

NBER WORKING PAPER SERIES

TOWARDS A METHODOLOGY FOR
MEASURING RENTAL PROPERTY OWNERSHIP IN THE UNITED STATES

Stephanie Kestelman
Rebecca Diamond
John Eric Humphries
Kate Pennington
Winnie van Dijk
John L. Voorheis

Working Paper 35258
<http://www.nber.org/papers/w35258>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau, Arnold Ventures, or the National Bureau of Economic Research. We are grateful to Sonya Porter and Oren Ziv for helpful comments on an earlier draft. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7511151; Disclosure Authorization Number CBDRB-FY25-CES014-021, CBDRB-FY26-CES014-002, CBDRB-FY26-CES014-006, CBDRB-FY26-CES023-005, and CESFRC-26-22). Humphries and Van Dijk gratefully acknowledge financial support from the Yale Tobin Center for Economic Policy and the Cowles Foundation for Research in Economics at Yale University. Kestelman and van Dijk acknowledge support from the Harvard Joint Center for Housing Studies. Noah Friedlander, Sean Kersey, and Lucas Marron provided excellent research assistance.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Stephanie Kestelman, Rebecca Diamond, John Eric Humphries, Kate Pennington, Winnie van Dijk, and John L. Voorheis. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Towards a Methodology for Measuring Rental Property Ownership in the United States
Stephanie Kestelman, Rebecca Diamond, John Eric Humphries, Kate Pennington, Winnie
van Dijk, and John L. Voorheis
NBER Working Paper No. 35258
May 2026
JEL No. H0, L0, R0

ABSTRACT

Roughly one-third of U.S. households rent their homes, yet measuring who owns rental property is difficult: ownership is frequently obscured by LLCs, partnerships, and other intermediary entities that separate legal from economic control. We develop a method that traces ownership through administrative records — combining deeds and property assessments with the Census Bureau's Business Register, IRS Schedule K-1 filings, and SEC filings on REITs — to identify ultimate owners and construct property portfolios across the full landlord size distribution. Applying the method to 11 large CBSAs, we find that individual landlords own a large majority of rental units, though their share varies meaningfully across markets. We also show that the widely used mailing-address aggregation approach both under- and over-states portfolio size in systematic ways. The method is designed to scale to national coverage and to support measurement of landlord identity, portfolio composition, and ownership concentration in U.S. rental markets. We also discuss the method's current limitations and outline directions for refinement and validation.

Stephanie Kestelman
Arnold Ventures
stephaniekestelman@gmail.com

Kate Pennington
U.S. Census Bureau
kate.pennington@berkeley.edu

Rebecca Diamond
Harvard University
Department of Economics
and NBER
rdiamond@fas.harvard.edu

Winnie van Dijk
Yale University
Department of Economics
and NBER
winnie.vandijk@yale.edu

John Eric Humphries
Yale University
Department of Economics
and NBER
johneric.humphries@yale.edu

John L. Voorheis
U.S. Census Bureau
Department of Economics
john.l.voorheis@census.gov

Towards a Methodology for Measuring Rental Property Ownership in the United States

Stephanie Kestelman, Rebecca Diamond, John Eric Humphries,
Kate Pennington, Winnie van Dijk, John Voorheis*

May 2026

Abstract

Roughly one-third of U.S. households rent their homes, yet measuring who owns rental property is difficult: ownership is frequently obscured by LLCs, partnerships, and other intermediary entities that separate legal from economic control. We develop a method that traces ownership through administrative records—combining deeds and property assessments with the Census Bureau’s Business Register, IRS Schedule K-1 filings, and SEC filings on REITs—to identify ultimate owners and construct property portfolios across the full landlord size distribution. Applying the method to 11 large CBSAs, we find that individual landlords own a large majority of rental units, though their share varies meaningfully across markets. We also show that the widely used mailing-address aggregation approach both under- and over-states portfolio size in systematic ways. The method is designed to scale to national coverage and to support measurement of landlord identity, portfolio composition, and ownership concentration in U.S. rental markets. We also discuss the method’s current limitations and outline directions for refinement and validation.

1 Introduction

Approximately one-third of U.S. households rent their homes and annual rent expenditures exceed \$700 billion. Rental housing is thus a central component of the U.S. economy, yet it remains difficult to measure who owns it. Property ownership is frequently obscured by limited liability companies (LLCs), partnerships, trusts, and other intermediary entities that separate legal ownership from economic control. Administrative property records therefore

*Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau or Arnold Ventures. We are grateful to Sonya Porter and Oren Ziv for helpful comments on an earlier draft. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7511151; Disclosure Authorization Number CBDRB-FY25-CES014-021, CBDRB-FY26-CES014-002, CBDRB-FY26-CES014-006, CBDRB-FY26-CES023-005, and CESFRC-26-22). Humphries and Van Dijk gratefully acknowledge financial support from the Yale Tobin Center for Economic Policy and the Cowles Foundation for Research in Economics at Yale University. Kestelman and van Dijk acknowledge support from the Harvard Joint Center for Housing Studies. Noah Friedlander, Sean Kersey, and Lucas Marron provided excellent research assistance.

typically identify only the immediate legal entity that nominally owns a property rather than the individuals or firms making investment and pricing decisions. A single legal entity may correspond to multiple individual owners or businesses, and an individual or business could exert control over multiple legal entities. This gap between legal and economic ownership therefore complicates measurement of ownership concentration, limits the ability to distinguish between small and large landlords, and makes it difficult to study how different types of landlords respond to economic shocks and policy changes.

Measuring who owns rental housing matters for a range of questions in housing economics and policy. Researchers have sought to understand how landlord characteristics affect rent setting and housing prices—questions at the heart of understanding market power, competition, and the spatial structure of urban housing markets (e.g., [Gurun et al., 2023](#); [Watson and Ziv, 2025b](#); [Barbieri and Dobbels, 2025](#); [Gorback et al., 2025](#); [Wang and Zhai, 2025](#)). Related work examines how changes in institutional ownership affect neighborhood composition and homeownership rates, as well as how ownership concentration affects tenant- and neighborhood-level outcomes such as eviction rates, demographic composition, and housing quality (e.g., [Travis, 2019](#); [An, 2024](#); [Coven, 2025](#); [Fesko, 2025](#); [Chang, 2025](#); [Harwood et al., 2025](#)). These same issues motivate policy design, including whether different types of landlords respond differentially to rent regulation, how tax policy might account for portfolio size, and how exposed landlords are to financial and macroeconomic shocks (e.g., [Levy, 2022](#); [Cocco et al., 2024](#); [Samuels, 2025](#)).

However, existing approaches to measuring rental property ownership have important limitations. Some jurisdictions have high-quality local data that permit high-quality measures of ownership (see, e.g., [Harwood et al., 2025](#); [Watson and Ziv, 2025b,a](#)), but these data are available only for a handful of places and cannot capture portfolios that span jurisdictions not covered by these local datasets. Surveys, such as the Residential Housing Finance Survey (RHFS), can be nationally representative, but they have small sample sizes, are subject to response error (especially since the respondent may be a property manager rather than the beneficial owner), and do not describe an owner’s full portfolio. Lastly, methods that aggregate portfolios by shared mailing addresses in tax assessments or deeds records can misidentify rental properties and both under- and over-estimate portfolio size. We review the benefits and drawbacks of these three approaches in more detail in Section 2.2.

This chapter introduces a new approach that has the potential to address the limitations outlined for these existing methods. Unlike the local data approach, our method can be applied nationally, capturing nearly all jurisdictions and constructing property portfolios across the entire United States. Our approach may also yield more accurate measures of landlord scale and ownership concentration, compared to methods based on survey data or mailing addresses, because it traces ownership through administrative records, following the owner listed on a property record through a chain of business owners to a final set of shareholders, while avoiding aggregation based solely on shared mailing addresses or intermediary agents.

We combine deeds and property assessment records with the Census Bureau’s Business Register, the Internal Revenue Service Schedule K-1 partnership data, and information on Real Estate Investment Trusts (REITs) and their publicly traded subsidiaries from the Securities and Exchange Commission Form 10-K to trace ownership through corporate structures and identify final owners. Starting from the legal owner listed in property records, we use these data sources to follow ownership chains through layers of LLCs, partnerships,

trusts, and other intermediary entities. Properties held under different legal names or entities are connected when they share a common underlying owner. Our approach assigns properties to “Final Owner Networks,” which correspond more closely to the economic decision-maker than the legal entity listed on a deed. Identifying ownership at this level allows landlord portfolios to be measured consistently across ownership forms and enables classifying owners into economically meaningful categories, including unincorporated individuals, incorporated individuals, individual trusts, business entities, and public owners.

At the same time, this approach faces several challenges, which this chapter discusses in detail. We are still working to resolve these challenges for the full set of rental properties in the United States. As we document in Section 3, there are opportunities to improve the identification of individual landlords and landlords who hold properties through special-purpose LLCs that do not file tax forms. Another limitation is that we currently identify fewer rental units than the total count implied by the American Community Survey (ACS) in every Core Based Statistical Area (CBSA) in our sample, indicating an undercount in our rental property identification.

We apply our methodology to 11 large CBSAs spanning diverse housing markets and regulatory environments: Atlanta, Charlotte, Chicago, Denver, Hartford, Houston, Jacksonville, Los Angeles, Phoenix, Seattle, and Washington, DC. In Section 4, we present descriptive findings for the subset of properties in these CBSAs that appear in the ACS. While the undercount relative to the ACS described above raises a concern about the accuracy of our identification of rental units, we are confident that rental properties are correctly identified in this subset (up to reporting error in the ACS survey data). We leave descriptives based on a larger sample of rental properties—including more CBSAs as well as properties in our CBSAs that do not appear in the ACS—to a future data build that has resolved our coverage challenges. This sample provides proof of concept for the types of analysis our methodology is designed to enable going forward, but we caution that they do not necessarily generalize to the full rental stock.

Using this subset, we document several statistics about the supply side of the rental housing markets captured in our sample. First, individual landlords own the majority of rental units: approximately 61 percent of rental units are owned by unincorporated individuals and an additional 7.6 percent by incorporated individuals. Second, rental ownership is widely dispersed among the 6.1 million properties in our sample: we identify more than 558,000 unique individual and business landlords corresponding to these properties. Third, rental income is an important component of individual landlords’ portfolios. Nearly one third of individual landlords earned no wage or salary income in 2019, with nearly half of unincorporated individuals earning less than \$100,000. These patterns indicate that a substantial portion of rental housing supply reflects the investment decisions of a large number of individuals and households who earn limited or no labor income. Lastly, we document that the share of unincorporated individual owners varies meaningfully across markets in our sample—ranging from roughly 50 percent in Houston to over 70 percent in Chicago. To the extent that policies affecting rental markets have different effects depending on the organization of ownership, this suggests there may be important heterogeneity in responses to such policies by location.

We also compare our ownership measures to those generated by the mailing address approach, the most commonly used existing approach, and show that existing methods can substantially misclassify both ownership type and portfolio size.

1.1 Related literature

Research on rental property ownership has expanded in recent years. Studies differ in how they recover landlord portfolios from administrative data, and these choices shape both which segment of the rental market can be observed and what analysis is possible. We organize the literature along two dimensions: coverage (both geographic and across types of landlords) and methodology. We briefly note the benefits and limitations of each approach; Section 2 provides a more detailed evaluation.

Coverage. Studies vary in both geographic scope and the types of landlords they capture. Most construct landlord portfolios in individual cities, states, or regions (Travis, 2019; Austin, 2022; Linger et al., 2022; An, 2024; An et al., 2024; Ramoutar, 2024; Torres, 2024; Barbieri and Dobbels, 2025; Fesko, 2025; Hangen and O’Brien, 2025; Watson and Ziv, 2025a,b; Harwood et al., 2025; Wang and Zhai, 2025), with some covering multiple markets (Ganduri et al., 2023; Gurun et al., 2023; Chang, 2025) and others providing national or near-national coverage (Garriga et al., 2022; Lambie-Hanson et al., 2022; Lee and Wylie, 2024; Coven, 2025; Gorback et al., 2025; Handbury et al., 2025; Samuels, 2025). Even among studies with national coverage, to our knowledge, none construct ownership classifications across the full landlord size and type distribution. Much of the literature focuses on large institutional landlords in the single-family rental market (Mills et al., 2019; Ganduri et al., 2023; Gurun et al., 2023; An, 2024; Lee and Wylie, 2024; Calder-Wang and Kim, 2024; Barbieri and Dobbels, 2025; Chang, 2025; Coven, 2025; Fesko, 2025; Gorback et al., 2025) or instead on large multifamily properties (Ramoutar, 2024; Calder-Wang and Kim, 2024; Damen et al., 2025; Handbury et al., 2025). Data from the American Community Survey shows that roughly two-thirds of renter households live in structures with more than one unit, a segment omitted by studies focused on single-family rentals. Similarly, studies focused on large multifamily properties do not fully characterize ownership for the roughly two-thirds of renter households living in structures with fewer than 10 units.

Methodology. The dominant approach constructs landlord portfolios directly from deed or tax assessor records by aggregating properties that share a common owner name and/or mailing address, typically after string standardization. This approach scales to large geographic areas and has been widely adopted (Lambie-Hanson et al., 2022; Ganduri et al., 2023; Gurun et al., 2023; An, 2024; An et al., 2024; Lee and Wylie, 2024; Ramoutar, 2024; Barbieri and Dobbels, 2025; Coven, 2025; Chang, 2025; Fesko, 2025; Samuels, 2025; Harwood et al., 2025; Watson and Ziv, 2025a,b; Wang and Zhai, 2025). However, name–address aggregation relies on consistent recording of ownership across properties, and professional registered agents or intermediaries may appear in ownership records rather than true owners. Measurement error may therefore arise in both directions: properties may be listed as owned by separate legal entities when they have the same economic owner, or different economic owners may list the same intermediary agent for their properties. Moreover, errors may vary with institutional arrangements and local recording practices. Section 2.2 discusses these sources of error in more detail and Section 4 illustrates the bidirectional measurement error in the mailing address approach using our dataset.

A smaller literature supplements parcel data with corporate registration records to link legally distinct entities that may share common control (Torres, 2024; Hangen and O’Brien,

2025). These approaches use shared officers, registered agents, or business addresses to infer ownership connections and can improve identification of fragmented institutional ownership in specific settings. In practice, however, corporate registries are maintained at the state level and differ in structure and reporting requirements, and shared professional agents can generate spurious links. Existing implementations remain limited to single-city applications.

A third strand of work identifies institutional investors using external information sources without attempting to recover landlord portfolios for the full market. Mills et al. (2019), Garriga et al. (2022), Calder-Wang and Kim (2024), Damen et al. (2025), Gorbach et al. (2025), and Handbury et al. (2025) identify large institutional investors using SEC filings, industry reports, or curated firm lists and match these firms to transaction data. This approach provides precise measurement for large institutional landlords, but these studies do not recover the broader distribution of landlord ownership.

Finally, outside the United States, a smaller literature relies on administrative data systems in which ownership can be observed directly through linked registries or tax records. Levy (2022) uses French administrative data linking property ownership and tax records, while Cocco et al. (2024) identify landlords through borrower information in U.K. buy-to-let mortgage data. Damen et al. (2025) use administrative datasets for Belgium and the Netherlands, and Moktan (2025) studies rental market structure using U.K. administrative ownership information. Because ownership linkage is embedded in the underlying data, these studies avoid many of the measurement challenges present in U.S. property records.

Across approaches, existing work using U.S. data faces a tradeoff between geographic scope, coverage across the landlord size distribution, and the ability to recover ultimate ownership. We aim to fill this gap by providing a framework for constructing landlord portfolios at the national level for the United States, covering both individual and institutional landlords across the building size distribution, and using entity resolution methods that do not rely solely on mailing address matching. Although the data used to illustrate our methods in this chapter does not yet offer national coverage, our approach can be scaled to achieve such coverage over time.

2 Comparing methods for measuring rental property ownership

2.1 What are reasonable goals for comprehensive ownership measurement?

The specific measurement requirements for studying rental housing ownership depend on the research question, but for many of the questions discussed in the introduction—about market power, the effects of ownership concentration, and the design of housing regulation—there are arguably four natural goals.

First, it is useful to have a source of nationwide data. Rental markets vary substantially across the United States in their regulatory environments, building stock, and possibly their ownership structures. A measurement framework limited to particular cities or states cannot characterize these differences or leverage them to analyze the impact of policy changes. Additionally, it cannot construct complete portfolios for landlords who own property in multiple jurisdictions. National coverage also helps ensure that findings are not driven by the

idiosyncrasies of a particular regulatory environment or data source.

Second, it is important to accurately identify residential rental properties. There is no centralized source of information on whether a property is renter- or owner-occupied at a given point in time. Property assessment records sometimes include an owner-occupancy flag, but these flags can be unreliable or incomplete. The distinction is not always straightforward: a single building may contain both owner-occupied and renter-occupied units, and tenure status can change over time. Misclassification of tenure can bias estimates of rental housing supply and distort the composition of measured landlord portfolios. For example, if owner-occupied properties are incorrectly classified as rentals, measured individual landlord ownership rates will be inflated, and if rental properties are missed entirely, the data will underrepresent certain segments of the market.

Third, for many questions it is valuable to identify the final owner of each property—the individual or entity that controls investment and pricing decisions. When ownership is held through intermediary entities such as LLCs or partnerships, the legal owner recorded in property records may be several steps removed from the economic decision-maker. Without tracing through these structures, it is difficult to distinguish incorporated individual landlords from large institutional investors, to classify owner types in a meaningful way, or to connect an owner’s characteristics (such as income, demographics, or other business activities) to the properties they control.

Fourth, many questions require constructing complete property portfolios for each final owner. Questions about market concentration, economies of scale, and the effects of portfolio size on tenant outcomes all depend on knowing not just who owns a given property, but also how many and which other properties that owner controls. Portfolio construction matters because an owner who holds a single property and one who holds hundreds may appear identically in property records if ownership is fragmented across legal entities. Because landlords may hold properties through multiple legal entities and across multiple jurisdictions, portfolio construction requires linking ownership across entity structures and geographies. Portfolio construction also involves conceptual choices: when two entities share an investor, the researcher must decide whether they constitute “the same landlord” and where to draw the boundary between individual and corporate ownership. The mailing address approach largely sidesteps these definitional questions—portfolios are simply the set of properties sharing an address—but at the cost of accuracy (see Section 2.2). Survey-based approaches like the RHFS cannot construct portfolios at all. Our approach requires explicit choices about shareholder count thresholds and network definitions, which we discuss in Section 3.

Meeting all four of these goals simultaneously has proven difficult. Existing approaches in the literature each have limitations along at least one dimension. Table 1 summarizes how the leading approaches perform along each of these four dimensions.

2.2 To what extent can existing approaches meet these goals?

Survey-based methods, such as the Residential Housing Finance Survey (RHFS), can be nationally representative, but the sample size is usually too small to support precise estimates for local markets, detailed cross-tabulations (e.g., by owner type within a metro), or analysis of the tails of the portfolio size distribution. For example, the RHFS samples roughly 100,000 properties nationally. The identification of final owners for these properties is noisy because the respondent may be a property manager rather than the beneficial owner. Moreover,

survey designs sample individual properties, making it impossible to construct an owner’s full portfolio or measure concentration. Lastly, the RHFS does not provide information on the individuals or other business entities behind LLCs, so it cannot distinguish between an incorporated individual landlord and a large institutional investor, nor can it connect jointly owned properties into portfolios.

The mailing address approach identifies landlords and their portfolios by aggregating properties that share a common owner mailing address in deed or tax assessor records. Rental properties are identified as those whose property address differs from the mailing address, final owners are identified as a unique mailing address, and portfolios are measured as the set of properties whose tax information is mailed to that address. This approach provides national coverage, but introduces measurement error in identifying rental properties, final owners, and portfolios. The use of address mismatch as a proxy for rental status can misclassify tenure in both directions: rental properties may be missed when an absentee landlord’s mailing address happens to match the property address (e.g., a landlord who lives in one unit of a multi-unit building), and owner-occupied properties may be incorrectly classified as rentals when the owner’s mail goes to a PO box, accountant, or another address that differs from the property. The mailing address approach can also generate portfolios that are either artificially small or artificially large. Portfolio sizes may be understated when commonly owned properties use different mailing addresses, for example when ownership is fragmented across legal entities or when correspondence is sent to multiple locations. Conversely, portfolio sizes may be overstated when unrelated properties share a mailing address through property managers, leasing agents, attorneys, accountants, or other intermediaries. In Section 4, we quantify the extent of these problems.

Local data sources can provide high-quality information about ownership within a particular locality. Cities like New York can impose local reporting requirements that identify rental properties and reveal landlord identities, and several studies exploit such data to study ownership and market power ([Harwood et al., 2025](#); [Watson and Ziv, 2025a,b](#)). However, these approaches are not national in scope and underestimate portfolios for landlords that own property in other jurisdictions.

National data on selected owner and building types can sometimes supplement parcel records with corporate registration data, linking legally distinct entities that may share common control ([Torres, 2024](#); [Hangen and O’Brien, 2025](#)). These datasets can be national in scope, but they typically include information about only a subset of rental properties and owner types. For example, corporate registries are maintained at the state level and differ in structure and reporting requirements. Shared professional agents can generate spurious links, and existing implementations remain limited to single-city applications, reflecting the difficulty of scaling corporate registry linkage across jurisdictions. Similarly, approaches that identify institutional investors from SEC filings, industry reports, or curated firm lists ([Mills et al., 2019](#); [Garriga et al., 2022](#); [Calder-Wang and Kim, 2024](#); [Gorback et al., 2025](#); [Handbury et al., 2025](#)) provide precise measurement for known large landlords but do not recover the broader distribution of ownership, including the individual landlords who, as we will show, supply the majority of rental housing.

Our approach uses administrative records to address many of these challenges. We begin from the near-universe of property records compiled from deeds and tax assessments by Black Knight. While this chapter discusses a prototype build using only 11 CBSAs, our data allow us to scale to a national dataset including all property ownership spells, uniquely identified by owner name, property, and year of purchase and sale. We identify rental properties by comparing owner PIKs to occupant PIKs in the MAF-ARF. To identify final owners, we trace ownership chains through the Business Register and K-1 filings. We construct Final Owner Networks linking properties through shared beneficial owners. Finally, we construct portfolios as the set of properties and shares owned by each final owner. Administrative data sources are national, so portfolios span jurisdictions by construction.

This approach has its own limitations, some of which are shared by any method that uses deed and assessment records. The underlying administrative records were not designed for longitudinal analysis: owner names are recorded inconsistently across sources and over time, records contain duplicates, and there are no stable person-level identifiers such as SSNs in the property records themselves. These problems affect any approach that constructs ownership from these records, including the mailing address method. We use Census-assigned PIKs where available and name-cleaning procedures to mitigate these issues, but many records lack PIKs, and overlapping and adjacent spells remain a challenge for constructing clean ownership histories. Beyond data quality, missing or misassigned PIKs introduce noise in the identification of both owners and rental properties. Properties owned by individuals through businesses may be misclassified as rentals when the business entity, rather than the individual, holds the deed, and the occupant’s PIK does not match the entity’s PIK. We cannot identify owners of non-filing LLCs or C-corp subsidiaries that do not appear in the Business Register or Schedule K-1 filings. And the construction of Final Owner Networks requires the choices about shareholder count cutoffs and network definitions discussed in Section 3.2. We discuss all of these issues in detail in Section 3.

3 Our approach: Using administrative records to identify landlords and their portfolios

We propose to use administrative records to address each of these challenges in new ways, each with its own strengths and drawbacks.

We start by constructing a national dataset from the near-universe of property records for the United States compiled from deeds and tax assessments by Black Knight. Final owners are identified by tracing owner names through a chain of data sources that reveal ownership relationships among business entities. However, we can only identify business owners who appear in our administrative records, and noisy string matching may introduce noise.

We identify residential rental properties using an owner-occupancy flag from property tax assessment data and by comparing property owner PIKs to occupant PIKs in Census administrative records. While this strategy can make determinations at both the property and housing unit level, it is sensitive to missing PIKs and to PIK misassignment.

We construct portfolios as the set of rental properties and shares owned by each final owner. Administrative data sources are national, so portfolios span jurisdictions by construction.

We end our exercise by categorizing final owners into a set of economically meaningful owner types that can be used for analysis.

3.1 Create unique ownership spells from property records

We begin by constructing data on ownership spells, defined as a continuous period during which a property is held by the same set of owners. Constructing ownership data at the spell level allows us to easily extract cross-sections or build panels at the property or owner level. These spells are constructed from information from three types of property records.

3.1.1 Property record data

Black Knight deed records, property tax assessments, and parcel shapefile. Black Knight (BK) compiles information from recorded deeds in a cumulative deed file, and compiles annual property tax assessment records. We use the 2023 cumulative deed file and property tax assessments for 2013-2022. The legal owner of each property is recorded in the buyer name(s) listed on the deed or the owner name(s) listed on the property assessment files. The U.S. Census Bureau has assigned a protected identification key (PIK) to individual property owners based on their first and last names and address, allowing us to link them to other Census data sets.¹

Since properties are individually assessed for tax purposes, these data include one observation each for single-family homes, multi-unit apartment buildings, and individual units in condominiums. The BK parcel shapefile allows us to identify all properties located at the same parcel (e.g., condo units in the same building). It also allows us to attach Census's property unit ID, the Master Address File identifier (MAFID), to each property using a spatial merge of properties and MAFIDs that fall in the same parcel.

Later, we will also use the assessment files to help identify which properties are rented versus owner-occupied and to attach property and unit characteristics, including owner mailing addresses.

Form 1098 tax returns. Form 1098 is a tax form that lenders use to report how much mortgage interest a borrower paid during the year. We use this form as a high-quality supplementary source of information about owner PIKs matched to properties, since it includes richer identifying information than the deed and property tax assessment records that helps Census successfully assign PIKs.

3.1.2 Building ownership spells

To build ownership spells, we integrate annual property tax assessment files with recorded deeds, using the Form 1098 data to fill in missing PIKs. Then we consolidate records of the same owner and property into a single observation uniquely identified by owner name(s), property, and the start and end years of ownership.

The challenge is that the same owner's name is not always recorded identically across files, and dates of purchase and sale may be missing in some files and not others. Furthermore, because the property records only include owner name and address without other identifying

¹The Census Bureau PVS system assigns PIKs to individuals in the Black Knight data using name and address, prioritizing mailing address. If a name-mailing address combination does not yield a match, then PVS attempts a match using name-property address. PIK assignment fails for some individuals whose identifying information does not clearly map to a unique person. For a detailed overview of the PVS process and why some individuals are not assigned PIKs, see [Layne et al. \(2014\)](#) and [Wagner and Layne \(2014\)](#).

information like birth year or SSN, Census is not always able to assign a PIK to each individual owner. Indeed, the same owner may be assigned a PIK in some files and not others, or even different PIKs in different files, depending on the completeness of the name record.

Consequently, we begin by cleaning and standardizing owner names in the deed and assessment records. Deeds and property assessments can each record up to two possible owners. Owner names are reported in four fields: the first name of owner 1, the last-and-middle names of owner 1, the first name of owner 2, and the last-and-middle names of owner 2. The fields for owner 2 are missing when there is only one owner. Business names are entered in the last-and-middle name fields, so that records with non-missing first names are likely to be individuals. Names are converted to uppercase and stripped of punctuation and extra whitespace. For likely individuals, we standardize the order of name components so that variations such as “JOHN W SMITH” and “SMITH JOHN W” are treated as equivalent.

To address the issue of missing PIKs, we compile a property-specific lookup that links standardized owner names to person identifiers (PIKs). If a given property-name pair is only ever associated with one non-missing PIK, we assign this PIK to all observations of that property-name pair. If a property-name pair is linked to multiple PIKs, we treat those records as distinct ownership spells. Next, we use 1098 tax returns for single-family homes to fill in remaining missing PIKs. We do this only for single-family homes because there is no ambiguity about which PIK should be assigned to which housing unit within a property.

Next, we append the cleaned deed and assessment records into a single table and reconcile duplicate observations within each property–purchase year. Records from different sources that refer to the same owner-property observation are merged, and owner fields are harmonized so that joint ownership is not split across separate records. Minor spelling differences are aligned within property–year when necessary.

Finally, we define ownership spells by ordering records within each property by start year and forming contiguous runs with identical owner sets. The stop year for each spell is set equal to the next observed start year. We also consolidate adjacent spells when one owner set is a subset of the other (for example, JOHN SMITH and JOHN SMITH; NANCY SMITH), provided that the records correspond to likely individuals rather than business entities. The resulting spells represent the longest continuous intervals during which a stable ownership set is observed and retain merged information from both deeds and property assessments. When two spells overlap, we retain the spell with the most non-missing information.

The ownership spell data include a spell id, property id, fields for the first and last name of owner 1 and owner 2, year start, and year stop. For the exercises in this chapter, we study the cross section of properties in 2019 by subsetting to spells that overlap with that year.²

3.2 Identify a final owner for each ownership spell

The main obstacle to identifying property owners is that it is common to hold properties through pass-through entities like Limited Liability Corporations (LLCs). We trace the owner name in an ownership spell through a series of administrative data sets that reveal the relationships between businesses that may own property and other business entities that may

²PIK assignment quality could be improved by expanding our use of Form 1098 data beyond single-family homes. This requires unit-level assignment of PIKs within multi-unit properties, but the richer identifying information in these filings would meaningfully increase the PIK rate.

in turn own them. In this way, we are able to follow an ownership spell from the LLC on the deed through several holding companies until we arrive at a final set of shareholders.

The first step is to find source data that covers the universe of property-owning businesses. We assemble records from several administrative data sources that together represent the majority of businesses that may hold property.

3.2.1 Match data sources

Business Register (BR). This dataset is a comprehensive registry of businesses operating in the U.S., containing information on ownership structures, firm identifiers (FIRMID), and industry codes. This allows us to match property owners listed in BK to registered firms and trace ownership hierarchies across business entities. We use information from 1) the employer-unit BR, which includes non-farm establishments with at least one paid employee, subject to unemployment insurance reporting, and 2) the SSN-unit BR, which includes sole-proprietorships with payroll.

Internal Revenue Service (IRS) Schedule K-1 Data. Schedule K-1 filings provide information on the ownership structure of pass-through entities. K-1s are filed by partnerships, S-corporations, and estates and trusts to report how income and other tax attributes flow from an issuing entity to its owners. Crucially, K-1s identify both the issuing entity (via its employer identification number, EIN) and each owner’s tax identifier, which may itself be another EIN when ownership is held by a business entity rather than an individual. We construct directed ownership links from each pass-through entity (owned EIN) to each reported partner/shareholder identifier (owner EIN or PIK). We then trace the chain of owned and owner entities until they terminate at an individual or at an entity that does not file a K-1, such as a C-corporation. To limit the expansion of ownership chains when entities have many reported owners, we cap chain length at 12 links and prune expansions when an entity has more than 10 owners, flagging pruned/truncated cases. The resulting output maps each owned EIN to one or more final owners and the minimum chain length. These ownership chains are central to constructing property portfolios for business owners, as the partners of S-corporations listed on K-1s may include the LLCs observed in property records.

Real Estate Investment Trusts (REITs) and their subsidiaries. We compile information on publicly traded Real Estate Investment Trusts (REITs) and their subsidiaries from firms’ publicly-available annual Form 10-K filings submitted to the U.S. Securities and Exchange Commission. These filings report the legal entities that make up each REIT’s corporate structure. We extract subsidiary names and parent–subsidiary relationships. Because REITs typically hold properties through numerous subsidiaries, identifying these entities helps link properties that share a common publicly traded parent.

Longitudinal Business Database (LBD). This dataset is a census of business establishments with paid employees. We use the LBD to identify parent companies for subsidiaries that appear in the BR or the K-1 data, allowing us to trace the corporate ownership chain a step further. Future work could leverage the LBD to document new facts about corporate landlords.

One category of business is not captured by administrative data: special-purpose LLCs which are separate legal entities distinct from their owners/shareholders but disregarded entities for tax purposes. Since they do not file business tax forms, they do not appear in the Business Register or Schedule K-1s. This type of LLC is sometimes used by developers to purchase land on which they will construct a multi-family building. This special-purpose LLC can own and manage the property, or the developer can sell the special-purpose LLC so that it becomes a wholly owned subsidiary of the new owner. We may be able to improve our match rate by incorporating additional data on these special-purpose disregarded LLCs from a source like OpenCorporates (see [Gorback et al., 2025](#), for a description of this dataset).

3.2.2 Matching owner names to match data sources

Next, we identify final owners for each spell by cascading through potential owner match types, removing each spell from the unmatched pool once a match is found. Because a spell may match multiple sources, this ordering reflects our judgment about the highest-quality matches and can affect the resulting owner types. Figure 1 illustrates the order in which we move through sources.

We begin by identifying government and unincorporated individual owners, who can be identified directly from the property record. Government owners can be identified using keywords in owner names (e.g., “Housing Authority” or “HUD”). Unincorporated individual owners can be identified if the owner name 1) has been matched to a PIK, or 2) has not been matched to a PIK, but the first and last name fields are non-missing. Again, the property record data inputs individuals’ first names into a first name field, while individuals’ last names and business names go into a last name field. This means that entries with a non-missing first name are likely to be individual owners. However, we observe occasional errors with data entry. To make sure that we identify only individual owners, we add two additional keyword requirements. First, the first name field must contain a first name that is associated with a PIK for at least one property record in our sample. This helps weed out erroneous entries into the first name field, without relying on a measure of how common the name is which might introduce bias against certain types of names. Second, the owner name cannot contain a business keyword like “Inc” or “Co.” Table 2 reports examples of government-related keywords chosen without looking at the owner names; to protect privacy, we cannot disclose the full set keywords selected by manual inspection of the data. In these cases, we take the named individual or government agency as the final owner.

For the remaining spells, we apply a fuzzy name matching procedure. In each step, we match owner names to candidate entities in a source dataset using a Jaccard-based inner join on character 6-grams implemented with locality-sensitive hashing (30 bands by default). Candidate pairs are generated using a similarity threshold of 0.80. We drop candidates with first-letter mismatches and retain the highest-scoring candidate for each owner name. If there is a tie, we select the candidate registered in the zip code closest to the zip code of the property. For each source dataset, we run the algorithm first on exact names and then on name variants that remove or add common legal suffixes (e.g., LLC, INC, and CORP).

We first apply this procedure to the employer-unit Business Register, imposing the additional requirement that the business was registered before the ownership spell began to make sure that we are only matching to businesses that existed at the time of purchase. Then we fuzzy match to the SSN-unit Business Register following the same rules.

Third, we fuzzy match to names of owned entities reported in Schedule K-1 tax filings, which include both the entity’s Employer Identification Number (EIN) and the partner’s tax identification information. These matches expand our ability to identify individuals and firms with ownership stakes in the entity named in the deed.

Next, we use keywords to identify individual trusts among the remaining unmatched ownership spells. These owner names must contain a trust keyword (e.g., TRUST, FT, RET) and a PIKed first name.

Fifth, we fuzzy name match to REITs and REIT subsidiaries identified from 10-K filings. Finally, we attempt another fuzzy name match to the employer-unit and SSN-unit BR without imposing the time overlap requirement, recognizing that there may be noise in both the ownership spell start year and the business registration year.

For spells matched to the BR, Schedule K-1 entities, or REIT subsidiaries, we trace ownership further using K-1 filings to construct ownership chains that link each owned entity to its ultimate owner(s), identified by EINs and/or shareholder PIKs. When an entity is traced to more than two final owners (e.g., shareholders or partners), we designate the entity—rather than the shareholders—as the final owner, to reflect the collective nature of control. We merge the final owner identified with the K-1 filings with the LBD and REITs, so that for each ownership spell involving a business owner, we arrive at the end of the observable chain. In these cases, the final owners are either a set of individual shareholders or a business entity that is not owned by any other business entity that appears in the LBD, Schedule K-1s, or list of publicly traded REITs.

Although the remaining unmatched spells cannot be matched to a specific individual or business owner, we can use string patterns in the owner name to categorize them as likely businesses or likely individuals. Likely businesses contain a business keyword and do not contain a common first name, while likely individuals contain a common first name and do not contain a business keyword.

Residual unmatched spells likely include sole proprietorships (which typically do not file Schedule K-1 and may appear in the BR under the owner’s name), subsidiaries of C-corporations that do not file Schedule K-1, and noisy or erroneous names that fail to meet our conservative similarity criteria.

To make this concrete, consider the examples in Table 3 (we omit public entities and individual trusts from the example because they are classified directly using the keyword rules described in Table 2). Column 2 contains the owner name listed on the deed or property assessment. Property 1 is owned by an individual, Jane Doe. Property 2 lists 123 Main St LLC as the property owner. Our methodology identifies John Smith as the owner of the business, and therefore classifies him as the final owner. Similarly, John Smith owns the business 125 Main St LLC and is the final owner of property 3. Property 4 is owned by 789 1st Ave LLC, which our algorithm identifies as a subsidiary of City Realty Inc. City Realty Inc has one hundred shareholders, including John Smith. Since the business has more than 2 shareholders, we assign City Realty Inc as the final owner of property 4. Property 5 is owned by 101 Broadway Ave LLC, which our procedure identifies as jointly owned by two individuals, Jane Doe and Taylor Lee; we therefore treat both as final owners of Property 5, each with a 50 percent share. Property 6 is owned by Landlord Inc, a business which is not owned by any other business entity in the Business Register, K-1s, or REITs; we therefore treat Landlord Inc itself as the final owner.

Table 4 reports, averaged across CBSAs, the share of ownership spells with at least one

match to each source in the identification cascade. Roughly 32 percent of spells are attributed to PIKed individuals, 31 percent are matched to the Business Register, and 8.8 percent are linked through Schedule K-1 filings. Nearly 3 percent are unPIKed individuals, and another 22 percent are classified as likely unPIKed individuals based on keyword patterns. Keywords identify public owners, likely business owners, or individual trusts for less than 1.5 percent of ownership spells. Only 1.8 percent remain fully unmatched after all steps of the cascade. Because an ownership spell can be matched to multiple sources when an entity appears in a multi-link ownership chain, these shares do not sum to one.

3.3 Identify residential rental properties

We have now created a data set of ownership spells matched with final owners for all owned properties that appear in the deed or assessment data. The next step is to identify which of these are residential rental properties. However, there is no national database of rental vs. owner-occupied properties.

We develop a strategy for combining information from property tax assessments with Census address history records to identify residential rentals. The property tax assessment data include a flag for owner occupancy, derived by Black Knight based on homeowner tax exemptions that are only available for owner-occupied primary residences and on whether the mailing and property addresses match.

We augment the owner-occupied flag with insights from Census data on address histories from the MAF-ARF. To do this, we first need to create a crosswalk between the taxable properties observed by Black Knight and Census’s internal property unit identifier, the MAFID. We use the Black Knight parcel shapefile to identify all properties in the same structure, or “site,” and spatially merge MAFIDs to these sites. For single-family buildings, there is a single property and a single MAFID. For-rent apartment buildings have multiple MAFIDs at a single property, and multifamily condominiums (“condos”) have multiple property assessments and multiple MAFIDs at the same parcel.

We compare owner PIKs with occupant PIKs in the MAF-ARF to determine if a housing unit is owner-occupied. We designate a property as a rental if it contains more units than owner-occupied units. In multiunit properties, one MAFID within a large residential building might be owner-occupied while the others are rentals. We therefore retain all MAFIDs for parcels that include at least one rental unit. This allows us to conduct analyses at both the unit and property levels. We drop MAFIDs that were unoccupied from 2017 to 2020 according to the MAF-ARF, as these are likely inactive, vacant, or second-homes that should be excluded from the active rental housing stock.

Using all three pieces of information, we identify rental units as those which 1) are not in properties flagged as owner-occupied in the 2019 Black Knight assessment file; 2) the resident PIK in the MAF-ARF does not match the owner PIK; and 3) were associated with at least one resident PIK from 2017-2020 in the MAF-ARF. We are confident that this procedure does well at avoiding false positives, but it undercounts rental units relative to the 2019 ACS 1-year estimates.³

³Future work may be able to improve on our identification of rental units with prediction methods, using the ACS subsample as the training data.

3.4 Construct final owner networks and their portfolios

With the final owner information and rental/owner-occupancy classification in hand, we next construct ownership portfolios and Final Owner Networks (FONs) that account for joint ownership relationships. A FON groups together owners who share ownership stakes in at least one property, allowing us to capture situations in which individuals co-own properties and may coordinate decisions across them.

We implement this using a graph representation in which nodes represent final owners (individuals or firms) and edges connect owners who jointly own a property. Each disconnected component of this graph defines a Final Owner Network with shared property interests. If a property has more than two final owners, we assume that it is managed independently and therefore treat the firm entity as the final owner. For each owner (node), their *portfolio* consists of the set of properties in which they have a stake and their corresponding ownership share. Because we do not observe precise equity stakes, we assume equal ownership shares among co-owners.

In the example records shown in Table 3, Jane Doe owns 100% of property 1 and 50% of property 5, so her portfolio is $\{(1, 100\%), (5, 50\%)\}$. John Smith’s portfolio is $\{(2, 100\%), (3, 100\%)\}$, City Realty Inc’s portfolio is $\{(4, 100\%)\}$, Taylor Lee’s portfolio is $\{(5, 50\%)\}$, and Landlord Inc’s portfolio is $\{(6, 100\%)\}$. Final Owner Networks capture connections between owners who share ownership stakes in the same property. In this example, Jane Doe and Taylor Lee jointly own property 5 and therefore belong to the same Final Owner Network, even though their individual portfolios differ. The FON links owners who may coordinate decisions across their jointly owned properties.

For this reason, our treatment of corporate entities depends on the number of individual owners identified in the ownership chain. Because City Realty Inc has one hundred shareholders, including John Smith, it is likely that the corporation rather than the set of shareholders is the decision-making entity, so we treat the corporation as the final owner of property 4. If City Realty Inc had one or two shareholders, we would instead assign ownership shares to the individual shareholders. In that case, John Smith’s portfolio would become $\{(2, 100\%), (3, 100\%), (4, 50\%)\}$.

3.5 Classifying Ownership Types

Next, we classify these Final Owner Networks into economically meaningful ownership types:

1. *Unincorporated Individuals*: individuals who own property directly and are not linked to any business entity. Properties in this category have owners with an assigned PIK, or owners with non-missing first and last names on the property records.
2. *Incorporated Individuals*: business entities with one or two individual shareholders. Although the property records list a corporate owner, the final owners are individuals with assigned PIKs.
3. *Individual Trusts*: trusts whose beneficiaries are individuals. The owner name listed on the property records contains terms such as “Family Trust” as well as a common first name.
4. *Business Entities*: firms or partnerships with three or more shareholders, which may include individuals or corporate entities. This category includes legal entities such as

LLCs and REITs, and may be organized as pass-through entities (e.g., partnerships or S-corporations) or as C-corporations.

5. *Public*: housing authorities and other government agencies. The owner name listed on the property records contains terms such as “Housing Authority”, “Housing Agency”, “City of”, “Municipality of”, etc.
6. *Unmatched*: owners that cannot be linked to a specific individual or firm through our matching procedure. These likely include sole proprietorships (which typically do not file Schedule K-1 and may appear in the BR under the owner’s name), subsidiaries of C-corporations that do not file Schedule K-1, and noisy or erroneous owner names that fail to meet our conservative similarity criteria.

This typology enables more meaningful comparisons than relying on LLC labels alone, which can obscure differences between family-owned operations and institutional landlords.⁴

In the examples shown in Table 3, the owner of Property 1 is an unincorporated individual. Properties 2, 3, and 5 are listed in the property records as owned by businesses, but our ownership-linking procedure reveals that these entities are owned by one or two individuals; we therefore classify these properties as owned by incorporated individuals. Property 4 is owned by a business with more than two shareholders and is thus classified as a business entity. For Property 6, we are unable to identify the underlying shareholders of the listed business owner, so we classify it as a business entity as well.

4 Descriptive statistics based on the subset of properties we can validate with the ACS

We close with descriptive statistics for the cross-section of 2019 ownership in 11 CBSAs for properties surveyed in the 2019 ACS. Because the ACS is a roughly one-percent sample, the statistics reported here reflect a small fraction of the rental stock in each CBSA. However, because the ACS independently identifies whether a housing unit is renter-occupied, we are confident that rental properties are correctly classified in this subsample (up to reporting error in the ACS survey data). We focus on this validated subset and leave descriptive results for the full universe of rental properties to a future data build that addresses the coverage challenges discussed in Section 3.3.

4.1 Ownership type composition

Figure 2 shows the distribution of final owner types for all properties and for rental properties separately. Individual owners own the majority of properties in both samples, but the composition shifts meaningfully when restricting to rentals. Among all properties, unincorporated individuals own 82.3 percent, reflecting the prevalence of owner-occupied single-family homes. Among rental properties, unincorporated individuals own 60.8 percent and incorporated individuals—business entities with one or two shareholders—account for an additional 7.6

⁴One category our typology does not separate is intra-family rental arrangements, in which an individual owner rents to family members. In principle, family relationships in 1040 filings (spouses, claimed dependents) could be used to flag these cases. We leave this for future work.

percent. Trusts hold 0.3 percent of rental properties, businesses with more than two shareholders hold 25 percent, 5.6 percent remain unmatched and 0.7 percent are publicly-owned. These patterns indicate that individual landlords, whether operating in their own names or through LLCs, supply the large majority of rental housing in our sample.

Figure 3 breaks down ownership composition by building type. Individual owners are most prevalent among single-family rentals and smaller multifamily structures. Business entities and incorporated individuals account for a larger share of 5+ unit multifamily properties, consistent with the expectation that larger buildings are more likely to be held through formal corporate structures. The unmatched share also varies across building types, suggesting that the likelihood of successful entity matching depends in part on the complexity of ownership arrangements typical of each segment.

4.2 Ownership composition across metropolitan areas

Figure 4 shows how the composition of rental ownership varies across the 11 CBSAs. While unincorporated individual owners are the largest category in every market, their share ranges from roughly 50 percent in Houston to over 70 percent in Chicago. Incorporated individuals who own rental properties via LLCs are particularly common in Los Angeles, Phoenix, and Denver. The combined share of public, trust, and unmatched owners also varies across CBSAs, reflecting both differences in local ownership practices and data quality variation, for example due to the prevalence of ownership structures that are difficult to trace through administrative records.

4.3 Characteristics of individual landlords

Table 5 reports demographic and income characteristics of PIKed individual landlords, who can be matched to their 2019 IRS 1040 filings. We compare characteristics for three types of individual landlords: unincorporated individuals, whose name is on the deed; incorporated individuals, who own a property through an LLC either alone or with one other individual; and individual business owners, who co-own property through an LLC or other business with at least two other individuals. Table 5 has a smaller sample of landlords relative to the full sample of landlords of ACS rental units. More specifically, there are fewer unincorporated individuals in Table 5’s sample, as unincorporated individuals are relatively less likely to be PIKed than incorporated individuals and individual business owners.

We find that unincorporated individual landlords have substantially lower incomes than their incorporated counterparts: mean adjusted gross income (AGI) is \$481,300 for unincorporated individuals compared to \$930,400 for incorporated individuals. Business owners have the highest mean AGI at \$1,730,000. Incorporated individuals are more likely to be male (75.6 percent) than unincorporated individuals (56.1 percent) and earn substantially higher W-2 wages (\$458,400 vs. \$102,000). Nearly all incorporated and business-owner landlords file a Schedule E (97.1 and 95.3 percent, respectively), compared to 78.8 percent of unincorporated individuals. The AGI distribution reveals that 18.6 percent of incorporated individuals and 25.3 percent of business owners have AGI above \$1 million, compared to only 4.3 percent of unincorporated individuals. At the other end of the distribution, 43.3 percent of unincorporated individual landlords have AGI below \$100,000.

These differences suggest that incorporation status is correlated with the economic scale

of a landlord’s operations and overall wealth. Incorporated individuals and business owners appear to be wealthier, earn more from wages, and have income profiles consistent with more diversified economic activity.

Table 6 reports mean income measures separately by CBSA and incorporation status. There is substantial geographic variation, particularly in AGI. Among unincorporated individuals, mean AGI ranges from \$43,530 in Denver to \$903,800 in Los Angeles—a pattern likely driven by differences in property values and capital gains across markets. Schedule income shows a similar pattern, with unincorporated individuals in Los Angeles reporting mean Schedule E rental real estate and royalty income of \$1,804,000 compared to \$36,410 in Hartford. Among incorporated individuals, mean AGI ranges from \$390,800 in Hartford to \$1,815,000 in Atlanta. Wage and salary income is more evenly distributed across CBSAs for unincorporated individuals but varies more widely for incorporated individuals, from \$158,900 in Washington, DC to \$581,000 in Charlotte.

4.4 Comparison with the mailing address approach

Figure 5 provides evidence that the commonly used mailing address approach can both understate and overstate portfolio sizes. Panel (a) focuses on rental properties whose final owner holds a multi-property portfolio and reports the share of these properties for which the property tax assessment is mailed to the property address. Among incorporated individuals with 1–2 shareholders, 4.6 percent of properties receive their assessments at the property itself. For businesses, the corresponding figure is 5.0 percent, and for incorporated individuals with 10+ shareholders, it is 5.6 percent. In these cases, the mailing address approach would classify the property as owner-occupied and exclude it from the owner’s portfolio, leading to an undercount of portfolio size.

Panel (b) shows the potential for overestimation. Among mailing addresses that differ from the property address, 38.3 percent receive assessments associated with two or more distinct final owners, and 24.5 percent are shared by five or more final owners. When unrelated owners share a mailing address—for example, through a common property management company, accountant, law firm, or registered agent—the mailing address approach incorrectly places their properties in the same portfolio, inflating the measure of concentration.

5 Directions for benchmarking and validity testing

No single data source provides a definitive measure of rental property ownership that we can use as ground truth. However, comparing our approach to multiple independent sources—each with different strengths and limitations—can help assess the reliability of our methodology and inform researchers about the sensitivity of conclusions to methodological choice. We organize future work along two dimensions: diagnostics on the entity matching procedure itself, and comparisons of the resulting ownership measures to independent data sources. Where comparisons reveal discrepancies, they can guide specific improvements to the matching process, including adjustments to matching thresholds, incorporation of additional data sources such as state corporate registries or OpenCorporates, and extensions to handle family ownership structures.

Matching diagnostics. Several exercises can assess the quality of our entity matching without requiring external data. First, varying the Jaccard similarity threshold used in the fuzzy string matching would reveal the sensitivity of match rates, portfolio sizes, and ownership concentration to this parameter. A high threshold is not unambiguously “conservative:” while it reduces false matches (incorrectly linked entities), it increases false non-matches (missed links between entities that are in fact commonly owned). Examining how the unmatched rate, portfolio size distribution, and tract-level Herfindahl–Hirschman Index (HHI) change as the threshold moves across a range of reasonable values would help assess the importance of this choice. If the main descriptive statistics about ownership structure are stable across a range of thresholds, that would suggest our conclusions are robust to this particular methodological choice; if they are not, identifying which findings are most sensitive would help prioritize further investment in matching quality.

Second, our current methodology matches on entity names but does not exploit family relationships among individual owners. Census data, including the MAF-ARF and 1040 joint filing records, contain household structure information that could link family members who co-own properties through separate entities. For example, if one spouse owns a property through an LLC and the other spouse owns a different property in their own name, our current approach treats these as separate owners. Incorporating family linkages is primarily a methodological extension rather than a validity test per se, but comparing portfolio size distributions and ownership type classifications with and without family linkages would reveal how much the current approach misses due to within-family fragmentation of ownership, and would help address known limitations of string matching for names that vary across family members or reflect naming conventions in different ethnic communities.

Cross-source comparisons. We can also compare our ownership measures directly to independent data sources. Across all such comparisons, the most informative exercises go beyond comparing aggregate portfolio sizes or ownership type distributions and instead examine ownership links at the property-pair level. For each pair of properties, we can classify whether both approaches link them to a common owner, only one approach links them, or neither does. When one source can serve as a ground truth reference, this decomposition yields estimates of the false match rate (the share of links we identify that the reference source does not confirm) and the false non-match rate (the share of links in the reference source that our approach misses), following standard terminology in the record linkage literature (??). Since no source is definitive, these rates should be computed in both directions—our approach relative to the comparison source, and vice versa—to avoid treating either as ground truth. They can be further broken down by property type, owner type, and Jaccard match score to identify where discrepancies are concentrated.

A natural starting point is New York City, where multifamily landlords are required to register annually with the Department of Housing Preservation and Development (HPD), providing contact information for individuals with ownership stakes exceeding 25 percent. [Watson and Ziv \(2025b\)](#) and [Harwood et al. \(2025\)](#) construct ownership portfolios and landlord classifications from this registry. Matching our administrative-record-based measures to the NYC registry for the overlapping geography would allow a direct comparison of owner type classifications, portfolio sizes, and ownership concentration at the tract or neighborhood level. Because the NYC registry captures ownership through a fundamentally different channel—mandatory self-reporting rather than entity matching through tax records—agreement between

the two approaches would provide strong evidence that both recover economically meaningful ownership. Systematic disagreements would require careful interpretation: they could reflect limitations of our approach (e.g., unmatched entities that the registry correctly identifies) or limitations of the registry (e.g., incomplete compliance for small buildings, coarse self-reported corporate designations that do not distinguish incorporated individuals from large business entities, or inability to trace ownership through tiered LLC structures). Examining the direction and pattern of disagreements—by building size, owner type, and corporate status—would help characterize the strengths and weaknesses of both approaches and point toward specific remedial steps. For example, if the registry identifies ownership links that our approach misses for entities near the Jaccard threshold, that would support lowering the threshold or supplementing the matching with additional data sources. If instead our approach identifies links that the registry misses—for instance, by tracing K-1 ownership chains across entities that register separately—that would provide evidence for the value added of administrative record linkage over local reporting requirements. [Harwood et al. \(2025\)](#) independently construct landlord portfolios for NYC multifamily buildings by linking registry contacts through geocoded business addresses and names, providing a portfolio-level comparison that does not rely on our entity matching. [Watson and Ziv \(2025b\)](#) construct tract-level ownership concentration measures from the same registry, enabling a direct comparison of HHI measures.

Two additional CBSAs in our sample—Jacksonville and Atlanta—offer comparison opportunities through state corporate registries. [Torres \(2024\)](#) uses the Florida business registry combined with tax parcel data to construct ownership portfolios in Jacksonville via a graph-based model that links subsidiary companies through shared officers and registered agents. Because this approach traces corporate structures through state registry linkages rather than federal BR and K-1 matching, it provides an independent check on our methodology using a fundamentally different identification channel. Similarly, [An et al. \(2024\)](#) use the Georgia Secretary of State business registry to link shell companies and subsidiaries to parent firms in Atlanta, combining registry data with name standardization algorithms applied to tax parcel records. Their approach is primarily focused on identifying large institutional investors, but for the subset of entity-owned properties they study, a comparison to our classifications would reveal whether our BR/K-1 matching recovers the same ownership links that state registries identify, and whether our approach captures additional connections through K-1 ownership chains that state registries cannot observe. As with the NYC comparison, applying the property-pair decomposition described above would provide the most granular assessment of where the approaches agree and diverge.

Finally, as documented in Section 4, the mailing address approach can both understate and overstate portfolio sizes relative to our methodology, and can produce misleading measures of ownership concentration. Extending this analysis to property-pair-level comparisons of owner type and portfolio assignment would provide a more granular picture of exactly which types of properties and owners are most affected by the choice of methodology, helping researchers in other settings assess whether the mailing address approach is adequate for their particular application.

6 Conclusion

This chapter introduces a methodology for measuring rental property ownership in the United States that uses administrative records to trace ownership through corporate structures and identify the individuals and firms that ultimately control rental housing. By linking property records to the Census Bureau’s Business Register, IRS Schedule K-1 filings, and Form 10-K filings, we follow ownership chains through LLCs, partnerships, and other pass-through entities to identify final owners and construct property portfolios. This approach addresses several limitations of existing methods: it does not rely on mailing address matching to define portfolios, it distinguishes incorporated individuals from large business entities, and it can be scaled to national coverage.

We apply our methodology to 11 large CBSAs and report descriptive findings for the subset of properties surveyed in the 2019 American Community Survey, for which we can independently validate rental status. Several findings emerge. Individual landlords—both unincorporated and incorporated—own the large majority of rental properties, accounting for nearly 70 percent of the rental stock in our sample. Ownership composition varies meaningfully across building types and metropolitan areas, with incorporated individuals more prevalent in larger multifamily structures. Individual landlords who hold property through corporate entities are substantially wealthier than those who own directly, with higher AGI, higher wage income, and a greater likelihood of filing Schedule E. We also demonstrate that the mailing address approach, while widely used, can both understate and overstate portfolio sizes: assessments mailed to the property address cause portfolios to be undercounted, while shared mailing addresses from intermediaries such as property managers and law firms cause unrelated owners to be grouped together.

Our approach has important limitations that are discussed throughout the chapter. We cannot trace ownership through tax-disregarded entity LLCs or C-corporation subsidiaries that do not appear separately in administrative records. Noisy name matching may introduce both false matches and missed links. Our current rental identification strategy undercounts rental units relative to the ACS, suggesting room for improvement in how we classify tenure status from administrative data. The descriptive results presented here are based on a one-percent subsample and may not generalize to the full rental stock.

Despite these limitations, the methodology provides a foundation for future work along several dimensions. As discussed in Section 5, validation tests using local data sources such as the New York City landlord registry and state corporate registries can help assess matching quality and guide improvements. Incorporating additional data, such as state business registries and family relationship information from tax records, could reduce the unmatched share. Extending coverage beyond the ACS subsample to the full universe of properties, and from 11 CBSAs to the entire United States, would enable nationally representative analysis of ownership concentration, portfolio size distributions, and the relationship between owner characteristics and housing market outcomes. We view this chapter as a proof of concept for a measurement infrastructure that can support a broad research agenda on rental housing ownership in the United States.

References

- An, Brian Y.**, “The Influence of Institutional Single-Family Rental Investors on Homeownership: Who Gets Targeted and Pushed Out of the Local Market?,” *Journal of Planning Education and Research*, 2024, 44 (4), 2231–2250.
- , **Andrew Jakabovics, Anthony W. Orlando, Seva Rodnyansky, and Eunjee Son**, “Who Owns America? A Methodology for Identifying Landlords’ Ownership Scale and the Implications for Targeted Code Enforcement,” *Journal of the American Planning Association*, 2024, 90 (4), 627–641.
- Austin, Neroli**, “Keeping Up with the Blackstones: Institutional Investors and Gentrification,” 2022. SSRN Working Paper.
- Barbieri, Felipe and Gregory Dobbels**, “Market Power and the Welfare Effects of Institutional Landlords,” 2025. Working paper.
- Calder-Wang, Sophie and Gi Heung Kim**, “Algorithmic Pricing in Multifamily Rentals: Efficiency Gains or Price Coordination?,” 2024. SSRN Working Paper.
- Chang, Konhee**, “Diversifying the Suburbs: Rental Supply and Spatial Inequality,” 2025. Working paper.
- Cocco, João, S. Lakshmi Naaraayanan, and Jagdish Tripathy**, “Individual Landlords in the Mortgage Market,” Technical Report, Working Paper 2024.
- Coven, Joshua**, “The Impact of Institutional Investors on Homeownership and Neighborhood Access,” 2025. Working paper.
- Damen, Sven, Matthijs Korevaar, and Stijn Van Nieuwerburgh**, “An Alpha in Affordable Housing?,” 2025. SSRN Working Paper.
- Fesko, Luke**, “Landlord Corporatization and Eviction in North Carolina,” 2025. SSRN Working Paper.
- Ganduri, Rohan, Steven Chong Xiao, and Serena Wenjing Xiao**, “Tracing the Source of Liquidity for Distressed Housing Markets,” *Real Estate Economics*, 2023, 51 (2), 408–440.
- Garriga, Carlos, Rodolfo E. Manuelli, and Adrian Peralta-Alva**, “The Economic Effects of Real Estate Investors,” *Real Estate Economics*, 2022, 50 (6), 1451–1483.
- Gorback, Caitlin, Franklin Qian, and Zipei Zhu**, “Impact of Institutional Owners on Housing Markets,” 2025. SSRN Working Paper.
- Gurun, Umit G., Jiabin Wu, Steven Chong Xiao, and Serena Wenjing Xiao**, “Do Wall Street Landlords Undermine Renters’ Welfare?,” *The Review of Financial Studies*, 2023, 36 (1), 70–121.
- Handbury, Jessie, Samuel K. Hughes, and Benjamin J. Keys**, “Segmentation and Returns on Rental Housing,” 2025. Working paper.

- Hangen, Forrest and Daniel T. O'Brien**, “Linking Landlords to Uncover Ownership Obscurity,” *Housing Studies*, 2025, *40* (4), 940–965.
- Harwood, Katharine W. H., Ingrid Gould Ellen, and Katherine O'Regan**, “The Rise of Corporate Landlords: An Examination of Behavioral Differences in the Multifamily Market,” *Real Estate Economics*, 2025.
- Lambie-Hanson, Lauren, Wenli Li, and Michael Slonkosky**, “Real Estate Investors and the U.S. Housing Recovery,” *Real Estate Economics*, 2022, *50* (6), 1425–1461.
- Layne, Mary, Deborah Wagner, and Cynthia Rothhaas**, “Estimating Record Linkage False Match Rate for the Person Identification Validation System,” CARRA Working Paper 2014-02, U.S. Census Bureau 2014.
- Lee, Keyoung and David Wylie**, “Institutional Investors, Rents, and Neighborhood Change in the Single Family Residential Market,” Working Paper 24-13, Federal Reserve Bank of Philadelphia 2024.
- Levy, Antoine**, “Housing Policy with Home-Biased Landlords: Evidence from French Rental Markets,” 2022. Working paper.
- Linger, Jacob, Hal J. Singer, and Ted Tatos**, “Does a Lack of Competition Exacerbate Inflation? A Case Study of Florida Rental Markets,” 2022. SSRN Working Paper.
- Mills, James, Raven Molloy, and Rebecca Zarutskie**, “Large-Scale Buy-to-Rent Investors in the Single-Family Housing Market: The Emergence of a New Asset Class,” *Real Estate Economics*, 2019, *47* (2), 399–430.
- Moktan, Sidharth**, “An Empirical Equilibrium Model of the Markets for Rental and Owner-Occupied Housing,” 2025. Working paper.
- Ramoutar, Fern**, “Market Power in Residential Real Estate: Evidence from Chicago Rental Properties,” 2024. Working paper.
- Samuels, Peleg Ronen**, “The Effect of Local Landlords on Rental Housing Prices and Returns,” 2025. SSRN Working Paper.
- Torres, Renz**, “Who Owns Our Homes? Methods to Group and Unmask Anonymous Corporate Owners,” *Cityscape*, 2024, *26* (1), 339–362.
- Travis, Adam**, “The Organization of Neglect: Limited Liability Companies and Housing Disinvestment,” *American Sociological Review*, 2019, *84* (1), 142–170.
- Wagner, Deborah and Mary Layne**, “The Person Identification Validation System (PVS): Applying the Center for Administrative Records Research and Applications’ (CARRA) Record Linkage Software,” CARRA Working Paper 2014-01, U.S. Census Bureau 2014.
- Wang, Zhichun and Daojing Zhai**, “The Expansion and Dynamic Equilibrium Effects of Institutional Landlords,” 2025. Working paper.
- Watson, C. Luke and Oren Ziv**, “Is the Rent Too High? Land Ownership and Monopoly Power,” 2025. Working paper.

- **and** —, “A Test for Pricing Power in Urban Housing Markets,” *The Review of Economics and Statistics*, 2025, pp. 1–33.

7 Tables

Table 1: Comparison of methodologies used to identify rental housing ownership

Methodology	Example	Features			
		National	Identify final owners	Identify rentals	Construct portfolios
Survey-based	Residential & Housing Finance Survey	Nationally representative subsample	Sometimes	Yes	No
Mailing address	CoreLogic property assessments	Yes	Can mismeasure	Can mismeasure	Can mismeasure
Local sources	New York City registries	No	Yes	Yes	Yes — within that locality
National data on selected owner and building types	OpenCorporates	Yes	Subset	Subset	Subset
Administrative record approach	Property records + business tax returns	Yes	Yes	Yes	Yes

Table 2: Keywords used to identify government and individual trust owners

Owner type	Keywords used in classification
Government	Borough of, City of, County of, Commonwealth of, Department of, Town of, Township of, State of, Village of, District of Columbia, City Dept/Department, Fire Dept/Department, Highway authority, Highway dept, Highway division, Housing Authority, HUD, Municipal authority, Parks Dept/Department, Police Dept/Department, Right of way, School district, Board of education, school board, USDOT,...
Unincorporated individual	First names associated with PIKs, such as Michael, John, Andrea,...
Individual trust	A first name + Family trust, Trust, Land trust, FT, (EST)
Business	LLC, Inc, Co, Mgmt,...

Notes: This table reports a subset of the keywords used to identify government owners, unincorporated individuals, individual trusts, and businesses. These examples were identified without looking at the owner names; to protect privacy, we cannot disclose the list of keywords we selected based on observing the owner names. To identify a government owner, a government keyword must be present. To identify an unincorporated individual, two conditions must be met: owner name 1) includes a first name associated with a PIK on at least one property record in the sample; 2) does not include a business keyword. To identify an individual trust, owner name must 1) include a first name; 2) include a trust keyword. To identify a likely business, owner name must 1) contain a business keyword; 2) not contain a first name.

Table 3: Synthetic data illustrating the final owner identification procedure

Property ID	Owner name (deed or assessment)	Mailing address (assessment)	Final firm owner (after linking to BR and K-1)	Final owner names (after linking to BR and K-1)
1	Jane Doe	8 Common Rd		Jane Doe
2	123 Main St LLC	123 Main St		John Smith
3	125 Main St LLC	125 Main St		John Smith
4	789 1st Ave LLC	3 Legal Rd	City Realty Inc	John Smith
4	789 1st Ave LLC	3 Legal Rd	City Realty Inc	Shareholder 2
:	:	:	:	:
4	789 1st Ave LLC	3 Legal Rd	City Realty Inc	Shareholder 100
5	101 Broadway Ave LLC	3 Legal Rd		Jane Doe
5	101 Broadway Ave LLC	3 Legal Rd		Taylor Lee
6	Landlord Inc	3 Legal Rd	Landlord Inc	Landlord Inc

Notes: This table presents synthetic data used to illustrate the methodology for identifying final owners, constructing portfolios, and classifying ownership types. The owner name on the deed or assessment is shown alongside the final firm owner and final individual owner name(s) after matching to the Business Register and Schedule K-1 filings. See Section 3.2 for a detailed walkthrough.

Table 4: Match shares by owner data source

Source	Mean share	SD
Government keyword	0.0096	0.0039
PIKed individual	0.3169	0.0578
UnPIKed individual	0.0278	0.0142
Business Register	0.3054	0.0615
Schedule K-1	0.0880	0.0513
Indiv trust keywords	0.0083	0.0070
REITs/10-Ks	0.0033	0.0020
Likely business keywords	0.0010	0.0003
Likely UnPIKed individual keywords	0.2218	0.0366
Unmatched	0.0180	0.0047

Notes: This table reports the mean and standard deviation across CBSAs of the share of ownership spells matched to different sources of information about owners, following the methodology described in Section 3.2. Note that the mean share does not sum to 1 because the same ownership spell can be matched to multiple sources of business ownership data if it is owned by a multi-link chain.

Table 5: Demographics and income sources of PIKed individual landlords

	Unincorporated individuals	Incorporated individuals	Business owners
% Male	56.1	75.6	70.0
<i>Mean income (\$)</i>			
AGI	481,300	930,400	1,730,000
W-2 wages	102,000	458,400	214,500
Wage & salary	88,880	326,300	178,800
Schedule E	860,800	108,500	470,800
Gross rent charged (\$)	1,418	1,277	1,384
<i>Income filing shares</i>			
% with Schedule E	78.8	97.1	95.3
% with Schedule C	28.1	30.8	25.0
% with Sched. E or C	83.7	97.6	95.8
<i>AGI distribution (mutually exclusive brackets)</i>			
AGI < \$0	3.8	4.5	7.8
\$0 ≤ AGI < \$100k	43.3	15.1	16.2
\$100k ≤ AGI < \$1M	48.6	61.7	50.7
AGI ≥ \$1M	4.3	18.6	25.3
<i>Wage & salary distribution (mutually exclusive brackets)</i>			
W&S = \$0	34.8	31.8	29.6
\$0 < W&S ≤ \$100k	36.8	24.0	30.4
\$100k < W&S ≤ \$1M	27.9	39.3	37.1
W&S > \$1M	0.5	4.9	2.9
<i>N</i>	27,500	531,000	800

Notes: This table reports characteristics of PIKed individual landlords pooled across 11 CBSAs in 2019. We report three categories of individuals: unincorporated, incorporated (1–2 individuals who own a property through an LLC or other business), and ‘business owner’ individuals who co-own property with at least 2 other individuals. Income information is linked to individuals by PIK from 2019 IRS 1040 filings. Schedule E includes positive rental real estate and royalty income. Gross rent charged is the gross rent for a given unit, as reported in the ACS. For details on ownership classification, see Section 3.5.

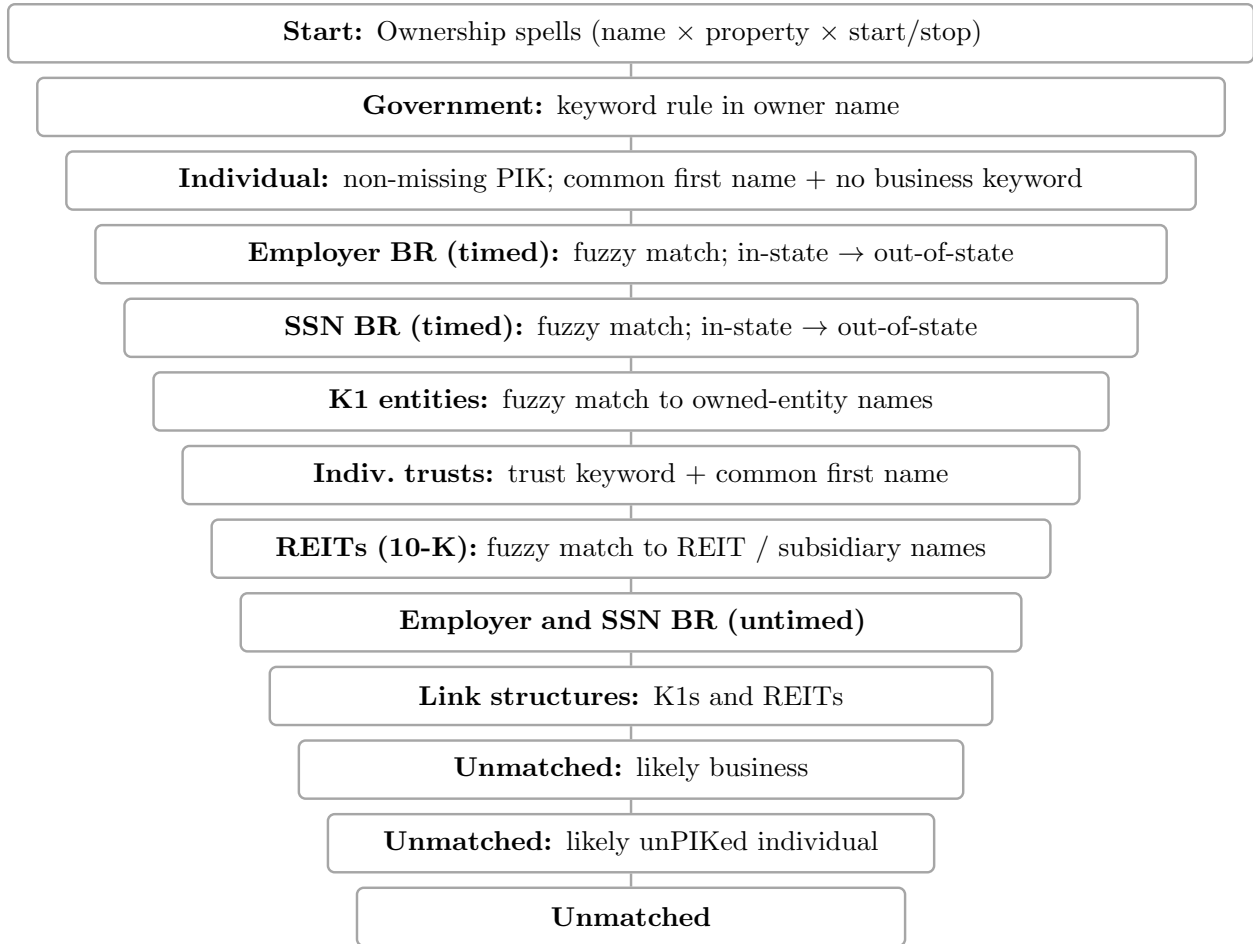
Table 6: Income characteristics of individual landlords by CBSA

CBSA	Unincorporated individuals				Incorporated individuals			
	AGI	Wage/Sal.	W-2	Sched. E	AGI	Wage/Sal.	W-2	Sched. E
Atlanta, GA	116,700	85,710	86,860	39,670	1,815,000	254,000	309,500	41,420
Charlotte, NC	154,500	91,810	90,030	49,010	1,008,000	581,000	809,800	33,270
Chicago, IL	173,700	90,360	98,040	128,900	608,600	271,900	344,000	152,700
Denver, CO	43,530	95,540	102,700	308,400	644,800	277,600	402,800	60,520
Hartford, CT	101,500	64,990	67,790	36,410	390,800	173,400	228,200	64,100
Houston, TX	172,000	101,000	97,260	71,970	1,624,000	183,900	296,800	115,100
Jacksonville, FL	126,100	68,920	71,650	50,120	968,700	306,500	435,700	83,440
Los Angeles, CA	903,800	76,090	104,700	1,804,000	1,151,000	491,900	764,700	100,700
Phoenix, AZ	220,800	88,660	94,520	677,100	893,400	276,900	355,900	138,500
Seattle, WA	284,100	138,800	141,400	204,100	1,186,000	284,900	478,200	74,950
Washington, DC	268,400	125,200	117,500	74,470	564,300	158,900	235,600	128,900

Notes: This table reports mean income measures for individual landlords, separately by CBSA and incorporation status. All income measures are in thousands of 2019 dollars, drawn from 2019 IRS 1040 filings. For details on ownership classification, see Section 3.5.


8 Figures

Figure 1: Ownership identification cascade used to determine final owners of property ownership spells




Notes: This figure illustrates the order in which we apply matching rules to identify the final owner of each ownership spell. We begin by identifying government and unincorporated individual owners directly from the property record. For remaining spells, we apply a cascading fuzzy name matching procedure through the Business Register, Schedule K-1 filings, individual trust keywords, and REIT/10-K filings. Unmatched spells are categorized as likely businesses, likely unPIKed individuals, or fully unmatched. For details, see Section 3.2.

Figure 2: Share of all properties vs. rental properties by final owner type




Notes: This figure displays the final owner type identified by our methodology for properties surveyed in the 2019 American Community Survey in 11 sampled CBSAs. For details on the identification of final owners and classification of ownership types, see Sections 3.2 and 3.5.

Figure 3: Rental property ownership by building type




Notes: This figure breaks out rental property owner type by building type. The sample includes rental properties reported in the 2019 ACS for 11 selected CBSAs. For details on the matching procedure, see Section 3.2.

Figure 4: Owner type by CBSA



Notes: This figure shows the share of rental units by owner type, separately for each of the 11 CBSAs in 2019. Three categories (public, individual trust, and unmatched) are pooled to achieve a sample size large enough to disclose. For details on owner identification, see Section 3.5.

Figure 5: Potential under- and over-estimation of portfolio size using the mailing address approach



Notes: Panel (a) shows, among multi-property portfolios owned by business and incorporated individuals, the share of properties and units whose tax assessment is mailed to the property address—an underestimate source for the mailing-address approach. Panel (b) shows the share of mailing addresses (different from the property address) that receive assessments for at least 2 and at least 5 distinct final owners—an overestimate source. The sample is restricted to business- and incorporated individual-owned rental properties in 2019. For details, see Section 4.