords raise the rent because incomes in their city go up, landlords in neighborhoods with high levels of concentrated ownership raise the rent more than landlords in less concentrated neighborhoods. This preliminary indication of market power deserves further study with better data and in a broader cross section of cities. The differences in the income shock regressions may be due to the differences in the meaning of the dependent variable — asking rents as compared to paid rents. The differences may also be attributable to the neighborhoods in which Zillow data is available compared to the much broader coverage of the American Community Survey. Indeed, among the 20 ZIP codes with the highest levels of concentration among the seven cities in the study, Zillow only has data for 7 of these ZIP codes in 2014. The average HHI for rental ZIP codes for which Zillow has data in 2014 is only 698, compared to an average HHI of 1,244.

The findings here suffer from three limitations that future studies ought to address. The first two limitations relate to the quality of rent data. First, ZORI asking rent data is available for a selective set of geographies. While Zillow data is disproportionally available in larger cities and in ZIP codes with a greater number of renters, this selection of these rental heavy neighborhoods requires substantial qualification of the meaning of the results. Second, while prior research has found that ZIP codes provide a reasonable definition for housing submarkets (Goodman and Thibodeau 2003), there is reason to question whether ZIP codes represent a rental housing submarket where one would reasonably expect landlords to be able to exhibit pricing power, even if higher levels of concentration were found. Both of these limitations would be addressed with better data on rents and rental housing search tendencies. Rae and Sener (2016) obtained data from a data provider similar to Zillow in the UK, in which users searching for for-sale housing could draw the own boundaries of their search. In an ideal world, a repeat-rent index comprised of point locations of rental properties could be combined with a more flexible geography to examine the relationship between concentration and rent at geographies that better represent housing submarkets unmoored by administrative boundary definitions.

The third limitation results from the quality of the rental ownership and registration data. While the cities included in this study all appear to have fairly accurate coverage of their housing market, they do not cover the entire rental housing market. In most of these cities, somewhere between 10 and 20 percent of the rental market is not present in the rental registry. Landlords in the "informal" market may behave differently towards tenants, meaning that the effects of concentration may be skewed by only examining the formal property market (Samuel, Schwartz, and Tan 2021). Additionally, while care was taken to clean the data and remove information about rental property managers, it is likely that the steps taken both missed some management companies and improperly differentiated among entities that appeared distinct but are actually owned and managed by the same corporation. Under the letter of the law for many of these rental registries, management companies should have been a non-issue, as owners and managers are expected to be listed separately. Better enforcement regarding the scope of the rental registries, and the accuracy of the information provided to the city, would allow for greater certainty with regards to the findings of this study.

The study of the concentration in the rental ownership market has taken strides in recent years, owing to new data sources and new techniques to link different entities together, through machine learning, natural language processing, and the incorporation of a variety of data sources. Yet significantly more work needs to be done. The work of qualitative researchers on the practices of landlords has provided valuable insights into the functioning of rental markets at the scale of the transaction and the landlord (Garboden and Rosen 2022; Rosen, Garboden, and Cossyleon 2021). More work is needed to understand the relationship between rental management companies as well. When considering the effects of concentration on the rental market — whose concentration matters? Ought we differentiate between rental management companies and rental owners? Who decides rents? These are questions to which we do not currently know the answer. While this study chooses to ignore the role of management companies, future research could examine the relationship between concentration of rental management companies and rents. This would continue a thread of research that examines different ownership configurations (Rose and Harris 2021), and the effect of on-site management (Carswell 2018).

Additionally, this study only looks at one indicator of the effects of concentration: rent levels and increases. Hangen and O'Brien (2022) also look at the relationship between concentration and housing issues, and concentration and eviction rates. Future work should consider whether differing levels of ownership concentration can result in differences among housing outcomes other than rents. This may be especially important in cities where rent control or rent stabilization are in effect, perhaps limiting the effects of concentration on rents, but not on other measures. Among the cities considered in this study, Washington, DC is the only one with rent stabilization, though the exact number of units subject to rent stabilization is unknown.[40]

The current findings indicate that higher levels of concentration are associated with higher levels of rent and may contribute to larger increases in rent at the neighborhood scale. This research contributes to an ongoing literature studying the effects of corporate landlords on a range of tenant outcomes (Ellen, Harwood, and O'Regan 2022 ; Rose and Harris 2021; Travis 2019), the effects of institutional investors on tenant outcomes (Fields and Vergerio 2022; Immergluck et al. 2020; Raymond et al. 2018), the effects of rental concentration on rents (Gomory 2021; Linger, Singer, and Tatos 2022; Watson and Ziv 2022), and a methodological literature seeking to uncover ownership structures despite data limitations (Gomory 2021; Hangen and O'Brien 2022). Taken together, these studies indicate that ownership structure, landlord size, and ownership concentration can all affect tenant outcomes, including asking rents and eviction rates. The rental housing market in the United States lends itself to further scrutiny given these findings.

[40] Using Census data for age of structure (which is of questionable validity, see Molfino et al. 2017), I estimate that *no more than* 57% of the DC housing stock is subject to rent stabilization, but that number is likely much lower. Other estimates put the number between 38% and 50% (Austermuhle 2020).

CHAPTER 4 | IS THIS A RENTAL? COMPARING METHODS FOR IDENTIFYING RENTAL UNITS

Abstract

Researchers and practitioners alike regularly attempt to identify individual housing units as either owner-occupied or renter-occupied. But the data sources available to do so are rarely purpose-built for answering that question. While the US Census Bureau provides rental housing unit counts, it does not provide building-level information. This paper explores the most common approaches used in the literature to identify rental properties in the United States, namely by identifying properties based on characteristics listed within a tax assessment database. Repurposing tax assessment data has some major limitations. This study shows the possible problems associated with the current approaches undertaken in the literature to identify rental properties based on homestead exemptions or address matching. An underutilized data source — local rental registries — are introduced as a possible alternative in the cities that have them. Differences between rental registries and tax assessment databases are discussed, and the number, count, and type of rental units are compared at the city-wide and sub-geographic levels. I identify possible sources of disagreement between tax assessment databases and rental registries. I also suggest methods for validation between rental registries and tax assessment databases, where possible. This paper cautions researchers who opt to use tax assessment databases, or proprietary data sources, to identify rental units. At the same time, it recognizes the limited utility and scope of rental registries, and encourages broader implementation and stronger enforcement of rental registries.

§4.1 — Introduction

The United States has no central register of rental properties. If one wants to identify whether a property is owner-occupied or renter-occupied, researchers typically rely on local or proprietary data sources that provide parcel-level information. The property tax assessment database is the most typical source used to identify rental properties (see, e.g., Ferrer 2021; Freemark, Noble, and Su 2021; Gomory 2021; Hangen and O'Brien 2022; Rose and Harris 2021; Travis 2019), yet it is not purpose-built to identify rental properties. The use of proxy data sources to identify underlying phenomena requires us to investigate their accuracy. However, no study has assessed whether the literature has heretofore accurately distinguished between rental properties and other residential properties.

There are important questions that scholars are trying to answer about America's rental housing market today. Who owns rental properties in the US, and do different types of owners differ in how they treat tenants (Gomory 2021; Immergluck et al. 2020; Ellen, Harwood, and O'Regan 2022; Robinson and Steil 2021; Travis 2019)? Do some owners have concentrated holdings (Hangen and O'Brien 2022; Linger, Singer, and Tatos 2022; Tapp and Peiser 2022; Watson and Ziv 2022)? Where are corporate investors buying properties (Charles 2020b; Dowdall et al. 2021)? In order to correctly answer these questions, scholars need to accurately identify whether or not each property in their study is a rental or not. At best, inaccurate estimates means an undercount; at worst, misidentifying a rental or an owner-occupied property equates to mislabeling which units are in the treatment and control groups. Yet it is not clear that the data we have available is suitable to answer these questions.

If one wants to know how many rental properties there are in a given place, the first data source used is typically the American Community Survey (ACS), which provides a count of housing units by tenure type and geography: within a certain area, how many units are vacant, owner-occupied or renter-occupied. Similar information is also provided by the American Housing Survey, with greater detail about housing conditions but at a coarser geographic scale. These federal data sources are helpful if a researcher hopes to understand the general patterns of ownership and rental rates by geography. They are not helpful, however, in answering the following question: is *this particular* property owner-occupied or renter-occupied? Researchers often compare the number of rental properties identified by their data source to the number of rental properties reported by the ACS; yet similar numbers do not necessarily guarantee accurate identification.

In some cities, another data source exists: rental registries, in which the city requires landlords to register their rental properties and provide information about the properties themselves. While relatively few reports or studies have used rental registries (Coulton et al. 2020; Ellen, Harwood, and O'Regan 2022; Haider 2021; Watson and Ziv 2022), this paper demonstrates the usability of rental registries in cities that have them and in which they are well-implemented. It shows how using rental registries can lead to different estimates than the estimates from property tax assessment databases. By comparing estimates of rental properties from tax assessment databases and rental registry databases, this paper shows the areas of agreement and disagreement among different methods for identifying rental units. Additionally, it discusses differences among different rental registration ordinances that, in turn, lead to differences in rental registries' coverage extent. Rental registration databases represent an actual sample of the rental units within a jurisdiction, as compared to estimates derived from tax assessment databases, the latter of which represent the researcher's best estimate of the universe of rental properties. Both sources are imperfect, and can be complementary. Table 4.2 in Section 4.4 shows the differences between rental registries and tax assessment databases, while Table 4.3 in Section 4.5 discusses the pros and cons of different methods to identify rentals.

This paper provides an overview of existing studies that seek to identify rental units in order to study rental housing markets. I compare the different ways they identify rental units, differentiating between the data sources, the methods used on those data sources, and the subset of the residential properties. I examine the differences between tax assessment databases and rental registries: their different purposes, advantages, and limitations. I then use the two most common approaches from the literature — comparing mailing addresses and examining homestead exemptions — to identify rental properties in five cities: Columbus, OH; Minneapolis, MN; Nashville, TN; Philadelphia, PA; and Washington, DC. I compare these estimates to estimates from locally collected rental registries. These estimates are compared against the ACS estimates for number of rental units at the city level, first among all units, then specifically among single-family rentals. Additionally, I compare owner addresses between the rental registries and the tax assessment databases, finding some level of agreement, but far from universal. Finally, I investigate sub-city counts of single-family rentals in Philadelphia, finding little geographic variation among the different methods. I use Philadelphia as a case study because, among the cities studied here, it has the highest proportion of single-family rentals,[41] and its registry data coverage is rather extensive.

While rental registries are introduced as a possible data source for identifying rental units, they are currently only available in a limited number of cities. Of the 50 most populous US cities, 30 of them have rental registries, but I was only able to obtain registration information from 19 of them. Obtaining that information was sometimes easy — if posted on an open data portal — and sometimes impossible — if the city did not respond to the public records request. See Appendix A for more information about rental registries and the information collected. The coverage of the data varies widely due to differing exemptions and compliance: in some cities, it appears that almost all rentals are actually registered; in others, it is estimated that as few as 1 in 6 rental properties are actually registered with the city (Lynch 2022). These drawbacks, in addition to the benefits of rental registration databases, are discussed in Section 4.4. In the next section, I discuss the methodological literatures related to the use and accuracy of administrative and novel data. Section 4.3 discusses the differences among the literature studying rental housing, and how different researchers have tried to identify rental properties. Sections 4.5, 4.6, 4.7, and 4.8 compare the universe of rental properties that can be uncovered through tax assessment and rental registration approaches.

§4.2 — Why Compare Data Sources

Within the fields of policy and planning, it is rare for the researcher to have a perfect source of data that perfectly measures the construct of interest. Instead, we must either create our own data sources — through surveys and experiments — or we must rely on existing data sources that contain the information we seek. Unfortunately, the information we seek is rarely found precisely as we wish it to be. Instead, we use proxies: we model daily mobility through mobile phone traces (Saxon 2020; Yabe et al. 2023); we look for displacement through heightened rates of mobility or

[41] Approximately 40% of rental units in Philadelphia are in single-family attached or detached houses, according to the 2017-2021 ACS.

involuntary moves (Ellen and O'Regan 2011; Freeman et al. 2023); or we use credit scores to stand in for income (Daepp 2022; Ding, Hwang, and Divringi 2016). Some proxies are better than others, and without stepping back to examine the benefits and deficiencies of a data source, we can sometimes lose sight of the fact that we are using a proxy at all.

There is a wide literature concerned with assessing the accuracy of novel data sources. As "big data" entered the lexicon and the number of data sources and volume of data increased, scholars sought to understand whether big data would be more accurate and timely than traditional data sources. Indeed, some argued that because of the *velocity*, *variety*, and *volume* of big data, its *veracity* would naturally follow (Boeing and Waddell 2017, 468). Arribas-Bel (2014) argued that, though data quality would be an issue, it would be but another step in the research process to consider the problems and usefulness of big datasets. Yet others pointed out that some of these data sources would be "sticky:" good in the context for which they were created, and considerably worse when applied to new contexts (Offenhuber 2017). Limitations are apparent: Folch, Spielman, and Manduca (2018) compared local administrative data on restaurant locations to Yelp data, using the US Census Bureau County Business Pattern data as a ground truth. They found that though Yelp and the local data had similar counts of establishments, both under-counted the US Census, and only one-third of the restaurants in each set were actual matches.

On the use and accuracy of administrative data, some researchers point out the substantial uncertainty associated with American Community Survey tables at smaller geographic scales (Fowler et al. 2020; Spielman, Folch, and Nagle 2014; Spielman and Folch 2015). Scholars often take the point estimates from the ACS to be completely accurate; however, at smaller geographies, some tables from the ACS have margins of error that are as large as the point estimates themselves. Focusing on the local government level, researchers at the Urban Institute created a "bias assessment tool" for government open data, noting that the representativeness of many data sets showed a bias towards certain parts of the city (Narayanan and MacDonald 2019). Importantly for this context, the authors noted that both citizen-generated data such as 311 calls,[42] and data on

[42] Also known as Volunteered Geographic Information (VGI) in the geography context

local government service provision have a spatial bias within the city (see also, for instance Trounstine 2018, ch. 5).

Specifically in the housing context, scholars have tried to capture the limits and possibilities of using alternative data sources to study local housing markets. For instance, Boeing, Wegmann, and Jiao (2020) looked at the differences in rental price information between the ACS, Fair Market Rents from the US Department of Housing and Urban Development (HUD), recent-mover rents from AHS, and Craigslist nationwide. They found that the federal data sources often estimated significantly lower rents than Craigslist. In some cases, this was expected, given the difference between asking rents and contract rents, yet in other cases, such as in fair market rents, this difference presents a problem. Molfino et al. (2017) compared ACS data, local tax assessment databases, and a propriety data source to understand which data source seemed more accurate in understanding residential property value, property type, unit counts, bedroom counts, and property taxes. They found that the tax assessment data proved a better source in estimating year built, house value, and property taxes paid than the ACS. In an approach that is similar to this present study, Wilson et al. (2023) demonstrated that a commonly-used dataset among housing researchers — locations of Low Income Housing Tax Credit (LIHTC) units provided by the HUD — was often inaccurate. This study is situated among this literature, as it seeks to understand the accuracy and limits of using property tax assessment data and rental registries.

§4.3 — Different Methods for Defining and Identifying Rentals

Most researchers looking for unit-level information about rental properties in the United States have relied on property tax assessment databases, an example of which can be seen in Figure 4.1. However, there is no unified approach to their use. Table 4.1, below, shows the different samples and the different approaches that a variety of studies take to identify rental properties in the United States. This section is not exhaustive of all studies focused on the housing market, but illuminates the many different ways that a set of representative studies have taken to identify rentals. These approaches can be broadly distinguished by the different ways they identify rental properties:

comparing mailing addresses, relying on the presence of a homestead exemption,[43] relying on a proprietary data source, or relying on a rental registry.

Open Minneapolis

Showing 25 of 130,772 rows

PIN	ASMTYEAR	TAXYEAR	HOUSE_NO	STREET_NAME	UNIT_NO	ZIP_POSTAL	FORMATTED_ADDRESS	OWNERNM	TAXPAYER1	TAXPAYER2	TAXPAYER3
p0102824130046	2,022	2,023	3313	25TH AVE S		55406-2405	3313 25TH AVE S	ROBERT JOHN REICHGELD II	ROBERT JOHN REICHGELD II	3313 25TH AVE S	MINNEAPOLIS MN 55406
p0102824130047	2,022	2,023	3317	25TH AVE S		55406-2405	3317 25TH AVE S	JLMPW HOLDINGS LLC	JLMPW HOLDINGS LLC	MATTHEW WAGNER	3317 25TH AVE S
p0102824130048	2,022	2,023	3321	25TH AVE S		55406-2405	3321 25TH AVE S	CITY/LAKES COMMUNITY L...	JACKSON J ARCOREN	TERESA E ARCOREN	3321 25TH AVE S
p0102824130049	2,022	2,023	3323	25TH AVE S		55406-2405	3323 25TH AVE S	TIMOTHY GODSILL	TIMOTHY GODSILL	6011 XERXES AVE S	MINNEAPOLIS MN 55410
p0102824130050	2,022	2,023	3327	25TH AVE S		55406-2405	3327 25TH AVE S	BRANDON V CULBERSON	BRANDON V CULBERSON	3327 25TH AVE S	MINNEAPOLIS MN 55406
p0102824130051	2,022	2,023	3331	25TH AVE S		55406-2405	3331 25TH AVE S	TYLER E KUKOWSKI	TYLER E KUKOWSKI	3331 25TH AVE S	MINNEAPOLIS MN 55406
p0102824130052	2,022	2,023	3335	25TH AVE S		55406-2405	3335 25TH AVE S	GIDGET HOULE	GIDGET HOULE	3335 25TH AVE S	MINNEAPOLIS MN 55406
p0102824130053	2,022	2,023	3339	25TH AVE S		55406-2405	3339 25TH AVE S	KARLSON LIVING TRUST	DAVID KARLSON	BEVERLY KARLSON	3339 25TH AVE S
p0102824130057	2,022	2,023	3323	24TH AVE S		55406-2403	3323 24TH AVE S	DAVID C KRAML	DAVID C KRAML	3323 24TH AVE S	MINNEAPOLIS MN 55406
p0102824130058	2,022	2,023	3327	24TH AVE S		55406-2403	3327 24TH AVE S	JAMES R BERG	JAMES R BERG	3327 24TH AVE S	MINNEAPOLIS MN 55406
p0102824130059	2,022	2,023	3331	24TH AVE S		55406-2403	3331 24TH AVE S	D J LINDSAY & K V FISHER	DANIEL JOSEPH LINDSAY	KELLY VERA FISCHER	3331 24TH AVE S
p0102824130060	2,022	2,023	3333	24TH AVE S		55406-2403	3333 24TH AVE S	J L HELWEGEN & K A MCNEIL	K A MCNEIL & J L HELWEGEN	3333 24TH AVE S	MPLS MN 55406
p0102824130061	2,022	2,023	3339	24TH AVE S		55406-2403	3339 24TH AVE S	GOLDEN POINT ESTATE LLC	GOLDEN POINT ESTATE LLC	229 MINNETONKA AVE S #843	WAYZATA MN 55391
p0102824130062	2,022	2,023	3338	25TH AVE S		55406-2406	3338 25TH AVE S	A T SPLINTER & J L NAJERA	ANNE THERISA SPLINTER	JOSE LUIS NAJERA	3338 25TH AVE S
p0102824130063	2,022	2,023	3334	25TH AVE S		55406-2406	3334 25TH AVE S	E D GORTON & K A GORTON	ETHAN D GORTON	KATHERINE A GORTON	3334 25TH AVE S
p0102824130064	2,022	2,023	3330	25TH AVE S		55406-2406	3330 25TH AVE S	J M BOWKER & L BOWKER	JACQUELINE MARIE BOWKER	LYNDA BOWKER	3330 25TH AVE S
p0102824130065	2,022	2,023	3326	25TH AVE S		55406-2406	3326 25TH AVE S	BRETT ERICKSON	BRETT ERICKSON	3326 25TH AVE S	MINNEAPOLIS MN 55406
p0102824130066	2,022	2,023	3322	25TH AVE S		55406-2406	3322 25TH AVE S	DORA E KAUFMAN TRUST	DORA E KAUFMAN	3322 25TH AVE S	MINNEAPOLIS MN 55406
p0102824130067	2,022	2,023	3320	25TH AVE S		55406-2406	3320 25TH AVE S	TIM J EAGEN	TIM J EAGEN	8721 ST CROIX TR S	HASTINGS MN 55033
p0102824130068	2,022	2,023	3314	25TH AVE S		55406-2406	3314 25TH AVE S	JOE G ZEBRO	JOE G ZEBRO	3314 25TH AVE S	MINNEAPOLIS MN 55406

Figure 4.1: An Example Property Tax Assessment Database

Note: This screen capture comes from Minneapolis Open Data (https://opendata.minneapolismn.gov/). PIN is the Parcel Identification Number, ASMTYEAR is Assessment Year.

Scholars who rely on mailing addresses look at the difference between the property address and the mailing address of the owner — I call this the "mailing address" approach. Travis (2019), who studies the relationship between code complaints and property owner type, used Milwaukee tax assessment records. He defined any property that met the following criteria as a rental property: a "non-owner occupied" single-family home, duplex, or condominium, as well as any multi-family or mixed-used property with a non-owner-occupied residential unit. In turn, Milwaukee defined owner-occupancy based on a comparison between a property's mailing address and its property address. The number of rental units identified by Travis is comparable to the number of rental units enumerated by the ACS.

[43] A homestead exemption is a property tax reduction for a subset of owner-occupied housing units. These may apply to all owner-occupied units, or they may be restricted to a subset of units, such as those occupied by senior citizens. See Langley (2015) for greater detail on differences in homestead exemption implementation.

Similarly, Linger, Singer, and Tatos (2022), studying property markets in Florida, define any property as a rental property if it is a multi-family building or if the property address and the mailing address do not match. Rose and Harris (2021) do the same, but they appear to exclude properties with more than six units. While Rose and Harris do not report a comparison to the American Community Survey, I estimate that they only identify 74 percent of rental properties in Rochester. Raymond et al. (2018) limit their study exclusively to single-family rentals, by comparing mailing addresses and property addresses among homes with a land use code indicating a single-family property. In the Canadian context, St-Hilaire, Brunila, and Wachsmuth (2023) compare owner and building addresses to identify rental properties in Montreal. They analyzed their accuracy at the Census Tract level, finding a 0.94 correlation between the number of rental units identified by their method, as compared to the Canadian Census.

Both Hangen and O'Brien (2022) and Gomory (2021) rely on an "owner occupancy" status provided by the City of Boston. Unlike Milwaukee, as studied by Travis (2019), Boston's owner occupancy status is determined based on whether the homeowner has a residential exemption on their property taxes — this is the "homestead exemption" approach. Though both Hangen and O'Brien (2022) and Gomory (2021) study Boston during a similar time period, they use a somewhat different approach to exclude some properties: Hangen and O'Brien exclude condominiums, while Gomory excludes mixed-use residential and commercial properties. In the Los Angeles context, Ferrer (2021) also relies on the presence of a homestead exemption to determine whether a property is owner occupied or a rental property. Freemark, Noble, and Su (2021) study property ownership in Hennepin County, home to Minneapolis. Because most municipalities in Hennepin County do not have rental registries, they rely on a measure based on whether a property has a homestead exemption, but they validated their measure against rental registry data in Minneapolis. As discussed in Section 4.6, the use of homeowner exemptions to define rental properties and owner-occupied properties raises serious concerns.

Proprietary data sources to study rental housing markets often allows the researcher to move from a local case-study to a national sample. While the studies discussed thus far have all examined a specific city or state, research by Gurun et al. (2022), Mills, Molloy, and Zarutskie (2019), and Tapp and Peiser (2022) are all national in scope. These authors rely on data from Zillow

(Gurun et al. 2022), CoreLogic (Charles 2020b; Mills, Molloy, and Zarutskie 2019), and Real Capital Analytics (Tapp and Peiser 2022). These proprietary data sources are often *derived from* county tax assessor databases, but they are standardized and cleaned. The data provider may also provide additional information — for example, by incorporating information from the United States Postal Service, property transaction history, or property characteristics from a previous sale listing. However, as Molfino et al. (2017) found in their study comparing ACS data to CoreLogic data, when a national data provider standardizes and cleans local data sources, they also homogenize that data, losing some information that may change the accuracy of an individual county's data. This makes the use of a proprietary data source a trade-off between breadth, depth, and accuracy. Additionally, some data providers have a limited scope: Real Capital Analytics, used by Tapp and Peiser (2022), only contains information on multifamily buildings that transacted for more than $2 million.

Shelton (2021) and Stiman (2019) explore questions related to ownership of second homes, and their methods serve as an illuminating comparison to the other papers examined here.[44] Both authors define second homes using property tax assessment databases, where a property owner who owns only one home with a mailing address outside of the city is defined as an owner of a "second" home, rather than a landlord. While Stiman augments her quantitative approach with interviews and survey methods, Shelton relies on his unique case site — Starkville, MS, a college town home to Mississippi State University — combined with Census Data, to differentiate between second homes and absentee owners.

Finally, four articles and reports have relied on rental registries to identify rental properties. In Philadelphia, Haider (2021) identified rental properties as those with an active rental license — or those with an expired rental license where the owner had not changed since the rental license expired. Though they estimated that 30 percent of units in Philadelphia were not currently licensed, they were able to identify approximately the same number of rental units as estimated by the ACS because of the inclusion of expired units. In Cleveland, Coulton et al. (2020) combined the local rental registry with tax assessment data to identify rental properties as those that either were

[44] I exclude them from Table 4.1, however, because they do not explicitly try to identify rental properties.

registered with the city or did not claim a homestead exemption. Both Watson and Ziv (2022) and Ellen, Harwood, and O'Regan (2022) used the New York City rental registry to study the rental property market there; Watson and Ziv (2022) study property ownership concentration, and Ellen, Harwood, and O'Regan (2022) corporate ownership. However, New York City's ordinance only requires registration from property owners with three or more units, limiting the scope of their studies to multifamily buildings of a certain size.[45]

Table 4.1: Different Methods and Different Samples of Rental Properties

	Mailing Address Approach	**Homestead Exemption Approach**	**Proprietary Data Source**	**Rental Registry**
All Rental Properties	Travis (2019); Linger, Singer, and Tatos (2022); St-Hilaire, Brunila, and Wachsmuth (2023)	Ferrer (2021); Coulton et al. (2020)		Haider (2021); Coulton et al. (2020)
Single Family Only	Raymond et al. (2018)	Freemark, Noble, and Su (2021)	Gurun et al. (2022); Charles (2020b); Mills, Molloy, and Zarutskie (2019)	
Single Family & Small Multifamily	Rose and Harris (2021)			
Multifamily Only			Tapp and Peiser (2022)	Ellen, Harwood, and O'Regan (2022); Watson and Ziv (2022)
Excludes Condominiums		Hangen and O'Brien (2022)		
Excludes Mixed-Use Buildings		Gomory (2021)		

§4.4 — Rental Registration Ordinances and Tax Assessment Databases

Rental registration ordinances, landlord licenses, proactive inspections, and other local government interventions in the rental market are currently on the rise in the United States. These ordinances have been heralded by the press (Demsas 2021), tenant activist groups (Beaty and Shankar 2022), and researchers (Dowdall et al. 2021) as a first step in the solution towards providing tenant relief — such as through the COVID-19 Emergency Rental Assistance Program — or to identify

[45] Ellen, Harwood, and O'Regan (2022) also combine the rental registry data with data on landlords who participate in the Housing Choice Voucher Program.

negligent or absentee landlords. Yet because these ordinances try to solve different problems, the data collected by local governments about their local rental markets vary widely. This heterogeneity in rental ordinance purpose, scope, and availability, results in significantly uneven coverage and effectiveness of rental registration ordinances (RROs).

On the other hand, tax assessment databases are widely available throughout the United States. Typically administered at the county level (though occasionally at a different level of government) tax assessment databases consist of a complete universe of real property within a jurisdiction. Such databases often provide information about the land uses[46] of a parcel, its ownership, and its assessed value. This section provides an overview of the different ways to use tax assessment databases and rental registries in order to identify residential rental properties. Table 4.2 outlines some of the differences, benefits, and limitations to using property tax assessment databases and data from RROs for identifying rental properties.

Table 4.2: Comparison Between Tax Assessment Databases and Rental Registries for Identifying Rental Properties

	Tax Assessment Databases	**Rental Registries**
Purpose	To assess the value of real property in the jurisdiction and determine the amount of property tax owed	To identify the location of rental properties and contact information of landlords, in pursuit of local government health and safety duties
Implementing entity	County Tax Assessor or Auditor	City Department of Licenses, City Housing Department, or City Buildings Department
Properties Included	All real property in a county	Properties that are rental-occupied or vacant for-rent. Often excludes properties that are occupied by the owner or owner's family. May be limited to only multi-family properties
Temporal Resolution	Annually	Annually or more frequently
Key Information Provided	Property location, assessed value, and use	Property location, owner contact information, management contact information
Unit Counts	Sometimes	Almost always
Ease of access	Easy. Publicly available or available for purchase	Difficult. Occasionally publicly available, often requires a right-to-know request.
Major Limitations	Data quality on mailing addresses, homestead exemptions, and unit counts may be poor. May require the exclusion of mixed-use, condominium, and cooperative buildings.	Data quality on mailing addresses may be poor. May have poor enforcement, limiting the sample of rental properties. City governments may or may not be willing to share all or any of what they collect.

[46] Such as "Residential," "Commercial," "Mixed Use," "Industrial," etc.

§4.4.1 — Tax Assessment Databases

Property tax assessment databases are typically created by the county Assessor's office, based on a combination of data from the county Recorder of Deeds and local tax assessors.[47] Property ownership information is recorded at the county level, at time of deed transfer. The purpose of a property tax assessment database is to provide information on the assessed value of real property within the taxing jurisdiction; all additional information collected by the assessor *may* support that function, but the presence and accuracy of other fields differs state-to-state and county-to-county. Additionally, not all states require re-inspection of the property after the initial assessment, meaning that details on the condition or occupancy of the property may not be accurate. In extreme cases, mailing addresses may be completely missing from tax assessment databases: in Philadelphia, PA, the tax assessor does not report the mailing address for 64 percent of residential properties.[48] (For an in-depth history on the property tax in the United States in general, though with a focus on Kansas, including its administration at the local level, see Fisher 1996).

One other facet of the property tax is worth mentioning: the homestead or homeowner exemption. The property tax is a strange tax: though the property is purchased only once, the tax is paid annually for the duration of that property holding.[49] Some homeowners can afford the purchase price of the property, but cannot easily pay continued property taxes on the property, especially as the assessed value of the property rises. This is why many states and localities have homestead exemptions, wherein a homeowner (homesteader) is exempted from part of their property tax bill. Langley (2015) provides an overview of homestead exemptions: 26 states provide a homestead exemption to nearly all owner-occupiers; 18 states provide special homestead

---

[47] Because counties and local governments are "creatures of the state" there is significant heterogeneity in the way each state, and in turn, each local county and local government, administers the property tax. This section is true to the typical case, but see "State-by-State Property Tax at a Glance" (2022) for a detailed state-by-state guide.

[48] In Philadelphia, the default mailing address is the property address, which likely explains the high rate of missing data: it can be assumed if the mailing address is missing, the property address is the mailing address. However, this is not universally true, even within Philadelphia: a manual check against the city property tax website shows that some properties with missing mailing addresses in the database have a mailing address online, which differs from the property address.

[49] The property tax is separate from an additional tax that is sometimes assessed on the value of the purchase price, known as a transfer tax, and paid at the time of transfer.

exemptions to seniors; and 31 states provide homestead exemptions to veterans.[50] These statewide exemption requirements exist in addition to locally-mandated homestead exemptions: all together, Langley (2015) estimates that there are 185 different homeowner property tax relief programs in the United States, the majority of which are locally funded. This means that, state to state or even city-to-city, the "homestead exemption" in a property tax assessment database does not mean the same thing. In Minneapolis, the homestead exemption applies to properties that are owned by individuals and that are occupied by themselves or a family member, but not if they own the property in the form of a trust or corporation, and not if their property is worth more than $413,000.[51] In Tennessee, however, the homestead exemption is limited to seniors or those with disabilities with an income below $24,000, or disabled veterans.[52]

Using tax assessment databases — either via the homestead exemption approach or the mailing address approach — is not particularly precise: not all homeowners who are eligible to take a homestead exemption take them: a study of neighborhood-level take-up rates of the homestead exemption in Florida found that, while 90% of homeowners apply for a homestead exemption, take-up rates were substantially lower in neighborhoods where incomes are lower and where residents are non-white (Ihlanfeldt 2021). When using mailing addresses to identify rental properties, the researcher runs the risk of identifying a rental property that is actually a second home (Shelton 2021; Stiman 2019), a home currently being "flipped," or a home where property tax bills are sent to a financial manager. On the other hand, rentals where the owner sends the property tax bill to the property would not be flagged as a rental — a practice that may be more common among landlords living abroad (Sakong 2021). Both the homestead exemption approach and the mailing address approach run the risk of identifying vacant, not-for-rent properties as rentals. Nationally, 11 percent of housing units are vacant, though that number varies substantially: in

---

[50] These are not mutually exclusive groupings of states.

[51] See https://www.minneapolismn.gov/resident-services/property-housing/property-values-taxes/homestead/regular-relative/

[52] See https://www.lincolninst.edu/research-data/data-toolkits/significant-features-property-tax/access-property-tax-database/property-tax-in-detail?field_tax_state_tid=6531&field_tax_year_tid=7036#gwipp_residential_tax_programs

Seattle, 7 percent of units are vacant, compared to 16 percent in Baltimore.[53] Though a *plurality* of these vacancies are vacant for-rent homes, many of these vacancies are vacant for-sale, vacant for "seasonal, recreational, or occasional use," or "vacant for other reasons."

Assessment databases also may not include crucial information about rental properties, such as the number of units in a building. Indeed: Regrid, a proprietary data source of tax assessment data (used in this paper) provides unit counts for only 159 counties out of the 3,215 for which they have other data. In Massachusetts, the state database of property tax records reports that only 50 percent of local governments collect unit counts for residential parcels. These three challenges with using property tax assessment databases: researchers' inaccurate exclusion of some rental properties, researchers' inaccurate inclusion of some non-rental properties, and property tax assessments' missing property-level information means that researchers should be cautious before using tax assessment databases (or propriety databases based on tax assessment information) to study rental housing markets. This is especially true for studies that aim to understand unit-level effects, like the relationship between ownership concentration and rent.

§4.4.2 — Rental Registration Ordinances

The need for rental registration ordinances (RROs) arises due the limitations of using tax assessment databases to track rental properties and landlords. RROs also provide additional information beyond what is collected by a tax assessment database, such as owner phone number and email addresses. Additionally, ownership is only part of the rental housing story in modern day America. With larger landlords who can more easily operate from a distance, management companies play an important role in the modern renting experience. Ownership information is not enough to understand the current rental market landscape. This poses a challenge for both RROs and property tax assessment data. Many property owners will list their management company for their mailing address with the tax assessor. While RROs typically differentiate between owner address and management company address, compliance with this requirement is uneven among

[53] These numbers come from the 2017-2021 ACS. The ACS differentiates between different types of vacancies, allocating vacant properties into the following categories: for rent; rented, not occupied; for sale; sold, not occupied; for seasonal use; for migrant workers; and other vacant. In the cities in this study, "vacant, for rent" and "other vacant" are among the top two types of vacancy.

different local governments and different management companies: some diligently differentiate between owner and manager information, while other owners provide manager information even when they should not.

In the cities that have them, RROs can fill the gaps left behind by using only assessment data to understand the extent of rental properties in the United States. As compared to tax assessment databases, then, RROs are significantly more limited in availability. Additionally, it is quite likely that not all cities within a metropolitan area would have a rental registry, limiting the utility of RROs to the city boundaries, rather than the county-level data provided by a tax assessment database. Even among cities that have RROs, they may not share them publicly; researchers may need to file a local right-to-know request, which cities may or may not fulfill. Depending on the jurisdiction, the researcher may or may not have recourse should the right to know request be denied. Figure 4.2 shows an example rental registry from Minneapolis.

Open Minneapolis

Showing 25 of 22,936 rows

apn	licenseNumber	category	milestone	tier	status	issueDate	expirationDate	address	ownerName	ownerAddress1	ownerAddress2	ownerCity
2102824320092	LIC385097	CONV	Active	Tier 1	Active	11/7/2019, 5:03 AM	2/29/2024, 7:00 PM	5924 OLIVER AVE S	Amy A Smith	1701 W 143rd St. #414		Burnsville
1211821240019	LIC375543	CONV	Active	Tier 1	Active	1/7/2019, 3:11 AM	2/29/2024, 7:00 PM	5013 ALDRICH AVE N	Dwayne Meier	605 Hwy 169 N	Ste 1050	Plymouth
0302824430056	LIC408633	CHOWNEXMPT	Active	Tier 1	Active	5/22/2023, 9:38 AM	2/29/2024, 7:00 PM	3635 1ST AVE S	Jeniffer N Kuria	15915 27th Pl N		Plymouth
2102924330058	LIC392462	CHOWN	Delinquent	Tier 1	Active	12/23/2020, 11:57 AM	2/28/2023, 7:00 PM	2001 GLENWOOD AVE	RYAN S DAVIS	2001 GLENWOOD AVENUE #2		MINNEAPOLIS
1602924110043	LIC390040	CONV	Active	Tier 1	Active	6/25/2020, 11:19 AM	2/29/2024, 7:00 PM	2311 LYNDALE AVE N	Alexander Gladkov	2905 STRATTON CIRCLE		MINNETONKA
1902923240006	LIC390318	CONV	Active	Tier 1	Active	7/27/2020, 5:11 AM	2/29/2024, 7:00 PM	923 25TH AVE SE	Gonzalo J Bellido	4631 Parksidge Dr		Eagan
1211821340060	LIC392622	CHOWN	Active	Tier 1	Active	1/25/2021, 8:55 AM	2/29/2024, 7:00 PM	4647 BRYANT AVE N	Rada Varshavskaya	5808 Olinger Blvd		Edina
1502824340041	LIC406459	CONV	Active	Tier 1	Active	1/25/2023, 9:42 AM	2/29/2024, 7:00 PM	216 VALLEYVIEW PL	Steve J Imhoff	117 Minnehaha Pkwy		Minneapolis
0702823220032	LIC393445	CHOWN	Active	Tier 1	Active	3/25/2021, 8:17 AM	2/29/2024, 7:00 PM	3117 38TH ST E	Daniel Riley	3330 Edmund Blvd		Minneapolis
2002924140133	LIC382697	CHOWN	Delinquent	Tier 2	Active	10/4/2019, 7:29 AM	2/28/2023, 7:00 PM	627 QUEEN AVE N	Victoria Yepez	7900 Telegraph Rd		Bloomington
0402924330210	LIC393735	CONV	Active	Tier 1	Active	4/12/2021, 10:33 AM	2/29/2024, 7:00 PM	3601 NEWTON AVE N	Christopher Luehr	3411 Yates Ave N		Crystal
0802824340171	LIC388870	ShrtTrmReg	Delinquent		Active	3/27/2020, 6:14 AM	2/28/2023, 7:00 PM	4517 ABBOTT AVE S	Andrew Root	4517 Abbott Ave S		Minneapolis
1302924220017	LIC391504	CHOWN	Delinquent	Tier 1	Active	10/7/2020, 11:51 AM	2/28/2023, 7:00 PM	1704 1/2 FILLMORE ST NE	MRK Properties LLC	2701 Georgia Ave S		St Louis Park
2602924230162	LIC391821	CONVEXEMPT	Active	Tier 2	Active	10/28/2020, 8:23 AM	2/29/2024, 7:00 PM	1014 PORTLAND AVE	JENNA T GILDAY	5865 AYRSHIRE LANE		SHOREWOOD
1202824310052	LIC392228	CHOWN	Active	Tier 1	Active	12/4/2020, 5:40 AM	2/29/2024, 7:00 PM	4201 21ST AVE S	Jasmine Radue	3412 Meridian Dr		Robbinsdale
1602924410027	LIC382346	CONV	Active	Tier 1	Active	8/6/2019, 10:35 AM	2/29/2024, 7:00 PM	1826 DUPONT AVE N	Oscar Vicente Jara	2120 14th Ave S		Minneapolis
0402824440081	LIC394020	CHOWN	Active	Tier 1	Active	4/30/2021, 8:51 AM	2/29/2024, 7:00 PM	3640 ALDRICH AVE S	Jesse Olson	3033 Excelsior Blvd - #100		Minneapolis
1602924320087	LIC394217	CONV	Delinquent	Tier 2	Active	5/24/2021, 8:42 AM	2/28/2023, 7:00 PM	1700 PENN AVE N	Bobbie Evans	2220 Queen Ave N		Minneapolis
0202824420025	LIC380210	CONV	Active	Tier 1	Active	6/10/2019, 5:15 AM	2/29/2024, 7:00 PM	1513 35TH ST E	Ephraim Eusebie	1521 35th St E		Minneapolis
1202924240167	LIC371777	CHOWN	Active	Tier 1	Active	10/30/2018, 12:01 PM	2/29/2024, 7:00 PM	2520 PIERCE ST NE	Philip Chan	3120 Hennepin Ave	Apt #107	Minneapolis

Figure 4.2: An Example Rental Registry Database

Note: This screen capture comes from Minneapolis Open Data (https://opendata.minneapolismn.gov/). APN is the Assessor Parcel Number, which can be matched to the PIN in Figure 4.1.

Ostensibly, RROs exist so that the state can keep better track of the location and ownership of rental properties. When paired with required inspections, as in Baltimore or Detroit, the registries also function to protect tenants from unsafe dwelling units. In Seattle, the name of a registered “contact person” is listed online, so that tenants can always have a name of someone to call in case of a problem at the property. Thus, RROs are put in place by local governments in order to identify the location of rental properties, and perhaps to protect tenants and nearby landowners from negligent landlords. They also serve as a source of revenue, albeit a minor one: nearly all RROs that I encountered assess a fee (per-building or per-unit, sometimes uniform and sometimes on a scale proportional to the number of units) on the registrant. These fees are often fairly nominal; for instance, Dallas assesses a $6 per unit fee, while Philadelphia assesses a $56 per unit fee.

However, RROs are far from uniform, and their uneven implementation leaves much to be desired. This unevenness appears in two ways: what data is collected by cities themselves, and the enforcement and actual coverage of the data. Regarding the former, who collects the data, and what they collect, differs significantly from ordinance to ordinance. At one extreme, entities that are not typically in the business of collecting housing information are tasked with enforcement. For example, the Charlotte-Mecklenburg Police Department and the Kansas City Health Department are charged with operating rental registries is their respective cities. In Portland (OR) the Revenue Division maintains the rental registry, and in San Diego, the City Treasurer maintains the rental registry, as in both of these cities it is a rental property “tax” that is being collected.[54] Most often, it is a code enforcement division that maintains rental registries, though licensing departments are also often the custodian of these records. These differences appear to arise from different statewide political conditions. Local governments are only able to enact ordinances within the scope of powers granted to them by the state; among cities in states that are prone to preempt local governments or otherwise restrict their actions, this means finding a way to enact a registry that does not run afoul of state law. That sometimes means a rental registry is maintained by the health department, even if that is otherwise not an obvious choice for such an enforcing agency.[55]

---

[54] The San Diego City Treasurer and Portland Revenue Division, notably, do *not* collect property tax, as that is collected at the county level.

[55] This is no hypothetical. Pittsburgh, PA has tried on multiple occasions to enact a rental registry, only to be found in violation of state law. In 2014, Oklahoma passed a law to prohibit Tulsa from enforcing its rental registry, HB 2620. In 2016, North Carolina similarly passed a law restricting rental registration ordinances, SB 326.

This heterogeneity in who collects the information affects what information is collected. If one were to design an RRO with the intention of identifying the ownership and management of rental properties for tenants and city administrators, one would collect information about the building itself (address, parcel ID, number of units rented), the ownership of the parcel (name, contact information), the management of the unit (name, contact information), and any beneficial ownership of the parcel, should the legal owner be a corporate entity. Some places collect all this information, but very few: Washington, DC, Philadelphia, PA, and Minneapolis, MN, appear to meet all of those requirements. Even then, Philadelphia's enforcement of the ordinance does not match the language of the ordinance, meaning that many corporate owners do not have recorded beneficial owners.[56] Others only require owner name and address, regardless of the owner's corporate status. Some only target out-of-town owners, requiring a manager name if the owner does not reside within a certain radius of the city. While there is great variety in the content of RROs, nearly all RROs seem to allow for the identification of which addresses in a city are rental properties, provided one considers the data collected by the RRO to be complete.

Coverage and enforcement are perhaps the larger problem with RROs. If one has a list of rental addresses, it is possible to cross-reference that list with assessor data to identify owners. However, an RRO that is poorly enforced is one that is missing rental properties in the city. Some of the differences in coverage can be explained by legislative differences: in New York City, Houston, Fort Worth, Las Vegas, and Long Beach, only multi-family buildings (variously defined as more than 3, 4, or 5 units) are required to register with the city; see Appendix A. A greater problem is low compliance. Approximately 14% of landlords seem to be registered with the city of Detroit (Lynch 2022), while the Pew Charitable Trusts reports only 55% of rental properties, and 70% of rental units, are registered in Philadelphia (Haider 2021). Landlords that own fewer units, and owners of single-family rentals, seem to be less likely to register, since it is harder to identify an unregistered single-family home than it is to identify an unregistered 25 unit apartment building. This lack of enforcement in Philadelphia persists, despite a seemingly strong incentive to landlords: in order to carry out a formal eviction, the property must be licensed as a rental. Even

[56] Based on my correspondence with the city government, while the ordinance requires beneficial owner's names and mailing addresses, the city is only collecting that information for new licenses, not for license renewals.

with this undercounting, Haider's report relies on rental registration data — including expired rental licenses — rather than property tax assessments, since they found the registration data was still more accurate in identifying rentals and unit counts than property tax assessment data. Table 4.3 provides a summary of the pros and cons of the three approaches examined here.

Table 4.3: Pros and Cons of Different Methods to Identify Rentals

	Mailing Address Approach	**Homestead Exemption Approach**	**Rental Registry Approach**
Pros	• Near universal data-availability • Simple interpretation	• Based on local identification of owner-occupancy	• Unlikely to include any non-rental properties • Provides unit counts • May provide additional information, like management names
Cons	• Misses properties for landlords who send mail to their rentals • Includes homeowners who mail property bill elsewhere • Requires extensive data cleaning	• May not apply in some localities • Limited uptake may mean inaccurate identification of rentals • May be inaccurate or out of date	• Limited availability • May suffer from limited enforcement • Some rental properties may be exempt • Difficult to acquire

§4.5 — Estimating Rentals Three Ways

In this section, I estimate the number of rental properties in five cities: Columbus, OH; Minneapolis, MN; Nashville, TN; Philadelphia, PA; and Washington, DC. I use these five cities because they all provide the parcel number in the rental registry, meaning that I am able to match rental registry data precisely to tax assessment data. In order to identify rental properties, I rely on both tax assessment databases and on rental registries. When using the tax assessment databases, I follow *both* the "mailing address" approach and the "homestead exemption" approach. I then compare the number of rental properties and rental units identified by the tax assessment databases to the number of rental properties and rental units identified by the rental registries. I compare all three approaches to the number of rental units as estimated by the American Community Survey 2017-2021 5-year estimates.

There is significant heterogeneity among these five cities. Minneapolis has its own tax assessor (distinct from its surrounding county), and provides residential unit counts for all parcels. I retrieved this information directly from the Minneapolis Open Data portal. I rely on Regrid data for the other four cities. Regrid is a national data provider of parcel-level information, collected from tax assessment databases and augmented with proprietary data sources and systems. These cities do not provide unit counts; instead, I rely on a field generated by Regrid, "address count," which counts the number of unique addresses associated with a parcel.[57] To assess the accuracy of the "address count," I compare it to the unit count provided by Minneapolis. The median parcel with a residential use has an identical address count and unit count. However, taking the average difference between unit counts and address counts results in a difference of 0.22 fewer Regrid addresses than would be expected based on the unit count; this is likely attributable to the fact that three percent of parcels have a missing address count, which is coded as zero addresses.[58] Thus, it appears that the address count from Regrid is an *acceptable* substitute, though not a perfect one, for local government-collected unit counts. Because Minneapolis is the only city that provides residential unit counts for all parcels, they are the only city for which I am able to estimate rental units within properties that are designed as "mixed use." In Philadelphia, Columbus, Washington, and Nashville, I exclude mixed-use properties, following Gomory (2021).

Homestead exemptions are another point of heterogeneity among these five cities. Nashville does not provide any information on homestead exemptions; it is not clear that they offer any form of tax relief to homeowners. While Ohio limits homestead exemptions to only seniors and veterans, Franklin County (in which Columbus sits) offers an owner-occupancy credit. As previously discussed, Minneapolis homestead exemptions are limited to owner-occupied houses with value below $413,000, but are applicable to residences that are occupied by the owner's family. According to the Minneapolis tax assessor, 25% of single-family homes have an assessed value greater than $413,000.

I tried two different variations of the mailing address approach. The first relied on matching only the street address.[59] The second relied on matching the street address, secondary address, city,

[57] Regrid refers to these as "secondary addresses" meaning that they are able to differentiate between street addresses and unit-level addresses.

[58] There are also cases of extreme disagreement — for instance, 400 more units than addresses associated with a relatively new apartment building.

[59] That is, matching on "House Number, Street Name," such as "77 Massachusetts Ave."

state, and ZIP code.[60] While I did minimal cleaning when matching only the street address, I followed a substantial cleaning process when matching full addresses, including standardizing street types, street directions, and city spellings. Figure 4.3, below, shows the challenges with both variations of the mailing address approach. On the edges of the local government jurisdiction, some homeowners list a mailing address associated with another jurisdiction. This is especially pronounced in Columbus, a city which has grown aggressively through annexation. On the other hand, using only the street address can misclassify rentals: there are nearly 250,000 non-unique street addresses in the United States.[61] While it is perhaps unlikely that a landlord would live at the exact same street address — but at a different location — as their property, it is perhaps more likely that a landlord of a condominium unit could live in the same building as their rental property, meaning that one would need the secondary address. Because of these challenges, I matched on street address for Columbus, while I matched on the cleaned full address for the other four cities.

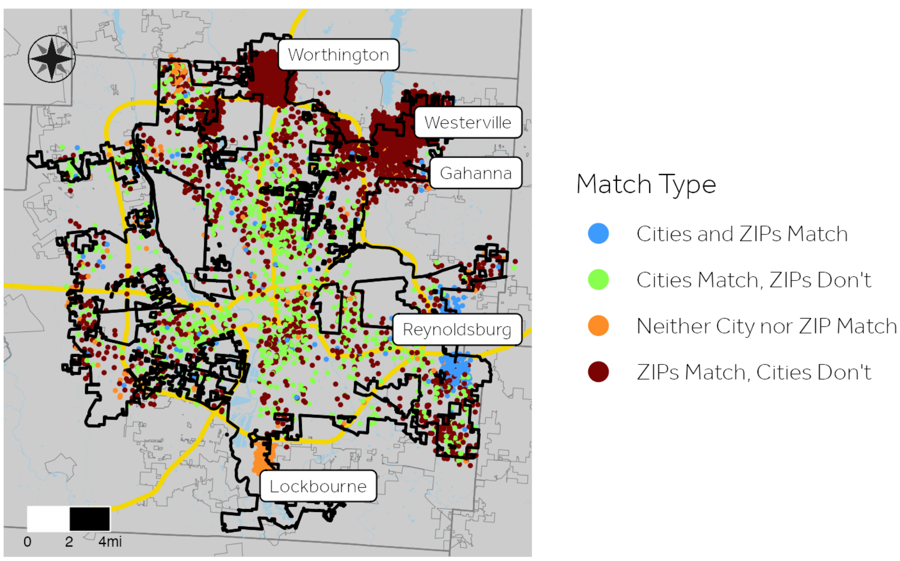

Figure 4.3: Challenges in Matching Addresses in Columbus, Ohio

Note: This map shows properties where the street address matches, but the full address does not. Commonly mismatched city names are labeled. Exact addresses are "jittered" to show density instead of precision.

[60] For example, "77 Massachusetts Ave, Cambridge, MA 02139."

[61] This figure was calculated using parcel data from Regrid. I aggregated the street addresses of every parcel in the United States and counted the number of street addresses that were duplicates.

In Columbus, the majority of the properties that have matching street addresses also have matching full addresses. However, 8,808 residential properties having matching street addresses and non-matching full addresses. These are mapped in Figure 4.3.[62] As is quite clear from the map, the problem of non-matching cities is limited to the edges. On the other hand, non-matching ZIP codes are spread throughout the city; sometimes people simply get their ZIP code wrong, but it is possible that there are multiple streets with the same name.[63] At the southern edge of the city, 132 addresses have neither a matching address nor a matching city. Perhaps these are actually distinct sets of addresses, or maybe there is an error in the data.[64] These are among the challenges with using the mailing address approach. In general, matching addresses is no easy task. There is lots of room for error. Researchers can use exact matching — are these two sets of words exactly the same — or they can use "fuzzy" matching approaches that allow for some room for error (see Jonge and Loo 2018, ch. 5), but that require a decision as to when two addresses are too different to match.

Below, Table 4.4 shows the different estimates for the number of rental properties and rental units obtained through the mailing address, homestead exemption, and rental registry approaches, as compared to the ACS. In Minneapolis and Nashville, estimating rental properties via the mailing address approach results in a significantly higher counts of rental units than the corresponding ACS estimate of housing units that are rental-occupied or vacant for-rent, while in Columbus and Washington, the estimated rental count is similar to the ACS estimate. In Philadelphia, looking for dissimilar addresses results in only identifying 72 percent of the number of rental units reported in the ACS, while using the homestead exemption approach results in identifying 100 percent of the ACS-reported number of rental units. Why is that? Likely because of the relatively low uptake in the homestead exemption in Philadelphia; it is estimated that 20 percent of eligible Philadelphians are not receiving a homestead exemption (Pananjady, Max

[62] There are 71 of the 8,808 properties that do not match for a reason other than the four reasons enumerated in the figure legend. They are not mapped.

[63] Boston, MA has at least three different locations with the address "30 Charles Street," in ZIP codes 02122, 02144, and 02136.

[64] It's not uncommon for suburban communities to have the same street names but with new numbering systems. This often causes problems. For instance, 13 Curtis St in Somerville, MA is located approximately one mile from 13 Curtis St in Medford, MA. These are two unique locations — with different city names and ZIP codes — but the same street address. Nationally, there are 12 different locations with the address "13 Curtis St."

Marin, and Lo 2022), increasing the number of housing units that are identified as rentals. Additionally, because 64% of Philadelphia parcels lack a mailing address, the ability to match on mailing addresses is extremely limited.

Table 4.4: Estimating Rental Properties via Tax Assessment Databases and Rental Registries

		Total Rentals via Mailing Address Approach	Total Rentals via Homestead Exemption Approach	Total In Registry	Total Possible ACS Rental units
Minneapolis	Property count	45,327	35,096	22,811	
	Unit Count	144,878 (146%)	134,634 (135%)	105,566 (106%)	99,471
Philadelphia	Property count	160,357	263,605	68,185	
	Unit Count	235,945 (72%)	338,267 (104%)	313,962 (96%)	325,805
Columbus	Property count	98,313	76,485	61,204	
	Unit Count	228,793 (105%)	208,819 (96%)	122,856 (56%)	217,742
Washington, DC	Property count	84,739	88,010	32,936	
	Unit Count	204,693 (105%)	210,675 (108%)	175,609 (90%)	195,264
Nashville	Property count	88,873	NA	15,064	
	Unit Count	160,504 (114%)	NA	194,52 (14%)	140,983

Note: Numbers in parentheses compare the unit counts from the local data estimates to the ACS estimate. The ACS estimate represents the sum of occupied rental units and "vacant, for rent" housing units. Nashville has no recorded homestead exemptions in the data provided, thus making those comparisons "NA".

Turning our attention to the rental registries, their coverage, relative to ACS unit counts, ranges from good — in Minneapolis, Philadelphia, and Washington — to acceptable — Columbus — to abysmal — Nashville. There are any possible number of reasons for this. In Nashville, rental registration is required by the state, which, in turn, directs the local government building code enforcement agency to register landlord information, and to fine landlords who do not register $50 per week.[65] On the other hand, in Philadelphia, a landlord without a rental license is not allowed to formally issue an eviction; in Minneapolis, landlords are not allowed to charge rent without a rental license; while in Washington, any unit not registered with the Rental Accommodations Division is automatically subject to rent control.[66] The length of time that these ordinances have existed also differs substantially. Tennessee only enacted rental registration in 2006; Minneapolis in 1990, Philadelphia has rental licenses going back to 1994, and Washington has rental licenses going back to 1999. There are also differences in the number of single-family rentals (SFRs), which

[65] 2010 Tennessee Code Title 66, Chapter 28, Part 1, §66-28-107 - Residential landlord registration.

[66] This includes units that otherwise would not be subject to rent control if they had registered with the city.

might lead to differences in enforcement, since it is difficult for a compliance officer to identify a renter-occupied as compared to an owner-occupied single-family home. One-in-four rentals in Nashville and Columbus are SFRs, compared to only one-in-seven in Minneapolis. This, however, does not explain all of the variation, as nearly 40 percent of Philadelphia rentals are SFRs, and the city appears to have high compliance.

Columbus, sitting in the middle in terms of coverage, is an unusual case. Rental registration is the responsibility of the county auditor — the same entity responsible for property tax assessment. Like Nashville, the rental registration requirement is the result of state law. Like Nashville, the penalty for non-compliance is paltry: a fine between $50 and $150. Yet Columbus's coverage is much more extensive than Nashville's. More qualitative work is needed to understand the types of rental housing in these two cities, the enforcement practices of the local governments, and the other political-economic contexts in which these two cities operate that may explain different apparent compliance with these two otherwise similar requirements.

Table 4.5 provides a more detailed comparison among the three approaches than can be found in Table 4.4. For each approach, it identifies how many units are included or excluded by the other two approaches. This allows us to see that between 4% and 20% of the rentals in the rental registry are not identified as a rental by either method using tax assessment data. In Philadelphia, this is likely because of the extensive missing data related to mailing addresses. In Minneapolis, the fact that one-in-eight rental properties would not be identified as a rental is striking, as the estimated number of rental units from the tax assessment database is so high.

Table 4.5: Areas of Overlap and Disagreement Among Tax Assessment Databases and Rental Registries

		Columbus		Minneapolis		Nashville		Philadelphia		Washington, DC	
		Property count	Unit Count	Property count	Unit Count	Property count	Unit Count	Property count	Unit Count	Property count	Unit Count
	Total Rentals via Mailing Address Approach	**98,313**	**228,793 (105%)**	**45,327**	**144,878 (146%)**	**88,873**	**160,504 (114%)**	**160,357**	**235,945 (72%)**	**84,739**	**204,693 (105%)**
	Total Rentals via Homestead Exemption Approach	**76,485**	**208,819 (96%)**	**35,096**	**134,634 (135%)**	**NA**	**NA**	**263,605**	**338,267 (104%)**	**88,010**	**210,675 (108%)**
	Addresses Match, No Homestead Exemption	11,106	10,473	7,299	7,270	NA	NA	115,720	114,233	22,378	20,178
	Addresses Don't Match, Has a Homestead Exemption	32,934	30,447	17,530	17,514	NA	NA	12,472	11,911	19,107	14,196
	Addresses Don't Match, No Homestead Exemption	65,379	198,346	27,797	127,364	88,873	160,504	147,885	224,034	65,632	190,497
Not In Registry	Addresses Match, No Homestead Exemption	5,608	5,178	5,938	5,928	NA	NA	102,968	101,579	13,865	13,297
Not In Registry	Addresses Don't Match, but has a Homestead Exemption	31,241	28,819	16,911	16,895	NA	NA	10,524	9,976	29,285	22,021
Not In Registry	Addresses Don't Match, No Homestead Exemption	13,883	31,811	9,112	43,337	77,264	141,963	75,328	90,882	54,775	74,707
In Registry	**Total In Registry**	**61,204**	**122,856 (56%)**	**22,811**	**105,566 (106%)**	**15,064**	**19,452 (14%)**	**68,185**	**313,962 (96%)**	**32,936**	**175,609 (90%)**
In Registry	Not Identified as a Rental by Tax Assessor	2,438	**8,339 (4%)**	1,761	**12,308 (12%)**	2,972	NA	26,524	**59,089 (18%)**	9,860	**20,858 (11%)**
In Registry	Addresses Match, No Homestead Exemption	4,914	4,092	1,369	1,349	NA	NA	13,152	13,850	3,587	4,229
In Registry	Addresses Don't Match, but has a Homestead Exemption	1,767	1,651	634	634	NA	NA	1,985	2,086	1,703	1,719
In Registry	Addresses Don't Match, No Homestead Exemption	52,085	108,774	19,047	91,275	12,092	19,452	26,524	238,937	17,786	148,803
	ACS Estimate of Rental Units		208,644		94,741		129,737		307,740		181,384
	ACS Estimate of Vacant For-Rent Rental Units		9,098		4,730		11,246		18,065		13,880
	Total Possible ACS Rental units		**217,742**		**99,471**		**140,983**		**325,805**		**195,264**

Note: Percentages in parentheses compare the number of estimated rental units to the total possible ACS rental units, bottom row. ACS estimates taken from the 2017-2021 five-year estimates, at the "place" level. For Philadelphia, Columbus, and Nashville, and Washington, D.C., I use "address counts" from Regrid to estimate unit counts for the tax assessment approaches. Nashville has no recorded homestead exemptions in the data provided, thus making those comparisons "NA".

Of particularly concerning note, there are hundreds, and in some cases thousands, of rental units that have a homestead exemption; as we can see in Table 4.6, this includes many single-family rentals. In Philadelphia, over 5,000 single-family homes with a homestead exemption also have a rental license,[67] while 5 percent of rental properties in Washington have a homestead exemption. Even in Columbus, where the same entity oversees rental licenses and homestead exemptions, nearly 2,000 properties seem to have both. Importantly, these comparisons are limited only to single-family properties. While it is possible that these properties have accessory dwelling units, basement apartments, or other rental units on-site while remaining in the tax assessment as a "single-family" unit, it is surprising how overly inclusive the homestead exemption approach seems to be. On the other hand, over 4,000 registered rental units in Washington and Columbus have matching location and mailing addresses. These findings show that neither approach using tax assessment data truly captures an accurate picture of the rental market.

Taking these observations all together, it seems that, in cities that have a rental registration ordinance, researchers should calibrate their methods against those rentals that are identified in the rental registry. If the researchers choose to use a tax assessment database to identify rental units, the most conservative approach would be to call a property a rental property if and only if the mailing address and building addresses did not match, and the housing unit lacked a homestead exemption. On their own, mismatching addresses are not enough to identify the full universe of rental licenses, while homestead exemptions are messy, owing to different uptake rates, incentives for landlords to claim a homestead exemption when they are not eligible, and differing state and local practices.

While coverage varies among rental registries among these five cities, it seems unlikely that many properties in the rental registry *are not* rentals. While it is conceivable that a landlord could move into their rental and let the license lapse, this would seem to be a minor issue compared to the extensive disagreement among the tax assessment database methods. While the rental registries may not represent the entire universe of rental units, they represent a conservative sample that is

---

[67] This number includes the 1,985 properties listed in the table as having a homestead exemption, plus the additional properties that had a homestead exemption and a matching building and home address.

less likely to include second homes, vacant not-for-rent properties, or properties where the owner has not taken a homestead exemption for which they are eligible. They are also able to identify rentals in mixed-use buildings, something otherwise only possible if the tax assessment database includes counts of residential units within a mixed-use building. The next section turns to Single Family Rentals (SFR), a subset of the rental housing market that has both garnered significant attention and allows us to overcome some challenges with the data sources.

§4.6 — Identifying Sources of Disagreement Among the Methods

Unfortunately, there is no ground truth to identify the true source of the discrepancies between the ACS unit counts, the tax assessment unit counts, and the rental registry counts. It is not possible to say whether the units captured by the methods using the tax assessment databases are actually rentals, or even the probability that they are rentals. One challenge in comparing the ACS unit counts to the other, local, sources of information is that the ACS asks respondents how many other units are in the building in which they live. That means rental units counts for mixed-tenure multifamily buildings may not align between ACS and local data sources.

To eliminate this source of variation between the ACS and local sources of information, I restrict rental units in the next table, Table 4.6, to single-family units only. The ACS provides a count of rental-occupied housing units that are single-family attached or detached homes, while all of the cities below have a set of land use codes associated with single-family dwellings.

What is immediately apparent is that, across all five cities, it is possible to construct a method using the tax assessment database that results in a similar number of units to the single-family rental (SFR) count from the ACS. However, which method gives you the most similar count differs among the cities. In Minneapolis, estimates via matched addresses results in number that is nearly double the count of estimated SFRs from the ACS. Based on a manual inspection, it appears to be due to the fact that street directions in Minneapolis are sometimes included in the mailing address, but not in the building address, or vice versa.[68] While this is not an insurmountable problem, it does require the researcher to have an intimate knowledge of street naming conventions

[68] For example, a building may have an address of "123 Main Street," but an owner mailing address of "123 Main Street S."

in the area of study. While it seems likely that two buildings with the same street address except for the street direction would be the same, that is a subjective decision by the researcher to ignore street directions.

On the other hand, using the homestead exemption approach is not an option in Nashville, and results in a huge overestimate of single-family rentals in Philadelphia. As discussed in the previous section, this is likely due to limited homestead exemption uptake in Philadelphia. This means that, to conduct a national study of rental markets, using a one-size-fits-all approach may not work. In some cities, the homestead exemption appears to be a more accurate approach for SFR counts, while in other cities, matching addresses seems to do a better job.

Table 4.6: Estimating Single Family Rentals from Tax Assessment and Rental Registry Data

	Columbus	Minneapolis	Nashville	Philadelphia	Washington, DC
Estimated SFRs, ACS Data	56,235	13,727	35,449	122,730	22,279
Estimated SFRs, Street Address Approach	53,334	22,161	46,690	123,133	20,357
Estimated SFRs, Mailing Address Approach	55,648	23,005	44,809	123,052	20,466
Estimated SFRs, Homestead Exemption Approach	39,648	13,681	NA	218,452	29,840
Estimated SFRs, Rental Registry	28,392	7,123	3,380	74,648	26,456
Estimated SFRs in Rental Registry, Not Identified as a Rental in the Tax Assessment Databases	359	622	708	3,506	2,170
Estimated SFRs in Rental Registry, Not Identified as a Rental via Address Mismatch[69]	4,860	1,546	711	15,850	4,181
Estimated SFRs in Rental Registry, Not Identified as a Rental via homestead exemption	1,761	1,050	NA	5,415	2,554

However, Table 4.6 also shows that rental registries often fall short when it comes to registering single-family rentals. Even in Minneapolis, which appeared to have accurate coverage in Table 4.5, actually only has 51% of the SFRs as identified in the ACS. Washington, DC has a higher count of SFRs to the ACS, the only city that comes close to what appears to be adequate

[69] This number is based on the more conservative address matching approach, above.

coverage. As was evident in the prior section, Nashville has extremely limited coverage of rental properties, SFRs included.

On the other hand, the bottom three rows paint a challenging picture for using the tax assessment databases. Even with the over-inclusion of rentals in Philadelphia and Minneapolis, somewhere between 2.8 and 4.5 percent of single-family homes in the rental registry are not identified using any of the tax assessment methods. In Washington, DC, that number is closer to 10 percent. In general, it seems like the mailing address approach misses a substantial portion of the SFR market: some landlords seem to keep their mailing address the same as the property address, even when they are renting it out. While the homestead exemption approach performs "better," there is still a concerning number of registered SFRs that have a homestead exemption. Given that homestead exemptions are typically only granted to owner-occupied homes, but most cities exempt owner-occupiers from registering as a rental property, it seems like a grave disagreement between the tax assessment and the rental registration ordinances.[70] None of these three approaches promise to identify all SFRs in a city. Again, however, rental registration ordinances are the most conservative: they include a clearly-identified subset of the rental market, without over-inclusion of any non-rental single-family homes.

§4.7 — Comparing Owner Addresses Among Methods

Both tax assessment databases and rental registries have a data column labeled "Owner's Address." However, it is not uncommon for that column to be filed with something other than the owner's address. Some landlords, when filling in the rental registration form, will provide the mailing address of their management company. The same appears to be true for tax assessment databases. Nonetheless, it is worth exploring whether landlords provide the same mailing address in the tax assessment database and in the rental registry.

[70] The one exception here is Minneapolis, which grants a "relative" homestead exemption to immediate family members of the owner, but still requires a rental license. Nonetheless, it seems unlikely that all 1,000 registered rental units with a homestead exemption are occupied by owner's relatives, and a spot-check online confirms that many of these properties do not have a "relative" homestead exemption.

Table 4.7 shows the fraction of mailing addresses that are the same in the tax assessment database and the rental registry, among all rentals that are registered with the city. Instead of using an exact match, I use the "cosine similarity" between the owner address in the registry and the mailing address in the tax assessment database. Cosine similarity compares how similar two sets of words or characters are by looking at how many overlapping words or characters there are between the two sets. A cosine similarity of one means that the two sets are identical, whereas a cosine similarity of zero means that the two sets share nothing in common.[71]

Hangen and O'Brien (2022) also use the cosine similarity, to account for the fact that misspellings are common. I opt to use cosine similarity here because I am matching from two data sets — the tax assessment database and the rental registry database — and would expect differences in how the addresses are recorded. I use a cutoff of 0.9 based on a manual inspection of sensible matches and mismatches.

Table 4.7: Similarity in Owner Addresses, Tax Assessment and Rental Registry Databases

Columbus	Minneapolis	Nashville	Philadelphia	Washington, DC
100.0%	69.7%	65.1%	53.9%	55.4%

Most of the time, the mailing addresses and the owner addresses match. Yet, with the exception of Columbus, where the rental registry is part of the tax assessment database, the level of agreement is rather tepid. This presents yet another issue for researchers studying rental markets, for if one were to try and understand, say, the residence of landlords relative to their rental properties, the results would be quite different if you were to use the tax assessment database as compared to the rental registration database.

§4.8 — Sub-City Variation in Rental Locations

Thus far, I have only examined city-level counts of the number of rental units. However, it is possible that some neighborhoods within a city have higher compliance with the rental registry (city services often have a spatial bias, see Narayanan and MacDonald 2019), higher or lower usage

[71] Cosine similarity does not count for order, so "77 Massachusetts Ave", "Massachusetts 77 Ave", and "77 Mustachesssat Ave" would all have a cosine similarity of 1.

of the homestead exemption (Ihlanfeldt 2021) or that there are unique issues with using address matching (as in Columbus, see Figure 4.3).

To test this, I map the number of SFRs at the census tract level in Philadelphia. I use Philadelphia because of its large number of SFRs and high coverage of its rental registry. I adopt the fitness ratio from Molfino et al. (2017):

$$Fitness = \frac{\text{ACS Estimate - Local Estimate}}{\text{90\%ACS MOE}}$$

Fitness is thus defined as whether the difference between the local estimate and the ACS estimate are within with 90% margin of error (MOE) from the ACS. A fitness ratio between -1 and 1 indicates a similar estimate, while values larger than |1| indicate disagreement beyond the uncertainty associated with the ACS. Figure 4.4 shows the number of SFRs in Philadelphia, and the associated fitness ratios relative to the ACS.

Mapping the fitness ratios shows remarkable agreement between the three methods. All three methods for estimating the number of SFRs from local data show higher-than-ACS numbers in Southwest Philadelphia and North Philadelphia. Many of the neighborhoods where the ACS estimate is lower than the local estimate are the same across all three maps. Generally speaking, it does not seem like any of the three methods perform better in some neighborhoods than others: all three methods perform similarly across the city, with similar areas of positive fitness ratios and similar areas of negative fitness ratios.

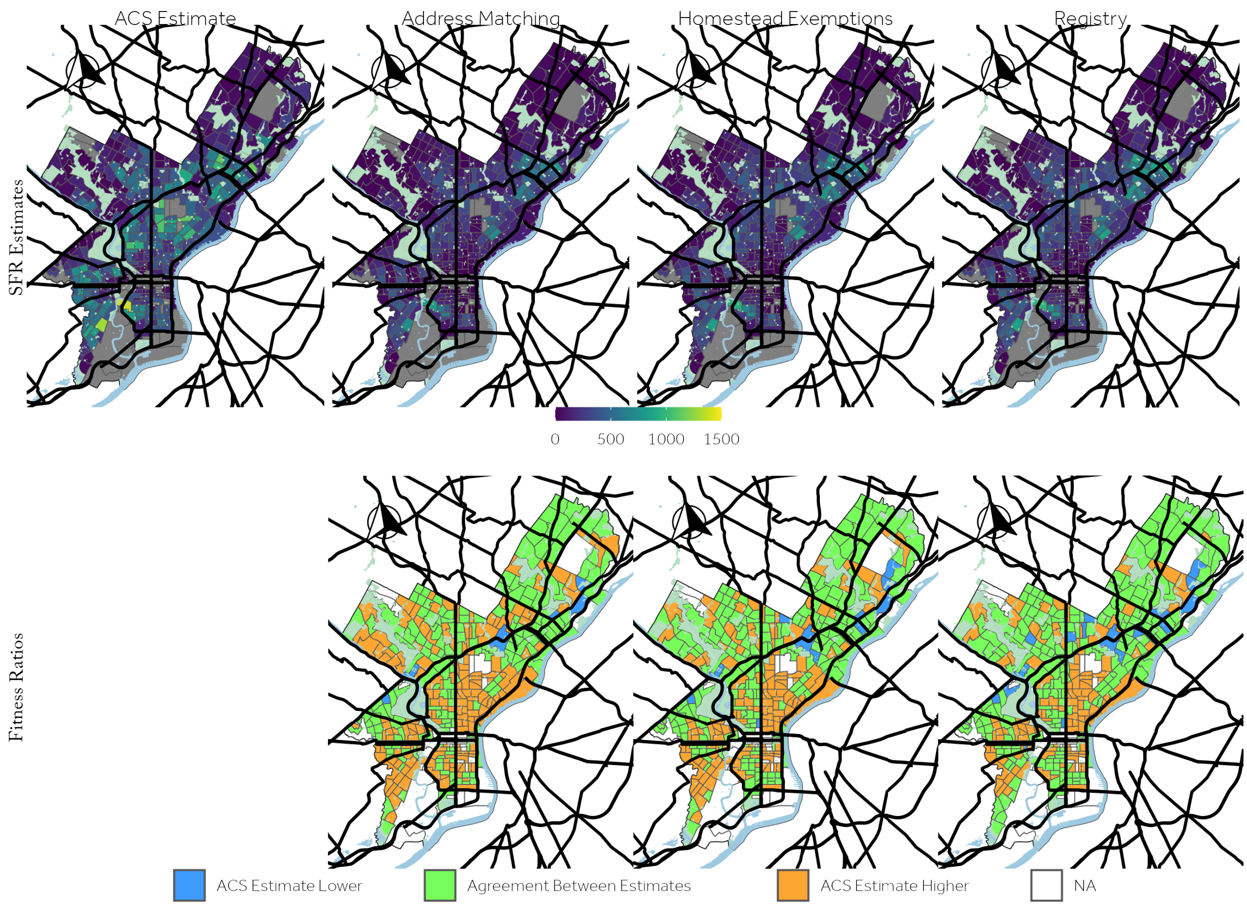

Figure 4.4: Number of SFRs in Philadelphia, by Census Tract

Note: ACS Estimates come from 2017-2021 5 Year ACS Estimates.

Table 4.8 shows the value across census tracts of various socioeconomic variables grouped by fitness ratios. A few patterns emerge from this table. For example, among the census tracts that overestimate the number of SFRs, relative to the ACS, the median income is higher and there is a smaller proportion of the population under 200% of the poverty line. In Philadelphia, many of the census tracts with the highest median incomes are in Center City and environs, which are dominated by multifamily buildings; it is thus possible that these approaches are picking up condominiums instead of single-family homes in these areas. Among census tracts that have similar ACS estimates and local-approach estimates, the demographics seem to generally align to the city-wide average. Similarly, and surprisingly, no pattern emerges based on the median gross rent. All categories tend to be lower than the city-wide median, suggesting that areas that are missing single-family rental estimates tend to have higher rents.

Table 4.8: Socioeconomic Characteristics of Census Tracts Based on Local Data Fitness

		Homestead Fitness			Address Fitness			Registry Fitness		
	Citywide Average	ACS Estimate Lower	Agreement Between Estimates	ACS Estimate Higher	ACS Estimate Lower	Agreement Between Estimates	ACS Estimate Higher	ACS Estimate Lower	Agreement Between Estimates	ACS Estimate Higher
% Non-Hispanic white	34%	43%	34%	30%	34%	35%	29%	42%	34%	31%
% Black	40%	24%	41%	42%	33%	41%	42%	24%	42%	41%
% Hispanic	15%	22%	14%	18%	20%	13%	19%	25%	14%	17%
% Over 25 with a College Degree	33%	30%	30%	33%	29%	30%	33%	28%	30%	32%
% Under 2x the Poverty Line	43%	32%	41%	48%	37%	41%	49%	31%	40%	46%
Median Household Income	$ 52,649	$ 62,970	$ 53,137	$ 46,728	$ 64,944	$ 52,856	$ 48,645	$ 56,732	$ 53,206	$ 45,055
Median Rent	$ 1,149	$ 1,150	$ 1,132	$ 1,111	$ 1,048	$ 1,116	$ 1,132	$ 1,125	$ 1,138	$ 1,099

Note: Fractions are based on cumulative populations from each census tract in a given category. Median household incomes and median rents calculated using Pareto interpolation from ACS household income and rent brackets. All data from 2017-2021 ACS estimates.

§4.9 — Conclusion

In this paper, I compared different methods for identifying rental properties through local administrative data in the United States. I reviewed the current literature that aims to study rental housing, and demonstrated that there is little agreement on the universe of rental units to study, nor the method to identify those rental units. I demonstrated significant differences in the number, location, and type of properties that are identified as rental properties through tax assessment databases. I compared these approaches to the universe of rental units as identified by rental registration ordinances. While RROs are fairly limited in the United States, numerous cities have them, and many more have proposals to create them. They rarely capture the entirety of the rental market, but they represent a true sample of the rental market.

There are other ways to identify rental properties beyond those identified here. The US Department of Housing and Urban Development keeps track of rental properties in the Housing Choice Voucher program, allowing for scholars to study a crucial rental submarket (Early, Carrillo, and Olsen 2019; Ellen, Harwood, and O'Regan 2022). The Internal Revenue Service asks those who receive income from rent to report the addresses of their rental properties.[72] Those who study institutional landlords have often looked for owner names that explicitly match the known subsidiaries from public company disclosures (Charles 2020b; Mills, Molloy, and Zarutskie 2019). Proprietary data sources that aggregate rental listings, such as Zillow or Apartment List, are another avenue. However, access to these sources of data is extremely limited.

This paper has two audiences: policymakers and researchers. I would hope that the policymakers would take from this study the need to implement a rental registration ordinance in their own jurisdiction. Without such an ordinance, our ability to identify rentals — and therefore target interventions on behalf of renters or landlords — is extremely limited. Such an ordinance should require all rental units within a jurisdiction to be registered, and to disclose the number of units, the name and address of the landlord, and the name and address of the management company. If the landlord is a corporation, the names of the human owners associated with that

[72] See IRS Schedule E.

corporation should be disclosed. At time of registration, the parcel number of the rental property should be linked to the tax assessment database, both to make tracking easier, but also to ensure that rental properties are not improperly receiving a homestead exemption. The information from these rental registries should be public, easily accessible, and updated regularly.

Researchers should be cautious about using tax assessment databases for research that requires precise information about unit-level information about rental status. While the approaches commonly used within the literature may result in the correct *number* of rental properties, they may not actually identify actual rental units. The different approaches used in the literature — be it the mailing address approach or the homestead exemption approach — work to varying degrees in different cities, depending on local data quality, accuracy, and legal institutions. The most conservative approach, when using a tax assessment database, would be to identify as a rental property only those residential properties that have *neither* a matching address *nor* a homestead exemption.

An alternative approach, as explored here, is to use data from rental registration ordinances in the cities that have them. Coverage and availability vary widely, so it is not always a feasible approach. But some rental registration ordinances appear to contain the majority of the rental units within a city, and should not face the same problem as tax assessment databases with regards to over- or under-inclusion depending on the method used by the researchers.

There may be an opportunity to use these two sets of data in tandem. Machine learning approaches could be used to identify probable rentals based on the characteristics of the properties within the tax assessment database. For now, though, we are left with few good choices in studying local rental markets.

CHAPTER 5 | CONCLUSION

This dissertation has examined rental housing markets in the United States, including the location of landlords, concentration of ownership of rental housing, and the different ways the literature has identified rental units for the purpose of studying rental markets. These papers have investigated the rental market at a moment-in-time, namely, when the data was collected, in 2022. However, as has been stated throughout the dissertation, rental housing markets are changing in numerous ways. There are larger actors, new actors, and new technologies. While the United States is building more multifamily units than it has in decades, the country is still short millions of affordable units. Towards that end, then, this concluding chapter identifies three future areas ripe for investigation. First, as an increasing number of large, professionalized landlords expand their market share, what happens when the rental market becomes more efficient? Second, what is the role of management companies, and what is the impact of technology, on rental markets? Third, how does the phenomenon of renting and owning rental properties interact with other secular changes in the US economy which have contributed to inequality? I conclude by reflecting on implications for policy and practice.

§5.1 — Findings & Contributions

I begin by summarizing the findings of these three papers. Owing to the nature of the structure of a three-paper dissertation, the questions answered here are precisely drawn and narrow in focus. Nonetheless, they are interrelated insofar as they all share a focus on rental properties and landlords in the United States. They draw from a common source of data — rental registries collected at the local level — and I used many of the same methods across the papers. Additionally, they all sought to illuminate areas of the rental market about which information heretofore has been limited: the locations of landlords, their holdings, and even the identification of rental properties.

Paper One, *Where the Landlords Are: A Network Approach to Landlord-Rental Locations,* studies the location of landlords relative to the properties that they own. I find that while most landlords are local, a significant share of landlords have mailing addresses far from their rental properties. Landlords with residential mailing addresses have addresses in neighborhoods that are richer, whiter, and have more college graduates than the neighborhoods in which they own property. Thus, landlords — via the capital that they are able to invest — extend their reach far beyond "their" neighborhoods and into poorer neighborhoods throughout the United States. This mobility of capital, it turn, constrains the mobility of labor, as low-income individuals are unable to move to neighborhoods of opportunity (Chetty, Hendren, and Katz 2016), and unable to move to cities with better job prospects (Glaeser and Gyourko 2018). Given that most landlords are local, but live in different types of neighborhoods than where they own property, this paper contributes to the understanding that the landlord-rental process is one of extraction that connects low-income and high-affluence areas within a region (Shelton 2018).

Paper Two, *The Relationship Between Local Rental Market Ownership Concentration and Rent*, studies the holdings of landlords. At the scale of the city, the ZIP code, the census tract, and the elementary school attendance zone, I calculate the market shares of landlords. I find that there are many sub-city geographies in the cities that I study where landlords have very large market shares. Large enough, in fact, that these neighborhoods meet the definition of "moderate" and "high" concentration. Higher levels of concentration are correlated with higher rents. Additionally, neighborhoods with higher levels of concentration see larger rent increases than neighborhoods with lower levels of concentration, conditional on a city experiencing an increase in wages.

Paper Three, *Is This a Rental? Comparing Methods for Identifying Rental Units*, aims to demystify the process of differentiating between rental units and owner-occupied units for the purposes of academic research. While it may be hard to believe, it is not an easy process to identify rental units from among the universe of all residential parcels in a given city. While academic researchers have used tax assessment databases to identify rental units through a variety of different approaches, these approaches have significant shortcomings. I compare these approaches to the sample of rental units identified by rental registries collected at the local level, and I find that, while rental registries and tax assessment databases can identify many of the same properties as rental properties, all approaches have benefits and drawbacks. In this paper, I also suggest the "ideal" data collected by a rental registration ordinance, and outline why tax assessment databases suffer the drawbacks I identify.

Across three papers, this dissertation contributes to literatures in urban studies, urban economics, and geography. It studies the current composition of the rental housing market in the United States. In contrast to research that typically studies low-income populations as "the other" (Fine 1994), this dissertation turns an eye towards property owners. It does so in concert with a stream of other studies in recent years that have also scrutinized these actors in society (Gomory 2021; Hochstenbach 2022; Rosen and Garboden 2022; Shiffer–Sebba 2020). Paper 1 contributes to the economic geography literature by focusing on the flows of capital between renter and landlord, and neighborhood to neighborhood. Paper 2 contributes to the urban economics literature by studying the extent to which local rental markets are competitive. Paper 3 contributes to the urban studies literature by investigating the utility of locally-collected administrative data, pointing the way for practitioners to implement tools that would allow them and academic researchers to better understand the rental market. This dissertation is only part of a broader research agenda that seeks to understand how inequality is produced and reproduced in the United States through differences in property ownership and wealth. The following sections highlight areas that are left for future research, informed by the current findings, yet still in need of deeper investigation.

§5.2 — A Research Agenda

§5.2.1 — Rental Market (In)Efficiences

In interviews with "circumstantial" landlords, Shiffer–Sebba (2020) describes market participants who are not acting rationally, in an economics sense. He describes landlords who forget to raise the rent, or choose not to because they have good tenants. On the other hand, Abood (2018) describes "Wall Street Landlords" who use algorithms to set rent increases. Near the end of the rental tenancy, tenants receive a notice from their landlord, offering them 16 different rental rates, depending on if they choose to sign a month-to-month lease, a 12 month lease, or a lease lasting between 4 and 18 months. Investigative reporting from ProPublica describes landlords using software that recommends rent increases upwards of 10% (Vogell, Coryne, and Little 2022), even though property managers would push back on such a large increase. In addition, a number of studies have found that larger landlords are more likely to evict tenants, as "serial filing" of evictions are often part of a fines-and-fees business model (Gomory 2021; Immergluck et al. 2020; Robinson and Steil 2021). Additionally, Gurun et al. (2022) found that larger landlords are explicitly able to extract greater surplus from renters.

At present, the majority of landlords remain "mom and pop" landlords who own only a handful of properties. Yet as large landlords continue to expand their holdings, and as technology companies continue to innovate in the "PropTech" sector (Charvel 2023; DelPrete 2023; Fields 2023; Goodspeed 2023), we might expect these changes to only accelerate. The tools described by Abood (2018) and Vogell, Coryne, and Little (2022) do not necessarily mean that landlords are "beating" the market, or exerting pricing power, only that they are able to bring their current tenants — or prospective tenants — up to market prices. This increased market efficiency may mean less slack for tenants, who, absent rent stabilization and anti-eviction policies, are required to accept steep rent increases or find another place to live.

Thus, there is reason to wonder if the days of an inefficient rental market are over. By this, I mean, we may see fewer rentals that are priced under the market, with renters finding it more difficult to find a "good deal" on rent. As landlords are better able to access information about asking prices and other landlords' contract rents, we may see more rent increases, higher rent burdens, and more evictions.[73] Predicting the future is hard, and these concerns may not come to pass — especially with policy interventions that could be implemented to mitigate them. Nonetheless, a more "efficient" rental market may mean more challenges for renters, given the continued appreciation of rent, especially in America's superstar cities. An inefficient market is not necessarily a "fair" market, — not all who need a good deal on rent find it, and some who find good

[73] Including no-fault evictions, where landlords opt to not renew a lease term.

deals could afford to pay more — but an efficient market may bring disaster for tenants in high priced cities with few tenant protections.

There are other changes in the rental housing market that may exacerbate these challenges. AirBnB and other short-term rental services have distorted the housing market, transitioning rental housing from long-term to short-term leases (Coles et al. 2018; Garcia-López et al. 2020). In subsidized housing where subsidy provisions have expired,[74] heated fights have erupted over redevelopment. Many Low Income Housing Tax Credit developments will age-out of their affordability provisions in the coming decade, raising alarm about the long-term prospects for these developments (Lens and Reina 2016).

While much of the academic literature has been focused on the expansion of institutional investors in the single-family rental housing market, large landlords have also expanded in the multifamily market (Vogell 2022). In the United States, only 5 percent of multifamily construction is in the form of condominiums; the rest are rentals (M. Neal and Goodman 2022). By design, that means that new multifamily units are almost exclusively owned and operated by large landlords. As multifamily construction is at a multi-decade peak, that means that large landlords will only continue to expand within the multifamily market, continuing to push towards a more efficient market.

Thus, future research needs to address the following questions. How (in)efficient is the rental market, right now? Do most tenants see annual rent increases, or do many landlords let years pass without increasing rents — even when they could increase rents if they chose to do so? Are there places — cities or neighborhoods, in the US or abroad — that have relatively efficient markets? If so, can they provide insights as to what might happen if and when the rental market in the United States becomes more efficient? Are there tools that policymakers can implement that would mean efficiency does not coincide with tenant displacement? When Massachusetts ended rent control, researchers found robust price appreciation in the years that followed (Autor, Palmer, and Pathak 2014). As rental markets become more efficient, it behooves us to understand the effects

[74] Such as projects funding with Low Income Housing Tax Credits, or HUD Section 202 or Section 811 Developments

that such efficiency may have on tenant welfare, real estate finance, and demand-side subsidies such as the Housing Choice Voucher Program.

§5.2.2 — Other Actors in the Rental Market

The three papers in this dissertation focused on landlords. In this case, landlords were defined as the owners of residential property that was rented to a tenant in exchange for rent. Yet landlords are not the only actors in the rental market. Approximately one in five rental properties are under the day-to-day management of a management company or a management agent paid by the owner (US Census Bureau 2021). I excluded information about management companies from my study because of the relative opacity of their role and impact on rent-setting and rent collection. While some studies have investigated the role of professional management (Carswell 2018; Garboden and Rosen 2022; Shiffer–Sebba 2020), there is still more that needs to be learned about their role in the rental market, especially the setting of rents.

The role of management is particularly important in the context of the new "laptop landlord," (Parker and Friedman 2022). Over the past decade, a number of property management and investment companies have launched, enabled by technology, that focus on out-of-town investors. These companies, such as "Evernest," "Arrived," and "Fundrise," specifically advertise the benefits of real estate investing to investors while attempting to "democratize" the market. In particular, Arrived and Fundrise are crowdsourcing platforms, allowing individuals to invest in real estate with "as little as $100." Evernest, on the other hand, allows individuals to buy rental properties through their platform, and then offers to manage those properties on behalf of the new owners.

If, in the past, landlords liked to be close to their rentals because knowing the market provided them with important information that allowed them to make wise investing decisions (Crook, Ferrari, and Kemp 2012; D'Lima and Schultz 2021), then these new technologies represent a different logic regarding real estate investment than the landlords of the past. Their existence — and whether they succeed or fail — raises questions about the role of technology in enabling the mobility of capital, as well as the role of management companies as a mediator of that capital. If real estate investors are able to identify, procure, and manage properties without ever seeing the properties themselves, then the old logics associated with "mom and pop" landlords may no longer apply. Indeed, portfolio size may not be the appropriate differentiator, but instead, the investment goals and decisions made, the intermediaries used, and the choices made by those investors. Like the "post-industrial widgets" of the subprime lending crisis, technology and capital have intertwined to invent a new way to extract value from low-income occupants in the housing market (Newman 2009).

Thus, future work is needed to understand the role of technology and management companies in the process of market making. While size, spatial scales, and investment objectives have all previously been used in the literature as differentiating factors (Özogul and Tasan-Kok 2020), it may be time to add new dimensions to the ways in which scholars categorize different types of landlords. We need a better understanding of the logics of investors who choose real estate mediated by technological platforms. There are questions about why technology platforms choose certain markets and not others — including explanations focused on supply constraints on the one hand, and lack of tenant protections on the other. Additionally, more information is needed to understand whether these new technologically-driven management companies differ in approach from more traditional management companies, in terms of how they raise rent, interact with tenants, and interact with owners.

§5.2.3 — Renting and Inequality

This dissertation found that landlords live in neighborhoods that are different from the neighborhoods in which they own property, even if the majority of landlords live in the same city or MSA as their rental properties. Thus, the geography of where landlords live and the geography of where renters live are different. Across the United States, nearly one in three neighborhoods are so-called "rental deserts," where there are few, if any, rental properties (Airgood-Obrycki and Wedeen 2022).[75] Research from abroad tells us that landlords are often higher-income households, with many of the wealthiest in society engaging in the act of being a landlord (Hochstenbach 2022; Statistics Canada 2022).

[75] Specifically, Airgood-Obrycki and Wedeen (2022) define a rental desert as a census tract where less than 20% of the housing units are renter-occupied or vacant for-rent.

While this dissertation uncovered some facts about the locations of landlords and the impact of concentrated holdings by landlords, it is far from a complete picture of how private rental markets generate inequality. Research from the National Association of Realtors found that only 17 percent of white renters, and 9 percent of Black renters, have the incomes needed to afford the median-priced home in the state where they live (National Association of Realtors 2023). If rent is such a large part of many renters' monthly expenditures, it could crowd out savings for retirement, or for a downpayment in order to purchase a home. As rent has increased faster than renters' wages for the past two decades, these costs may continue to contribute to micro-level inequality, where renters are further squeezed by the cost of rent.

Additionally, there is a need to better conceptualize how the payment of rent contributes to inequality in society. Of course, there has long been a Marxist treatment of "The Housing Question," (see Engels 1954; Harvey 2009). These accounts note the limits of homeownership, and private property in general, in addressing the challenges faced by workers in the capitalist urban context. Additional, Harvey expressly articulates that renters are primarily concerned with the use value of a property, whereas landlords are concerned about its exchange value; this disconnect leads to different experiences within urban space, as Harvey (2009, 171) writes: "we therefore arrive at the fundamental conclusion that the rich can command space whereas the poor are trapped in it." In Canada, most landlords receive very little profit from their rental properties[76] (Statistics Canada 2022). Nonetheless, Hochstenbach (2022) found that among the top percentile of Dutch households by wealth, one in three were landlords. What is the cumulative transfer from renters to landlords? To what extent does this transfer represent the cost of providing rental housing, and to what extent is this renters' surplus that is being captured by landlords?

At the macro level, privately-owned rental housing may generate inequality through numerous pathways. How are rental properties valued, for the purposes of tax assessments — especially when some property owners "milk" their properties in order to maximize profits by minimizing expenses (Mallach 2010)? Do these differences in assessed value result in cumulative

[76] Median family rental income was $2,750 CAD; among landlords who reported positive income, it was $4,880 CAD.

differences in local government revenues, given the importance of property taxation for local government resources? How do differences in rental and homeownership rates at the regional level contribute to intra-regional inequality, in terms of differences in services and local government revenues?

§5.3 — Implications for Policy & Practice

This dissertation has attempted to answer a set of research questions that are of great interest to the general and policymaking public in the United States today. The findings of these papers can be used to inform policy and practice at the local and state levels. Below, I connect my findings to the current state of the practice.

The findings from Paper One regarding the location of landlords relative to their properties means that practitioners may wish to find ways to cultivate local landlords, or to transition properties from the private rental market to other tenure types. Absentee landlords have been found to produce worse outcomes for tenants and properties (Robinson and Steil 2021; Rose and Harris 2021), while corporate landlords are also found to produce worse outcomes for tenants (Travis 2019). While cities are unlikely to be able to enact blanket bans on absentee or corporate landlords, there are policies already enacted in some local jurisdictions that encourage different paths in the future. Responding to the specific concerns about absentee landlords, cities can also implement rental regulations that require a local contact for tenants living in rental housing. Other renter rights, such as the ability to withhold payment for property upkeep, can correct some of the power imbalances that renters face when dealing with landlords who are slow to repair property. Code enforcement is a two-edged sword, since some landlords may choose to abandon the property rather than invest in its maintenance (Bartram 2022). Nonetheless, some cities have implemented programs that provide small landlords with grants to improve property conditions, in exchange for a period of time in which the landlord agrees to rent to low-income renters.[77]

The findings from Paper Two, regarding the size of landlords' holdings, bring to light the challenges with larger landlords. While other studies have highlighted larger landlords' practices

[77] For instance, see the Small Landlord Fund operated by the Urban Redevelopment Authority of Pittsburgh.

with regards to evictions (Immergluck et al. 2020; Robinson and Steil 2021), I find that locally large landlords — i.e., landlords that have a large market share — may contribute to higher and faster-rising rents. This points to the need for policies to counter the effects of the practices associated with larger landlords. Again, cities are unlikely to be able to enact blanket limits on the size of landlords. Nonetheless, there are policies that can stem the tide of increasing portfolio size. In Washington, DC, the "Tenant Opportunity to Purchase Act," and its companion, the "District Opportunity to Purchase Act" allows tenants, and the local government, a right of first refusal when landlords sell their holdings. In Massachusetts, state law 40T gives the Commonwealth a right of first refusal when subsidized multifamily developments approach the end of their affordability covenants, and when owners try to sell those developments. Massachusetts, Rhode Island, Vermont, and New Hampshire lead the nation in providing tenants of manufacturing housing developments the opportunity to purchase the parks themselves (National Consumer Law Center 2022). "Opportunity to purchase" bills have been introduced in Massachusetts, New York, and California, and they represent one avenue for allowing renters to become owners. In the DC case, tenants are allowed to assign their purchase rights to a non-profit, thus converting market-rate rental housing into rental housing that is more likely to remain affordable.

Finally, the findings from Paper Three highlight the extent to which policymakers and practitioners are currently flying blind with regards to understanding their rental market. It is nearly impossible to identify large landlords who operate using multiple different corporate names (see Chapter 3, this dissertation, Gomory 2021; Hangen and O'Brien 2022). In states that do not release information about owners and agents for corporations, this process is even more difficult,[78] as researchers and practitioners have no way to know who is behind "123 Main Street LLC." States have the ability to rectify this shortcoming, by releasing more information about corporations and their beneficial owners. Local governments can also enact rental registries that address the shortcomings of their states. These rental registries would, ideally, require owners of all rental properties to register with the local government, and to provide information about themselves and

[78] OpenCorporates, whose data is used in Paper Two of this dissertation, provides an "Open Company Data Index," which scores jurisdictions on how open their data is regarding corporations. No US state scores above an 80, out of 100 — worse than the entire United Kingdom. 30 states score lower than 25 out of 100. See https://opencorporates.com/registers

the management of the rental property. Local governments can work with their utility providers to identify unlawful building-level subdivisions[79] that may indicate informal rental properties. Greater transparency about the ownership and management of rental properties would mean that tenants and local governments could resolve maintenance issues more easily, and could more readily identify tenants and landlords in need of assistance. This became particularly apparent during the COVID-19 pandemic, when localities had difficulty in deploying emergency rental assistance due to the inability to accurately identify renters and landlords (Demsas 2021).

Taken together, then, states and localities can use these findings to implement tenant protections that advance the interest of renters and local governments themselves. These policies include better tenants' rights in responding to building upkeep, pathways for tenant rights-of-first-refusal, and rental registries to better identify rental properties and landlords. These would represent modest steps in the direction of addressing the underlying concerns about inequalities in the rental market, and the changes to rental market efficiency that may mean more strain for renters in the future. The time to act on these policies is now, so that we may better prepare for continued changes in the rental market tomorrow.

[79] Scraton, PA, has tried to identify buildings with a "single-family" use code that have multiple gas or electric meters.

APPENDIX A | RENTAL REGISTRIES

This appendix summarizes the status and availability of rental registries among the 50 most populous cities in the United States. Cities without a rental registry are not listed. The Notes column denotes whether I was able to obtain the registry, and, if applicable, why it was not used. In general, I did not use any registry that was predominantly focused on multifamily.

Table A.1: Rental Registries Among 50 Most Populous US Cities

Name	Registry Title	Availability	Exclusion	Notes
New York city, New York	Multiple Dwelling Registration	Available online through NYC OpenData	Owner-occupied one/two unit rentals	See Ellen, Harwood, and O'Regan (2022) and Watson and Ziv (2022). I excluded NYC due to its focus on multifamily.
Houston city, Texas	Multi-Family Rental Property Registration	PDF list of all buildings available online.	One/two unit rentals	Obtained via Right-to-Know request. Did not use do to exclusion of one- and two-unit properties
Phoenix city, Arizona	Rental Properties in Phoenix must be registered with County Assessor	For sale - $425.	Family occupied excluded	Opted not to pay to obtain this dataset, in part because of limited coverage in Tucson (Pima County)
Philadelphia city, Pennsylvania	Rental License	OpenData Philly	Family occupied excluded	Manager information is not provided online.
San Diego city, California	Rental Unit Business Tax	Not publicly available	Owner/family occupied excluded	Did not pursue do to its collection via tax office.
Dallas city, Texas	Single Family Rental Program and Multi-Tenant Rental Registration	Open records request	Family occupied excluded	Obtained via Right-to-Know request.
Indianapolis city (balance), Indiana	Landlord Registration Program	Available for download online, but very incomplete	None	Obtained via Right-to-Know request, but compliance is extremely low.
Columbus city, Ohio	Rental Registration	Available online	None	Obtained via Franklin County Auditor Website, via "WebReporter."
Fort Worth city, Texas	Rental Registration	Submitted public records request	One/Two Family Buildings with no code violations	City claimed that they had no responsive records.
Charlotte city, North Carolina	Rental Registration	Submitted public records request	None	City denied request. It is possible that the City no longer collects this

				information, per SB 326, passed in 2016.
Seattle city, Washington	Rental Registration and Inspection Code	Online/Submitted public records request.	Short term rentals	City provided all requested information, including phone numbers and email addresses of registrants.
Detroit city, Michigan	Rental Certificate of Compliance	Submitted public records request	None	City would not provide unit counts.
Washington city, District of Columbia	Rental Housing Business License	Open Data DC	None	Unit counts and corporate ownership information was acquired via scraping of "scout.dcra.gov," and linked to OpenCorporates data.
Boston city, Massachusetts	Rental Registration	Acquired via contact with the city government.	None	
Nashville-Davidson metropolitan government (balance), Tennessee	Landlord Permit	Partial information available online. Submitted public records request.	None	Tennessee resident required to submit request. Did not provide parcel number, so scraped that information from website. Provides phone numbers and email addresses of registrants.
Portland city, Oregon	Residential Rental Registration Program	Submitted public records request.	None	Due to private nature of tax records, the city provided me the information about rental unit locations and 9-digit zip codes of landlords. Did not include parcel numbers.
Las Vegas city, Nevada	Apartment License	Available for download online	4 or fewer units	Did not use due to exclusion of smaller rentals.
Baltimore city, Maryland	Rental Property Registration and Licensing	Submitted public records request	None	City partially responded to public records request, but would not provide unit counts.
Louisville/Jefferson County metro government (balance), Kentucky	Rental Housing Registry	Submitted public records request.	Units where rent is not paid	Required Kentucky resident to submit public records request. City provided addresses for rental properties, but not unit counts or owner information.

Milwaukee city, Wisconsin	Property Registration Program	Submitted public records request.	Owner-occupied	City did not respond to records request.
Tucson city, Arizona	Rental Registry	Via Pima County Assessor, available for download	Family occupied excluded	Doesn't seem to include number of units. Extremely limited availability.
Fresno city, California	Rental Registration	Submitted public records request.	None	City provided list of rental addresses and unit counts, but no parcel numbers or owner information.
Sacramento city, California	Rental Housing Inspection Program	Submitted public records request.	Those regularly inspected by another agency, or new properties (less than 5 years) exempt.	City denied public records request.
Mesa city, Arizona	Rental Properties in Mesa must be registered with County Assessor	Same as Phoenix above	Family occupied excluded	Part of same information from Phoenix; did not acquire.
Kansas City city, Missouri	Healthy Homes Rental Inspection Program	Submitted public records request.	Voucher-holder inhabited homes exempted	City responded to public records request. Did not include full mailing addresses (only city, state, and ZIP codes).
Long Beach city, California	Residential Property Rental Business License	Submitted public records request.	Fewer than 4 units exempted.	City provided list of rental addresses and unit counts, but no parcel numbers or owner information.
Omaha city, Nebraska	Rental Registration	Submitted public records request.	None	Includes those units within 3 miles of the border of Omaha. City responsive of request.
Minneapolis city, Minnesota	Rental License	Open Minneapolis	None	Obtained and used throughout dissertation. Incudes manager information, and email address and phone numbers for both owners and managers.
Cleveland city, Ohio	Residential Rental Property Disclosure	Submitted public records request	None.	City did not provide unit counts.
Arlington city, Texas	Multifamily and Extended Stay License	Submitted public records request.	Excludes single-family	City claimed it had no responsive records.

APPENDIX B | APPENDIX FOR CHAPTER 3

§ B1: ZORI Coverage

It should be noted that Zillow provides rents for different ZIP codes during different data downloads. Thus, the construction of these figures and the estimates of rents come from four different downloads of Zillow data, two in 2022 (which contain rents for 2014) and two in early 2023 (which contain rents for 2022). Including all four of these downloads increases the geographic coverage substantially.

Zillow coverage is not evenly distributed across cities in various ways. First, as shown in Table B1.1, ZORI has better or worse coverage in different cities. While all cities have a majority of ZIPs in ZORI for 2022, that percentage ranges from 65% of ZIPs in Kansas City, to 96% of ZIPs in Washington, DC. In Boston, Seattle, and Washington, DC a majority of ZIPs have ZORI data requisite for the income shock regressions, meaning that there is ZORI data for both 2014 and 2022.

Table B1.1: Zillow Coverage by City

City	ZORI in Both 2014 and 2022	ZORI in 2022 Only	No ZORI coverage	Total
Boston	21 (65.6%)	6 (18.8%)	5 (15.6%)	32 (100.0%)
Columbus	7 (17.5%)	27 (67.5%)	6 (15.0%)	40 (100.0%)
Dallas	24 (46.2%)	19 (36.5%)	9 (17.3%)	52 (100.0%)
Kansas City	9 (20.9%)	19 (44.2%)	15 (34.9%)	43 (100.0%)
Minneapolis	9 (40.9%)	9 (40.9%)	4 (18.2%)	22 (100.0%)
Seattle	19 (67.9%)	6 (21.4%)	3 (10.7%)	28 (100.0%)
Washington, DC	15 (65.2%)	7 (30.4%)	1 (4.3%)	23 (100.0%)

Note: Counts are of ZIPs with ZORI data.

Additionally, ZORI coverage lacks balance along the lines of HHI, number of renter-occupied housing units, median household income, and distance to the CBD, as shown in Figures B1.1, B1.2, B1.3, and B1.4. In general, ZORI has better coverage in areas with lower HHI, with more renter-occupied housing units, and closer to the CBD. No clear patterns emerge with regards to the relationship between ZORI coverage and household income. In general, this means that the findings of regressions based on the geography with ZORI coverage will differ from all ZIPs, though the balance figures below do not indicate that ZORI coverage is biased towards places with higher HHI, the main concern in this paper.

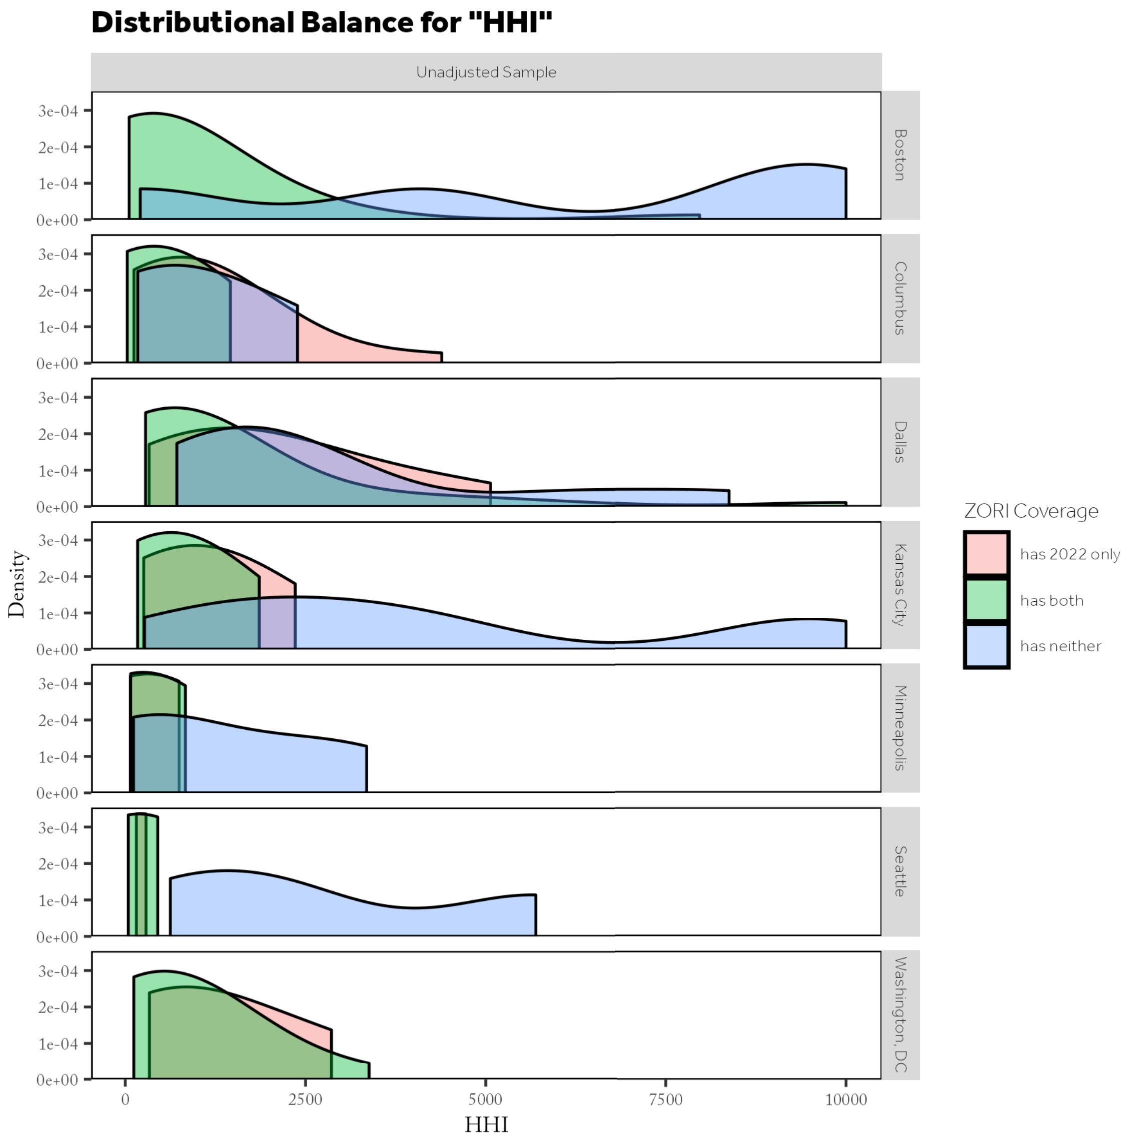

Figure B1.1: ZORI HHI Balance

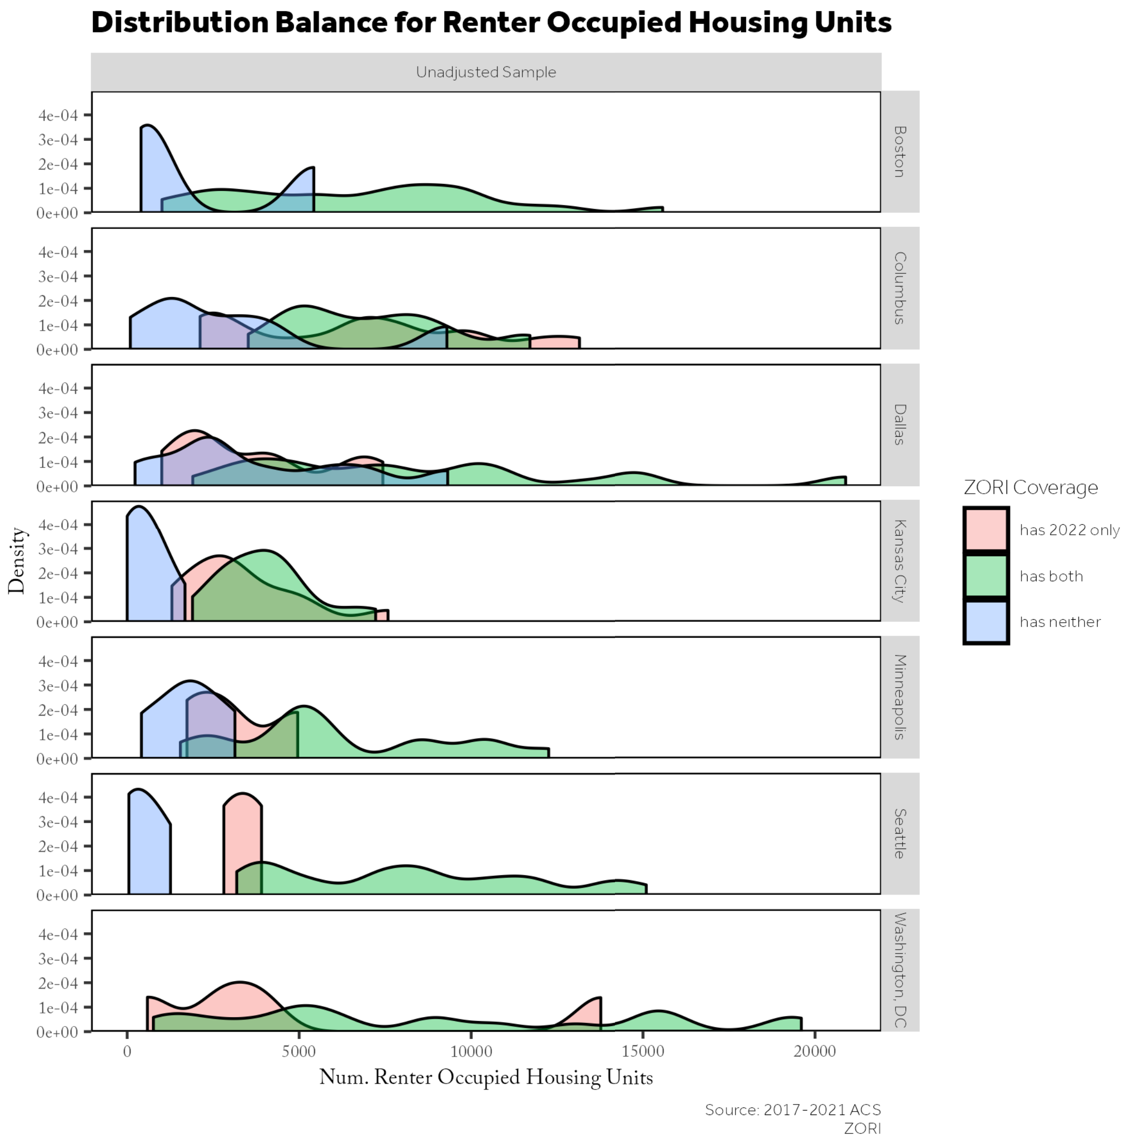

Figure B1.2: ZORI Housing Unit Balance

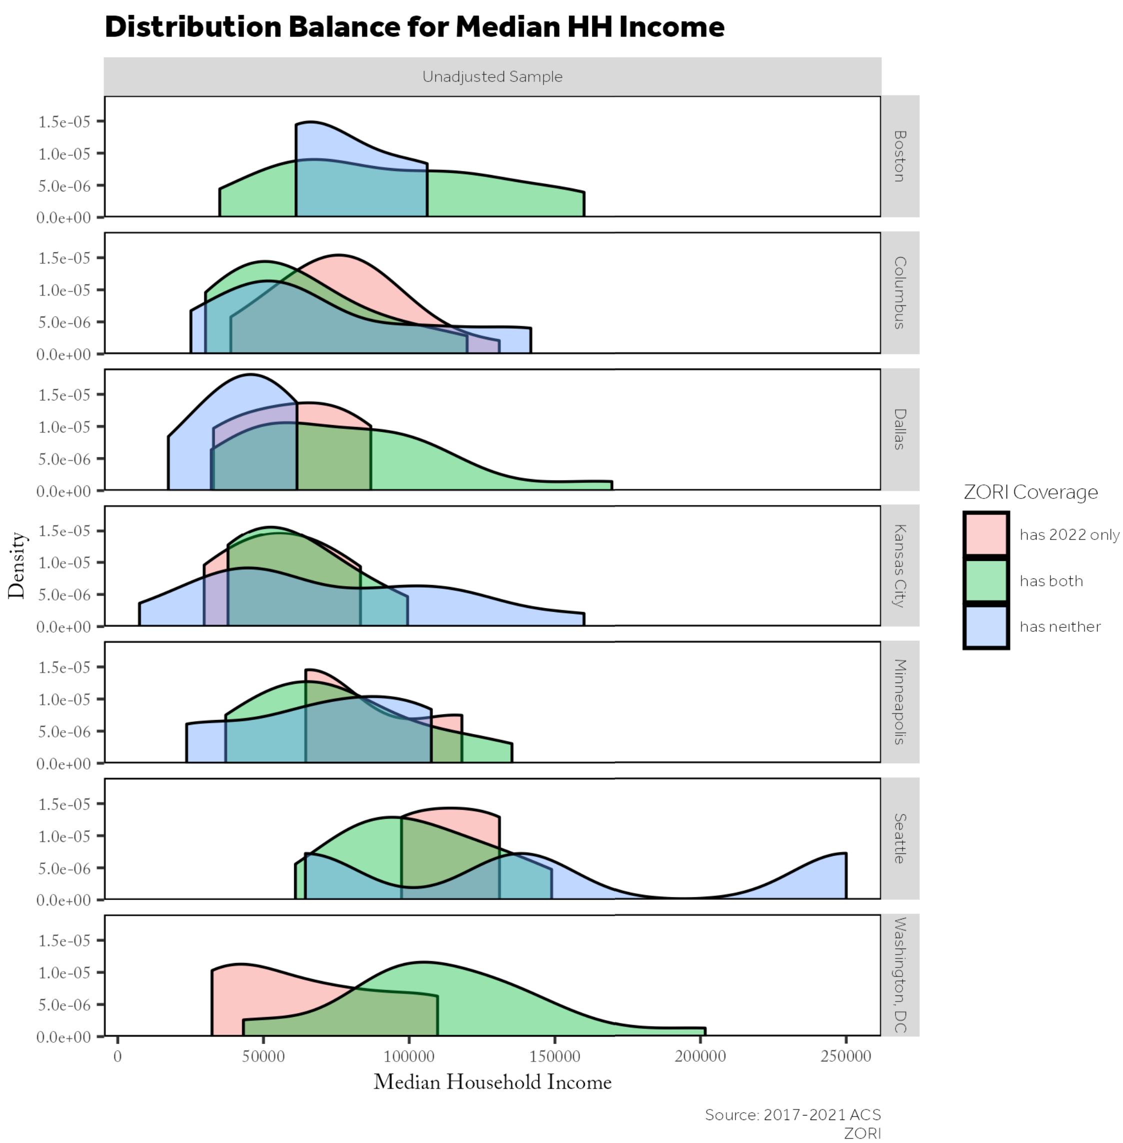

Figure B1.3: ZORI Household Income Balance

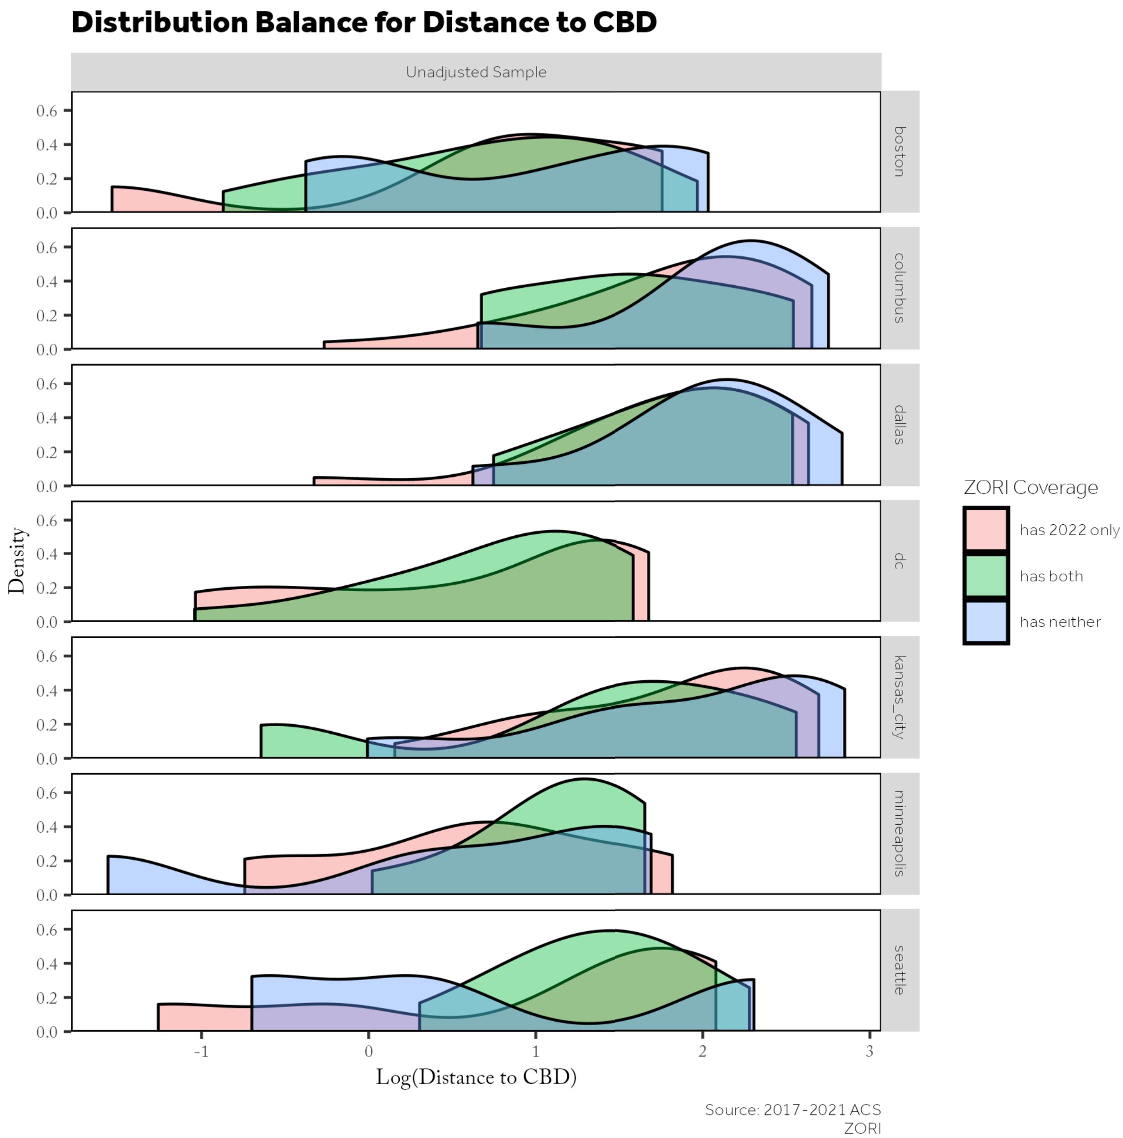

Figure B1.4: ZORI Distance Balance

§ B2 Identification of Landlord Grouping

Figure B2.1 shows the impact of the various means used to group distinct rental registry records into a coherent set of landlords. The figure shows the relative number of landlords identified at each stage of the deduplication process — where each successive stage of the deduplication process shows the number of "unique" landlords relative to the registered rental properties. As is quite clear, the bulk of the deduplication process is done through simple groupings: if all of the owner-related information is the same between two entries, then they are grouped and treated as one entity. Cleaning the data generally provides only a small improvement in the number of grouped landlords. Interestingly, in Boston, the number of unique landlords actually increases after the data has been cleaned. I attribute this to the fact that the data cleaning process involves the removal of management company information, leading to groups that had appeared to be controlled by the same entity to be split apart, as new information is incorporated and management information is removed. In general, the final stage of the process — Dedupe — provides a modest but important improvement to the deterministic grouping process based on exact information. In most cities, Dedupe reduces the number of unique landlord groupings by approximately 20 percent. At the low end, Dedupe reduces the number of unique landlords by only 7 percent in Kansas City. On the other hand, Dedupe further reduces the number of unique landlords in Columbus by 26 percent.

It is also important to note, briefly, that the use of a relative scale in Figure B2.1 obscures the significant differences in the number of landlords. For instance, in Columbus, there are 35,680 unique landlords identified at the first grouping stage, compared to only 5,360 landlords in Dallas. This means that, in Columbus, there are 1.7 rental properties and 3.5 rental units for every registered landlord, compared to 2.03 rental properties and 46 rental units for every registered landlord in Dallas. These disparities are more clear when examining rental concentration at the city level.

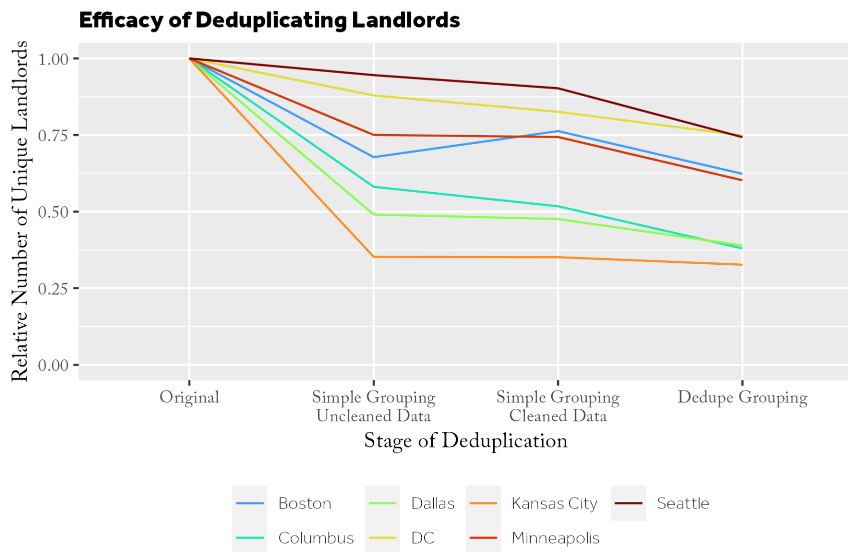

Figure B2.1: Deduping Efficacy

§ B3 Income Shock Robustness

Tables B3.1, B3.2, and B3.3 repeat the regressions from Table 3.6. Each table includes an additional interaction term. Table B3.1 includes an additional interaction term between expected income growth and owner occupancy rate; B3.2 between expected income growth and fraction of the population with a bachelor's degree; B3.3 between expected income growth and median house value. In addition to the discussion in Section 3.4.3, I note here that ZIPs with higher-value houses, and ZIPs with more college grads, have lower rent hikes when wages increase.

Table B3.1: Income Shock Regressions, Owner Occupancy Interaction

		Rent Change	
	Model VII: Delta ZORI	Model VIII: Delta ACS	Model IX: Delta ACS
Log(HHI)	-0.471	0.392	-1.398***
	-1.108	-1.367	-0.328
Log(HHI)*Expected Income Growth	1.978	-1.698	6.409***
	-4.831	-6.193	-1.495
Expected Income Growth	-21.033	-36.698	-96.781***
	-28.688	-38.559	-10.999
Frac. Owner-Occupied Housing Units*Expected Income Growth	28.12	27.38	68.301***
	-21.93	-31.603	-23.714
Frac. White	0.189	0.066	-0.033
	-0.156	-0.126	-0.1
Frac. 25+ w/ Bachelor Degree	-0.634***	-0.440***	-0.509***
	-0.169	-0.135	-0.179
Log(Median Household Income)	-0.018	0.151*	0.154*
	-0.093	-0.086	-0.091
Log(Number of Renter-Occupied Housing Units)	-0.024	0.028	-0.009
	-0.037	-0.035	-0.012
Log(Distance to CBD)	0.127***	0.001	-0.017
	-0.032	-0.023	-0.025
Frac. Vacant Housing Units	-0.004	0.004	-0.283
	-0.431	-0.359	-0.244
Frac. Owner-Occupied Housing Units	-6.548	-6.295	-15.327***
	-5.093	-6.874	-5.178
Log(Median House Value)	0.013	0.126***	0.159**
	-0.067	-0.037	-0.062
Constant	5.693	5.152	18.429***
	-6.603	-8.815	-2.213
N	104	104	228
R^2	0.616	0.449	0.384
Adjusted R^2	0.565	0.376	0.35
Residual Std. Error	0.140 (df = 91)	0.123 (df = 91)	0.147 (df = 215)
F Statistic	12.151*** (df = 12; 91)	6.173*** (df = 12; 91)	11.167*** (df = 12; 215)
Note:	***p < .01; **p < .05; *p < .1		
	Standard Errors Clustered At the City Level		

Table B3.2: Income Shock Regressions, College Graduate Interaction

	Rent Change		
	Model I: Delta ZORI	Model II: Delta ACS	Model III: Delta ACS
Log(HHI)	-0.787	0.836	-1.522^{**}
	-0.927	-1.538	-0.687
Log(HHI)*Expected Income Growth	3.37	-3.702	7.003^{**}
	-4.041	-6.988	-3.137
Expected Income Growth	-18.979	17.885	-44.877^{*}
	-33.502	-39.79	-25.957
Frac. 25+ w/ Bachelor Degree*Expected Income Growth	5.328	-60.469^{*}	-52.107^{**}
	-13.661	-33.159	-23.557
Frac. White	0.219	0.014	-0.02
	-0.17	-0.105	-0.096
Frac. 25+ w/ Bachelor Degree	-1.871	12.951^{*}	10.934^{**}
	-3.177	-7.374	-5.292
Log(Median Household Income)	-0.052	0.167^{**}	0.158^{*}
	-0.092	-0.076	-0.088
Log(Number of Renter-Occupied Housing Units)	-0.027	0.017	0.002
	-0.035	-0.027	-0.01
Log(Distance to CBD)	0.119^{***}	-0.01	-0.039
	-0.033	-0.018	-0.024
Frac. Vacant Housing Units	-0.076	0.188	0.022
	-0.493	-0.337	-0.251
Frac. Owner-Occupied Housing Units	-0.045	-0.232	-0.275^{**}
	-0.174	-0.162	-0.122
Log(Median House Value)	0.022	0.107^{**}	0.154^{**}
	-0.06	-0.048	-0.066
Constant	5.471	-6.741	6.877
	-7.534	-8.597	-5.364
N	104	104	228
R^2	0.608	0.498	0.382
Adjusted R^2	0.556	0.432	0.347
Residual Std. Error	0.142 (df = 91)	0.118 (df = 91)	0.148 (df = 215)
F Statistic	11.747^{***} (df = 12; 91)	7.529^{***} (df = 12; 91)	11.061^{***} (df = 12; 215)

Note: $^{***}p < .01$; $^{**}p < .05$; $^{*}p < .1$
Standard Errors Clustered At the City Level

Table B3.3: Income Shock Regressions, Median House Value Interaction

	Rent Change		
	Model I: Delta ZORI	Model II: Delta ACS	Model III: Delta ACS
Log(HHI)	-0.633	2.07	-0.124
	-0.742	-1.357	-0.843
Log(HHI)*Expected Income Growth	2.696	-9.315	0.628
	-3.229	-6.175	-3.853
Expected Income Growth	74.696	501.279***	412.139***
	-47.081	-135.547	-121.576
Log(Median House Value)*Expected Income Growth	-6.985**	-37.979***	-35.292***
	-3.099	-9.012	-8.296
Frac. White	0.186	0.003	-0.006
	-0.163	-0.123	-0.088
Frac. 25+ w/ Bachelor Degree	-0.598***	-0.302**	-0.501***
	-0.174	-0.124	-0.179
Log(Median Household Income)	-0.035	0.164**	0.177**
	-0.096	-0.069	-0.074
Log(Number of Renter-Occupied Housing Units)	-0.027	0.018	0.008
	-0.04	-0.028	-0.011
Log(Distance to CBD)	0.115***	-0.021	-0.045**
	-0.029	-0.013	-0.02
Frac. Vacant Housing Units	0.177	0.617**	0.509*
	-0.476	-0.303	-0.27
Frac. Owner-Occupied Housing Units	-0.007	-0.16	-0.217*
	-0.163	-0.156	-0.128
Log(Median House Value)	1.627**	8.447***	7.894***
	-0.752	-1.975	-1.808
Constant	-16.269	-112.936***	-93.669***
	-11.418	-29.849	-26.744
N	104	104	228
R^2	0.616	0.576	0.456
Adjusted R^2	0.565	0.52	0.425
Residual Std. Error	0.140 (df = 91)	0.108 (df = 91)	0.139 (df = 215)
F Statistic	12.160*** (df = 12; 91)	10.311*** (df = 12; 91)	14.998*** (df = 12; 215)
Note:	***p < .01; **p < .05; *p < .1		
	Standard Errors Clustered At the City Level		

This page intentionally left blank.

REFERENCES

Abood, Maya. 2018. "Wall Street Landlords Turn American Dream into a Nightmare." ACCE Institute; Americans for Financial Reform; Public Advocates. www.acceinstitute.org/wall_street_landlords_turn_american_dream_into_american_nightmare.

Airgood-Obrycki, Whitney, and Sophia Wedeen. 2022. "Rental Deserts Perpetuate Socioeconomic and Racial Segregation." Joint Center for Housing Studies. August 4. www.jchs.harvard.edu/blog/rental-deserts-perpetuate-socioeconomic-and-racial-segregation.

Andris, Clio, Xi Liu, Jessica Mitchell, James O'Dwyer, and Jeremy Van Cleve. 2021. "Threads Across the Urban Fabric: Youth Mentorship Relationships as Neighborhood Bridges." *Journal of Urban Affairs* 43 (1): 77–92. doi:10.1080/07352166.2019.1662726.

Arnott, Richard. 1995. "Time for Revisionism on Rent Control?" *Journal of Economic Perspectives* 9 (1): 99–120. doi:10.1257/jep.9.1.99.

Arribas-Bel, Daniel. 2014. "Accidental, Open and Everywhere: Emerging Data Sources for the Understanding of Cities." *Applied Geography*, The New Urban World, 49 (May): 45–53. doi:10.1016/j.apgeog.2013.09.012.

Austermuhle, Martin. 2020. "A Tool To Keep Track Of Rent-Controlled Housing In D.C. Has Been Delayed For Years." *WAMU NPR*, December 22. www.npr.org/local/305/2020/12/22/949179415/a-tool-to-keep-track-of-rent-controlled-housing-in-d-c-has-been-delayed-for-years.

Autor, David H., Christopher J. Palmer, and Parag A. Pathak. 2014. "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts." *Journal of Political Economy* 122 (3). The University of Chicago Press: 661–717. doi:10.1086/675536.

Bartik, Timothy J. 1991. "Who Benefits from State and Local Economic Development Policies?" Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. ideas.repec.org//b/upj/ubooks/wbsle.html.

Bartram, Robin. 2022. *Stacked Decks: Building Inspectors and the Reproduction of Urban Inequality*. Chicago ; Illinois: The University of Chicago Press.

Baum-Snow, Nathaniel, and Fernando Ferreira. 2015. "Chapter 1 - Causal Inference in Urban and Regional Economics." In *Handbook of Regional and Urban Economics*, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 5:3–68. Handbook of Regional and Urban Economics. elsevier. doi:10.1016/B978-0-444-59517-1.00001-5.

Beaty, Andrea, and Vishal Shankar. 2022. "Tenant Organizers Call On Biden To Tackle Rent Inflation." *Revolving Door Project*, October 12. therevolvingdoorproject.org/tenant-organizers-call-on-biden-to-tackle-rent-inflation/.

Boeing, Geoff, and Paul Waddell. 2017. "New Insights into Rental Housing Markets Across the United States: Web Scraping and Analyzing Craigslist Rental Listings." *Journal of Planning Education and Research* 37 (4): 457–76. doi:10.1177/0739456X16664789.

Boeing, Geoff, Jake Wegmann, and Junfeng Jiao. 2020. "Rental Housing Spot Markets: How Online Information Exchanges Can Supplement Transacted-Rents Data." *Journal of Planning Education and Research*, February, 0739456X2090443. doi:10.1177/0739456X20904435.

Briz-Redón, Álvaro, Francisco Martinez-Ruiz, and Francisco Montes. 2020. "Reestimating a Minimum Acceptable Geocoding Hit Rate for Conducting a Spatial Analysis." *International Journal of Geographical Information Science* 34 (7): 1283–1305. doi:10.1080/13658816.2019.1703994.

Brown, Peter J. B., and Stephen Hincks. 2008. "A Framework for Housing Market Area Delineation: Principles and Application." *Urban Studies* 45 (11): 2225–47. doi:10.1177/0042098008095866.

Bruch, Elizabeth, and Joffre Swait. 2019. "Choice Set Formation in Residential Mobility and Its Implications for Segregation Dynamics." *Demography* 56 (5): 1665–92. doi:10.1007/s13524-019-00810-5.

Bureau of Labor Statistics. 2021. "Average Annual Expenditures and Characteristics of All Consumer Units, Consumer Expenditure Surveys, 2021-2021." www.bls.gov/cex/tables.htm#topline.

Carrillo, Laura, Mary Pattillo, Erin Hardy, and Dolores Acevedo-Garcia. 2016. "Housing Decisions Among Low-Income Hispanic Households in Chicago." *Cityscape* 18 (2). US Department of Housing and Urban Development: 109–50. www.jstor.org/stable/26328259.

Carswell, Andrew T. 2018. "Living Where You Work: Determining the Value-Added Nature of the on-Site Residential Property Manager." *Facilities* 36 (5/6). Emerald Publishing Limited: 258–71. doi:10.1108/F-12-2016-0106.

Charles, Suzanne Lanyi. 2020a. "A Latent Profile Analysis of Suburban Single-Family Rental Housing (SFR) Neighborhoods." *Housing Policy Debate* 30 (2): 205–27. doi:10.1080/10511482.2019.1657927.

———. 2020b. "The Financialization of Single-Family Rental Housing: An Examination of Real Estate Investment Trusts' Ownership of Single-Family Houses in the Atlanta Metropolitan Area." *Journal of Urban Affairs* 42 (8): 1321–41. doi:10.1080/07352166.2019.1662728.

Charvel, Roberto. 2023. "The Digital Entrepreneur's Future Impact on the Housing Cycle." In *Bringing Digitalization Home Symposium*. Cambridge, MA: Harvard Joint Center for Housing Studies. www.jchs.harvard.edu/research-areas/working-papers/digital-entrepreneurs-future-impact-housing-cycle.

Cherry, R., and E. J. Ford Jr. 1975. "Concentration of Rental Housing Property and Rental Housing Markets in Urban Areas." *Real Estate Economics* 3 (1): 7–16. doi:10.1111/1540-6229.00129.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106 (4): 855–902. doi:10.1257/aer.20150572.

Chetty, Raj, Matthew O. Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B. Fluegge, Sara Gong, et al. 2022. "Social Capital II: Determinants of Economic Connectedness." *Nature* 608 (7921, 7921). Nature Publishing Group: 122–34. doi:10.1038/s41586-022-04997-3.

Coles, Peter, Michael Egesdal, Ingrid Gould Ellen, Xiaodi Li, and Arun Sundararajan. 2018. "Airbnb Usage Across New York City Neighborhoods: Geographic Patterns and Regulatory Implications*." In *The Cambridge Handbook of the Law of the Sharing Economy*, edited by John J. Infranca, Michèle Finck, and Nestor M. Davidson, 108–28. Cambridge Law Handbooks. Cambridge: Cambridge University Press. doi:10.1017/9781108255882.009.

Collinson, Robert, John Eric Humphries, Nicholas S. Mader, Davin K. Reed, Daniel I. Tannenbaum, and Winnie van Dijk. 2022. "Eviction and Poverty in American Cities." Working Paper. Working Paper Series. August. doi:10.3386/w30382.

Cosman, Jacob, and Luis Quintero. 2021. "Fewer Players, Fewer Homes: Concentration and the New Dynamics of Housing Supply." SSRN Scholarly Paper ID 3303984. Rochester, NY: Social Science Research Network. doi:10.2139/ssrn.3303984.

Cossyleon, Jennifer E., Philip M. E. Garboden, and Stefanie DeLuca. 2020. "Recruiting Opportunity Landlords: Lessons from Landlords in Maryland." Poverty & Race Research Action Council. www.jstor.org/stable/resrep27291.

Coulton, Claudia, April Urban, Francisca García-Cobián Richter, Michael Henderson, Michael Schramm, Jeesoo Jeon, and Robert Fischer. 2020. "Characteristics of Rental Properties and Landlords in Cleveland." Case Western Reserve University Center on Urban Poverty and Community Development. case.edu/socialwork/povertycenter/sites/case.edu.povertycenter/files/2020-10/Landlords_09022020r_accessible%20%281%29.pdf.

Crook, A. D. H. (Tony), Ed Ferrari, and Peter A. Kemp. 2012. "Knowing the Area: The Management of Market and Business Risks by Private Landlords in Scotland." *Urban Studies* 49 (15): 3347–63. doi:10.1177/0042098012442415.

D'Lima, Walter, and Paul Schultz. 2021. "Residential Real Estate Investments and Investor Characteristics." *The Journal of Real Estate Finance and Economics* 63 (3): 354–93. doi:10.1007/s11146-020-09771-8.

Daepp, Madeleine I. G. 2022. "Small-Area Moving Ratios and the Spatial Connectivity of Neighborhoods: Insights from Consumer Credit Data." *Environment and Planning B: Urban Analytics and City Science* 49 (3): 1129–46. doi:10.1177/23998083211051742.

Dash Nelson, Garrett, and Alasdair Rae. 2016. "An Economic Geography of the United States: From Commutes to Megaregions." Edited by Joshua L Rosenbloom. *PLOS ONE* 11 (11): e0166083. doi:10.1371/journal.pone.0166083.

DelPrete, Mike. 2023. "Emerging Models in Residential Real Estate or How to Lose a Billion Dollars in Real Estate Tech." In *Bringing Digitalization Home Symposium*. Cambridge, MA: Harvard Joint Center for Housing Studies. www.jchs.harvard.edu/research-areas/working-papers/emerging-models-residential-real-estate-or-how-lose-billion-dollars.

Demsas, Jerusalem. 2021. "One Way the US Could Have Prevented the Fight over the Eviction Moratorium." *Vox*, August 4. www.vox.com/2021/8/4/22606530/eviction-moratorium-rent-relief-rental-registry.

Desmond, Matthew, and Nathan Wilmers. 2019. "Do the Poor Pay More for Housing? Exploitation, Profit, and Risk in Rental Markets." *American Journal of Sociology* 124 (4): 1090–1124. doi:10.1086/701697.

Dezember, Ryan, and Laura Kusisto. 2017. "Meet Your New Landlord: Wall Street." *Wall Street Journal: Markets*, July 21. www.wsj.com/articles/meet-your-new-landlord-wall-street-1500647417.

Ding, Lei, Jackelyn Hwang, and Eileen Divringi. 2016. "Gentrification and Residential Mobility in Philadelphia." *Regional Science and Urban Economics* 61 (November): 38–51. doi:10.1016/j.regsciurbeco.2016.09.004.

Dowdall, Emily, Ira Goldstein, Bruce Katz, and Benjamin Preis. 2021. "Investor Home Purchases and the Rising Threat to Owners and Renters: Tales from 3 Cities." Nowak Metro Finance Lab at Drexel University. drexel.edu/nowak-lab/publications/reports/investor-home-purchases/.

Early, Dirk W., Paul E. Carrillo, and Edgar O. Olsen. 2019. "Racial Rent Differences in U.S. Housing Markets: Evidence from the Housing Voucher Program." *Journal of Regional Science* 59 (4): 669–700. doi:10.1111/jors.12422.

Ellen, Ingrid Gould, Katharine WH Harwood, and Katherine O'Regan. 2022. "Equity in the Housing Choice Voucher Program: The Role of Corporate Landlords." presented at the 2022 APPAM fall research conference, Washington, D.C.

Ellen, Ingrid Gould, and Katherine M. O'Regan. 2011. "How low income neighborhoods change: Entry, exit, and enhancement." *Regional Science and Urban Economics* 41 (2): 89–97. doi:10.1016/j.regsciurbeco.2010.12.005.

Engels, Friedrich. 1954. *The Housing Question*. Moscow: Progress Publishers.

Ferrer, Alexander. 2021. "Beyond Wall Street Landlords." Strategic Actions for a Just Economy. www.saje.net/wp-content/uploads/2021/03/Final_A-Just-Recovery-Series_Beyond_Wall_Street.pdf.

Fields, Desiree. 2018. "Constructing a New Asset Class: Property-led Financial Accumulation After the Crisis." *Economic Geography* 94 (2). routledge: 118–40. doi:10.1080/00130095.2017.1397492.

———. 2022. "Automated Landlord: Digital Technologies and Post-Crisis Financial Accumulation." *Environment and Planning A: Economy and Space* 54 (1). SAGE Publications Ltd: 160–81. doi:10.1177/0308518X19846514.

———. 2023. "Big Capital, Precision Technologies, and Structural Power in the US Housing Market." In *Bringing Digitalization Home Symposium*. Cambridge, MA: Harvard Joint

Center for Housing Studies. www.jchs.harvard.edu/research-areas/working-papers/big-capital-precision-technologies-and-structural-power-us-housing.

Fields, Desiree, Rajkumar Kohli, and Alex Schafran. 2016. "The Emerging Economic Geography of Single-Family Rental Securitization." Working Paper 2016-02. Federal Reserve Bank of San Francisco. www.frbsf.org/community-development/publications/working-papers/2016/january/emerging-economic-geography-single-family-rental-securitization/.

Fields, Desiree, and Manon Vergerio. 2022. "Corporate Landlords and Market Power: What Does the Single-Family Rental Boom Mean for Our Housing Future?" UC Berkeley. escholarship.org/uc/item/07d6445s.

Fine, Michelle. 1994. "Working the Hyphens: Reinventing Self and Other in Qualitative Research." In *Handbook of Qualitative Research*, 70–82. Thousand Oaks, CA, US: Sage Publications, Inc.

Fisher, Glenn W. 1996. *The Worst Tax? A History of the Property Tax in America*. Studies in Government and Public Policy. Lawrence, Kan: University Press of Kansas.

Folch, David C., Seth E. Spielman, and Robert Manduca. 2018. "Fast Food Data: Where User-Generated Content Works and Where It Does Not." *Geographical Analysis* 50 (2): 125–40. doi:10.1111/gean.12149.

Fowler, Christopher S., Nathan Frey, David C. Folch, Nicholas Nagle, and Seth Spielman. 2020. "Who Are the People in My Neighborhood?: The 'Contextual Fallacy' of Measuring Individual Context with Census Geographies." *Geographical Analysis* 52 (2): 155–68. doi:10.1111/gean.12192.

Freeman, Lance, Jackelyn Hwang, Tyler Haupert, and Iris Zhang. 2023. "Where Do They Go? The Destinations of Residents Moving from Gentrifying Neighborhoods." *Urban Affairs Review*, April. SAGE Publications Inc, 10780874231169921. doi:10.1177/10780874231169921.

Freemark, Yonah, Eleanor Noble, and Yipeng Su. 2021. "Who Owns the Twin Cities?" The Urban Institute. www.urban.org/sites/default/files/publication/104355/who-owns-the-twin-cities_2.pdf.

Freemark, Yonah, Justin Steil, and Kathleen Thelen. 2020. "Varieties of Urbanism: A Comparative View of Inequality and the Dual Dimensions of Metropolitan Fragmentation." *Politics & Society* 48 (2). SAGE Publications Inc: 235–74. doi:10.1177/0032329220908966.

Garboden, Philip M. E., and Eva Rosen. 2018. "Talking to Landlords." *Cityscape* 20 (3). US Department of Housing and Urban Development: 281–91. www.jstor.org/stable/26524885.

———. 2022. "By the Book: Considering the Costs and Benefits of Professional Management." presented at the Urban Affairs Association, Washington, DC, April 13.

Garcia-López, Miquel-Àngel, Jordi Jofre-Monseny, Rodrigo Martínez-Mazza, and Mariona Segú. 2020. "Do Short-Term Rental Platforms Affect Housing Markets? Evidence from Airbnb in Barcelona." *Journal of Urban Economics* 119 (September): 103278. doi:10.1016/j.jue.2020.103278.

Garriga, Carlos, Pedro Gete, and Athena Tsouderou. 2023. "The Economic Effects of Real Estate Investors." *Real Estate Economics* 51 (3): 655–85. doi:10.1111/1540-6229.12427.

Gilderbloom, John I. 1985. "Social Factors Affecting Landlords in the Determination of Rent." *Urban Life* 14 (2). SAGE Publications: 155–79. doi:10.1177/0891241685014002002.

———. 1989. "Socioeconomic Influences on Rentals for U.S. Urban Housing: Assumptions of Open Access to a Perfectly Competitive 'Free Market' Are Confronted with the Facts." *The American Journal of Economics and Sociology* 48 (3). [American Journal of Economics and Sociology, Inc., Wiley]: 273–92. www.jstor.org/stable/3487363.

Gilderbloom, John I., and Richard P. Appelbaum. 1987. "Toward a Sociology of Rent: Are Rental Housing Markets Competitive?*." *Social Problems* 34 (3): 261–76. doi:10.2307/800766.

Glaeser, Edward L., and Joseph Gyourko. 2018. "The Economic Implications of Housing Supply." *Journal of Economic Perspectives* 32 (1): 3–30. doi:10.1257/jep.32.1.3.

Gomory, Henry. 2021. "The Social and Institutional Contexts Underlying Landlords' Eviction Practices." *Social Forces*, June, soab063. doi:10.1093/sf/soab063.

Gonzalez, Libertad, and Francesc Ortega. 2013. "Immigration and Housing Booms: Evidence from Spain*." *Journal of Regional Science* 53 (1): 37–59. doi:10.1111/jors.12010.

Goodman, Allen C., and Thomas G. Thibodeau. 2003. "Housing Market Segmentation and Hedonic Prediction Accuracy." *Journal of Housing Economics* 12 (3): 181–201. doi:10.1016/S1051-1377(03)00031-7.

Goodspeed, Robert. 2023. "Digitalization and the Incomplete Disruption of the US Housing Industry." In *Bringing Digitalization Home Symposium*. Cambridge, MA: Harvard Joint Center for Housing Studies. www.jchs.harvard.edu/research-areas/working-papers/digitalization-and-incomplete-disruption-us-housing-industry.

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst. 2013. "Endogenous Gentrification and Housing Price Dynamics." *Journal of Public Economics* 100 (April): 45–60. doi:10.1016/j.jpubeco.2013.02.001.

Gupta, Arpit, Vrinda Mittal, Jonas Peeters, and Stijn Van Nieuwerburgh. 2022. "Flattening the Curve: Pandemic-Induced Revaluation of Urban Real Estate." *Journal of Financial Economics* 146 (2): 594–636. doi:10.1016/j.jfineco.2021.10.008.

Gurun, Umit G, Jiabin Wu, Steven Chong Xiao, and Serena Wenjing Xiao. 2022. "Do Wall Street Landlords Undermine Renters' Welfare?" Edited by Tarun Ramadorai. *The Review of Financial Studies*, March, hhac017. doi:10.1093/rfs/hhac017.

Haider, Elinor. 2021. "Rental Code Enforcement in Philadelphia." The Pew Charitable Trusts. pew.org/3qMPpN2.

Hangen, Forrest, and Daniel T. O'Brien. 2022. "Linking Landlords to Uncover Ownership Obscurity." Preprint. December 3. doi:10.31235/osf.io/anvke.

Harvey, David. 2009. *Social Justice and the City*. Rev. ed. Geographies of Justice and Social Transformation 1. Athens: University of Georgia Press.

Herbert, Christopher, Daniel McCue, and Rocio Sanchez-Moyano. 2014. "Is Homeownership Still an Effective Means of Building Wealth for Low-Income and Minority Households?" In *Homeownership Built to Last*, edited by Eric S. Belsky, Christopher E. Herbert, and Jennifer H. Molinsky, 50–98. Balancing Access, Affordability, and Risk After the Housing Crisis. Brookings Institution Press. www.jstor.org/stable/10.7864/j.ctt6wpj09.6.

Hochstenbach, Cody. 2022. "Landlord Elites on the Dutch Housing Market: Private Landlordism, Class, and Social Inequality." *Economic Geography* 98 (4): 327–54. doi:10.1080/00130095.2022.2030703.

———. 2023. "Networked Geographies of Private Landlordism: Mapping Flows of Capital Accumulation and Rent Extraction." *Housing Studies* 0 (0). routledge: 1–26. doi:10.1080/02673037.2023.2174255.

Hwang, Sungsoon. 2015. "Residential Segregation, Housing Submarkets, and Spatial Analysis: St. Louis and Cincinnati as a Case Study." *Housing Policy Debate* 25 (1). routledge: 91–115. doi:10.1080/10511482.2014.934703.

Ihlanfeldt, Keith. 2021. "Property Tax Homestead Exemptions: An Analysis of the Variance in Take-Up Rates Across Neighborhoods." *National Tax Journal* 74 (2). The University of Chicago Press: 405–30. doi:10.1086/714168.

Immergluck, Dan, Jeff Ernsthausen, Stephanie Earl, and Allison Powell. 2020. "Evictions, Large Owners, and Serial Filings: Findings from Atlanta." *Housing Studies* 35 (5). routledge: 903–24. doi:10.1080/02673037.2019.1639635.

Immergluck, Dan, and Jonathan Law. 2014. "Speculating in Crisis: The Intrametropolitan Geography of Investing in Foreclosed Homes in Atlanta." *Urban Geography* 35 (1): 1–24. doi:10.1080/02723638.2013.858510.

Jones, Colin. 2002. "The Definition of Housing Market Areas and Strategic Planning." *Urban Studies* 39 (3): 549–64. doi:10.1080/00420980220112829.

Jones, Colin, and Mike Coombes. 2013. "An Assessment of Tenure-Specific Housing Market Areas for Housing Planning." *Housing Studies* 28 (7): 993–1011. doi:10.1080/02673037.2013.783201.

Jonge, Edwin de, and Mark van der Loo. 2018. *Statistical Data Cleaning with Applications in R*. Hoboken, NJ: John Wiley & Sons.

Krysan, Maria. 2008. "Does Race Matter in the Search for Housing? An Exploratory Study of Search Strategies, Experiences, and Locations." *Social Science Research* 37 (2): 581–603. doi:10.1016/j.ssresearch.2007.06.001.

Langley, Adam H. 2015. "How Do States Spell Relief?" *Land Lines, Lincoln Institute of Land Policy*, April. www.lincolninst.edu/publications/articles/how-do-states-spell-relief.

Lee, Hyojung. 2017. "Who Owns Rental Properties, and Is It Changing?" Housing Perspectives, Joint Center for Housing Studies. August 18. www.jchs.harvard.edu/blog/who-owns-rental-properties-and-is-it-changing.

Lennartz, Christian. 2014. "Market Structures of Rental Housing: Conceptualising Perfect Competition in Mixed Local Rental Markets." *International Journal of Housing Policy* 14 (1): 56–78. doi:10.1080/14616718.2013.877686.

Lens, Michael C., and Vincent Reina. 2016. "Preserving Neighborhood Opportunity: Where Federal Housing Subsidies Expire." *Housing Policy Debate* 26 (4-5). routledge: 714–32. doi:10.1080/10511482.2016.1195422.

Linger, Jacob, Hal J. Singer, and Ted Tatos. 2022. "Does a Lack of Competition Exacerbate Inflation? A Case Study of Florida Rental Markets." SSRN Scholarly Paper. Rochester, NY. December 6. doi:10.2139/ssrn.4295000.

Lynch, Edward. 2022. "Understanding the Rental Landscape." Detroit Future City. detroitfuturecity.com/wp-content/uploads/2022/08/FINAL-Understanding-the-Rental-Landscape.pdf.

Maclennan, Duncan, and Anthony O'Sullivan. 2012. "Housing Markets, Signals and Search." *Journal of Property Research* 29 (4). routledge: 324–40. doi:10.1080/09599916.2012.717102.

Mallach, Alan. 2010. "Meeting The Challenge Of Distressed Property Investors In America's Neighborhoods." lisc. www.lisc.org/media/filer_public/4f/46/4f462994-10b0-4eb0-8855-ddafc6e82d37/08142018_resources_distressed_property_investors.pdf.

Manduca, Robert. 2020. "The Spatial Structure of US Metropolitan Employment: New Insights from Administrative Data." *Environment and Planning B: Urban Analytics and City Science*, June, 2399808320934821. doi:10.1177/2399808320934821.

Manson, Steven, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles. 2021. "National Historical Geographic Information System: Version 16.0." Minneapolis, MN: IPUMS. doi:10.18128/D050.V16.0.

Medwell, Jamie. 2021. "Abolish Landlords." *Jacobin Magazine*, July. jacobin.com/2021/07/abolish-landlords-cancel-rent-eviction-homelessness.

Mills, James, Raven S. Molloy, and Rebecca E. Zarutskie. 2015. "Large-Scale Buy-to-Rent Investors in the Single-Family Housing Market: The Emergence of a New Asset Class?" *Finance and Economics Discussion Series* 2015 (84): 1–58. doi:10.17016/FEDS.2015.084.

Mills, James, Raven Molloy, and Rebecca Zarutskie. 2019. "Large-Scale Buy-to-Rent Investors in the Single-Family Housing Market: The Emergence of a New Asset Class." *Real Estate Economics* 47 (2): 399–430. doi:10.1111/1540-6229.12189.

Molfino, Emily, Gizem Korkmaz, Sallie A. Keller, Aaron Schroeder, Stephanie Shipp, and Daniel H. Weinberg. 2017. "Can Administrative Housing Data Replace Survey Data?" *Cityscape* 19 (1). US Department of Housing and Urban Development: 265–92. www.jstor.org/stable/26328312.

Narayanan, Ajjit, and Graham MacDonald. 2019. "Toward an Open Data Bias Assessment Tool." The Urban Institute. www.urban.org/sites/default/files/publication/99844/toward_an_open_data_bias_assessment_tool_1.pdf.

National Association of Realtors. 2023. "Snapshot of Race and Home Buying in America." www.nar.realtor/research-and-statistics/research-reports/a-snapshot-of-race-and-home-buying-in-america.

National Consumer Law Center. 2022. "Promoting Resident Ownership of Communities." www.nclc.org/resources/promoting-resident-ownership-of-communities-2/.

National Multifamily Housing Council. 2022. "NMHC 50 Archive." www.nmhc.org/research-insight/the-nmhc-50/nmhc-50-archive/.

Neal, Michael, and Laurie Goodman. 2022. "The Housing Market Needs More Condos. Why Are So Few Being Built?" Urban Wire, Urban Institute. January 31. www.urban.org/urban-wire/housing-market-needs-more-condos-why-are-so-few-being-built.

Neal, Zachary P. 2022. "Backbone: An R Package to Extract Network Backbones." *PLOS ONE* 17 (5). Public Library of Science: e0269137. doi:10.1371/journal.pone.0269137.

Neal, Zachary P., Rachel Domagalski, and Bruce Sagan. 2022. "Analysis of Spatial Networks From Bipartite Projections Using the R Backbone Package." *Geographical Analysis* 54 (3): 623–47. doi:10.1111/gean.12275.

Newman, Kathe. 2009. "Post-Industrial Widgets: Capital Flows and the Production of the Urban." *International Journal of Urban and Regional Research* 33 (2): 314–31. doi:10.1111/j.1468-2427.2009.00863.x.

Offenhuber, Dietmar. 2017. "Sticky Data: Context and Friction in the Use of Urban Data Proxies." In *Data and the City*. routledge.

Owens, Ann. 2020. "Unequal Opportunity: School and Neighborhood Segregation in the USA." *Race and Social Problems* 12 (1): 29–41. doi:10.1007/s12552-019-09274-z.

Özogul, Sara, and Tuna Tasan-Kok. 2020. "One and the Same? A Systematic Literature Review of Residential Property Investor Types." *Journal of Planning Literature* 35 (4): 475–94. doi:10.1177/0885412220944919.

Pananjady, Kasturi, Max Marin, and Jasen Lo. 2022. "Homeowners in City Are Losing Millions in Available Tax Relief." *Philadelphia Inquirer*, November 30. www.proquest.com/newspapers/homeowners-city-are-losing-millions-available-tax/docview/2741269638/se-2?accountid=12492.

Parker, Will, and Nicole Friedman. 2022. "Everyone's a Landlord—Small-Time Investors Snap Up Out-of-State Properties." *Wall Street Journal: US*, August 28. www.wsj.com/articles/everyones-a-landlordsmall-time-investors-snap-up-out-of-state-properties-11661705278.

Rae, Alasdair, and Ebru Sener. 2016. "How Website Users Segment a City: The Geography of Housing Search in London." *Cities* 52 (March): 140–47. doi:10.1016/j.cities.2015.12.002.

Raymond, Elora Lee, Richard Duckworth, Benjamin Miller, Michael Lucas, and Shiraj Pokharel. 2018. "From Foreclosure to Eviction: Housing Insecurity in Corporate-Owned Single-

Family Rentals." *Cityscape* 20 (3). US Department of Housing and Urban Development: 159–88. www.jstor.org/stable/26524878.

Reid, Carolina. 2013. "To Buy or Not to Buy? Understanding Tenure Preferences and the Decision-Making Processes of Lower-Income Households." In, 28. Harvard Business School, Boston: Joint Center for Housing Studies.

Richardson, Todd. 2018. "Landlords." *HUD EDGE*, June. www.huduser.gov/portal/pdredge/pdr-edge-frm-asst-sec-061118.html.

Robinson, David, and Justin Steil. 2021. "Eviction Dynamics in Market-Rate Multifamily Rental Housing." *Housing Policy Debate* 31 (3-5): 647–69. doi:10.1080/10511482.2020.1839936.

Robsky Huntley, Eric, La-Brina Almeida, Shep Heaton, and Andrea Moon Park. 2022. "Housing Justice Beyond the Emergency." Homes for All Massachusetts. www.homesforallmass.org/covid-evictions-report/hfamass-covid-evictions-report-short-03-22-screen-rfs.pdf.

Rose, Geoff, and Richard Harris. 2021. "The Three Tenures: A Case of Property Maintenance." *Urban Studies*, July, 004209802110292. doi:10.1177/00420980211029203.

Rosen, Eva. 2014. "Rigging the Rules of the Game: How Landlords Geographically Sort Low–Income Renters." *City & Community* 13 (4). SAGE Publications: 310–40. doi:10.1111/cico.12087.

Rosen, Eva, and Philip M. E. Garboden. 2022. "Landlord Paternalism: Housing the Poor with a Velvet Glove." *Social Problems* 69 (2): 470–91. doi:10.1093/socpro/spaa037.

Rosen, Eva, Philip M. E. Garboden, and Jennifer E. Cossyleon. 2021. "Racial Discrimination in Housing: How Landlords Use Algorithms and Home Visits to Screen Tenants." *American Sociological Review* 86 (5): 787–822. doi:10.1177/00031224211029618.

Royuela, Vicente, and Miguel A. Vargas. 2009. "Defining Housing Market Areas Using Commuting and Migration Algorithms: Catalonia (Spain) as a Case Study." *Urban Studies* 46 (11): 2381–98. doi:10.1177/0042098009342600.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Danika Brockman, Grace Cooper, Stephanie Richards, and Megan Schouwiler. 2023. "IPUMS USA: Version 13.0." Minneapolis, MN: IPUMS. doi:10.18128/D010.V13.0.

Saiz, Albert. 2007. "Immigration and Housing Rents in American Cities." *Journal of Urban Economics* 61 (2): 345–71. doi:10.1016/j.jue.2006.07.004.

Saiz, Albert, and Susan Wachter. 2011. "Immigration and the Neighborhood." *American Economic Journal: Economic Policy* 3 (2): 169–88. doi:10.1257/pol.3.2.169.

Sakong, Jung. 2021. "Effect of Ownership Composition on Property Prices and Rents: Evidence from Chinese Investment Boom in US Housing Markets." Working Paper. doi:10.21033/wp-2021-12.

Samuel, Andrew, Jeremy Schwartz, and Kerry Tan. 2021. "Licensing and The Informal Sector in Rental Housing Markets: Theory And Evidence." *Contemporary Economic Policy* 39 (2): 325–47. doi:10.1111/coep.12501.

Sassen, Saskia. 1990. *The Mobility of Labor and Capital: A Study in International Investment and Labor Flow*. Transferred to digital printing. Cambridge: Cambridge Univ. Press.

———. 2005. "The Global City: Introducing a Concept Global Cities." *Brown Journal of World Affairs* 11 (2): 27–44. heinonline.org/HOL/P?h=hein.journals/brownjwa11&i=279.

Saxon, James. 2020. "The Local Structures of Human Mobility in Chicago." *Environment and Planning B: Urban Analytics and City Science*, August, 239980832094953. doi:10.1177/2399808320949539.

Serrano, M. Ángeles, Marián Boguñá, and Alessandro Vespignani. 2009. "Extracting the Multiscale Backbone of Complex Weighted Networks." *Proceedings of the National Academy of Sciences* 106 (16). Proceedings of the National Academy of Sciences: 6483–88. doi:10.1073/pnas.0808904106.

Shelton, Taylor. 2018. "Rethinking the RECAP: Mapping the Relational Geographies of Concentrated Poverty and Affluence in Lexington, Kentucky." *Urban Geography* 39 (7): 1070–91. doi:10.1080/02723638.2018.1433927.

———. 2021. "Gameday Homes: Mapping Emerging Geographies of Housing Speculation and Absentee Ownership in the American South." *Cities* 115 (August): 103230. doi:10.1016/j.cities.2021.103230.

Shiffer–Sebba, Doron. 2020. "Understanding the Divergent Logics of Landlords: Circumstantial Versus Deliberate Pathways." *City & Community* 19 (4): 1011–37. doi:10.1111/cico.12490.

Small, Mario L. 2010. *Unanticipated Gains Origins of Network Inequality in Everyday Life*. Oxford; New York: Oxford University Press.

Smith, Neil. 1979. "Toward a Theory of Gentrification A Back to the City Movement by Capital, Not People." *Journal of the American Planning Association* 45 (4): 538–48. doi:10.1080/01944367908977002.

Spielman, Seth E., and David C. Folch. 2015. "Reducing Uncertainty in the American Community Survey Through Data-Driven Regionalization." *PLOS ONE* 10 (2). Public Library of Science: e0115626. doi:10.1371/journal.pone.0115626.

Spielman, Seth E., David Folch, and Nicholas Nagle. 2014. "Patterns and Causes of Uncertainty in the American Community Survey." *Applied Geography* 46 (January): 147–57. doi:10.1016/j.apgeog.2013.11.002.

"State-by-State Property Tax at a Glance." 2022. Lincoln Institute of Land Policy and George Washington Institute of Public Policy. www.lincolninst.edu/research-data/data-toolkits/significant-features-property-tax/state-state-property-tax-glance.

Statistics Canada. 2022. "A Portrait of Landlords and Their Rental Income: A Tax Data Perspective." www150.statcan.gc.ca/n1/daily-quotidien/221102/dq221102b-eng.htm.

St-Hilaire, Cloé, Mikael Brunila, and David Wachsmuth. 2023. "High Rises and Housing Stress." *Journal of the American Planning Association* 0 (0). routledge: 1–15. doi:10.1080/01944363.2022.2126382.

Stiman, Meaghan. 2019. "Speculators and Specters: Diverse Forms of Second Homeowner Engagement in Boston, Massachusetts." *Journal of Urban Affairs* 41 (5). routledge: 700–720. doi:10.1080/07352166.2018.1550348.

Tapp, Renee, and Richard Peiser. 2022. "An Antitrust Framework for Housing." *Environment and Planning A: Economy and Space*, November, 0308518X2211356. doi:10.1177/0308518X221135612.

Travis, Adam. 2019. "The Organization of Neglect: Limited Liability Companies and Housing Disinvestment." *American Sociological Review* 84 (1): 142–70. doi:10.1177/0003122418821339.

Trounstine, Jessica. 2018. *Segregation by Design: Local Politics and Inequality in American Cities*. 1st ed. Cambridge University Press. doi:10.1017/9781108555722.

US Census Bureau. 1996. "Who Owns the Nation's Rental Properties?" SB/96-1. www2.census.gov/library/publications/1996/demographics/sb96-01.pdf.

———. 2021. "Rental Housing Finance Survey." www.census.gov/programs-surveys/rhfs.html.

US Department of Education National Center for Education Statistics. 2016. "School Attendance Boundary Survey." nces.ed.gov/programs/edge/sabs.

Vigdor, Jacob L. 2004. "Other People's Taxes: Nonresident Voters and Statewide Limitation of Local Government." *The Journal of Law and Economics* 47 (2): 453–76. doi:10.1086/425061.

Vogell, Heather. 2022. "When Private Equity Becomes Your Landlord." *Propublica*, February 7. www.propublica.org/article/when-private-equity-becomes-your-landlord?token=203ntxGVDUHDrjJiIJU-odgk4aPCUinL.

Vogell, Heather, Haru Coryne, and Ryan Little. 2022. "Rent Going Up? One Company's Algorithm Could Be Why." *Propublica*, October 15. www.propublica.org/article/yieldstar-rent-increase-realpage-rent.

Watson, C. Luke, and Oren Ziv. 2022. "Is the Rent Too High? Land Ownership and Monopoly Power." Working Paper. clukewatson.github.io/research/nyc-monopoly/monopoly_watsonziv.pdf.

Wilson, Nicole E., Michael Hankinson, Asya Magazinnik, and Melissa Sands. 2023. "Inaccuracies in Low Income Housing Geocodes: When and Why They Matter." *Urban Affairs Review*, March. SAGE Publications Inc, 10780874231165767. doi:10.1177/10780874231165767.

Yabe, Takahiro, Bernardo García Bulle Bueno, Xiaowen Dong, Alex Pentland, and Esteban Moro. 2023. "Behavioral Changes During the COVID-19 Pandemic Decreased Income Diversity of Urban Encounters." *Nature Communications* 14 (1, 1). Nature Publishing Group: 2310. doi:10.1038/s41467-023-37913-y.