Segmentation and Returns on Rental Housing*

Jessie Handbury[1], Samuel K. Hughes[2], and Benjamin J. Keys[1]

[1]The Wharton School and NBER
[2]Federal Reserve Board

December 2025

Preliminary Draft – Please do not cite without permission

Abstract

This paper studies rents and landlords' returns across the property quality distribution. We establish three key facts: (1) asking rents increased more on low-quality properties post-Great Recession; (2) realized internal rates of return (IRR) also rose more in the lower-quality segment; and (3) most new housing supply was added in higher-quality segments. These facts are closely related. New supply reduces rent growth and returns for existing high-quality housing. Using a city-level shift-share design capturing local demand growth, we show that low-end rents and returns respond more strongly to demand shocks, and find this differential response is largely due to the new supply channel.

*Thank you to Ailey Fang, Kathy Bi, Ezeriki Emetonjor, Shusheng Zhong & Hongrui He for innovative and high-quality research and data construction assistance. We thank seminar and session participants at the NBER Real Estate Summer Institute, the Wharton School, Bank of England, and UEA for thoughtful suggestions. Handbury and Keys thank the research sponsors of Wharton's Zell/Lurie Real Estate Center for support. The content is the sole responsibility of the authors. The views expressed in this paper are solely those of the authors and do not necessarily reflect the opinions of the Federal Reserve Board, or the Federal Reserve System.

1 Introduction

Since the Great Recession, concerns about housing affordability have remained persistent. Over the past two decades, real median rents increased nearly 20% faster than household incomes—a trend that disproportionately affects lower-income renters.[1] While a popular narrative attributes rising rents to 'greedy' landlords, it's unclear whether these increases reflect growing profits or are merely offsetting higher acquisition and operating costs. Even if landlord returns have risen, the source of their increased market power remains uncertain. Activists often point to the proliferation of new 'luxury' apartment buildings as displacing affordable housing, but landlords argue that rising rents stem from heightened cash flow risks—such as tenant non-payment and regulatory pressures like rent control—which, in turn, drive up their cost of capital.

This paper asks how rents have changed since the Great Recession, whether landlords have captured these gains, and why. We focus especially on differences across the quality segments of the U.S. multifamily rental market. Our first contribution is to assemble detailed new data on landlord returns, rents, and new supply in each tercile of the per-unit rental income distribution (relative to other properties in the same city). With these data, we establish new facts: since the Great Recession, low-quality units experienced more growth in rents and prices; this outsized growth was captured by greater landlord returns and sustained by limited new supply of low-quality units. We reconcile these facts in a model of a segmented rental housing market. The model predicts that high-quality housing supply and, conversely, low-quality rents and prices, will be more elastic to variation in aggregate demand. We test these predictions using cross-CBSA variation in rents, prices, and local labor market shocks. Local demand shocks induce more rent and return growth for low- compared to high-quality rental units, but less so once we control for new rental housing supply. We conclude that the effect of new supply explains 60 to 70% of the differential return responsiveness of low- versus high-quality units to demand shocks, and about 60% of the differential rent growth.

We first show that, in a panel of large market-rate apartment buildings in the 48 largest cities in the U.S., asking rents have risen disproportionately for lower-quality buildings relative to higher-quality buildings between 2005 and 2019. Asking rents grew in all quality segments, but 11 percentage points faster in the lowest quality tercile, where nominal rents grew by 48 percent, than in the top quality tercile, where nominal rents grew 37 percent. This divergence grew the most during the Great Recession and in the late 2010s.

To study landlords' returns, we construct a property-level panel dataset on apartment buildings from two main sources. First, we construct a panel of annual net operating incomes from commercial mortgage-backed securities (CMBS) reporting information, obtained from Trepp. Second, we observe transaction values for properties placed in CMBS at both the origination of a mortgage term, and the end of the mortgage term (at prepayment or maturity) from apartment transaction data obtained from Yardi. The panel construction allows us to measure property-specific returns over landlords' holding periods, which we specify as a mortgage term in Trepp or the years between transactions in

[1] Joint Center on Housing Studies, 2024.
https://www.jchs.harvard.edu/blog/high-housing-costs-are-consuming-household-incomes.

Yardi for each property. Using the net income flows and prices at the beginning and end of holding periods, we construct realized internal rates of return (IRR) at the property-by-holding period level.

With this panel of annual net operating incomes and transaction data, we measure landlords' returns over time and across the quality distribution. Generally, lower-quality apartments have higher IRRs, partly reflecting the lower cost basis of these buildings at the beginning of their holding period. Further, while IRRs increased on average between 2005 and 2019, we find that they increased most for lower-quality buildings. The average return for a property in the lowest-quality tercile grew about 5 percentage points more than the top-quality tercile. In part, the overall rise in IRRs reflects a decline in capitalization rates (an inverse indicator of how much capital markets value cash flows) driven in part by declining interest rates. A decomposition shows, however, that the widening quality gap in returns is driven equally by outsized low-end growth in net operating income combined with outsized growth in how capital markets' valued those cash flows. This result leads us to conclude that there were real economic shocks disproportionately affecting lower-quality buildings relative to higher-quality buildings, not simply a differential impact of national-level cap rate compression on segment-specific IRRs. We find outsized rent and return growth at the low end, even after accounting for potential differences in propensity to renovate, loan delinquency, and differential demand growth.

A plausible explanation for the outsized rent and return growth at the lower end of the market is limited supply growth in response to the sustained long-run demand growth in rental markets experienced between the Great Recession and the COVID-19 pandemic. The final descriptive fact that we highlight is that most new rental housing indeed enters at the top of the existing rent distribution. Anecdotally, this fact is well-recognized but there is little work quantifying this phenomenon or attempting to understand its causes and consequences. We use American Community Survey data from the Census Bureau to study where newly constructed rental units enter the rent distribution over time and across cities. Since the Great Recession, 50 to 60% of new rental units enter the top quartile of the unconditional city rent distribution. This skew is even higher among market-rate buildings: We find that 80% of new market rate units in large cities enter the top quartile of the rent distribution.

We reconcile these facts in a model of a segmented rental housing market. In the model, housing supply is determined by developer decisions and depreciation. Heterogeneous developers produce housing with a fixed land input and unit capital requirement that is greater for high- than low-quality housing. Housing units per unit of land are capped by zoning regulations. We demonstrate that more developers will expect non-negative profits and build in the high-quality segment than in the low-quality segment as long as the expected premium paid for high-quality units exceeds the excess unit capital requirement to build high-quality units. The model predicts that high-quality housing supply will be more elastic to variation in aggregate demand. The high-quality sector will have more new supply, attenuating price and rent growth, in response to positive demand shocks. By contrast, the low-quality sector will have more new supply (and greater price and rent declines) when there are negative demand shocks that stem all new development, because depreciated high-quality housing will continue to add to the low-quality supply through filtering.

We test these predictions using cross-CBSA variation in rents, prices, and local labor market shocks. Are rental incomes and returns on lower-quality buildings more responsive to local labor

market shocks than those on higher-quality buildings? We find that NOI from low-quality properties increases by 30 percent more than high-quality properties for the same increase in city-level employment.

This differential responsiveness may reflect confounding factors that drive city-level employment growth and also explain differential rent growth. To identify the causal effects of local employment growth, we use a Bartik-style shift-share instrument as a proxy for local employment growth and find that the elasticities are both biased towards zero (coefficients on the Bartik measure are 2.2 pp and 1.7 pp at the low- and high-end respectively) but the differential across the quality distribution persists. The quality differential is even larger when measuring the responsiveness of IRR to employment growth: Low-end property returns are 1.2 pp higher with each 1 percent higher employment growth, while high-end property returns are only 0.1 higher. In order to argue that the differences in returns are caused by labor demand shocks, we show that the results in the low-high return gap is comparable using Bartik shift-share in employment.

We then estimate how much of the differential rent and return responsiveness is due to the more elastic supply of higher-quality rental housing. Specifically, we study how much of the elasticity differentials estimated above persist when controlling for new construction and the share of new construction in the higher quality segment. We find that new rental housing supply is negatively associated with rental income growth and returns for existing owners. Controlling for this correlation disproportionately lowers returns for higher-quality properties, but has very little effect on the low-quality segment. These new housing supply variables absorb 57% of the differential effects of employment growth on rental income and 62% of the differential effects on returns.

As a product of this analysis, we show that city-level new housing supply is negatively correlated with subsequent rental income growth and returns, but primarily (or exclusively, in the case of returns) for higher quality housing. We also show that the share of new housing entering the top of the existing rent distribution is positively related to rental income and returns for lower-quality properties.

Beyond differential supply elasticities, other factors that could explain the outsized returns at the low end of the rental housing market include differential changes in asset markets. Ghent (2021), for example, shows that market liquidity and transaction volume varies across cities in a way that is related to delegated investors (or the market signals dynamics of developers in Büchler, Van de Minne, et al. 2020). It is possible that the higher returns at the low-end of the market result from higher returns demanded in capital markets due to weak growth expectations or greater perceived risk. We allow for granular controls for changes in the cost of capital and find little evidence that these costs can explain a meaningful portion of the differential returns. We also consider that capital improvements could play a role in outsized returns. Using a proxy for substantive quality improvements—property renovations—shows that the improvement channel plays little or no role in the dynamics of the difference in returns by segment.

This paper presents new evidence on differential returns to real estate investment that offers a potential explanation for the severity of the affordability crisis among the lowest-income renters. Previous work establishes that commercial rental yields and returns vary across rental housing types (Adams and Verbrugge 2021), across geographic markets (Amaral et al. 2024; Eichholtz et al. 2021;

Korevaar 2023; Plazzi et al. 2010; Samuels 2024) and over time (Peng and Thibodeau 2020; Reher 2021). We compile timely new data to document variation in rents and returns across the quality distribution and find that returns are higher on properties with lower-priced units. This finding contrasts with results from related work studying single-family house prices over the same period. Peng and Zhang (2021) found that higher-priced homes appreciated more between 2000 and 2015, while Hartman-Glaser and Mann (2017) showed that returns are higher in high-income areas, where prices also tend to be less volatile.

In related work measuring rent growth across segments, Molloy (2024) finds that lower-income households reported higher rent growth in the American Housing Survey between 1985 and 2021. Halket et al. (2023), meanwhile, find that single-family homes in low-income submarkets had higher yields and returns. Controlling for heterogeneity by looking at rental and price appreciation within property, we find that the low-end of the multifamily market – the focus of the debate over tenant affordability and "landlord greed" – has also seen outsized rental growth and returns since the Great Recession.

In contemporaneous work, Damen et al. (2025) document that this fact holds internationally – in the U.S., Belgium, and the Netherlands – and using complementary measurement approaches. Their primary measurement approach uses hedonic valuation indices based on values and income at underwriting to measure average total returns. We show that these results are robust when returns are measured using transaction prices and net income realized over full hold periods. Leveraging these rich data, we additionally turn to examining geographic variation in market dynamics and the role of new supply.

Our paper focuses on the ability of rental market fundamentals to generate outsized high-end returns. We develop a theoretical framework showing that fundamentals can drive more development of high-end than low-end properties in growing markets. The relative inelasticity of supply at the low-end of the market generates outsized rent growth. Damen et al. (2025), in contrast, pursue the complementary question of why these return differentials are not eliminated through arbitrage in financial markets—in other words, why capital does not flow into low-end property investments in a way that moderates returns. Their analysis highlights the importance of investor segmentation, with reputational concerns deterring large investors from participating in low-end property investments. Taken together, the papers suggest that both developers' and investors' orientation toward the high end of the market contribute to outsized returns at the low end.

The differential elasticity of rental income at the low vs. the high end of the market is consistent with results from Watson and Ziv (2025b), who show market power in the multifamily rental market. Our results suggest that this market power derives from limited new supply of lower-quality rental units, consistent with Watson and Ziv (2025a)'s evidence that zoning constraints generate market power and increase rents (see also Anagol et al. (2025)). Most research on the effect of new supply on rents studies the effect of a new building on surrounding asking rents (e.g., Asquith et al. (2023), Li (2022), and Pennington (2021)). Using data on city-level rents and prices, Büchler, Ehrlich, et al. (2021) argue for a cap rate amplification channel for housing supply—if an increase in rental income subsequently increases prices by more than one-for-one, rent increases will result in more housing supply than a similar-sized increase in prices alone. Our paper is among the first to show how new

supply relates to rental income as well as returns.

Finally, we provide novel evidence consistent with findings on quality segmentation derived from assignment models in the housing market (Anenberg and Kung 2020; Landvoigt et al. 2015; Nathanson 2023; Wang 2023). We present new results on rental market segmentation consistent with endogeneity in the quality-segment-specific supply-side response, consistent with models of endogenous supply in the owner-occupied housing market as in Wang (2023) and Ma (2025).

2 Data

We rely on property-level financial data for multifamily rental properties collected by servicers on mortgages placed in commercial mortgage-backed securities and maintained by Trepp. The financial data are collected quarterly by servicers, though most borrowers only report updates to their financial outcomes on an annual basis. There are 105,000 multifamily properties in the Trepp dataset, covering 2000–2020. The elements of financial data relevant for this paper are the total annual revenue and operating costs at the property level, the term of each mortgage, and the appraised or transacted value at the time of mortgage origination. We limit attention to borrower-reported full fiscal year financials. Revenue line items can include gross rents, vacancy and credit losses, and parking, laundry, and vending revenue. Operating expenses include management fees, utilities, taxes, maintenance and repairs, insurance, payroll and benefits, advertising and marketing, professional fees, and general and administrative costs. Landlords report the year of completion for any renovation on each property.[2]

The two pieces of information needed to construct landlord returns are property value at the beginning and end of their holding period, and a panel of net income during each holding period. At the beginning of each hold period, we use the appraised or transacted values reported in Trepp. The mortgage data generally does not include valuation at the end of hold periods when a property's mortgage prepays or matures (e.g., by selling or refinancing). We therefore complement the Trepp data with multifamily property transaction values from Yardi. The Yardi data includes transaction values for 49,000 properties, 21,604 of which we can match to the Trepp data using either exact or fuzzy matching on address strings.[3] For these matched properties, we use the Yardi transaction value at the end of each hold period. With this merged data, we observe actual price appreciation over our sample of property-hold periods.

[2]We use this information to explore the role of renovation as a proxy for capital expenditures which are not well-captured in the CMBS datasets. We conduct two further analyses: 1) we include distributed leads and lags around the year of renovation in our main regression results; and 2) in Appendix C.1, we estimate our main regressions in samples with only renovated or only non-renovated properties to show whether the main coefficients change because of renovations. Hughes 2024 shows the average effects of renovation. In those specifications, the average renovation raises annual revenues by 4 to 5%, suggesting that it is only in rare cases that a renovated property will move from "Low" to "High" quality segment, where high-end revenues are about 60-70% higher, on average.

[3]We exactly match 16,112 properties in the Yardi and Trepp datasets based on address strings, and match a further 4,784 properties through fuzzy string matching. We believe this is not an unreasonable match rate given that Trepp only includes properties in CMBS data, which pre-Great Recession was around 20% of the overall CRE mortgage market (Ambrose et al. 2022). Fannie Mae and Freddie Mac multifamily securitization expanded after the Great Recession, so our highest match rate is during the 2010–2019 period. As Yardi sources data from a mix of surveys and voluntary contributions from property managers and owners, it is unsurprising that some property transactions recorded in Trepp (particularly refinancings) are not captured in Yardi.

Even among properties that consistently report financial data to their CMBS servicer, there is often at least one year of missing net operating income at the beginning or ending of their hold period (e.g., landlords may not be required to report their annual NOI after they sell the property mid-year). Other properties are often missing NOI for some years in the middle of their hold period. We interpolate missing NOI between observations for each property when possible, and extrapolate NOI over missing observations when necessary using the national average NOI growth rate by quality segment (winsorizing at the 5th and 95th percentiles). We only keep property hold-period observations where NOI is not missing for more than 50% of the hold period. After all of these restrictions, we have a sample of 4,990 property-hold periods (4,830 properties).

2.1 Other Datasets

We supplement the Trepp and Yardi data with additional data from Moody's REIS platform, which collects quarterly asking rents and other information from property owners and managers. Our REIS dataset is a set of the near-universe of market-rate apartment buildings in the 50 largest cities in the U.S. over the 2005–2019 period. We study within-property changes in asking rents using the Q4 values of market asking rents.

To study geographic variation, we merge in the CBSA-level housing supply elasticity measures from Baum-Snow and Han (2024). We also use data from the Quarterly Census of Employment and Wages (QCEW) on average wage and employment growth across counties, and employment, population, income, and housing supply measures from the Census Bureau's American Community Survey (ACS) and Decennial Census.[4]

2.2 Data Description

Table 1 shows summary statistics for the sample of property-level holding periods. The average property has a per unit monthly rental income (i.e., annual revenue per unit divided by 12) of $1,029, and per unit monthly operating expenses of $472. In the average hold period, net operating income (NOI) grows by 4.2% annually.[5] The average property begins with an average value of $23.5 million, and ends their hold period with an average value of $27 million, growing 15% over an average period of 4.3 years. The combination of net operating income and growing property values leads to an average internal rate of return (IRR) of 15% during our period of analysis.

To study the role of quality segmentation, we create three categories for quality based on each property's average monthly revenue per unit in the first year of their hold period. In that first year, we assign a quality designation of Low, Mid, or Top quality based on whether the property's rental income per unit is in the bottom, middle, or top tercile within their city's revenue per unit distribution. The distribution of rental income per unit is taken using all existing multifamily properties in CMBS loans reporting their financial performance.

[4]Steven Ruggles, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Renae Rodgers, and Megan Schouweiler. IPUMS USA: Version 15.0 [dataset]. Minneapolis, MN: IPUMS, 2024. https://doi.org/10.18128/D010.V15.0

[5]As we impute some values of NOI, the average revenue net of expenses in the raw data does not exactly match the NOI reported in the table (which is used to calculate returns).

Table 1 shows that this method of measuring quality segmentation relates meaningful variation. Lower-quality properties are older on average relative to the top quality properties (built in 1976 vs. 1993), and relatively smaller (215 vs. 261 units). As we would expect, lower-quality properties have lower average values at origination ($12 million vs. $36 million).

From this descriptive table, we can see already that high-quality properties have higher operating incomes and higher operating margins (NOI divided by revenue), but lower average IRRs. Higher-quality properties take 57% of revenues as net income (before debt service), while lower-quality properties take 50% of revenues as net income. We observe no evidence that this quality gap is related to differential hold period length or due to differential probability of renovations around the beginning or during the hold period.

We briefly discuss the data representativeness and selection into CMBS. Where multifamily properties do or do not borrow from lenders who securitize their mortgages is not the main object of this paper, but is relevant insofar as it influences the interpretation and generalizability of our results. Figure 1a shows that, for the average CBSA, the Trepp data includes 19% of the rental units reported in the Census American Community Survey (ACS). Many of those units are in buildings financed by subsidized (Ginnie Mae) mortgages, which do not report financial performance to Trepp. Even among those that do report financial performance and valuations, we drop observations for data quality reasons described above. Our estimation sample (where we feel comfortable computing IRRs) includes 2.2% of the ACS rental unit count in the average CBSA (Figure 1b). This dataset provides better coverage of the overall rental market in areas with fewer project-based subsidies and in areas with more large multifamily properties or rental communities, meaning less single-family, smaller multifamily (e.g., 2-4 unit structures), or condo-based rental housing.

Consistent with the geographic coverage, the properties in the estimation sample are concentrated in the middle of the unconditional national rent distribution. We use Census ACS data on reported rents by tenants to construct the national rent distribution in Figure 2. The Census rent distribution has a thick left tail, likely comprised of subsidized buildings and tenants (though the ACS does not ask about rent subsidies). The Census rent distribution also has a large right-tail where 2-3% of rents are top-coded. The estimation sample of CMBS properties does not have the lumpy left-hand side of the distribution because the sample primarily consists of market-rate apartment buildings. There also may be some properties with very high rents that match the long right-tail in the Census data, but those observations are smoothed because the estimation sample averages rental income across all types of floor plans and apartment sizes within the property. In sum, combining multiple data sources yields a novel sample with detailed information on property cash flows and profitability that covers a sizable share of large market rate apartment buildings in the U.S.

2.3 Constructing Returns

Our main measure of returns for landlords is the internal rate of return (IRR) on an unlevered property investment during each "holding period," defined as the period of time between adjacent capital events. In our setting, capital events include any time the landlord realizes the change in their property value. This might be when they buy or sell the property, or when they prepay or

Table 1: Description of Main Sample

Revenue Per Unit	1029.38	782.08	990.49	1292.66
	(508.11)	(242.04)	(337.45)	(670.95)
Expense Per Unit	472.83	394.00	454.09	562.71
	(219.68)	(102.61)	(133.05)	(312.27)
NOI Per Unit	560.26	391.79	538.73	736.42
	(334.60)	(182.12)	(245.93)	(418.68)
NOI Per Unit Growth	0.042	0.054	0.042	0.030
	(0.100)	(0.128)	(0.084)	(0.079)
IRR	0.15	0.18	0.15	0.13
	(0.10)	(0.15)	(0.07)	(0.07)
Value, Origination (millions)	23.53	12.40	21.52	35.79
	(26.30)	(14.92)	(21.61)	(32.74)
Value, End Period (millions)	26.99	14.72	25.13	40.20
	(28.92)	(17.91)	(24.47)	(34.96)
Years of Hold Period	4.30	4.03	4.43	4.43
	(2.66)	(2.52)	(2.71)	(2.71)
Share Renovated Before Hold Period	0.248	0.271	0.285	0.192
	(0.432)	(0.445)	(0.452)	(0.394)
Share Renovated During Hold Period	0.056	0.064	0.057	0.047
	(0.23)	(0.25)	(0.23)	(0.21)
Number of Units	240.37	214.73	243.39	260.52
	(162)	(183)	(160)	(138)
Year Built	1984	1976	1982	1993
	(18.72)	(17.16)	(16.04)	(18.74)
Employment (thousands), CBSA	1398.37	1361.41	1416.63	1414.82
	(1389.367)	(1346.144)	(1492.793)	(1325.258)
Employment Growth, CBSA	0.018	0.019	0.017	0.018
	(0.016)	(0.015)	(0.017)	(0.017)
Bartik, CBSA	-0.002	-0.002	-0.002	-0.002
	(0.002)	(0.002)	(0.002)	(0.002)
Share of New Units (%), CBSA	0.992	1.019	0.944	1.011
	(0.623)	(0.648)	(0.588)	(0.629)
Share of New Units >75pctile (%), CBSA	58.603	59.322	58.965	57.611
	(14.797)	(15.564)	(14.231)	(14.541)
Observations	4990	1617	1632	1738

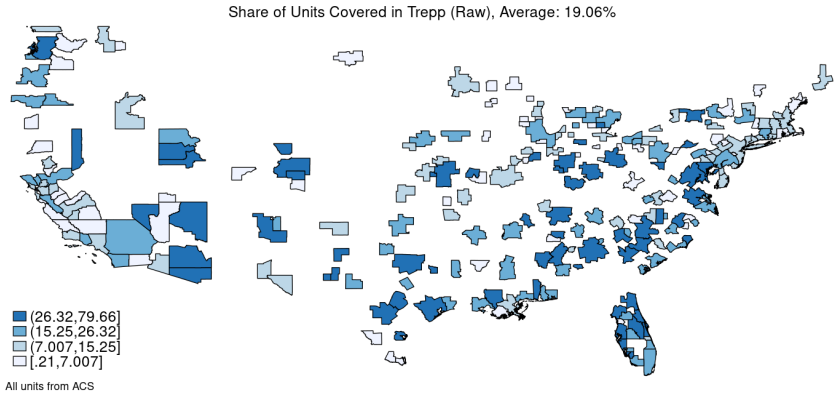
Note: Table shows summary statistics of financial variables, property characteristics, and city-level variables for the estimation sample. "All" indicates pooled observations for the entire estimation sample. "Low," "Mid," and "Top" signify lower, middle, and higher quality properties where quality is measured based on the distribution of the average revenue per unit for each property in the first year it enters the dataset within the city-by-year distribution. Observations are the number of property-hold periods observed. Data is collapsed by property-hold period. Summary statistics are averaged over those property-hold period observations (the average hold period is 4.3 years).

refinance the mortgage on the property.

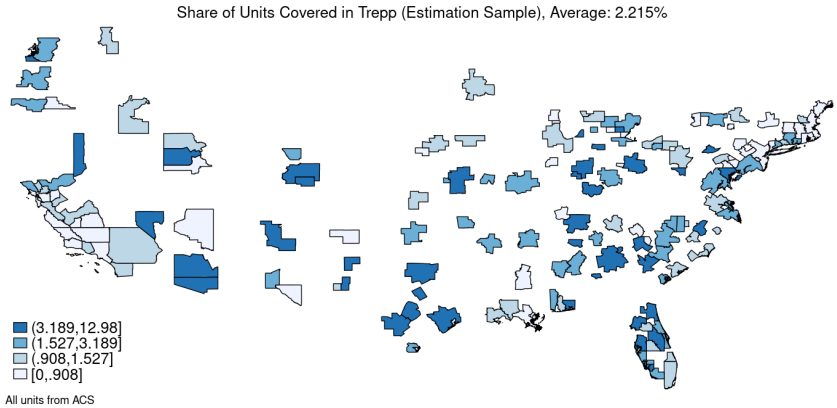
To calculate the IRR we first construct a series of annual cash flows, CF_n, over each individual property-holding period, $n \in \{0, 1, .., N\}$. The cash flow at year 0, CF_0, is the purchase price, $-PurchasePrice$, and for each year through $N-1$, the cash flow, CF_n, is the annual net operating

Figure 1: Map of CMBS (Trepp) Coverage as of 2019

(a) All Properties



(b) Estimation Sample



Note: Figure shows the share of rental housing units covered by the Trepp multifamily mortgage dataset across cities. Total rental units are measured in the Census American Community Survey. Top panel shows the Trepp dataset comprises 19% of total rental units, while the bottom panel shows the cleaned estimation sample includes just over 2% of total rental units.

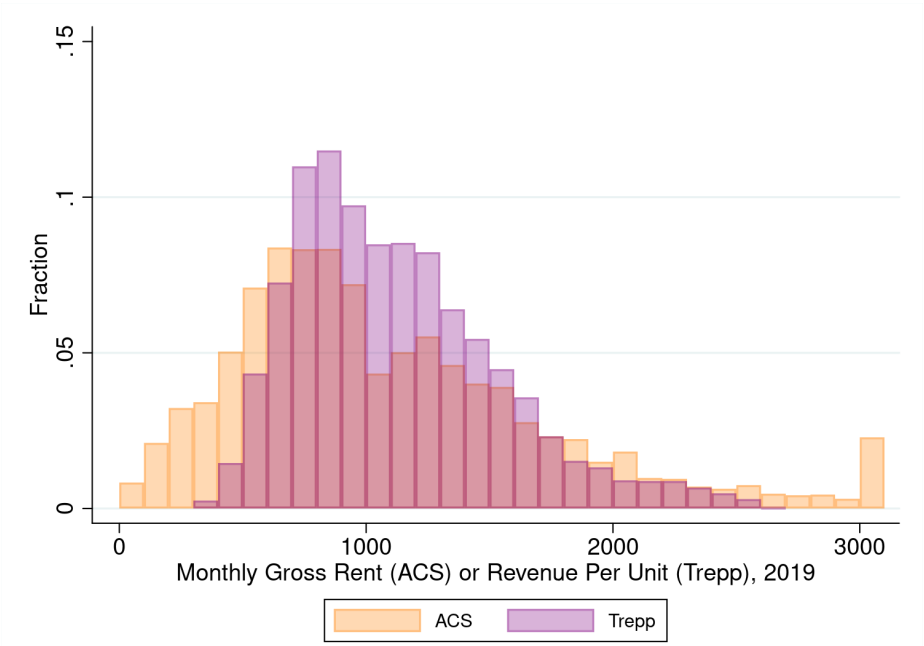
income of the property, NOI_t. In year N, the year of the capital event, the cash flow is the sum of annual NOI and the sale price, $SalePrice + NOI_N$.The internal rate of return IRR of this sequence of cash flows is the yield at which the landlord's discounted present value is equal to zero:

$$\sum_{t=0}^{N} \frac{CF_t}{(1+IRR)^t} = 0 \tag{1}$$

The financial interpretation of the IRR is the percentage rate of return earned on each dollar invested over N years with variable cash flows. This method to calculate returns is particularly useful in the context of lumpy cash flows in real estate. For example, buying an apartment building requires a large up-front investment and yields noisy annual operating incomes over a fairly long investment horizon.

Figure 3 summarizes the distribution of IRRs in our estimation sample in each year from 2005

Figure 2: Census Contract Rents vs. CMBS Monthly Rental Revenue Per Unit

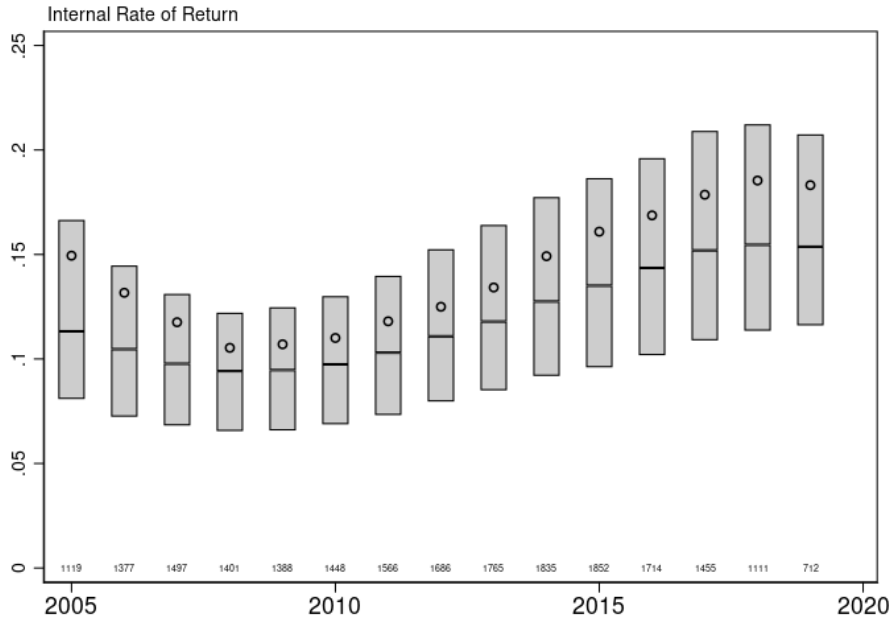


Note: Figure shows a histogram with the 2019 gross rent distribution in the Census American Community Survey in orange, and the revenue per unit per month distribution from Trepp in purple (overlapping areas of this distribution are in an purplish-orange color). Revenue per unit provides a measure of average total rental income per property, and includes revenue items like fees, parking, laundry, and vending income.

to 2019, assigning the annual IRR over the entirety of each hold period to every year represented therein. The figure presents a box plot of IRRs for properties actively operating in each year. The average IRR is shown in circular dots, with the interquartile range in gray box and the median value as the black line in the center of the box. The median IRR is generally between 7% and 20%. The calendar time dynamics suggest that IRRs fell during the housing boom in the mid-2000s. There was a subsequent increase in returns from 2010 to 2018. We compare average IRRs across quality segments in the next section.

Our approach to averaging IRRs is worth discussing in more detail. It is conceptually different from many 'total return' calculations that measure the return made on a dollar invested assuming that a property is purchased at the beginning of the year and sold at the end of the year, plus the rental income divided by the assumed purchase price. For example, Eisfeldt and Demers (2015) argue that measurements of total returns map nearly linearly onto IRRs, assuming constant duration, disposition costs, and financing conditions. We note that the assumptions about duration or horizon are potentially important, given that, descriptively, deal volume declined sharply during the Great Recession, and more generally appears to be sensitive to interest rate fluctuations. Ours is among the first papers to measure changes in IRRs over time using actual holding periods, horizons, values, and operating income. Using these actual components of IRRs allows us to answer questions about

Figure 3: Summary of IRR Averaged Over Active Holding Periods



Note: Figure shows internal rate of return for multifamily properties over time, including the average value (black dot), the median (black line), and 25th to 75th percentile of the distribution (bottom and top of gray box). After calculating the IRR, we assign that IRR to every year of the property hold period. The IRR summarized in the figure is the average return from holding and operating a multifamily property in each year. The bottom of the figure shows the number of properties contained in the underlying dataset for each year.

whether landlords' returns have changed and how much landlords gain in response to demand shocks, which we explore later in the paper.

3 Three New Facts on Segmented Rental Markets

3.1 Fact 1: Asking Rents Grew More in the Lower Quality Segment

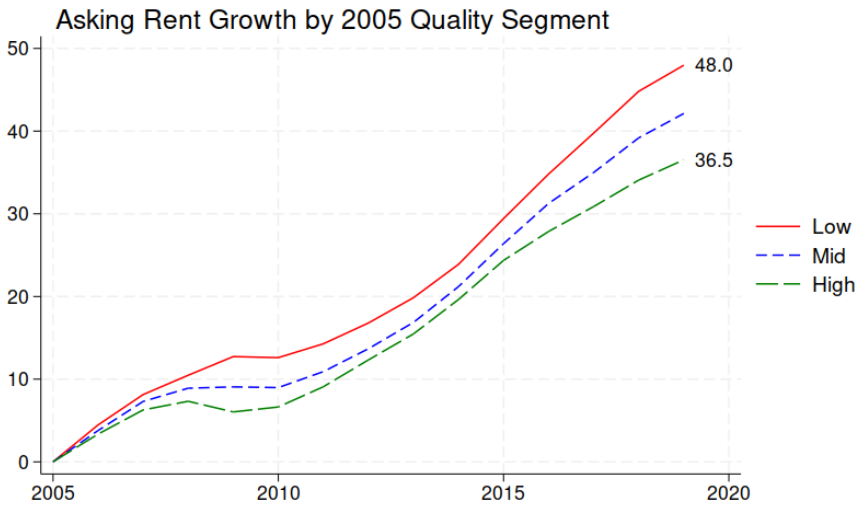
The first fact we present is that nominal asking rents grew more in lower-quality apartment buildings. To study rent growth by quality, we analyze a constant sample of market-rate properties that have rent data from 2005 through 2019 from Moody's REIS platform. The buildings we study are constructed before 2005 and continue operating through the end of our sample period.

To study quality segmentation, we define quality in the beginning of our time period, 2005, by computing each property's percentile of average asking rent in their building within their CBSA. We fix this quality segment from 2005 for the entire period, assigning based on the within-city rent tercile (low, middle, and high). We find minimal movement across quality terciles over time within this set of apartment buildings. Asking rent growth is measured within building starting from the

2005 value of asking rent, and averaged across buildings by quality segment.

Figure 4 presents the results of this exercise, where the cumulative growth starts at zero in 2005 and shows the average growth over time for each segment. Properties in the lowest quality segment experience the highest cumulative rent growth over this period with noticeably higher growth around the Great Recession and late 2010s.[6] The difference in cumulative growth is 11.5 percentage points between lower-quality properties whose rents grew by 48% over 14 years, and higher-quality properties where rents grew 36.5%. In sum, this exercise shows that proportional rent growth in the lower-quality segment was meaningfully larger over the period 2005–2019.

Figure 4: Asking Rents Grew More in Lower Quality Segments



Note: Figure shows average cumulative growth in nominal asking rents from Moody's Analytics (REIS) sample of market rent apartment buildings. "Low," "Mid," and "Top" signify lower, middle, and higher quality properties where quality is measured based on the distribution of the asking rent in the first year of the holding period within each city-by-year distribution.

3.2 Fact 2: Returns Grew More in the Lower Quality Segment

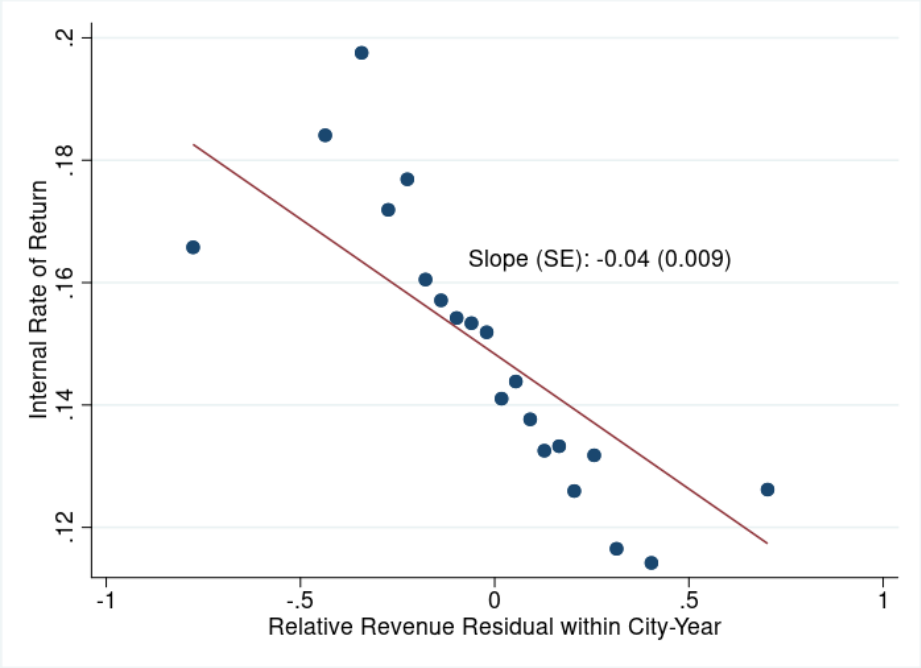
This section examines the IRRs discussed in Section 2.3, comparing IRRs across quality segments, first in the cross-section and then over time. Figure 5 shows the cross-sectional variation in IRRs over the entire time period in a binned scatterplot of IRRs on the y-axis against the relative rent on the x-axis. The relative rent measure is the residual of a property's log of revenue per unit within the city-specific distribution of revenue per unit in the first year of the property's holding period. This metric is used to assign properties into quality terciles throughout the remainder of the paper.

The main takeaway from Figure 5 is that higher-quality apartment buildings have lower average returns over our period of analysis. On average, a property with ten percent higher rent than the

[6]The figure focuses on growth rates, but because the nominal dollar value of rents begin from a much higher level in the high-quality segment, the high segment actually grew more in raw dollar terms than the lower-quality segments.

city mean has a 0.4 percentage point lower IRR.

Figure 5: Returns Are Higher in Lower Quality Segments

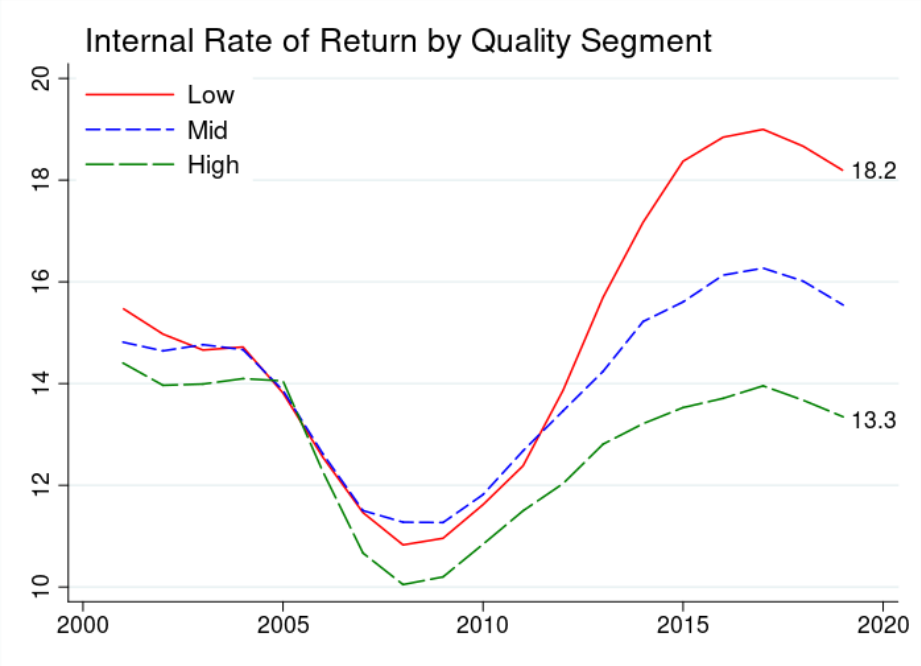


Note: Figure shows binned scatter plot of IRR on y-axis and relative rent distribution (as the proxy for measuring property quality) on x-axis. Relative rent is defined by each property's log revenue per unit residualized at the beginning of its hold period within the city-by-year revenue per unit distribution. Line and reported slope shows a linear fit of the relationship between IRR and relative rent.

To match our fact on rental returns to the fact on asking rent growth, we need to examine the change in IRRs over time. To do so, we average IRRs by property segment and year across all properties with an active hold period in each year. This calculation provides a measure of the average return on a dollar invested in a property in each quality segment during that year, averaging across all of the years that the dollar could have been originally invested or the years that the holding period could end.

Figure 6 reports the average IRR for the properties in each quality segment with active hold periods in each year. In the early 2000s, lower quality properties had slightly higher but mostly average returns. Those returns declined and converged in the mid-2000s. Returns bottomed in 2008–2009 with particularly low IRRs among higher-quality properties. From 2010 to 2017, IRRs grew in all quality segments, but differentially more in the lower-quality segment. In 2019, lower-quality buildings had returns of 18.2%, while higher-quality properties had an average IRR of 13.3%. The difference between the two peaked in 2016–2017 when lower-quality properties experienced average returns of close to 19 percent, and higher-quality properties returns remained below 14 percent. In sum, property-level returns grew disproportionately more in the low-quality segment of the distribution, especially following the Global Financial Crisis.

Figure 6: Returns Grew More in Lower Quality Segments

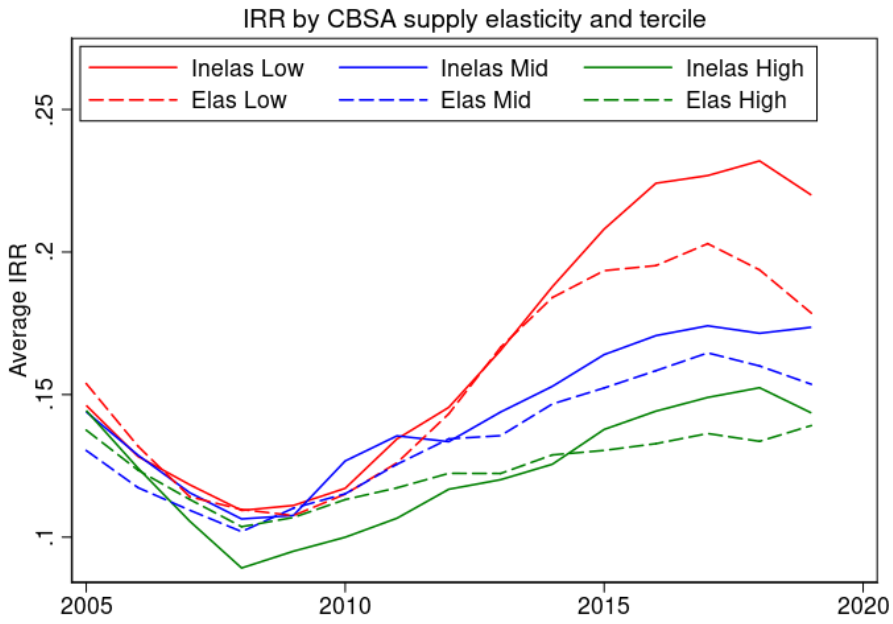


Note: Figure shows average internal rates of return for multifamily properties over time by quality segment. "Low," "Mid," and "High" signify lower, middle, and higher quality properties where quality is measured based on the distribution of the average revenue per unit for each property in the first year it enters the dataset within the city-by-year distribution. The IRR summarized in the figure is the average annual return from holding and operating a multifamily property in each year.

We present one addendum to this fact about differential expansion in returns among lower-quality properties: the expansion in IRRs among lower-quality properties occurred nationwide, but was particularly large in cities with the most inelastic housing supply. Figure 7 uses housing supply elasticity measures from Baum-Snow and Han (2024) to show that in the late 2010s, cities with below-median supply elasticities have higher IRRs than more elastic cities. This phenomenon appears to be a departure from the 2000s, when returns did not vary strongly across elastic or inelastic cities.[7] Higher-quality properties in inelastic cities appear to have had lower IRRs than high-quality properties in more elastic cities. Because the rich data needed to produce these calculations is generally not available before the early 2000s, we cannot speak to whether the divergence of returns by quality or by city-level housing supply elasticity is unique to the recovery from the Great Recession or is a more general business cycle or credit cycle phenomenon. In sum, the IRR expansion in the early 2010s affected returns in inelastic cities relatively more than elastic cities.

[7]While returns and rents were higher in supply inelastic cities in the 2010s, we found little evidence that employment growth caused higher rent or return responses in supply inelastic cities, similar to Louie et al. (2025).

Figure 7: Returns Grew More in Inelastic Cities & Lower Quality Segments



Note: Figure shows average internal rates of return for multifamily properties over time by quality segment for elastically and inelastically-supplied cities. Housing supply elasticities come from Baum-Snow and Han (2024) where inelastic cities are those with below median supply elasticities. "Low," "Mid," and "High" signify lower, middle, and higher quality properties where quality is measured based on the distribution of the average revenue per unit for each property in the first year it enters the dataset within the city-by-year distribution. The IRR summarized in the figure is the average annual return from holding and operating a multifamily property in each year.

3.3 Fact 3: Most New Rental Housing Has Entered the Top of the Quality Distribution

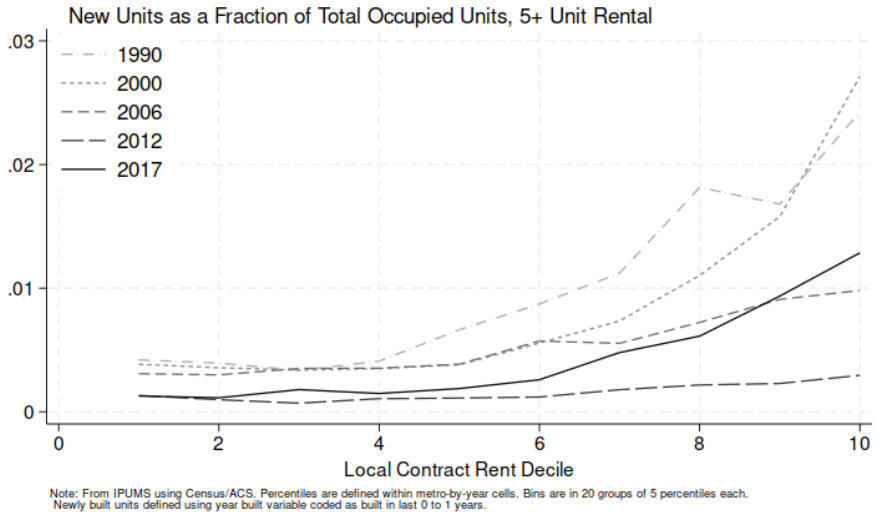
The simplest explanation for why prices or returns may not rise as much in regions with elastic housing supply is that supply expands in response to positive demand shocks. One of the core insights from our analysis of landlords' returns above is that lower-quality housing became more profitable than higher-quality housing in the 2010s. We investigate that same logic of supply growth, but applied to segmentation within a city.

In this section, we use Decennial Census and ACS data on the universe of rental housing to establish that, since 1990, most new rental housing entered the top of the quality distribution. As in previous sections, we assign percentiles to each occupied rental unit based on the rent distribution in each city and year. We focus exclusively on units in "large buildings," i.e., those with 5 or more units.

Figure 8 shows the share of existing rental housing in each decile of the rent distribution that was newly built in each Census year. We define "newly built" as either completed in the same year as the survey or one year prior. Going back to the 1990s, new rental housing is heavily skewed towards

the top of the rent distribution. This relationship is nearly monotonic, and appears to be consistent regardless of whether the U.S. was building a substantial amount of housing overall (as in 1990 and 2000), or very little (as in 2012 or 2017).

Figure 8: Most New Housing Enters at the Top of the Quality Distribution



Note: Figure shows the share of multifamily rental housing units classified as newly built across the local rent distribution in various Census years (using decennial Census data for 1990 and 2000, and ACS data for 2006, 2012, and 2017). The x-axis represents the city rent decile, and the y-axis shows the share of rental units in each decile that were built within the same year or the year prior.

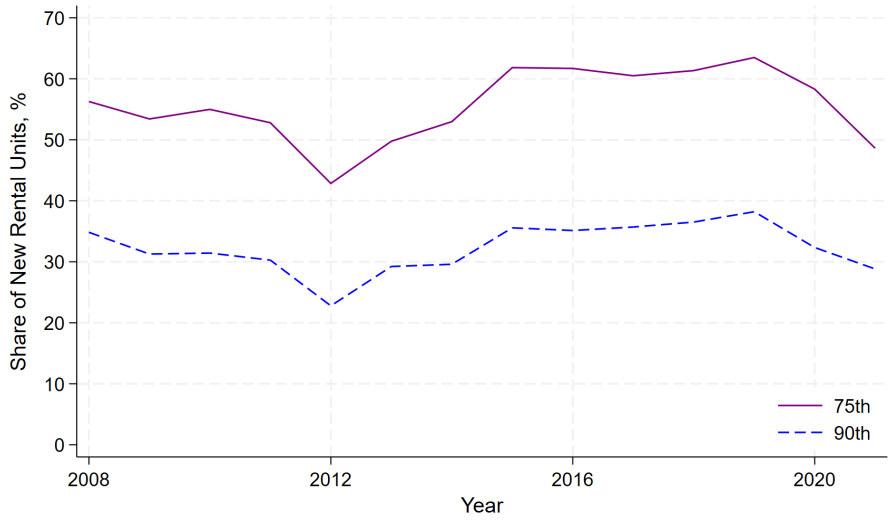
To match this observation to our prior facts, Figure 9 shows the share of newly-built housing in each year whose starting rents fell in the top quartile of the within-city rent distribution for that year. From 2008–2011, 55% of new housing entered the top quartile of the local rent distribution. This percentage dipped in 2012, then rose to more than 60% from 2015–2019. Approximately 30-40% of new housing entered the top ten percent of the local rent distribution over this period with dynamics matching the measure using the top quartile. Taken together, Figures 8 and 9 show that over our time period, most new rental units were more expensive than the average existing unit.

4 Theoretical Framework Linking These Facts

This section discusses a theory consistent with the facts above. We show conditions under which housing supply is more elastic to aggregate demand shocks at the high-end than the low-end of the market and, in turn, rents and prices are more responsive to demand shocks at the low-end of the market. The next section uses CBSA-level Bartik shocks to provide causal evidence of this result.

The appendix presents a model of a segmented housing market. There are two types of housing, high-quality and low-quality ($k \in \{H, L\}$), provided in horizontally-differentiated buildings, that in equilibrium each provide as many units of housing as allowed by zoning constraints ($\bar{h}$). Housing

Figure 9: The Share of New Housing Entering the Top of the Local Distribution Has Remained High Across the Cycle



Note: Figure shows the share of new rental housing units entering the local rent distribution above the 75th or 90th percentile using one-year ACS data from 2008 to 2021.

supply adjusts over time through filtering (a fraction δ^H of high-quality housing becomes low-quality, and a fraction δ^L of low-quality housing becomes unusable) and new construction. Housing construction requires one unit of land per building and K^k units of capital per unit of housing, with $K^H > K^L$. Developers enter the market each period and draw a productivity parameter ϕ_j that determines the amount of capital they will have to purchase in order to cover the capital requirements of construction. There is free entry, and expected profits dissipate into land prices, which adjust to clear land markets. There is a one-period construction lag. Developers have rational expectations and choose to build in t if the expected price of the housing—the expected present value of its future cash yield—exceeds development costs, given land prices (P_t^l) and their productivity draw. We normalize the cost of capital to one.

A population M_t of households have unit housing demand and nested logit preferences over buildings in the two quality segments and an outside option. Indirect utility for household i from a unit of housing in building j of quality k is

$$V_{ij} = v_j + \varepsilon_{ij}, \quad v_j = \alpha^k (y_i - p^k),$$

where α^k is a quality k amenity shifter, p^k is the rental price of quality-k housing units, and y_i is household income distributed $F_t(y)$. ε_{ij} are iid nested-logit shocks with dissimilarity parameter $\tau_k \in (0, 1]$ in nest k. In any given period the supply of buildings in each segment k is fixed at S_t^k

and rents adjust to clear each market, satisfying

$$M_t \int \lambda^k(y)(p^H, p^L) dF_t(y) = S_t^k \bar{h}$$

where $\lambda^k(y)$ is the share of households earning income y that select housing of quality k. Buildings are traded on an asset market at a price equal to an increasing function of rents, $P_t^k = \Phi_t^k(p_t^k)$.

In equilibrium, there is more development of high-quality housing if the expected price premium for high-quality housing (net of land requirements), exceeds the marginal capital requirement of high-quality development:

$$\frac{\mathbb{E}\left[P_{t+1}^H\right] - P_t^l/\bar{h}}{\mathbb{E}\left[P_{t+1}^L\right] - P_t^l/\bar{h}} > \frac{K^H}{K^L}.$$

Note, the entry gap will increase in land costs P_t^l and decrease in the stringency of zoning $\bar{h}$. This result mimics the intuition in Ahlfeldt and Barr (2022), where developers pick building height, rather than quality, and the optimal skyscraper height increases with the amenity value of a location.

The model further predicts that rents and prices are more elastic to demand shocks at the low-end of the market. High-end supply responds more to an unexpected aggregate demand shock (ΔM), when the elasticity of the premium on high-quality housing exceeds the high-to-low quality ratio of capital intensities. In the subsequent period, the outsized elasticity of supply in the high-quality segment will offset the rental and price growth of high-quality housing. In a more realistic situation where demand shocks are autocorrelated, the excess elasticity of supply at the high-end will mitigate continued growth of rents and prices. In the event of a negative demand shock, meanwhile, high-end housing supply will contract by more than low-end housing supply. If all development ceases, both types of housing will depreciate, but low-end depreciation will be offset by the filtering of high-end housing into the low-end sector. The relative contraction of high-end housing supply will mitigate the collapse in high-end rents and prices.

The model and its implications are described in more detail in Appendix Section B. The subsequent section tests these predictions.

5 Rental Income and IRR Response to Demand Shocks

In this section, we investigate whether rental incomes and returns respond differently to local labor demand shocks by quality segment. To do so, we regress returns and NOI growth on city-level employment growth, allowing the effect of employment growth to vary by segment. We find that lower-quality apartments' experience larger increases in NOI and returns compared to medium- and high-quality rentals in response to similar employment growth changes. These results hold using a proxy for labor demand growth based on an industry-mix shift-share variable (Bartik 1991).

We then attempt to understand why quality segments respond to demand shocks differently by studying how rental incomes and returns respond to new supply. To the extent that (a) labor demand shocks result in a subsequent increase in new supply and (b) new supply is tilted towards higher-quality apartments, it could be the case that high-quality returns respond less to demand shocks because high-quality new supply is more elastic (and low-quality new supply is relatively

inelastic). We test this hypothesis by regressing NOI growth and returns on annual new housing supply, allowing effects to vary by quality segment. We find that new supply is strongly related to lower returns for high-quality housing, but less so for existing lower-quality buildings. In addition, directly controlling for new supply explains 40 to 70% of the differential response of NOI and returns to labor demand shocks.

5.1 Empirical Strategy

To measure the response to local demand shocks, we estimate the following equation:

$$y_{it} = \alpha_t + \alpha_c + \beta * \Delta E_{ct} + \epsilon_{it} \tag{2}$$

where y_{it} is an outcome variable for property i located in city c in year t. The coefficient β gives the average effect of a change in city employment E_{ct} on the outcome variable y_{it}. The outcome variable in the main regressions will be either year-over-year growth in NOI (ΔNOI_{it}) or realized returns over the holding period (IRR_{it}).[8] We include city and year fixed effects to isolate the variation coming from labor demand shocks, accounting for the national change within years, and variation specific to particular cities. Standard errors are clustered by city (CBSA).

We also estimate specifications that proxy for growth in city-level labor demand with a Bartik shift-share variable, ΔB_{ct}, to identify causal effects of labor demand shocks. Our main concern for identifying labor demand shocks is omitted variable bias: an unobservable shock might drive both employment growth and increase rents or property values via an independent channel (e.g., an urban renewal project that is paired with new zoning regulations may boost employment, shifting demand, but also induce new supply). The Bartik instrument isolates the variation in employment that is predicted by national shifts in labor demand and pre-existing local industry employment shares, making it orthogonal to such idiosyncratic policy shifts by design.

Our main regression specification modifies the above equation by interacting a quality segment indicator variable with the labor demand shock variable to recover a segment-specific responsiveness of rental income or returns to local demand:

$$y_{it} = \alpha_t + \alpha_c + \beta^q * \Delta E_{ct} * \mathbf{1}[i \in q] + \epsilon_{it} \tag{3}$$

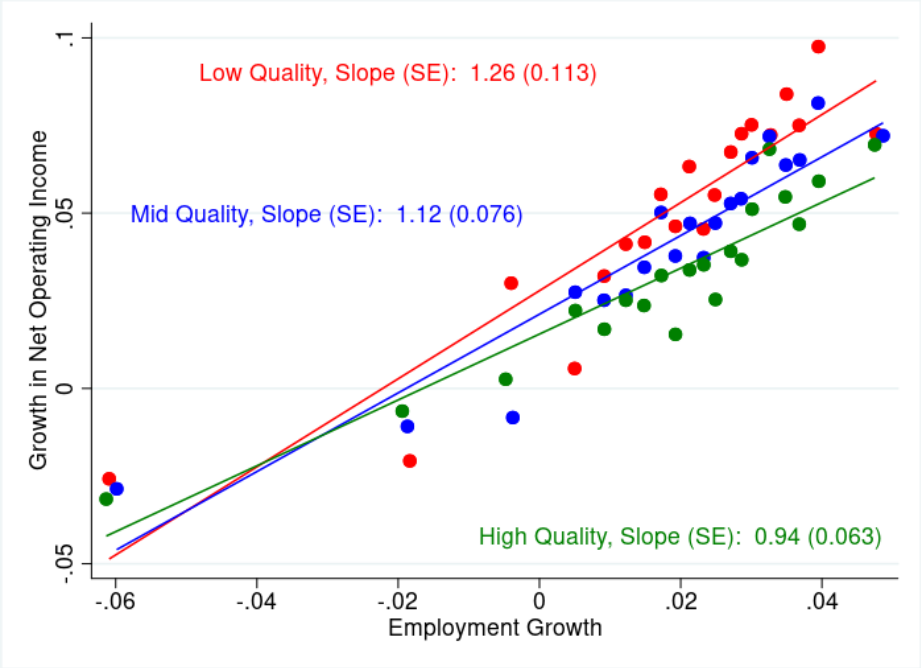
This modification can be represented by an indicator function $\mathbf{1}[i \in q]$, which signifies whether property i is in the bottom, middle, or top tercile of their CBSA's revenue per unit distribution in the year entering the hold period. The quality segment-specific coefficient, β^q, tells us the low-, middle-, and high-quality segment's average change in rental income or returns from a same-size shift in city labor demand.

[8] Because the dependent variable is either annual growth in NOI or the average annual returns (IRR) over the hold period, we do not include property fixed effects in these regressions, as residualizing by property fixed effects would absorb the variation of interest (annual rent growth and landlords returns) and instead measure the change in growth rates. Any property-specific variation should be absorbed by differencing net rental incomes, and by calculating a property-specific realized annual return over the hold period (essentially giving us the annual change in the realized return over the holding period).

5.2 Responsiveness of Property Cash Flows and Returns to Local Demand Shocks

To visualize the variation we use in our main specification, we present binned scatterplots in Figures 10 and 11. The bins group the data in quantiles separately by quality segment across the x-axis, which is average annual CBSA-level employment growth, and show the variation in the dependent variable on the y-axis. Figure 10 shows the relationship between labor demand growth and annual changes in net operating income, by low (red), middle (blue), and high (green) quality segments. Figure 11 plots the IRR on the y-axis. The linear slope of an ordinary least squares regression line for each segment is reported in the figure.

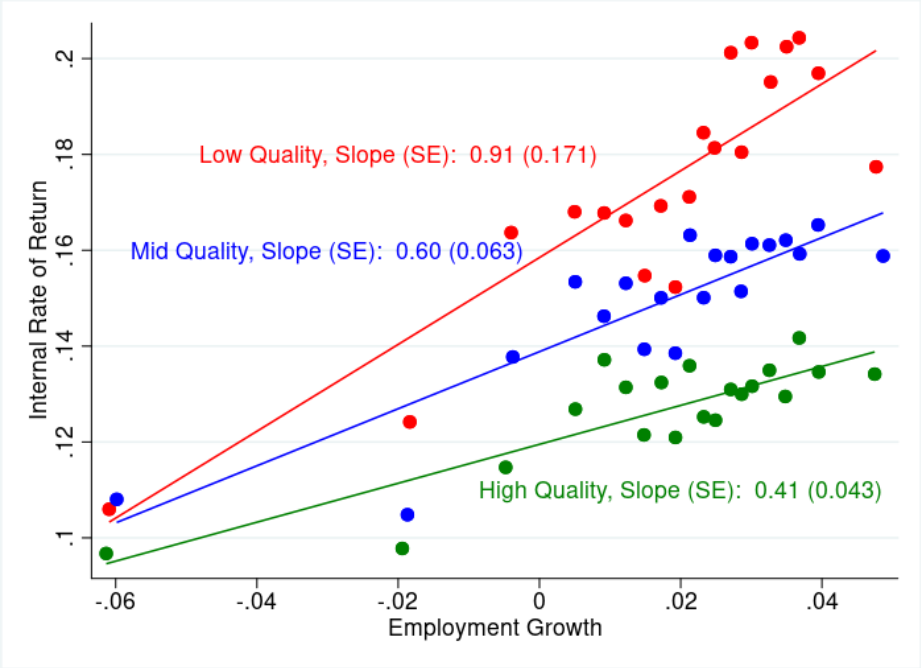
Figure 10: Net Rental Income vs. CBSA-Level Employment Growth



Note: Figure shows binned scatter plot of net rental income (revenue net of expenses) on the y-axis and annual employment growth on the x-axis. Change in NOI and employment measured in year-over-year change in logs. Binned scatters for low, middle, and higher quality properties are plotted separately in red, blue, and green, respectively. The linear slope coefficient and standard error is reported for each group.

The primary takeaway from Figures 10 and 11 is that both rental income (net of actual expenses) and returns are responsive to local demand growth. The next section explores the variation in these responses by quality segment. For now we note that, in both graphs, the low-quality segment has a steeper slope than the middle- and high-quality segments, indicating that lower-quality apartment buildings' rental incomes and returns are more responsive to labor demand shocks.

Figure 11: Internal Rate of Return (IRR) vs. CBSA-Level Employment Growth



Note: Figure shows binned scatter plot of internal rates of return on the y-axis and annual employment growth on the x-axis. Binned scatters for low, middle, and higher quality properties are plotted separately in red, blue, and green, respectively. The linear slope coefficient and standard error is reported for each group.

We next quantify the relationships between changes in employment growth and NOI growth or IRR. Table 2 shows that one percentage point higher employment growth rate (without distinguishing property quality segments) is associated with 1.5 percent higher growth in net operating income. When accounting for endogeneity using Bartik shift-share variables in a two-stage least squares instrumental variable regression, we estimate that one percentage point higher employment growth causes 1.7 percent higher NOI growth.

Table 3 presents the same regressions with IRR as the dependent variable. The results indicate that over the average holding period, returns are 0.5 ppt higher in response to one percent higher employment growth. Similar to the NOI regression on Bartik shocks, the IRR regression has a larger coefficient on the shift-share for employment. The instrumental variables specification yields a slightly smaller effect relative to the version without an IV, such that one percent higher employment growth causes a 0.4 ppt larger internal rate of return.

Table 2: Average Response of NOI to Labor Demand

	NOI Growth (1)	NOI Growth (2)	Emp Growth (3)	NOI Growth (4)
Employment Growth	1.471***			1.692***
	(0.166)			(0.357)
Bartik (3-Digit, Log)		1.840***	1.091***	
		(0.548)	(0.264)	
Spec	OLS	RF	FS	IV
N	27819	27819	27819	27819
F Stats	11.98	3.51	17.11	
CBSA	163	163	163	163
Adjusted R-sq	0.05	0.04	0.85	0.01
Year FE	Yes	Yes	Yes	Yes
CBSA FE	Yes	Yes	Yes	Yes

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: *** p<0.01, ** p<0.05, * p<0.10. Standard errors are clustered by CBSA. Each specification includes year and CBSA fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Change in NOI and employment measured in year-over-year change in logs. Columns show ordinary least squares, reduced form, first stage, and instrumental variables specification. Bartik shift-share variable is constructed using 3-digit NAICS industry codes where shifts are annual changes in log employment by industry.

Table 3: Average Response of IRR to Labor Demand

	IRR	IRR	Emp Growth	IRR
	(1)	(2)	(3)	(4)
Employment Growth	0.534***			0.412*
	(0.162)			(0.234)
Bartik (3-Digit, Log)		0.448*	1.091***	
		(0.256)	(0.264)	
Spec	OLS	RF	FS	IV
N	27819	27819	27819	27819
F Stats	10.36	8.43	17.11	
CBSA	163	163	163	163
Adjusted R-sq	0.12	0.12	0.85	0.03
Year FE	Yes	Yes	Yes	Yes
CBSA FE	Yes	Yes	Yes	Yes

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: *** p<0.01, ** p<0.05, * p<0.10. Standard errors are clustered by CBSA. Each specification includes year and CBSA fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Columns show ordinary least squares, reduced form, first stage, and instrumental variables specification. Bartik shift-share variable is constructed using 3-digit NAICS industry codes where shifts are annual changes in log employment by industry.

These empirical specifications capture the effect of contemporaneous demand shocks on NOI and IRR. We may be concerned that these estimates are driven by differential trends. To test this hypothesis, we can regress lags and leads of the demand shock variables on rents and returns to measure the dynamics of how a marginal change in demand passes through to landlords' rents and returns.

The regression of NOI growth or annual returns on leads and lags of the employment growth variable is specified as:

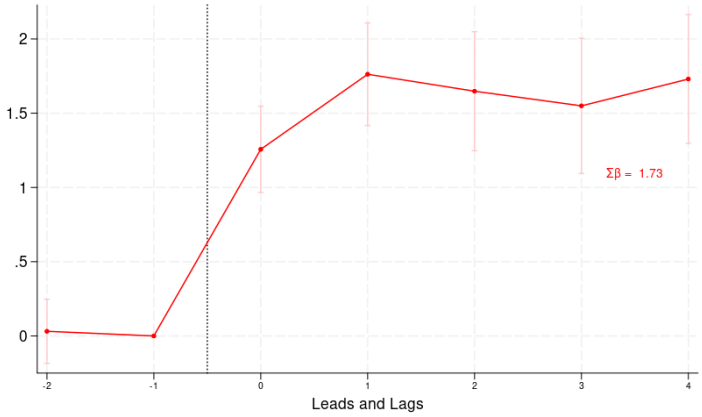
$$y_{it} = \alpha_t + \alpha_c + \sum_{s=-2}^{4} \beta_s * \Delta E_{c;t-s} + \epsilon_{it} \tag{4}$$

Figure 12 shows the dynamic effects of employment growth on NOI and IRR. Coefficients from these regressions are displayed in the graphs by cumulatively summing the regression coefficients and are reported in the graph relative to the leading year ($s = -1$). The conditional correlations of both NOI and IRR with leads of employment growth (representing growth 1 and 2 years in the future, plotted at $s = -2$ and $s = -1$ respectively) are indistinguishable from zero — that is, net income and returns do not predict employment growth. By contrast, net rents and returns are both increasing in contemporaneous marginal employment growth. The magnitude of this contemporaneous conditional correlation (depicted at $s = 0$) and the cumulative coefficient one year after the shock (at $s = 1$) are similar to the average effects estimated in Tables 2 and 3 above. The estimates for IRR (Figure 12b) are less precise than the estimates for NOI growth (Figure 12a) since we include data on NOI for all years in each hold period, but only include IRR in the final year of the hold period, when it is realized.[9] In sum, we find strong and persistent responsiveness to multifamily investors' NOI and IRR from local employment growth.

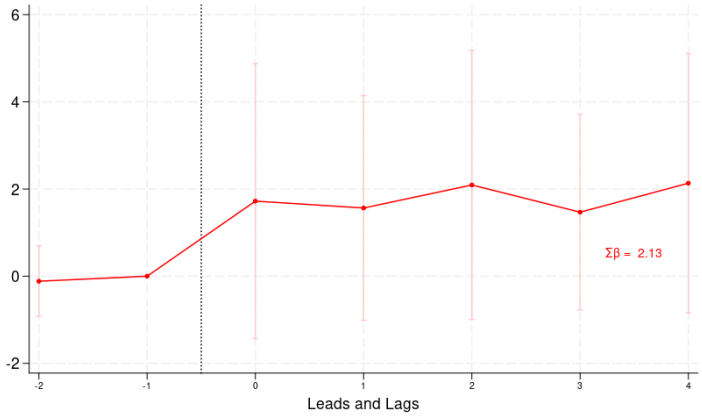
[9]Our other analyses smooth the IRR across all years of the hold period. This smoothed IRR would mechanically predict any impact of the employment shocks that drive the NOI in the later years of the hold period as well as the sales price.

Figure 12: Dynamic Effects of Employment Growth on NOI Growth and IRR

(a) Dependent Variable: NOI Growth



(b) Dependent Variable: IRR



Note: *** p<0.01, ** p<0.05, * p<0.10. Standard errors are clustered by CBSA. Each specification includes year and CBSA fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Change in NOI and employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Panel (a) includes data from all years in each hold period. Panel (b) only includes data from the final year of the hold period. Plotted coefficients and 95% confidence intervals from specification with two leads and four lags of employment growth (coefficients are cumulatively summed and reported relative to the first lead). Sum of lead and lag coefficients is reported in the top right corner of each plot.

5.3 Differential Responsiveness to Local Demand Shocks by Quality Segment

In this subsection, we investigate whether the growth in rental income and returns varies across the quality distribution. We have shown a descriptive version of this test in Figures 10 and 11. Table 4 formalizes this analysis following Equation 3, which interacts an indicator variable for each quality

segment (Low, Mid, and High quality) with the employment growth varying at the city-by-year level.

We observe in Column (1) of Table 4 that city-level employment growth increases net rental income by 1.8 percent for properties in the lower-quality segment, while higher-quality properties increase by 1.3 percent. Below the coefficients, Table 4 reports the difference between the low- and high-quality coefficients, as well as a p-value from a t-test that the coefficients are equal. This specification implies that the same employment growth rate has a 30 percent smaller effect on rental income growth for higher-quality properties, and this difference is highly statistically significant. Using the Bartik shock version of the specification leads to an estimated differential responsiveness of 23 percent between low- and high-quality properties, remaining statistically significant (Appendix Table D.3).

Studying the differential response of returns to demand shocks in Column (1) of Table 5 shows that the average relationship in Table 3 masks substantial heterogeneity by quality segment. One percent higher employment growth increases average IRR by 1.2 percentage points for lower-quality properties, but by only 0.1 for higher-quality properties. The high-quality coefficient is 90 percent smaller than the low-quality coefficient, a difference that is highly statistically significant. Appendix Table D.4 shows analogous regressions using Bartik shocks as proxies for local demand. The effects are quite similar: Lower-quality property IRRs are about 90 percent more responsive to demand shocks than higher-quality property returns.

In Appendix Figures D.1 andD.2, we show the dynamic versions of these regressions, allowing the leads and lags of labor demand shocks to vary by quality segment. The effects of employment growth are comparable across quality segments in the first year of a demand shock, but differences grow over the four subsequent years. Consistent with the idea of a quality segment-specific supply response, which we explore in the next section, the differential effects on high quality properties' rental income fade three to four years after a shock.

Table 4: NOI-employment correlation by quality segment, adding controls

	NOI					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Employment Growth	1.790***	1.575***		1.603***		1.629***
	(0.187)	(0.193)		(0.192)		(0.191)
Mid Qual X Employment Growth	1.516***	1.501***		1.514***		1.519***
	(0.161)	(0.162)		(0.162)		(0.165)
Top Qual X Employment Growth	1.250***	1.345***		1.312***		1.365***
	(0.171)	(0.160)		(0.160)		(0.164)
Low Qual X New Rental Share		0.021	−0.008	0.186		
		(0.289)	(0.237)	(0.251)		
Mid Qual X New Rental Share		−0.778***	−1.017***	−0.685***		
		(0.238)	(0.243)	(0.238)		
Top Qual X New Rental Share		−0.834***	−1.625***	−1.077***		
		(0.290)	(0.265)	(0.273)		
Low Qual X New Units >75pctile		0.006			0.021***	0.012**
		(0.006)			(0.007)	(0.005)
Mid Qual X New Units >75pctile		0.003			0.005	−0.000
		(0.006)			(0.006)	(0.005)
Top Qual X New Units >75pctile		−0.007			−0.010	−0.012**
		(0.005)			(0.006)	(0.005)
Low-High Quality Difference	0.540	0.230		0.291		0.264
P-value Estimates are Different	0.000	0.017		0.001		0.003
Pct Change from Baseline		57.4%		46.1%		51.1%
N	27819	27819	27819	27819	27819	27819
CBSA	163	163	163	163	163	163
Adjusted R-sq	0.05	0.06	0.04	0.06	0.04	0.05

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is annual log change in NOI. Change in NOI and employment measured in year-over-year change in logs. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1.

Table 5: IRR-employment correlation by quality segment, adding controls

	IRR					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Employment Growth	1.173***	0.762***		0.881***		0.804***
	(0.218)	(0.189)		(0.201)		(0.188)
Mid Qual X Employment Growth	0.526***	0.527***		0.510***		0.545***
	(0.155)	(0.151)		(0.148)		(0.154)
Top Qual X Employment Growth	0.103	0.351**		0.269*		0.356**
	(0.172)	(0.157)		(0.155)		(0.159)
Low Qual X New Rental Share		0.173	1.168***	0.877***		
		(0.311)	(0.305)	(0.278)		
Mid Qual X New Rental Share		−0.267	−0.563***	−0.370**		
		(0.166)	(0.164)	(0.164)		
Top Qual X New Rental Share		−0.889***	−1.933***	−1.474***		
		(0.277)	(0.310)	(0.281)		
Low Qual X New Units >75pctile		0.022***			0.036***	0.028***
		(0.007)			(0.008)	(0.006)
Mid Qual X New Units >75pctile		−0.003			−0.001	−0.002
		(0.004)			(0.004)	(0.004)
Top Qual X New Units >75pctile		−0.018***			−0.030***	−0.026***
		(0.005)			(0.005)	(0.004)
Low-High Quality Difference	1.070	0.411		0.612		0.448
P-value Estimates are Different	0.000	0.005		0.000		0.001
Pct Change from Baseline		61.6%		42.8%		58.1%
N	27819	27819	27819	27819	27819	27819
CBSA	163	163	163	163	163	163
Adjusted R-sq	0.14	0.15	0.14	0.15	0.14	0.15

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1.

5.4 Does Housing Supply Help Explain the Differential Responsiveness to Local Demand Shocks?

In this subsection, we unpack the differential response of lower and higher-quality rental properties to local demand shocks. The key issue to address is whether low- and high-quality properties respond similarly to new housing entering the market in their city — whether the differences in overall new supply can explain (statistically) why demand shocks affect lower-quality properties' rental income and returns disproportionately more.

Column (3) of Table 4 interacts the quality indicator variables with a city-by-year level variable, 'New Rental Share,' which captures how much new rental housing is entering the market. The new rental share variable is constructed as the number of occupied rental housing units that were built over the current or prior year divided by the total number of occupied rental housing units in the city. In the average city, one percent of rental units are newly constructed (see Table 1).

New occupied rental supply has strong negative effects on rental income, although this effect is concentrated among higher-quality properties. One percent more new rental supply is related to 1.6 percent lower NOI growth for high-quality properties, yet it has little effect on lower-quality properties. In Column (5) of Table 4, we observe a positive relationship between the share of new supply entering the top of the quality distribution and the NOI growth of existing lower-quality buildings.

From this analysis, if we assume that directly controlling for supply captures the channel by which demand shocks could be differentially affected through new housing, then the gap between low versus high quality employment coefficients narrows considerably. We estimate that the gap falls by 57%. This exercise is repeated in Appendix Table D.3 using the Bartik shock. Similar to Table 4, low-quality properties are more sensitive to demand shocks, and over half of this differential responsiveness can be explained by new rental housing supply.

In Tables 5, we study the effect of local demand on landlords' returns. We regress the property-by-holding period IRR on annual changes in labor demand. The coefficients can be interpreted as the effect of a percentage point increase in average employment growth over the hold period on landlords' internal rate of return. The differential responses are pronounced whether we use employment growth directly or Bartik shocks (see Appendix Table D.4). On average, one percent higher employment growth is correlated with 1.2 percentage points higher IRR for lower-quality properties, while it is associated with only 0.1 pp higher returns on higher-quality properties (90% smaller than the lower-quality segment coefficient). The effect on the higher-quality segment is not statistically significant, while the lower-quality segment effect is highly significant.

Thus, we conclude that new rental supply has large negative effects on returns for higher-quality properties and, if anything, positive effects on returns for lower-quality properties. Controlling for the new rental supply in Column (2) of Table 5 explains 62% of the differential response to demand shocks, reducing the size of the lower quality segment coefficient and increasing the higher quality segment coefficient.

These effects are consistent with new supply moderating the effects of labor demand shocks. New supply entering the local rental market competes down current rent growth, mostly on high-

quality properties. New supply thereby reduces returns on high-quality properties. In contrast, given the positive coefficients on new rental supply on low-quality IRR, new supply seems to relatively *increase* expected future rent growth on low-quality properties. This potentially surprising finding could be due to a loss of option value on irreversible investment—realized entry of a new high-quality property on a redevelopable lot means that a new low-quality property cannot enter there in the future (Glancy et al. 2024).

6 Robustness

Our main facts and supply-side explanation are robust to various alternative specifications, controls, and interpretations. In Appendix C, we show that our results are not driven by a quality bias in either renovations or delinquencies; outsized responsiveness of low-end cap rates to falling interest rates; or segment-specific local demand shocks.

Approximately 30% of the sample reports a renovation (in Trepp's CMBS servicing data) in the 2 years before and the years during the hold period, and this subsample reports higher returns (17.3% relative to 14.6% for the hold periods without property upgrades). If these renovations are more likely to occur for low-end properties in high-demand growth areas, under-reporting of the capital expenses associated with these upgrades would push low-end returns upwards and drive a correlation between these returns and local demand shocks. Appendix C.1 shows, however, that our results persist in both subsamples when we segment the hold periods that include renovations and those that do not.

One also might be concerned that differential changes in returns between low- and high-quality properties are due to differences in the share of loans in distress. Property owners with delinquent loans might respond differently to local demand shocks (e.g., liquidating or selling in a weak market), and be more likely to report their financial information to CMBS servicers. Appendix C.2 shows that excluding distressed properties does little to change our overall estimates.

Our main specifications estimate the elasticity of segment-specific property performance to aggregate local demand shocks. If the growth in employment we measure and instrument with Bartik shocks is driven by low-wage workers who cannot afford high-end apartments, it is natural that they would tend to drive larger performance gains in low-end properties. Instead, Appendix C.3 shows that low-quality property NOI increases *more* in CBSAs with relatively more high-income population growth. High-end NOI growth actually increases slightly less in markets with disproportionate high-income population growth. Consistent with our main explanation—quality differentials in demand shocks are moderated by new supply—the effects of growth in high-income households in the city has little or no differential effect after controlling for new high-end housing supply.

The model in Section 4 implies that the relative inelasticity of low-end housing supply to local demand growth compensates low-end property investors for differential risks associated with low-end investments. An alternative explanation is that these investors benefited from outsized cap rate compression due to falling interest rates. That is, they purchased properties at higher cap rates that had further to fall as interest rates dropped. We instead find that properties with higher entry cap rates see less pass-through of interest rate declines to returns. Our main results are also robust to

directly including controls for this alternative mechanism (see Appendix C.4).

7 Discussion

In this paper, we have shown that lower end properties saw higher growth in net rental income and realized more appreciation in the growing multifamily market between the Great Recession and the COVID pandemic. The combination of greater asset appreciation and rental income led to differentially higher returns on investment for low-end property landlords. We hypothesize that these differentials are due to relatively inelastic supply of low-end units. Demand growth was met with new high-end supply, mitigating rent growth and appreciation at the high end of the market, but not the low end.

We use variation in these dynamics across markets to develop our case. We show that the lower end of the market saw more operating income growth and higher returns than high-end properties in markets with higher local employment growth, even when instrumented with Bartik shocks. This differential effect on net rental income falls by 60% after controlling for local new supply and the share of new supply at the high-end. The differential effect on landlord returns closes by 60 to 70% accounting for new supply.

A natural next step in this research is to better understand the reason for limited low-end supply and barriers to affordable development. One explanation relates to the cost of capital. In efficient capital markets, the outsized returns generated by supply frictions in the asset market must compensate investors for taking on more risk or making a contrarian bet. A recent literature has considered the role of rent non-payment risk in adjusting landlord behavior (Abramson and Van Nieuwerburgh 2024; Bezy et al. 2024). Other sources of risk to low-end landlords include policy risk (such as the potential introductions of rent controls or eviction moratoria) or long-run price volatility, which itself may be exacerbated by supply inelasticities. Additionally, barriers to building affordable multifamily units due to regulatory or permitting hurdles or higher fixed costs of construction are worthy areas of further investigation (Baum-Snow and Han 2024; Gorback and Keys 2020; Saiz 2010). Which of these factors is most important in explaining limits to new low-end supply will help direct policy to address the imbalances that we identify in this paper.

References

Abramson, B., & Van Nieuwerburgh, S. (2024). *Rent guarantee insurance* (tech. rep.). National Bureau of Economic Research.

Adams, B., & Verbrugge, R. (2021). Location, location, structure type: Rent divergence within neighborhoods.

Ahlfeldt, G. M., & Barr, J. (2022). The economics of skyscrapers: A synthesis. *Journal of Urban Economics*, *129*, 103419.

Amaral, F., Dohmen, M., Kohl, S., & Schularick, M. (2024). Interest rates and the spatial polarization of housing markets. *American Economic Review: Insights*, *6*(1), 89–104.

Ambrose, B. W., An, X., & Lopez, L. A. (2022). Credit market spillover effects to the real economy: Evidence from rental housing evictions.

Anagol, S., Ferreira, F. V., & Rexer, J. (2025). Estimating the economic value of zoning reform. *NBER Working Paper*, (w29440).

Anenberg, E., & Kung, E. (2020). Can more housing supply solve the affordability crisis? evidence from a neighborhood choice model. *Regional Science and Urban Economics*, *80*, 103363.

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

Bartik, T. J. (1991). Who benefits from state and local economic development policies?

Baum-Snow, N., & Han, L. (2024). The microgeography of housing supply. *Journal of Political Economy*, *132*(6), 000–000.

Bezy, T., Levy, A., & McQuade, T. (2024). *Insuring landlords* (tech. rep.).

Büchler, S., Ehrlich, M. v., & Schöni, O. (2021). The amplifying effect of capitalization rates on housing supply. *Journal of urban economics*, *126*, 103370.

Büchler, S., Van de Minne, A., & Schöni, O. (2020). Intrinsic real option value: Empirical evidence from commercial real estate investors. *MIT Center for Real Estate Research Paper*, (10).

Damen, S., Korevaar, M., & Van Nieuwerburgh, S. (2025). An alpha in affordable housing? *Working Paper*.

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

Eisfeldt, A., & Demers, A. (2015). *Rental yields and hpa: The returns to single family rentals*. National Bureau of Economic Research.

Eriksen, M. D., & Orlando, A. W. (2022). Returns to scale in residential construction: The marginal impact of building height. *Real Estate Economics*, *50*(2), 534–564.

Ghent, A. C. (2021). What's wrong with pittsburgh? delegated investors and liquidity concentration. *Journal of Financial Economics*, *139*(2), 337–358.

Glancy, D., Kurtzman, R., & Loewenstein, L. (2024). Cre redevelopment options and the use of mortgage financing. *Real Estate Economics*.

Gorback, C. S., & Keys, B. J. (2020). *Global capital and local assets: House prices, quantities, and elasticities* (tech. rep.). National Bureau of Economic Research.

Halket, J., Loewenstein, L., & Willen, P. S. (2023). *The cross-section of housing returns* (tech. rep.). Working paper.

Hartman-Glaser, B., & Mann, W. (2017). Collateral constraints, wealth effects, and volatility: Evidence from real estate markets. *Working Paper*.

Hughes, S. (2024). Debt rollover in rental housing. *Working Paper*.

Korevaar, M. (2023). Reaching for yield and the housing market: Evidence from 18th-century amsterdam. *Journal of Financial Economics*, *148*(3), 273–296.

Landvoigt, T., Piazzesi, M., & Schneider, M. (2015). The housing market (s) of san diego. *American Economic Review*, *105*(4), 1371–1407.

Li, X. (2022). Do new housing units in your backyard raise your rents? *Journal of Economic Geography*, *22*(6), 1309–1352.

Louie, S., Mondragon, J., & Wieland, J. (2025). Supply constraints do not explain house price and quantity growth across us cities. *NBER Working Paper*, (w33576).

Ma, L. (2025). Build what and for whom? the distributional effects of housing supply. *Working Paper*.

Molloy, R. (2024). Jue insight: Differences in rent growth by income from 1985 to 2021 and implications for inflation. *Journal of Urban Economics*, *142*, 103669.

Nathanson, C. G. (2023). Trickle-down housing economics. *Manuscript*.

Peng, L., & Thibodeau, T. G. (2020). Interest rates and investment: Evidence from commercial real estate. *The Journal of Real Estate Finance and Economics*, *60*(4), 554–586.

Peng, L., & Zhang, L. (2021). House prices and systematic risk: Evidence from microdata. *Real Estate Economics*, *49*(4), 1069–1092.

Pennington, K. (2021). Does building new housing cause displacement?: The supply and demand effects of construction in san francisco.

Plazzi, A., Torous, W., & Valkanov, R. (2010). Expected returns and expected growth in rents of commercial real estate. *The Review of Financial Studies*, *23*(9), 3469–3519.

Reher, M. (2021). Finance and the supply of housing quality. *Journal of Financial Economics*, *142*(1), 357–376.

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, *125*(3), 1253–1296.

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

Wang, X. B. (2023). Housing supply inequality and house price premium: An assignment model approach. *Manuscript*.

Watson, C. L., & Ziv, O. (2025a). Is the rent too high? land ownership and monopoly power. *Working Paper*.

Watson, C. L., & Ziv, O. (2025b). A test for pricing power in urban housing markets. *Review of Economics and Statistics*, *Forthcoming*.

Appendix

A Data Appendix

A.1 Financial Data

The primary dataset used in the paper is property-level financial data for multifamily rental properties collected by servicers on mortgages placed in commercial mortgage-backed securities and maintained by Trepp. There are 112 thousand multifamily properties in the Trepp dataset, covering 2000–2020.

The Trepp data includes three primary tables: Prop, Loan, and Loan2. Trepp provides property and loan IDs that link properties across tables (though they do not link properties across sequential mortgages). Prop contains monthly reporting on performance and characteristics, though the financial information we care about is only updated on average once per year. The variables we care most about are the property value at securitization, and the current (or recent) revenue and net operating income. Quality segments are defined in the Trepp data using the city-year distribution in revenue per unit (using the revenue data in the first year a property enters the data and the distribution of all other properties reporting in that year). The Loan file includes monthly information about loan payments and performance, while the Loan2 file reports the eventual disposition of the loan. Both Loan and Loan2 have relevant static and dynamic loan characteristics (including things like LTV, DSCR, maturity, extension, or prepayment terms).

In many cases, there is missing net operating income data at the beginning or ending of a property's hold period (e.g., landlords may not be required to report their annual NOI after they sell the property mid-year). Other properties are often missing NOI for some years in the middle of their hold period. We interpolate or extrapolate NOI over these missing observations using the national average NOI growth rate by quality segment (winsorizing at the 5th and 95th percentiles). We only keep property-hold period observations where NOI is not missing for more than 20% of the years where hold periods are greater than 5 years or where NOI is not missing for more than 50% of the time for hold periods less than 5 years (i.e., not missing for more than 1 year of 5 year hold periods, or non-missing for more than 2 of 4 year hold periods). After all of these restrictions, we have a sample of 10,756 property-hold periods.

A.2 Transaction Data

To study price appreciation, we use multifamily transaction data from Yardi, which includes transaction values for There are 49,436 properties with non-missing sales information (with around 111,124 thousand property records overall). We exactly match 16,112 properties in the Yardi and Trepp datasets based on address strings, and match a further 4,784 properties fuzzy string matching. Matching uses address strings and zip codes.

A.3 Supply Data

We construct new rental housing supply data using microdata from the American Community Survey. New rental units are aggregated by core-based statistical area of which there are between 260 and 265 over the 2005-2020 period. New supply is defined using the year built variable, taking new rental units as households indicating that their structure was built this year or last year. We can only construct the new units measure using data starting in the 2008 ACS because of changes in the year built coding from 2005 to 2007. We take the rent distribution in each CBSA-year cell and calculate the share of new rental units entering the rent distribution above the 75th and above the 90th percentile.

A.4 Asking Rent Data

Using property-level asking rent data from Moody's REIS platform, we analyze a constant sample of market-rate properties for the all market-rate buildings in the 50 largest cities from 2005 to 2019. The overall sample includes 37 thousand apartment properties. The buildings we include in the sample for our asking rent exercise are constructed before 2005 and continue operating through the end of our sample period. Quality is fixed in the beginning of our time period, 2005, by computing each property's percentile of average asking rent in the building relative to all other buildings in their metropolitan area. This gives us a within-city rent tercile (low, middle, and high). Asking rent growth is measured within building starting from the 2005 value of asking rent, and averaged across buildings by quality segment.

B Model

This appendix provides a theory consistent with the facts above. We show conditions under which housing supply is more elastic to aggregate demand shocks at the high-end than the low-end of the market and, in turn, rents and prices are more responsive to demand shocks at the low-end of the market.

B.1 Rental Housing Market

Housing is provided in buildings (varieties), indexed with j, each supplying h_j units of housing. Buildings are vertically differentiated and can be high- or low-quality ($k \in \{H, L\}$). Within quality types, buildings are also horizontally differentiated.

B.1.1 Demand

A mass M_t of households each demand one unit of housing and choose either a unit in a building j from one of the two quality types k, or an outside good. Indirect utility for household i from a unit of housing in building j of quality k is

$$V_{ij} = v_j + \varepsilon_{ij}, \quad v_j = \alpha^k (y_i - p^k),$$

where α^k is a quality-specific amenity shifter, p^k is the rental price of quality-k housing units, and y_i is household income, distributed $F_t(y)$. ε_{ij} are iid nested-logit shocks with dissimilarity parameter $\tau_k \in (0, 1]$ in nest k. Indirect utility from the outside good is normalized to zero for all consumers. Under these assumptions, the market share of quality-k housing among households earning income y is:

$$\lambda^k(y) = \frac{E^k}{1 + E^L + E^H}$$

where $E^k = (S^k)^{\tau^k} e^{\alpha^k (y - p^k)}$ denotes the inclusive value over a choice set of S^k type-k buildings. Integrating over household income, the aggregate market share of type k housing is

$$\int \lambda^k(y) dF_t(y)$$

.

B.1.2 Supply

Let S_t^k denote the supply of buildings of quality k on the rental market at time t. This supply adjusts period-to-period through new construction and filtering.

New Construction Each period, developers buy land and capital to construct new housing in each segment. This decision is described in more detail below. For now, let C_t^k denote the number of quality-k buildings whose construction started and was paid for in period t. There is a one-period construction lag, so C_t^k buildings are added to the quality-k housing supply in period $t + 1$.

Filtering A fixed share of the housing stock in each quality segment (δ^k) depreciates in any given period. When high-quality housing depreciates, it adds to the low-quality housing stock. When low-quality housing depreciates, it becomes unusable.

The housing supply in each segment therefore adjusts over time according to the following formulas:

$$S_t^H = (1 - \delta^H) S_{t-1}^H + C_{t-1}^H$$

$$S_t^L = (1 - \delta^L) S_{t-1}^L + C_{t-1}^L + \delta^H S_{t-1}^H$$

The developer decisions and filtering that govern this adjustment process are described in more detail in the two subsections below.

B.1.3 Equilibrium Rents

In each period, rents clear rental markets with no vacancies. The market-clearing condition in each quality segment is:

$$M_t \int \lambda^k(y)(p^H, p^L) dF_t(y) = S_t^k h_t^k,$$

where h_t^k denotes the number of housing units provided in each building of quality k in market t. These two equations implicitly define market-clearing rents p_t^{H*} and p_t^{L*}.

Developers sell properties on a competitive asset market where buildings trade at prices P_t^k equal to the discounted expected rental yield a property will generate, net of operating costs. For now, we simply assume that property prices in each quality segment P_t^k are an increasing function $\Phi_t^k(\cdot)$ of unit rents p_t^k in that segment.

B.2 Land Market and Development

Each period t, a continuum of developers enter each segment of k of the market and decide (i) whether to buy a unit of land and (ii) how many units to build on that land. Developers face a zoning cap under which they can build at most $\bar{h}$ housing units–of either quality type–on each unit of land.[10] Land is supplied with elasticity η and sold at a unit price P_t^l. The construction of high- and low-quality housing respectively require K^H and K^L units of capital per unit of housing, with $K^L < K^H$. Developers sell newly-completed units on the property market at time $t+1$ for a quality-specific price P_{t+1}^k.

Developers vary in their efficiency at delivering capital to the project. Each developer j must purchase $\phi_j \geq 1$ units of capital for each unit of capital that contributes to housing production. The developer-specific unit capital requirement for housing production in segment k is therefore $\phi_j K^k$. K^k units contribute to housing production, while $(\phi_j - 1)K^k$ units are wasted (similar to an iceberg cost). Developers draw their idiosyncratic inefficiency (or inverse productivity) parameter ϕ_j from a distribution G before deciding whether to buy land and build. Each developer can only build one building per draw, so ϕ_j could also be thought of as a building or project-specific productivity draw.

We solve the developer problem iteratively. At the beginning of time t, a developer first chooses which segment to enter, then draws a productivity term ϕ_j, and finally decides whether to buy land and how much to build on that land.

A developer who has entered segment k and drawn productivity ϕ_j chooses how many units h to develop to maximize expected profits at $t+1$:

$$\mathbb{E}\left[\pi_{j,t+1}^k(h)\right] = h\mathbb{E}\left[P_{t+1}^k\right] - \left(h\phi_j K^k + P_t^l\right) \quad s.t. \quad 0 \leq h \leq \bar{h}$$

Developers do not internalize the elasticity of prices with respect to new construction, so each builds as many units as possible as long as they expect some margin:

$$h_j^{k*} = \begin{cases} \bar{h}, & \text{if } \mathbb{E}\left[P_{t+1}^k\right] - \left(\phi_j K^k + P_t^l/\bar{h}\right) > 0 \\ 0, & \text{otherwise,} \end{cases}$$

The resulting maximum expected profit for a developer j in segment k is

$$\mathbb{E}\left[\pi_j^{k*}\right] = \bar{h}\max\{\left(\mathbb{E}\left[P_{t+1}^k\right] - \phi_j K^k\right), 0\} - P_t^l.$$

For each quality k, there is a cutoff ϕ_t^k that bounds the range of productivity draws $[1, \phi_t^k]$ over

[10] An alternative to the zoning cap could be costs increasing with height – see section 3.4 of Baum-Snow and Han 2024 and Eriksen and Orlando 2022.

which development is profitable:

$$\phi_t^k = \frac{\mathbb{E}\left[P_{t+1}^k\right] - P_t^l/\bar{h}}{K^k} \tag{B.1}$$

Only those developers with productivity draws $\phi_j \leq \phi_t^k$ build quality-k housing.

Normalizing the mass of entrants in each of the quality segments to one, the total flow of quality-k units started in period t (and completed in period $t+1$) is

$$C_t^k = \Pr\left(\phi_j \leq \phi_t^k\right)\bar{h} = G\left(\frac{\mathbb{E}\left[P_{t+1}^k\right] - P_t^l/\bar{h}}{K^k}\right)\bar{h} \tag{B.2}$$

With free entry, land prices P_t^l adjust to drive the expected profit in each quality segment k to zero, such that developers are indifferent between entering one quality segment or the other. The zero expected profit condition is

$$\mathbb{E}\left[\pi_j^{k*}\right] = \int \left[\bar{h}\left(\mathbb{E}\left[P_{t+1}^k\right] - \left(\phi K^k + P_t^l/\bar{h}\right)\right)^+\right] dG(\phi) = 0$$

The total number of developers who build and, therefore, the total land demand in period t is:

$$L^d(P_t^l, \mathbb{E}\left[P_{t+1}^H\right], \mathbb{E}\left[P_{t+1}^L\right]) = \left[G\left(\frac{\mathbb{E}\left[P_{t+1}^H\right] - P_t^l/\bar{h}}{K^H}\right) + G\left(\frac{\mathbb{E}\left[P_{t+1}^L\right] - P_t^l/\bar{h}}{K^L}\right)\right]$$

There is perfect competition between atomistic landowners who sell land to developers. Aggregate land supply is isoelastic:

$$L^s(P_t^l) = \left(\frac{P_t^l}{\bar{P}^l}\right)^\eta$$

where $\bar{P}^l$ is a scale parameter. Land prices are then implicitly defined by the market-clearing condition $L^d(P_t^l, \mathbb{E}\left[P_{t+1}^H\right], \mathbb{E}\left[P_{t+1}^L\right]) = L^s(P_t^l)$.

B.3 Asset Market and Depreciation

Property prices are determined on the asset market, where developers (and current landlords) trade buildings. Assuming this market is competitive, properties will trade at prices P_t^k equal to the discounted expected rental yield a property will generate, net of operating and capital costs, over its depreciable lifetime.

We model filtering, or property depreciation, as a Markov process. That is, a share δ^H of randomly-selected high-quality buildings convert into low-quality buildings in each period and a share δ^L of randomly-selected low-quality buildings become defunct in each period. For simplicity, we assume that landlords invest sufficient capital to offset general wear-and-tear, but they are not able to make major capital upgrades to offset the depreciation process described in section .

We assume that this capital expenditure and operating expenses together equal a fixed proportion of rents, such that properties yield a net operating margin on rents of $NOI^k = \kappa^k p^k$. Rents are expected to grow at a quality-specific constant growth rate of g^k.

Under these assumptions the expected discounted value of future cash flow from a low-quality

property at a quality-specific discount rate r^L is:

$$V^L = \sum_{s=0}^{\infty} \left[\left(1-\delta^L\right)^s \cdot \left(\frac{1+g^L}{1+r^L}\right)^s \cdot \kappa^L p_t^L \left(1+g^L\right)^s \right] = \frac{\left(1+r^L\right)}{\left(r^L-g^L\right)+\delta^L\left(1+g^L\right)} \cdot \kappa^L p_t^L \tag{B.3}$$

In each period, a high-quality property becomes low-quality with probability δ^H. Discounting with a rate r^H, the value of a high-quality property is:

$$V^H = \frac{\left(1+r^H\right)}{\left(r^H-g^H\right)+\delta^H\left(1+g^H\right)} \cdot \kappa^H p_t^H + \frac{\delta^H}{r^H+\delta^H} \cdot V^L \tag{B.4}$$

The first term reflects the value of the high-quality property cash flows expected prior to depreciation, while the second reflects the continuation value from eventually becoming a low-quality asset.

B.4 Equilibrium

A competitive equilibrium at time t is a set of prices and quantities, $\{\{p_t^k, P_t^k, C_t^k, S_t^k, \lambda_t^k(y)\}_{k\in\{H,L\}}, P_t^l\}$, such that households and developers optimally make housing and construction decisions; housing and land markets clear; and there is no arbitrage between rental and asset markets. In the rental market, the competitive equilibrium is characterized by aggregate housing-type market shares (λ_t^k) and rental prices (p_t^k) that, given the income distribution, mass of renters, and quantity of housing stock, clear housing markets (given utility maximization). In the development and land market, the short-run equilibrium is characterized by quantities of new construction (C_t^k) and a land price (P_t^l) that, given the developer productivity distribution, building unit cap, land supply elasticity, and expectations for $t+1$ rental and, accordingly, asset prices, clear land markets and generate zero maximum expected profits for developers.

The long-run competitive equilibrium is a short-run equilibrium in which prices are stable, so rational price expectations ($\mathbb{E}\left[P_{t+1}^k\right] = P_t^k$) are realized. That is, new supply exactly meets new demand.

In a steady-state equilibrium, the stock of housing in each segment will be stable ($S_t^H = S_{t-1}^H$ and $S_t^L = S_{t-1}^L$). This implies that new high-quality supply exactly offsets depreciation ($C_t^H = \delta^H S_t^H$) and the new low-quality supply, through new construction and filtering/depreciation of high-quality housing, exactly offset low-quality housing depreciation ($\delta^H S_t^H + C_t^L = \delta^L S_t^L$). This implies that a steady state requires $C_t^L = \delta^L S_t^L - \delta^H S_t^H \geq 0$, or $S^L \geq (\delta^H/\delta^L) S_t^H \Leftrightarrow \frac{S^L}{S^L+S^H} \geq \frac{\delta^H}{\delta^L+\delta^H}$.

A balanced growth path, where rental stock, construction, and demand grow at a constant rate g, is feasible when $\frac{S^L}{S^L+S^H} \geq \frac{\delta^H}{\delta^H+\delta^L+g}$. There is no low-quality construction when these conditions hold with equality. There will be more high-quality construction than low-quality when $\frac{S^L}{S^L+S^H} < \frac{2\delta^H+g}{\delta^L+2\delta^H+2g}$; that is, when the high-quality stock is sufficiently large that, in order to maintain stable stock (or constant growth), high-quality new construction (and depreciation, $C_t^H = \delta^H S_t^H$) contributes most of the new low-quality supply required to offset the deprecation of low-quality housing ($\delta^L S_t^L$).

B.5 Differential Supply Elasticity

We are interested in the elasticity of supply in each market segment S_t^k in response to income-neutral aggregate demand shocks (i.e., changes in the mass of renters M_t, holding the income distribution $F_t(y)$ fixed). We first consider the case of a positive demand shock, and then return to discuss the case of a negative demand shock.

1. Elasticity of supply in response to aggregate demand increases An unexpected increase in aggregate population ΔM at period t will increase housing demand and, in turn, rental prices (p_t^k) and, in turn, developer asset price expectations. and new construction. The supply response to a one-time, unexpected increase in aggregate population ΔM at period t in a given sector k will be a product of the responsiveness of construction to asset prices, that of asset prices to rental prices, and that of rents to demand. Using equation (B.2), (B.4), and (B.3):

$$\begin{aligned}\frac{\partial C^L}{\partial M} &= \bar{h}\frac{G'\left(\phi^L\right)}{K^L}\left(\Omega^L\frac{\partial p_t^L}{\partial M}\right)\\ \frac{\partial C^H}{\partial M} &= \bar{h}\frac{G'\left(\phi^H\right)}{K^H}\left(\Omega^H\frac{\partial p^H}{\partial M}+\Omega^L A^H\frac{\partial p^L}{\partial M}\right)\end{aligned}$$

where $\Omega^k = \frac{\kappa^k\left(1+r^k\right)}{(r^k-g^k)+\delta^k(1+g^k)}$ and $A^H = \frac{\delta^H}{r^H+\delta^H}$.

High-quality supply is more elastic to demand if $\frac{\partial C^H}{\partial M} > \frac{\partial C^L}{\partial M}$, or

$$\frac{G'\left(\phi^H\right)}{K^H}\left(\Omega^H\frac{\partial p^H}{\partial M}+\Omega^L A^H\frac{\partial p^L}{\partial M}\right) > \frac{G'\left(\phi^L\right)}{K^L}\left(\Omega^L\frac{\partial p_t^L}{\partial M}\right) \tag{B.5}$$

Assuming that productivities are drawn from a uniform distribution over $[0,\bar{\phi}]$, and we are at an interior equilibrium where both H and L quality units are built, this condition simplifies to:

$$\left(\frac{\Omega^H}{\Omega^L}\right)\left(\frac{\partial p^H/\partial M}{\partial p^L/\partial M}\right)+A^H > \frac{K^H}{K^L}$$

If the relative elasticity of the high-quality asset price premium exceeds the relative capital requirements of high-quality construction.

2. Elasticity of supply in response to aggregate demand decreases Reductions in demand will reduces entry into construction. If development continues in both quality sectors and condition (B.5) holds, then the high-quality sector will see a larger reduction in development as long as the quality premium exceeds the relative capital intensity. Even in the extreme case where development stalls in both quality segment ($C_{t-1}^H = C_{t-1}^L = 0$), an important asymmetry arises in which the low-quality supply is partially backfilled by filtering of depreciated high-quality units:

$$S_t^H = (1-\delta^H)S_{t-1}^H + C_{t-1}^H$$

$$S_t^L = (1 - \delta^L) S_{t-1}^L + C_{t-1}^L + \delta^H S_{t-1}^H$$

This back-filling will slow the decline of the low-quality housing stock in response to weakened demand.

C Robustness: Other Explanations for Outsized Low-End Property Performance

In this Appendix, we assess the robustness of our main facts and results to various alternative explanations. In particular, we show that the outsized NOI growth and returns that we measure at the low-end of the market, or the demonstrated role of high-end biased new supply in generating these differentials, are not explained by a quality bias in either renovations or delinquencies; outsized responsiveness of low-end cap rates to falling interest rates; or segment-specific local demand shocks.

C.1 Renovations

Are the changes in returns driven by quality/capital improvements of existing rental housing rather than the differential responses of new supply? In this section, we test this potential confounding mechanism. One might be concerned that the differential aggregate IRR growth, and the differential responsiveness of low-end returns to local demand growth, is due to under-reporting of capital expenditures in the CMBS data we use to measure returns. If such capital expenditures are correlated with local demand shocks and those quality changes differ across the quality distribution, they could confound our main estimates. The proposed channel of this confounder is that a positive demand shock could cause low-quality landlords to renovate or improve their building, leading to a larger rent increase or cap rate compression, which we attribute to the differential demand response channel.

We address this confounder using an indicator for renovations that we observe in the Trepp data. Annual NOI growth is nearly identical during hold periods that include renovations and those that do not (4.18 and 4.15 percent, respectively), but IRRs are about 2.7 percentage points higher during renovation hold periods (17.3 percent) than during non-renovation hold periods (14.6 percent). The return differential suggests that our results could be explained by property upgrades whose costs are not reported in our data.

To illustrate whether capital or quality improvements are driving our main results, we split the estimation sample into properties that experience a renovation during or in the 2 years before a hold period versus properties that have no record of a renovation or quality improvement in their financial reporting. In Table C.1, columns (1)-(2) replicate the main results, while columns (3)-(4) show the results with only non-renovated properties and columns (5)-(6) include only properties with a recorded renovation. Approximately 30% of the sample has a renovation in the years just before and during a hold period—this is not surprising given the reporting about apartment quality improvements during the 2010s (e.g., Reher 2021).

We find that the differential demand response is not driven by quality improvements. Non-renovated properties if anything show an even larger differential demand response in column (3), while renovated properties respond somewhat less to local demand shocks. This is not entirely surprising if new renovations take several years to complete, have higher rent levels on the lease-up of new units, and perhaps have additional frictions or constraints associated with meeting target returns on investment based on the newly renovated units. Table C.2 shows that these same results are broadly consistent using the Bartik shock measure of local labor demand. In both cases, adding

the supply controls closes or even reverses the differential demand responsiveness of high- and low-quality apartments. This suggests that new supply, and adding new supply when demand grows matters quite a lot, even for newly renovated apartment buildings.

Table C.1: IRR-employment correlation by quality segment, properties with and without renovations

	IRR					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Employment Growth	1.228***	0.795***	1.109***	0.744***	1.340***	0.839***
	(0.226)	(0.196)	(0.225)	(0.187)	(0.309)	(0.277)
Mid Qual X Employment Growth	0.541***	0.525***	0.576***	0.550***	0.475*	0.479*
	(0.161)	(0.156)	(0.181)	(0.182)	(0.255)	(0.248)
Top Qual X Employment Growth	0.059	0.336**	0.115	0.334*	0.048	0.374
	(0.182)	(0.161)	(0.186)	(0.181)	(0.313)	(0.273)
Low Qual X New Rental Share		0.204		0.316		−0.183
		(0.336)		(0.302)		(0.468)
Mid Qual X New Rental Share		−0.214		−0.081		−0.629**
		(0.164)		(0.176)		(0.286)
Top Qual X New Rental Share		−0.909***		−0.991***		−0.564
		(0.291)		(0.245)		(0.457)
Low Qual X New Units >75pctile		0.023***		0.020***		0.024*
		(0.008)		(0.008)		(0.013)
Mid Qual X New Units >75pctile		−0.002		−0.001		−0.006
		(0.004)		(0.004)		(0.008)
Top Qual X New Units >75pctile		−0.019***		−0.010**		−0.039***
		(0.006)		(0.005)		(0.011)
Low-High Quality Difference	1.169	0.459	0.994	0.410	1.292	0.465
P-value Estimates are Different	0.000	0.003	0.000	0.007	0.000	0.028
Pct Change from Baseline		60.7%	15.0%	65.0%	−10.5%	60.3%
N	27819	27819	19594	19594	8224	8224
CBSA	163	163	158	158	128	128
Adjusted R-sq	0.11	0.13	0.13	0.15	0.14	0.15

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects. Standard errors are clustered by CBSA. Columns 1 and 2 shows pooled sample with main estimating data. Columns 3 and 4 only include properties with no renovation reported during the hold period or in the 2 years prior to the hold period start. Columns 5 and 6 only include properties with a renovation reported during or in the two years before a hold period begins. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2-6 show the difference in this difference compared to the baseline specifications in columns 1.

C.2 Delinquent Loans

Are the differential changes in returns driven by differences in mortgage distress? Delinquency, default, and foreclosure could confound our main estimates through two different mechanisms. First, if distressed landlords in low-quality segments are differentially less likely to report their financial outcomes, the IRR construction could have differentially more imputed values. This could be considered a form of non-classical measurement error, potentially biasing our estimates of low-quality IRRs upwards by applying average NOI growth or value appreciation to distressed property. Second, dis-

Table C.2: IRR Bartik effect by quality segment, properties with and without renovations

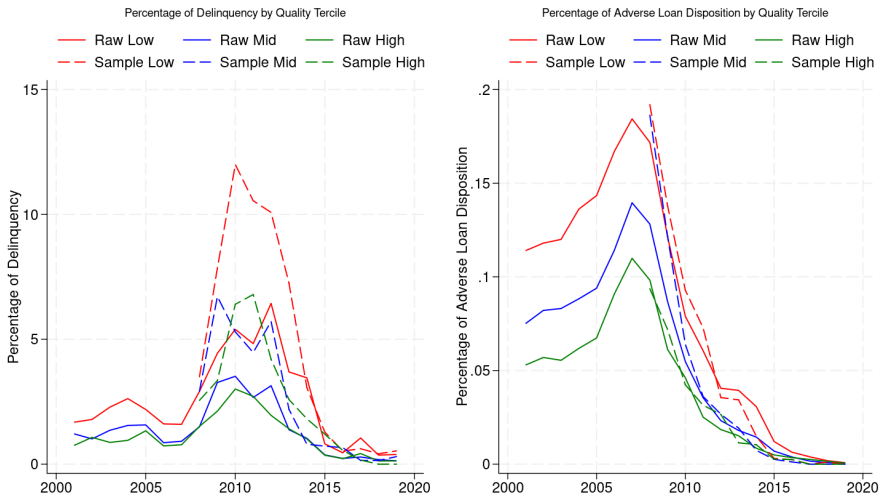
	IRR					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Bartik (Log)	1.166***	0.484*	0.743**	0.210	1.929***	0.991**
	(0.295)	(0.278)	(0.330)	(0.343)	(0.478)	(0.406)
Mid Qual X Bartik (Log)	0.539**	0.363	0.279	0.135	1.132**	0.870*
	(0.244)	(0.224)	(0.235)	(0.211)	(0.517)	(0.489)
Top Qual X Bartik (Log)	0.079	0.162	−0.114	−0.053	0.598	0.646
	(0.272)	(0.244)	(0.267)	(0.242)	(0.516)	(0.479)
Low Qual X New Rental Share		0.225		0.313		−0.168
		(0.346)		(0.318)		(0.498)
Mid Qual X New Rental Share		−0.332**		−0.190		−0.775***
		(0.152)		(0.167)		(0.272)
Top Qual X New Rental Share		−1.141***		−1.212***		−0.834*
		(0.278)		(0.227)		(0.469)
Low Qual X New Units >75pctile		0.027***		0.024**		0.030**
		(0.009)		(0.009)		(0.013)
Mid Qual X New Units >75pctile		−0.001		0.001		−0.006
		(0.004)		(0.004)		(0.009)
Top Qual X New Units >75pctile		−0.018***		−0.010**		−0.037***
		(0.006)		(0.005)		(0.011)
Low-High Quality Difference	1.086	0.322	0.857	0.263	1.330	0.344
P-value Estimates are Different	0.000	0.096	0.000	0.261	0.000	0.154
Pct Change from Baseline		70.4%	21.1%	75.8%	−22.4%	68.3%
N	27819	27819	19594	19594	8224	8224
CBSA	163	163	158	158	128	128
Adjusted R-sq	0.10	0.13	0.12	0.14	0.13	0.15

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects. Standard errors are clustered by CBSA. Columns 1 and 2 shows pooled sample with main estimating data. Columns 3 and 4 only include properties with no renovation reported during the hold period or in the 2 years prior to the hold period start. Columns 5 and 6 only include properties with a renovation reported during or in the two years before a hold period begins. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for the shift-share, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2-6 show the difference in this difference compared to the baseline specification in column 1.

tressed properties may behave differently in responding to local demand shocks. It is unclear what direction this would bias our estimates: if landlords abandon managing the property, then those distressed apartments may become extremely unresponsive to local demand shocks; on the other hand, this form of distress might heavily drive changes in returns if distressed properties are forced to sell or liquidate in a down market.

Figure C.1 shows the time series trends in the percent of outstanding mortgages delinquent in each given month and the percent that experience an adverse event. The adverse loan disposition is inclusive of foreclosure, liquidation, and forced sales. It is clear from these graphs that distress and forced sales are more common among low-quality properties.

Figure C.1: Percent of Delinquent Loans and Adverse Loan Dispositions in Raw Trepp Data and Estimation Sample, by quality segment



Note: Figures show the share of loans that are recorded as delinquent on their mortgage payments (left) and that ever have an adverse loan disposition (right). Delinquency includes any properties whose mortgage is classified as greater the 30 days delinquent. Adverse loan disposition is inclusive of foreclosure, liquidation, and forced sales. Figures show these delinquency and distress by quality segment (Low, Mid, and High) and in the raw Trepp data versus in the main estimation sample.

In the next figure, we study whether excluding properties that ever experience distress changes our main facts about returns for landlords. Figure C.2 shows that overall excluding distressed properties does little to change the overall story: low-quality properties' returns are higher and they rise much faster post-Great Financial Crisis. It is possible that the inclusion and imputation of NOI and returns for lower quality properties leads use to overestimate returns in 2010-2014 and underestimate returns from 2015-2019. It appears likely that the differences in returns are mostly driven by appreciation and realization, because Figure C.2 shows that NOI is higher, on average, in the sample excluding delinquent properties during 2010-2015.

Figure C.2: Property Performance for Delinquent and Non-Delinquent Loans, by quality segment



Note: Figures show the average IRR (left) and average NOI growth (right). Figures show these IRR and NOI growth by quality segment (Low, Mid, and High) and in the main estimation sample versus the properties in the estimation sample that never have a recorded delinquency over the hold period.

C.3 Outsized Low-End Demand Growth

To study differential demand shocks by segment, we have created several different versions of variables representing high- versus low-end rental demand at the city-level. Our primary approach uses variables constructed at the CBSA-level by taking year-over-year employment growth for the top 100-paying occupation codes and bottom 100-paying occupation codes (defined using the occupation-income distribution for the entire U.S.). We then interact high-paying occupation growth and low-paying occupation growth each with the three property quality indicators.

Appendix Table C.3 and C.4 shows these results. In specifications with these segment-specific demand variables and the average employment growth, we see that segment-specific demand does capture some of the average and differential effects of employment growth. The segment-specific household demand measures also have significant differential effects by property quality. More low-paying and high-paying jobs raise net rents for low-quality properties, but not high-quality properties. Why might this be the case? As our main results illustrate, new supply moderates the effect of local demand growth. Columns (5) and (6) of Appendix Table C.3 and C.4 add the controls for new rental housing supply. Exactly what we would expect based on the main results, the differential effects of segment-specific demand (and overall demand growth using the Bartik shocks) are mitigated or eliminated once we account for new housing supply.

We also have used income or wage growth, or changes in total payrolls (instead of growth in household or employment counts). The results are similar.

Table C.3: NOI Bartik effect by quality segment, adding controls for segment-specific demand

	NOI					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Bartik (Log)	2.241***	1.878***	2.041***			1.805***
	(0.521)	(0.515)	(0.500)			(0.492)
Mid Qual X Bartik (Log)	1.940***	1.763***	1.923***			1.743***
	(0.516)	(0.510)	(0.505)			(0.497)
Top Qual X Bartik (Log)	1.719***	1.655***	1.754***			1.611***
	(0.523)	(0.511)	(0.510)			(0.497)
Low Qual X New Rental Share		−0.171			−0.333	−0.201
		(0.297)			(0.288)	(0.299)
Mid Qual X New Rental Share		−1.022***			−1.122***	−0.991***
		(0.253)			(0.254)	(0.251)
Top Qual X New Rental Share		−1.173***			−1.257***	−1.176***
		(0.307)			(0.280)	(0.308)
Low Qual X New Units >75pctile		0.011			0.011	0.010
		(0.007)			(0.007)	(0.007)
Mid Qual X New Units >75pctile		0.009			0.010	0.009
		(0.006)			(0.006)	(0.006)
Top Qual X New Units >75pctile		−0.001			−0.003	−0.002
		(0.006)			(0.005)	(0.005)
Low Qual X Δ Low-Paying Employment Growth			0.042	0.054	0.013	0.008
			(0.047)	(0.048)	(0.047)	(0.045)
Mid Qual X Δ Low-Paying Employment Growth			0.045	0.046	0.042	0.042
			(0.036)	(0.036)	(0.036)	(0.035)
Top Qual X Δ Low-Paying Employment Growth			0.056	0.045	0.086**	0.091**
			(0.036)	(0.037)	(0.036)	(0.035)
Low Qual X Δ High-Paying Employment Growth			0.186***	0.216***	0.136***	0.119**
			(0.046)	(0.045)	(0.047)	(0.046)
Mid Qual X Δ High-Paying Employment Growth			0.041	0.047	0.050	0.044
			(0.033)	(0.033)	(0.035)	(0.036)
Top Qual X Δ High-Paying Employment Growth			−0.012	−0.033	0.047	0.055*
			(0.031)	(0.030)	(0.030)	(0.031)
Low-High Quality Difference	0.522	0.223				0.194
P-value Estimates are Different	0.000	0.114				0.183
Pct Change from Baseline		57.3%				62.8%
N	27819	27819	27819	27819	27819	27819
CBSA	163	163	163	163	163	163
Adjusted R-sq	0.04	0.05	0.04	0.04	0.04	0.05

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects. Standard errors are clustered by CBSA.

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is annual log change in NOI. Change in NOI and Bartik shift-share variable measured in year-over-year change in logs. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for the shift-share, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2 and 6 show the difference in this difference compared to the baseline specification in column 1. Columns 3-6 control for the quality segment indicators interacted with the year-over-year employment growth in ACS for the top 100-paying occupation codes and bottom 100-paying occupation codes (defined using the occupation-income distribution for the U.S. overall).

Table C.4: IRR Bartik effect by quality segment, adding controls for segment-specific demand

	IRR					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Bartik (Log)	1.080***	0.439	0.766***			0.397
	(0.286)	(0.273)	(0.257)			(0.261)
Mid Qual X Bartik (Log)	0.507**	0.360	0.461*			0.333
	(0.241)	(0.224)	(0.233)			(0.222)
Top Qual X Bartik (Log)	0.106	0.173	0.312			0.212
	(0.265)	(0.242)	(0.250)			(0.239)
Low Qual X New Rental Share		0.175			0.068	0.118
		(0.320)			(0.316)	(0.321)
Mid Qual X New Rental Share		−0.382**			−0.462***	−0.423***
		(0.155)			(0.154)	(0.156)
Top Qual X New Rental Share		−1.109***			−1.026***	−1.044***
		(0.267)			(0.254)	(0.264)
Low Qual X New Units >75pctile		0.026***			0.026***	0.025***
		(0.009)			(0.008)	(0.008)
Mid Qual X New Units >75pctile		−0.002			−0.002	−0.003
		(0.004)			(0.004)	(0.004)
Top Qual X New Units >75pctile		−0.017***			−0.017***	−0.016***
		(0.005)			(0.005)	(0.005)
Low Qual X Δ Low-Paying Employment Growth			0.131***	0.152***	0.065**	0.060*
			(0.043)	(0.041)	(0.029)	(0.034)
Mid Qual X Δ Low-Paying Employment Growth			0.025	0.021	0.029*	0.027
			(0.020)	(0.018)	(0.015)	(0.017)
Top Qual X Δ Low-Paying Employment Growth			−0.087***	−0.104***	−0.025**	−0.018
			(0.022)	(0.021)	(0.013)	(0.015)
Low Qual X Δ High-Paying Employment Growth			0.195***	0.244***	0.066	0.053
			(0.046)	(0.060)	(0.045)	(0.033)
Mid Qual X Δ High-Paying Employment Growth			0.036*	0.028	0.054***	0.050***
			(0.021)	(0.019)	(0.017)	(0.017)
Top Qual X Δ High-Paying Employment Growth			−0.148***	−0.185***	−0.031**	−0.016
			(0.028)	(0.032)	(0.015)	(0.016)
Low-High Quality Difference	0.974	0.267				0.184
P-value Estimates are Different	0.000	0.160				0.327
Pct Change from Baseline		72.6%				81.1%
N	27819	27819	27819	27819	27819	27819
CBSA	163	163	163	163	163	163
Adjusted R-sq	0.12	0.15	0.13	0.13	0.15	0.15

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for the shift-share, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2 and 6 show the difference in this difference compared to the baseline specification in column 1. Columns 3-6 control for the quality segment indicators interacted with the year-over-year employment growth in ACS for the top 100-paying occupation codes and bottom 100-paying occupation codes (defined using the occupation-income distribution for the U.S. overall).

C.4 Falling Interest Rates

One explanation for cap rate and return dynamics over the 2010s is that as interest rates declined towards zero, the cap rate distribution compressed. Naturally (the story goes), a marginal interest rate change has a larger effect on cap rates (and therefore prices and returns) for lower-quality buildings because lower-quality buildings' cap rates had further to fall. We propose several potential tests of this hypothesis.

First in Table C.5, we regress internal rates of return (average annual returns over a hold period) on the change in 10-year Treasury yields over the hold period. To allow the effect of yield compression to vary, we interact the change in Treasury yields with the own-property entry cap rate. We find that changes in our proxy for risk-free rates are positively correlated with IRRs. The interaction between entry cap rate and change in yields is negative, suggesting that at relatively higher cap rates, Treasury yields actually pass through less to returns. The second column shows that this interaction effect essentially disappears when controlling for the entry cap rate. The final columns show suggestive evidence that part of the differential responsiveness of IRRs to local demand shocks may be related to differences in initial costs of the investment. However, this empirical approach may be problematic because the entry cap rate is a component of the IRR (the cost basis for the investment), generating a mechanical (one-for-one) relationship with returns.

Appendix Tables D.9 through Table D.12 replicate Table C.5 using four different specifications to capture initial cost of capital: national-level Treasury yields, which only vary by year; CBSA-level cap rates which vary over CBSA and years; national quality segment-specific cap rates which vary by segment and year; and CBSA-level quality segment-specific average entry cap rates, varying at the CBSA-by-segment-by-year level. Because we impute prices and IRRs on properties currently operating at the end of our dataset in 2020, we drop those properties in this specification so as not to overweight the unusual aspects of the 2020 COVID economy. Each of these specifications shows that the change in Treasury yields is positively related to average returns conditional on year and CBSA fixed effects. In the case of segment or CBSA-specific cap rates, controlling for the initial cost basis seems to have a one-for-one or 1.1 relationship with IRR.

There is essentially no evidence that an outsized response of low-end property values to declining Treasury yields explains the outsized low-end return growth. If higher cap rate segments of the market were more responsive to falling interest rates, then the d(Treasury yield) x entry cap coefficient would be positive: properties with larger initial cap rate should see greater return growth when Treasury rates fall. Instead, nearly every specification shows the opposite result: Properties with higher initial cap rates see less return growth in response to changes in Treasury rates. It appears that controlling for these proxies for changes in Treasury yields and the initial cost basis (entry cap rate) explain some of the differential responsiveness to local demand shocks across segments. However, none of the cap rate specifications fully explain the low- versus high-quality gap, and the housing supply variables still explain a portion of the different responses separately from the cost of capital controls.

Table C.5: Segment-specific IRR-employment correlation, controlling for Own-Entry Cap

	IRR (1)	IRR (2)	IRR (3)	IRR (4)	IRR (5)	IRR (6)	IRR (7)
d(10yr Treasury Yield)	0.064***	0.032***		0.063***	0.062***	0.031***	0.031***
	(0.005)	(0.005)		(0.005)	(0.005)	(0.004)	(0.004)
d(10yr Yield) X Entry Cap Rate	-0.460***	-0.028		-0.458***	-0.457***	-0.029	-0.031
	(0.045)	(0.054)		(0.045)	(0.045)	(0.050)	(0.047)
Entry Cap Rate		1.100***				1.091***	1.082***
		(0.111)				(0.104)	(0.096)
Low Qual X Employment Growth			1.455***	1.317***	0.969***	1.144***	0.851***
			(0.258)	(0.224)	(0.188)	(0.188)	(0.152)
Mid Qual X Employment Growth			0.824***	0.683***	0.721***	0.646***	0.670***
			(0.212)	(0.189)	(0.200)	(0.162)	(0.173)
Top Qual X Employment Growth			0.467**	0.411**	0.504**	0.329*	0.421**
			(0.217)	(0.196)	(0.194)	(0.169)	(0.169)
Low Qual X New Rental Share					-0.187		-0.060
					(0.391)		(0.356)
Mid Qual X New Rental Share					-0.077		-0.138
					(0.252)		(0.208)
Top Qual X New Rental Share					-0.687**		-0.659**
					(0.341)		(0.283)
Low Qual X New Units >75pctile					0.039***		0.028***
					(0.010)		(0.010)
Mid Qual X New Units >75pctile					-0.009		-0.008
					(0.007)		(0.006)
Top Qual X New Units >75pctile					-0.011*		-0.013**
					(0.006)		(0.006)
Low-High Quality Difference			0.988	0.906	0.465	0.815	0.430
P-value Estimates are Different			0.000	0.000	0.001	0.000	0.001
Pct Change from Baseline				8.3%	53.0%	17.5%	56.5%
N	13026	13026	13026	13026	13026	13026	13026
CBSA	136	136	136	136	136	136	136
Adjusted R-sq	0.26	0.41	0.18	0.27	0.28	0.42	0.43

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Includes independent variable for change in 10-year Treasury yields over the hold period, and the change in Treasury yields interacted with the property's on entry cap rate. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution.

D Additional Tables and Figures

Table D.1: Average Response of NOI to Labor Demand

	NOI Growth (1)	NOI Growth (2)	Emp Growth (3)	NOI Growth (4)
Employment Growth	1.444***			1.618***
	(0.166)			(0.366)
New Rental Share	-0.552**	-0.819***	-0.193**	-0.507**
	(0.228)	(0.248)	(0.080)	(0.256)
New Units >75pctile	0.000	0.005	0.004**	-0.001
	(0.005)	(0.006)	(0.002)	(0.005)
Bartik (3-Digit, Log)		1.710***	1.060***	
		(0.543)	(0.263)	
Spec	OLS	RF	FS	IV
N	27819	27819	27819	27819
F Stats	14.87	4.13	8.61	
CBSA	163	163	163	163
Adjusted R-sq	0.05	0.04	0.85	0.01
Year FE	Yes	Yes	Yes	Yes
CBSA FE	Yes	Yes	Yes	Yes

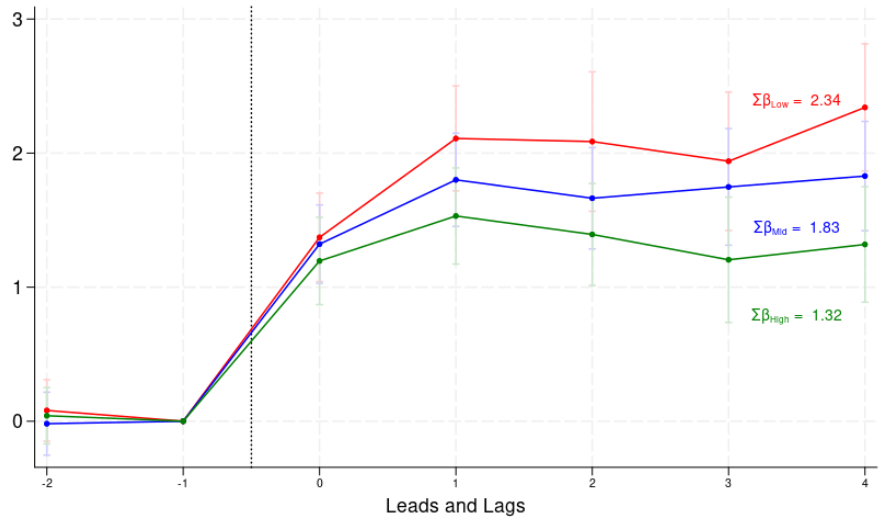
Note: *** p<0.01, ** p<0.05, * p<0.10. Standard errors are clustered by CBSA, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Each specification includes year and CBSA fixed effects. Change in NOI and employment measured in year-over-year change in logs. Columns show ordinary least squares, reduced form, first stage, and instrumental variables specification. Bartik shift-share variable is constructed using 3-digit NAICS industry codes where shifts are annual changes in log employment by industry. Includes controls for new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution.

Table D.2: Average Response of IRR to Labor Demand

	IRR	IRR	Emp Growth	IRR
	(1)	(2)	(3)	(4)
Employment Growth	0.519***			0.363
	(0.160)			(0.239)
New Rental Share	-0.296**	-0.407***	-0.193**	-0.337**
	(0.141)	(0.136)	(0.080)	(0.151)
New Units >75pctile	-0.000	0.002	0.004**	0.000
	(0.003)	(0.004)	(0.002)	(0.004)
Bartik (3-Digit, Log)		0.383	1.060***	
		(0.256)	(0.263)	
Spec	OLS	RF	FS	IV
N	27819	27819	27819	27819
F Stats	9.91	8.57	8.61	
CBSA	163	163	163	163
Adjusted R-sq	0.12	0.12	0.85	0.03
Year FE	Yes	Yes	Yes	Yes
CBSA FE	Yes	Yes	Yes	Yes

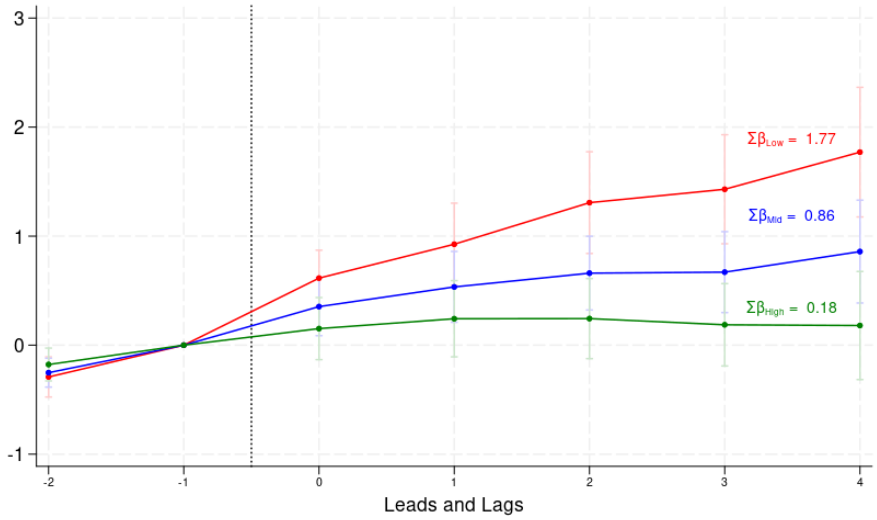
Note: *** p<0.01, ** p<0.05, * p<0.10. Standard errors are clustered by CBSA, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Each specification includes year and CBSA fixed effects. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Columns show ordinary least squares, reduced form, first stage, and instrumental variables specification. Bartik shift-share variable is constructed using 3-digit NAICS industry codes where shifts are annual changes in log employment by industry. Includes controls for new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution.

Figure D.1: Dynamics of Employment Growth on NOI Growth with Distributed Leads & Lags by Segment



Note: *** p<0.01, ** p<0.05, * p<0.10. Standard errors are clustered by CBSA. Each specification includes year and CBSA fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Change in NOI and employment measured in year-over-year change in logs. Plotted coefficients and 95% confidence intervals from specification with two leads and four lags of employment growth (coefficients are cumulatively summed and reported relative to the first lead). Sum of lead and lag coefficients is reported in the top right corner of each plot.

Figure D.2: Dynamics of Employment Growth on IRR with Distributed Leads & Lags by Segment



Note: *** p<0.01, ** p<0.05, * p<0.10. Standard errors are clustered by CBSA. Each specification includes year and CBSA fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Plotted coefficients and 95% confidence intervals from specification with two leads and four lags of employment growth (coefficients are cumulatively summed and reported relative to the first lead). Sum of lead and lag coefficients is reported in the top right corner of each plot.

Table D.3: NOI Bartik effect by quality segment, adding controls

	NOI					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Bartik (Log)	2.241***	1.878***		1.927***		1.992***
	(0.521)	(0.515)		(0.526)		(0.517)
Mid Qual X Bartik (Log)	1.940***	1.763***		1.797***		1.889***
	(0.516)	(0.510)		(0.514)		(0.514)
Top Qual X Bartik (Log)	1.719***	1.655***		1.630***		1.791***
	(0.523)	(0.511)		(0.515)		(0.517)
Low Qual X New Rental Share		−0.171	−0.008	0.010		
		(0.297)	(0.237)	(0.266)		
Mid Qual X New Rental Share		−1.022***	−1.017***	−0.899***		
		(0.253)	(0.243)	(0.254)		
Top Qual X New Rental Share		−1.173***	−1.625***	−1.418***		
		(0.307)	(0.265)	(0.282)		
Low Qual X New Units >75pctile		0.011			0.021***	0.019***
		(0.007)			(0.007)	(0.006)
Mid Qual X New Units >75pctile		0.009			0.005	0.005
		(0.006)			(0.006)	(0.006)
Top Qual X New Units >75pctile		−0.001			−0.010	−0.008
		(0.006)			(0.006)	(0.006)
Low-High Quality Difference	0.522	0.223		0.297		0.201
P-value Estimates are Different	0.000	0.114		0.020		0.142
Pct Change from Baseline		57.3%		43.1%		61.5%
N	27819	27819	27819	27819	27819	27819
CBSA	163	163	163	163	163	163
Adjusted R-sq	0.04	0.05	0.04	0.05	0.04	0.04

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is annual log change in NOI. Change in NOI and Bartik shift-share variable measured in year-over-year change in logs. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for the shift-share, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1.

Table D.4: IRR Bartik effect by quality segment, adding controls

	IRR					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Bartik (Log)	1.080***	0.439		0.636**		0.501*
	(0.286)	(0.273)		(0.269)		(0.270)
Mid Qual X Bartik (Log)	0.507**	0.360		0.359		0.428*
	(0.241)	(0.224)		(0.228)		(0.225)
Top Qual X Bartik (Log)	0.106	0.173		0.093		0.269
	(0.265)	(0.242)		(0.248)		(0.243)
Low Qual X New Rental Share		0.175	1.168***	1.019***		
		(0.320)	(0.305)	(0.292)		
Mid Qual X New Rental Share		−0.382**	−0.563***	−0.520***		
		(0.155)	(0.164)	(0.170)		
Top Qual X New Rental Share		−1.109***	−1.933***	−1.762***		
		(0.267)	(0.310)	(0.300)		
Low Qual X New Units >75pctile		0.026***			0.036***	0.035***
		(0.009)			(0.008)	(0.008)
Mid Qual X New Units >75pctile		−0.002			−0.001	−0.001
		(0.004)			(0.004)	(0.004)
Top Qual X New Units >75pctile		−0.017***			−0.030***	−0.027***
		(0.005)			(0.005)	(0.005)
Low-High Quality Difference	0.974	0.267		0.543		0.232
P-value Estimates are Different	0.000	0.160		0.004		0.205
Pct Change from Baseline		72.6%		44.3%		76.1%
N	27819	27819	27819	27819	27819	27819
CBSA	163	163	163	163	163	163
Adjusted R-sq	0.12	0.15	0.14	0.14	0.14	0.14

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for the shift-share, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1.

Table D.5: NOI employment effect by quality segment, adding volatility

	NOI					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Employment Growth	1.797***	1.580***		1.812***		1.552***
	(0.189)	(0.196)		(0.190)		(0.183)
Mid Qual X Employment Growth	1.520***	1.507***		1.492***		1.470***
	(0.161)	(0.163)		(0.153)		(0.150)
Top Qual X Employment Growth	1.255***	1.352***		1.241***		1.345***
	(0.173)	(0.163)		(0.166)		(0.155)
Low Qual X New Rental Share		0.008			−0.363	−0.029
		(0.289)			(0.279)	(0.285)
Mid Qual X New Rental Share		−0.800***			−1.253***	−0.861***
		(0.242)			(0.265)	(0.237)
Top Qual X New Rental Share		−0.849***			−1.327***	−0.882***
		(0.297)			(0.285)	(0.286)
Low Qual X New Units >75pctile		0.004			0.017**	0.008
		(0.006)			(0.008)	(0.006)
Mid Qual X New Units >75pctile		0.002			0.005	−0.000
		(0.006)			(0.006)	(0.006)
Top Qual X New Units >75pctile		−0.008*			−0.008	−0.011**
		(0.005)			(0.005)	(0.005)
Volatility			−0.017	−0.021	−0.053**	−0.043**
			(0.018)	(0.016)	(0.020)	(0.018)
Low-High Quality Difference	0.542	0.228				0.208
P-value Estimates are Different	0.000	0.019				0.030
Pct Change from Baseline		58.0%				61.7%
N	27709	27709	27709	27709	27709	27709
CBSA	155	155	155	155	155	155
Adjusted R-sq	0.05	0.06	0.04	0.05	0.05	0.06

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is annual log change in NOI. Change in NOI and employment measured in year-over-year change in logs. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2 and 6 show the difference in this difference compared to the baseline specification in column 1. Controls for volatility calculate standard deviation of annual growth in log NOI by CBSA and year.

Table D.6: IRR employment effect by quality segment, adding volatility

	IRR					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Employment Growth	1.182***	0.770***		1.161***		0.765***
	(0.221)	(0.191)		(0.222)		(0.191)
Mid Qual X Employment Growth	0.532***	0.534***		0.570***		0.528***
	(0.158)	(0.154)		(0.164)		(0.152)
Top Qual X Employment Growth	0.107	0.355**		0.127		0.354**
	(0.174)	(0.160)		(0.178)		(0.159)
Low Qual X New Rental Share		0.175			0.122	0.169
		(0.312)			(0.317)	(0.312)
Mid Qual X New Rental Share		−0.267			−0.442***	−0.277*
		(0.169)			(0.152)	(0.167)
Top Qual X New Rental Share		−0.878***			−1.095***	−0.884***
		(0.278)			(0.255)	(0.277)
Low Qual X New Units >75pctile		0.022***			0.029***	0.022***
		(0.007)			(0.009)	(0.008)
Mid Qual X New Units >75pctile		−0.003			−0.002	−0.003
		(0.004)			(0.004)	(0.004)
Top Qual X New Units >75pctile		−0.018***			−0.020***	−0.018***
		(0.005)			(0.005)	(0.005)
Volatility			0.047***	0.029***	−0.012	−0.007
			(0.012)	(0.007)	(0.008)	(0.007)
Low-High Quality Difference	1.075	0.414				0.411
P-value Estimates are Different	0.000	0.005				0.005
Pct Change from Baseline		61.4%				61.7%
N	27709	27709	27709	27709	27709	27709
CBSA	155	155	155	155	155	155
Adjusted R-sq	0.14	0.15	0.12	0.14	0.15	0.15

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2 and 6 show the difference in this difference compared to the baseline specification in column 1. Controls for volatility calculate standard deviation of annual growth in log NOI by CBSA and year.

Table D.7: NOI employment effect by quality segment, adding LIHTC share

	NOI					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Employment Growth	1.790***	1.575***		1.705***		1.577***
	(0.187)	(0.193)		(0.191)		(0.191)
Mid Qual X Employment Growth	1.516***	1.501***		1.512***		1.502***
	(0.161)	(0.162)		(0.158)		(0.160)
Top Qual X Employment Growth	1.250***	1.345***		1.319***		1.351***
	(0.171)	(0.160)		(0.168)		(0.159)
Low Qual X New Rental Share		0.021			−0.327	0.016
		(0.289)			(0.284)	(0.288)
Mid Qual X New Rental Share		−0.778***			−1.186***	−0.790***
		(0.238)			(0.259)	(0.233)
Top Qual X New Rental Share		−0.834***			−1.252***	−0.797***
		(0.290)			(0.287)	(0.295)
Low Qual X New Units >75pctile		0.006			0.012	0.005
		(0.006)			(0.008)	(0.007)
Mid Qual X New Units >75pctile		0.003			0.008	0.002
		(0.006)			(0.007)	(0.006)
Top Qual X New Units >75pctile		−0.007			−0.001	−0.006
		(0.005)			(0.005)	(0.005)
Low Qual X Sh. New LIHTC Units			2.983***	1.422	0.918	0.007
			(1.103)	(1.009)	(1.196)	(1.108)
Mid Qual X Sh. New LIHTC Units			0.253	−0.358	0.919	0.089
			(1.075)	(1.001)	(1.240)	(1.087)
Top Qual X Sh. New LIHTC Units			−2.552**	−2.347***	−0.205	−0.887
			(1.050)	(0.894)	(1.041)	(0.933)
Low-High Quality Difference	0.540	0.230				0.226
P-value Estimates are Different	0.000	0.017				0.021
Pct Change from Baseline		57.4%				58.2%
N	27819	27819	27819	27819	27819	27819
CBSA	163	163	163	163	163	163
Adjusted R-sq	0.05	0.06	0.04	0.05	0.04	0.06

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is annual log change in NOI. Change in NOI and employment measured in year-over-year change in logs. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2 and 6 show the difference in this difference compared to the baseline specification in column 1. Controls for volatility calculate standard deviation of annual growth in log NOI by CBSA and year. Controls for LIHTC are calculated as the share of new LIHTC units relative to overall rental units by CBSA and year.

Table D.8: IRR employment effect by quality segment, adding LIHTC share

	IRR					
	(1)	(2)	(3)	(4)	(5)	(6)
Low Qual X Employment Growth	1.173***	0.762***		1.050***		0.781***
	(0.218)	(0.189)		(0.208)		(0.188)
Mid Qual X Employment Growth	0.526***	0.527***		0.515***		0.535***
	(0.155)	(0.151)		(0.151)		(0.150)
Top Qual X Employment Growth	0.103	0.351**		0.219		0.361**
	(0.172)	(0.157)		(0.169)		(0.157)
Low Qual X New Rental Share		0.173			0.203	0.256
		(0.311)			(0.315)	(0.308)
Mid Qual X New Rental Share		−0.267			−0.446***	−0.280*
		(0.166)			(0.152)	(0.167)
Top Qual X New Rental Share		−0.889***			−1.084***	−0.875***
		(0.277)			(0.261)	(0.285)
Low Qual X New Units >75pctile		0.022***			0.030***	0.024***
		(0.007)			(0.009)	(0.007)
Mid Qual X New Units >75pctile		−0.003			−0.004	−0.005
		(0.004)			(0.004)	(0.004)
Top Qual X New Units >75pctile		−0.018***			−0.021***	−0.019***
		(0.005)			(0.006)	(0.005)
Low Qual X Sh. New LIHTC Units			3.973***	1.644**	−1.711*	−2.178**
			(0.876)	(0.822)	(0.876)	(0.877)
Mid Qual X Sh. New LIHTC Units			−0.903	−0.767	0.080	−0.206
			(0.636)	(0.577)	(0.592)	(0.587)
Top Qual X Sh. New LIHTC Units			−5.738***	−4.260***	−0.534	−0.668
			(1.240)	(0.935)	(0.795)	(0.806)
Low-High Quality Difference	1.070	0.411				0.419
P-value Estimates are Different	0.000	0.005				0.004
Pct Change from Baseline		61.6%				60.8%
N	27819	27819	27819	27819	27819	27819
CBSA	163	163	163	163	163	163
Adjusted R-sq	0.14	0.15	0.13	0.14	0.15	0.15

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1. Controls for LIHTC are calculated as the share of new LIHTC units relative to overall rental units by CBSA and year.

Table D.9: Cap Rate Compression Using Treasury Yield

	IRR	IRR	IRR	IRR	IRR	IRR	IRR
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
d(10yr Treasury Yield)	0.056***	0.056***		0.054***	0.053***	0.054***	0.053***
	(0.010)	(0.010)		(0.010)	(0.010)	(0.010)	(0.010)
d(10yr Yield) X Entry Yield	-0.006***	-0.006***		-0.006***	-0.006***	-0.006***	-0.006***
	(0.002)	(0.002)		(0.002)	(0.002)	(0.002)	(0.002)
Entry 10yr Treasury Yield		0.000				0.000	0.000
		(.)				(.)	(.)
Low Qual X Employment Growth			1.455***	1.313***	0.974***	1.313***	0.974***
			(0.258)	(0.240)	(0.205)	(0.240)	(0.205)
Mid Qual X Employment Growth			0.824***	0.677***	0.700***	0.677***	0.700***
			(0.212)	(0.199)	(0.210)	(0.199)	(0.210)
Top Qual X Employment Growth			0.467**	0.381*	0.489**	0.381*	0.489**
			(0.217)	(0.205)	(0.204)	(0.205)	(0.204)
Low Qual X New Rental Share					-0.038		-0.038
					(0.414)		(0.414)
Mid Qual X New Rental Share					-0.042		-0.042
					(0.286)		(0.286)
Top Qual X New Rental Share					-0.771**		-0.771**
					(0.346)		(0.346)
Low Qual X New Units >75pctile					0.035***		0.035***
					(0.011)		(0.011)
Mid Qual X New Units >75pctile					-0.007		-0.007
					(0.007)		(0.007)
Top Qual X New Units >75pctile					-0.013*		-0.013*
					(0.006)		(0.006)
Low-High Quality Difference			0.988	0.932	0.485	0.932	0.485
P-value Estimates are Different			0.000	0.000	0.001	0.000	0.001
Pct Change from Baseline				5.7%	50.9%	5.7%	50.9%
N	13026	13026	13026	13026	13026	13026	13026
CBSA	136	136	136	136	136	136	136
Adjusted R-sq	0.20	0.20	0.18	0.22	0.23	0.22	0.23

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1. Controls for Treasury yields take the change in the 10-year yield over the hold period, and interact it with the Treasury yield at the beginning of the hold period.

Table D.10: Cap Rate Compression Using CBSA-Level Entry Cap Rate

	IRR (1)	IRR (2)	IRR (3)	IRR (4)	IRR (5)	IRR (6)	IRR (7)
d(10yr Treasury Yield)	0.053***	0.025***		0.052***	0.051***	0.025***	0.025***
	(0.006)	(0.008)		(0.006)	(0.006)	(0.008)	(0.008)
d(10yr Yield) X Entry Cap Rate, CBSA	-0.328***	-0.006		-0.324***	-0.324***	-0.007	-0.011
	(0.077)	(0.095)		(0.075)	(0.071)	(0.093)	(0.090)
Entry Cap Rate, CBSA		0.867***				0.852***	0.842***
		(0.146)				(0.140)	(0.138)
Low Qual X Employment Growth			1.455***	1.323***	0.977***	1.259***	0.930***
			(0.258)	(0.237)	(0.202)	(0.217)	(0.182)
Mid Qual X Employment Growth			0.824***	0.697***	0.735***	0.647***	0.691***
			(0.212)	(0.200)	(0.211)	(0.186)	(0.198)
Top Qual X Employment Growth			0.467**	0.411**	0.508**	0.344*	0.452**
			(0.217)	(0.206)	(0.205)	(0.191)	(0.191)
Low Qual X New Rental Share					-0.089		0.017
					(0.405)		(0.384)
Mid Qual X New Rental Share					0.047		0.143
					(0.265)		(0.264)
Top Qual X New Rental Share					-0.811**		-0.732**
					(0.340)		(0.328)
Low Qual X New Units >75pctile					0.037***		0.034***
					(0.011)		(0.011)
Mid Qual X New Units >75pctile					-0.010		-0.012*
					(0.007)		(0.007)
Top Qual X New Units >75pctile					-0.010		-0.012*
					(0.007)		(0.006)
Low-High Quality Difference			0.988	0.912	0.469	0.914	0.479
P-value Estimates are Different			0.000	0.000	0.001	0.000	0.001
Pct Change from Baseline				7.7%	52.6%	7.5%	51.5%
N	13026	13026	13026	13026	13026	13026	13026
CBSA	136	136	136	136	136	136	136
Adjusted R-sq	0.22	0.24	0.18	0.23	0.24	0.25	0.26

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1. Controls for Treasury yields take the change in the 10-year yield over the hold period, and interact it with the average entry cap rate in the CBSA at the beginning of the hold period.

Table D.11: Cap Rate Compression Using National-Level Quality Segment Entry Cap Rate

	IRR (1)	IRR (2)	IRR (3)	IRR (4)	IRR (5)	IRR (6)	IRR (7)
d(10yr Treasury Yield)	0.066***	0.023***		0.065***	0.065***	0.032***	0.044***
	(0.006)	(0.008)		(0.006)	(0.006)	(0.008)	(0.008)
d(10yr Yield) X Entry Cap Rate, Q	-0.486***	0.028		-0.480***	-0.489***	-0.083	-0.232**
	(0.067)	(0.104)		(0.068)	(0.069)	(0.102)	(0.099)
Entry Cap Rate, Q		1.393***				1.082***	0.702***
		(0.219)				(0.185)	(0.186)
Low Qual X Employment Growth			1.469***	1.342***	0.987***	1.159***	0.910***
			(0.261)	(0.237)	(0.203)	(0.222)	(0.199)
Mid Qual X Employment Growth			0.838***	0.687***	0.730***	0.754***	0.776***
			(0.216)	(0.200)	(0.210)	(0.204)	(0.212)
Top Qual X Employment Growth			0.480**	0.426**	0.526**	0.558***	0.589***
			(0.219)	(0.204)	(0.203)	(0.210)	(0.206)
Low Qual X New Rental Share					0.017		0.027
					(0.430)		(0.439)
Mid Qual X New Rental Share					-0.015		-0.014
					(0.274)		(0.273)
Top Qual X New Rental Share					-0.770**		-0.702**
					(0.347)		(0.352)
Low Qual X New Units >75pctile					0.037***		0.032***
					(0.011)		(0.011)
Mid Qual X New Units >75pctile					-0.010		-0.010
					(0.007)		(0.007)
Top Qual X New Units >75pctile					-0.011		-0.008
					(0.007)		(0.007)
Low-High Quality Difference			0.989	0.916	0.462	0.602	0.321
P-value Estimates are Different			0.000	0.000	0.001	0.000	0.015
Pct Change from Baseline				7.3%	53.3%	39.2%	67.5%
N	13015	13015	13015	13015	13015	13015	13015
CBSA	136	136	136	136	136	136	136
Adjusted R-sq	0.21	0.22	0.18	0.23	0.24	0.23	0.24

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1. Controls for Treasury yields take the change in the 10-year yield over the hold period, and interact it with the average entry cap rate in the quality segment across the overall sample at the beginning of the property's hold period.

Table D.12: Cap Rate Compression Using CBSA-Level Quality Segment Entry Cap Rate

	IRR (1)	IRR (2)	IRR (3)	IRR (4)	IRR (5)	IRR (6)	IRR (7)
d(10yr Treasury Yield)	0.067***	0.024***		0.065***	0.065***	0.025***	0.026***
	(0.005)	(0.005)		(0.005)	(0.005)	(0.005)	(0.005)
d(10yr Yield) X Entry Cap Rate, Q-CBSA	-0.482***	0.043		-0.477***	-0.478***	0.029	0.013
	(0.052)	(0.056)		(0.051)	(0.051)	(0.054)	(0.054)
Entry Cap Rate, Q-CBSA		1.345***				1.300***	1.262***
		(0.129)				(0.119)	(0.120)
Low Qual X Employment Growth			1.469***	1.310***	0.958***	0.953***	0.717***
			(0.261)	(0.229)	(0.191)	(0.170)	(0.139)
Mid Qual X Employment Growth			0.838***	0.672***	0.715***	0.565***	0.612***
			(0.216)	(0.191)	(0.202)	(0.162)	(0.175)
Top Qual X Employment Growth			0.480**	0.411**	0.504**	0.371**	0.445**
			(0.219)	(0.198)	(0.196)	(0.169)	(0.173)
Low Qual X New Rental Share					-0.109		0.103
					(0.401)		(0.345)
Mid Qual X New Rental Share					0.013		0.079
					(0.263)		(0.229)
Top Qual X New Rental Share					-0.744**		-0.495*
					(0.333)		(0.289)
Low Qual X New Units >75pctile					0.039***		0.024**
					(0.011)		(0.010)
Mid Qual X New Units >75pctile					-0.010		-0.009*
					(0.007)		(0.006)
Top Qual X New Units >75pctile					-0.010		-0.008
					(0.006)		(0.006)
Low-High Quality Difference			0.989	0.899	0.454	0.582	0.273
P-value Estimates are Different			0.000	0.000	0.001	0.000	0.009
Pct Change from Baseline				9.1%	54.1%	41.1%	72.4%
N	13015	13015	13015	13015	13015	13015	13015
CBSA	136	136	136	136	136	136	136
Adjusted R-sq	0.24	0.33	0.18	0.26	0.27	0.33	0.34

Note: *** p<0.01, ** p<0.05, * p<0.10. All specifications include CBSA and year fixed effects, as well as a set of binary indicator variables capturing the timing of a renovation including one for the year of renovation, 4 lags for years after the renovation, and 6 leads for the years preceding a renovation (where the first lead and last lag capture years more than 4 years after and more than 6 years before the renovation). Standard errors are clustered by CBSA. Dependent variable is the internal rate of return. Change in employment measured in year-over-year change in logs, and IRR is the internal rate of return for the property over its hold period. Independent variables are created by interacting a binary indicator for each quality segment with the continuous variables for employment growth, new rental housing units as a share of the existing rental stock, and the share of new rental housing entering above the 75th percentile of the local rent distribution. Middle panel reports the difference between the low quality coefficient and high quality coefficient, where columns 2, 4, and 6 show the difference in this difference compared to the baseline specification in column 1. Controls for Treasury yields take the change in the 10-year yield over the hold period, and interact it with the average entry cap rate in the CBSA and quality segment at the beginning of the hold period.