An Alpha in Affordable Housing?*

Sven Damen Matthijs Korevaar Stijn Van Nieuwerburgh

July 17, 2025

Abstract

Residential properties with the lowest rent levels provide the highest investment returns to their owners. Using detailed rent, cost, and price data from the United States, Belgium, and The Netherlands, we show that this phenomenon holds across housing markets and time. If anything, low-rent units hedge business cycle risk. We also find no evidence for differential regulatory risk exposure. We document segmentation of investors, with large corporate landlords shying away from the low-tier segment possibly for reputational reasons. Financial constraints prevent renters from purchasing their property and medium-sized landlords from scaling up, sustaining excess risk-adjusted returns. Low-income tenants ultimately pay the price for this segmentation in the form of a high rent burden.

Keywords: affordable housing, rental market, risk and return in housing, market segmentation

JEL: G5, R2, R3

*Sven Damen is at the University of Antwerp (Sven.Damen@uantwerpen.be). Matthijs Korevaar (korevaar@ese.eur.nl) is at the Erasmus School of Economics, Erasmus University Rotterdam. Stijn van Nieuwerburgh is at Columbia Business School, Columbia University (sgv2110@columbia.edu). For this project, Korevaar has benefited from funding of the Netherlands Organisation for Scientific Research (grant number VI.Veni.231E.026). Van Nieuwerburgh thanks the Paul Milstein Center for funds to purchase the Koeter operating expense data. This paper includes own calculations based on non-public microdata from Statistics Netherlands (project 9817). We are grateful for excellent research assistance from Zi Wang, Dayou Xi, and Ziyi Qiu. We thank Boaz Abramson, Jiakai Chen (discussant), Leila Collins, Joost Driessen (discussant), Mark Francke (discussant), Pedro Gete (discussant), Michael Lacour-Little (discussant), Tim McQuade (discussant) and seminar and conference participants at the University of Cape Town, HEC Liège, University of Antwerp, the AREUEA National Meeting, the Affordable Housing conference in Madrid, the Bank of England conference on household finance and housing, the ERES national meeting, the AREUEA international conference, and the NBER Summer Institute real estate meeting for comments.

1 Introduction

Lack of affordable housing is one of the most pressing policy issues of our time. While homeowners may also face rising housing cost burdens on the back of rising house prices and mortgage rates, the problem is most pressing for renters who tend to be lower-income households. The lack of affordable housing is a global crisis. Renters made up anywhere from 15% of the population in Spain, 29% in Belgium, 36% in the Netherlands and the United States, to 60% in Switzerland in 2015. The median renter's share of gross household income that is spent on rent varies from 19% in Switzerland and Austria, 26% in the Netherlands, 29% in Belgium, 31% in the United States, and 32% in Spain. The share of severely rent-burdened households, those who spend more than 50% of their gross income on rent, is 29% in the US, 27% in Spain, 18% in Belgium, 9% in the Netherlands, and only 4% in Switzerland (Carliner and Marya, 2016). Moreover, rent burdens have been rising considerably over the past decade (JCHS, 2024).

While some countries have sizable public housing systems, policymakers in other countries provide incentives to private-market actors to increase the stock of affordable housing. However, there is limited data on the returns and risks that landlords face when investing in rental housing, and on the rents their tenants pay. Is the under-provision of rental housing to lower-income tenants due to low average returns earned by landlords in this segment of the housing market? Or maybe it is due to the high levels of risk? Are there frictions that stifle entry in the lower-rent segments of the rental market? Understanding risk and return in the various segments of the rental market is crucial to devise successful affordable housing policy.

This paper studies the cross-section of investment returns across different segments of the rental market. Our main finding is that housing units in the low-rent tier of the rental market consistently earn significantly higher returns compared to properties in the high-rent tier. Using new, large-scale, high-quality micro data from Belgium, The Netherlands, and the United States, we show that this return difference is present in different housing market settings and for different time periods.

We match the rental data to property sales transactions to compute the gross rental yield, the ratio of annual rental revenues to property value or the rent-to-price ratio. The annual gross rental yield is 2.1% higher for the first decile of the rent distribution than for the tenth decile in Belgium. In The Netherlands, the gross yield gap is 2.7% per year.

The net rental yield, which subtracts out costs associated with maintaining and managing the property, property taxes, and revenue loss from tenant turnover and tenant non-payment, is 0.9% higher in the first than in the tenth decile in Belgium and 1.1% in The Netherlands. In the United States, we directly observe the net yield on multifamily rental properties, absolving us from the need to measure the various expenses. The net yield gap in the US is 0.6% per year. So, while expenses are higher for low-rent properties due to differences in both property and tenant characteristics, they do not eliminate the

gap in rental yields.

The capital gain yield, measured as the average annual price appreciation over a five-year holding period computed from a rich hedonic model, is also substantially higher for low-rent than for high-rent properties. The capital gain yield differential between the first and tenth deciles of the rent distribution is 0.84% per year in Belgium, 3.0% in the Netherlands, and 3.25% in the United States.

Combining net yields with capital gain yields, we arrive at the total return on rental properties. Annual returns are 1.74% higher in Belgium, 4.08% in the Netherlands, and 3.86% in the United States for low-rent than high-rent properties. Investors in low-end properties earn significantly higher returns.

Next, we explore what factors can explain these cross-sectional differences in returns across properties, exploiting the relative strengths of our American, Belgian, and Dutch datasets. The natural question is whether the return differential compensates investors for differential risk exposure. We explore a myriad of risk-based explanations.

The most traditional one is systematic risk, measured as covariance with the business cycle. Using data on the evolution of properties' net cash flows, we find that low-rent properties have a *lower* beta than high-rent properties. Low-rent properties' net cash flows fall by less, or even rise, in recessions. Intuitively, having an affordable rental becomes more valuable in bad times, when tenants' income may suffer adverse shocks and some tenants may be looking to move from higher to lower tiers of the rental market. Since tenants need a minimum quantity and quality of housing, affordable housing is an inferior good. This counter-cyclical property of rent makes low-tier properties a hedge from investors' perspective. Hedges should earn average returns that are low, not high.

A second source of systematic risk is regulatory risk. Specifically, regulations that affect tenant protections and landlords' ability to increase rents could disproportionately affect the lower segments of the rental market, and make investing in those segments riskier. We take advantage of substantial regulatory and political heterogeneity across US States to explore this possibility. We use GenAI tools to construct a tenant protection index from state-level laws. We also look at state-level measures of economic policy uncertainty and at states that are politically dominated by Democratic governments. We find no evidence that states with stronger tenant protections, more policy uncertainty, or a Democratic government have higher returns or more steeply downward-sloping return profiles. Evidence from The Netherlands that exploits geographic variation in how binding rent control regulations are for private rentals arrives at the same conclusion.

Next, we turn to idiosyncratic risk. Using ownership data, we show that many rental investors are under-diversified in The Netherlands and Belgium. Given this under-diversification, idiosyncratic risk may be compensated in equilibrium. The argument requires that low-rent properties display more idiosyncratic risk than high-rent properties. Studying detailed unit-level and tenant characteristics for Belgium and the Netherlands reveals a mixed picture. Low-rent properties tend to be older, have younger tenants, with

lower income, and slightly higher turnover and non-payment rates. These characteristics could be associated with higher idiosyncratic cash flow risk at the unit level. On the other hand, left-tail tenant income risk is lower in low-rent properties in the Dutch data.

Using US data, we can study idiosyncratic risk for multifamily properties, which partly diversify the idiosyncratic risk present at the unit level. We do not find evidence for higher mortgage default rates for low-rent properties. Low-rent properties have only modestly higher cash-flow growth volatility over reasonable holding periods of 3-5 years. The Belgian and Dutch ownership data show that the typical investors in low-rent units own multiple properties, providing further gains from diversification. The conclusion that the excess return earned by landlords of low-rent units solely reflects compensation for residual idiosyncratic risk is implausible.

We are left with the most plausible interpretation of the empirical results, namely that investors in low-rent housing earn positive risk-adjusted returns; that there is a positive alpha in affordable housing. This interpretation begs the question: If low-tier rentals earn excess returns, why do these returns not get competed away? What are the frictions, the limits to arbitrage that prevent more capital from flowing into the low-tier segment of the rental market? We offer a three-pronged answer to this question.

The first potential source of additional capital is renters. If rents are excessively high in the low-tier segment, why do the renters not purchase their property to take advantage of the mispricing? The evidence shows that the renters in these properties are on average younger, lower-income, lower-wealth households. They are financially constrained and less likely to move to owner-occupied property. Also, many low-rent units are in apartment buildings which are not easily bought by (collections) of households. In the data, most of these units are sold as part of a portfolio.

The second potential source of additional capital is very large, well-diversified, institutional landlords. These investors have access to outside capital on public or private equity markets. We show that very large incorporated landlords do not invest in low-rent properties. Rather, their portfolios are concentrated in the high-rent tiers of the rental market where net yields and returns are the lowest. We also show that this type of landlord is the only type that is geographically diversified. What prevents them from reallocating capital to the lower tiers of the markets where they are already active? We hypothesize that they face reputational risk when going down-market. Their managers and capital backers, often public pension funds or public equity investors considerate of their corporate social responsibility, do not fancy the risk of being stigmatized as "slumlords." We also provide evidence that the very largest landlords may run into diseconomies of scale in the management of low-income housing.

The third potential source of additional capital comes from the investors already active in the low-rent segment. We show that these investors are typically not incorporated. They have medium-sized portfolios which display a strong local bias. These investors lack access to external equity capital. They must finance the downpayment (the equity com-

ponent) for the next unit from the retained earnings on the existing portfolio. Financing frictions prevent them from scaling up more rapidly. Additionally, the informational or behavioral frictions that underly their local bias, prevent investors like them in other regions from reallocating capital spatially.

This constellation of frictions results in an equilibrium where unincorporated, medium-sized investors active in areas with a preponderance of low-income rentals persistently earn high risk-adjusted returns. We present some evidence from capital flows that, if and when additional capital gets invested in low-rent properties, the return-rent slope flattens.

Related Literature Our paper contributes to a growing stream of the housing literature that measures the returns to housing investments. This literature increasingly explores the heterogeneity in returns across geographies and demographic groups, e.g., Bracke (2015); Desmond and Wilmers (2019); Chambers et al. (2021); Amaral et al. (2021); Eichholtz et al. (2021); Demers and Eisfeldt (2022); Goldsmith-Pinkham and Shue (2023); Halket et al. (2023); Kermani and Wong (2021); Diamond and Diamond (2024); Cocco et al. (2024); Colonnello et al. (2024); Samuels (2024); Nicholas and Spaenjers (2025). Our contribution to this literature is twofold.

First, we make significant progress on measurement. A key challenge in this literature is that large-scale and representative matched data on rents, costs, and prices is not available, complicating inference on property-level returns. Most papers in this literature work either with smaller and geographically narrow samples where rents and prices are available for the same properties, or apply imputations or matching algorithms to larger cross-sections to infer rents and sale prices. Rental data often comes from listings rather than actual contracts. Typically, data on costs is lacking entirely, based on self-reported values, or imputed using simple estimates.[1] In this paper, we make use of new data sources that allow us to observe both rents and property values administratively for the entire population (Belgium) or for a sizable part of the rental market (US and The Netherlands). We observe rental cash flows either after adjusting for all costs (US) or we use granular estimates from appraisers that we link to administrative data to infer how expected costs differ across the cross-section of properties (Belgium and The Netherlands). Taken together, this allows us to build reliable estimates of rates of return across different segments of the rental market.

Our second main contribution is an in-depth investigation of the factors that drive cross-sectional differences in rental housing returns. The existing literature has focused on establishing return differences across specific geographies or demographic groups, as mentioned above. Our starting point is that there is significant heterogeneity in returns across properties which we correlate with characteristics of the properties, their tenants, and their landlords. To draw implications from these correlations, we need to understand

[1] One exception is Chambers et al. (2021).

the origins of that heterogeneity. We critically examine risk-based explanations of differences in returns and present evidence to support a segmentation hypothesis along the quality spectrum that is sustained by limits to arbitrage. This explanation points to the importance of landlord heterogeneity. For a complementary perspective on the role of investors in real estate markets, see recent work by Badrinza et al. (2022); Garriga et al. (2023); Hanson (2024); Coven (2025); Koijen et al. (2025); Moktan (2025); Cornaggia et al. (2025).

Our work connects to a larger literature on housing affordability and insecurity (Desmond, 2012; Desmond and Gershenson, 2017; Fowler et al., 2015; Collinson et al., 2024b) and policies aimed to address it, for example rent control (Glaeser and Luttmer, 2003; Autor et al., 2014; Diamond et al., 2019; Favilukis et al., 2023), zoning (Glaeser and Gyourko, 2003), tax credits for developers (Baum-Snow and Marion, 2009; Diamond and McQuade, 2019), rental assistance (Kling et al., 2005; Collinson and Ganong, 2018; Abramson, 2023; Collinson et al., 2024a), and rent guarantee insurance (Abramson and Van Nieuwerburgh, 2024; Bezy et al., 2024). Our contribution to this literature is to point to a complementary set of policies aimed at encouraging more investor entry into the low-rent segment of the housing market. This could take the form of an information campaign that touts higher returns or corrects potential misconceptions around risks, an investment tax credit or mortgage subsidy for investors in low-rent properties, or a subsidy that helps renters to purchase their home.

Finally, our findings have implications for the measurement of housing consumption, a large component of the consumption basket and a key input in the measurement of inflation (e.g., Ambrose et al., 2023; Abramson et al., 2024). A large share of housing consumption consists of imputed rents for owner-occupiers. These imputations typically apply a constant rent-price ratio to owner-occupied units. Assuming that owner-occupied housing units are more similar to the higher-rent segment of the rent distribution, applying the average rent-price ratio from rentals would lead one to overstate housing consumption for owner-occupiers and in the aggregate. These imputations also fail to remove expense categories such as property management, and vacancy and credit losses, which increase rents but are not relevant for owners. This leads to an overstatement of housing consumption that is more pronounced for low-income renters.

The rest of this paper is organized as follows. In Section 2 we describe our data and the rental market context. Section 3 lays out our measurement approach. Section 4 reports the main stylized facts on gross and net yields and total returns across the rent spectrum. Section 5 investigates a series of risk-based explanations and provides evidence for the limits to arbitrage argument. Section 6 concludes and discusses the policy implications.

2 Data and context

This paper utilizes detailed data on housing rents, costs, and prices from the United States (US), Belgium (BE), and The Netherlands (NL). Before discussing the data from each country, we provide an overview of the benefits of a cross-country approach and the main differences across datasets and housing market contexts.

First, the nature and ownership contexts of the rental markets in these countries differ significantly, enabling us to investigate whether differences in returns across property types are specific to the market context or hold more generally. In the US, our analysis is based on data derived from mortgage loans on multifamily rental properties. These properties are predominantly owned by institutional rather than retail investors. While there is interesting variation in regulation within the US which we will exploit below, rent regulations and tenant protections are *on average* less stringent in the US than in Europe.

Belgium and The Netherlands, despite being geographically close, have markedly different rental markets. BE has a relatively unregulated rental market that is dominated by retail investors. In contrast, the NL private rental market is smaller, as a share of the overall rental market, and includes both retail and institutional investors. Most rental housing in NL is provided by non-profit housing associations, offering public housing at low rents. In addition, private market rentals are subject to rent regulations. The latter were relaxed significantly in the 2010s but tightened again in 2024.

The second advantage of a cross-country approach is the ability to leverage the unique strengths of each dataset to explore the mechanisms behind our findings. We cannot investigate all candidate explanations for any one of our countries. Table 1 summarizes the characteristics of the main samples we use.

Table 1: Overview of Data Across Countries

	US	**BE**	**NL**
Sample Characteristics			
Coverage	Multi-family loans	All private rentals	Most private rentals
Level	Building	Unit	Unit
Time period	2001-2024	2007-2022	2018-2023
Property characteristics	Limited	Extensive	Average
Tenant characteristics	None	Limited	Extensive
Investor characteristics	None	Limited	Extensive
Return Components			
Property value	Sale price or appraisal	Sale price	Tax value and sale price
Measured at	*Loan origination*	*Moment of sale*	*January 1 or moment of sale*
Rent	Underwritten NOI	Contractual rent	Contractual rent and list price
Measured at	*Loan origination*	*Moment of signing*	*January 1 or moment of listing*
Cost estimate	n/a	Appraisal (50 types)	Appraisal (150 types)

In the US data, we observe net operating income per unit, which measures revenues

minus costs. This absolves us from having to separately measure costs, but also prevents us from computing gross yields. The US data also allows us to exploit state-level variation in regulation and its longer coverage period can provide insights into how returns vary with the business cycle.

The BE rent data comes from a unique administrative dataset covering all rental contracts in the country over a long period of time.[2] Rental contracts are registered with the government; there is a universal registration requirement. In many other datasets, there is limited coverage of properties in the lower-rent tiers of the rental market, which are key focus of this paper, since low-quality rentals are often allocated in informal markets.

The NL data has the virtue of enabling a direct link between rents and prices to the characteristics of property owners and tenants, obtained from administrative tax records. The tenant data allows us to connect rentals to the characteristics of their tenants. The landlord data allows for a detailed analysis of ownership patterns of sorting across rental market segments and geographies.

Costs are not directly observed in the BE and NL data, but we obtain high quality maintenance cost data for a wide variety of property types from a company specialized in such cost estimates, a trusted source of data to the NL government and business alike.

2.1 United States

For the US, we focus on data from the multifamily rental market. Fannie Mae and Freddie Mac dominate this market, accounting for approximately half of the total stock of multifamily loans. Our dataset includes loan-level information covering 123,623 loans from January 2001 to May 2024. We supplement this with 60,100 multifamily loans originated and sold into private-label commercial mortgage-backed securities (CMBS) and commercial real estate collateralized loan obligations (CRE CLOs), compiled by the data aggregator CredIQ. Combined, this dataset covers the majority of multifamily loans. Multifamily properties constitute the majority of rental units in the US (63.5%) and most multifamily housing units are rentals (89%). Multi-family rental properties are nearly universally financed with a mortgage.

Key variables include the loan-to-value (LTV) ratio and the debt service coverage ratio (DSCR), both measured at loan origination. The LTV allows us to infer property value as $V = L/LTV$, where L is the loan balance. The DSCR, defined as net operating income (NOI) divided by debt service, provides a measure of asset performance. NOI, calculated as DSCR times debt service, reflects rent revenues minus operating expenses. Both prices and NOI are expressed per unit and adjusted to 2023 real dollars. Outliers are removed. Further details are provided in Appendix B.

While property values are observed at loan origination only, we also obtain dynamic loan performance data, which include a time series for actual realized annual net cash

[2]To the best of our knowledge, the only other countries with a rent registry are Turkey (but only as of July 2023), Italy, Malta, and the United Arab Emirates.

flow of the property. We will make extensive use of this variable.

2.2 Belgium

For BE, we combine several administrative data sources. As rental contracts must be registered by law, we have access to the universe of rental contracts provided by the Federal Public Service Finance (FPS Finance). Failure to register the rental contract gives tenants the right to unilaterally cancel the contract without penalty. The data includes rental contracts from 2007 to 2022. Contracts are observed at origination. The data set includes the start date, lease length, monthly rent, a housing unit identifier, and anonymized landlord and tenant identifiers.

We merge the rental contracts with the universe of housing transactions from 2007 to 2022 provided by FPS Finance. The transaction data include the sales price, date of sale, detailed property characteristics, the exact address and a unique anonymous identifier of the buyer and seller. We also merge the rental contracts with the universe of energy performance certificates (EPC database), which provides additional detail about the energy efficiency.

To estimate maintenance costs, we make use of the services of a maintenance cost appraiser, Koeter Vastgoed Adviseurs (henceforth, Koeter). At our request, Koeter computed the maintenance costs for the 50 most common categories of rental properties in BE, as defined by us in terms of property age, property type, and energy efficiency. For every property type, Koeter then computed the average annual maintenance expenses necessary to keep the property at constant quality for at least 99 years. Details on the BE data are available in Appendix D.

2.3 The Netherlands

For NL, we use data from Statistics Netherlands, primarily the Woonbase database. The data covers the period 2018–2023. This database includes property characteristics, tenant demographics, and housing costs of the entire NL population as of January 1 each year. Property values are based on tax values. The latter are estimated based on transaction prices of similar properties in a window of twelve months before and after the reference date.[3]

Rental prices are administratively observed for housing association properties, representing two-thirds of rental properties. For private market rentals, rents are derived by Statistics Netherlands from various sources, including fiscal data, listing platforms, and housing surveys, with direct rent observations available for 50% of properties. We restrict the sample to properties whose residents moved in on or after January 1, 2016, and live

[3]By law, the tax value of a property needs to equal the expected market price of a property on the reference date (if sold free of use and without any liens). While individual municipalities have some discretion over the exact model they use to determine these values, an independent authority (Waarderingskamer) enforces that tax values in every municipality accurately capture market values.

in one-household units. We do so because rent estimates are less reliable for long-term tenants, who often benefit from lower, regulated rents, and are unavailable when multiple households share a housing unit.

Since the Woonbase database only cover a short time period and does not identify realized sales prices or rents for new contracts, we augment our analysis with listings data for 2007 to 2023 from the listing platform with the largest coverage in the Netherlands, as well as with administrative data on all housing transactions from 2009 to 2023.

We link the Woonbase data to ownership and landlord portfolio information, differentiating between retail (personal) and corporate investors (businesses) as well as the sizes of their portfolios. We also link to registry data on residential mobility and on individuals with problematic debts (arrears).

Koeter estimates maintenance costs for 147 NL property categories, using the same method as for BE. Details on the NL data are available in Appendix C.

3 Methodology

We aim to study differences in expected returns across the distribution of rental values. To infer this relationship from the data, we take the following approach.

We start from the assumption that the observed rent (R) and sales price (P) of a property i can be modeled as the sum of a vector of relevant property characteristics $\mathbf{X}$ (e.g., location, size, age), the time-varying "prices" associated with these characteristics in the sales market (β) and rental market (Γ), and residuals (ε and u):

$$\log P_t^i = \beta_t' \mathbf{X}_t^i + \varepsilon_t^i \tag{1}$$

$$\log R_t^i = \Gamma_t' \mathbf{X}_t^i + u_t^i \tag{2}$$

We refer to $\exp(\Gamma_t' \mathbf{X}_t^i)$ as the fundamental rental value, or *rental value* for short. Based on observed rents and prices, we infer how yields and total returns differ across the distribution of rental values.

To understand what causes property-level returns to differ, we can take a log-linear approximation to the one-period holding period return to property i, ignoring costs for the time being:

$$\text{Ret}_{t,t+1}^i \approx \left[(\Gamma_{t+1} - \beta_{t+1})' \mathbf{X}_{t+1}^i + (u_{t+1}^i - \varepsilon_{t+1}^i)\right] + \left[(\beta_{t+1}' \mathbf{X}_{t+1}^i - \beta_t' \mathbf{X}_t^i) + (\varepsilon_{t+1}^i - \varepsilon_t^i)\right] \tag{3}$$

The first bracketed term is the rental yield, the second term the capital gain yield.

Differences in returns across properties arise when characteristics are priced differently in the rental than in the sales market ($\Gamma_{t+1} \neq \beta_{t+1}$). For instance, old low-quality properties may have relatively high rents compared to sales prices, while new luxury properties in prime locations may command a larger premium in the sales market. Returns also vary

due to changes in the pricing of characteristics over time ($\beta_{t+1} \neq \beta_t$) or due to changes in the characteristics themselves ($\mathbf{X}^i_{t+1} \neq \mathbf{X}^i_t$), for example due to aging or renovation. Finally, the idiosyncratic components of rents (u^i) and prices (ε^i) also influence property-level returns. For example, if a landlord overpays for a particular property ($\varepsilon^i_t > 0$), the property is more likely to have a low realized rental yield and capital gain yield relative to other properties with the same fundamental value. If a luxury property is rented for an unusually low observed rent ($u^i < 0$), it has a lower yield and return than it should given its characteristics.

This framework points to two challenges when measuring returns across the distribution of rental values. First, prices, rents, and all relevant costs are not observed for every property at every point in time. As a result, estimating the return relies on extrapolation or estimation to an extent that varies across data sets.

Second, observed prices incorporate both fundamental value and residual components per equations (1) and (2). Although our interest is in the fundamental component, disentangling the two is empirically challenging, especially when property characteristics in the data are limited. In the latter case, residuals may reflect unobserved characteristics rather than idiosyncratic noise that is irrelevant to prices.

In summary, we face a trade-off: either we use observed prices which include idiosyncratic noise, or we use a hedonic model to estimate rental and property values, which may omit relevant property characteristics. The optimal approach depends on the nature and strength of the available datasets, which varies between our three countries.

3.1 Rents and Rental Yields

A key challenge is to combine observed rents and sales prices, potentially measured at different points in time and with idiosyncratic noise, to infer rental yields and returns, and how they vary across the distribution of rental values.

3.1.1 United States

In the US sample, the net rental yield is the reported net operating income (NOI) of a property relative to the property value, as inferred from the mortgage data at the time of loan underwriting. We examine the distribution of net yields and total returns across the distribution of reported NOI per housing unit (apartment), grouping properties into ten deciles.

Unlike in the BE and NL samples, we opt to group observations based on the reported NOI instead of an estimated rental value. This choice is motivated by two considerations. First, we observe only a limited number of property characteristics in the US data (e.g., age, number of units, and renovations). Combined with the smaller number of loans per geography, this makes it difficult to capture sufficient variation in rental values using a hedonic model. Second, idiosyncratic variation in net rents (NOI) is likely more limited in

the US sample, reducing potential bias. Measurements of underwritten NOI and property value are based on assessments of professionals, either from actual transaction prices or appraisals. These valuations are likely to vary less than the observed prices in markets where retail investors dominate, such as BE. Since these valuations pertain to properties with multiple (sometimes hundreds of) units, rent and price noise are smaller than when rents and prices are measured at the unit level, as they are in BE and NL.

The drawback of this approach is that, if the observed NOI includes a mean-zero noise component u_t^i orthogonal to rental value, then there will be a downwardly biased slope estimate (decile 1 minus decile 10) of the relationship between rental value and rental yield or return. Sorting on rents (NOI) instead of rental values makes our estimates conservative. We have confirmed that this is the case in the US data.

3.1.2 Belgium

To study the relationship between yields and rental values in the BE data, we group properties in deciles based on the predicted rental value of a property, while using observed rents and prices to estimate average returns in each of these groups. We model the log monthly rent as a function of a rich set of characteristics as in (2). For apartments, the hedonic model includes the variables: useful area, the number of floors, the number of habitable rooms, number of bathrooms, number of garages, the presence of central heating, age, years since renovation, energy efficiency score (EPC), neighborhood fixed effects, and year fixed effects.[4] For single-family homes, we add: land area, built area, the presence of a habitable attic, and construction type (detached, semi-detached, terraced).

For all properties, the gross yield is defined as the ratio of the observed rent to the market value of the property. To calculate the market value of a property we use the sales price that is closest in time to the date of the rental contract and adjust the sales price to the start year of the rental contract using a hedonic model for sales prices. The latter contains the same hedonics as the rental model, and adds an indicator variable for whether the property is renter-occupied.

3.1.3 The Netherlands

In the NL data, rent pertains to the monthly rent for individual housing units owned by retail and corporate investors on January 1 of each year. This is slightly different from the BE data where we observe the rental contracts at the time of execution rather than on a fixed date.

Similar to BE, we group properties into deciles based on the predicted rental value of a property. To determine the predicted rental value, the hedonic model for the log monthly rent includes the following characteristics: number of square meters, construction vintage,

[4]We require at least 30 observations to estimate the neighborhood fixed effect. For neighborhoods with fewer observations we estimate municipality (ZIP code) fixed effects instead.

building vintage, construction type, number of units and neighborhood fixed effects. The decile classification is based on predicted rental values for all private rental properties, including properties for which we do not observe the rent. However, we use observed rents to compute the average yield and total return within a decile.

The gross yield is defined as the ratio of the observed rent to the market value of the property. Our main return measure is based on the set of properties for which we have both the rent of a recent mover and the tax value as a measure of the price (and capital gain). As a robustness check, we compute yields based on properties purchased and listed for rent in the same year and annualized real capital gains based on transaction prices.

3.2 Costs

Unlike in the US data, we do not directly observe net operating income, measured as revenues minus operating costs, in the BE and NL datasets. We need to infer expected costs using other sources.

3.2.1 Maintenance Cost

We use bottom-up maintenance cost estimates provided by Koeter for 147 common NL housing types and for 49 common BE housing types. Their cost estimates represent the annual expenses an owner must incur to maintain the property at constant quality, i.e., to offset depreciation. To extrapolate these costs to the entire housing stock, we regress the maintenance costs of these common unit types on the unit characteristics. Since the available characteristics differ, we use a slightly different model for BE and NL. The results are presented in Appendix Table A1. A simple model based on area (in square meters), property type, and building age explains over 75% of the variation in observed maintenance costs across the 147 NL housing types and close to 70% of the variation across the 49 BE housing types. We proceed with the model reported in Columns 3 and 6, respectively. We use the corresponding regression coefficients to predict maintenance costs for all properties in the NL and BE samples. Koeter also provides index coefficients to link maintenance costs across years, which we apply to link up the timing of the rent, price, and cost observations.

Our BE and NL cost data assume that properties are maintained to keep quality constant. For properties with low rental values, the maintenance expense required to offset depreciation may not be financially feasible or desirable to the landlord. This would result in us underestimating net rental yields for low-rent properties, and hence the slope between low- and high-rent units, rendering our conclusions conservative. Excess depreciation would also empirically manifest itself as a lower capital gain yield, which we would overestimate. The total return, the sum of the net yield and the capital gain yield, should not be affected.

3.2.2 Other Costs

In addition to maintenance, property investors face various other costs. These include property taxes, turnover costs, cost of non-payment of rent, and property management costs.

Taxes and insurance: Annual property taxes and other levies to be paid by the owner are observed administratively in both the NL and BE data. In NL, we include the cost of mandatory building insurance. This cost is born by the tenant in BE. We express them as a fraction of property value, i.e., as a tax rate.

Turnover costs: When a tenant leaves, it may take time to find a new tenant. In the interim, the property does not generate rental income. Identifying a new tenant may also require paying a broker fee. Turnover events are identified administratively using newly signed rental contracts in the BE data and household mobility records in the NL data, which we link to rental listings to estimate vacancy duration (see Appendix C). For each turnover, we assume costs equal to the estimated number of vacant days between successive tenants' contract end and start dates times the daily rent.

Tenant default costs: Landlords may also face credit losses from non-payment of rent due to tenant delinquency. We collect information on the stock of rent arrears and irrecoverable rent amounts in the NL public housing sector, which we apply to the NL private rental sector.[5] In 2017, housing associations reported total rent arrears of 1.2% of rental income, of which they deemed 0.3% irrecoverable. We use the 0.3% of irrecoverable rent arrears and add a back rent recovery cost of 0.6% of total private rental revenue.[6] We distribute these costs equally over all tenants with registered problematic debts. While rent default itself is not observed, we can administratively observe non-payment of tax bills, health insurance premia or student loans. Most tenants in rent arrears have such debts. Appendix C.4 provides the details.

We assume the same 0.9% of credit losses (arrears plus recovery) for BE in the absence of more specific data. We use survey data from the 2008-2022 waves of the EU-SILC Survey in Belgium on the share of households that missed two consecutive rent payments to allocate the credit losses across rent deciles.

Management costs: We set management costs following industry standards in both NL and BE. External rental agencies that manage properties on behalf of landlords typically either charge a percentage of rent or a flat fee, both of which decrease in the

[5]Since the public housing sector covers the lowest tiers of the rental market in terms of tenant income and income stability, this assumption renders our results conservative.

[6]This assumes that landlords incur costs equal to 50% of the stock of rent arrears to recover rent or evict tenants. Lacking observable data on these expenses, we adopt this conservative assumption to avoid understating the cost of tenant default.

number of units managed. In our baseline model for NL, we incorporate a management fee of 4% to 6% of rent per unit subject to a minimum fee of 32-42 euros per month (2023 prices). We discuss sources and assumptions in Appendix C. In BE, the usual management fee is 7% + VAT.

3.3 Capital Gains

Estimating capital gains for housing poses challenges because these gains are generally unobserved, and the subset of (rental) properties for which both purchase and sale prices are observed is both small and likely non-random.

To measure capital gains in the US and BE, we primarily rely on a hedonic methodology, estimating versions of (1) from realized transaction prices and observed property characteristics. Starting from the market price at time t, the start date of the rental contract, the hedonic model is used to estimate the expected value of the property at time $t+T$. Capital gains may differ across properties due to changes in the prices assigned to characteristics or due to changes in the characteristics themselves.

In the US, the hedonic model for the log property price per unit includes property age, age squared, NOI, and zip-code average household income.[7] These covariates are time-varying, allowing for heterogeneous capital gains across properties over time. Other variables include the number of units in the property, a renovation dummy, deal type, and state and time fixed effects. These covariates allow for heterogeneity in log price levels, but not in returns, which are log price changes.

In BE, we use a separate hedonic model for apartment and single-family units. The covariates in the model for apartments are: the usable area, the number of floors, the number of habitable rooms, number of bathrooms, number of garages, a central heating dummy, age, and age squared. The model for single-family units adds: land area, built area, the presence of a habitable attic, and construction type dummies (detached, semi-detached, terraced). Both models include neighborhood-by-year fixed effects for neighborhood-years with at least 30 sales and municipality-by-year fixed effects otherwise.

For both the US and BE, we need to make assumptions about the nature of the observed price residuals at the end of the holding period over which we wish to compute the capital gain yield, $\varepsilon_{i,t+T}$, since we do not observe a price at that time. We assume that $\varepsilon_{i,t+T} = A\varepsilon_{i,t}$, with $A \in [0,1]$. This results in a residual term in the capital gain yield of $(A-1)\varepsilon_{i,t}$. In the special case of $A = 1$, the residual in the initial transaction $\varepsilon_{i,t}$ solely reflects omitted characteristics. The omitted characteristics affect both the log price at t and $t+T$ and, assuming the valuation of the omitted characteristics did not change, drop out of the capital gain yield. In the special case of $A = 0$, $\varepsilon_{i,t}$ solely reflects idiosyncratic noise, reflecting over- or underpaying in the initial purchase. This noise is

[7] The property price is the one implied by the loan to value ratio and the loan balance at the time of origination. Some of these prices arise from transactions while others from a professional appraisal.

irrelevant for the value at time $t+T$, and hence the initial under- or over-payment fully passes through into the capital gain yield. We calibrate the value for A such that the resulting distribution of capital gain yields from the hedonic model aligns as closely as possible with the distribution of capital gain yields in the subsample of repeat sales. This results in a value of A=0.5 in the US data and A=0.25 in the BE data. Since we have more covariates for the BE data and hence fewer omitted covariates, a smaller value for A in BE is intuitive. Our results are robust to other choices for A. Appendices B.3 for the US and D.3 for BE show that hedonic and repeat-sales capital gains deliver very similar results.

In the NL data, our main measure is based on the growth of the tax value during the year which is available for every property. Since these valuations assume that the property's quality is held constant, they align with our assumption that the landlord makes the required maintenance expenses to keep up the property's quality. As a robustness check, we use realized capital gains based on repeat-sales and capital gains from a hedonic model that vary by unit size, building age, unit type, and region. The capital gain yield-rent slope is very similar to the baseline estimate in both cases.

4 Stylized Facts on Rents and Returns

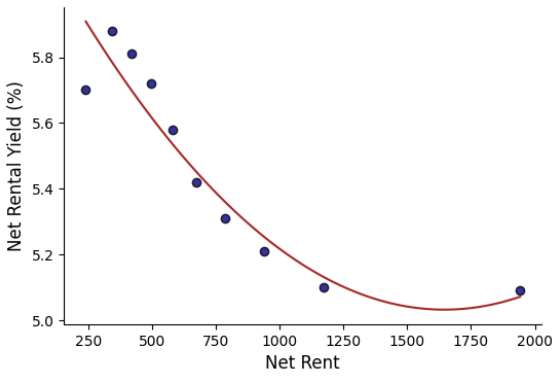
4.1 Main Fact

Applying the methodology from the prior section, we are now ready to present the main stylized facts on net rental yields and total returns across the distribution of rental value for the US, BE, and NL. Figure 1 presents the paper's main empirical finding. The left panels plot the net yield, the right panels the total return. Each dot represents the middle of a decile of the distribution of rental value, where decile 1 (D1) pertains to the 10% most affordable properties and decile 10 (D10) to the 10% highest rental-value properties.

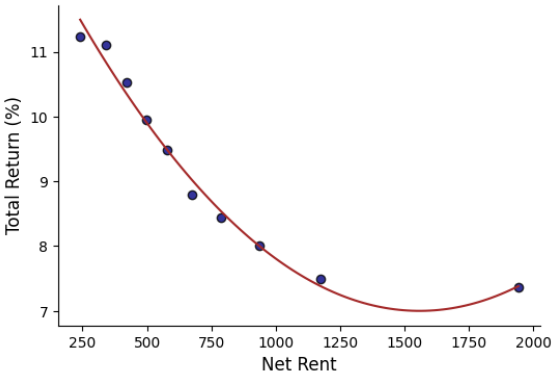
In all three countries, we observe a clearly downward-sloping relationship between rental values and both net yields and total returns. In terms of net yields, low-rent properties in D1 have yields that are between 0.61% and 1.10% points higher per year compared to high-rent properties in D10. This pattern is quite similar across countries, given different housing market settings and time periods.

In each country, we also find evidence for substantial differences in average capital gain returns between low- and high-rent units in the same direction as net yield differences. D1 units experience annual house price appreciation that is 84 bps higher than D10 units in BE, 298 bps higher in NL, and 325 bps higher in the US. The pattern is monotone in BE and the US, but in NL capital gain yields are highest in D2. In Appendix C.6 we find a monotone relationship using alternative methods based on repeat-sales and hedonic cpaital gains. Much of the decline in NL is concentrated from D9 to D10. Still,

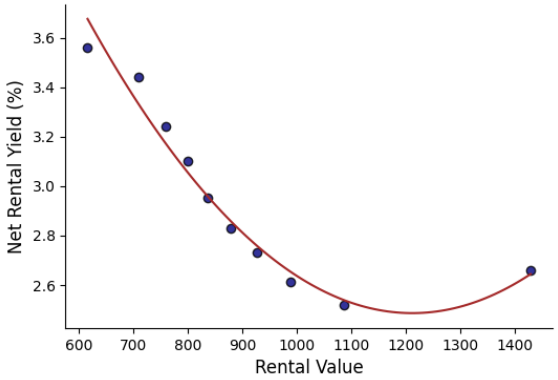
Figure 1: Net Yield and Total Return by Rent Value Decile



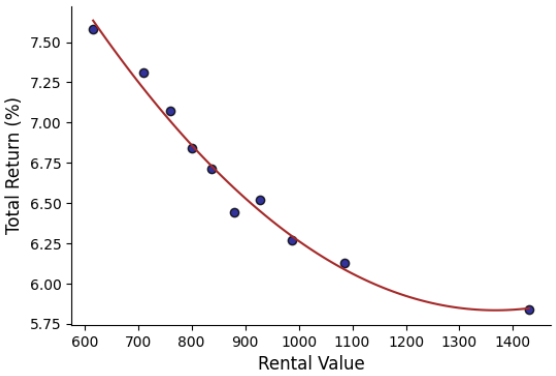
(a) Net Rental Yield US



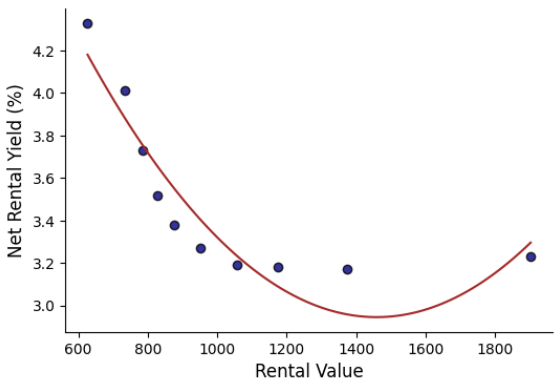
(b) Return US



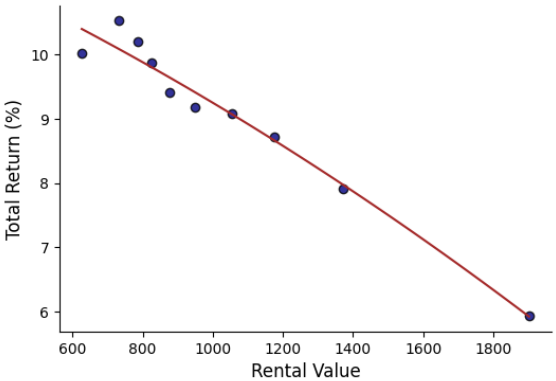
(c) Net Rental Yield BE



(d) Return BE



(e) Net Rental Yield NL



(f) Return NL

Notes: Each panel plots the relation between the net rental yield (left panels) and total return (right panels) and the net rent or rental value of a property. Every point corresponds to the average return in a rent decile, order from D1 (left) to D10 (right). The top panel is based on net rents for the US, the bottom two panels for rental value in BE and NL.

the D2-D9 capital gain yield difference is 176 bps.[8]

When adding capital gain yields to net rental yields to construct annual total returns, the difference between low- and high-rent properties increase to 174 bps in BE, 408 bps in NL, and 386 bps in the US.

In the US data, we find that net yields are slightly higher in D2 than in D1. This could reflect an idiosyncratic component in rents (prices) that is more negative (positive) on average for properties in D1 than in D2. Consistent with the presence of omitted characteristics, we find that the relationship between rents and returns becomes more linear as we increase the parameter A.[9]

4.2 Accounting for Cost Differences

Before delving into possible explanations, we use our BE and NL data to decompose the net yield into a gross yield and a cost yield. Table 2 shows the gross yield, monthly rent level, cost components and capital gains for every decile of the rental value distribution. All nominal amounts are deflated by the consumer price index and expressed in 2023 euros or 2023 dollars. We recall that in the US data, we only observe the net yield and cannot do this decomposition. In Appendix Table A2 we provide rents, property values and cost components expressed in euros. Gross yields are downward sloping in predicted rent levels with a D1–D10 difference of 2.13% in the BE data and 2.66% in the NL data.

This gap shrinks considerably after we account for the various expenditures, indicating that expenditures represent a larger share of property value for low- than for high-tier rentals. Low-tier rentals tend to be older and tend to be located in lower property-value locations. Both factors increase maintenance costs relative to property value. The maintenance cost differential eliminates 0.78% and 0.96% of the gross yield gap in BE and NL, respectively. Costs fall monotonically with rent values.

Other cost factors play a more modest role. Property taxes are a larger burden for low-tier properties. In BE, this is not surprising since property taxes are levied based on the "cadastral income" of the unit which reflects rental rather than property value. In NL, property taxes are levied on property value, but municipalities with more high-end

[8] As an aside, average capital gain yield levels are the highest in the NL due to the sample period. Data from Eurostat confirm that NL house prices indeed grew rapidly from 2018 until 2023. The same Eurostat data confirms that BE house price appreciation is also strong during this period, but somewhat lower than in NL. Similarly, house price growth was exceptionally high in the US during this period. This period includes the Covid-19 pandemic house price boom which triggered an aggregate demand increase for housing due to work from home and health concerns Mondragon and Wieland (2022); Gupta et al. (2022). This also helps explain the sharply lower capital gains in NL in D10 relative to D9, since properties in Amsterdam, the largest city in NL with a dense urban core, are overrepresented in this decile. In Appendix Figure C2 we provide longer-term evidence in the NL based on listings data (2008-2023), which lead to a smaller gap between D9 and D10 while retaining a large return premium.

[9] Our analysis focuses on unlevered returns. However, investors may employ leverage, and there may be differences in the cost of debt across the rent distribution. Appendix B.5 shows that the cost of mortgage debt in the US between the bottom quintile and the top quintile of the rent distribution differs by only 9 basis points once standard underwriting variables are controlled for.

properties are able to levy lower tax rates, while still raising the required tax revenue. Moreover, taxes such as for water and sewer systems and cost of property insurance, which are included in property taxes in NL, are directly proportional to property value.

Turnover costs are somewhat higher in the low-tier segment, but overall this contributes less than 10 bps to the difference in yields in gross returns. In general, costs for turnover are higher in the NL sample because turnover events are more frequent. This difference might be related to both housing market factors and the measurement approaches. Relative to BE, there is a large social rental sector (2/3 of rentals) in NL offering lower rent levels than in the private sector. In the owner-occupied sector, there are generous mortgage interest deduction subsidies and lower transaction taxes. Hence, Dutch private-sector renters are likely to move quickly to social housing when they are able to get off the waiting list and to owner-occupied housing when they can obtain a mortgage. Our measurement in NL focuses on renters who moved fairly recently. Recent movers are statistically more likely to move in the future. Renters move less frequently in BE; the standard rental contract is for 3, 6, or 9 years. We measure turnover in BE when we observe a new rental contract for the same unit with a different tenant.[10] BE has much more limited mortgage interest deductibility, slowing transitions from rentership to ownership.

Unsurprisingly, the costs of non-payment are highest in low-tier rentals. However, even under our conservative assumptions, the differential costs of rental delinquency can only explain about 9 basis points of the gross yield gap in NL. The spread is 7bps for BE.

Finally, given that management costs are a fixed fraction of rent, additionally subject to a minimum fee in NL, they weigh more heavily on properties with low rental values. The gap is 18 bps in BE and 29 bps in NL.[11] If low-rent properties are more likely to be self-managed, our net rental yield difference between D1 and D10 is understated. Put differently, we assign a financial value to the effort put into self-management of the property.

Taken together, the cost yield is 2.77% for D1 properties (44% of rent revenue) and 1.53% for D10 properties (36% of rent revenue), accounting for 1.24% of the 2.13% spread in BE. In NL, the cost yield ranges from 2.51% in D1 (37% of rent revenue) to 0.95% in D10 (23% of rent revenue), accounting for 1.56% of the 2.66% gap in gross yields in NL.[12] We find that there remains a sizable difference in the net yield between low- and high-tier properties of 1.12% in NL and 0.89% in BE after carefully accounting for the various costs.

[10]While this could miss some subleasing activity, subleasing does not represent a loss to the landlord and hence does not affect our cost estimates.

[11]We report management costs under alternative scenarios in Appendix Table A2. In NL, a model with only fixed costs would reduce the slope by 9bp, a variable cost model would increase it by 15bp.

[12]The lower cost yield for D10 properties in NL than in BE is to a large extent driven by the fact that D10 properties have higher average prices in NL than in BE both in absolute terms and relative to D1 properties (Table A2).

Table 2: Returns by rent decile (2023 prices)

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D1-D10
						Panel A: Belgium					
Gross yield	6.32	6.07	5.69	5.39	5.13	4.88	4.68	4.46	4.24	4.19	2.13
Cost yield											
Maintenance cost	1.59	1.58	1.48	1.38	1.30	1.22	1.17	1.09	0.99	0.81	0.78
Property taxes	0.49	0.40	0.37	0.36	0.35	0.34	0.33	0.32	0.31	0.32	0.17
Turnover	0.06	0.05	0.04	0.04	0.04	0.03	0.03	0.03	0.03	0.02	0.04
Tenant default	0.09	0.09	0.07	0.05	0.05	0.05	0.03	0.03	0.03	0.02	0.07
Management	0.54	0.51	0.48	0.46	0.43	0.41	0.40	0.38	0.36	0.36	0.18
Net yield	3.56	3.44	3.24	3.10	2.95	2.83	2.73	2.61	2.52	2.66	0.89
Cap. gain yield	4.02	3.87	3.83	3.74	3.76	3.62	3.79	3.66	3.60	3.18	0.84
Return	7.58	7.31	7.07	6.84	6.71	6.44	6.52	6.27	6.13	5.84	1.74
						Panel B: Netherlands					
Gross yield	6.84	6.24	5.77	5.36	5.03	4.74	4.52	4.40	4.26	4.18	2.66
Cost yield											
Maintenance cost	1.36	1.22	1.13	1.10	0.90	0.80	0.71	0.63	0.53	0.40	0.96
Property taxes	0.34	0.33	0.32	0.30	0.26	0.25	0.23	0.21	0.19	0.16	0.18
Turnover	0.23	0.20	0.18	0.17	0.16	0.16	0.15	0.15	0.16	0.17	0.06
Tenant default	0.11	0.09	0.08	0.07	0.05	0.04	0.03	0.03	0.02	0.02	0.09
Management	0.48	0.38	0.33	0.29	0.26	0.23	0.21	0.21	0.19	0.19	0.29
Net yield	4.33	4.01	3.73	3.52	3.38	3.27	3.19	3.18	3.17	3.23	1.10
Cap. gain yield	5.69	6.52	6.48	6.34	6.01	5.91	5.90	5.54	4.76	2.71	2.98
Return	10.02	10.54	10.21	9.87	9.41	9.19	9.09	8.72	7.92	5.94	4.08
						Panel C: United States					
Net yield	5.70	5.88	5.81	5.72	5.58	5.42	5.31	5.21	5.10	5.09	0.61
Cap. gain yield	5.53	5.23	4.72	4.23	3.91	3.38	3.13	2.79	2.39	2.28	3.25
Return	11.23	11.11	10.53	9.95	9.49	8.80	8.44	8.00	7.49	7.37	3.86

Notes: Gross yields, cost yields, net yields, capital gains, and the corresponding total returns per decile of rent for Belgium (BE, panel A), The Netherlands (NL, panel B), and the United States (US, panel C).

These calculations highlight the importance of taking into account heterogeneity in costs when examining return differences across properties. Often, given limited data, the literature assumes that costs are a fixed fraction of either rental revenue or property value. Neither assumption is appropriate throughout the cross-section.

Cost differences also have important implications for the measurement of the housing consumption and returns of owner-occupiers (Coffey et al., 2022; Diamond and Diamond, 2024). Together, turnover, default, and management costs are sizable, particularly at the low end of the rent distribution. In BE (NL) they account for 0.69% (0.82%) in D1, compared to 0.40% (0.38%) in D10. Landlords increase rents to recuperate such costs in a competitive rental market (e.g. Abramson, 2023). Owner-occupiers of similar properties, however, do not bear these costs since these particular expenses are unrelated to housing consumption. This leads to a discrepancy between observed rental expenditures of renters and unobserved housing consumption of owners. Failing to adjust rents for such costs would lead one to overstate housing consumption and returns, and more so for low-value properties.

Because many of these costs are uncertain, such as the required amount of maintenance and the likelihood of turnover or non-payment, the cash flows of low-tier rental units may be more risky. However, idiosyncratic expense risk at the unit level can be partially diversified away at the building level (for multifamily properties) or by holding a portfolio of (single- or multi-family) units. We return to this discussion in the next section when we explore risk-based explanations of the net rental yield and return gap.

4.3 Renovation Measurement

One potential concern is that buyers of low-tier properties disproportionately engage in major renovations following the purchase. Our cost measurement in the U.S. does not capture major capital expenditures since NOI excludes one-off capex. The Dutch and Belgian maintenance data impute the average annual capex.[13] The major capital improvements allow the owner to increase the rent, relative to the pre-renovation purchase price. Hence, unmeasured renovation expenditures result in high measured gross and rental yields. They also result in high measured capital gains since the resale value reflects the improved property.

We explore this concern using data on renovations for the US (in Appendix B.4) and BE (Appendix D.4). We find the opposite pattern: properties in the lowest rent segments are less likely to ever have been renovated, and conditional on a renovation, the renovation occurred a longer time ago.[14] We conclude that unmeasured renovation

[13]It is worth noting that the Koeter maintenance model for BE and NL does include the major renovations that have to be made over the next 99 years to continue operating the property, which vary with the features of the property, but it imputes the annualized cost. If those renovations are undertaken shortly after purchase, this approach understates the realized capital expenditures.

[14]In addition, repeat-sales capital gains would be higher to reflect renovation expenditures whereas capital gains from hedonic valuation models usually would not, or at least not to the same extent. If

expenditures cannot explain our facts. If anything, they make the findings stronger.

5 What Can Explain the Return Gap?

Having documented a new stylized fact that properties with low rent levels earn much higher, this section delves into potential risk-based explanations. Ultimately, we conclude that these explanations cannot fully account for the return differences, leaving a positive alpha on low-rent relative to high-rent properties. We provide arguments and evidence for the economic forces that sustain this alpha in equilibrium.

5.1 Systematic Risk

We start our inquiry with the standard risk-based asset pricing explanation: Maybe low-rent properties have higher returns because they have a commensurately higher amount of systematic risk, or beta? We measure bad aggregate states by low GDP growth and ask whether cash flow growth among low-rent properties is lower in bad aggregate states of the world than cash flow growth for high-rent properties.

To explore this possibility, we take advantage of a unique feature of the US data. For a subsample of properties (loans originated by Fannie Mae), we observe multiple years of net cash flows (NCF). The poperty-level NCF can be extracted from time series information on the DSCR at the property level and information on debt payments. We also have a fairly long sample in the US.

Figure 2a plots median NCF growth among properties in each net rent decile. The dashed line represents real GDP growth. The graph illustrates that, if anything, cash-flow growth of low-rent properties is less pro-cyclical than that of high-rent properties. During the Covid-19 recessions, NCFs even increased for low-rent properties while they decreased for high-rent properties. Similarly, during the 2008 financial crisis, properties in the lowest decile demonstrated positive growth, while properties in higher deciles experienced declines. This result is not specific to the United States: Figure 2b reports the same relationship in estimated rental growth rates over the 2008-2023 period in NL and Figure 2c for Belgium.

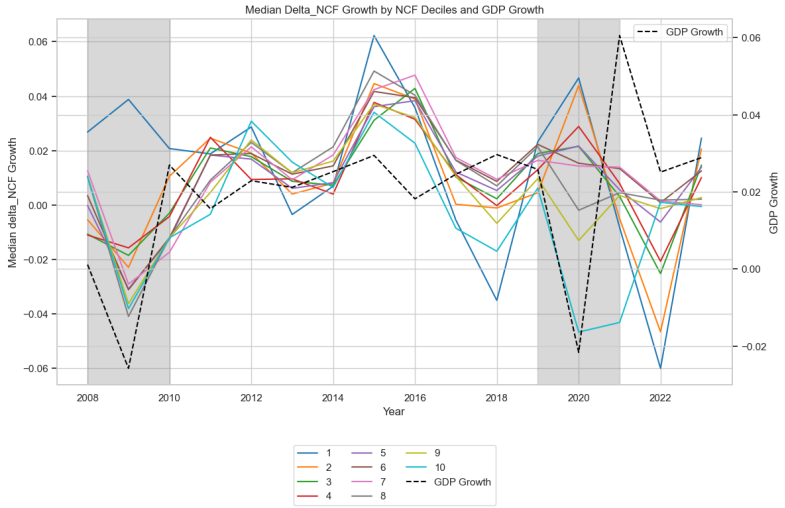
Returning to the US, Figure A1 uses the lagged stock return as the measure of the state of the economy and finds the same relationship with cash flow growth and total returns on rental housing. There is no evidence that low-rent housing is riskier by these metrics.

Shelter is a necessity. It has a consumption floor: there is a minimum house size/quality that is required to satisfy basic shelter needs. Guvenen et al. (2014) show that the negative cross-sectional skewness of labor income growth becomes more pronounced in recessions.

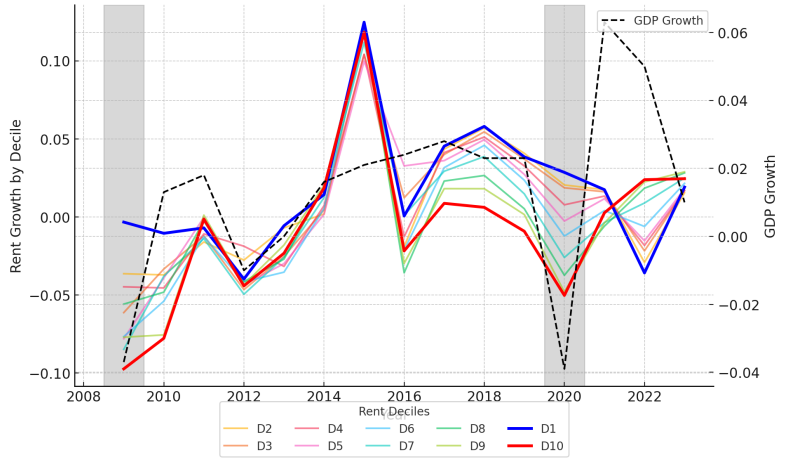
low-rent properties were renovated more often, the difference between the repeat sales capital gain and the hedonic capital gain would be higher in the lower- than in the higher-rent tiers. Figure B4 shows that this is not the case in the US and Table C2 shows this is not the case in NL either.

Figure 2: Covariance of Cash Flow Growth With GDP Growth

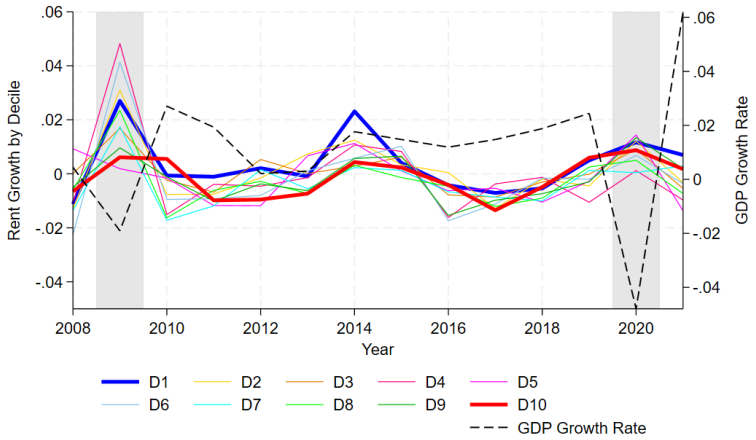
(a) US



(b) NL



(c) BE



Notes: For the US, the figure plots the median annual NCF Growth on rental properties sorted by underwritten real NCF per unit alongside real US GDP growth. NCF growth is computed using the Fannie Mae Performance Data sample. For NL, the figure plots the average rent growth by rent decile alongside real Dutch GDP growth. Rent growth is computed using Realstats listings sample. For BE, the figure presents real rent growth from a repeat rent index by rent decile alongside real Belgian GDP growth.

In other words, labor income reductions are larger for lower-income households in downturns. Low-income households' desire to reduce housing expenditures may lead to an increase in demand for low-rent housing. Affordable housing is an inferior good. This could explain why its price, the rent, is not falling as much, or even rising, in recessions. From the investor's perspective, affordable housing is a good recession hedge.[15]

In sum, the traditional systematic risk explanation seems to go the wrong way to explain the pattern in rental housing returns along the quality dimension. Low-rent housing seems to have a higher CAPM alpha than high-rent housing.

5.2 Regulatory Risk

Maybe the above model is missing important sources of aggregate risk? One prominent candidate is regulatory risk that affects the ability of landlords to charge market-conform rents. While existing renter protections are already reflected in observed rents and prices, potential future tightening of such protections represents a source of risk to rental housing investments. If such risks materialize, they may result in a disproportional loss of cash flow and value for low-rent properties.[16] Ex-ante, investors may require a return premium to invest in low-tier properties.

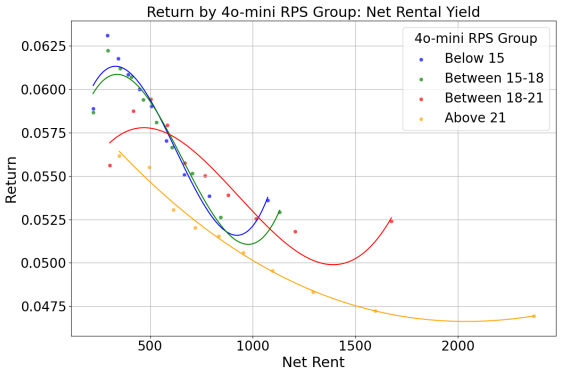
We use our US sample to investigate this possibility. Appendix E provides the details. We exploit cross-state variation along three dimensions of regulatory risk. First, we create our own renter protection score, "GPT RPS," based on 15 aspects of tenant protections present in State laws. We use GenAI for this classification, corroborating it with extant renter protection indices (e.g., McCollum and Milcheva (2023)) and showing robustness across various GenAI methods. Figure 3a shows the relationship between net rent and total return for four groups of US States, sorted by strength of tenant protection laws where higher scores denote stricter tenant protections. Second, we classify States based on the average level of their State-level economic policy uncertainty index from Baker et al. (2023). Many state laws pertain to housing. Figure 3b shows the relationship between net rent and total return for four groups, sorted from low to high levels of local EPU. Third, Figure 3c classifies States in groups based on the concentration of political power in the three branches of state-level government. A State is in the Democratic Trifecta group in a given year if the House of Representatives, the Senate, and the Governor are controlled by Democrats in that year. Intuitively, a Democratic trifecta increases the risk of stronger tenant protections since Democrats have been more supportive of such policies in the past.

[15] Anecdotal evidence from industry reports confirms the notion of affordable housing as a "recession resistant inflation hedge".

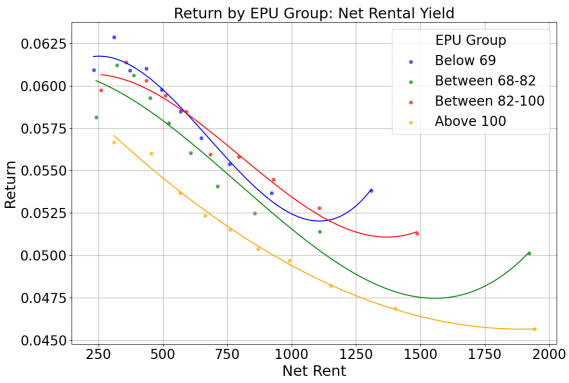
[16] As a case in point, the Housing Stability and Tenant Protection Act of 2019 introduced significant changes to New York State's rent regulations, aiming to enhance tenant protections. Recent data indicates a notable increase in mortgage delinquency rates in New York, particularly affecting rent-regulated multifamily properties. As of the second half of 2023, the delinquency rate for such properties nearly doubled from 2019 levels, after remaining relatively stable since 2001.

Figure 3: Return by Regulatory Risk

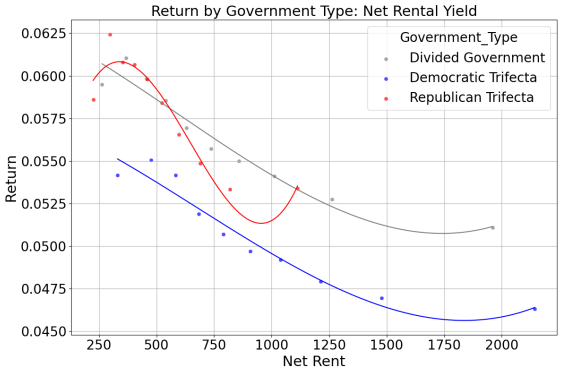
(a) Return by Renter Protection Score, US



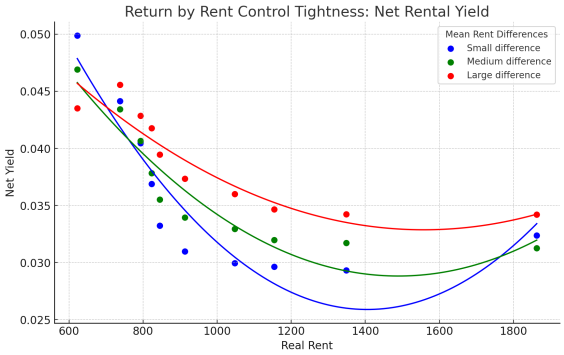
(b) Return by Policy Uncertainty, US



(c) Return by Political Control, US



(d) Return by Exposure To Rent Ceilings, NL



Notes: Every panel sorts states (US) and NUTS-3 regions (NL) into three or four groups based on the US renter protection score (panel 3a), US economic policy uncertainty (Panel 3b, US political control (Panel 3c) and NL tightness of local rent ceilings (Panel 3d), and plots the yield slope for each of these groups.

If the risk of stricter tenant protections were a key factor in generating the negative relationship between rent levels and returns, we would expect a stronger negative slope in locations more exposed to such risks. The data do not bear this out. If anything, we find *flatter* rather than steeper relationships in areas with stronger tenant protections, more economic policy uncertainty, and Democratic political control. Moreover, net yields and total returns tend to be *lower* for low-rent properties in such areas. This is exactly the opposite of what the regulatory risk theory predicts.

We can also investigate the regulatory risk hypothesis in NL. There, private rentals are subject to rent control. The government sets ceilings on rent based on a property-level score that reflects the property's characteristics. The rent ceiling applies to properties with a score below the "rent liberalization" threshold. About 50% of private rentals in our sample fall in this category. However, this rent ceiling was not strictly enforced. Many landlords charged rents (well) above of the ceiling, essentially market rents.[17] The

[17]Tenants have the right to appeal to a rent committee. A successful appeal lowers their rent to the ceiling. However, tenants must appeal within six months of signing a new lease otherwise the contract

key risk they faced is that of stricter future enforcement of the rent ceiling regulation. Different regions are differentially exposed to that risk, since they vary in the share of rentals with scores below the threshold and in the prevalence of rent ceiling violations. In fact, this risk materialized in July 2024, after the end of our sample period, when rent controls were declared binding and their reach extended to 90% of the private rental sector.[18]

We use our micro data to estimate the average difference between the current rent and the rent ceiling using property characteristics, with details provide in Appendix E. We then sort the 40 NL regions (COROP) into three groups based on the average distance between contract rents and the rent ceiling across properties in each rent decile. Figure 3d plots the relationship between rent levels and returns for these three groups of regions. We find that the slope is steepest in regions where rent controls were the least binding (if they were to be enforced). This is the opposite of what we would expect if this regulatory risk were a primordial concern.

In sum, the body of evidence from the US and NL suggests that regulatory risk, resulting in a reduced ability to increase rents or recover costs in the lower-rent segment, is not a key driver of our results.

5.3 Idiosyncratic Risks

Real estate markets contain many frictions, so it might not be too surprising that standard systemic risk explanations fail to explain cross-sectional differences in real estate returns. Notably, real estate portfolios may not be well-diversified, thereby exposing investors to significant idiosyncratic risk. Such exposure could be compensated in equilibrium (e.g. Eiling et al., 2019; Amaral, 2024a; Levy, 2022). If low-tier properties are exposed to more idiosyncratic risk than high-tier properties, then differential idiosyncratic risk exposure could potentially explain the observed return premium. The plausibility of this explanation depends on the idiosyncratic risk differences across tiers of the rental market and on the extent of portfolio diversification.

5.3.1 Unit-Level Risk

We start by examining risk at the level of an individual housing unit, making use of our unit-level data for BE and NL. While we do not have data on the actual net cash flows for these units, we can investigate whether tenant and property characteristics differ significantly across deciles, and whether such difference could be plausibly related to revenue or expense risk. Second, investors might also be concerned about liquidity risk.

rent becomes legally binding. Due to a general lack of awareness about this appeal process, only 664 private tenants appealed their initial rent in 2022. See 2022 Report Dutch Rent Committee.

[18] Using a back-of-the envelope calculation, we estimated the paper loss of rental revenue from the July 2024 reform by rent decile. We find the highest effect in D10 (-22.5%) and the lowest in D4 (-10.0%). In D1, the effect is -13.5%

If low-tier units are less liquid or experience more uncertain capital gains, that might lead to lower prices (e.g. Amaral, 2024b). Table 3 reports various property, tenant, and risk characteristics.

In both countries, low-rent properties are smaller, older, and more likely to be apartments. The age variation is rather limited in the bottom five deciles (even more so in NL than in BE), and apartments tend to have lower maintenance costs than houses. From these statistics it is not obvious that low-rent properties face more maintenance risk.

Taking advantage of the merge with the income tax data in NL, we find that low-rent properties tend to have tenants with lower income. On the other hand, the risk of a large negative income shock of 25% or more is low for D1 tenants and rises from D1 to D7. Tenants in low-tier are more likely to have problematic debts, but this does not differ markedly between D1 and D5, where most of the slope concentrates. Tenants are on average younger and more likely to be between the ages of 18 and 25, a category that includes most higher education students. D10 properties have a much higher share of expats, new residents in the NL. New residents in the country make up only a slightly higher share of D1-D2 tenants than of D3-D4 tenants. In BE, we use neighborhood income as a proxy for tenant income. We also find lower incomes in the low-rent segment, lower educational achievement, and a higher share of foreigners. There also is a higher share of expats in the D10 group in BE, but the difference with D9 is much less pronounced than in NL. While lower incomes make it harder to afford the rent, and could lead to higher risk of non-payment, the lower risk of sharp income drops goes in the opposite direction.

Turning to unit-level risk characteristics, tenant turnover is higher for D1 properties in both BE and NL. This is in part driven by a higher share of students who have elevated churn. Turnover is declining in rent levels in BE, whereas it is highest in D10 in the NL, likely due to the very high share of expats in NL. Vacancy duration is both lower and less volatile in D1 than in D10 in NL, and similar across rent groups in BE. The differences in vacancy duration are very small in general. The bottom three deciles have higher risk of non-payment in BE, but delinquency rates are low.[19]

In terms of liquidity risk, we use changes in ownership to measure asset liquidity. D1 properties are the most liquid, with 6.4% of properties sold per year, while D10 properties the least liquid, with a 4.5% sales rate. While the volatility of capital gains is largest within D1, the difference between D1 and D10 is small, and smaller compared to the difference in average capital gains.

[19]Furthermore, landlords in Belgium have the option of renting their property to a Sociaal Verhuurkantoor (SVK), a non-profit social rental agency that leases private rental dwellings from landlords and sublets them to low-income tenants. SVK provides guaranteed rent payments to the landlord, even if the property is temporarily vacant, handles property maintenance (excluding normal wear and tear), property management, and all tenant relations. It also provides an annual tax advantage to landlords and reduces the registration duties at purchase by 5 percentage points. In return for giving up on average 15% of the rent (Van den Broeck and Winters, 2018), the landlords can eliminate all expenses and all expense *risk*. The presence of this option (i) provides an upper bound on the expenses for low-tier properties, (ii) a lower bound on the D1-D10 net rental yield slope, and (iii) rules out expense risk as a plausible explanation of our findings in BE.

Table 3: Characteristics by rent decile (2023 prices)

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D1-D10
Panel A: Belgium											
Property characteristics											
Square meters (mean)	64.2	92.0	102.3	109.2	116.1	122.0	129.4	137.2	146.7	162.0	-97.8
% built pre-1940	37.7	53.8	51.3	46.3	41.6	37.7	32.9	27.1	23.7	21.4	16.2
% built post-2000	4.9	4.7	5.2	5.4	6.5	7.3	8.7	10.7	14.0	13.9	-8.98
% apartments	65.27	42.2	39.7	40.1	40.0	38.7	37.4	34.6	32.8	28.77	36.51
Neighborhood characteristics											
Monthly income	2752	2816	2887	2965	3022	3089	3167	3286	3432	3850	-1097
% Lower education	38.8	40.0	39.3	38.1	37.2	35.8	34.6	32.6	30.0	24.2	14.58
% Foreign population	12.7	10.6	10.2	10.0	10.01	10.5	10.6	11.5	13.4	18.07	-5.35
Risk characteristics											
Turnover rate	11.2	9.4	8.8	8.0	8.0	7.9	7.1	7.3	7.2	6.5	4.70
Vacancy duration (mean, days)	32.6	31.6	35.1	29.3	34.3	24.8	26.7	24.5	25.88	32.0	0.5
% 2 or more nonpayments	4.1	4.9	4.4	3.4	3.7	4.0	2.2	3.1	3.3	2.7	1.4
Panel B: Netherlands											
Property characteristics											
Square meters (mean)	40.2	55.1	67.4	76.8	83.5	89.4	94.1	95.7	98.2	92.1	-51.9
% built pre-war	35.7	35.5	30.0	30.0	30.7	27.3	22.9	20.5	23.5	41.3	-5.6
% built post-2000	10.6	10.2	8.3	9.4	12.8	19.7	31.1	42.2	44.0	32.0	-21.4
% apartments	80.2	73.6	68.3	62.2	61.5	60.1	59.3	62.4	67.3	81.9	-1.7
Tenant characteristics											
Monthly income	2046	2248	2445	2622	2923	3324	3950	4423	4975	5801	-3755
% with 25% income loss	7.1	7.4	7.8	8.1	9.0	9.5	9.9	9.6	8.8	7.0	0.1
% with debt problems	11.8	12.7	13.1	12.3	11.1	9.4	8.2	7.5	6.8	5.1	6.7
Age (mean)	31.8	33.8	35.1	36.2	36.7	37.9	38.5	38.9	38.4	37.1	-6.3
% age 18-25	42.1	30.5	25.5	21.4	19.4	17.0	13.2	10.4	8.9	8.2	-33.9
% new resident	15.4	15.0	12.5	12.5	14.1	14.7	16.5	19.9	28.8	51.5	-36.1
Risk characteristics											
Turnover rate	29.6	26.3	24.0	22.9	22.5	22.8	22.7	22.8	25.8	32.5	-2.9
Vacancy duration (mean, days)	27.1	28.0	28.9	29.3	30.3	31.2	30.6	30.9	31.9	33.9	-6.8
Vacancy duration (SD, days)	26.6	27.3	27.6	27.2	27.1	28.0	26.8	27.1	27.1	28.4	-2.2
Sale liquidity	6.4	5.7	5.8	5.8	5.6	6.0	6.0	5.3	5.1	4.5	1.9
Capital gains (SD)	7.5	7.2	6.8	6.7	6.7	6.4	6.2	6.1	6.3	6.8	0.7

Notes: Property and tenant characteristics at the housing unit-level for Belgium (BE, panel A) and The Netherlands (NL, panel B).

On balance, these statistics do not suggest meaningfully higher idiosyncratic revenue or expense risk for low-rent compared to high-rent properties, and similarly do not point to higher liquidity risk for low-tier properties.

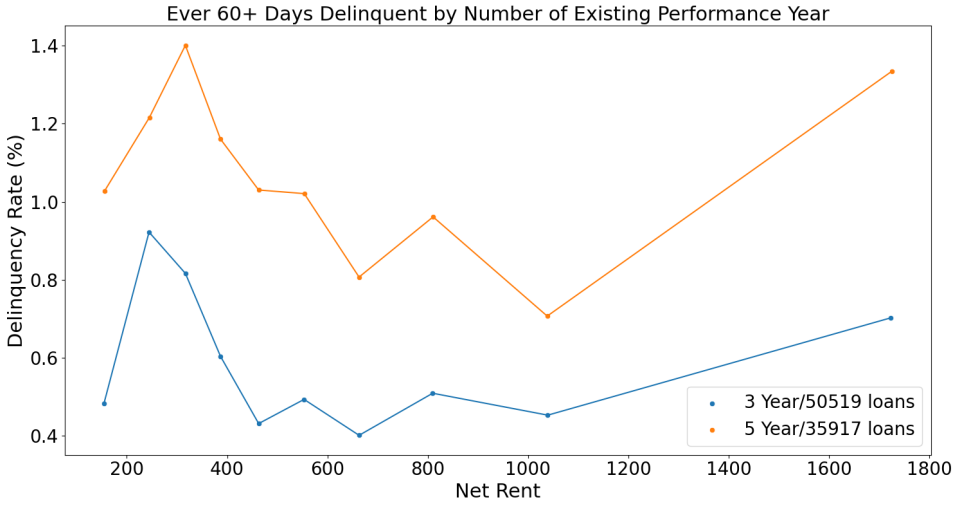
5.3.2 Property-Level Risk

The extent to which higher cash flow risk at the unit level, if present, affects valuations depends on the extent to which property owners are able to diversify this risk across units in a given property or across properties.

We turn to the US data, which contains multifamily properties with multiple units, to investigate cash flow risk at the property level.

If low-rent multifamily properties had severely more left-tail net cash flow risk than higher-rent properties, we would expect this to translate into higher mortgage delinquency rates. To examine this possibility, we use the US Fannie Mae mortgage performance data which dynamically track the mortgage payment status. Figure 4 plots the share of mortgages that are ever 60 days or more late in the three (five) years since origination in blue (orange). We make two observations. First, multifamily delinquency rates are low. Only about 0.5% (1%) of mortgages become delinquent in the first 3 (5) years. Second, the pattern of delinquency is irregular in the level of rents. The lowest-rent properties have low delinquency rates. The delinquency rate then ticks up for D2 and D3 and comes back down for the middle-rent deciles. D10 has higher delinquency rates than D9 and D1. There is little difference in the 60-day delinquency rate over three (five) years between the first decile and the top-half of the NOI distribution. This evidence is inconsistent with substantially higher idiosyncratic left-tail cash-flow risk at low-tier properties.

Figure 4: Mortgage Delinquency Rate by NOI Level



Notes: Bin-scatter plot of Percent of Delinquency vs Real NOI Per Unit.

Default on a multifamily loan is a tail risk event. It also depends on the the bal-

ance sheet strength of the landlord, which could differ across rental market segments. An additional and maybe more direct way to study idiosyncratic cash flow risk at the property level is to measure the volatility of cash flow growth rates. Higher net cash flow growth volatility could stem from more volatile rents, operating expenses, and capital expenditures, or from higher rates of tenant turnover, longer spells of vacancy conditional on turnover, and higher rates of missed rent payments, partial payments, or rent forbearance.

Since we observe the actual evolution of net cash flows (NCF) at the property level in our US data, we can investigate whether low-rent properties have higher cash flow risk. Since the Fannie Mae performance data tracks a loan's DSCR over time, we can compute multiple years of NCF for the same property. To deal with potentially negative or zero NOI values, we define the annualized growth rate of NCF between year t and $t+T$, Δ^i_{t+T}, as follows:

$$\Delta^i_{t+T} = \left(\frac{NCF^i_{t+T} - NCF^i_t}{\frac{1}{2}\left(NCF^i_{t+T} + NCF^i_t\right)} + 1 \right)^{\frac{1}{T}} - 1$$

Year t is the actual net cash flow in the first year after loan origination.[20]

We define the future cash-flow risk of a property that is currently in rent decile i as the *cross-sectional* standard deviation of NCF growth among all properties in decile i. The idea is that the same factors, such as particularly negative NCF due to high maintenance expenses or particularly high turnover vacancy in a given year that affect some properties in decile i, may affect other properties in that same decile in some future year.

Figure 5 plots the volatility of NCF growth plotted against net rent for horizons $T = 1, \cdots, 4$. The relationship is decreasing between D1 and D7, but increasing modestly from D7 to D10. Properties in the middle of the rent distribution have the most stable cash flows, while properties on both extremes of the rent distribution experience more cash flow growth volatility. Quantitatively, the difference between D10 and D1 is not very large for realistic holding periods of 3-4 years. The D1-D10 volatility gap is 5.5% at the 1-year horizon but shrinks to 1.5% at the 3- and 4-year horizons.[21] To justify a 3% point return gap based on a 1.5% point volatility gap, one would need a representative investor with constant relative risk aversion preferences to have a risk aversion coefficient of 267.[22] Standard values for risk aversion of, say, 5 can only justify return difference of 5.6 (75.6) basis points given volatility differences of 1.5% (5.5%).

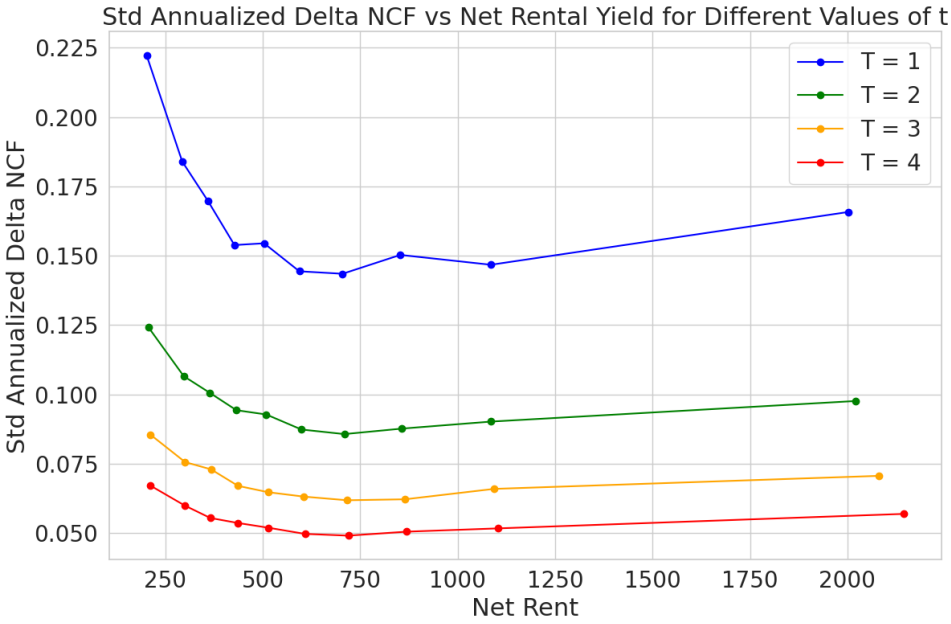
We also note that, in the US data, net cash flow growth rates are 150 to 200 basis points per year higher in D1 than in D10 (see Figure A2). Hence, higher cash-flow risk

[20] In other words, this calculation does not rely on underwritten NCF, only on the actual NCFs that materialized and that reflect actual revenues and expenses.

[21] We note that low-NCF properties tend to have more units in the data, and hence better diversification of idiosyncratic risk across units within the same property, than properties with higher NCF per unit. This implies that the slope of the volatility of NCF growth curve is more steeply downward-sloping at the unit level than at the property level.

[22] Equalizing the certainty equivalent return between two assets for an investor with risk aversion γ gives: $\mu_1 - \frac{\gamma}{2}\sigma_1^2 = \mu_2 - \frac{\gamma}{2}\sigma_2^2$. Setting $\mu_2 = \mu_1 + .03$ and $\sigma_2 = \sigma_1 + 0.015$ results in $\gamma = 266.67$.

Figure 5: Cash Flow Growth Volatility



Notes: Bin-Scatter Plot of the cross-sectional standard deviation of net cash-flow growth rates over horizon T, $Std(\Delta^i_{t+T})$. Computed using the Fannie Mae Performance Data sample.

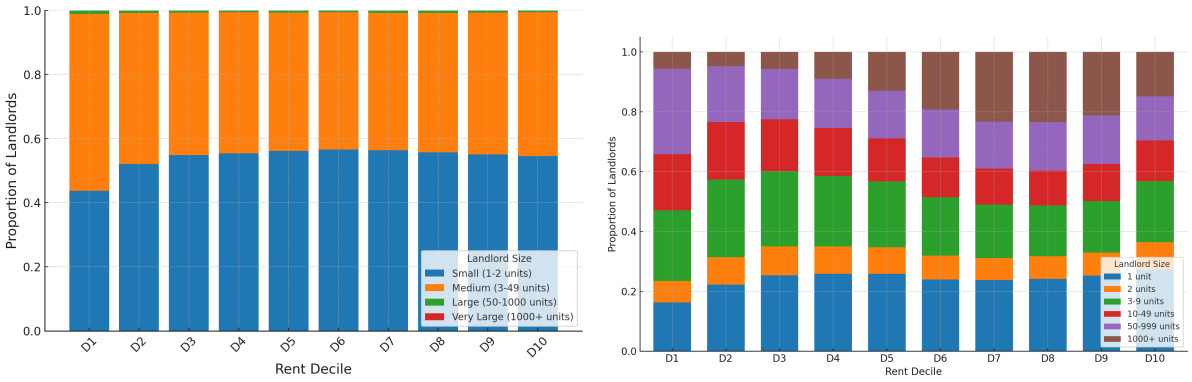
does not necessarily translate into a higher likelihood of negative cash flow growth, which would otherwise increase the probability of delinquency. The higher cash-flow growth result for D1 also holds for BE (Figure A3) and NL (Figure C2c), where cash flows are measured as rent revenues.

In short, while the US data provide some evidence that cash flows on low-tier properties are more volatile, the additional volatility is modest over typical holding periods and cannot explain the observed return differences for reasonable values of risk aversion. It also does not translate into materially higher mortgage default rates, which might be due to higher cash flow growth. Furthermore, the fact that default and cash flow risk are also higher for high-rent properties calls into question the notion that idiosyncratic cash flow risk is a primordial driver of the return differences between low-rent and high-rent properties. Finally, as we discuss below, landlords in the lowest-rent segment tend to hold a substantial number of properties, substantially diversifying the idiosyncratic risk, and lessening its impact on portfolio-level volatility.

5.3.3 Landlord Portfolio Size

The extent to which idiosyncratic cash flow risk is compensated by the market in the form of a higher average return depends on the extent of portfolio diversification. While we cannot directly observe risk at portfolio level, we can use our NL and BE data to investigate differences in portfolio composition across rental market segments and explore the relationship with average returns.

Figure 6: Landlord Size by Rent Decile



(a) BE (b) NL

Notes: Distribution of landlords across rent deciles by their portfolio size. Portfolio size is observed administratively in the Dutch data, while it is inferred from rent contract data in BE, which excludes properties that were not leased since 2011.

We define small landlords as those who own one or two rental properties. Medium landlords own between 3 and 49 properties, large investors own between 50 and 999 properties, and very large investors own 1,000 properties or more. Most small and medium landlords own property as a natural person (personally), whereas those owning over 50 units typically own properties through a corporate entity (through a corporation). Figure 6 plots the distribution of of the number of units owned by landlord size across the rent distribution. In contrast with NL, corporations only own 5% of units in BE.

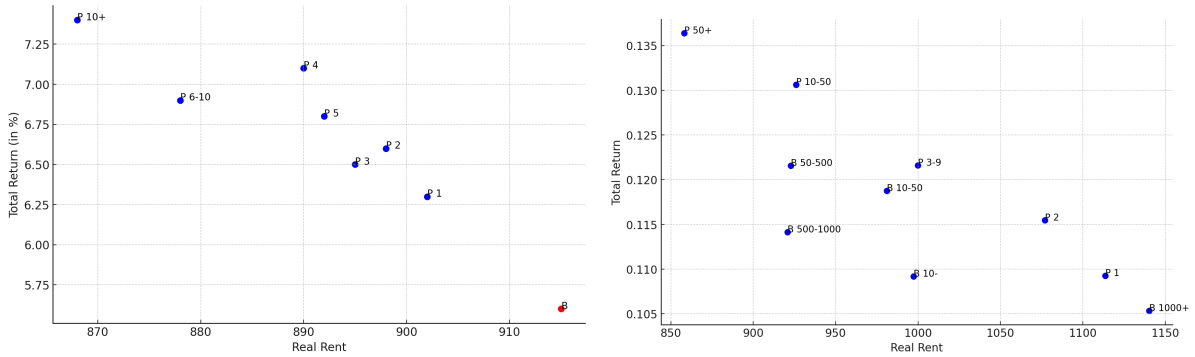
The low-rent segment of the market is dominated by medium-size investors, especially in NL. Both small investors and very large investors shy away from the lowest tier of the market in NL. Very large investors are most active in the top-third of the rental distribution. Small investors are also prominent in D9 and D10. This shows that those owning property in the lowest-rent segment are fairly well-diversified against the potentially higher idiosyncratic cash flow risk in that segment. It also shows that different types of investors have their own preferred habitats.

We cannot easily generalize this analysis to the US. Our US sample includes only investors who take loans for multifamily properties, which, by definition, are larger investors. That said, Figure A5 uses Real Capital Analytics data to show that even among landlords investing in properties worth more than $2.5 million, the largest investors tend to concentrate in the higher-rent tiers.

This segmentation in the rental market results in different types of landlords earning different returns. Figure 7 plots the relationship between rent and total return by landlord type in BE and NL. In this plot, we use a finer split by portfolio size and by ownership structure (P for personal, B for business/corporate). We find a strong negative relationship between rent levels charged and total returns earned.[23] In NL, the highest

[23]We find exactly the same pattern in net rental yields (not reported).

Figure 7: Total return by predicted rent and landlord type



(a) BE (b) NL

Notes: This plot shows the relation between total return and rental value, with every point corresponding to the average observed rent and rental yield per owner type. We distinguish personal (P) and corporate (B) owners and group them by portfolio size.

total returns are earned by personal investors owning more than 50 units, followed by personal investors who own between 10 and 50 units. In BE, the highest returns are similarly earned by personal investors with more than 10 properties. Small personal investors who own 1 or 2 units are active in the higher-rent segments and earn lower returns. In the NL, where we have more variation in the sizes of corporate owners than in BE, very large corporations charge the highest rents and earn the lowest returns.

5.4 Limits To Arbitrage

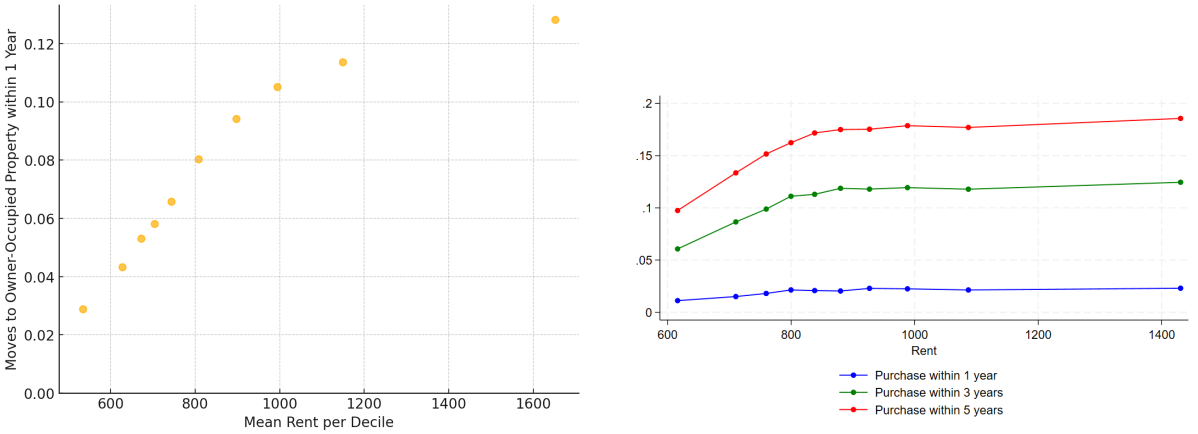
Having cast doubt on a risk-based explanation for the downward-sloping relationship between rent levels and returns (or net yields), we are left with the possibility that investors in low-tier rental housing earn positive abnormal returns. This begs the question what forces sustain this situation as an equilibrium outcome. What prevents capital from flowing into the low-rent market segment and eliminate the excess returns? We argue that there are three necessary conditions and that all three are plausibly satisfied.

5.4.1 Low-income Renters Cannot Buy

The first source of additional capital that could eliminate the arbitrage opportunity comes from the households that are currently renting in the low-rent segments. Since their rent is high compared to the value of the property they rent, and they seem to be missing out on high capital gains, they may be better off buying their property. They could take out a mortgage to finance the purchase.

Financial Constraints The first problem with this argument is of course that financial constraints are much more likely to be binding for D1 tenants. Purchasing their property requires coming up with the required downpayment to overcome the loan-to-value con-

Figure 8: Transition to owner-occupied properties by rent decile, NL+BE



(a) NL (b) BE

Notes: This figure plots the likelihood that a tenant living in a rental property moves to an owner-occupied property by current rent decile. In NL, this is measured over a one-year horizon for all households living in a private rental on January 1. In BE, this is measured from the date of signing a rental contract, over horizons of one, three, and five years.

straint and with sufficient (and sufficiently stable) income to overcome the debt-to-income constraint. In the NL data, we can match rentals to their tenants, and have shown that D1 renters have the lowest incomes. We find that fewer than 3% of tenants that live in a D1 property move to owner-occupied housing within one year. In Figure 8a, we plot the share of tenants that rent privately on January 1 but live in an owner-occupied property on December 31. Transitions from renting to owning are most common among D10-tenants, a segment of the population where financial constraints are much less binding.[24] Figure 8b reports the same relationship for BE. In BE, we measure transitions starting from the date of signing a new rental contract, and BE rental contracts are usually three, six, or nine years in length. As a result, one-year transition rates are very low. However, over a 3- to 5-year horizon, we clearly see that D1 tenants are less likely to buy property.

Collective Action Frictions Furthermore, there is a preponderance of apartment properties in the D1 segment (recall Table 3). Many of these are large properties which cannot easily be bought by the tenants as it would require a willingness to sell on the part of the landlord and a substantial coordination effort on the part of the tenants. The tenants would need to set up a cooperative or condominium association, bargain with the landlord on the price and with each other on each tenant's financial contribution, vote on the proposed arrangement, agree on a set of bylaws for the coop, etc. Collective action

[24]NL also has a large Housing Association sector, with affordable units for which many D1 tenants qualify based on their income. However, these properties have long waiting lists and therefore do not provide an immediate alternative to renting privately, particularly for young people. Only five percent of D1 tenants move to an affordable property within one year, and this share reduces to one percent in D10. Hence, this channel is unlikely to put downward pressure on rents in D1.

frictions abound.[25]

In NL, we can determine for each rental unit whether it is sold individually — which is always the case when sold to owner-occupants — or as part of a portfolio (e.g., the sale of an entire multi-family building). Appendix Figure A4 shows the share of units sold as part of a portfolio. Consistent with stronger conversion frictions, 57% of existing private rental units in the lowest decile (D1) are sold in a portfolio. This share declines monotonically to 25% in the highest decile (D10). By contrast, there appear to be fewer frictions in forming portfolios: the share of units sold in large deals exceeding 10 million euros remains similar across all deciles.

5.4.2 Large Corporate Investors Decide To Stay Away

Why do medium-sized investors tilt towards the low-rent, high-yield segment, while very large, usually corporate, investors stay away from this segment? We put forward two potential explanations that may work in conjunction with one another.

Reputational risk: Unlike medium-size personal landlords, large corporate investors are by virtue of their size and corporate structure subject to more scrutiny from their shareholders, employees, customers, and regulators. For example, large corporate landlords invest sizable amounts of pension fund money in NL.

Properties in the lowest tiers of the rental market tend to be of lower quality and cater to an economically more vulnerable group of tenants. Consistent with the notion of elevated reputational risk associated with investing in low-rent properties, Figure A6 shows that properties catering to lower incomes experience more evictions (US data) and Department of Buildings violations (New York City data). Appendix B.6 shows that low-rent properties are located in areas with higher rates of property crime.

Anecdotal evidence from large landlords amplifies this reputational risk concern.[26] The capital providers to these large institutional investors may steer management away from the low-rent segment. In contrast, non-corporate investors with not-too-large portfolios can operate under the radar and away from public scrutiny. Yet, they have sufficient scale to diversify away the bulk of idiosyncratic cash-flow risk.[27]

[25] That said, governments in several countries have at times pursued large-scale privatizations of public housing, for example in Sweden (Sodini et al., 2023). In Washington DC, the Tenant Opportunity to Purchase Act requires that landlords first offer the tenants residing in that property the opportunity to purchase it before selling that property to a third-party. Tenants then have the option to convert it to a market rate or affordable condominium or cooperative (Lawton, 2012).

[26] Reputational concerns are not new: Sternlieb (1966) studies low-tier landlords and writes: "The high rate of current return demanded by investors in slum tenements can be summarized as a compound of ... [*various reasons*] ..., and in substantial part, the pejoratives which society heaps upon the 'slum lord'." In modern times, large NL investors indicate that they generally do not serve the riskier segment of the rental market because they perceive a conflict between evicting tenants and the norms surrounding socially responsible investment (Verwey-Jonker Institute, Retrieved January 15, 2025.)

[27] Consistent with this, the 2024 list of 100 worst landlords in New York City, published by the City's Public Advocate, contains almost exclusively medium-sized personal investors.

Diseconomies of scale: Another possible reason why very large corporate investors may shy away from the low-tier segment is diseconomies of scale. While small-scale investors might benefit from economies of scale in management and maintenance expenditures as they scale-up in the low-tier segment, managing very large portfolios of low-tier rental property might prove challenging. For example, low-rent tenants might require more individual attention and flexibility, and lower-quality units may face more expense risk. Keeping track of a more complex property and tenant base may prove uneconomical in very large portfolios.

We can try to test this hypothesis using data on average operating expenses per unit for all major Housing Associations in NL, the large social housing segment of the market. Appendix Table A3 regresses the log average operating expenses per unit on the size category of the Housing Association (Column 1), after controlling for urban municipalities, average rents, and the rent ceiling. The omitted category is Housing Associations with size between 10,000 and 25,000 units. We find evidence for significant diseconomies of scale. Operating expenses per unit are between 14% (Column 1) and 20% (Column 2) lower for associations with fewer than 1,000 units than for the baseline category.

Although we do not have definitive evidence for the reasons that lead large corporate landlords to stay away from the lower-rent segment, these two factors are also mentioned by professionals (e.g., Oxley et al., 2015). More importantly, the fact that they do stay away reduces the flow of capital invested in low-tier rental housing.

5.4.3 Medium-size Landlords Cannot Scale Up

The third source of additional capital could come from the very landlords that are already active in this market segment: medium-size personal investors. Why do they not scale up? We advance two potential explanations that may work in conjunction with one another.

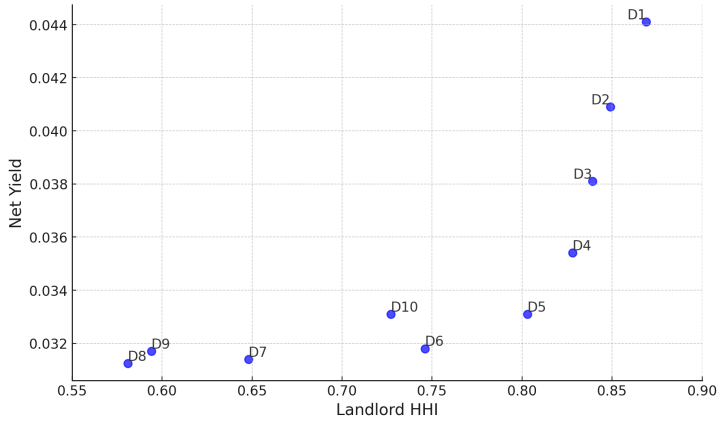
Financial Constraints As personal investors, these medium-size landlords do not enjoy the same access to external public or private equity capital as very large investors. For each additional property they want to invest in, they get a bank loan (usually with a LTV ratio below that of a traditional homeowner) and must finance the (sizable) equity component from retained earnings on the existing portfolio. Debt and especially equity financing constraints prevent them from scaling up more rapidly. This is a case of slow-moving (equity) capital.

Local Bias The NL data allow us to compute the extent of geographic concentration of investors' rental property portfolio. Using portfolio shares in the 40 NL regions, we compute a Herfindahl-Hirschmann Index (HHI) of 0.15 for corporate investors with over 1,000 properties. The HHI increases to 0.50 for investors with 500-1,000 properties, and

is above 0.70 for all other investor groups. The average landlord's HHI is 0.90. While it may not be surprising that small investors do not diversify geographically (e.g. Levy, 2022), even medium and large investors fail to do so.[28] In other words, medium (and small) investors display a strong local bias. This local bias impedes the flow of capital across space. Since at least some of return spread between low- and high-rent properties is accounted for by geography, local bias is an additional force sustaining cross-regional differences in excess returns.

The lack of geographic diversification results in a strong positive relationship between the investor HHI and net yield. Figure 9 shows that properties in D1 are held by landlords that hold highly geographically-concentrated portfolios. This is true despite the fact that the low-tier segment has the fewest small landlords, as we showed in the right panel of Figure 6. Figure A7 shows that the same relationship holds among larger multifamily investors in the US.

Figure 9: Yield by rent decile and average landlord HHI

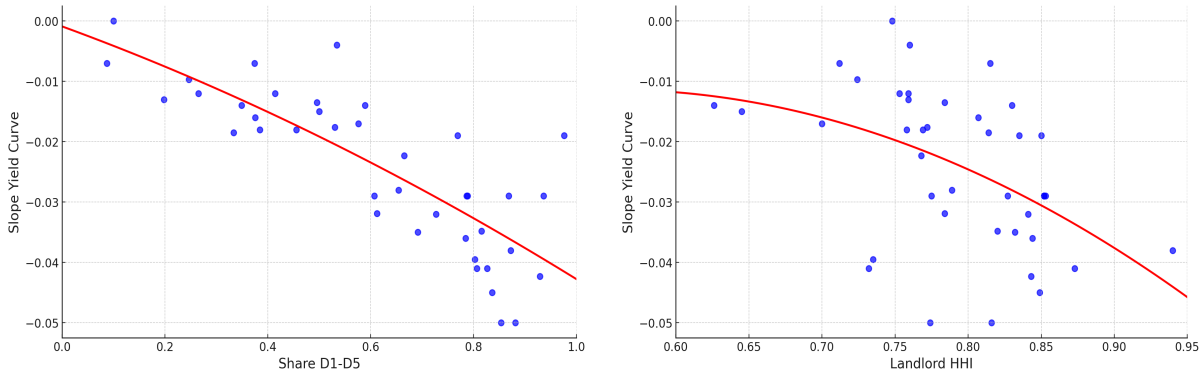


Notes: This plot shows the relation between net rental yield and the landlord HHI in NL, aggregated per rent decile.

To investigate the role of local bias further, we estimate the relationship between the rent and the net yield separately for each of the 40 regions in NL. Figure 10 links the 40 estimated slopes to the fraction of D1-D5 properties in each region (left panel) and to the average landlord HHI in each region (right panel). The slope estimate is negative in all but one region, and hence a pervasive fact even conditional on location. However, the *magnitude* of the slope is heterogeneous and largest in absolute value in regions with a high share of properties in the bottom half of the rent distribution (left panel) and the most local bias among landlords (right panel).

[28] Local bias is also prevalent in other settings such as the stock market (Coval and Moskowitz, 1999), and even among professional investors.

Figure 10: Yield-Rent Slope by Low-Tier Supply and Landlord HHI



(a) Slope and Share Below-Median Property

(b) Slope and Landlord HHI

Notes: This plot shows the relation between the slope of the yield-rent curve and the share of properties in D1 to D5 (left panel) and the geographic concentration of landlord portfolios. Every point corresponds to the relation estimated in one of the 40 NUTS-3 regions in NL.

5.4.4 Summing Up Capital Flow Frictions

Why does capital not flow from high-rent to low-rent market segments? Very large landlords, who are geographically diversified, stay away from the low-tier possibly out of concern for their reputational and/or diseconomies of scale. Medium-size landlords from other regions with few low-rent properties have local bias, and may not be aware of the excess returns. Medium-size landlords already active in areas with more low-tier rentals may lack the external equity capital to scale up meaningfully. Finally, tenants of low-rent properties lack the financial resources to purchase the homes they rent. This constellation of frictions results in an equilibrium where unincorporated medium-sized investors active in regions with a prevalence of low-rent units persistently earn high risk-adjusted returns.

5.4.5 Capital Flows Analysis

The final piece of analysis provides direct support for the capital flows mechanism by showing that, if and when (large, institutional) investors actually manage to overcome the capital flow frictions, the net rental yield and return slopes between low- and high-rent properties flatten. Appendix B.7 for the US and D.5 for BE analyze the impact of flows into the bottom 20% and top-20% of the rent distribution. We find that more purchases in the bottom-20% (top-20%) flatten (steepen) the net rental yield profile and the capital gain profile, with the latter effect being the strongest. We also find in the U.S., that large investors and institutional investors exert a stronger influence compared to smaller and other types of investors. This is consistent with the notion that large, institutional investors help reduce the higher returns on low-rent properties when they scale up their investments in said properties.

6 Conclusion

We document a strong negative relationship between the level of rents and the return to rental housing. An investment in houses in the cheapest ten percent of the market earns returns that are substantially higher than an investment in the most expensive ten percent. Using high-quality micro data for Belgium, the Netherlands, and the United States, we show that this relationship holds robustly across housing market settings and time periods. It is present in gross and net rental yields as well as in capital gain yields.

We argue that natural risk-based explanations such as systemic risk, regulatory risk, or idiosyncratic cash-flow risk associated with underdiversified portfolios are unlikely to explain the entire return gap. Rather, there appear to be limits to arbitrage that plague the low-tier rental market segment. Most landlords display strong local bias and lack the resources to scale up due to financial constraints. The largest institutional landlords who could arbitrage away the returns are held back by reputation concerns and potential diseconomies of scale. Tenants in low-rent properties lack the resources to buy the homes they rent. This equilibrium allows for medium-scale geographically-concentrated investors to earn excess returns, at the expense of low-income tenants.

What are the implications for housing affordability? Our analysis points in the direction of policies that work to alleviate the frictions preventing arbitrage. One set of policies could focus on stimulating the flow of institutional-quality capital towards the lower tiers of the rental market. The US has seen a recent increase in private investment in affordable and "workforce" housing stimulated by federal, state, and local policies such as Opportunity Zones (Theodos et al., 2023), Freddie Mac's Workforce Housing Preservation Initiative, HUD Section 542 guarantees, and Florida's 2023 Live Local Act, just to name a few. In many European countries, such as The Netherlands and Belgium, additional investment in existing low-rent housing is often done through the government or the non-profit sector (Housing Associations).

A second set of policies could focus on reducing financial frictions among medium-size investors. For example, a country's sovereign wealth fund could set up a nationally-diversified housing fund that explicitly targets equity investments in local landlords that own rental properties in the bottom of the rent distribution. Information campaigns providing investors with detailed intelligence on which neighborhoods and properties have the highest returns might help to reduce local bias and stimulate the flow of capital from small and medium investors.

A third set of policies could make it easier for renters in the low-rent segment to purchase their own home. Since providing more and cheaper debt tends to increase house prices and increase households' and banks' financial fragility, programs that provide equity may be preferable. Examples are shared equity programs such as the U.K.'s Help To Buy program (Benetton et al., 2022), or first-time homebuyer tax credits (Mabille, 2023), but would need to target tenants in the low-rent segment. Laws that give tenants

a right of first refusal to purchase their rental unit could also spur this source of capital. Those policies are likely to be more effective when supply is more elastic (Carozzi et al., 2024).

Finally, adding supply in the low-rent tier of the market would also lower returns. One issue here is that new construction of affordable housing is rather expensive. In the US, large federal, state, and local subsidies are paramount for the private development of affordable housing. Adding supply in higher-rent market segments may also help, as it can filter down to the lower-rent tier through moving chains (Bratu et al., 2023; Asquith et al., 2023).

References

Abramson, B. (2023). The welfare effects of eviction and homelessness policies. Technical report, Columbia Business School.

Abramson, B., De Llanos, P., and Han, L. (2024). Monetary policy and rents. Technical report, Working Paper Columbia Business School.

Abramson, B. and Van Nieuwerburgh, S. (2024). Rent guarantee insurance. Technical report, National Bureau of Economic Research.

Amaral, F. (2024a). Price uncertainty and returns to housing. Technical report, Working Paper.

Amaral, F. (2024b). Price uncertainty and returns to housing.

Amaral, F., Dohmen, M., Kohl, S., and Schularick, M. (2021). Superstar returns.

Ambrose, B., Coulson, E., and Yoshida, J. (2023). Housing rents and inflation rates. *Journal of Money, Credit, and Banking*, 55:975–992.

Asquith, B. J., Mast, E., and Reed, D. (2023). Local effects of large new apartment buildings in low-income areas. *Review of Economics and Statistics*, 105(2):359–375.

Autor, D. H., Palmer, C. J., and Pathak, P. A. (2014). Housing market spillovers: Evidence from the end of rent control in cambridge, massachusetts. *Journal of Political Economy*, 122(3):661–717.

Badrinza, C., Ramadorai, T., and Shimizu, C. (2022). Gravity, counterparties and foreign investment. *Journal of Financial Economics*, 145(2):132–152.

Baker, S. R., Davis, S. J., and Levy, J. A. (2023). State-level economic policy uncertainty. *Journal of Monetary Economics*, 132:81–99.

Baum-Snow, N. and Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics*, 93:654–666.

Benetton, M., Bracke, P., Cocco, J. F., and Garbarino, N. (2022). Housing consumption and investment: evidence from shared equity mortgages. *The Review of Financial Studies*, 35(8):3525–3573.

Bezy, T., Levy, A., and McQuade, T. (2024). Insuring landlords. *Working Paper*.

Bracke, P. (2015). House prices and rents: Microevidence from a matched data set in central london. *Real Estate Economics*, 43(2):403–431.

Bratu, C., Harjunen, O., and Saarimaa, T. (2023). Jue insight: City-wide effects of new housing supply: Evidence from moving chains. *Journal of Urban Economics*, 133:103528.

Carliner, M. and Marya, E. (2016). Rental housing: An international comparison. Technical report, Working Paper Joint Center for Housing Studies.

Carozzi, F., Hilber, C. A. L., and Yu, X. (2024). On the economic impacts of mortgage credit expansion policies: Evidence from help to buy. *Journal of Urban Economics*, 139:103611.

Chambers, D., Spaenjers, C., and Steiner, E. (2021). The rate of return on real estate: Long-run micro-level evidence. *The Review of Financial Studies*, 34(8):3572–3607.

Cocco, J., Naaraayanan, S. L., and Tripathy, J. (2024). Individual landlords in the mortgage market. Technical report, Working Paper.

Coffey, C., McQuinn, K., and O'Toole, C. (2022). Rental equivalence, owner-occupied housing, and inflation measurement: Microlevel evidence from ireland. *Real Estate Economics*, 50(4):990–1021.

Collinson, R., DeFusco, A. A., Humphries, J. E., Keys, B. J., Phillips, D. C., Reina, V., Turner, P. S., and van Dijk, W. (2024a). The effects of emergency rental assistance during the pandemic: Evidence from four cities. Technical report, National Bureau of Economic Research.

Collinson, R. and Ganong, P. (2018). How do changes in housing voucher design affect rent and neighborhood quality? *American Economic Journal: Economic Policy*, 10(2):62–89.

Collinson, R., Humphries, J. E., Mader, N. S., Reed, D. K., Tannenbaum, D. I., and Van Dijk, W. (2024b). Eviction and poverty in american cities. *Quarterly Journal of Economics*, 139(1):57–120.

Colonnello, S., Marfè, R., and Xiong, Q. (2024). Housing yields. *Available at SSRN 3970610*.

Cornaggia, J., Spaenjers, C., and Steiner, E. (2025). Who should manage impact investments? evidence from affordable housing. Working paper, SSRN. Latest SSRN version, abstract_id=4874577.

Coval, J. D. and Moskowitz, T. J. (1999). Home bias at home: Local equity preference in domestic portfolios. *The Journal of Finance*, 54(6):2045–2073.

Coven, J. (2025). The impact of institutional investors on homeownership and neighborhood access. Working paper, NYU Stern & SSRN.

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

Desmond, M. (2012). Eviction and the reproduction of urban poverty. *American Journal of Sociology*, 118(1):88–133.

Desmond, M. and Gershenson, C. (2017). Who gets evicted? assessing individual, neighborhood, and network factors. *Social science research*, 62:362–377.

Desmond, M. and Wilmers, N. (2019). Do the poor pay more for housing? exploitation, profit, and risk in rental markets. *American Journal of Sociology*, 124(4):1090–1124.

Diamond, R. and Diamond, W. F. (2024). Racial differences in the total rate of return on owner-occupied housing. Technical report, National Bureau of Economic Research.

Diamond, R. and McQuade, T. (2019). Who wants affordable housing in their backyard? an equilibrium analysis of low-income property development. *Journal of Political Economy*, 127(3):1063–1117.

Diamond, R., McQuade, T., and Qian, F. (2019). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review*, 109(9):3365–94.

Eichholtz, P., Korevaar, M., Lindenthal, T., and Tallec, R. (2021). The total return and risk to residential real estate. *The Review of Financial Studies*, 34(8):3608–3646.

Eiling, E., Giambona, E., Lopez Aliouchkin, R., and Tuijp, P. (2019). Homeowners' risk premia: Evidence from zip code housing returns. *Available at SSRN 3312391*.

Favilukis, J., Mabille, P., and Van Nieuwerburgh, S. (2023). Affordable housing and city welfare. *Review of Economic Studies*, 90:293–330.

Fowler, K. A., Gladden, R. M., Vagi, K. J., Barnes, J., and Frazier, L. (2015). Increase in suicides associated with home eviction and foreclosure during the us housing crisis: findings from 16 national violent death reporting system states, 2005–2010. *American journal of public health*, 105(2):311–316.

Garriga, C., Gete, P., and Tsouderou, A. (2023). The economic effects of real estate investors. *Real Estate Economics*, 51(3):655–685.

Glaeser, E. L. and Gyourko, J. (2003). The impact of zoning on housing affordability. *Economic Policy Review*, 9(2):21–39.

Glaeser, E. L. and Luttmer, E. F. (2003). The misallocation of housing under rent control. *American Economic Review*, 93(4):1027–1046.

Goldsmith-Pinkham, P. and Shue, K. (2023). The gender gap in housing returns. *The Journal of Finance*, 78(2):1097–1145.

Gupta, A., Mittal, V., Peeters, J., and Van Nieuwerburgh, S. (2022). Flattening the curve: Pandemic-induced revaluation of urban real estate. *Journal of Financial Economics*, 146(2):594–636.

Guvenen, F., Ozkan, S., and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660.

Halket, J., Loewenstein, L., and Willen, P. S. (2023). The cross-section of housing returns. Technical report, Working paper.

Hanson, S. (2024). Institutional investors in the market for single-family housing: Where did they come from, where did they go? Working paper, Stanford University Graduate School of Business.

JCHS (2024). America's rental housing. *Joint Center for Housing Studies*.

Kermani, A. and Wong, F. (2021). Racial disparities in housing returns. Technical report, National Bureau of Economic Research.

Kling, J. R., Ludwig, J., and Katz, L. F. (2005). Neighborhood effects on crime for female and male youth: Evidence from a randomized housing voucher experiment. *The Quarterly Journal of Economics*, 120(1):87–130.

Koijen, R., Shah, N., and Van Nieuwerburgh, S. (2025). The commercial real estate ecosystem. Working paper, SSRN. Available at SSRN abstract_id=5120847.

Lawton, J. D. (2012). Tenant purchase as a means of creating and preserving affordable homeownership. *Georgetown Journal on Poverty Law & Policy*, 20:55.

Levy, A. (2022). Housing policy with home-biased landlords: Evidence from french rental markets.

Mabille, P. (2023). The missing homebuyers: Regional heterogeneity and credit contractions. *The Review of Financial Studies*, 36(7):2756–2796.

McCollum, M. and Milcheva, S. (2023). How "bad" is renter protection for institutional investment in multifamily housing? *Journal of Housing Economics*, 59(A).

Moktan, S. (2025). An empirical equilibrium model of the markets for rental and owner-occupied housing. Working paper, London School of Economics. Latest version available online.

Mondragon, J. A. and Wieland, J. (2022). Housing demand and remote work. Working Paper 30041, National Bureau of Economic Research.

Nicholas, T. and Spaenjers, C. (2025). Race, rental yields, and housing decay in manhattan. Working paper, SSRN. Latest version available on SSRN.

Oxley, M., Tang, C., Lizieri, C., Mansley, N., Mekic, D., Haffner, M., and Hoekstra, J. (2015). Prospects for institutional investment in social housing. *IPF, London*.

Samuels, P. R. (2024). Spatial dispersion in returns to rental housing: A decomposition of local rent-price ratios. Working paper, Harvard University.

Sodini, P., Van Nieuwerburgh, S., Vestman, R., and von Lilienfeld-Toal, U. (2023). Identifying the benefits from homeownership: A swedish experiment. *American Economic Review*, 113(12):3173–3212.

Sternlieb, G. (1966). The tenement landlord.

Theodos, B., Meixell, B., and McManus, S. (2023). What we do and don't know about opportunity zones. *Urban Wire*. Accessed: 2025-01-29.

Van den Broeck, K. and Winters, S. (2018). Kosteneffectiviteit en efficiëntie van sociale huisvestingsmaatschappijen en sociale verhuurkantoren. *Steunpunt Wonen*.

A Supplementary Tables and Figures

Table A1: Maintenance Costs Models for NL and BE

	NL			BE		
	(1)	(2)	(3)	(4)	(5)	(6)
Type: Corner house	0.050	−1.332**	−1.330**			
	(0.076)	(0.594)	(0.584)			
Type: Row house	−0.053	−1.417**	−1.418**	−0.101	−0.144	−1.713*
	(0.048)	(0.587)	(0.576)	(0.119)	(0.116)	(0.990)
Type: Semi-detached house	0.092	−1.317**	−1.314**	0.046	0.032	−1.601
	(0.070)	(0.602)	(0.592)	(0.132)	(0.127)	(1.014)
Type: Detached house	0.463***	0.555	0.408	0.266**	0.251**	−1.413
	(0.073)	(0.649)	(0.631)	(0.123)	(0.120)	(1.017)
Type: Apartment w/elevator				0.020	0.021	0.009
				(0.096)	(0.095)	(0.100)
log(m2)	0.512***	0.682***	0.675***	0.482***	0.503***	0.421***
	(0.053)	(0.079)	(0.078)	(0.098)	(0.091)	(0.119)
log(m2) × Detached		−0.298**	−0.297**			
		(0.131)	(0.129)			
log(m2) × Apartment		−0.292**	−0.263**			
		(0.119)	(0.118)			
Building year	−0.001**	−0.001**		−0.002*	−0.003	−0.002*
	(0.001)	(0.001)		(0.001)	(0.002)	(0.001)
Pre-war			0.124***			
			(0.045)			
1970-2009					0.036	
					(0.110)	
Post 2009					0.156	
					(0.194)	
log(EPC)				−0.110	−0.100	−0.168*
				(0.067)	(0.071)	(0.088)
log(EPC) × House						0.142
						(0.101)
Constant	8.276***	8.494***	6.038***	10.450***	12.870***	11.070***
	(1.127)	(1.111)	(0.438)	(2.594)	(4.663)	(2.496)
Observations	147	147	147	49	49	49
R^2	0.764	0.777	0.782	0.676	0.685	0.691
Adjusted R^2	0.754	0.764	0.769	0.621	0.612	0.620

Table A2: Rents, Sales Prices and Costs in BE and NL (in 2023 €)

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D1-D10
Panel A: Belgium											
Rent and Property Value											
Rent (mean, €/month)	616	710	760	800	838	880	927	988	1086	1430	-814
Rent per square meter (€/m²)	12.0	9.1	8.7	8.5	8.4	8.4	8.3	8.3	8.5	9.7	2.4
Property value (mean)	134,542	160,369	181,809	199,595	217,153	238,358	260,366	289,680	333,842	438,546	-304,004
Annual Expenses (for mean-value property)											
Maintenance cost (€/year)	1863	2220	2364	2461	2566	2662	2791	2922	3074	3290	-1427
Turnover cost (€/year)	74	71	71	68	72	71	70	75	80	96	-22
Property tax owner (€/year)	578	581	619	660	700	742	799	874	984	1348	-770
Default cost (€/year)	102	126	115	92	100	109	62	87	99	90	12
Management cost (€/year)	626	722	772	813	851	894	942	1004	1104	1453	-827
Panel B: Netherlands											
Rent and Property Value											
Rent (mean, €/month)	629	747	827	899	991	1058	1152	1274	1472	2008	-1379
Rent per square meter (€/m²)	15.6	13.5	12.3	11.7	11.9	11.8	12.2	13.3	14.9	21.8	-6.2
Property value (mean)	112,075	142,339	165,575	188,681	217,123	246,500	279,650	319,492	386,497	555,500	-443,425
Annual Expenses (for mean-value property)											
Maintenance cost (€/year)	1524	1737	1871	2075	1954	1972	1986	2013	2048	2222	-698
Turnover cost (€/year)	258	285	298	321	347	394	419	479	618	944	-686
Property tax cost (€/year)	381	470	530	566	565	616	643	671	734	889	-508
Default cost (€/year)	123	128	132	132	109	99	84	96	77	111	12
Management cost (€/year)	538	541	546	547	565	567	587	671	734	1055	-517
Alternative scenarios											
Management cost (fixed)	539	530	536	539	537	532	518	510	520	547	-8
Management cost (variable)	379	446	487	516	543	580	624	685	815	1182	-803

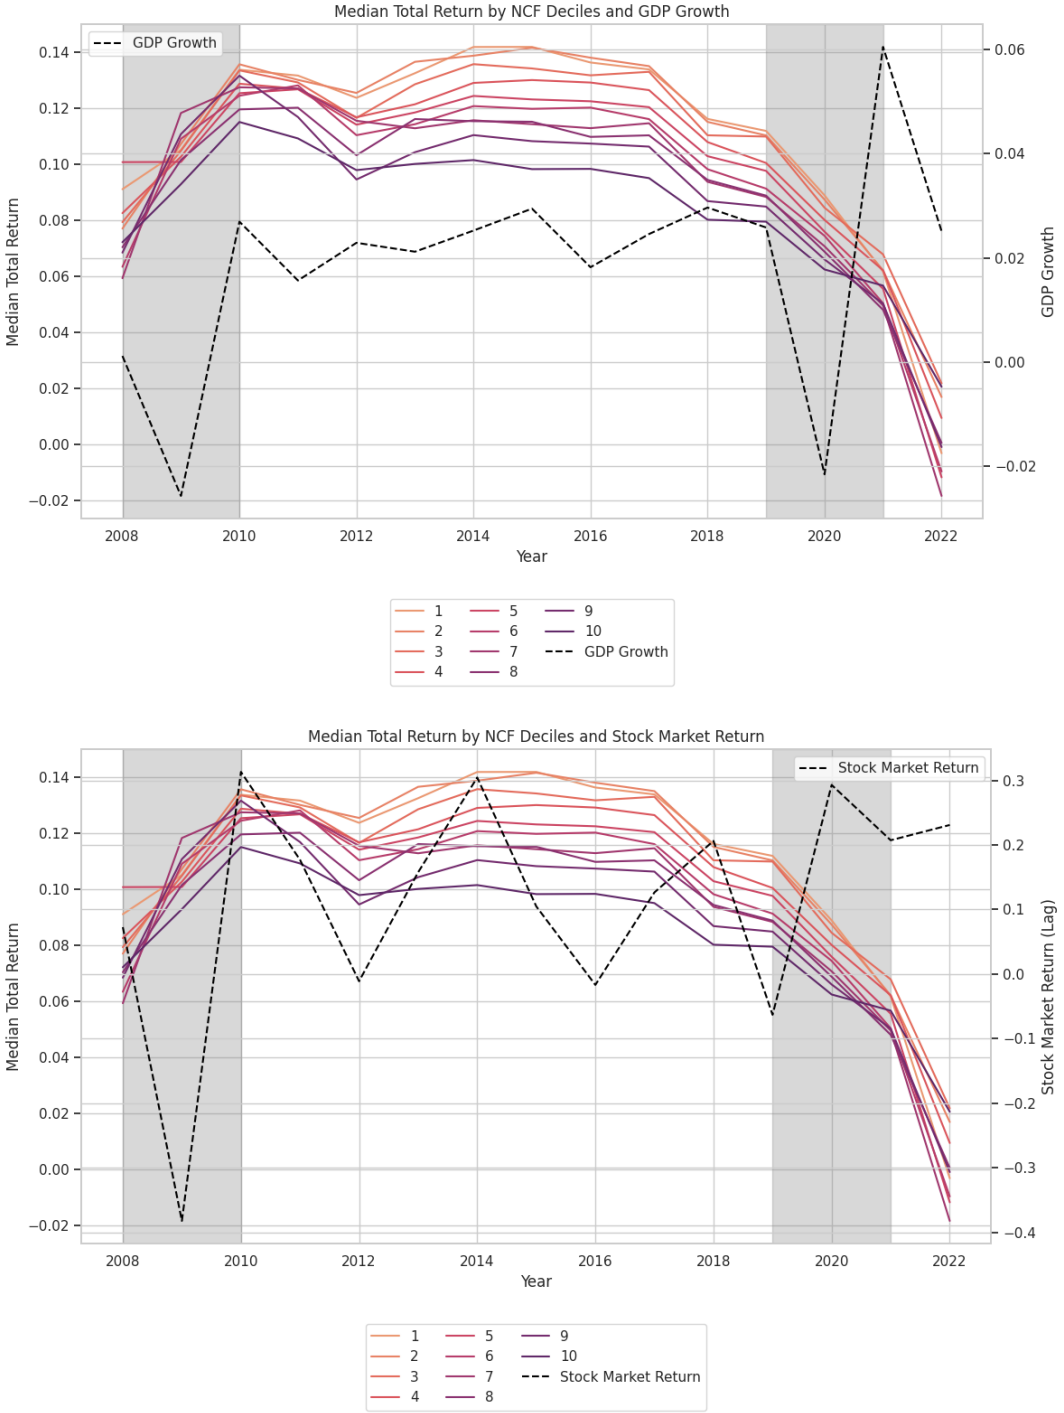
Table A3: Expenses Per Unit and Association Size

	Dependent variable:	
	log(Operating Expenses per Unit)	
	(1)	(2)
< 1,000 units	-0.136^{**}	-0.196^{***}
	(0.056)	(0.059)
1,000-2,500 units	-0.128^{***}	-0.173^{***}
	(0.043)	(0.047)
2,500-5,000 units	-0.133^{***}	-0.172^{***}
	(0.044)	(0.047)
5,000-10,000 units	−0.067	-0.084^{**}
	(0.042)	(0.043)
10,000-25,000 units	–	–
25,000+ units	−0.005	−0.021
	(0.062)	(0.078)
Top-4 city		0.002
		(0.072)
Top-44 city		−0.047
		(0.035)
log(Mean Rent)		0.273
		(0.196)
log(Rent Ceiling)		−0.234
		(0.184)
Constant	7.979^{***}	7.311^{***}
	(0.033)	(1.369)
Observations	260	258
R^2	0.058	0.084

*p<0.1; **p<0.05; ***p<0.01

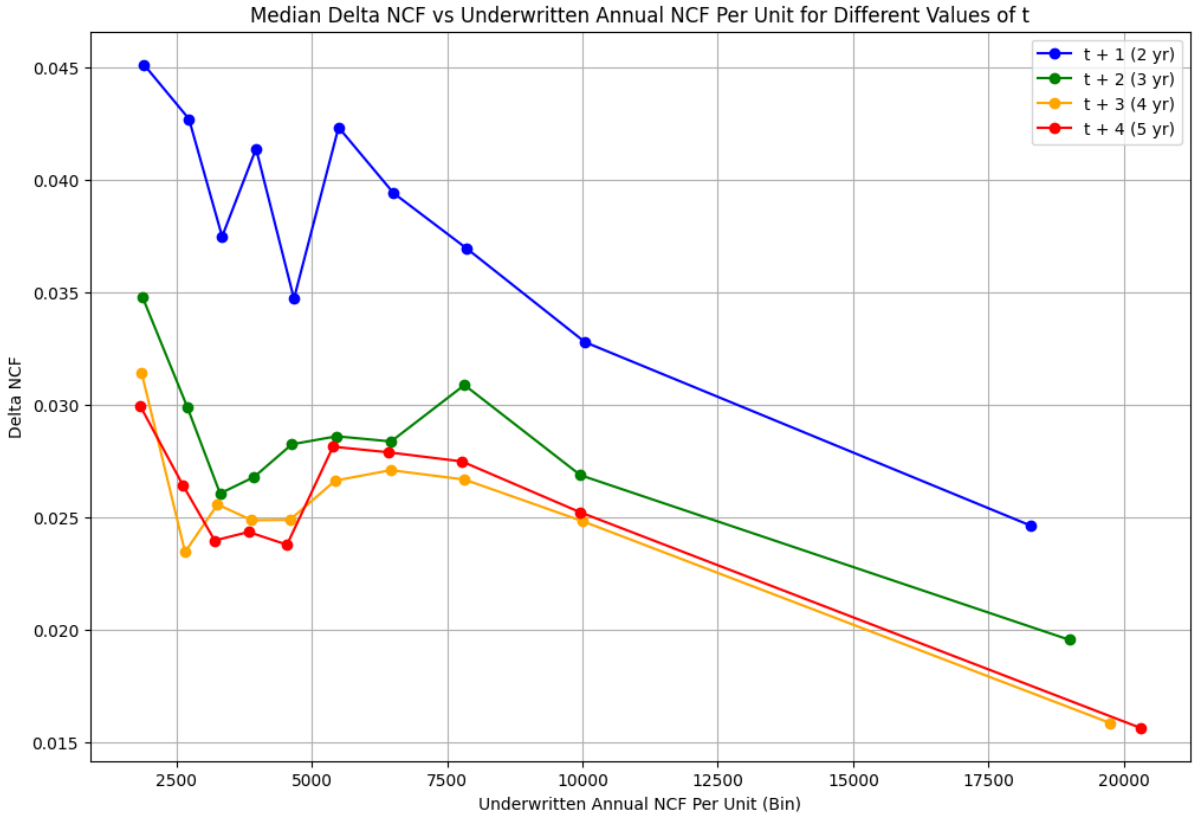
Notes: These data are provided by *Aedes*, the Dutch Association of Housing Associations.

Figure A1: US Rental Housing Returns, GDP Growth, and Stock Returns



Notes: This figure plots the median total housing return by year in each net rent decile (colored lines) alongside GDP growth (top panel) and the annual stock market return (bottom panel). We lag the stock market return by one calendar year since it is a leading indicator of economic activity. The stock return is the real value-weighted return on the S&P 500 index.

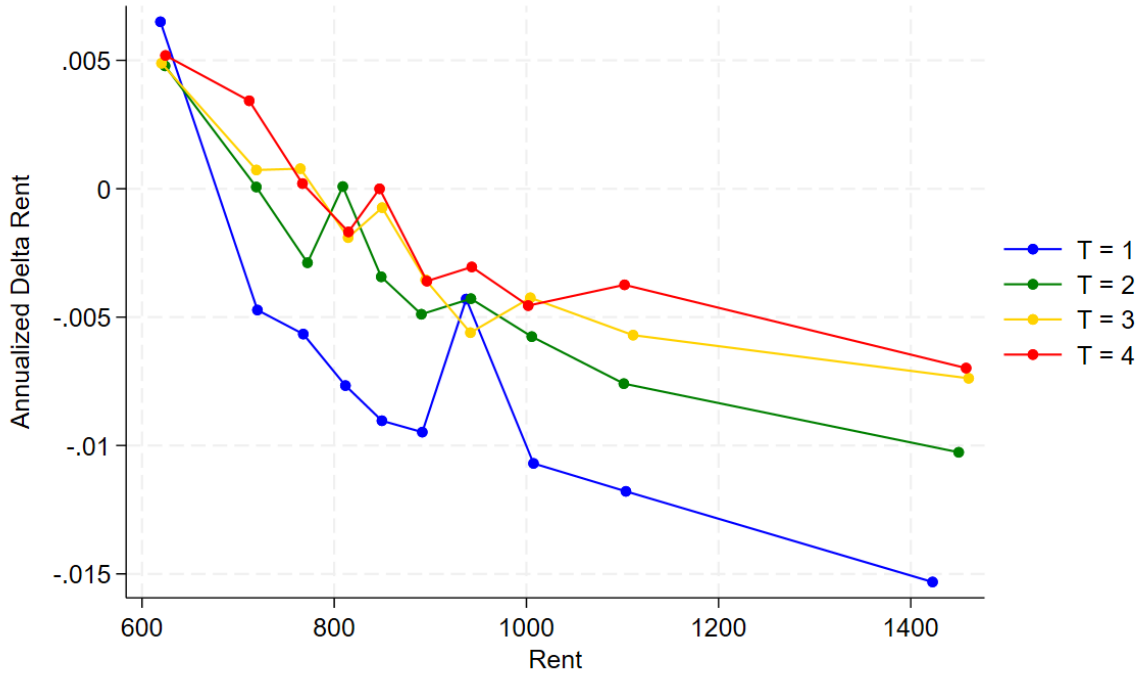
Figure A2: Cash Flow Growth by Rent Decile, US



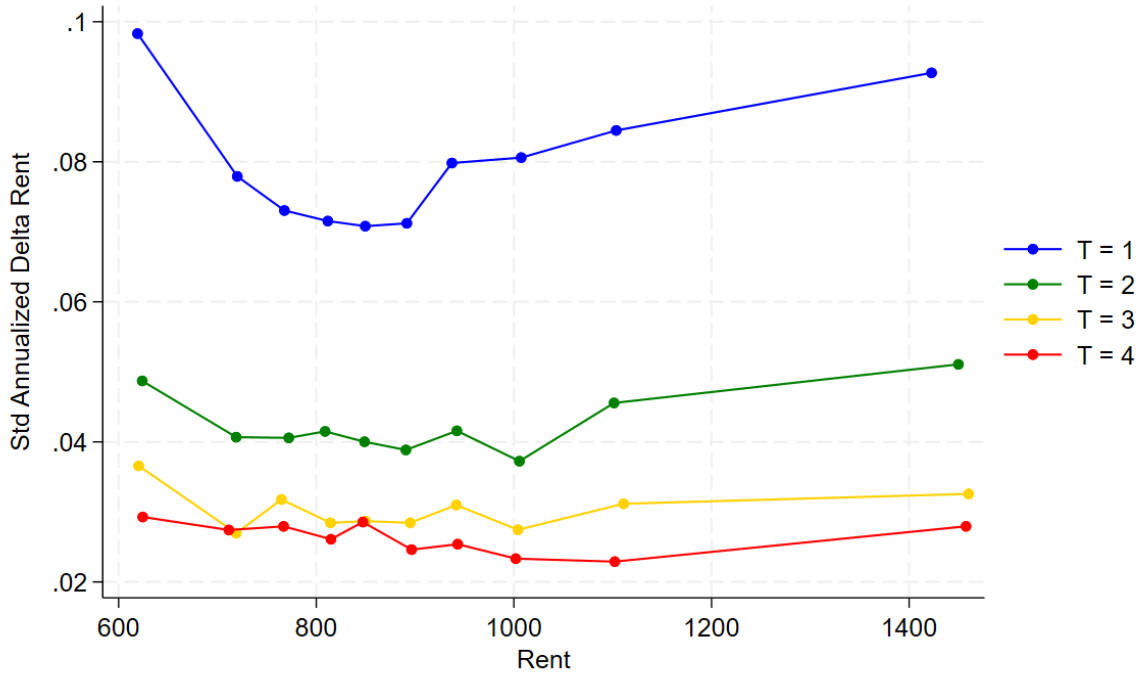
Notes: Bin-Scatter Plot of the cross-sectional median of net cash-flow growth rates over horizon T, $Std(\Delta^i_{t+T})$, with properties sorted by net rent decile. Computed using the Fannie Mae Performance Data sample.

Figure A3: Cash Flow Growth Mean and Volatility by Rent Decile, Belgium

(a) Predicted Cash Flow Growth, by decile

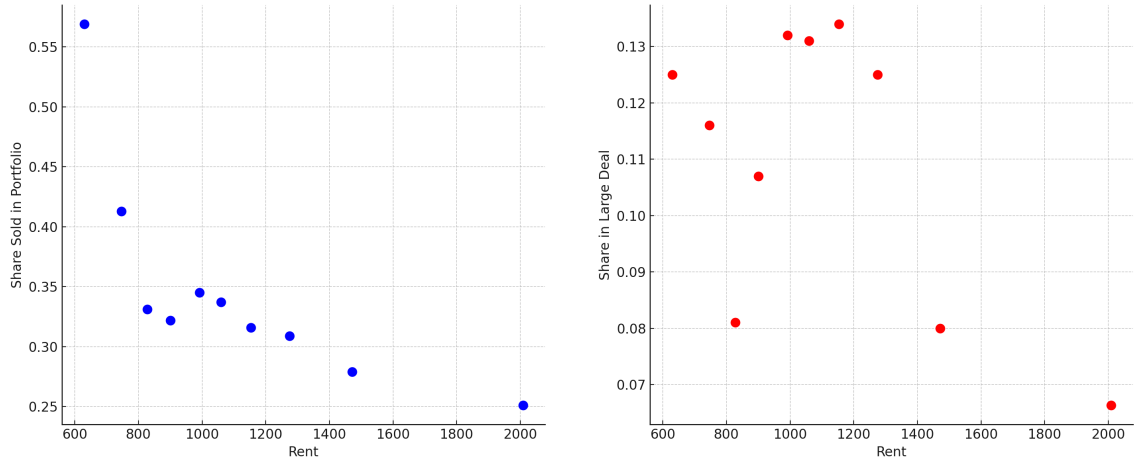


(b) Volatility of Cash Flow Growth, by decile



Notes: Using repeated rental contract from the Belgian rent registry, we calculate the annualized realized rent growth between time t and $t+T$, where $t+T$ corresponds to the time when a new rental contract is registered. T varies from one to four years. We then compute the cross-sectional average (panel a) and standard deviation (panel b) of these growth rates in each rent decile.

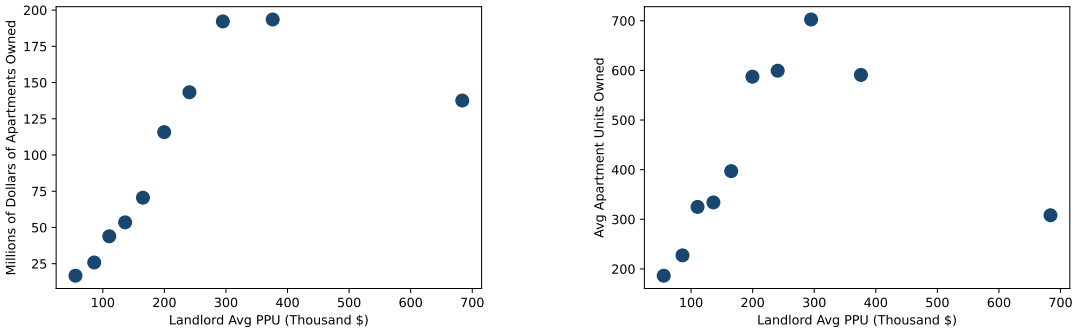
Figure A4: Share of NL private rental units sold in large or portfolio deals



(a) Share of units sold in portfolio (if sold) (b) Share of units sold in 10m+ deal (if sold)

Notes: These figures use Dutch Land Registry data to compute the share of rental units that is sold in portfolio (left) or sold as part of a deal of 10 million euros or more (right), conditional on the unit being sold in a given year. Consistent with frictions between rental units and owner-occupied units, most private rental units in low tiers are sold in a portfolio. Despite low-tier properties being cheaper per unit than high-tier properties, we do not find that low-tier units are less likely to be sold as part of very large real estate deals. This suggests investors do not face significant difficulty in aggregating property portfolios.

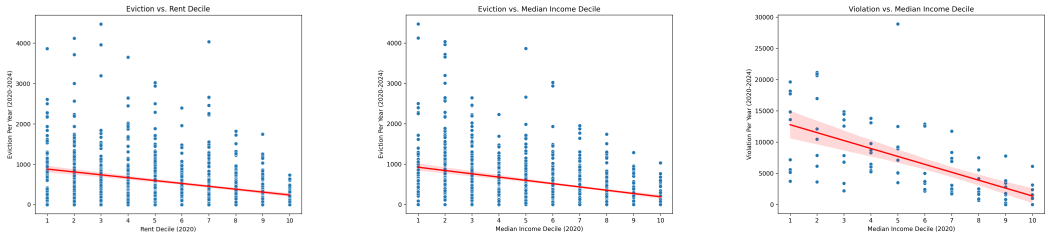
Figure A5: Portfolios of U.S. Multifamily Investors by Rent Decile, US



(a) Value of Multifamily Portfolio (b) Size of Multifamily Portfolio

Notes: These figures use RCA transactions data to construct multifamily portfolio sizes of landlords at the end of 2023 in terms of (a) dollar value and (b) number of units. We then sort landlords into deciles by average property value per unit. The data only includes transactions above $2.5M and is hence biased toward investors with larger holdings. Despite this, we observe a remarkably similar pattern as in the Netherlands (Figure 6), with the largest investors concentrating in D6 or higher in both size and value.

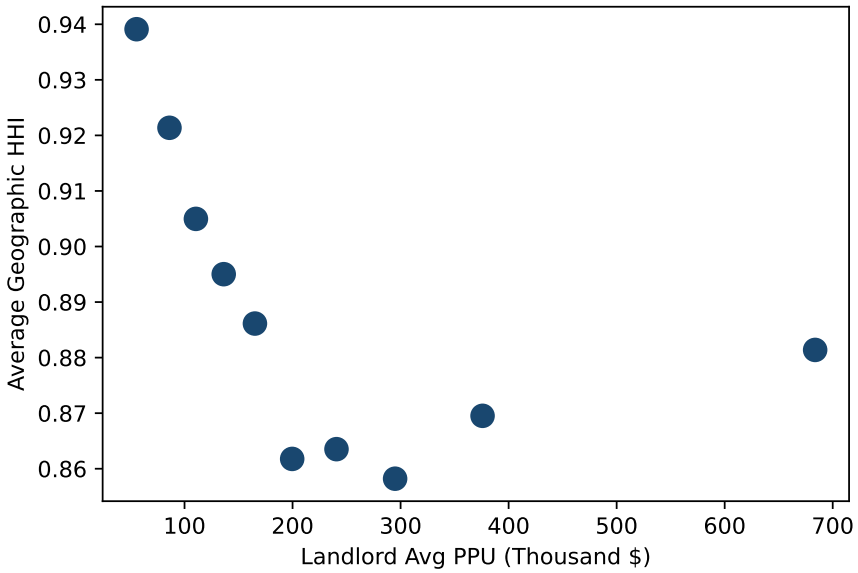
Figure A6: Evictions and Violations by Rent and Income Decile, US



(a) Evictions by Rent (b) Evictions by Income (c) Violations by Income

Notes: We use US data from the Princeton Eviction Lab on average annual evictions in 2020–24 by ZIP code, match these ZIP codes to properties in our database, and sort these properties by the rent deciles they belong to. This results in a cross-section of 923 ZIP codes with non-missing rent and eviction data (blue dots). We observe a strong negative relationship between evictions and rent levels (panel a). In panel (b), we sort properties by ZIP-code level median income decile and observe a similar pattern. In panel (c), we use data from NYCOpenrecords on the number of Department of Buildings violations per year for New York City properties and link this to the income decile of each ZIP code (89 ZIP codes). For NYC, we cannot reliably match to rent deciles in our data, as data coverage per decile is limited.

Figure A7: Geographic HHI by Landlord and Decile, USA (RCA Data)



Notes: We use data from MSCI Real Capital Analytics (RCA) to reconstruct the portfolio sizes of landlords at the end of 2023. We then compute the geographic Herfindahl index (HHI) of their units, where each U.S. state is treated as one region, and sort landlords into deciles by their average property value per unit. RCA includes all apartment transactions above $2.5 million. We sort properties into 60 markets (geographies), making small modifications to the geography variable "RCA market" provided by RCA. The market roughly corresponds to a city for the largest cities, and to collections of smaller cities in a wider region for the reamining locations.

B Data: United States

B.1 Data Description

We extract property information from mortgage loans made against multifamily rental properties. Fannie Mae and Freddie Mac occupy a dominant position in the multifamily mortgage market. Fannie and Freddie accounted for 58% of all new multifamily mortgage originations in the United States in 2023. They have $1.0 trillion in multifamily loans outstanding, out of a total $2.09 trillion market (Mortgage Bankers Association). Fannie and Freddie each publish loan-level data that contain detailed loan and property characteristics. The Fannie Mae data set contains 66,487 loans, spanning the period January 2001 until May 2024. The Freddie Mac data set contains 57,209 loans, spanning the same period. We also collect data from Ginnie Mae, the securitizer of multifamily mortgages originated by the Federal Housing Administration and the Veterans Administration. We supplement the agency loans with 60,100 multifamily loans that are originated and sold into private-label commercial mortgage-backed securities (CMBS) and commercial real estate collateralized loan obligations (CRE CLOs). These data come from the securitization trusts, as compiled by the data aggregator CredIQ. The multifamily properties we are missing are those that do not have mortgages (unusual), and those whose loans are held on the balance sheets of banks, insurance companies, and mortgage REITs in the form of whole loans (not securitized).[29]

For each loan, we use two key variables: the loan-to-value (LTV) ratio and the debt service coverage (DSCR) ratio, from which we derive the NOI. Note that how we do this depends on the type of loan, since the debt service will vary depending on the loan type. For a fully amortizing mortgage, it is a constant periodic payment consisting of principal and interest that pays down the loan over its maturity. For the (common) balloon mortgage, the term is shorter than the maturity and only a modest amount of principal has been paid down when the loan matures. For an interest only mortgage, the debt service contains only interest payments. There are several intermediate cases. Since we know the type of loan and the interest rate, we can compute the debt service accurately for each loan. We compute the NOI as the DSCR times the debt service. Given that multifamily buildings have multiple housing units, we express both the price and the NOI per unit. We express all historical prices and NOI in 2023 real dollars. Since the Ginnie Mae data does not contain the LTV or DSCR fields, we cannot use those loans and associated properties in the analysis.

[29] At the end of 2023, commercial banks held $612 billion (29 percent) of multifamily mortgages, life insurance companies held $235 billion (11 percent), state and local governments held $116 billion (6 percent), and CMBS, CDO and other ABS issues held $67 billion (3 percent).

B.2 Data Processing

Table B1 shows the data filters we employ to arrive at the final USA sample. In addition to the loss of Ginnie Mae data, the largest loss of data is for loans that are missing either LTV, DSCR, or loan amortization information, which prevents us from computing price and/or NOI. We further trim the data for implausible values of LTV, NOI per unit, and cap rates. Our final sample contains 105,923 mortgages on 80,608 unique multifamily properties with a combined 11,759,738 housing units.

Table B1: Final Sample Selection

Filter	**Sample Size**	**% Dropped**	**# Dropped**
Full sample of loan observations	159,148		
Eliminating Ginnie Mae loans	146,069	8.22%	13,079
Eliminating loans if LTV or NOI are missing	109,842	24.80%	36,227
Eliminating LTVs above 100%	109,550	0.27%	292
Eliminating real NOI per Unit below $1,200 and above $48,000	106,677	2.62%	2,873
Eliminating cap rates below 1% and above 18%	105,923	0.71%	754
Final Sample of Loans	**105,923**		
Number of Properties	**80,608**		
Number of Units	**11,759,738**		

Figure B1 shows the make-up of the final sample in terms of deal composition: Fannie Mae loans, Freddie Mac loans, and loans in private-label securitizations (CMBS or CRE CLO).

Figure B1: Deal Type Distribution

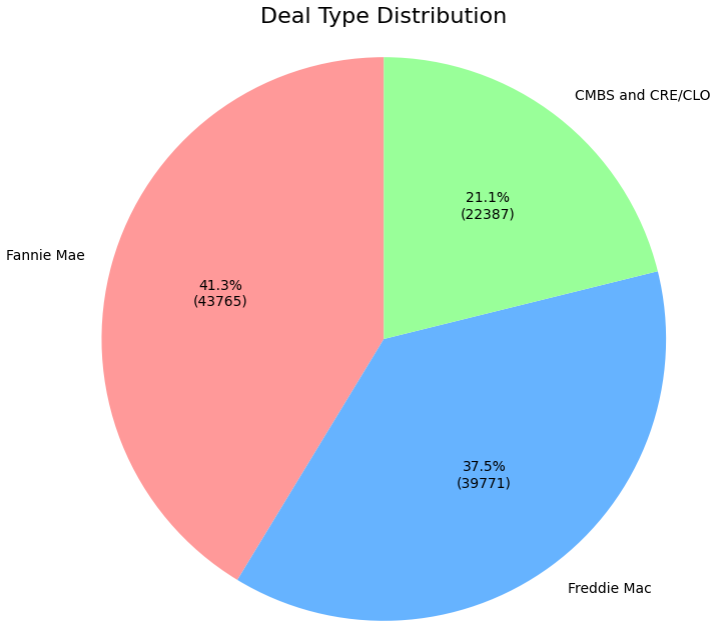
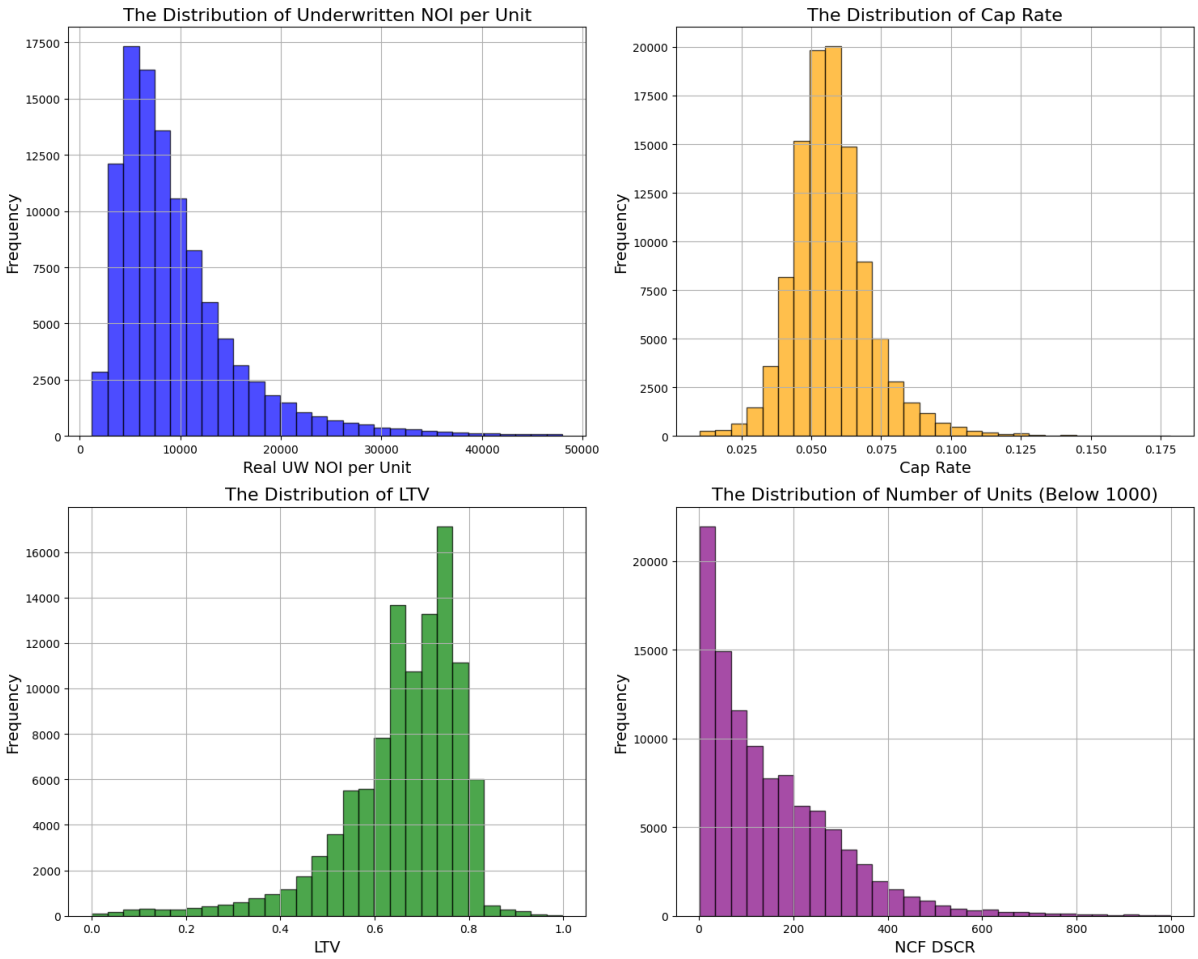


Figure B2 shows the distributions of annual NOI per housing unit, cap rate, LTV ratio, and the number of units per property in the final sample. Net rents earned by

landlords on US apartments are mostly between $500 and $1000 per month ($6,000 and $12,000 per year) per unit, with a long right tail. Most net rental yields are between 3% and 8%. Most complexes have fewer than 250 units, and most LTVs at mortgage origination are between 60% and 80%.

Figure B2: Distribution of Key Variables in Final Sample



B.3 Capital Gain Returns

B.3.1 Hedonic Approach

We use a hedonic method for computing the capital gain yield, comparing the results to a repeat-sales method below. We estimate the following model for the log property value (in real dollars) per housing unit, measured at the time of origination:

$$\begin{aligned} \ln P_{i,t} &= \alpha_s + \alpha_t + \alpha_d + \beta_1 \ln Y_{i,t} + \beta_2 \text{age}_{i,t} + \beta_3 \text{age}_{i,t}^2 \\ &+ \beta_4 \text{Renovated}_i + \beta_5 \ln \text{NOI}_{i,t} + \beta_6 \ln S_i + \varepsilon_{i,t}, \end{aligned} \tag{4}$$

where $Y_{i,t}$ is the median household income in year t in the ZIP code where property i is located, expressed in real dollars, age is the age of the property, the Renovated flag

indicates whether the property has been renovated, NOI is the underwritten NOI for the property at the time of mortgage origination in real dollars, and $S_{i,t}$ is the number of units in the property. Table B2 shows the results form the hedonic valuation model. The omitted deal type is Fannie Mae loans. The model explains 89% of the cross-sectional variation in log apartment unit prices.

Table B2: Hedonic Valuation Model for Log Price Per Unit

Variable	**Coefficient**	**Std. Err.**	**t-Statistic**
Constant	1.2751***	0.031	40.582
Log Median Income (ZIP-Level)	0.1032***	0.002	42.376
Building Age	-0.0018***	8.5e-05	-21.151
Building Age$_t^2$	1.683e-05***	7.26e-07	23.164
Renovated	-0.0135***	0.002	-6.342
Log UW NOI per Unit	1.0051***	0.002	541.058
Log Number of Units	0.0054***	0.001	6.457
Deal Type: Conduit	0.0441***	0.003	13.694
Deal Type: CRE CLO	0.2062***	0.005	38.888
Deal Type: Freddie Mac	0.0247***	0.002	12.520
State Fixed Effects (FE)	Y		
Origination Year FE	Y		
R-squared	0.894		
Number of Observations	85,014		

We use the estimated coefficients of the model in (4) to construct a valuation of property i at time $t+T$ (in real dollars) as follows:

$$\begin{aligned} \ln P_{i,t+T} &= \widehat{\alpha}_s + \widehat{\alpha}_{t+T} + \widehat{\alpha}_d + \widehat{\beta}_1 \ln Y_{i,t+T} + \widehat{\beta}_2 \text{age}_{i,t+T} + \widehat{\beta}_3 \text{age}^2_{i,t+T} \\ &+ \widehat{\beta}_4 \text{if_renov}_i + \widehat{\beta}_5 \ln \widehat{\text{NOI}}_{i,t+T} + \widehat{\beta}_6 \ln S_i + A\varepsilon_{i,t} \end{aligned} \tag{5}$$

Note that we include a scaled version of the time-t residual $A\varepsilon_{i,t}$ in the time-$t+T$ valuation. The idea is that this residual could pick up omitted property attributes that affect the valuation, e.g., exact location, quality of construction, etc. However, the residual could instead reflect mispricing (including good or bad negotiation skills) at time t, which is irrelevant for the valuation at time $t+T$. We use three values for $A \in \{0, a, 1\}$ to allow for both extremes and one intermediate case. When $A = 1$, then $\varepsilon_{i,t}$ only picks up omitted characteristics. When $A = 0$, then $\varepsilon_{i,t}$ only picks up mispricing. When $A = a$, the residual is driven by a mix of omitted characteristics and initial mispricing. We will choose $a = 0.5$, a value for which the distribution of hedonic capital gain yields best fits the distribution of repeat-sales capital gain yields, as shown below.

The capital gain yield is then computed using the original appraisal value at t and

the predicted property value at time $t+T$:

$$\begin{aligned}\text{Cap Gain} &= \frac{1}{T}\left(\ln P_{i,t+T} - \ln P_{i,t}\right) \\ &= \frac{1}{T}\Big(\widehat{\beta}_1 \cdot (\ln Y_{i,t+T} - \ln Y_{i,t}) + \widehat{\beta}_2 \cdot T + \widehat{\beta}_3 \left(age^2_{i,t+T} - age^2_{i,t}\right) \\ &\quad + \widehat{\beta}_5 \cdot (\ln \widehat{NOI}_{i,t+T} - \ln NOI_{i,t}) + \widehat{\alpha}_{t+T} - \widehat{\alpha}_T + (A-1)\varepsilon_{it}\Big) \qquad (6)\end{aligned}$$

Capital gains in this model arise from five sources. The first four reasons are: (i) the building ages, (ii) the NOI of the building changes (in real dollars), (iii) the transaction year may have different macro-economic effects (origination year fixed effect is different), and (iv) the local economic conditions may change as measured by the local median income (in real dollars). The geographic fixed effect, deal fixed effect, renovation flag, and building size all drop out of the return. When $A = 1$, the ε_{it} term also drops out of the capital gain return: the omitted hedonics affect the log price at t and at $t+T$ equally and their effect disappears from the difference. However, when $A \neq 1$ because ε partly or wholly reflects mispricing, the term $(A-1)\varepsilon_{it}$ contributes to the capital gain, the fifth source of capital gains.

Note that the capital gain return in (6) includes the term $\log \widehat{NOI}_{i,t+T} - \log NOI_{i,t}$. Since we generally only observe the NOI at the time of loan underwriting and not thereafter, we do not observe $NOI_{i,t+T}$. However, for the Fannie Mae data, we do observe NOI for several years after loan origination from the loan performance files. This allows us to estimate an auxiliary model for the NOI growth rate between t and $t+T$. We propose the following model for NOI growth:

$$\begin{aligned}\ln NOI_{i,t+T} - \ln NOI_{i,t} &= \alpha + \alpha_s + \alpha_t + \beta_1(\ln Y_{i,t+T} - \ln Y_{i,t}) + \beta_2 age_{i,t} + \beta_3 age^2_{i,t} \\ &+ \beta_4 LTV_i + \beta_5 \ln S_i + \beta_6 \ln NOI_{i,t0} + \beta_7 Occup_{i,t0} + e_{i,t+1} \qquad (7)\end{aligned}$$

where α_s is a geographic fixed effect, α_t is a time (origination year) fixed effect. The other independent variables are the log of the ZIP-code median income growth, property age and age squared, the occupancy rate, and the log size of the property. We also allow the growth rate of NOI to depend on the log level of the property's underwritten NOI, which we label $NOI_{i,t0}$. Note that we do not use the underwritten NOI on the left-hand side of the NOI growth model since we do not want conservatively-underwritten properties to show up as having a very high NOI growth rate in the first year. We use the underwritten NOI on the right-hand side as that is the only property-level cash-flow variable that we have available for the *non-Fannie Mae* loan sample, and the variable is statistically important for NOI growth prediction.

We use $T = 5$ as our default horizon. Given that the Fannie Mae NOI data goes until 2023, we can compute NOI growth rates with $T = 5$ years for all transactions that occur in 2017 or earlier ($t0$ is 2017 or earlier, t is 2018 or earlier, $t+T$ is 2023 or earlier).

To maximize the number of observations in the auxiliary NOI growth prediction model, we compute NOI growth rates over shorter horizons for more recent transactions. For transactions in 2018, we set $T = 4$. For transactions in 2019, we set $T = 3$, etc. Table B3 reports the coefficient estimates. While predicting five-year NOI growth is difficult (R^2 is 6.9%), all of the predictors are statistically significant.

Table B3: Auxiliary Model For Annual NOI Growth

Variable	**coef**	**std err**	**t-stat**
Constant	0.2893***	0.025	11.584
Log Median Income Growth	0.1028***	0.012	8.840
Building Age_t	0.0016***	0.000	14.120
Building Age_t^2	-1.054e-05***	1.07e-06	-9.855
LTV	-0.0478***	0.007	-6.697
Log Number of Units	0.0106***	0.001	9.058
Log UW NOI per Unit	-0.0229***	0.002	-12.227
Occupancy Rate	-0.0013***	0.000	-8.774
State FE	Y		
Year FE	Y		
R-squared	0.069		
Number of Observations	53,634		

With the estimated coefficients in hand, we can form the predicted log change in NOI between t and $t+T$ using right-hand side variables measured at the time of underwriting as well as contemporaneous median income growth in the ZIP code measured between t and $t+T$:

$$\begin{aligned} \ln \widehat{NOI}_{i,t+T} - \ln NOI_{i,t} &= \widehat{\alpha} + \widehat{\alpha}_s + \widehat{\alpha}_t + \widehat{\beta}_1(\ln Y_{i,t+\tau+1} - \ln Y_{i,t+\tau}) + \widehat{\beta}_2 age_{i,t0} + \widehat{\beta}_3 age_{i,t0}^2 \\ &+ \widehat{\beta}_4 LTV_{i,t0} + \widehat{\beta}_5 \ln S_{i,t0} + \widehat{\beta}_6 \ln NOI_{i,t0} + \widehat{\beta}_7 Occup_{i,t0}. \end{aligned} \tag{8}$$

We can now return to the capital gain return measurement. For the Fannie Mae loans, we use the realized NOI growth rate in the capital gain return formula (6) since we observe it. For the non-Fannie loans, we use the imputed NOI growth rate (8). The first three columns of Table B4 report summary statistics of the capital gain yield for the full sample, obtained using the hedonic model under three assumptions on A, i.e., on the theory for the residual. In column (1), we set $A = 0$, which is the all-mispricing model, in column (2), we set $A = 0.5$, our benchmark model, and in column (3), we set $A = 1$, we is the all-omitted characteristics model.

Columns (4) to (6) report the moments of the total return, computed as the sum of the net rental yield plus the capital gain yield. We see that both the capital gain yield and the total return have broadly similar distributions under all three theories for the residual (values for A).

Table B4: Summary Statistics for Capital Gain and Total Returns

Statistic	**Capital Gain**			**Total Return**		
	A=0	A=0.5	A=1	A=0	A=0.5	A=1
Obs	81,239	81,239	81,239	81,239	81,239	81,239
Mean	3.66%	3.76%	3.86%	9.14%	9.24%	9.34%
Std	11.05%	8.80%	8.47%	11.78%	9.34%	8.63%
5%	-12.82%	-9.57%	-7.53%	-8.79%	-5.27%	-2.69%
25%	0.88%	1.63%	1.90%	5.80%	6.61%	7.11%
50%	4.77%	4.65%	4.55%	10.35%	10.21%	10.17%
75%	7.94%	6.98%	6.32%	14.10%	13.09%	12.30%
95%	15.48%	13.27%	13.61%	22.14%	19.54%	18.93%

B.3.2 Comparison Of Hedonic and Repeat-Sales Approaches

Starting from the final sample of 105,923 mortgage records, we identify 13,328 transaction pairs with two mortgages at different time points in time, based on matching by street address, city, state, and zip code. These, 13,328 transaction pairs represent 10,983 different properties since 1,737 properties have 3 mortgages and 263 properties have 4 mortgages in our data set. For the repeat-sales analysis we trim annualized capital gains at -5% and +24%, resulting in a final sample of 12,631 repeat-sales capital gain yield observations.

Each mortgage in the dataset includes a valuation at the time of underwriting. Properties with subsequent mortgages have a new valuation recorded at a later date. We use the first valuation at t and the second valuation at $t+T$ on the same property to construct the annualized capital gain return between t and $t+T$:

$$\text{Capital Gain} = \frac{1}{T}\left(\log P_{i,t+T} - \log P_{i,t}\right)$$

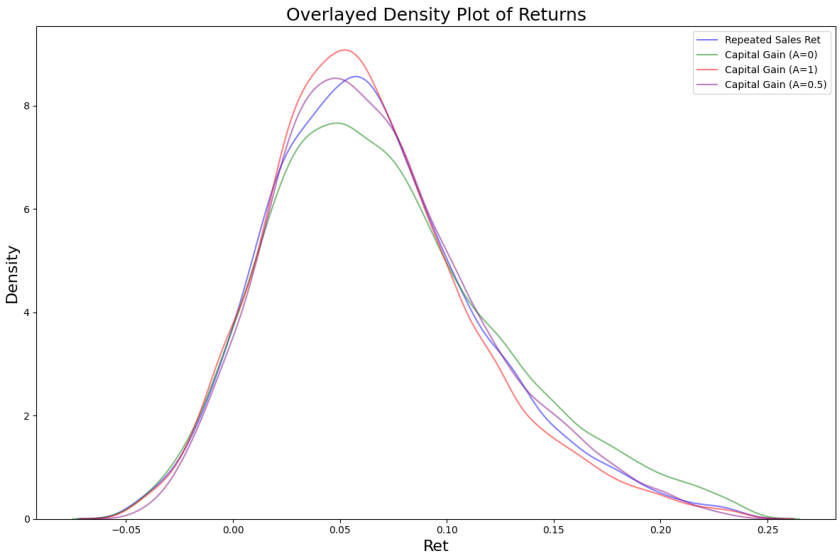
Table B5 and Figure B3 compare the capital gain yields obtained from the repeat-sales approach to those obtained from the hedonic approach for those observations for which we have a repeat-sales return. For this sample, the repeat-sale capital gain yield distribution lies in between the hedonic distributions for $A = 0$ and $A = 1$. Based on this evidence, we consider an intermediate case for the hedonic capital gain yield that sets $A = 0.5$ so that the hedonic model and the repeat-sales model have approximately the same mean capital gain yield as well as a similar dispersion.

Finally, Figure B4 plots the hedonic capital gain return and the repeat-sale capital gain return for the overlapping sample by NOI/decile. It shows that our main capital gains return (the case $A = 0.5$) is very close to the repeat-sales capital gain return along the NOI/unit distribution. The same is true for the hedonic capital gain returns for the cases $A = 1$ and $A = 0$. To the extent that hedonic capital gain returns may miss renovations that are captured by repeat-sales capital gain returns, this evidence suggests

Table B5: Summary Statistics for Capital Gain Yields

	Repeat-Sales Approach	**Hedonic Approach** **A=0**	**A=0.5**	**A=1**
Mean	6.83%	7.44%	7.06%	6.68%
Std	5.31%	7.25%	6.54%	6.61%
5%	-0.80%	-1.34%	-0.97%	-1.21%
25%	3.06%	3.04%	3.12%	2.93%
50%	6.17%	6.50%	6.24%	5.89%
75%	9.88%	10.83%	10.06%	9.45%
95%	17.00%	19.88%	18.13%	17.32%
Obs	12,631	12,613	12,613	12,613

Figure B3: Comparing Capital Gain Yield Distributions



that renovations do not differentially affect the different segments of the market, and hence are not an important concern.

B.3.3 Capital Gain Yield and Total Return in Full Sample

Returning to the main relationship of interest, Figure B5 plots the relationship between net rent and the capital gain yield in the full sample. The relationship is strongly downward sloping for all three values of A. In the main text, we use $A = 0.5$.

Given that both the net rental yield and the capital gain yield are downward sloping in net rent, the total return is even more strongly downward sloping in net rent, as shown in Figure B6. In the main text, we use $A = 0.5$.

Figure B4: Repeat Sales Return and Hedonic Capital Gain by NOI Decile

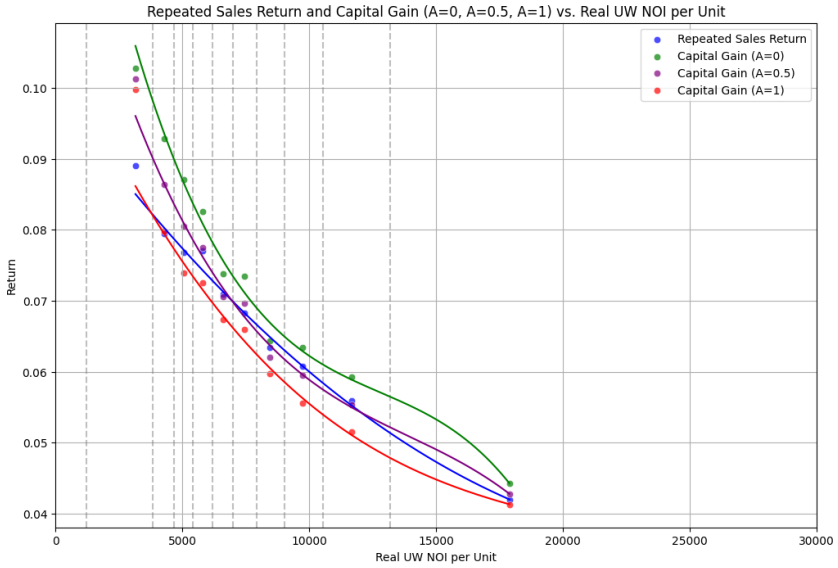


Figure B5: Relationship Between Net Rent and Capital Gain Return

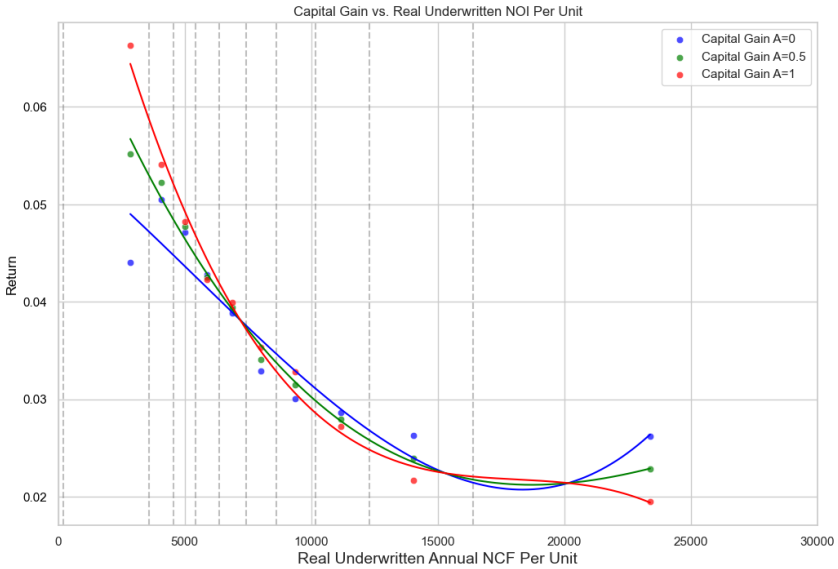
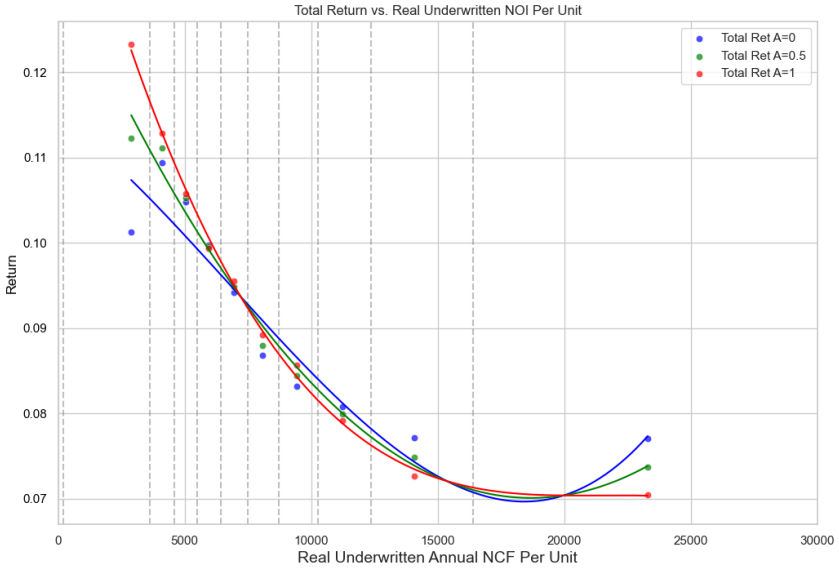


Figure B6: Relationship Between Net Rent and Total Return



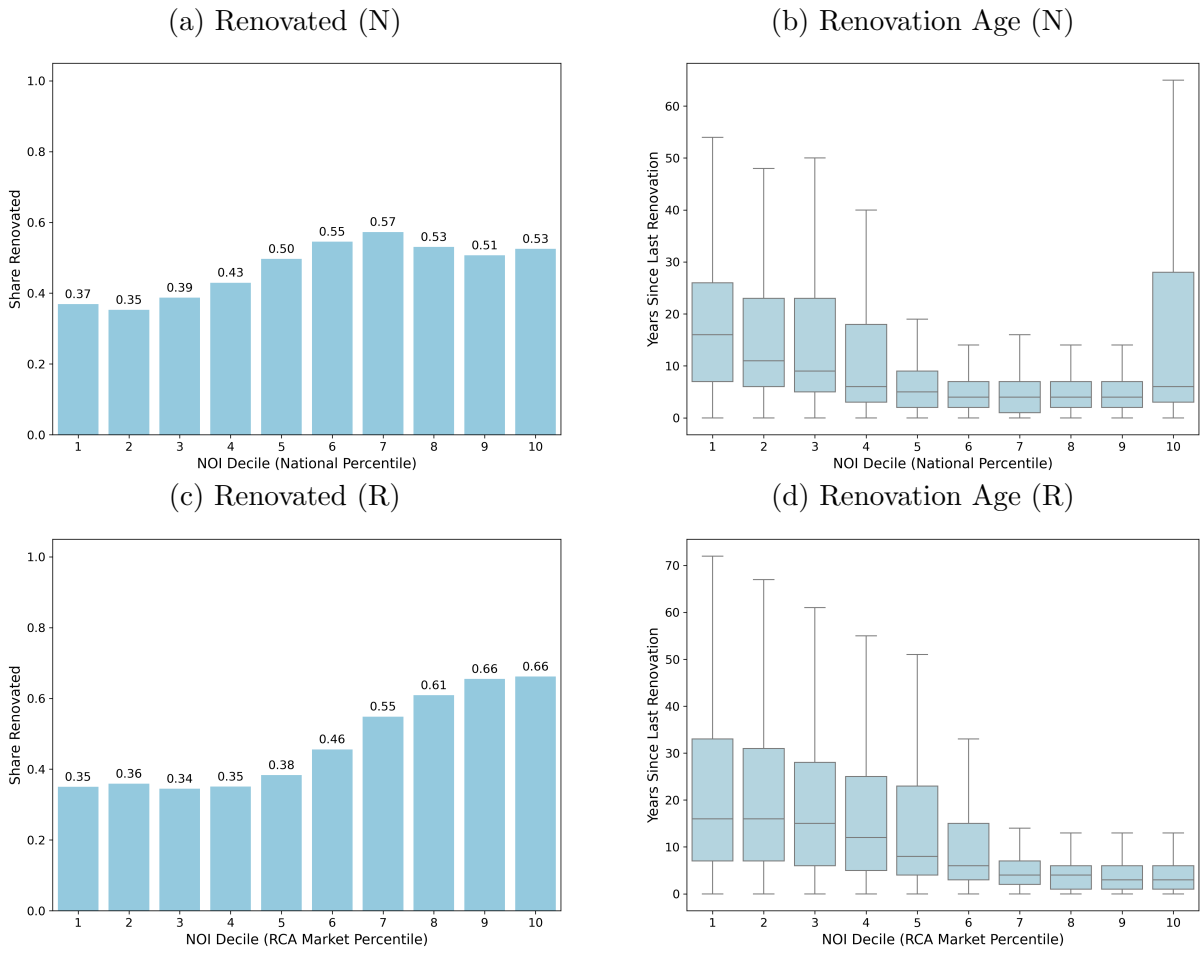
B.4 Renovation Expenditures

We conduct two exercises to assess the importance of renovation expenditures on the downward sloping relationship between net rental yields, capital gain returns, and total returns on the one hand and NOI/unit on the other hand.

First, repeat-sales capital gains are more likely to reflect incurred renovation expenditures than capital gains obtained from a hedonic valuation model since there are usually insufficiently rich hedonics to capture renovations. If low-rent properties were renovated more often, the difference between the repeat sales capital gain and the hedonic capital gain would be higher in the lower- than in the higher-rent tiers. Figure B4 shows that this is not the case.

Second, there is information on the year of renovation in the MSCI Real Capital Analytics data set, described below in Appendix B.7. For each apartment property that transacts, we observe the year of construction and the year of the most recent renovation. This allows us, in the last year of the sample, 2023, to determine whether a property has ever been renovated, *Renovated*, and how long ago the most recent renovation took place conditional on a renovation, *RenovationAge*. In order to plot these two renovation measures by NOI/unit, we have to assign each property in the RCA data to a NOI/unit decile. Appendix B.7.3 describes how we do this. The left panels of Figure B7 show the average value for *Renovated* across rent deciles. The right panels plot the distribution of the variable *RenovationAge* in each decile. The top panels sort properties into deciles by ranking their NOI/unit nationally whereas the bottom panels rank properties on their NOI/unit within each of 60 RCA markets. The results show the low-tier properties are substantially less likely to have been renovated, and that the time since the last renovation is substantially longer.

Figure B7: Renovations By Rent Deciles: U.S.



Notes: Sample is apartment properties that transacted between 2001 and 2023 in the MSCI Real Capital Analytics database. At the end of the sample, in December 2023, we compute whether a property has been renovated (*Renovated*) and how long ago in years the last renovation was (*Renovation Age*). Each property is assigned an NOI/unit as Appendix B.7.3. The top panels sort properties into NOI/unit deciles nationally (N), whereas the bottom panels sort properties into NOI/unit deciles within each of 60 RCA markets (R).

B.5 Borrowing Cost Differences Across Rent Deciles

This appendix uses mortgage interest rates, LTV, DSCR, property location fixed effects, and year fixed effects to investigate whether the cost of borrowing is higher for properties in the low-rent than in the high-rent deciles.

We assign mortgage loans into three groups based on their NOI/unit (expressed in 2023 dollars). We define loans with the highest 20% NOI as NOI^{High} loans, those with the lowest 20% NOI as NOI^{Low} loans, and the rest as NOI^{Mid} loans.

Table B6 shows the regression results. In column (1) we regress mortgage interest rates on NOI group indicators NOI^{High} and NOI^{Low}, omitting NOI^{Mid} as the baseline group. We find that the mortgage rates of high (low) NOI loans are 24 basis points lower (21 basis points higher) than the rates of NOI^{Mid} loans. The mortgage rate gap between low and high NOI properties is 45 basis points. In column (2), we include year fixed effects,

which narrows the gap to 31 basis points. Column (3) shows that the gap narrows further to 27 basis points after we control for the standard mortgage underwriting characteristics of loan-to-value and debt service coverage ratios. Including geographic fixed effects in columns (4) and (5) further narrows the gap. In the last column, the mortgage rate spread between low- and high-NOI/unit properties is a mere 9 basis points.

This analysis makes clear that financing cost differences cannot account for the higher (unlevered) returns on low NOI/unit properties. The raw interest rate differences are much smaller than the total return differences, and controlling for standard underwriting variables and basic fixed effects eliminates the cost of capital differences almost entirely.

Table B6: Mortgage Rate Model for NOI Groups

	(1)	(2)	(3)	(4)	(5)
NOI^{High}	-0.0024***	-0.0011***	-0.0012***	-0.0004***	0.0001
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
NOI^{Low}	0.0021***	0.0020***	0.0015***	0.0016***	0.0008***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
LTV			-0.0112***		-0.0142***
			(0.0003)		(0.0003)
DSCR			-0.0057***		-0.0058***
			(0.0001)		(0.0001)
Year FE	No	Yes	Yes	Yes	Yes
Market FE	No	No	No	Yes	Yes
N	82,633	82,633	82,633	82,633	82,633

B.6 Crime Exposure

To lend credence to the reputational risk mechanism, we explore whether properties in the lower NOI/unit tiers are disproportionately exposed to local crime. Higher crime rates in low-rent areas may discourage investment from institutional investors owners due to reputational concerns, as suggested in Section 5. If crime is perceived as a disamenity or a signal of neighborhood distress, properties in high-crime locations may attract less capital, resulting in persistent underinvestment and, consequently, persistently lower prices and higher rental yields.

Crime data comes from the FBI's Uniform Crime Reporting (UCR) system and is available at the police-station level at monthly frequency. We focus on two types of crime in 2023: property crime and total crime (which includes both property and violent crimes). We aggregate monthly crime counts to the annual level and compute crime rates per 100,000 population in 2023. We then spatially match each property in our sample to the nearest police station using the haversine distance and the latitude-longitude pairs

of the properties and the police stations. To ensure accuracy, we restrict the sample to matches where the nearest police station is located within 10 kilometers of the property.

We use two different sets of apartment properties for robustness. The left panels of Figure B8 use the apartment properties from our main CredIQ dataset. Their NOI/unit is the underwritten NOI/unit at loan origination. The right panels use apartment properties from the RCA dataset. That dataset and the procedure to impute a NOI/unit to those properties are described in Appendix B.7.

To assess the relationship between local crime exposure and net rent (NOI/unit), we residualize the relationship with respect to regional fixed effects (MSA dummies in the CredIQ sample and RCA market dummies in the RCA sample). We then group observations into 20 bins based on the residualized crime rate and plot the average NOI residual in each bin. All four panels of Figure B8 reveal a strong negative relationship between crime exposure and net rent levels, consistent with crime being a deterrent to capital investment in lower-tier locations. The null hypothesis that the slope is zero can be rejected at the 8% level (p-value = 0.079) for property crime and at the 6% level (p-value = 0.056) for total crime in the CredIQ sample. In the RCA sample, the slope is significant at the 2% level for property crime (p-value = 0.012) and at the 2% level for total crime (p-value = 0.021).

The figure also plots the best-fitting quadratic spline (dashed red line). It suggests that crime ticks back up in expensive neighborhoods.

B.7 Investor and Capital Flows Analysis

This appendix investigates whether inflows in the low-rent segment of the rental market affects net rental yields, capital gain returns, and total returns in that segment.

B.7.1 Transaction Data with Investor Information

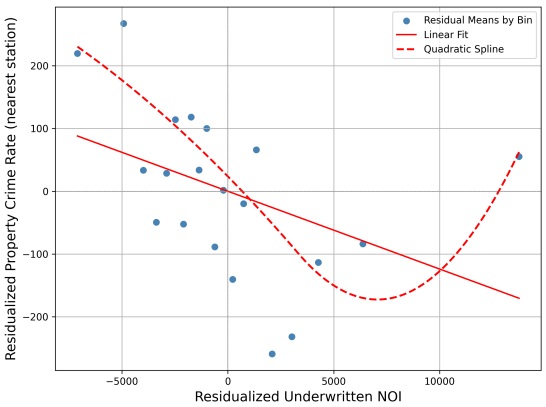
To do so, we obtain data from MSCI Real Capital Analytics, henceforth RCA, on multifamily property transactions between 2001 and 2023 in the United States over $0.5 million. We refer the reader to Koijen et al. (2025) for a detailed description of these data. This data set contains information on the identity of buyers and sellers, including what type of investor they are (e.g., REIT, Real Estate Private Equity Fund, Owner-Operator-Developer, etc.). The information on investor identities allows Koijen et al. (2025) to construct a measure of investor portfolio size from the observed transactions.

B.7.2 Descriptive Statistics

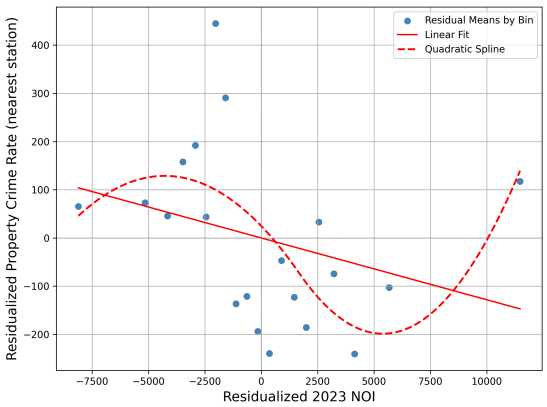
As a first look at the data, we compare multifamily properties in the RCA data to their counterparts in the CredIQ data in terms of the number of units in the building. Figure B9 plots the histograms of all multifamily properties in the RCA data set (orange) and overlays it on the histogram of properties in the main CredIQ dataset (blue). There is very

Figure B8: Crime Exposure and NOI

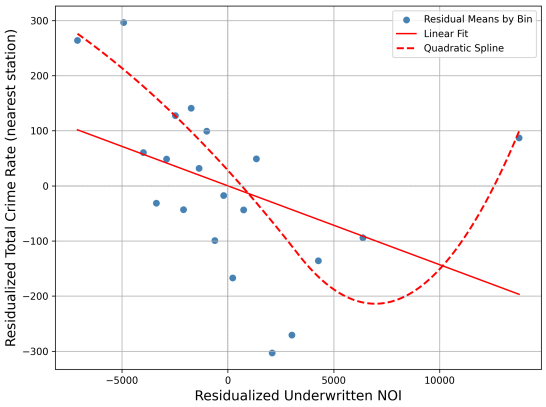
(a) Underwritten NOI (CredIQ) - Property



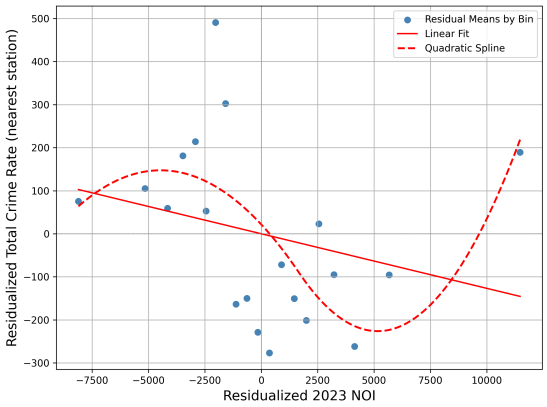
(b) 2023 NOI (RCA) - Property



(c) Underwritten NOI (CredIQ) - Total



(d) 2023 NOI (RCA) - Total

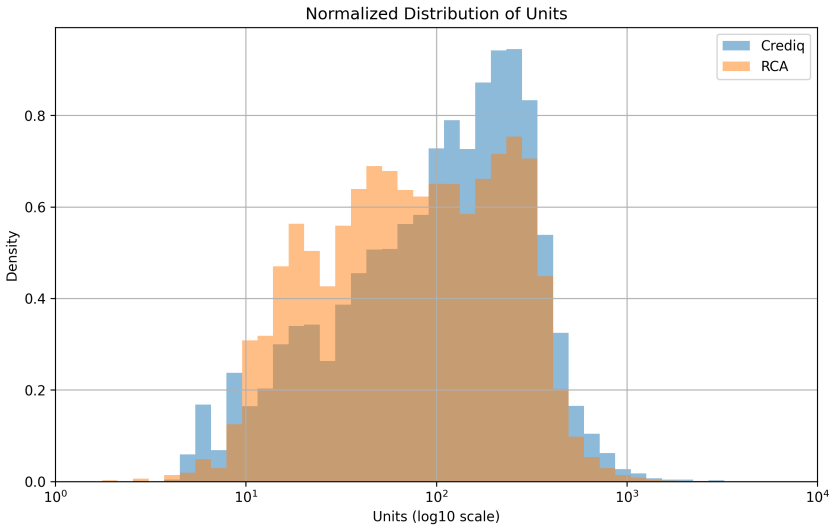


Notes: Each panel shows a binned scatterplot of NOI residuals and crime rate residuals after controlling for regional fixed effects (MSA or RCA market).Each property's crime exposure is defined as the annual crime rate of the corresponding type (per 100,000 population) of its nearest police station based on 2023 UCR data from the FBI. The top row uses property crime rates; the bottom row uses total crime rates (including property and violent crimes). The left panel uses underwritten NOI from our main CredIQ dataset, while the right panel uses 2023 NOI estimates from RCA dataset. All plots indicate a strong negative relationship between crime and net rental income.

large overlap between the two distributions. Since the RCA data is sourced from deeds records, it does not condition on any mortgage activity. Hence, it contains properties that are purchased with all equity (no mortgage) as well as properties purchased with a bank mortgage. The CredIQ data pertains to properties financed with a mortgage that is securitized (by Fannie, Freddie, or private-label CMBS and CRE CLOs). One might conjecture that securitized mortgages are associated with larger properties, but this turns out not to be the case. While both datasets are missing single-family rentals as well as some of the very smallest properties, we have a substantial number of buildings with few units.

Second, we look at the distribution of multifamily property ownership by investor type in the RCA data. Following Koijen et al. (2025), we split the largest investor type, owner-operator-developers into two subgroups: local (OOD_L) and national (OOD_N), where the former are investors that are only ever observed in one RCA market (there

Figure B9: Histogram of Building Size



are 60 geographies in the data) and the latter are observed in more than one market. The category "Users" combine companies, non-profits, and government institutions. The category "Institutional" combines pension funds, insurance companies, endowments, investment managers, open-ended funds, banks, finance companies, and CMBS. The data also contain the categories REITS, Real Estate Private Equity funds (REPE), Foreign investors, Individual investors (high net worth and non-traded REITs), and a small category of Unknown investor types. Figure B10 shows the composition of multifamily assets held at the end of 2023 in the RCA dataset, by dollar share of holdings of each investor type. The figure shows that there is substantial participation in the U.S. multi-family market by a wide variety of investor types. Local OODs account for nearly 15% of holdings while national OODs (OODs active in more than 1 market) account for 43.3% of transactions. On the other end of the spectrum, large institutional investors such as Institutional (12.6%), REPE (9.3%), REITS (5.9%), and Foreign investors (7.3%) collectively account for 35% of holdings.

Third, we split the universe of RCA multifamily transactions into quintiles, based on the number of units in the property. The left panel of Figure B11 shows the composition of investor types in each of these segments. The right panel shows the average investor size and its IQR for the same five property-size groups. We see that local and national OODs are more active in the smallest property size group whereas REITS, REPE, Institutional, and Foreign investors are much more active in the largest segment. In terms of size, we see positive assortative matching between building size and investor size.

Fourth, we sort all RCA multifamily transactions into deciles by their price per unit, expressed in real dollars. We then show the composition of buyers in each decile, by type on the left-hand side of Figure B12 and by investor size on the right-hand side of Figure B12. We see that the local and national OODs are more active in the low-value segment whereas the REITs, REPE, Institutional, and Foreign investors flock to the high-quality

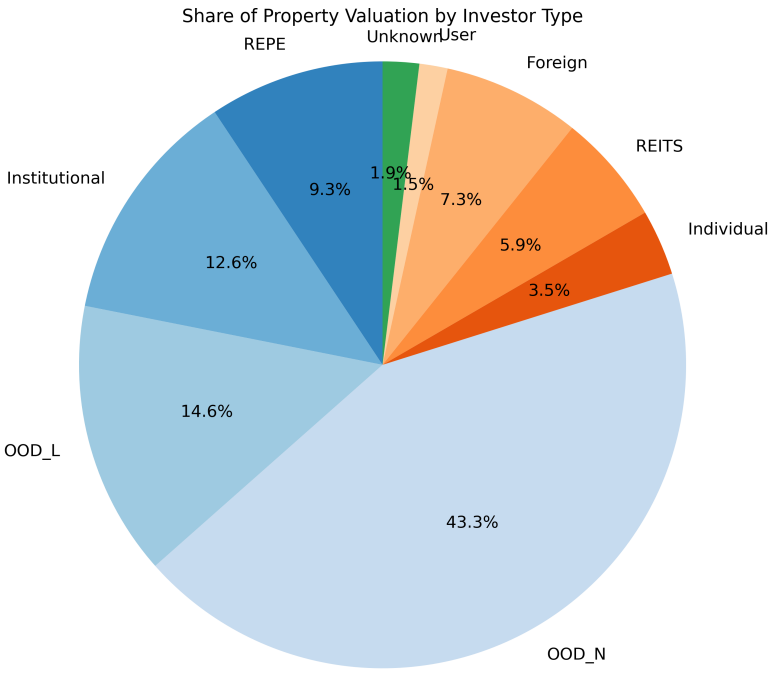


Figure B10: Share of Property Valuation by Investor Type

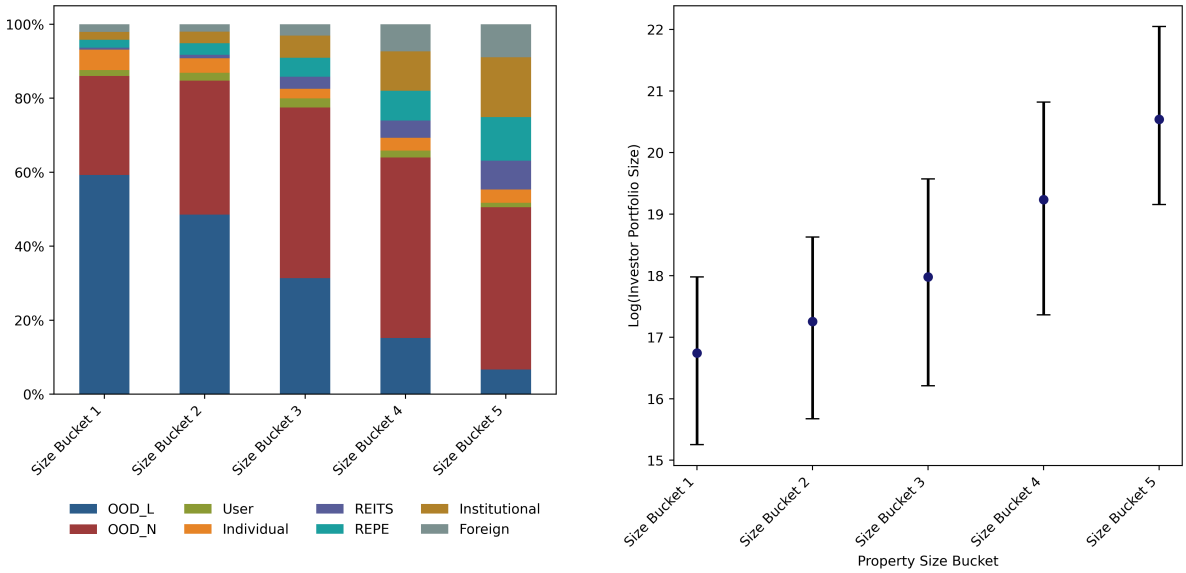
Notes: This figure shows the composition of apartment ownership by investor type, where ownership is measured at the end of 2023 in USD. Data source: RCA.

segment. In the lowest price-tier of the markets, the average size of investors is smaller. Higher priced apartments are owned by larger investors on average, with diminishing investor size effects in the top three quintiles.

B.7.3 Assigning RCA Properties to NOI deciles

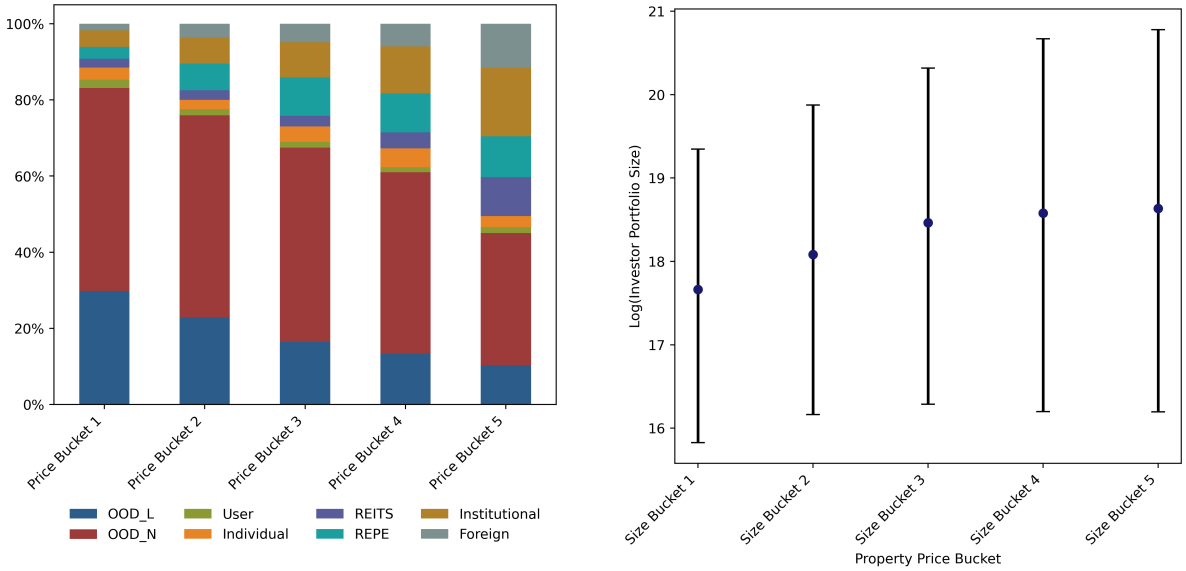
The RCA transaction data allows us to study investor flows, but RCA does not contain information on NOI/unit. Hence we aim to assign to each RCA property a value for the NOI/unit using the following procedure. For each RCA transaction, we first try to find a match in the CredIQ database. We pursue a fuzzy matching algorithm that uses information on the property address, its latitude and longitude, the transaction date, the transaction value, and the number of units. We construct a match score. If the match score is above a high threshold, we consider it a property-level match. We assign the weighted-average NOI/unit in the past three years from the CredIQ dataset to the RCA transaction. We weight more recent transactions more heavily. If we do not find a good enough match, we zoom out geographically and assign the weighted-average NOI/unit among CredIQ transactions in the neighborhood of the focal property. Our smallest

Figure B11: Investor Composition by Property Size



Notes: This figure shows investor concentration by the aforementioned property size buckets at the end of 2023 for investors that own at least one property. The left figure shows the distribution of investor types by the size buckets while the right figure shows the average and inter-quartile range of the log of investor portfolio size for each property size bucket. The cutoff points of property size buckets are 25, 54, 111, and 225 units. Data source: RCA.

Figure B12: Investor Composition by Property Price per Unit



Notes: This figure shows investor composition by property price per unit buckets at the end of 2023 for investors that own at least one property. The cutoff points of property per unit price buckets are \$104,604, \$157,491, \$223,220, and \$324,212 in 2023 dollars. The right figure shows the average and inter-quartile range of the log of investor portfolio size for each property price bucket. Data source: RCA.

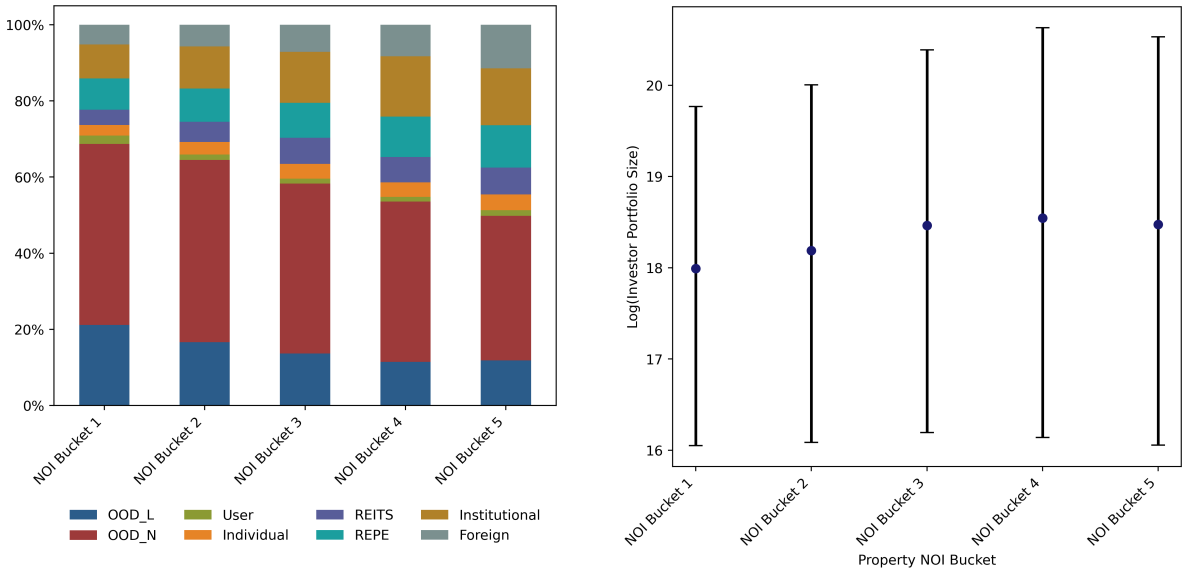
definition of neighborhood is a 0.6 by 0.6 mile square (1km by 1km). If we have at least 5 properties built prior to the focal transaction date in that area, then we assign the

weighted-average NOI/unit of those neighboring properties. If we do not have sufficient transactions in that small geography, we zoom out further and consider a 6 mile by 6 mile area (10km by 10km). When the weighted-average rent is still missing in this block, we fill it with the value in the 60 mile by 60 mile block. If necessary, we zoom out further to the RCA market. As a last resort, we assign the national average NOI/unit.

B.7.4 Investor Composition by NOI/unit Decile

Having assigned a NOI/unit to each RCA transaction, as discussed above, we can study the investor composition by NOI/unit decile. Figure B13 is similar to Figure B12 but sorts on NOI/unit rather than on price/unit. Its message is similar. There is substantial segmentation among investors with smaller and more local investors (esp. the OOD_L and OOD_N types) disproportionately active in the low-net rent segments.

Figure B13: Investor Composition by NOI per Unit



Notes: This figure shows investor composition by quintiles, computed from the NOI per unit distribution of apartment properties at the end of 2023 for investors that own at least one property. The left figure shows investor types while the right figure shows the average and inter-quartile range of the log of investor portfolio size. NOI per unit quintiles are formed within each RCA market. Results are similar whn the ranking is done nationally

B.7.5 Capital Flows

Next, we turn to the main question of interest: does increased investor activity in low-tier rental properties change the net rental yield, capital gain return, or total return of low-tier properties. Does it matter which investors invest in low-tier housing?

We run a series of regressions. The dependent variable is the slope of the relationship between Y_t^i and $NOI_t^i/unit$ estimated cross-sectionally across properties i for a given

year t, where Y_t is one of the net rental yield, the capital gain return, or the total return. We obtain one such slope estimate for each year t.

The independent variable is the investor flow. We define the variable *Flow Bottom 20* to be the purchases of apartment properties whose units are in the bottom two deciles of the NOI/unit distribution. Similarly, *Flow Top 20* are the purchases of properties in the top two deciles of the NOI/unit distribution. We lag the flow variables by one year to allow the market prices to adjust to the capital flows. Table B7 shows the results.

The first column pools all investors. It shows that in years in which there are more investor purchases in the bottom tiers of the rental market (large *Flow Bottom 20*), the net rental yield slope increases (panel A), the capital gain return slope increases (panel B) and the total return slope increases (panel C). An increase in the slope means that it becomes less negative, less steep. Conversely, more investor purchases of apartment units in the top tier of the market lower the slope, they make it more negative, steepen it. All six slope coefficients are measured precisely; the coefficients are statistically significant. most of the effect on total returns comes through capital gains. These results are consistent with the theory. When more investor capital flows into the bottom of the rental market, net rental yields, capital gain returns, and total returns all decrease in that segment, flattening the relationship of those variables with NOI/unit. Conversely, when more capital flows into the top segment, the relationships of net rental yields, capital gain returns, and total returns with NOI/unit steepen.

The remaining four columns break down the effect into that by subsets of investors. Columns "Inst" and "OOD" break down the universe of investors into two groups. The column 'Inst" combines the total purchases of the investor categories Institutional, REITs, REPE, and Foreign. The column "Other" combines the total purchases of all other investor categories (OOD_L, OOD_N, User, Individual, and Unknown). These are the smaller, more geographically concentrated, less sophisticated actors in the marketplace. The main finding is that the effects are stronger for Institutional investors than for the other investor types. The effects on capital gains remain significant for Institutional investors.

The last two columns look at investor flows by the largest one-third of investors ("Large") and the smallest one-third of investor ("Small"). Point estimates are significant for large investors but not for small investors in all panels. Again, this is consistent with the theory that when large investor types enter in the bottom tiers of the market, they have a larger impact, working to reduce return differences between the low and high tiers.

Table B7: Investor Flows

	All	Instit	Other	Large	Small
Panel A: Net Rental Yield					
Constant	-3.56	-2.13	-3.41	-3.77	-1.55
	(3.36)	(4.72)	(3.35)	(3.38)	(1.88)
Flow Bottom 20	**0.81**	0.29	0.35	**0.78**	0.13
	(0.38)	(0.33)	(0.55)	(0.38)	(0.34)
Flow Top 20	**-0.68**	-0.21	-0.23	**-0.64**	-0.07
	(0.36)	(0.21)	(0.57)	(0.35)	(0.40)
R^2	0.246	0.068	0.079	0.234	0.051
Panel B: Capital Gain Return					
Constant	13.10	-7.80	14.80	11.35	10.39
	(9.75)	(9.64)	(10.81)	(9.55)	(5.74)
Flow Bottom 20	**2.59**	**2.55**	1.28	**2.65**	0.37
	(1.10)	(0.67)	(1.77)	(1.07)	(1.05)
Flow Top 20	**-3.20**	**-2.21**	-1.99	**-3.19**	-0.98
	(1.04)	(0.42)	(1.85)	(0.99)	(1.21)
R^2	0.448	0.663	0.165	0.470	0.232
Panel C: Total Return					
Constant	9.53	-9.94	11.39	7.58	8.83
	(9.90)	(10.07)	(11.94)	(9.67)	(6.46)
Flow Bottom 20	**3.40**	**2.83**	1.63	**3.44**	0.50
	(1.12)	(0.70)	(1.95)	(1.08)	(1.19)
Flow Top 20	**-3.88**	**-2.41**	-2.22	**-3.83**	-1.06
	(1.06)	(0.44)	(2.04)	(1.00)	(1.36)
R^2	0.509	0.682	0.121	0.531	0.159

C Data: Netherlands

C.1 Data Description

The main database of Statistics Netherlands we use is the Woonbase, which covers the period from 2015 until 2023. For every year, this data provides for the universe of properties in the residential housing stock, the residents of these properties and the households they are part of, as well as the housing costs they face (from 2018 only). These costs include rent or mortgage payments, owner and user taxes, housing subsidies and energy consumption. These measurements are taken every year as of January 1. This implies that we can track which households move during the year and to which properties they move. We can link residents (persons and households) to their demographic characteristics, as well as to their income, wealth, and employment status. For properties, we know the neighborhood, property type, the building year, the surface area, and the number of units in the building in case of a multifamily property.

For the purpose of our paper, we use this data to build a panel of properties on January 1 of each year and the rents and sales prices of these properties. We explain the key measurements we use to define returns in more detail in this appendix.

Rents: The Woonbase provides either the actual rent or an estimate of rent for every housing unit in which one household resides. For households sharing housing (e.g. people renting a bedroom), it is not possible to link the rent payment to the unit accurately. Unlike Belgium, the Netherlands does not have a central rent registry, so the database pools rental data from various sources. First, about two-thirds of rental housing in the Netherlands is provided by non-profit housing associations that predominantly cater to relatively low incomes. Rents on these properties can be observed administratively. However, this paper focuses on the private rental market, the remaining one-third of rentals. These rents are pooled from various sources. First, renters with lower incomes in the private rental market are eligible for means-tested rent subsidies from the government, and their rent can be observed from fiscal sources. Second, the data includes rent payments from the Dutch housing survey and rent survey, which are used to evaluate the evolution of housing conditions and to estimate the rent CPI, and cover all types of landlords. Finally, Statistics Netherlands buys data from all major listing platforms in The Netherlands. If the listing year is different from the currently observed year and the resident has not changed, the rent payment is indexed. Statistics Netherlands is not permitted to reveal which source was used, so we cannot observe whether a property was listed or not. For households where the rent cannot be observed from any of the sources available to Statistics Netherlands, Statistics Netherlands provides an estimate of the rent.

Property value: To compute rental yields and estimate capital gains we make use of an annual measure of market value based on the tax value of properties, the so-called

WOZ-waarde. Formally, the tax value estimates the price a property could sell for in the open market, conditional on it being free-of-use and without any liens. The tax value is estimated based on realized sales of similar properties in the 12 months before and after the reference date for valuation, which is January 1 of the previous year for the current-year tax value. Municipalities estimate these values based on extensive data they have available on these properties and transactions of similar properties. By law, tax values on average have to equal estimated market value, and they tend to be highly accurate. As a case in point, some lenders even allowed using tax values to underwrite mortgages until European legislation in 2021 mandated the use of an appraiser. Changes in property value due to maintenance or gradual depreciation are typically not captured in tax values as these remain hard to observe.

Taxes: The Woonbase provides information on taxes that have to be paid by property owners and must not be passed through to tenants. These are property taxes paid to the municipality and the regional water management authority. Additionally, some property owners have to pay land leasing fees which vary at a granular geographic level. From the data, we obtain average tax rates per neighborhood for owner-occupiers and apply this to rented properties. Tax rates do not differ between owner-occupiers and landlords, but since the Woonbase takes costs from the user perspective, it only reports these taxes for owner-occupiers.

Supplementary data: We complement the data from the Woonbase with three additional datasets from Statistics Netherlands. First, using the dataset Eigendomwoningtab, we can identify the portfolio size of owners as well as the legal form of landlords in the private rental market. Private rental properties can be owned by businesses ("corporate landlords"), natural persons ("retail investors"), or other legal entities ("other investors"). The landlord information is only available from 2021 onward. For the earlier years, we use the 2021 information unless the owner of the property changed. If the owner changed and the property was sold, we can infer the type of owner and the size of its portfolio from transactions data of the Dutch Land Registry that we describe below.

Second, we use supplementary wealth data from the Vehtab dataset to get information about bank balances of households on January 1, the other assets they own as well as their debts. More specifically, we look at the category "other debts," which include consumer loans obtained from banks and any debt owed to the fiscal authorities as well as to health care providers.

Finally, to measure the vacancy duration for properties that experience tenant turnover, we make use of longitudinal data on the residential spells of the entire population of the Netherlands (dataset: Gbaadresobjectbus).

C.1.1 Transactions data

While tax values are administratively observed for all properties and time periods, we link them to sales prices to estimate investment flows across segments and compare "paper" capital gains to realized capital gains. We obtain administrative data from the Dutch Land Registry that contain the universe of property transactions from 2009 onward. For every transaction we observe the buyer and the seller, the transaction price, and the housing unit(s) involved in the transaction. We also observe the total property portfolio of the buyer and seller at the date of the transaction (excluding the transferred property). We merge this data to the Woonbase panel. For every housing unit in the Woonbase panel on January 1 of each year, we define whether it is sold in a given year and link the characteristics of the sale to the data on the housing unit itself. In the rare case that a unit is sold repeatedly within the same year, we only use the first sale. We also link this information to other administrative data sources. If units are sold in a portfolio, we observe the sales price for the entire deal and cannot distinguish the prices of individual units.

C.1.2 Maintenance cost data

The Woonbase does not provide estimates of maintenance costs for rental property. For this purpose, we make use of data from Koeter Vastgoed Adviseurs, a company specialized in making cost estimates for real estate investors. They provide estimates of maintenance costs under a variety of scenarios and investment horizons. We obtain information on the maintenance costs of 147 typical Dutch property types, representative of the Dutch housing stock overall. For every property type, Koeter Vastgoed Adviseurs determines what type of maintenance is needed and and what frequency. To estimate the costs of these activities, they maintain an extensive database with the current cost of the building materials that would be necessary as well as the costs for the building professionals that would be necessary to exectue the mainteannce.

To determine maintenance costs, we assume the costs that are necessary for long-term commercial exploitation of these properties. These are defined by Koeter as the annual costs necessary to continue operating a property for 99 years in current prices. In practice, rental investors may have shorter horizons. They may invest less in a property, and sell a property at a time that it needs significant capital expenditures. However, these investors would then also receive lower capital gains due to excess depreciation of the property. Since we cannot observe investors' decision on maintenance at property level, we use a consistent maintenance expenditure benchmark provided by Koeter. Our measure of capital gains is consistent with this measure, since the tax values in NL are not adjusted for depreciation, assuming that owners make the required capital expenditures to ensure that the property holds its value over time.

C.1.3 Data on non-payment

In the data, we cannot directly observe other costs made by landlords, including costs for non-payment and management. To proxy for the costs of non-payment, we use data on the realized stock of rent arrears, the number of tenants with arrears and actual evictions based on data from housing associations. Every year, the housing associations have to provide information to the government to assess their performance, the so called DVi.[30] In 2017, this assessment included data on non-payment. The associations provide this information separately for every individual municipality they are active in, providing us variation at the association-municipality level. Although housing associations charge lower rents, which might reduce non-payment, they also provide the bulk of housing to tenants at any risk of non-payment and due to their social function, and might accept more non-payment until proceeding to eviction compared to for-profit landlords. In short, it seems implausible that costs of non-payment risk are significantly higher in the for-profit sector. In recent years, non-payment and evictions of housing associations has declined significantly, as part of early-signaling programs. However, these declines had not yet materialized by 2017.

Overall, the housing associations reported that in 2017 8% of tenants were in rent arrears, that this led to non-payment equal to 1.2% of total rental revenue and 4,307 executed evictions (1.9 per 1000 properties). Personal default is not possible in The Netherlands, housing associations and other creditors try to recover rent arrears unless they cannot reach the tenant. In 2017, housing associations wrote-off 0.33% of total rental revenue as irrecoverable. However, this excludes any costs for evictions and administrative costs to recover late payments. Accordingly, the 0.33% serves as a lower bound on the cost of non-payment and the 1.2% as an upper bound.

C.1.4 Management costs

In the Netherlands, our measure of measurement costs covers the typical fee charged by agents for administrative and technical management of a property. Management companies either charge a fixed fee per unit per month or a percentage of total rental revenue. The typical variable fee for management of a property portfolio is 4% of rental revenue.[31] For most large portfolios, the exact management cost is negotiated directly between owners and managers so that we cannot observe this directly. We use the fee (excl. VAT) reported in the yearly guidelines for market-based valuation of rental property for Housing Associations.[32] However, management costs for an individual or small number of properties tend to be larger than the values reported by the Housing Associations. For management of a single property, a typical variable fee is 6% of rent per unit and a fixed

[30] Source: Verantwoordingsinformatie Woningcorporaties, DVi 2017

[31] For example, see the cost scenarios of Vastgoed Belang, the main advocacy group for (non-instiutional) private landllods in the Netherlands

[32] Guidebook Model-Based Market Valuation 2023, Retrieved January 23, 2025.

fee is typically 50 euros per month.[33]

In our baseline scenario, we consider a management cost of 4% per unit for owners of 10+ properties, 5% per unit for owners of 2-9 units, and 6% per unit for owners of a single unit. These groupings correspond to the categories we observe in the ownership data. We impose a minimum monthly fee of 32 euros (10+ units), 37 euros (2-9 units) and 42 euros (1 unit) in 2023 prices. In two alternative scenarios, we either impose only a variable fee or only a fixed fee.

C.1.5 Listings data

We supplement the rent data in the Woonbase with listings data from Realstats, which cover the period from 2006-2023. This dataset combines listings from Pararius and huurwoningen.nl, which together cover the large majority of the listed Dutch private rental market. We can link most of these listings to our administrative data, providing us with the same property characteristics. This database grows in scope over time, and data coverage is thin before 2008 and in 2015. For every property, we observe the list price, the number of square meters of the unit, and various other property characteristics. We observe both the date the property was listed and the date of de-listing, which typically happens after a tenant is found and a rental contract is about to be found. For properties that are listed multiple times within 90 days, we take the first listing date and the final de-listing date and remove all observations in between.

We use this data for two purposes. First, we use the listings data to calibrate the vacancy period of every property in the administrative data, which we discuss below. Second, we use the listings sample to expand the time horizon of our NL analysis and analyze the covariance of GDP growth and rent growth across deciles. Third, we use the listings samples to compute yields for the small set of properties that are purchased and subsequently listed and leased, so that we have both an actual transaction price and a recent rental price. In our main sample, we cannot distinguish whether a property was recently leased.

C.2 Data Processing

The key purpose of our analysis is to use the Woonbase and the linked data sources to estimate total returns to real estate investments for low-value, properties catering to low-incomes relative to more expensive properties with richer tenants.

C.2.1 Data Filters

The Woonbase data contains the universe of properties in the Netherlands on January 1 of each year from 2015 to 2023. It comprises about 8 million housing units observed

[33]For example, DHVC Property Management and GoHome Management

over 9 years. From this panel, we first remove observations with extreme, incomplete, or implausible values.

- Tax value unavailable or in the bottom or top 0.5%
- Square meters of housing unit below 20 or above 300 square meters
- Property type unknown
- Building year before 1500 or after 2024
- Property occupied by institutional households (e.g. prison, nursing home, youth care)

Because the data quality improves over time, the number of observations considered outliers drops from 7% in 2015 to 4% in 2023.

C.2.2 Main Sample

To determine investment returns for private rental properties, our main sample is restricted to rental properties owned by natural persons or business entities, which form the private rental market. We sort the properties in this sample into ten deciles based on their predicted market rents, following the method explained in the main text.

For these properties, we compute returns only for properties for which we have an actually observed rent and for which the tenant moved to the property within the last two years. For private rentals, 45% of rent payments cannot be observed, in particular for properties with long-time residents. The number of missing observations is most severe for landlords with small portfolios. By focusing on recent movers, we ensure that our sample is representative of new rental contracts, and thereby more consistent with the Belgian sample which solely relies on new contracts. For existing tenants, rent indexation is estimated and typically unobserved. This leaves us with 808,000 observations over six years.

Before estimating rental returns, we remove observations from the data with extreme or implausible rental values. First, we estimate a regression that models log rents and log rental yields as a function of property, location and household characteristics, as well as time, tenancy duration, type of owner, and owner portfolio size. We exclude neighborhoods with five or fewer observations, for which neighborhood fixed effects cannot accurately be defined. We estimate this regression separately for private rentals and housing association rentals. We remove extreme observations whose rents or yields are more than 0.693 log points away from predicted value, which indicate data errors or implausible values suggesting non-standard cases. This removes 4% of observations of our sample, leaving us with 775,490 observations. On this sample, we compute our main return measures per rental decile.

C.3 Measuring Turnover and Vacancy Duration

We identify a turnover event on a property in a given year when the primary tenant at the beginning of the year relocates to another property during the year. To estimate the vacancy period, we measure the time between the tenant's administrative deregistration and the registration of the subsequent tenant.

Although administrative records allow us to track all residents, this measure has two key challenges. First, while registration is mandatory, tenants do not always register immediately—or at all. Foreign residents staying in the Netherlands for less than four months are not required to register, and these tenants often occupy private rental properties. As a result, the vacancy period derived from administrative data is likely to be significantly overestimated relative to actual rental contracts.

Additionally, the tenant's deregistration date may not precisely coincide with the contract end date. Since tenants must vacate the property on the contract's termination date, they often move their belongings beforehand.

To adjust the administrative vacancy period, we link tenants who moved out in 2022 to our Realstats listings database. We identify a match if the property appears on a rental platform within three months of the tenant's deregistration—looking both backward and forward. We then determine when the property is delisted, assuming the new tenant's contract begins seven days after delisting.

If a property is delisted at least seven days before the previous tenant leaves, we assume a vacancy duration of zero days. In all other cases, the vacancy period is defined as the number of days between the delisting date plus seven days and the deregistration date of the previous tenant. We then regress this vacancy period on the administratively observed vacancy period, incorporating polynomial terms up to the fifth degree. This adjustment generally shortens the administrative vacancy period, aligning Dutch estimates more closely with Belgian estimates, where vacancy duration can be measured precisely.

Finally, we extrapolate this model to all turnover events in the dataset. If no new tenant is identified within two years, we assume turnover costs correspond to the median vacancy duration.

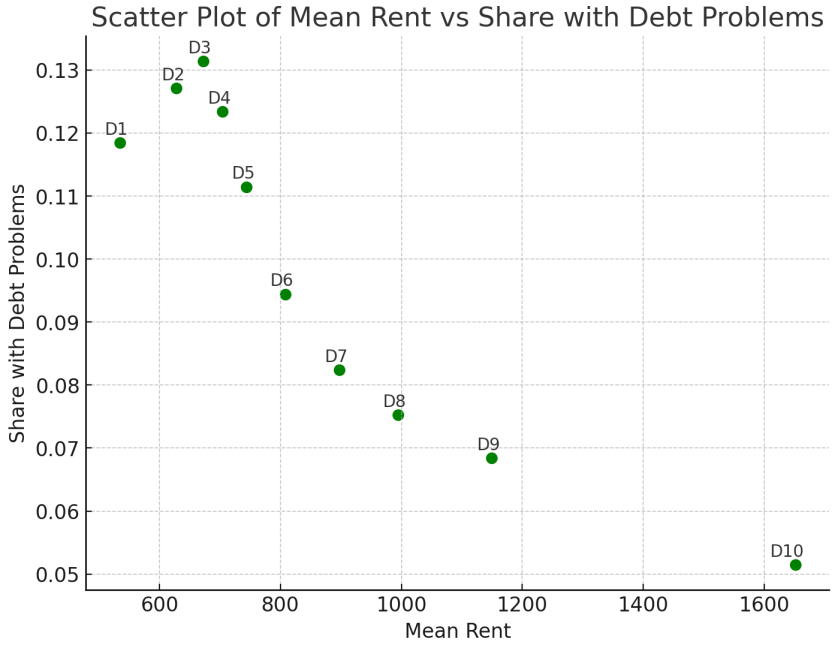
C.4 Measuring Costs of Non-payment

In the data, we do not directly observe non-payment at property-level, and there is no direct information on non-payment and evictions for private landlords. Instead, we develop a measure of the population that is at risk of non-payment and eviction and distributed the realized cost of non-payment nationwide over these tenants.

Non-payment risk measure We use administrative data that identify for every individual in The Netherlands whether this person at any point in time has problematic debts. An individual is registered as having problematic debts if it has an unpaid tax bill,

health insurance premium, student loan (re)payment, government subsidy repayment or fine for a period longer than the regular period for paying these bills with a payment plan. In The Netherlands, individuals can typically opt for a payment plan in case they have difficulty making payment, such that debt is only counted as problematic if they fail to pay after the regular period of a payment plan has expired.[34] Individuals under garnishment of income are included in the definition too. While this measure does not directly incorporate rent arrears, private landlords describe that most tenants in rent arrears have other problematic debts.[35]

Figure C1: Problematic debts by rent decile



We plot the risk measure by predicted rent decile in Figure C1. Non-payment risk concentrates in the bottom deciles of the predicted rent distribution, and our risk measure indicates about one in about eight tenants in the bottom five deciles has significant non-payment risk. This reduces significantly towards one in twenty tenants for the top decile. Overall, it appears plausible that non-payment risk is much smaller for tenants that can afford expensive properties.

Overall, the number of tenants with rent arrears is likely somehwat smaller than the number of tenants with any problematic debts. For example, data provided by the housing associations suggest about 8% of tenants has currently any arrears, whereas we estimated about 11% of these tenants have any problematic debts. It is not surprising that the latter number is higher: individuals might prefer not to pay their health insurance or tax bill before not paying rent, since not paying rent might lead to force them to move to another property or even formal eviction. We will therefore only use this measure as a

[34] See Statistics Netherlands, Dashboard Schuldenproblematiek for a more detailed discussion
[35] See Vroeg erop af in de particuliere verhuur?

proxy for the *relative* incidence of non-payment across rent deciles and the corresponding costs.

Average cost of non-payment To estimate how non-payment risk translates into actual costs, we use data on the realized stock of rent arrears, number of tenants with arrears and actual eviction based on data housing associations provided in 2017 to the housing association authority, the last year such information was asked.[36]

Overall, the housing associations reported that in 2017 8% of tenants were in rent arrears, that this led to non-payment equal to 1.2% of total rental revenue and 4,307 executed evictions (1.9 per 1000 properties). Personal default is not possible in The Netherlands, housing associations and other creditors try to recover rent arrears unless they cannot reach the tenant. In 2017, housing associations wrote-off 0.3% of total rental revenue as irrecoverable. However, this excludes any costs for evictions and administrative costs to recover late payments. Accordingly, the 0.3% serves as a lower bound on the cost of non-payment and the 1.2% as an upper bound, which would assume that all unpaid rent will be never repaid.

From default risk to estimated cost of non-payment The question is how we can translate our observed measure of tenants at risk of non-payment and the reported rent arrears of housing associations into an tangible measure on the cost of non-payment. This always requires assumptions, and we therefore engage in a baseline scenario and a more conservative scenario that will serve as an upper bound on the costs of non-payment across different deciles of predcited rent.

In our baseline scenario, we will assume that about 0.3% of total revenue is written-off each year, which is directly line with the observations of housing associations. On top of this, we make assumption on the cost of recovering any rent. Based on the stock of rent arrears of housing association of 1.2% per year, we assume costs of recovering rents or eviction that equal 0.6% of total rent. We assume that theses costs are the same in 2017 euros for every instance of non-payment.

In our second scenario, which is more conservative, we assume that over the entire private rental sector, the cost of non-payment equals 1.2% of rental revenue, similar to the stock of non-payment relative to rental revenue for housing associations. This effectively assumes that the cost of evicting tenants and recovering non-payment equal total rental revenue recovered. To be conservative, we assume that the cost of non-payment is the same in euros for every tenant that is at risk of non-payment. Hence, this means that non-payment does not only concentrate in the lowest deciles of predicted rent, but that the cost of non-payment relative to predict rent per instance of non-payment is also much higher in the low deciles. This is likely an overestimate, as Housing Associations provide the large majority of rental housing to low incomes and also executed the large majority

[36]Source: Verantwoordingsinformatie Woningcorporaties, DVi 2017

of evictions in this period.[37]

C.5 Hedonic Model for Rents

We estimate a hedonic model to explain the distribution of real 2023 rental prices in the Dutch data. We estimate the following model, explaining log rental prices with the log size of a property and its age, and fixed effects for location at the neighborhood level, the number of units, the property type and the building vintage.

$$\log R_i = \alpha_{location} + \alpha_{units} + \alpha_{vintage} + \alpha_{type} + \beta_1 \log m_i^2 + \beta_2 \text{age}_{i,t} + \varepsilon_{i,t}, \quad (9)$$

The estimation results are reported in Table C1. The average population of the neighborhoods in our sample is around 2,000 residents. Fixed effects for the number of units are in steps of one for properties with 1-30 units, in steps of 5 for 30-100 units, and finally a class for 100-150, 150-200 units and 200+ units. Building vintage fixed effects are for every decade starting in 1870 and for every century in the period before 1870. Construction types are: apartments, detached properties, semi-detached properties, corner properties, and terraced properties.

Table C1: Hedonic Model for Rental Value in NL

	log(Rent)
$\log m^2$	0.349***
	(0.006)
Building age	-0.000025*
	(0.0001)
Fixed Effects:	
Neighborhood (7,515 categories)	Yes
Number of units (47 categories)	Yes
Building vintage (18 categories)	Yes
Construction type (5 categories)	Yes
S.E. Clustering	*Neighborhood-Level*
Observations	717,901
R^2	0.664
Within R^2	0.219
Note:	*p<0.1; **p<0.05; ***p<0.01

[37] The Dutch Bailiff Association reported in its 2017 annual rapport that it executed 5,213 evictions in 2017, but does not provide the reason for eviction nor the type of tenancy (housing association, private landlord or owner-occupiers (by mortgage lenders)).

C.6 Robustness: Alternative Rent and Capital Gain Yields

In our baseline results, we estimate returns based on the rent paid by recent movers and use the tax-assessed value as a proxy for sales prices. To assess the robustness of our baseline rental yield and capital gain measures, we now employ several alternative definitions, reported in Table C2.

First, we construct a market-based measure of yield by focusing on the subset of housing units that were both purchased and listed for rent in the same year. For these properties, we observe both an actual transaction sale price and a listed rent in the same year. Naturally, the main drawback of this approach is that it is relatively rare for investment properties to be transacted and listed on *Pararius* in the same year, limiting the sample to 6,581 observations. Unsurprisingly, Table C2 shows that yields of newly-listed properties are higher, as rents for such properties are typically higher compared to units that are not listed or where tenants have already resided for longer. Aside from this level difference in rental yields, the slope across deciles remains very similar.

Second, we test the robustness of our capital gains measure using alternative methods. In our baseline approach, we compare the tax-assessed value of properties on January 1 of each year to the value on January 1 of the previous year to measure annual capital gains. We can compute this measure for the period from 2018 until 2023.

In our first alternative, we examine properties sold at least twice between January 1, 2018, and December 31, 2024, and compute the average annualized real capital gain for those sold at least two years apart. On average, this repeat-sale sample results in higher capital gains than the baseline, as property prices rose again toward the end of the sample period. However, the slope in capital gain returns between rent deciles remains similar, and is 44 basis points higher than for our baseline approach. One difference is that, capital gains in Decile 1 are lower than in Decile 2 in the baseline, whereas they are higher in Decile 1 under the repeat-sales method.

In a second alternative, we apply a hedonic method, following Equation (9). We include parameters for building size and age for each year separately and include property type-by-year and region-by-year fixed effects. Hence, capital gains of properties can vary across deciles due to diffences in size, age, type and location. We then compute each property's hedonic valuation at the midpoint of year t and year $t+1$, using data from 2018–2023 (thus producing capital gains for 2019–2024). This hedonic model also generates a clear return slope, which is again similar to that in the baseline model. This close similarity may arise due to a stronger overlap in timing and methodology. One difference is that the hedonic model suggests higher capital gains in Decile 1 relative to Decile 2, whereas the baseline shows the opposite.

In conclusion, all three approaches to capital gain measurement produce a consistent and quantitatively similar capital gain yield slope. There is some evidence that our baseline approach underestimates capital gains in the lowest decile, which would imply an even steeper slope in total returns, rendering our conclusions conservative. Some of

these differences likely reflect the timing of property valuation. Similar to the findings in Section 4.3, it does not seem that unobserved capital expenditures significantly affect the magnitude of the slope. If capital expenditures and renovations were more common in low- versus high-decile properties, we would expect a steeper slope in the repeat-sales approach (which includes the value of renovations) compared to the hedonic method. Instead, the return slope is similar across both the hedonic and repeat-sales approaches.

Table C2: Returns by Rent Decile, NL: Listed Rents and Actual Sales Prices (2023 Prices)

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D1-D10
Gross yield, baseline	6.84	6.24	5.77	5.36	5.03	4.74	4.52	4.40	4.26	4.18	2.66
Gross yield, bought+listed	7.47	7.18	6.71	6.14	6.03	5.71	5.64	5.29	5.05	4.73	2.74
Capital gain yield, baseline	5.69	6.52	6.48	6.34	6.01	5.91	5.90	5.54	4.76	2.71	2.98
Capital gain yield, repeat-sale	7.55	6.48	5.92	5.67	5.65	4.99	4.78	4.83	4.51	4.13	3.42
Capital gain yield, hedonic	6.41	6.21	5.97	5.61	5.36	5.27	5.26	4.83	3.97	2.63	3.78

Notes: Gross yields, capital gains, and total returns by rent decile for the Netherlands (NL). Gross yields compare baseline results to yields of properties bought and listed in the same year. Capital gains compare baseline results (based on annual tax value increases) to annualized gains from repeat-sales of investor-owned properties (2018–2024) and to hedonic predictions of annualized gains (2018–2024, mid-year). Differences in timing and sample selection imply that returns are not expected to be identical across methods.

C.7 Robustness: Longer Sample of Listings data

We use listings data to replicate the main facts for the Netherlands using a longer but less representative sample. We also use this data to study the relation between GDP growth and rent growth by decile in the Dutch data, replicating the analysis in the USA (Figure 2). The main advantage of the listings data is that it covers a much longer time period than the main dataset. The main drawback is that fewer low-tier properties tend to be listed, such that the sample is more expensive compared to our main sample. We restrict the sample to listings that we can match to our administrative data, and which do not contain extreme observations. Extreme observations are removed in the same way as for the main sample.

To identify rent growth per decile, we proceed in three steps. First, we estimate for every year individually the following rent model, with the list price as dependent variables and the following variables as independent variables:

$$\log R_{i,t} = \alpha_{location} + \alpha_{units} + \alpha_{vintage} + \alpha_{type} + \beta_1 \log m_i^2 + \beta_2 \log \text{list}m_{i,t}^2 + \beta_3 \text{age}_{i,t} + \beta_4 \log TV_{i,t} + \varepsilon_{i,t} \tag{10}$$

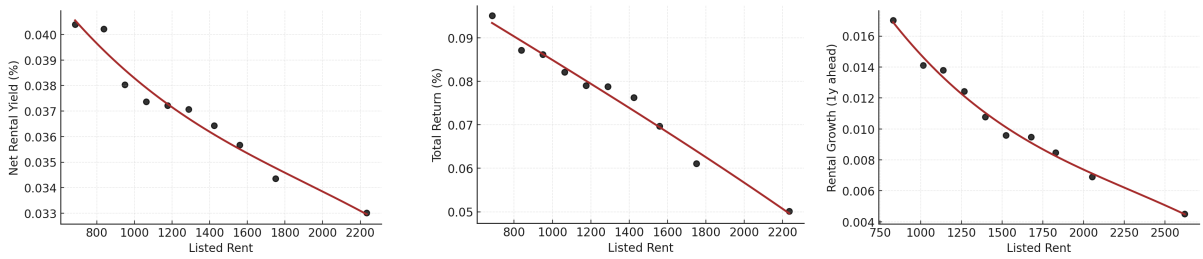
Relative to our main specification in equation (9), there are two differences. First, we include the log of the tax value of a property, $TV_{i,t}$, as an explanatory variable. This helps to reduce variation in case there are mismatches. Second, we add the square meters in the listing as an additional explanatory variable. The reason is that the measure of square meters in the administrative data employs a different definition than the square meter measure used to advertise properties. In practice, the latter is more predictive of

rental prices and thus improves our estimate.

We estimate the rent model for every calendar year from 2008 and 2023. Next, we sort properties in the sample at time t in deciles based on their predicted rental values based on the model estimated for listings at $t-1$. This ensures our decile classifications is insensitive to overfitting. For properties in 2008, we use an additional model based on data from 2006 and 2007 together, given the small number of observations in the early years. To compute cash flow growth, we compute the predicted rent of every property in our sample in every year.

We then compute the rental yield for every property by dividing the listed rent to the tax value, and aggregate by predicted rent decile. Since the list price includes the cost for furnishing, upholstering, and other service costs, while our main estimates are based on the "bare rent" for properties alone, we adjust rental prices for the average difference between the two. Since we do not have cost measures for the longer sample, we convert gross yields into net yields by matching properties in our main dataset to those in the listings database, and compute the average cost per rent decile in the matched sample. To compute total returns, we compute the average growth rate of the tax value between time t and $t+1$, and aggregate this per rent decile. Finally, we compute the predicted rental growth on a property by comparing the predicted rental price at time $t+1$ with the price at time t, and average this per rent decile.

Figure C2: Net Yield, Total Return and Cash Flow Growth by Decile (2008-2023)



(a) Net Yield, Listings, NL (b) Total Return, Listings, NL (c) Rent Growth, Listings, NL

Notes: Each panel plots the relation between the net rental yield (left panel), total return (middle panel), and cash flow growht (right panel) and the net rent or rental value of a property. Every point corresponds to the average return in a rent decile, order from D1 (left) to D10 (right). The deciles in the listings data differ from those in our main sample, as listed properties tend to concentrate more in higher deciles.

Figure C2 plots the estimated net yield, total return and predicted rental growth by rent decile. In line with our main findings, we see a strong slope both in yields and total returns. The D1-D10 difference is smaller for yields, but larger for total returns. We should acknowledge here that the first two rent deciles in the listings sample span the first five deciles in the larger sample, which might explain the smaller yield slope. Nonetheless, the large difference we find among total returns suggest they are not driven by the specific time period we look at in our main sample (2008-2022).

One reason why we find both higher yields and capital gains is because rent growth

appears significantly higher in the bottom deciles compared to the top deciles, with a D1-D10 difference of 125 basispoints per year.

C.8 Tracking New Investment Across Deciles

Are investors recognizing the excess returns in the low-rent tier and allocating additional capital to this segment?

C.8.1 Growth in the Housing Stock

A key advantage of the Dutch data is that we can track in detail all property sales as well as information on new construction, property transformation, and demolition. In this section, we examine the extent to which new investments are being made in the private rental segment. We consider four sources of investment:

- **Net Purchasing:** We sum the number of units bought by private investors from owner-occupiers or housing associations and subtract from it the number of sales by private investors to owner-occupiers and housing associations.
- **New Construction:** We count the number of newly constructed private rental units.
- **Property Transformation:** We count the number of new private rental units created by transforming existing properties into housing units (e.g., office-to-residential conversions).
- **Demolitions:** We count the number of demolished private rental units.

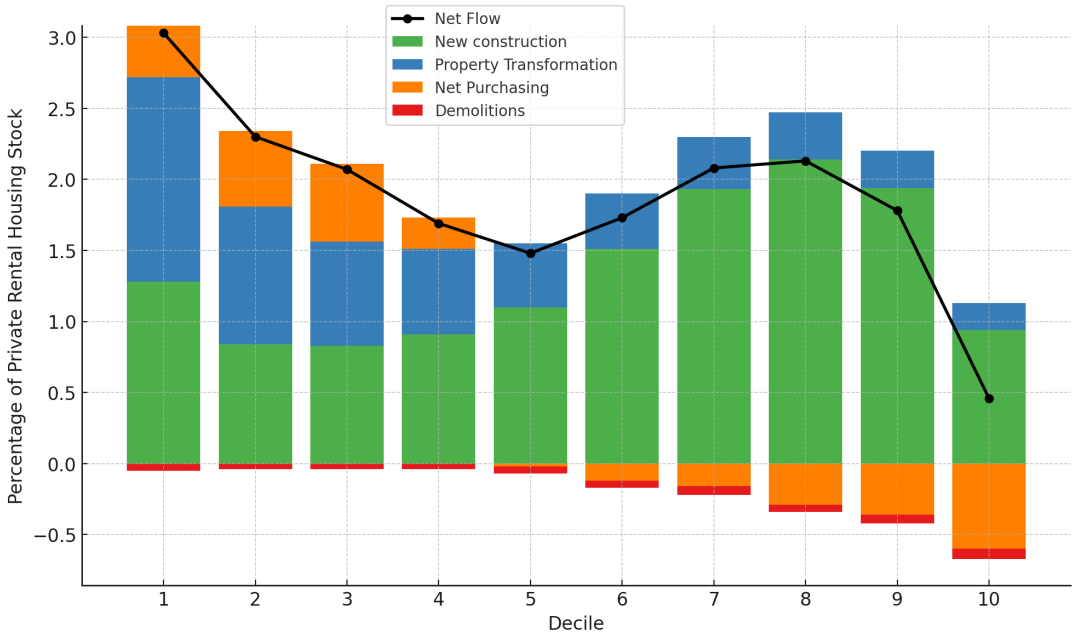
We aggregate these flows and express them as a fraction of the total private rental stock (expressed in number of units) in each decile, shown in Figure C3.

Between 2018–2023, the annual growth rate of the number of units in the bottom decile was about 3% per year — the highest among all deciles. In contrast, the growth rate was lowest in the top decile, at about 0.5% per year. Hence, it appears that the stock of rental housing in the lowest decile (low-rent tier) grew more rapidly compared to the highest decile (high-rent tier), suggesting that the market was moving capital in the right direction during this period.

The main factor contributing to growth in the low-rent tier was the conversion of existing stock. This occurred either through converting low-tier owner-occupied property into rental units (and vice versa for the top tier), or by transforming properties from other uses (e.g., offices) into low-tier rental housing.

Although there was substantial new construction in the lowest-rent decile, rates of new construction were higher in the mid-to-high rental tiers. From an economic perspective, this is sensible: new construction is typically of higher quality and tends to concentrate

Figure C3: Net Investment Flow by Decile as % of the Housing Stock, NL



Notes: The figure plots the annual net investment flow per segment, expressed as the total number of newly added units relative to the total private rental housing stock, from the bottom 10% (Decile 1) to the top 10% (Decile 10).

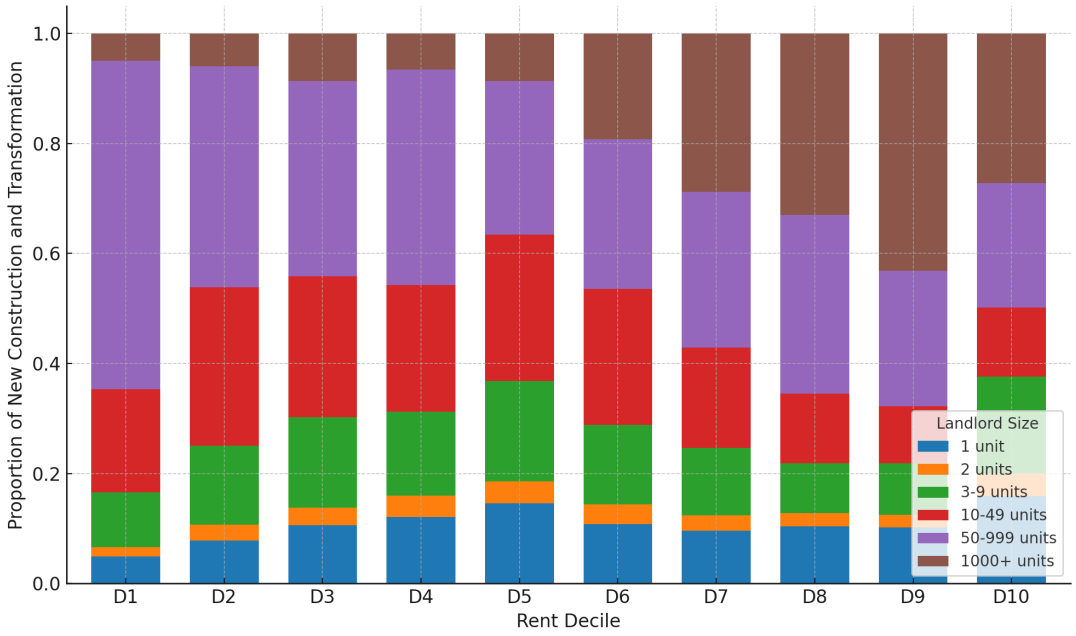
in more expensive areas, where property values are more likely to exceed the costs of new construction.

Some of the differences in flows across deciles may also reflect the preferred habitats of different types of investors. Most buy-to-rent (or rent-to-buy) conversions are carried out by small-scale investors who purchase or sell individual units. For larger investors, buying and selling individual units on the owner-occupied market typically does not allow them to deploy new capital quickly enough. Conversely, the construction and transformation of new units is mostly done by larger investors, as these activities require significant up-front capital investments. Even among these larger investors, there are clear differences in the rent tiers they target. Figure C4 plots the share of new construction and transformation in each decile by investors of different size. It is consistent with Figure 6 in the main text.

In the bottom decile, about 60% of transformation and construction was carried out by large investors owning between 50 and 1,000 units. Although most new units overall were built by large investors, they were most active in the bottom tier. By contrast, very large investors (owning over 1,000 units) tended to avoid the low tiers, focusing instead on constructing new properties in the top tiers and engaging little in property transformation. Small investors were most active in the middle tiers, while medium-sized investors concentrated their activity in deciles 2–6, investing less in both the lowest and highest tiers.

In summary, on average, more capital flowed to the low-rent tiers than to the high-

Figure C4: Net Construction and Transformation by Investor Size and Decile, NL



Notes: Distribution of new construction and transformation by landlords across rent deciles, categorized by their portfolio size. The figure plots the distribution from the bottom 10% (Decile 1) to the top 10% (Decile 10).

rent tiers. However, the highly segmented nature of private rental markets — by both segment and investor size — means that these differences were not as pronounced as one might expect. The key fact that stands out is that very large investors still made most of their additional investments by constructing new rental properties in higher tiers.

C.8.2 Capital Flows and Returns

When private investors buy properties, we expect this to lower rents and raise prices, thereby reducing rental yields in the specific segment where the purchases take place. Do such capital flows into the lower tiers put downward pressure on yields, thereby flattening the rent slope?

To investigate this, we first compute capital flows by rent decile at the regional level between January 1, 2018 and January 1, 2023. We focus on transaction-based investment flows, in line with Appendix B.7 for the US and D.5 for Belgium, counting the number of purchases and sales by investors in each region-decile combination and expressing them as a fraction of the total private rental housing stock in 2018 in that combination.

Second, we compute the average capital gain and rent growth for each region-decile, restricting the sample to those properties that changed hands both in 2017 and 2022. Changes in the relative supply of private rental and owner-occupied housing primarily affect rent levels for new residents, while rent changes for existing tenants are largely unaffected.

We then regress changes in yields, rents, and prices on the investment flows. The

results are in Table C3. For each variable, we consider three specifications: one that separately examines inflows (investors buying from owner-occupiers) and outflows (Columns 1, 4, and 7), one that looks at net flows (Columns 2, 5, and 8), and one that examines net flows while including fixed effects for every region and decile (Columns 3, 6, and 9).

Table C3: Investment Flows and Yields, Rent Growth and Capital Gains

	Δ Yield			Δ log(Rent)			Δ log(Price)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
% inflow	-0.166***			-1.377***			2.204**		
	(0.047)			(0.321)			(0.901)		
% outflow	0.321***			0.431			-3.826***		
	(0.084)			(0.778)			(1.274)		
% net flow		-0.156***	-0.195***		-1.432***	-1.547***		2.110**	1.324***
		(0.048)	(0.045)		(0.304)	(0.493)		(0.936)	(0.458)
Constant	-0.006***	-0.004***		0.197***	0.185***		0.461***	0.434***	
	(0.002)	(0.001)		(0.012)	(0.006)		(0.025)	(0.512)	
Fixed-Effects:									
Region	No	No	Yes	No	No	Yes	No	No	Yes
Decile	No	No	Yes	No	No	Yes	No	No	Yes
Clustering	Region	Region	Region	Region	Region	Region	Region	Region	Region
Obs.	392	392	392	396	396	396	396	396	396
R^2	0.082	0.057	0.392	0.066	0.055	0.312	0.058	0.046	0.524

Starting with yields, we find that a one percentage point increase in investor inflows reduces yields by about 17 basis points, while a one percentage point increase in outflow raises yields by 32 basis points. The effect of net inflow (Column 2) is very similar to that of gross inflow (Column 1), at around 16 basis points. This is important, as in similar analyses for Belgium (Appendix C.8) and the USA (Appendix B.7) we cannot precisely separate inflows and outflows. The effect of net inflow remains similar when including region and decile fixed effects. Overall, this suggests that increased investor flows toward the low-tier segment reduce yields.

This effect can result both from downward pressure on rents and upward pressure on prices. We replicate the specifications in Columns 1–3 for changes in log rental prices (Columns 4–6) and house prices (Columns 7–9). We find that a one percentage point (net) inflow into a region-decile lowers rents by about 1.4–1.5% and raises house prices by about 1.3–2.2%. In short, investment inflows affect yields through both reduced rents and increased prices.

D Data: Belgium

D.1 Data Description

For Belgium, our primary data sources consist of administrative records covering the entire universe of rental contracts, housing transactions, and energy performance certificates (EPCs).

Rents: Since 2007, Belgian law mandates the registration of all rental contracts within two months of signing. Registration is cost-free, and failure to comply grants tenants the right to unilaterally terminate the contract without penalty. As a result, we have access to the universe of rental contracts provided by the Federal Public Service Finance (FPS Finance) spanning from 2007 to 2022. Rental contracts are recorded at the time of origination, and the dataset includes detailed information such as the contract start date, signing date, lease duration, monthly rent, exact property address, and anonymized identifiers for both landlords and tenants.

Property values: To compute rental yields, capital gains, and retrieve property characteristics for rental contracts, we utilize the universe of housing transactions in the secondary market from 2007 to 2022, provided by FPS Finance. This dataset includes transaction details such as the sales price, date of sale, exact property address, and unique anonymized identifiers for both buyers and sellers. Additionally, it contains detailed property attributes, including useful surface area, plot size, built area, number of floors, rooms, garages, and bathrooms, as well as the presence of central heating, construction type (apartment or detached/semi-detached/terraced house), elevator availability, and construction year. To incorporate a measure of energy efficiency, we merge our dataset with the universe of EPCs. In Belgium, obtaining an EPC for rental properties has been mandatory since 2009 in Flanders, June 2011 in Wallonia, and November 2011 in Brussels. We match each rental contract with the most recent EPC issued before the lease start date, or, if unavailable, the closest EPC issued thereafter. To estimate the property value at the time of the rental contract, we adjust the matched sales price using a hedonic price index, ensuring that valuations reflect market conditions at the time of rent initiation. The methodology underlying this hedonic model is discussed in Section D.3.

Taxes: The annual property tax is determined based on the cadastral income, a government-assigned estimate of the property's annual rental value as it would have been in 1975. Given that this valuation is outdated, it is adjusted annually through indexation. The final property tax is calculated by multiplying the indexed cadastral income by a regional base rate, with additional surcharges imposed by provincial and municipal authorities.

Since our transaction dataset includes the cadastral income and precise location of each property, we can accurately compute the exact property tax.

Maintenance costs: Similar to the Netherlands, data on maintenance costs for Belgium are provided by Koeter Vastgoed Adviseurs. Koeter provided estimates of maintenance costs of 49 common Belgian property types, representative of the Belgian rental market as defined by us in terms of property age, property type, and energy efficiency. The cost estimates include replacement and maintenance activities with cycles up to 25 years. Consequently, high-cost items such as roof replacement have not been included, whereas expenses such as painting wooden window frames and replacing an electric boiler have been considered. To impute maintenance costs for all rental properties in our dataset, we estimate a regression model that is presented in Table A1.

Management costs: In Belgium, a landlord can use a *rentmeester* (property manager or estate manager) who acts on behalf of the landlord to oversee the management and administration of rental properties. Their responsibilities typically include rent collection, tenant communication, and ensuring compliance with lease agreements. Additionally, they handle property maintenance, coordinate repairs, and manage service contracts to preserve the value of the property. According to data from *Rentexpert*, a company that provides software for property managers, most companies charge a percentage fee between 6.5% and 7.5% plus VAT. Therefore, we assume a percentage fee of 7% plus VAT for management costs.

Turnover costs: To estimate tenant turnover, we utilize anonymous renter identifiers from the rent registry. A turnover event is defined in two ways: (i) when a renter appears in a new rental contract at a different address within a year of signing the initial contract, or (ii) when a new contract with a different tenant is registered at the same address within a year. When a landlord experiences a turnover, the associated costs consist of the broker fee for securing a new tenant and the expected vacancy period before the property is reoccupied. In Belgium, the typical broker fee for finding a new tenant is equivalent to one month's rent. To estimate the expected vacancy duration across rent deciles, we use data from ERA Belgium, the largest real estate brokerage franchise network in the country. This dataset provides information on the date a rental property becomes vacant. By merging these listings with the complete set of rental contracts, we can identify the exact start date of the subsequent lease. This allows us to precisely measure vacancy duration, which we compute separately for each rent decile.

Default costs: Since non-payment of rent is not directly observable in our dataset, we rely on external data sources to estimate default costs. To approximate the risk of non-payment, we use data from the 2008–2022 waves of the EU Statistics on Income and Living

Conditions (EU-SILC) Survey for Belgium. The data includes almost 21,000 households from Belgium who rent at market rents. The EU-SILC asks whether the household was unable to pay the rent on time in the past 12 months. Only situations when the household was unable to cover the costs due to financial difficulties are recorded. Instances where rent was covered through borrowing (e.g., from a bank, relatives, or friends) are not classified as late payments. We define a household as being at risk of default if it missed two or more rental payments within the past 12 months and compute default risk by rent decile. As no official data exist in Belgium on the share of total rent lost due to non-payment, we adopt the same assumptions as in the Netherlands, applying a 0.3% non-payment rate and 0.6% recovery cost rate in our baseline scenario. These costs are allocated across renters following the methodology used for the Netherlands, but using the share of renters in each rent decile that missed at least two rent payments, as reported in the EU-SILC data.

D.2 Data Processing

Our dataset from the rent registry includes all rental contracts that have been legally required to be registered since 2007. In addition to rental contracts for residential properties such as houses and apartments, the registry also contains contracts for garages and commercial leases, including office and retail spaces. Given that capital gains can be calculated up to 2021, we restrict our analysis to rental contracts from 2007 to 2021, encompassing 3,566,743 contracts. We exclude non-profit and public landlords to focus on the private rental market.

Since the rent registry does not provide detailed property characteristics or sufficient information to systematically identify houses and apartments for all years, we match rental contracts with the universe of housing transactions on the secondary market. This dataset is then merged with the universe of Energy Performance Certificates (EPCs) to incorporate energy efficiency information. Approximately 14% of observations are lost due to incorrect addresses or missing EPC records.

To remove outliers and implausible observations, we apply the following exclusion criteria: properties with more than 15 habitable rooms, 4 garages, 4 floors, or 4 bathrooms; as well as observations with rent, transaction price, or EPC values outside the P1–P99 range. The final sample is further restricted to contracts where 2023-equivalent rents fall between €300 and €3,000 per month and where rental yields range between 1% and 18%.

One potential challenge is that some landlords may rent out single rooms within a larger property, leading to cases where rental income is incorrectly matched with the sale price of the entire building. To address this issue and exclude rental contracts with implausible values, we estimate a regression model with log rents as a function of property and location characteristics. To capture location effects, we include neighborhood fixed effects if at least 30 observations are available within a neighborhood. If this threshold

is not met, we use municipality fixed effects, provided there are at least 30 observations within the municipality. Since this regression combines property characteristics from transaction data with rents from rental contracts, substantial deviations between the predicted and observed rent values are likely indicative of implausible observations. To filter out such cases, we exclude observations where the actual rent deviates by more than 0.3 log points from the predicted rent.

Table D1: Final Sample Selection

Filter	**Sample Size**
Full sample of rental contracts including commercial leases	3,566,743
Exclude non-profit and public landlords and business or public tenants	2,797,205
Matched with house or apartment sale	419,884
Matched with EPC	361,353
Exclude observations with implausible characteristics	331,796
Yield between 1-18%, rent between 300-3,000 euros	287,806
Rent residual within 0.3	198,351
Final Sample of Rent Contracts	**198,351**

D.3 Hedonic Model for Rents and Prices

To explain the distribution of rental prices in Belgium, we estimate a hedonic model with log rents as the dependent variable and detailed property and location characteristics as independent variables. We include neighborhood fixed effects when at least 30 observations are available within a given neighborhood. If this condition is not met, we apply municipality fixed effects, provided that at least 30 observations are available at the municipal level. The results with neighborhood fixed effects are shown in Table D2.

To estimate expected capital gains, we estimate a similar hedonic model with log price as dependent variable and property and location characteristics as independent variables. Again, we include neighborhood fixed effects when at least 30 observations are available within a given neighborhood or municipality fixed effects otherwise.

We apply the same procedure as for our US data and use T=5 as our default horizon to predict capital gains. For rental contracts in 2018, we set T=4, etc. Similar to the US, we match A such that the distribution of the hedonic capital gain lies closest to the repeat sales capital gains distribution. Figure D1 shows that $A = 0.25$ yields a distribution that is closest to the distribution from the repeat sales sample. Therefore, we use $A = 0.25$ in our baseline model for capital gains in Belgium.

Table D2: Hedonic valuation model for log(rent) in Belgium

	log(rent)	
	(1) Houses	(2) Apartments
log(useful surface)	0.110***	0.367***
	(0.00513)	(0.00266)
log(plot size)	0.0551***	
	(0.00181)	
log(built area)	-0.0294***	
	(0.00406)	
Number of floors	-0.0162***	
	(0.00149)	
Habitable attic dummy	0.00289	
	(0.00205)	
Habitable rooms	-0.00763***	0.0154***
	(0.000756)	(0.000869)
Garages	0.00781***	0.0233***
	(0.00219)	(0.00153)
Bathrooms	0.0136***	0.0487***
	(0.00216)	(0.00335)
Central heating dummy	0.0284***	0.00792***
	(0.00229)	(0.00297)
Age	-0.00470***	-0.00607***
	(0.000244)	(0.000136)
Age squared	0.0000353***	0.0000577***
	(0.00000239)	(0.00000159)
Construction year before 1850	-0.0430***	-0.0408***
	(0.00497)	(0.00869)
Construction year 1850-1874	-0.0463***	-0.0278***
	(0.00502)	(0.00958)
Construction year 1875-1899	-0.0546***	-0.0144**
	(0.00393)	(0.00597)
Construction year 1900-1918	-0.0396***	-0.00263
	(0.00377)	(0.00485)
Construction year 1919-1930	-0.0285***	-0.00774
	(0.00376)	(0.00496)
EPC	-0.000412***	-0.000139***
	(0.0000183)	(0.0000192)
EPC squared	0.000000201***	0.000000108***
	(1.85e-08)	(2.51e-08)
Construction type	X	
Year-Month FEs	X	X
Neighborhood FEs	X	X
Observations	80376	83244
R^2	0.534	0.675

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D3: Hedonic valuation model for log(price) in Belgium

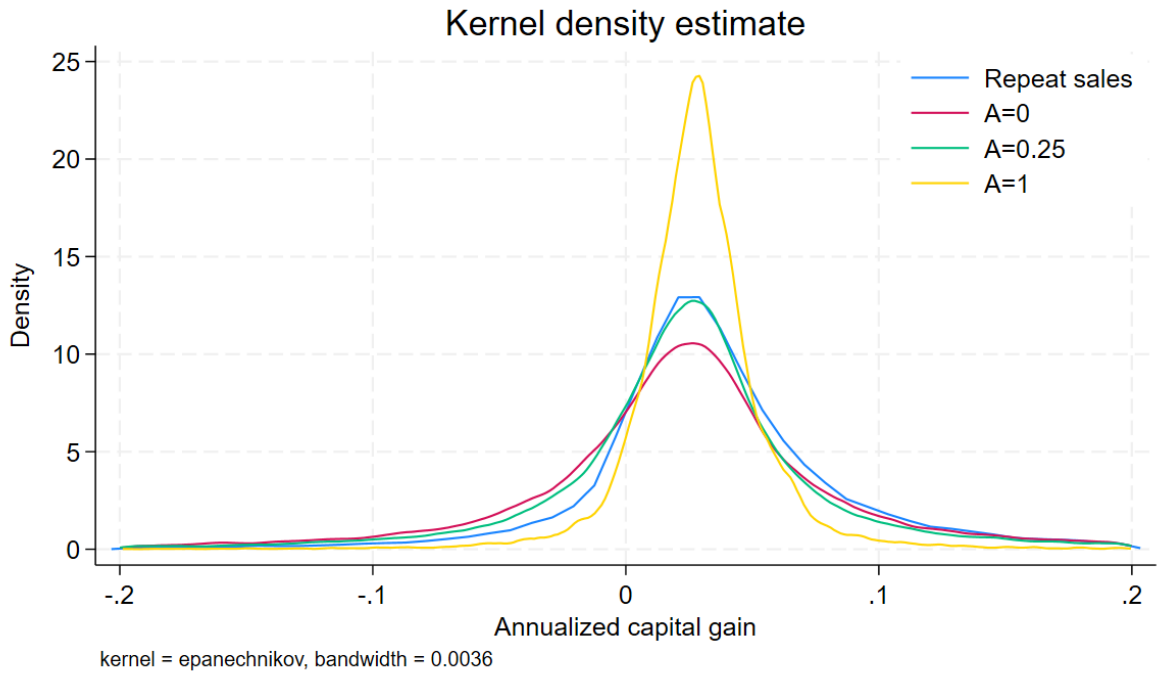
	log(sales price)	
	(1)	(2)
	Houses	Apartments
Rental property dummy	-0.0137***	-0.0350***
	(0.00106)	(0.00103)
log(useful surface)	0.377***	0.576***
	(0.00180)	(0.00160)
log(plot size)	0.138***	
	(0.000575)	
log(built area)	-0.0461***	
	(0.00140)	
Number of floors	0.000555	
	(0.000979)	
Habitable attic dummy	0.0308***	
	(0.000846)	
Habitable rooms	0.00505***	0.0407***
	(0.000263)	(0.000547)
Garages	-0.0000705	0.0396***
	(0.000686)	(0.00102)
Bathrooms	0.0479***	0.0722***
	(0.000828)	(0.00203)
Central heating dummy	0.112***	0.0525***
	(0.000823)	(0.00202)
Age	-0.0108***	-0.00574***
	(0.0000727)	(0.0000337)
Age squared	0.0000755***	0.00000305***
	(0.000000778)	(5.83e-08)
Construction year before 1850	-0.0401***	0.260***
	(0.00205)	(0.00607)
Construction year 1850-1874	-0.0533***	0.230***
	(0.00198)	(0.00638)
Construction year 1875-1899	-0.0434***	0.236***
	(0.00161)	(0.00402)
Construction year 1900-1918	-0.0263***	0.222***
	(0.00150)	(0.00310)
Construction year 1919-1930	-0.0190***	0.184***
	(0.00146)	(0.00316)
Construction types	X	X
Year FEs	X	X
Neighborhood FEs	X	X
Year FEs X Neighborhood	X	X
Observations	1024501	388391
R^2	0.775	0.753
Adjusted R^2	0.733	0.723

Standard errors in parentheses

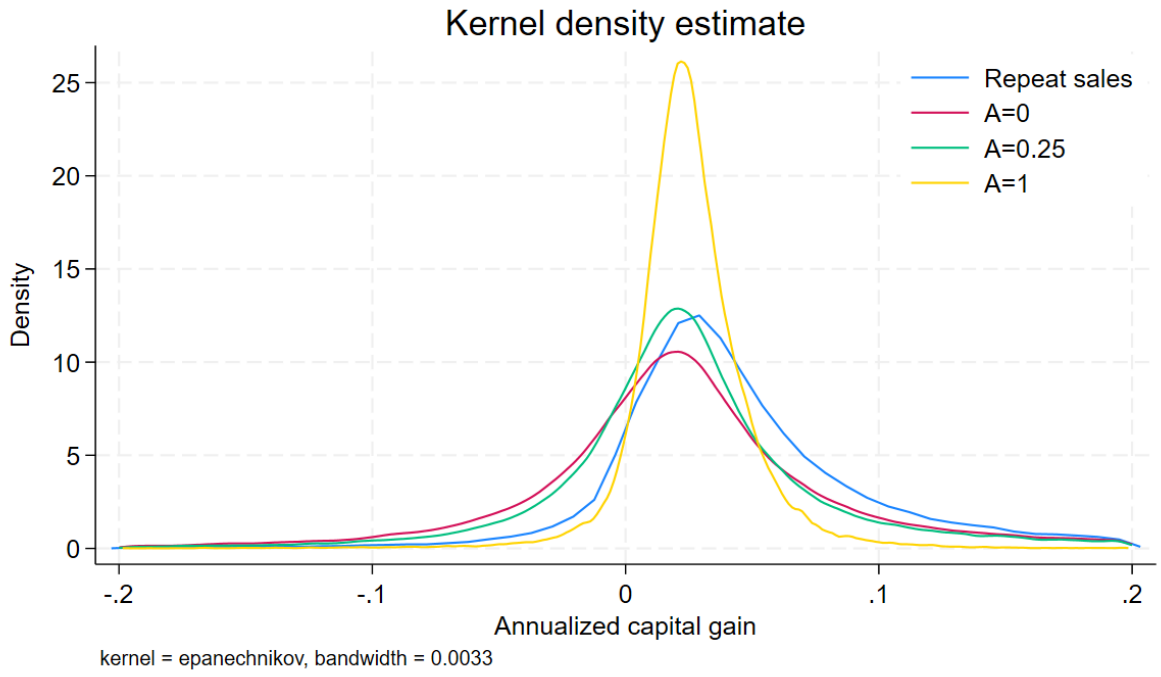
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure D1: Comparison of hedonic and realized repeat sales capital gain

a) Apartments



b) Single-family houses



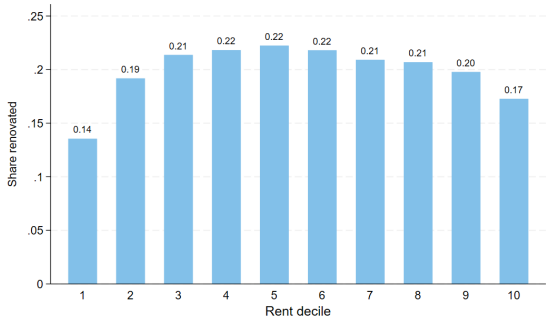
D.4 Renovation Expenditures

In this section, we study how major renovation expenditures differ between different rent segments in Belgium. In Belgium, each renovation that could impact the rental value of the house should be reported to the tax administration and renovations that require a building permit are automatically reported to the tax administration. We observe for each transaction the year of the most recent renovation that is reported to the tax administration. As we have matched the transactions with the rental contracts we can study how the occurrence of a renovation differs over different segments of the rental market and the years since the most recent renovation conditional on the property being renovated. The results are presented in Figure D2.

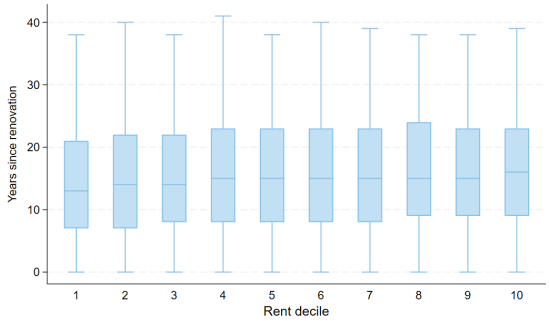
The results indicate that rental units in the lower segments are less likely to be renovated. Conditional on renovation, the number of years since the most recent renovation does not differ significantly between different segments.

Figure D2: Renovations By Rent Deciles: Belgium

(a) Renovated



(b) Renovation Age



D.5 Capital flows Analysis

Like we did for the U.S., we investigate how capital flows into the various segments of the rent distribution affect the relationship between the net rental yield, the capital gain return, and the total return on the one hand and the rent per unit on the other hand.

First, we estimate the slope of the relationship between Y_t^i and R_t^i (expressed in thousands of euros) where Y_t^i represents the net rental yield, capital gain return and total return. As the time span of our data is rather short in comparison to the US, we estimate the slope for each year t and district.

Second, the Belgian data allows us to calculate the net investor flows. The net investor flow is defined as the number of purchases for which we observe a rental contract after the transaction but not before minus the number of sales for which we observe a rental contract before but not after the transaction. We calculate the net investor flows in the

upper and lower segments of the rental market for each district and year. We lag these variables by one year.

Finally, we estimate a regression with the slope as the dependent variable and the lagged net investor flows in the bottom and top 20% of the rent distribution as the independent variables. The results are presented in Table D4. The results indicate that increased investments in the lower segments make the total return slope less negative whereas investments in the top segments make the slope more negative. Similar to the US, most of the effect on total returns comes through the capital gain return.

Table D4: Investor flows, Belgium

	(1) Net Rental Yield	(2) Capital Gain Return	(3) Total Return
Net Flow Bottom 20	0.00	0.21***	0.20***
	(0.01)	(0.05)	(0.06)
Net Flow Top 20	-0.04***	-0.04***	-0.08***
	(0.01)	(0.01)	(0.02)
Constant	-3.39***	-0.03	-3.76***
	(0.11)	(0.70)	(0.77)
Observations	426	426	426
R-squared	0.10	0.02	0.02

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

E Regulatory Risk

E.1 Renter Protection Score with GenAI

We construct a Renter Protection Score (RPS) that evaluates renter protection laws across US states, leveraging ChatGPT's analytical capabilities to evaluate and assign scores based on comprehensive state-by-state landlord-tenant law data from IProperty Management.[38] Our RPS consists of five components: Landlord-Tenant Rights, Security Deposit Returns, Eviction Process, Rent Control, and Early Lease Termination. Each component is worth a maximum of 6 points. Each component itself is broken down into three subcomponents, scored on a scale from 0 to 2 points each. Table E1 provides an overview of the scoring criteria. The LLM scores each state plus Washington DC on each of the 15 criteria. We aggregate to a numerical score that is between 0 and 30. States earn points based on the strength and effectiveness of their tenant protection measures. Higher scores indicate stronger renter safeguards, with points awarded based on standardized criteria applied consistently across all states.

[38] https://ipropertymanagement.com/

The results are consistent with anecdotal evidence. Three states score the maximum 30 points: California, Oregon, and New York. Washington D.C. scores 29, New Jersey 26. The weakest renter protections are in Arkansas, which scores 9, followed by Wyoming (11), Oklahoma (12), and West Virginia (13). Our RPS score has a correlation of 50% and a rank correlation of 55% with the renter protection score of McCollum and Milcheva (2023).

Table E1: Categories and Scoring Criteria for the Renter Protect Score System

Main Components (0-6 points each)	**Subcomponents** (0-2 points each)
Tenant Rights	- Clarity of Laws: Are laws clearly defined and understandable? - Accessibility: Are resources about tenant rights readily available? - Enforcement: Are tenant rights effectively enforced by agencies?
Security Deposits	- Deposit Limits: Are there reasonable caps on security deposits? - Timelines: Are strict deadlines enforced for deposit returns? - Enforcement: Are there penalties for non-compliance?
Eviction Process	- Notice Periods: Are generous notice periods required? - Legal Recourse: Can tenants contest evictions effectively? - Protections: Are anti-retaliation safeguards in place?
Rent Control	- Presence of Policies: Are rent control policies in place? - Effectiveness: Do policies prevent unjustified rent increases? - Scope: Are laws broadly applicable to housing types and areas?
Breaking a Lease	- Permissible Reasons: Are tenants allowed to break leases for valid reasons? - Penalty Severity: Are penalties for early termination lenient? - Ease of Process: Is the termination process tenant-friendly?

If areas with strong renter rights implied more risk for landlords, and this regulatory risk were compensated, either in the form of a higher rental yield or in the form of a higher capital gain return, we would expect the total return to be higher in the high RPS group than in the low RPS group. We would also expect the regulatory risk effect to be particularly pronounced for lower-rent properties, whose protection regulators may care about disproportionately. This would show up as a more steeply downward-sloping relationship between rent levels and net yields/total returns in high RPS states than in low RPS states.

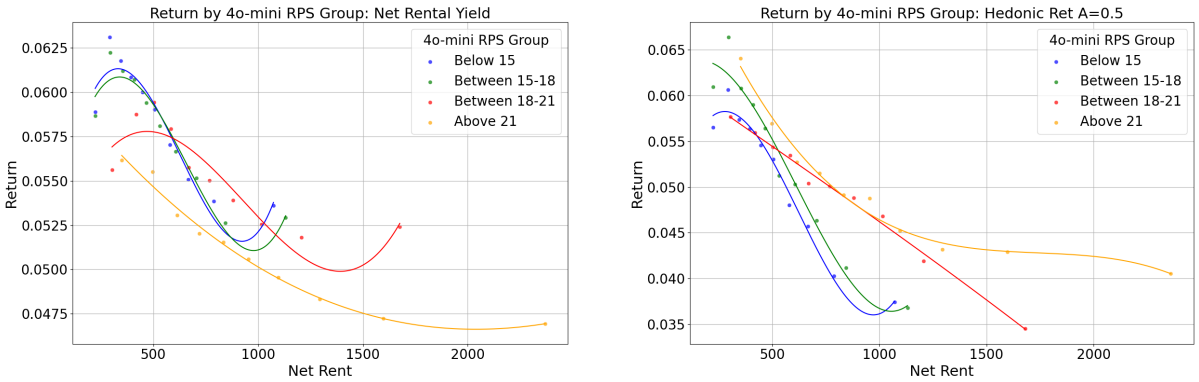
We sort states into 4 groups: 14 states with RPS below 15, 14 states with RPS between 15 and 18, 13 states with RPS between 18 and 21, and 10 states with RPS above 21. We plot the dividend yield, capital gain return, and total return by NOI/unit for each of these groups in Panels A through C of Figure E1.

We find substantial differences between low-RPS and high-RPS states. In terms of cap rates, Panel A shows lower cap rates in high-RPS states than in low-RPS states, holding

fixed NCF per unit. The D1-D10 spread in cap rates is about 1% in both the lowest two RPS groups and the highest RPS group. However, the cap rate curve is clearly steeper in the lowest two RPS groups. If high cap rates reflected high risk to landlords from strong tenant protections, as this risk was particularly strong in low-rent properties, we would expect to see the steepest cap rate curves in the high-RPS states. This is exactly the *opposite* of what we see in the data.

We obtain the same conclusion from the capital gain return in Panel B. Capital gain returns on low-income properties are only slightly higher in high-RPS states, and the capital gain curve is flatter than in the low-RPS states. The extra capital gain return on a unit with a NCF of $10,000 per year relative to a unit with a NCF of $5,000 per year is largest in the low-RPS states. This conclusion carries over to the total return curve in Panel C.

Figure E1: Net Rental Yield and Capital Gain Yield by Tenant Protection Groups



(a) Net Rental Yield

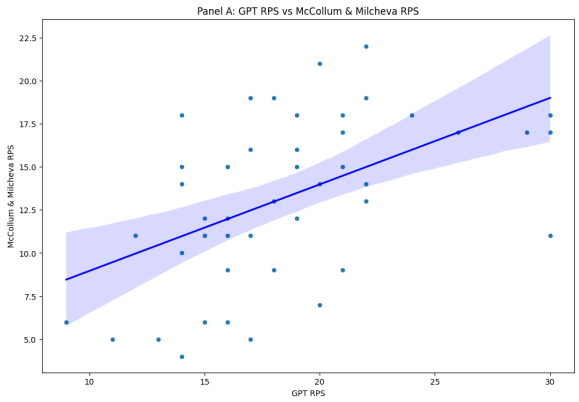
(b) Capital Gain Yield

Notes: Panel (a) shows the net rental yield and panel (b) shows the capital gain yield plotted against net rent, measured as real underwritten NOI per unit. The 50 US states plus Washington DC are sorted into four groups based on their Renter Protection Score (RPS). The RPS is created from state laws with the help of ChatGPT 4o-mini. The capital gain yield assumes the value $A = 0.5$.

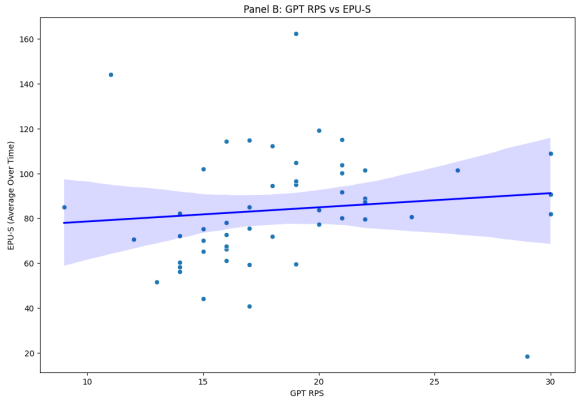
McCollum and Milcheva (2023) similarly argue that multifamily housing is *less* risky in high-RPS states, as reflected in a lower cap rate, a lower mortgage delinquency rate, and lower NOI growth volatility. While our own RPS score only has a correlation of 50% with theirs (and disagrees substantially on the level of tenant protections in states such as New York, South Dakota, Rhode Island, and to a lesser degree California), we ultimately arrive at a similar conclusion. Figure E2 shows scatter plots of our RPS score against the McCollum and Milcheva (2023) score, the EPU-S, cash flow growth volatility, and mortgage delinquency.

We explore robustness of our renter protection score, computed using ChatGPT 4o, to alternative GenAI algorithms: Claude, Gemini, and ChatGPT o1. All four methods deliver very similar results.

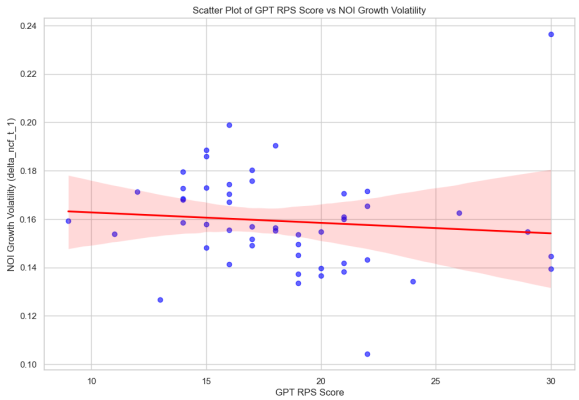
Figure E2: Comparing Our RPS to Other Tenant Risk Metrics



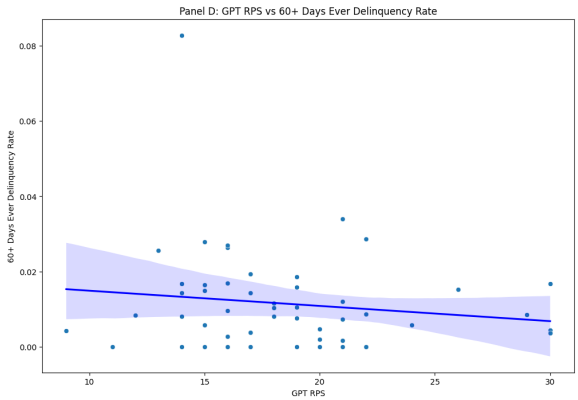
(a) Our RPS vs McCollum & Milcheva RPS



(b) Our RPS vs EPU-S



(c) Our RPS vs 1yr-NOI Growth Vol



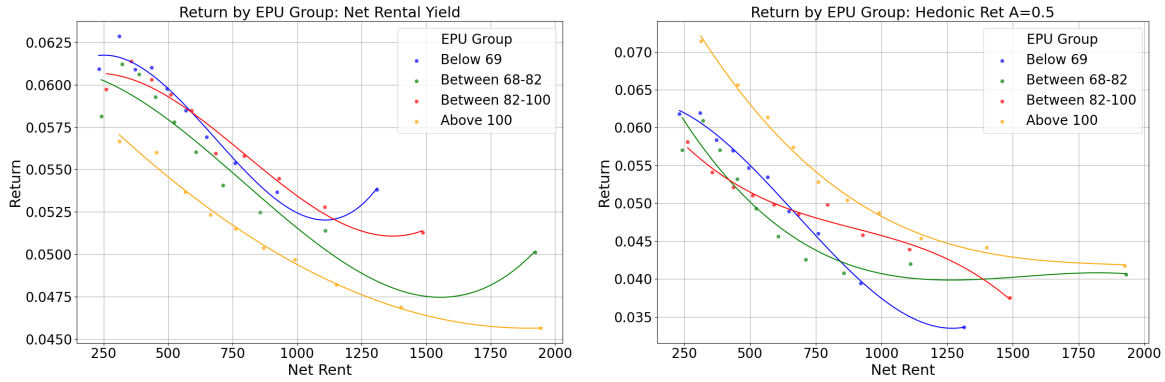
(d) Our RPS vs 60+ Day Delinq. Rate

E.2 State-level Economic Policy Uncertainty

An alternative measure of regulatory risk is the *state-level* economic policy uncertainty index of Baker et al. (2023). Based on articles in 3,500 local newspapers, these authors compile a time-series of economic policy uncertainty about local economic policy for each US state (their EPU-S series). Local economic policy uncertainty pertains to uncertainty around local elections, but also zoning policy is an important source of heterogeneity and dynamics. We form four groups of states based on the (time-series average) level of the state-level EPU.

We plot the dividend yield and capital gain yield by NOI/unit for each of these groups in Panels A and B of Figure E3. The total return graph is included in the main text. Overall, we see little difference in the level or slope of the rental yield, capital gain, or total return across these groups. It is possible that this measure is too broad to capture the rental risk that is relevant to landlords, but at least there is nothing in the measure that suggests that state-level policy uncertainty is associated with meaningfully higher returns or a steeper return gradient. If anything, the evidence points in the other direction.

Figure E3: Net Rental Yield and Return By Local Economic Policy Uncertainty



(a) Net Rental Yield

(b) Capital Gain Yield

Notes: Panel (a) shows the net rental yield and panel (b) the capital gain yield plotted against the net rent, measured as the real NOI/unit. The 50 US states plus Washington DC are sorted into four groups based on their time-series average level of local economic policy uncertainty (EPU-S). The EPU-S data are from Baker, Bloom, and Davis (2022). The capital gain yield assumes $A = 0.5$.

E.3 Risk of Pro-Tenant Regulation: Democratic Control

We next explore whether uncertainty about future regulatory changes affects returns on multifamily properties. The idea is that landlords may demand compensation for the risk that future policy changes could affect their property rights or cash flows. We use the political control structure of state governments at the time of loan origination to measure the likelihood and direction of future regulatory changes. Democratic party control of the state government is much more likely to result in pro-tenant regulatory change than Republican control.

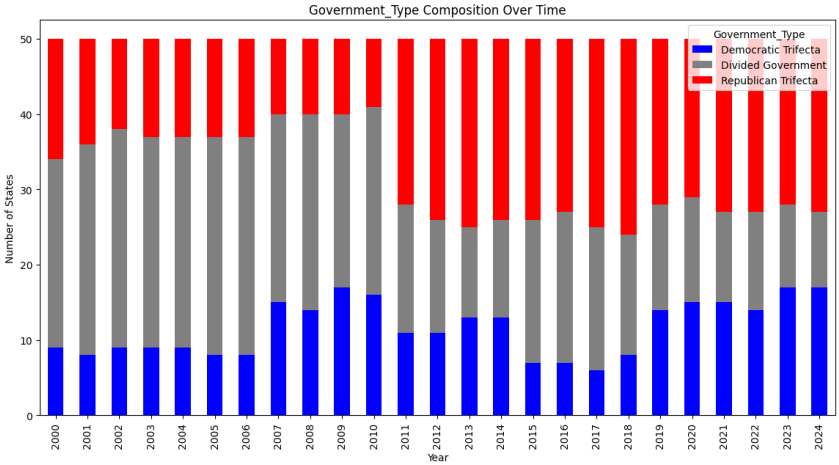
Data from Ballotpedia indicate whether Democrats or Republicans controlled the governorship, state senate, and state house in each state from 1994 to 2024.[39] We classify states into three groups: Democratic trifecta (Democratic control of executive and both legislative branches), Republican trifecta (complete Republican control), and divided government (split control between parties). Figure E4 shows the share of US states in each of these three groups every year from 2000 until 2024.

Figure E5 shows how our main graphs of interest across political regimes. Panel A shows that properties in Republican trifecta and divided government states have higher cap rates than those in Democratic trifecta states. The cap rate-NOI relationship is steepest in Republican trifecta states, while both divided government and Democratic trifecta states exhibit relatively flat relationships.

Panel B reveals that total returns are highest in Republican trifecta states, followed by divided government states, with Democratic trifecta states showing the lowest returns. The total return-NOI relationship is steepest in Republican trifecta states, while both

[39].https://ballotpedia.org/Historical_and_potential_changes_in_trifectas.

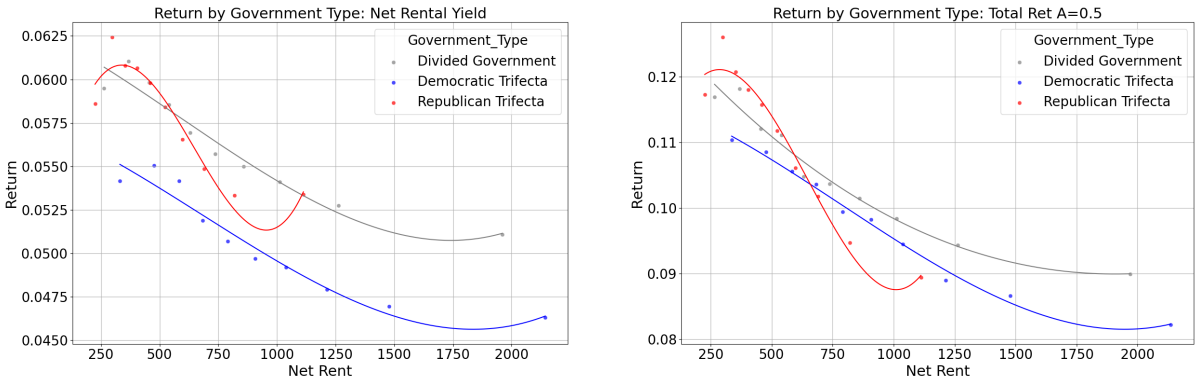
Figure E4: Composition of State Political Control from 2000 to 2024



divided government and Democratic trifecta states show similarly flat relationships.

These findings present a puzzle for the regulatory risk hypothesis. If Democratic control signals a higher probability of future tenant-friendly regulation, we might expect to see both a higher cap rate and return level and a more steeply declining relationship in rent levels, as low-NOI properties may face greater regulatory scrutiny. Instead, we observe the opposite pattern.

Figure E5: Cap Rates and Returns by Political Control Groups



(a) Net Rental Yield

(b) Total Return

Notes: Returns by NOI/Unit across political control groups. Panel A shows Cap Rate and Panel B shows Total Return for Democratic trifecta, Republican trifecta, and divided government states. States are classified based on political control at the time of loan origination. The curves show fitted splines through binscatter points.

E.4 Exposure to Rent Ceilings in The Netherlands

We finally explore variation in rent ceilings across regions in the Netherlands, in line with our cross-state analysis for the United States. In the Netherlands, private rentals are subject to rent control, although these regulations were typically not adhered to in the

private rental market due to limited enforcement. Additionally, landlords could legally charge higher market prices as long as the tenant did not appeal within six months of signing the contract.

The government sets rent ceilings based on a property-level score (*huurpunten*), which reflects the property's characteristics. The rent ceiling applies to properties with a score below the "rent liberalization" threshold, which covers around 50% of private properties. However, these scores are not officially registered and are based on property characteristics that can only be partially observed in administrative sources.

Given the lack of data on property scores in the private rental market, we use data on property scores from housing associations. In general, housing associations record the scores for all their properties. To estimate the scores for private rental properties, we develop a model that predicts the number of points assigned to a property using a rich set of property characteristics, which are also available for private rental units. These characteristics include the property's square meters (both linear and log), tax value (both linear and log), property age (both linear and log), the number of units, building vintage, and municipality. This model explains approximately 70% of the variation in observed property scores.

We use the model coefficients to extrapolate these scores to the private rental sector and translate the predicted property scores into the maximum legal rents (at 2023 prices). In our main measure, we aggregate the log difference between the rent ceiling and the observed rent per COROP region (NUTS-3) and categorize regions into three groups. Figure 3d in the main text plots the relationship between rent levels and returns for these three groups of regions.

In 2024, the rent ceilings based on this system were declared binding and extended to approximately 90% of the private rental sector. Note that in our main estimates, we do not incorporate the actual implemented reform. This reform was only proposed in 2023, and the possibility of a tightening of rent controls was first announced in December 2021 when a new Dutch government was formed. Since our data refer to rental contracts signed between January 2017 and December 2021, they could not possibly account for the actual reform. Various modifications were also introduced to the property score system, while the property scores in our data are still based on the old system. A rough back-of-the-envelope calculation nonetheless suggests the reform likely had the most significant impact on properties in D10. This confirms that the risk of rent regulation is unlikely to explain our slope.