Spatial risk premia in residential real estate

Felipe S. Iachan[1], Leandro Gorno[1], and Ana Paula Ruhe[2]

[1]FGV EPGE
[2]FEA USP

First Version: August 19, 2024
This Version: November 13, 2026

Abstract

Real estate returns differ systematically across locations. Within cities, cheaper or more distant units earn higher returns, and across cities, returns rise with exposure to aggregate risk. We analyze these patterns with a dynamic urban economy model. Rents inherit the risk of local economic conditions, which move with aggregate shocks, and investors price that risk. Because rents fall with distance from the employment center while operating costs fall less, distant, low-rent locations have thinner margins and greater exposure. Expected growth moves return between yields and appreciation, so risk premia can be measured only from total returns. Most housing data report one component but not both. We assess the model with NCREIF multifamily data, which report both return components and operating costs. A portfolio long high-beta and short low-beta cities earns 0.39 percentage points per quarter (s.e. 0.14). Within cities, excess returns rise by 2.31 basis points per quarter for each mile from the employment center (s.e. 0.80). Operating costs excluding property taxes rise with rents at an elasticity of 0.80 (s.e. 0.04). A calibration with either a fixed non-tax cost or this elasticity matches the income-return profile across distance quartiles and implies that the higher cap rates of distant locations compensate for risk rather than low expected growth.

Keywords: Residential real estate; Spatial heterogeneity; Risk premia; Urban economics; Real estate pricing.

JEL Classification: G12, G19, R30, R31.

1 Introduction

Housing is a large asset class with average returns comparable to those on equities (Jordà et al., 2019), and those returns differ systematically across locations. Within cities, cheaper housing earns higher rental yields and often higher total returns: single-family rentals in low-price ZIP codes earn higher net yields and higher total returns than those in high-price ZIP codes (Demers and Eisfeldt, 2022).[1] Across cities, yields and appreciation largely offset each other (Demers and Eisfeldt, 2022), and the return differences that remain are associated with how strongly housing returns covary with income growth (Amaral et al., 2025). Several of these studies point to risk, but none derives from the economics of cities which locations should carry higher risk premia.

We provide a risk-based explanation for this pattern, from a model that combines urban economics with macrofinance. Rents are set in spatial equilibrium: because workers move freely, the rent at each location reflects the local wage net of commuting costs and the city's amenities, relative to the utility workers can obtain elsewhere (Alonso, 1964; Muth, 1969; Mills, 1967; Rosen, 1979; Roback, 1982). Wages and that outside option move with aggregate shocks, so rents are risky, and investors price rent streams with a stochastic discount factor that compensates for exposure to aggregate risk. Cities differ in how strongly their wages respond to aggregate shocks. This difference in exposure translates into rents, yielding the prediction that cities whose rents load more on aggregate shocks earn higher expected excess returns.

Operating costs generate a second prediction, about returns within a city. Without them, all locations in a city share the same wage shocks and earn the same expected return. Maintenance and repairs are paid in local labor, so their cost follows the city wage, while gross rents fall with distance from the employment center. Peripheral properties therefore operate on thinner margins, and a common shock to rents moves their net operating income (NOI) by a larger percentage, the operating leverage that Novy-Marx (2011) relates to equity returns. When this amplified exposure carries a positive price of risk, expected returns rise with distance from the center. Because the spread scales with the price of risk, it should also widen when aggregate expected returns are high.

With constant risk premia and no operating costs, the model reduces to a risk-adjusted Gordon formula, and this special case provides lessons about measurement:

$$\mathbb{E}_t[R_{t+1}] \;=\; \underbrace{(r_f + \pi - g)}_{\text{cap rate}} \;+\; \underbrace{g}_{\text{expected appreciation}} \;=\; r_f + \pi,$$

where r_f is the risk-free rate, π is the risk premium, and g is expected rent growth. Risk premia appear

[1] The same pattern appears in other markets and periods. Rent-price ratios fall with property value and neighborhood affluence in London and across England (Bracke, 2015; Clark and Lomax, 2020), gross yields were persistently higher in the poorer neighborhoods of nineteenth- and twentieth-century Paris and Amsterdam (Eichholtz et al., 2021), and low-rent housing segments earn higher returns (Damen et al., 2026).

in yields. Expected growth lowers the yield and raises appreciation by the same amount, leaving the total return unchanged. A high income yield can therefore signal a risky location or a slow-growing one. Campbell et al. (2009) and Himmelberg et al. (2005) apply this logic to rent-price ratios over time. Plazzi et al. (2010) find that cap rates predict returns for apartments, retail, and industrial properties, but provide only limited evidence of rent-growth predictability, concentrated in office properties and in metropolitan areas with stringent land-use restrictions. Across locations, most housing data come from repeat-sales price indices, which measure appreciation well but do not capture yields, so on their own they cannot recover the risk premium.

Returns on institutional multifamily properties from the National Council of Real Estate Investment Fiduciaries (NCREIF) measure what the model needs. NCREIF computes both return components from property accounts: income return is given by NOI, and capital return is the change in appraised value, both in relation to the capital invested in the property. Its expense accounts report operating costs by category, so we can measure how costs scale with rents across locations. Returns are aggregated by core-based statistical area (CBSA) and by ZIP code, the two levels of the model's predictions.

Across cities, average returns are high and their composition follows the Gordon pattern. The mean excess total return is 1.71 percentage points (pp) per quarter, about 7.0% annualized, and income return supplies 1.36 pp of the 2.10 pp mean total return. Across CBSAs, a 1 pp higher average excess income return comes with a 1.06 pp lower average capital return (s.e. 0.30), the offset the Gordon case predicts when expected growth differs across cities. A return this high is consistent with a risk premium.

Cities more exposed to aggregate apartment-market risk earn higher returns. We sort cities by their beta on the national NCREIF apartment index. A portfolio long the highest tercile and short the lowest earns 0.39 pp per quarter (s.e. 0.14), about 1.6 percentage points a year, and Fama–MacBeth regressions (Fama and MacBeth, 1973; Newey and West, 1987) on portfolios and on individual cities give the same sign.

Within cities, excess returns rise with distance from the primary employment center. Fama–MacBeth regressions of ZIP-code returns on distance give a gradient of 2.31 basis points (bp) per mile per quarter (s.e. 0.80), and the farthest distance quartile earns about 60 bp per quarter more than the closest. In the expense accounts, non-tax operating costs rise with rents at an elasticity of 0.80 (s.e. 0.04), so margins are thinner where rents are lower. A calibration with either a fixed non-tax cost or this elasticity matches the income-return profile across distance quartiles and implies negative risk-adjusted growth of rents relative to wages, the condition under which cap rates rise with distance. Non-tax operating leverage interacted with the consumption-wealth ratio (Lettau and Ludvigson, 2001a) also predicts ZIP-code returns (coefficient 1.140, s.e. 0.442), a secondary, time-varying

restriction of the model.

Our main innovation is to derive the spatial pattern of housing risk premia from urban fundamentals and to measure it with both return components. The model prices any rental cash flow, so it applies to single-family rentals, to owner-occupied housing through imputed rents, and to other rental real estate. Unlike the total-return studies that document the pattern in these segments (Demers and Eisfeldt, 2022; Amaral et al., 2025; Eichholtz et al., 2021; Damen et al., 2026), we derive which locations should earn higher premia and test that restriction. The model also gives an urban foundation to tests of whether housing risk is priced, and with what sign (Crone and Voith, 1999; Cannon et al., 2006; Case et al., 2011; Eiling et al., 2019; Sinai and Souleles, 2005; Han, 2013). Macro-finance models of housing study how risk premia, credit, and collateral move house prices over time and across quality segments (Piazzesi and Schneider, 2016; Favilukis et al., 2017; Landvoigt et al., 2015), and dynamic spatial models explain the dispersion of house prices across cities (Van Nieuwerburgh and Weill, 2010). We join the urban rent model to a stochastic discount factor to study expected returns across and within cities.

Four limits bound the interpretation. NCREIF covers institutional multifamily properties, and its returns rest on appraisals, which smooth measured returns. The results need not extend to other segments, though the studies cited above find similar spatial patterns there. The cross-city tests rest on 48 metropolitan areas. Illiquidity may confound both the level and the spatial pattern of returns. Smaller housing markets are less liquid (Amaral et al., 2025), and illiquidity shapes the risk and return of individual properties (Sagi, 2021; Giacoletti, 2021). We cannot separate this channel from risk, so part of the measured premium may compensate for illiquidity. Finally, operating leverage is our interpretation of the distance gradient. The expense accounts show that properties are levered, and the income, total-return and time-series evidence is consistent with the model. Leverage is highest where rents are low relative to wages, however, and the neighborhoods and buildings in those locations differ in other ways. We have no variation in leverage that holds these differences fixed, so we cannot separate its effect from theirs.

Section 2 describes the data. Section 3 presents the cross-city and within-city evidence. Section 4 develops the model, Section 5 reports the quantitative assessment, and Section 6 concludes.

2 Data and Measurement

This section describes the data and spatial measures used in the empirical analysis. The central measurement requirement is to observe total returns together with their income and capital components. We then construct a common within-metropolitan distance measure that maps each property to its location relative to the primary employment center of its metropolitan area.

2.1 NCREIF data construction

The empirical analysis uses data from the National Council of Real Estate Investment Fiduciaries (NCREIF) Research database. NCREIF aggregates property- and fund-level performance data submitted by institutional real estate investment managers (its Data Contributor Members), primarily reflecting privately held U.S. commercial properties. Affiliated managers report quarterly appraisals, income, and expenses for each property, from which NCREIF computes income returns (net operating income divided by a common adjusted market-value denominator) and capital returns (the change in market/appraised value, net of capital expenditures and including partial sales, divided by that same denominator) at the property level.

We restrict the sample to multifamily apartment properties. This choice has two motivations. First, apartments are the closest analogue to residential housing available in institutional data. Second, restricting to a single property type avoids composition effects across retail, office, and industrial sectors, whose operating cost structures and lease terms differ substantially. Single-family housing is not represented in the NCREIF database, so our results apply directly to the multifamily rental market.

NCREIF's advantage over price-only datasets is that it separately reports the income and capital components of total return. Both components use a time-weighted denominator that adjusts beginning-of-quarter market value for capital expenditures, partial sales, and the timing of NOI. Income return equals NOI divided by this denominator, and capital return measures the change in appraised value, net of capital expenditures and including partial sales, divided by the same denominator. We estimate risk exposure from total returns, because realized appreciation reflects both expected growth and realized shocks, which need not average out over our sample.

Three measurement issues influence the interpretation. First, we use the standard appraisal-based NCREIF return series. NCREIF also contains transaction-based sale information, but those observations are sparser and are not the baseline return object. Appraisal-based returns provide repeated income and capital-return measurement for the geographic aggregates in our extract, while smoothing affects return second moments (Geltner, 1991; Couts et al., 2024). We therefore use the NCREIF appraisal-based returns as the baseline and report one-step AR(1) unsmoothing in Appendix A.8. The within-metro distance gradient remains positive after unsmoothing. Second, the NCREIF sample is institutional: properties are owned through investment fiduciaries and professionally managed (Sagi, 2021). The results therefore apply most directly to institutional multifamily residential assets, even though the model is more general. Third, NCREIF applies a confidentiality filter, requiring at least three properties from at least three managers per geographic-quarter cell. This filter limits coverage in early years and in smaller markets. These restrictions make NCREIF representative of one segment of the residential market, institutional multifamily, rather than of housing as a whole. Within that segment, its direct measurement of income, appreciation, and expenses is well suited to studying risk

premia.

Sample construction Our extracts come from the NCREIF Research Database, which reports quarterly returns aggregated by NCREIF to the Core Based Statistical Area (CBSA) and ZIP-code levels. We apply the following filters. We restrict the sample to 2000 Q1 onward, when NCREIF residential coverage becomes deeper and the panel aligns with the Decennial Census 2000 controls used below. For the cross-city analysis, we require each CBSA to have at least 40 quarters of data, yielding 48 CBSAs. For the within-city descriptive analysis, we require each CBSA to have at least 100 ZIP-quarter observations across at least 4 distinct ZIP codes, yielding 22 CBSAs and 312 ZIP codes. Returns are winsorized at the 5th and 95th percentiles of the pooled panel distribution to limit the influence of appraisal or other anomalies. Section 5.1 also uses the NCREIF expense accounts, which report rental income and expenses per square foot for each property, and NCREIF's metropolitan indices of rent and NOI growth. Table 7 summarizes both.

Excess returns We define excess returns as total returns minus the quarterly risk-free return. We construct the quarterly series from the monthly one-month Treasury-bill returns in Kenneth French's data library (French, 2026).

2.2 Distance definition

Within-metro analyses require a common measure of each ZIP code's location relative to the metropolitan employment core. Our baseline measure is distance to a data-driven primary employment center constructed from 2022 LEHD-LODES workplace employment. We aggregate jobs to ZIP Code Tabulation Areas, identify employment-density clusters using a modified Giuliano-Small procedure (Giuliano and Small, 1991), and define the primary center as the highest-employment cluster in the CBSA. Distance is measured from the ZIP centroid to that center. This reproducible measure closely matches distance to a hand-curated downtown ZIP code in each metropolitan area: among ZIPs covered by both definitions, the correlation between the two distances is 0.97. Appendix A.1 gives the construction details.

To accommodate polycentric cities, we construct an alternative nearest-center measure. The same algorithm identifies secondary employment centers in metros such as Dallas–Fort Worth, Miami–Fort Lauderdale, San Francisco–Oakland, Seattle–Tacoma, and Tampa–St. Petersburg. For this alternative measure, each ZIP is assigned its distance to the nearest identified employment center, not necessarily the primary one. We use primary-center distance in the main analysis. The nearest-center result remains positive and is somewhat steeper per mile. Appendix A.1 reports it on the common sample used to compare all three distance measures.

3 Empirical Facts

This section presents the paper's main empirical facts. We begin with cross-city return variation and then turn to within-city variation. The final subsection checks the distance gradient against neighborhood quality, appraisal smoothing and the pandemic.

3.1 Cross-city returns and risk pricing

Housing has earned high average returns, comparable to those on equities over long samples and across many countries (Jordà et al., 2019). Our sample shows the same. The average CBSA earned 1.71 pp per quarter over the Treasury bill rate (Table 1), against 1.64 pp for the U.S. stock market.[2] Returns also differ across cities, and part of the difference appears to compensate for risk. Case et al. (2011) find that house-price returns across U.S. metropolitan areas rise with their beta on the national index. Amaral et al. (2025) find that housing in large cities earns lower returns, and they attribute the gap to lower covariance with local income and greater liquidity. We start from average returns and their components, then sort cities on beta, and end with Fama–MacBeth regressions.

Table 1 reports quarterly returns pooled over CBSAs and quarters, winsorized at 5 percent in each tail. Income returns supply most of the average return, 1.36 of the 2.10 pp per quarter. The standard deviations in the table are computed over the whole panel, 2.21 pp for capital returns and 0.31 pp for income returns, so they combine movements over time with differences across cities.

The two sources of dispersion mean different things. Movements over time are risk, and persistent differences in average returns across cities are the heterogeneity that risk premia should explain. We measure risk by the standard deviation over time of the cross-city average return, and heterogeneity by the standard deviation across cities of each city's average return. Capital returns carry nearly all the risk, 1.74 pp against 0.27 pp for income returns, but their averages differ across cities only about twice as much, 0.35 pp against 0.16 pp. Persistent differences across cities appear in both components.

Table 1: Descriptive Statistics, NCREIF at CBSA aggregation

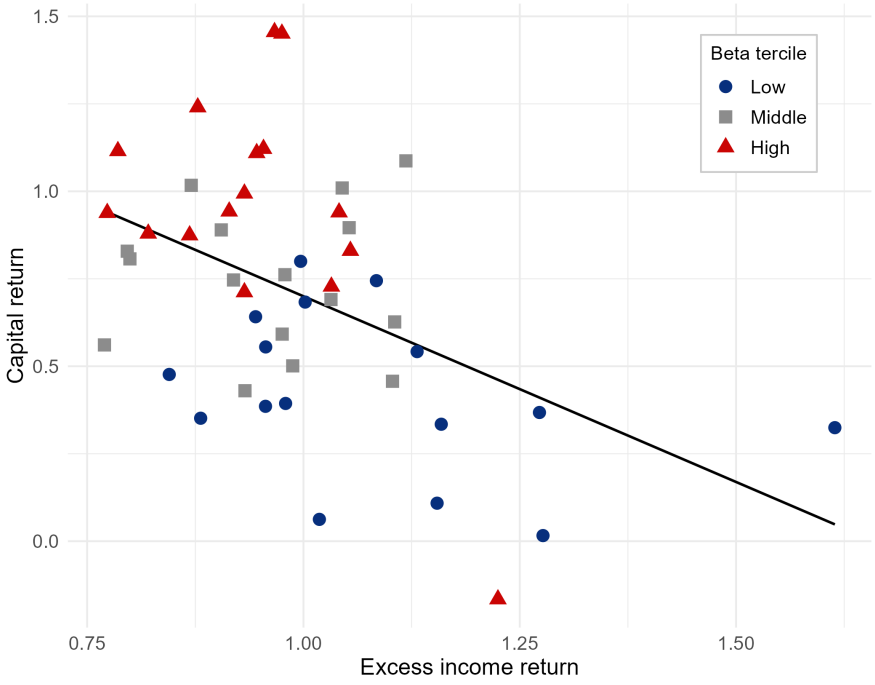
Variable	mean	sd	median	min	max
Total Return	2.0957	2.2794	2.1000	-2.9400	6.5125
Income Return	1.3589	0.3120	1.3100	0.9400	2.3000
Capital Return	0.7353	2.2138	0.7100	-4.2000	5.0350
Excess Return	1.7098	2.3626	1.6900	-4.2700	6.5125

Note: values are quarterly returns in percentage points.

[2] Fama–French market excess return, compounded within quarters, 2000Q1–2023Q3 (French, 2026).

Return components and risk Figure 1 plots each city's average capital return against its average excess income return. In the Gordon benchmark, which holds risk premia and expected growth constant in a CBSA over time, faster expected growth raises appreciation and lowers the income return one-for-one, holding the risk premium fixed. Across the 48 CBSAs, the slope of average capital returns on average excess income returns is −1.06 (s.e. 0.30), close to the one-for-one offset. Demers and Eisfeldt (2022) find a similar negative relation between yields and appreciation for single-family rentals across 30 cities.

Figure 1: Income and capital returns across CBSAs



	Fitted line
Slope	−1.06
	(0.30)
Intercept	1.76
	(0.30)
R^2	0.22
Residual s.d.	0.31
CBSAs	48

Note: Each point is one CBSA's average quarterly excess income return and capital return, in percentage points, over the quarters it is observed between 2000Q1 and 2023Q4. Returns are winsorized at 5 percent in each tail. The line is the least-squares fit of capital returns on excess income returns across the 48 CBSAs, weighted equally. The table reports its slope and intercept, with standard errors in parentheses, and the residual standard deviation in percentage points. Colors and markers give each CBSA's tercile of full-sample beta on the national NCREIF apartment index, the portfolios of Table 2.

Differences in expected growth move cities along a line of equal total return, while a higher risk premium raises that line. Because growth and risk exposure may be correlated across cities, the slope alone cannot separate the two.[3] In Figure 1, higher-beta cities tend to lie above the fitted line, with higher total returns. We next compare total returns across beta terciles.

[3]In the Gordon benchmark a city's income return is $r_f + \beta\lambda - g$ and its expected appreciation is g. Across cities, the slope of capital returns on income returns is then -1 when only g varies, and it turns positive when riskier cities also grow faster, $\text{cov}(\beta\lambda, g) > \text{var}(g)$. Beta does covary with pre-sample house-value growth and supply constraints, so the slope does not hold expected growth fixed.

Risk premia Average excess total returns rise across the full-sample beta terciles of Figure 1, from 1.49 pp per quarter in the low-beta tercile to 1.70 and 1.89 pp in the middle and high terciles (Panel A of Table 2). The high-beta tercile earns 0.39 pp per quarter more than the low-beta tercile (s.e. 0.14), or about 1.6 percentage points a year.

Fama–MacBeth regressions on these portfolios express the same comparison as a price of risk (Fama and MacBeth, 1973). In the first stage, we regress each tercile's quarterly excess return (the average of its CBSAs' excess returns) on the excess return of the national NCREIF apartment index, which gives post-ranking betas of 0.59, 0.70 and 0.77. In the second stage, for each quarter, we regress the three portfolio returns on these betas and average the slopes. The price of risk is 2.15 pp per quarter per unit of beta (s.e. 0.76), with an intercept of 0.22 (s.e. 0.40) (Panel B of Table 2). With three portfolios, this slope restates the high-minus-low spread per unit of beta. Finer portfolios have noisier betas, and the slope falls as the groups get finer (Appendix A.3).

Regressions on individual cities give the same sign with less precision. With unadjusted betas, the price of risk is 0.66 (s.e. 0.53). Each city's beta is estimated from at most 96 quarters of appraisal-based returns, and errors in these betas pull the slope toward zero. Shrinking each beta toward the cross-sectional mean in proportion to its sampling variance (Vasicek, 1973) gives a price of risk of 1.70 per unit of beta (s.e. 0.74).[4] Betas of MSA-level Zillow appreciation on national Zillow appreciation, a measure of exposure from a separate data source, give 0.42 (s.e. 0.18). [5] Because the three beta measures have different scales, the last row of Panel B reports the effect of a one-standard-deviation difference in beta: 0.09, 0.16 and 0.18 pp per quarter.

Across cities, higher exposure to national housing risk goes with higher average returns. The relation has the same sign in the beta sort and under every beta measure, and it is large in annual terms. The high-beta tercile earns about 1.6 percentage points a year more than the low-beta tercile, and a one-standard-deviation difference in shrunk or Zillow beta goes with 0.6 to 0.7 percentage points a year.[6]

[4]Baton Rouge has the least precise beta, estimated from 46 quarters, and a Vasicek weight of 0.29 on its unadjusted estimate. Excluding it gives a one-standard-deviation effect of 0.18 percentage points per quarter under unadjusted and shrunk betas.

[5]Zillow reports appreciation rather than total returns and is aggregated at the MSA level. The beta is a loading of monthly metropolitan price appreciation on national appreciation, built from the Zillow Home Value Index and attached to the same 48 cities. Appendix A.3 reports the construction.

[6]Without winsorization, the spread is 0.27 pp per quarter (s.e. 0.26) and the effect of a one-standard-deviation difference in shrunk beta is 0.14 pp (s.e. 0.12), against 0.16 (s.e. 0.07) in the baseline. Further results are available upon request.

Table 2: Excess returns and the price of beta risk across cities

Panel A: Beta-sorted portfolios summary				
	Low beta	Middle	High beta	High − low
Excess total return (pp)	1.49	1.70	1.89	0.39
				(0.14)
Post-ranking beta	0.59	0.70	0.77	0.18
CBSAs	16	16	16	
Panel B: Fama–MacBeth regressions				
	Portfolios	CBSAs		
	Tercile beta	Unadjusted beta	Shrunk beta	Zillow beta
Price of risk	2.15	0.66	1.70	0.42
	(0.76)	(0.53)	(0.74)	(0.18)
Intercept	0.22	1.27	0.53	1.23
	(0.40)	(0.29)	(0.35)	(0.34)
One-s.d. effect		0.09	0.16	0.18
		(0.07)	(0.07)	(0.08)
Test assets	3	48	48	48

Notes: CBSAs are sorted into terciles of their full-sample beta on the national NCREIF apartment index. Shrunk betas give the same terciles. Panel A reports time-series means of each tercile's equal-weighted excess total return, in percentage points per quarter, over 96 quarters from 2000Q1 to 2023Q4, and post-ranking betas from regressions of each tercile's excess return on the excess return of the national index. Panel B reports time-series means of quarterly cross-sectional regressions of excess returns on betas, with an intercept. The first column uses the three tercile portfolios and their post-ranking betas. The other columns use the 48 CBSAs with their own full-sample betas on the national index (unadjusted), the same betas shrunk toward the cross-sectional mean in proportion to their sampling variance (Vasicek, 1973) (shrunk), and betas of MSA-level Zillow appreciation on national Zillow appreciation (Zillow). The one-s.d. effect multiplies the price of risk by the cross-sectional standard deviation of the beta measure. Returns are winsorized at the 5th and 95th percentiles. Parentheses contain Newey–West standard errors with four lags. Conventional standard errors are 0.08 for the spread and 0.44, 0.33, 0.45 and 0.10 for the prices of risk. Appendix A.3 reports finer portfolios, winsorized betas, and estimates without Baton Rouge.

3.2 Within-city returns and the distance gradient

Within cities, cheaper locations earn higher rental yields. Demers and Eisfeldt (2022) document this across ZIP codes for single-family rentals, and Halket et al. (2023) find that lower-income ZIP codes earn higher yields and higher total returns within 21 metropolitan areas.[7] In the Gordon benchmark, a higher yield can reflect either more risk or slower expected growth. We first show that NCREIF reproduces the yield gradient across distance from the primary employment center, and then that the gradient carries over to excess total returns, whose sample average estimates the risk premium.

Summary statistics The within-city sample has 312 ZIP codes in 22 CBSAs. In the regressions below, we further restrict them to the 296 ZIP codes within 30 miles of their primary employment center. ZIP-level returns are more volatile than the CBSA aggregates for the same metropolitan areas, with a standard deviation of quarterly total returns of 2.79 pp against 2.40 pp, mostly through capital returns (Appendix A.2).

Distance gradient of returns Table 3 reports returns by quartile of distance from the primary employment center. Panel A reproduces the yield gradient: mean excess income return rises from 0.72 pp per quarter in the closest quartile (Q1) to 0.89 pp in the farthest (Q4), a spread of about 17 basis points. Panel B shows a larger gradient in excess total returns, which rise from 0.94 pp to 1.55 pp, a spread of about 60 basis points per quarter, or 2.4 percentage points annualized. In these data, the income return and capital return gradients go mostly in the same direction and do not offset each other in the total return. Both gradients rise across all four quartiles.

The ZIP panel is unbalanced, because the set of reporting ZIP codes varies across metropolitan areas and quarters. If distant ZIP codes report disproportionately in high-return metro-quarters, quartile means can rise with distance even when returns within a metro-quarter do not. Table 4 removes this composition effect by subtracting CBSA$\times$quarter means. Both gradients persist and now rise across all four quartiles: residualized excess income returns go from -0.05 pp in Q1 to $+0.06$ pp in Q4, and residualized excess total returns from -0.19 pp to $+0.21$ pp.

As a continuous gradient, a Fama–MacBeth regression of excess total returns on distance across the 296 ZIP codes gives 2.31 basis points per mile per quarter (s.e. 0.80).[8]

[7] Rental yields rise with distance from the center in German cities (Ahlfeldt et al., 2023) and in U.S. metropolitan areas (Gupta et al., 2022), and they fall with property value and neighborhood affluence in London and across England (Bracke, 2015; Clark and Lomax, 2020). Gross yields were persistently higher in the poorer neighborhoods of nineteenth- and twentieth-century Paris and Amsterdam (Eichholtz et al., 2021), landlords earn higher profits in poor U.S. neighborhoods (Desmond and Wilmers, 2019), and low-rent housing segments earn higher returns (Damen et al., 2026).

[8] The gradient is 1.69 basis points per mile per quarter (s.e. 0.75) when each CBSA has its own intercept in every quarter, so that only comparisons within a metropolitan area remain (Table A5 in Appendix A.5). The baseline regression has one

Table 3: Descriptive Statistics, NCREIF at ZIP code aggregation

Distance group	mean	sd	median	min	max	annualized mean
Panel A: Excess Income Return						
1	0.7231	0.4560	0.770	-0.600	1.670	2.9239
2	0.7676	0.4701	0.830	-0.600	1.670	3.1058
3	0.8522	0.5143	0.950	-0.600	1.670	3.4526
4	0.8943	0.5144	0.995	-0.540	1.670	3.6254
Over 30 miles	0.7915	0.5428	0.910	-0.470	1.660	3.2040
Panel B: Excess Total Return						
1	0.9407	2.8432	0.950	-5.477	7.550	3.8161
2	1.3066	2.8839	1.240	-5.477	7.550	5.3297
3	1.4729	3.0289	1.380	-5.477	7.550	6.0231
4	1.5522	2.9731	1.515	-5.477	7.550	6.3547
Over 30 miles	1.4197	2.9953	1.385	-5.377	7.550	5.8010

Note: values are quarterly returns in percentage points (except the last column), winsorized at 5% in both tails. Distance quartiles are pooled across the ZIP codes within 30 miles of their primary employment center, with each ZIP code assigned once; the cut points are 2.7, 6.4 and 13.2 miles.

Distance and estimated betas Controlling for estimated betas leaves the distance gradient unchanged. We estimate ZIP-level betas for the 66 ZIP codes in 18 CBSAs with at least 40 quarters of returns.[9] The national beta is a ZIP code's loading on the national NCREIF apartment index, and the local beta is its loading on the apartment index of its own CBSA. Their means are 0.80 and 0.96 (Table 5).[10] Each quarter we regress ZIP-level excess total returns on distance and the estimated betas,

$$R^e_{i,t} = a_t + \lambda^N_t \beta^N_i + \lambda^C_t \beta^C_i + \delta_t \, d_i + \epsilon_{i,t}, \tag{1}$$

where d_i is the distance in miles from ZIP code i to its primary employment center, and we average the quarterly slopes. In the 66-ZIP sample, distance alone gives 4.06 basis points per mile (s.e. 1.07), and adding the national beta, the local beta or both gives 4.06 to 4.09 (Table 6).

Distance is also measured more precisely than estimated betas. Betas are generated regressors, estimated from appraisal-based return series of at most 96 quarters, and the national and local betas

intercept per quarter and also compares ZIP codes across metropolitan areas.

[9]Boston, Orlando, Philadelphia and San Francisco are in the descriptive sample but drop out of the pricing sample because their ZIP codes lack 40 quarters of returns.

[10]Mean betas below unity reflect the fact that the national NCREIF apartment index is value-weighted across metropolitan apartment markets, not the equal-weighted average of the ZIP codes in our sample. Multifamily properties in these metros need not have unit exposure to the national index.

Table 4: Descriptive Statistics, Residuals of FE regressions

Distance group	mean	sd	median	min	max	annualized mean
Panel A: Residual Income Return						
1	-0.0535	0.1286	-0.0512	-0.7310	0.7144	-0.2136
2	-0.0096	0.1254	0.0000	-0.5775	0.4767	-0.0383
3	0.0187	0.1303	0.0120	-0.5914	0.6314	0.0746
4	0.0643	0.1236	0.0595	-0.5650	0.5419	0.2576
Over 30 miles	0.0335	0.1636	0.0346	-0.5500	0.7500	0.1340
Panel B: Residual Total Return						
1	-0.1905	1.3844	-0.1787	-5.3586	6.1165	-0.7597
2	0.0072	1.4776	0.0000	-9.5565	8.4062	0.0289
3	0.0438	1.5345	0.0000	-7.4663	7.1797	0.1752
4	0.2087	1.5430	0.1015	-7.8725	5.8481	0.8373
Over 30 miles	0.0949	1.6747	0.0000	-6.0853	5.8087	0.3801

Note: values are residuals of regressions of quarterly returns in percentage points against CBSA-time fixed effects. Distance quartiles as in Table 3.

Table 5: Summary Statistics of Beta Measures

Measure	mean	sd	median	min	max
National beta (alone)	0.80277	0.18678	0.81527	0.11133	1.21221
Local beta (alone)	0.95973	0.13206	0.97559	0.37074	1.20775
National beta (joint)	0.03061	0.44420	0.12652	-1.77334	0.86798
Local beta (joint)	0.93143	0.43286	0.86230	0.26222	2.46063

Notes: National beta is the ZIP loading on the national apartment index. Local beta is the ZIP loading on the ZIP's own CBSA apartment return. Joint rows come from regressions including both beta measures.

are nearly collinear when estimated jointly. Distance is observed directly, and under the mechanism we propose in Section 4.2 it proxies for exposure, because it moves with the rent gradient that governs operating leverage. Within CBSAs, the slope of estimated local betas on distance from the primary center is positive but imprecise: 0.045 per 10 miles (s.e. 0.022), with a wild bootstrap p-value of 0.13 (Appendix A.6). Demers and Eisfeldt (2022) find a related pattern for single-family rentals. ZIP-level appreciation loads more strongly on city appreciation in lower price tiers, and they read the higher returns of those tiers as possible compensation for risk. Across ZIP codes within U.S. metropolitan

Table 6: Within-city distance gradients

	Broad sample	At least 40 quarters of returns			
		Distance only	National beta	Local beta	Both betas
Distance (bp/mile/qtr)	2.31	4.06	4.09	4.09	4.06
	(0.80)	(1.07)	(1.05)	(1.07)	(1.02)
ZIPs / CBSAs	296 / 22	66 / 18	66 / 18	66 / 18	66 / 18
Quarterly cross-sections	96	94	94	94	94

Notes: Fama–MacBeth averages of quarterly distance slopes, with Newey–West standard errors with four lags in parentheses. Conventional standard errors are 0.61, 0.75, 0.74, 0.75 and 0.73. Every quarterly cross-section includes an intercept, whose estimates are not reported. The broad sample has all ZIP codes within 30 miles of the primary employment center; the other columns keep those with at least 40 quarters of returns, long enough to estimate ZIP-level betas, and the last three add the indicated betas as controls. In the last column, both betas come from a single first-stage regression of the ZIP return on the national and CBSA apartment indices. Returns are winsorized at 5/95 after the CBSA coverage screen. Distance is measured in miles to the primary employment center (LEHD-LODES). Appendix A.4 reports the beta coefficients and distance-group specifications, and Appendix A.5 adds a separate intercept for each CBSA and quarter.

areas, Eiling et al. (2019) find that exposure to the local housing market is priced in a sizable share of metros. The sign agrees with the mechanism, but estimated betas are too noisy to measure the gradient.

3.3 Robustness analysis

The distance gradient is robust to correcting for appraisal smoothing and to excluding the pandemic years. Census 2000 neighborhood controls reduce it by about half, but do not eliminate it.

Urban observables: Census 2000 controls We add predetermined Census 2000 characteristics to the beta-controlled regression for the 65 ZIP codes with complete data.[11] The estimates average the 74 quarters in which at least 16 of these ZIP codes are available (Appendix A.7). Median income, college share and population density absorb about two-fifths of the gradient, which falls from 4.32 to 2.52 bp/mile (s.e. 0.82). Adding the non-Hispanic White share and the poverty rate leaves it at 2.72 bp/mile (s.e. 0.88).

Neighborhood characteristics that vary with distance can absorb part of the gradient for two reasons. They can shift rents relative to operating costs, the channel Section 4.2 develops, or they can carry premia of their own, such as compensation for illiquidity or tenant credit risk. The Census regressions cannot separate the two. Separating them would require variation in neighborhood

[11] We lose one ZIP code from missing Census data.

characteristics that our data do not provide.

Measurement robustness: appraisal smoothing Desmoothing leaves the gradient positive. NCREIF returns rest on periodic appraisals, which induce positive autocorrelation and attenuate volatility and covariances. We apply the one-step AR(1) unsmoothing filter used in the real estate literature (Geltner, 1991, 1993; Couts et al., 2024), with a pooled smoothing parameter $\hat{\alpha} = 1 - \hat{\rho} = 0.484$ (the median ZIP-specific $\hat{\alpha}$ is 0.436). The desmoothed gradient is 2.90 bp/mile (s.e. 0.88), against 2.31 (s.e. 0.80) in the baseline (Appendix A.8).

Pre-2020 robustness: pandemic amplification The gradient predates the pandemic. After 2020, prices and rents rose in the suburbs relative to city centers (Gupta et al., 2022), so part of the gradient could reflect a one-time revaluation rather than a persistent risk premium. Excluding all quarters from 2020Q1 onward, the gradient falls from 2.31 to 1.62 bp/mile (s.e. 0.78). Appendix A.9 reports the full pre-2020 tables.

4 Model

In this section, we introduce a framework to interpret the empirical patterns previously discussed. Our model builds on classical urban economics but introduces three features needed to study asset pricing: (1) wages and outside options evolve stochastically over time, so future rents are uncertain; (2) investors are risk-averse and require compensation for bearing risk; and (3) risk exposure varies across locations, which restricts the cross-section of expected returns.

4.1 The baseline model

Time is continuous, indexed by $t \geq 0$, and extends over an infinite horizon. There are K metropolitan areas, or cities, each indexed by k. These cities are assumed to be monocentric and circular, with locations within each city k indexed by their distance $d \geq 0$ from the central business district (CBD). We also include an outside option location k_0 representing locations outside these metropolitan areas.

Agents The model incorporates two distinct types of agents: workers and real estate investors. Workers consume their entire income without saving (hand-to-mouth behavior). They are highly mobile, moving freely across locations to maximize their instantaneous welfare. Real estate investors are risk-averse, are not wealth-constrained, and are the marginal investors in real estate. To explore the

effects of limited financial integration, we introduce the possibility of restricting real estate investors to investing only within their own city k.[12]

Uncertainty The uncertainty in the model is driven by independent stochastic processes. We assume a two-shock structure that separates aggregate risk, common across cities, from city-specific risk. The process $\{B_t\}$ represents aggregate shocks. Each city k is additionally exposed to idiosyncratic shocks, encoded in the process $\{B_t^k\}$. These stochastic processes jointly determine local economic conditions, including wages and equilibrium rents.

For analytical tractability, we assume that $\{B_t\}$ and $\{B_t^k\}$ follow independent standard Brownian motions. This assumption implies that, for any time interval of length $s > 0$, the changes in the aggregate process $B_{t+s} - B_t$ and the local process $B_{t+s}^k - B_t^k$ are uncorrelated normal random variables with mean 0 and variance s.

Worker behavior and commuting Workers in the model derive utility from private consumption, denoted as c, the size (or quality) of their dwellings, s, and the amenities available in their city, a. Their utility function is given by

$$U(c, s, a) = c^{\alpha} s^{1-\alpha} a,$$

where $\alpha \in (0, 1)$ is a parameter describing the relative importance of private consumption. Workers' decisions are constrained by their budget, which determines the total consumption of a worker who lives in city k, at distance d from the CBD, at time t, as

$$c_{k,d,t} = w_{k,t} - r_{k,d,t} s - T_k(d).$$

Here, $w_{k,t}$ is the wage earned by workers in city k at time t, $r_{k,d,t}$ is the rent per unit of housing at distance d from the CBD in city k at time t, and $T_k(d)$ represents the commuting cost from location d to the CBD. For tractability, we assume that the commuting cost is linear in wages and distance:

$$T_k(d) = \tau_k \, w_{k,t} \, d,$$

where τ_k represents the opportunity cost of time for workers in city k. Substituting into the budget constraint, we obtain

$$c + r_{k,d,t} s = \kappa_{k,d} \, w_{k,t},$$

where $\kappa_{k,d} = \max\{1 - \tau_k d, 0\}$ represents the effective wage after accounting for commuting costs.

[12] See Levy (2021) for evidence of limited financial integration in residential rental markets in France.

Stochastic Processes Let Ψ_t denote the utility derived by workers in the outside option location k_0. We assume it follows a geometric Brownian motion given by

$$\frac{d\Psi_t}{\Psi_t} = \mu_\Psi dt + \sigma_\Psi dB_t, \tag{2}$$

where μ_Ψ and σ_Ψ are the drift and volatility coefficients, and dB_t is the aggregate shock. The baseline treats Ψ_t as exogenous. The Online Appendix studies a closed system of cities, in which the outside option is an endogenous index of city wages.

Similarly, wages in each city k follow a stochastic process described by

$$\frac{dw_{k,t}}{w_{k,t}} = \mu_{w,k} dt + \sigma_{w,k} dB_t + \varsigma_{w,k} dB_t^k, \tag{3}$$

where $\mu_{w,k}$ is a drift coefficient specific to city k, $\sigma_{w,k}$ is a volatility coefficient which measures the exposure to aggregate risk, and $\varsigma_{w,k}$ captures the risk specific to city k brought about by the idiosyncratic shock dB_t^k.

Finally, while we allow amenity endowments to vary across cities, we assume that they are constant over time: $a_{k,t} = a_k$.[13] This assumption simplifies the model by focusing the analysis on the interaction between wages, rents, and commuting costs.

Equilibrium Rent Determination Equilibrium rent is determined by wages, commuting costs, and worker mobility. Workers are mobile within cities, equalizing utility across all occupied locations. For location d in city k at time t, this spatial equilibrium condition is:

$$r_{k,d,t} = \zeta \left(\kappa_{k,d} w_{k,t}\right)^{\frac{1}{1-\alpha}} a_k^{\frac{1}{1-\alpha}} \Psi_t^{-\frac{1}{1-\alpha}},$$

where ζ is a constant, $\kappa_{k,d}$ adjusts the wage for commuting costs, a_k represents the amenity level in city k, and Ψ_t is the outside option. Taking the natural logarithm of both sides, the rent equation can be rewritten as:

$$\ln r_{k,d,t} = \ln \zeta + \frac{1}{1-\alpha} \ln \left(\kappa_{k,d} w_{k,t}\right) + \frac{1}{1-\alpha} \ln a_k - \frac{1}{1-\alpha} \ln \Psi_t.$$

Rent is positively related to wages and amenities, and negatively related to the outside option. The term $\frac{1}{1-\alpha}$ governs the sensitivity of equilibrium rent to changes in wages and the outside option: because housing's expenditure share is $(1-\alpha)$, a 1% increase in the wage–outside-option ratio raises rent by $\frac{1}{1-\alpha}$%. The constant expenditure share that underlies this link is an empirical regularity: Davis

[13] Allowing amenities to vary stochastically would introduce an additional source of rent risk. Empirically, identifying time-varying amenity shocks separately from wage or outside-option shocks is difficult because amenity proxies (school quality, crime rates, green space) move slowly and are endogenous to local economic conditions. The pricing results carry over when amenities respond to local income or population density, which changes how strongly rents respond to wage and outside-option shocks. Derivations are available upon request.

and Ortalo-Magné (2011) document that the share of expenditure on housing is essentially constant across U.S. metropolitan areas and over time.

The rent gradient reflects two forces: commuting costs τ_k, and the exponent $1/(1-\alpha)$, which governs how sharply rent responds to the wage net of commuting costs.

Rent dynamics In this model, rent evolves according to the stochastic processes that capture both aggregate and location-specific risks. The dynamics of rent at location d in city k at time t are given by:

$$\frac{dr_{k,d,t}}{r_{k,d,t}} = \mu_{r,k}dt + \sigma_{r,k}dB_t + \varsigma_{r,k}dB_t^k,$$

where $\mu_{r,k}$ is the drift term, $\sigma_{r,k}$ captures aggregate risk, and $\varsigma_{r,k}$ reflects location-specific risk. The volatility terms are given by:

$$\sigma_{r,k} = \frac{1}{1-\alpha}\sigma_{w,k} - \frac{1}{1-\alpha}\sigma_\Psi,$$

$$\varsigma_{r,k} = \frac{1}{1-\alpha}\varsigma_{w,k}.$$

Appendix B.1 derives these terms and the drift $\mu_{r,k}$.

Wage risk passes through to rents: higher wage volatility implies higher rent volatility. Because $\frac{1}{1-\alpha} > 1$, a wage increase generates a proportionally larger increase in rents, all else equal. The outside option Ψ_t, which comoves with aggregate shocks, partially offsets this amplification: when wages and the outside option rise together, the net effect on rent is smaller than the wage shock alone.

Savers and Pricing Savers price assets, including housing, through a stochastic discount factor (SDF). The SDF for savers in city k at time t is

$$\frac{d\Lambda_t^k}{\Lambda_t^k} = -r_f dt - \lambda^k dB_t - \xi^k dB_t^k,$$

where r_f is the risk-free rate, λ^k represents the price of aggregate risk, and ξ^k captures the price of location-specific risks. In the benchmark case, with frictionless markets, we assume $\lambda^k = \lambda$ and $\xi^k = 0$ for all k, implying no location-specific risk premium. Untabulated cross-city tests of a premium for city-specific residual volatility are imprecise and sensitive to specification, and details are available upon request. Section 4.3 relaxes the constant-λ assumption, allowing the price of risk to vary with a macroeconomic state variable.

Using this SDF, the price Q_t of any asset (including a claim on rent) that pays a dividend y_t is determined by the expectation of future discounted cash flows, so

$$0 = E_t\left[d\left(\Lambda_t^k Q_t\right)\right] + \Lambda_t^k y_t dt.$$

Specifically, the price of a rent claim at location d in city k is given by:

$$P_{k,d,t} = \frac{r_{k,d,t}}{r_f + \lambda^k \sigma_{r,k} + \xi^k \varsigma_{r,k} - \mu_{r,k}},$$

where the denominator includes the risk-free rate, aggregate risk premium $\lambda^k \sigma_{r,k}$, potential idiosyncratic risk premium $\xi^k \varsigma_{r,k}$, and expected rent growth $\mu_{r,k}$. This formula illustrates how various risk factors impact the valuation of rent claims, thereby linking asset prices to underlying economic conditions. Appendix B.2 derives it.

Pricing formulas and testable implications In this simple baseline, the pricing formula takes the form of a risk-adjusted Gordon growth model. In a standard Gordon model, the dividend yield equals the discount rate minus the growth rate. Here the discount rate includes a risk premium that compensates investors for bearing aggregate and, potentially, local rent risk.

Rearranging the pricing expression gives a risk-adjusted Gordon growth relation for the cap rate:

$$\frac{r_{k,d,t}}{P_{k,d,t}} = r_f + \underbrace{\left(\lambda^k \sigma_{r,k} + \xi^k \varsigma_{r,k}\right)}_{\pi_k} - \underbrace{\mu_{r,k}}_{g_{r,k}}.$$

That is, cap rates equal the risk-free rate, plus a risk premium π_k, minus expected rent growth $g_{r,k}$. Expected appreciation equals $g_{r,k}$, so the expected total return satisfies

$$\mathbb{E}_t[\mathrm{R}_{k,d,t+1}] = \frac{r_{k,d,t}}{P_{k,d,t}} + \mathbb{E}_t\left[\frac{P_{k,d,t+1} - P_{k,d,t}}{P_{k,d,t}}\right] = r_f + \pi_k.$$

Two implications follow. First, in this baseline model with constant risk premia, risk premia load into cap rates, while expected growth shifts cap rates and expected appreciation in opposite directions. Models with time-varying discount rates, learning about parameters, or state-dependent constraints can generate risk premia through expected appreciation as well, an extension developed in Section 4.3. Second, the cross section of expected returns is governed by the priced exposures $(\sigma_{r,k}, \varsigma_{r,k})$ and their prices (λ^k, ξ^k). This makes the model amenable to standard tests: relate average excess returns to measured exposures (and potentially distinguish aggregate from local risk).

Prediction 1 (Cross-city premium)**.** *Cities whose rents are more exposed to aggregate risk earn higher expected excess returns,* $\pi_k = \lambda^k \sigma_{r,k} + \xi^k \varsigma_{r,k}$.

In the constant-risk-premium benchmark, this risk premium is capitalized into income yields, while Table 2 tests whether measured exposure predicts average excess returns. The volatility terms $\sigma_{r,k}$ and $\varsigma_{r,k}$ in the model map directly to the measured beta coefficients: cities with higher covariance with aggregate shocks exhibit higher betas and, according to the model, should command higher risk premia. The empirical evidence in Section 3.1 is consistent with this prediction, documenting a positive relation between CBSA-level beta estimates and average excess returns.

The baseline model predicts constant returns within each city, in contrast with the NCREIF facts we establish and complementary stylized facts from the literature. Table 3 documents a systematic distance gradient: properties farther from the primary employment center exhibit higher excess income returns and higher total excess returns. This within-city pattern cannot arise from the baseline model, where all locations in city k share identical risk exposures $(\sigma_{r,k}, \varsigma_{r,k})$.

Gross rents decline with distance from the CBD, but the cost of running a building declines less. Properties farther out therefore have thinner margins, and the same percentage change in rents moves their net operating income by more. We now add these costs to the model.

4.2 A model with operating costs

The empirical puzzle The baseline model accounts for why properties in different cities earn different returns: cities whose wages are more exposed to aggregate shocks should earn higher returns. Section 3.1 shows this pattern in NCREIF, and Amaral et al. (2025) find that large U.S. cities, whose housing returns covary less with local income, earn lower total returns. This cross-city heterogeneity arises when locations have different risk exposures $(\sigma_{r,k}, \varsigma_{r,k})$ and different prices of risk (λ^k, ξ^k).

However, the baseline model predicts that all properties *within* the same city should earn identical returns. Properties in city k at different distances from the CBD face the same wage process $w_{k,t}$, the same outside option Ψ_t, and (given our assumption of fixed amenities) the same amenity level a_k. While gross rents vary with distance due to commuting costs, risk exposures remain constant across locations within city k.

The data reject this prediction. Properties in the quartile closest to the CBD earn average excess returns of 0.94 percent per quarter and those in the farthest quartile earn 1.55 percent, about 60 basis points more (Table 3). The gradient survives controls for each ZIP code's exposure to city-wide risk (Section 3.2), and external evidence on single-family rentals (Demers and Eisfeldt, 2022) shows a related pattern: within a city, cheaper ZIP codes earn higher total returns.

Every segment of the housing market, from institutional apartments to owner-occupied single-family homes, carries the costs of operating a property. We add these costs to the model.

Operating costs in practice Property owners pay for maintenance and repairs, property management, utilities not paid by tenants, insurance, and property taxes. The model treats two kinds of costs differently. Maintenance and repairs are labor-intensive, so their cost follows local wages and moves little with a property's rent. The model ties them to the wage through $\delta w_{k,t}$. Property taxes are levied on assessed value and move with the property's price, so the model adds them below as a claim proportional to value.

Assumption 1 (Operating costs)**.** *Non-tax operating costs rise less than proportionally with rents across locations within a city.*

The costs need not be fixed. The model's specification $\delta w_{k,t}$ is the special case in which they are the same across locations. The expense accounts support the assumption: Section 5.1 estimates an elasticity of 0.80 (s.e. 0.04), so a location with 10 percent lower rents has costs only 8 percent lower, and costs absorb a larger share of its income.

Consider two apartment buildings in the same city. A downtown building commands high gross rents due to superior access to employment and amenities. A suburban building, located 15 miles from the CBD, earns lower gross rents due to commuting costs. Yet both require similar maintenance labor: appliance repair, plumbing, landscaping, and tenant management. If local wages set the cost of this labor, both properties face similar operating costs per unit, despite their different gross rent levels.

This creates an *operating leverage* effect. To see this, consider a numerical example. Suppose operating costs in both properties are δ. The downtown property earns gross rent of 2δ per period, while the suburban property earns 1.5δ. The downtown property generates net operating income (NOI) of δ, while the suburban property generates NOI of only 0.5δ: suburban NOI is 50% lower despite suburban gross rents being only 25% lower.

Now suppose an aggregate shock reduces all gross rents in the city by 10%. The downtown property's NOI falls from δ to 0.8δ (a 20% decline), but the suburban property's NOI falls from 0.5δ to 0.35δ (a 30% decline). The suburban property experiences *amplified* sensitivity to rent shocks because its thinner operating margin (the ratio of NOI to gross rent) magnifies percentage changes. This amplification generates a larger cash-flow exposure and, when that exposure carries a positive premium, higher expected returns.

Modeling operating costs We formalize this intuition by introducing an operating cost δ (in units of local wages) that must be paid each period to maintain a developed property. Net operating income at location (k, d) and time t is then

$$\text{NOI}_{k,d,t} = r_{k,d,t} - \delta w_{k,t},$$

where $r_{k,d,t}$ is gross rent and $w_{k,t}$ is the local wage.

To price claims on this NOI stream, we work with an auxiliary state variable $x_{k,d,t}$ defined as the ratio of gross rent to wage:

$$x_{k,d,t} \equiv \frac{r_{k,d,t}}{w_{k,t}}.$$

Normalizing by wages is useful because, within a city, x varies only with location characteristics (distance, amenities) while the common wage level cancels. Cross-city comparisons of x are not meaningful. Beyond economic interpretation, the normalization reduces the pricing problem from two state

variables $(r_{k,d,t}, w_{k,t})$ to one $(x_{k,d,t})$, yielding the ordinary differential equation in Appendix B.3.[14] Properties closer to the CBD have higher x (higher gross rent per unit of wage) due to lower commuting costs, while distant properties have lower x.

With this notation, the dollar price and NOI can be written as

$$P_{k,d,t} = w_{k,t} \cdot p_k(x_{k,d,t}), \qquad \text{NOI}_{k,d,t} = w_{k,t}(x_{k,d,t} - \delta),$$

where $p_k(x)$ is the price function in wage-normalized units. The cap rate is then

$$\text{cap rate}_{k,d,t} = \frac{\text{NOI}_{k,d,t}}{P_{k,d,t}} = \frac{x_{k,d,t} - \delta}{p_k(x_{k,d,t})}.$$

Properties with lower x (farther from the CBD) have thinner operating margins $(x - \delta)/x$, which amplifies their exposure to shocks in x. We now make this precise.

Benchmark solution without options Consider first the case where properties are never abandoned and new development is exogenous (no real options). Adjusting for risk premia via the risk-neutral drift $\mu^{\mathbb{Q}}_{x,k}$ for the state variable x, the perpetual claim on NOI admits a closed-form solution:

$$p_k(x) = \frac{x}{\nu_{\text{Y},k} - \mu^{\mathbb{Q}}_{x,k}} - \frac{\delta}{\nu_{\text{Y},k}},$$

where

$$\nu_{\text{Y},k} = r_f + \lambda\sigma_{w,k} + \xi\varsigma_{w,k} - \mu_{w,k}, \qquad \mu^{\mathbb{Q}}_{x,k} = \mu_{r,k} - \mu_{w,k} - \left(\lambda\sigma_{x,k} + \xi\varsigma_{x,k}\right),$$

and x follows

$$\frac{dx}{x} = \mu_{x,k}\,dt + \sigma_{x,k}\,dB_t + \varsigma_{x,k}\,dB^k_t,$$

with $\sigma_{x,k} = \sigma_{r,k} - \sigma_{w,k}$ and $\varsigma_{x,k} = \varsigma_{r,k} - \varsigma_{w,k}$ (Appendix B.3). The rate $\nu_{\text{Y},k}$ discounts a claim to wages: the riskless rate plus the premium on wage risk, minus expected wage growth. The drift $\mu^{\mathbb{Q}}_{x,k}$ is the expected growth of rents relative to wages minus the premium on x risk. The price is a claim to rents minus a claim to wage-indexed costs. Rents are discounted at $\nu_{\text{Y},k} - \mu^{\mathbb{Q}}_{x,k} = r_f + \pi_k - \mu_{r,k}$, the cap rate of the baseline model, and costs are discounted at $\nu_{\text{Y},k}$.

Risk exposure depends on the *pricing elasticity*

$$\varepsilon_p(x) \equiv \frac{x \cdot p'(x)}{p(x)} = \frac{d\ln p}{d\ln x}.$$

[14] Gross rent still follows the geometric-Brownian process derived above. What changes is the cash flow being priced: net operating income, $r_{k,d,t} - \delta w_{k,t}$, is not itself geometric Brownian, so the proportional-price Gordon formula no longer applies. The wage normalization recovers tractability by making $x_{k,d,t} = r_{k,d,t}/w_{k,t}$ the state variable. The wage specification is a choice of convenience. The model accommodates more general costs, such as costs that follow their own index, but the price then depends on two state variables and solves a partial differential equation rather than an ordinary one. Any cost that is homogeneous of degree one in rent and the wage, including a cost level that differs across locations (Section 5.2), keeps the problem one-dimensional.

This elasticity measures the percentage change in price for a 1% change in the state variable x. Direct calculation yields

$$\varepsilon_p(x) = \frac{x}{x - \bar{\delta}_k}, \quad \text{where} \quad \bar{\delta}_k \equiv \delta \left(1 - \frac{\mu^{\mathbb{Q}}_{x,k}}{\nu_{Y,k}}\right).$$

Assumption 2 (Finite value)**.** $\nu_{Y,k} > \max\{\mu^{\mathbb{Q}}_{x,k}, 0\}$.

Values are finite when the wage-unit discount rate $\nu_{Y,k}$ is positive and exceeds the risk-adjusted growth of rents relative to wages. This is the analogue of $r > g$ in a Gordon growth model. We consider properties with positive value, $x > \bar{\delta}_k$.

Several features are immediate. First, for positive operating costs $\delta > 0$, $\varepsilon_p(x)$ is *decreasing* in x under Assumption 2. Properties farther from the CBD have lower x, hence higher elasticity. Second, $\varepsilon_p(x) = 1/(1 - \bar{\delta}_k/x)$ is the inverse of an effective operating margin. Because $\bar{\delta}_k$ incorporates risk-adjusted growth, this pricing elasticity differs from the accounting leverage ratio $x/(x - \delta)$. Thinner effective margins generate larger elasticities.

Third, and most importantly, this elasticity governs the risk loading in expected returns. The instantaneous expected excess return satisfies

$$\mathbb{E}_t[dR_{k,d,t}] - r_f\, dt = \left[(\lambda\sigma_{w,k} + \xi\varsigma_{w,k}) + \varepsilon_p(x_{k,d,t}) \cdot (\lambda\sigma_{x,k} + \xi\varsigma_{x,k})\right] dt,$$

where the first term captures direct exposure to wage shocks and the second term captures exposure to shocks in x (the rent-to-wage ratio), amplified by the pricing elasticity.

The cross-sectional prediction also requires the amplified exposure to carry a positive premium. Define $B_k \equiv \lambda\sigma_{x,k} + \xi\varsigma_{x,k}$. Because $\sigma_{r,k} = \sigma_{w,k} + \sigma_{x,k}$ and $\varsigma_{r,k} = \varsigma_{w,k} + \varsigma_{x,k}$, $B_k = \pi_k - \pi_{w,k}$, the difference between the baseline premium on rent risk and the premium on wage risk, $\pi_{w,k} \equiv \lambda\sigma_{w,k} + \xi\varsigma_{w,k}$. A property is a levered position, long a claim to rents that earns π_k and short a claim to wage-indexed costs that earns $\pi_{w,k}$. Its pricing elasticity is the value of the rent claim divided by the property's value. Expected excess returns can be rewritten as

$$\pi_k + (\varepsilon_p(x) - 1)(\pi_k - \pi_{w,k}),$$

the familiar formula for levered equity. Leverage raises expected returns if and only if rents carry more priced risk than wages, $B_k > 0$. Without operating costs, $\varepsilon_p = 1$ and expected excess returns equal the baseline premium π_k. Given that rent per unit of wage falls with distance, $\frac{dx}{dd} < 0$, $\varepsilon_p(x)$ rises with distance.

Prediction 2 (Total-return gradient)**.** *Within a city, expected excess total returns rise with distance from the center if and only if $B_k > 0$.*

In the diversified-owner benchmark, where $\xi = 0$ and $\lambda > 0$, this condition becomes $\sigma_{x,k} > 0$, or $\alpha\sigma_{w,k} > \sigma_\Psi$ under the rent dynamics above.

Operating leverage also shapes income returns. The cap rate equals $\nu_{Y,k} - \mu^{Q}_{x,k}\,\varepsilon_p(x)$ (Appendix C.2), so the income gradient has its own sign condition.

Prediction 3 (Income-return gradient)**.** *Within a city, cap rates and expected income returns rise with distance from the center if and only if* $\mu^{Q}_{x,k} < 0$.

Prediction 3 is the less obvious of the two. It requires negative risk-adjusted growth of rents relative to wages. Because $\mu^{Q}_{x,k} = (\mu_{r,k} - \mu_{w,k}) - B_k$, that happens either when rents are expected to fall relative to wages or when the premium on x risk exceeds their expected growth. The calibration in Section 5.2 points to the second.

The model implies that each location's full aggregate diffusion loading is $\sigma_{w,k} + \varepsilon_p(x_i)\sigma_{x,k}$. The second term varies across locations within a city, and over time as x_i evolves. Estimated ZIP-level betas are noisy proxies for this loading. Distance, measured with little error, captures the monotone relationship between location and operating leverage and can proxy for the risk-exposure gradient. The premium it earns in Table 6 is therefore consistent with the model.

When $B_k > 0$, properties with lower x have higher $\varepsilon_p(x)$ and earn higher expected returns. This rationalizes the distance gradient documented in Table 3. The mechanism works by amplification: all locations within city k face the same underlying wage and amenity shocks, and operating leverage magnifies the cash-flow sensitivity to these shocks for properties with thin margins.

Heterogeneity in x and the generality of the mechanism While our empirical analysis focuses on distance from the CBD (the classical source of rent variation in urban models), the operating leverage mechanism depends only on variation in x, not on its source. Neighborhoods with superior amenities (schools, parks, safety) command higher gross rents, generating higher x even at the same distance from the CBD. Properties in these high-amenity locations have thicker operating margins and lower pricing elasticity.

Consider two properties at distance d from the CBD, but in different sectors of the city. Property A is in a neighborhood with amenity level a_A, while property B faces $a_B < a_A$. From the equilibrium rent condition, property A's normalized rent is $x_A \propto a_A^{1/(1-\alpha)}$, while $x_B \propto a_B^{1/(1-\alpha)}$, so $x_A > x_B$. The elasticity formula immediately implies $\varepsilon_p(x_A) < \varepsilon_p(x_B)$: when $B_k > 0$, the lower-amenity property has greater priced exposure and earns higher expected returns, despite being equidistant from the CBD.

This generalizes the within-city gradient beyond the radial dimension. Operating leverage gen-

erates exposure heterogeneity whenever gross rents vary, whether through distance, amenities, or other neighborhood characteristics. It also helps reconcile our evidence with that of Demers and Eisfeldt (2022), who sort ZIP codes by price tier rather than by distance. A ZIP code can be cheap because it is far from the center, because it lacks amenities, or both, and in the model either reason lowers x and raises the pricing elasticity. Their finding that cheaper ZIP codes earn higher total returns is the same prediction, seen through a broader measure of x. Distance captures only the radial part of rent variation, and it does so with noise, since cities are not perfectly radial. Coastlines, transit corridors, and neighborhoods with better schools all shift rents at a given distance. This noise makes the distance gradient harder to detect. We use distance anyway because it is fixed by geography, while prices also reflect the expected returns we are trying to explain.

Property taxes Property taxes differ from the other operating costs because they scale directly with a building's assessed value. We model them as a proportional claim on contemporaneous value. The owner pays $\tau_P P_{k,d,t}$ per period, so the cash flow becomes $w_{k,t}(x_{k,d,t} - \delta) - \tau_P P_{k,d,t}$, and from here on NOI denotes this cash flow net of property taxes, as in the NCREIF data. Actual assessments adjust with lags, but taxing current value keeps the accounting simple. A tax proportional to price works like a higher discount rate, so the perpetual-claim solution is the one above with $\nu_{Y,k} + \tau_P$ in place of $\nu_{Y,k}$:

$$p_k(x) = \frac{x}{\nu_{Y,k} + \tau_P - \mu^{\mathbb{Q}}_{x,k}} - \frac{\delta}{\nu_{Y,k} + \tau_P},$$

where δ still excludes property taxes. The pricing elasticity has the same form as before,

$$\varepsilon_p(x) = 1/(1 - \theta_P e),$$

where $e = \delta/x$ is the non-tax cost share and $\theta_P = 1 - \mu^{\mathbb{Q}}_{x,k}/(\nu_{Y,k} + \tau_P)$. When rents are expected to fall relative to costs in risk-adjusted terms, a higher discount rate shrinks the present value of costs more than that of rents, because the costs extend further into the future. Property taxes therefore dampen but do not eliminate operating leverage. The after-tax cap rate is still $\nu_{Y,k} - \mu^{\mathbb{Q}}_{x,k}\,\varepsilon_p(x)$, which means the tax reaches income returns only through the elasticity.

Prediction 4 (Property taxes)**.** *Across locations within a city, property taxes rise more than one-for-one with rents.*

Within a city, the wage is common to all locations, so across locations the elasticity of value with respect to rent is the pricing elasticity, $d\ln P/d\ln r = d\ln p/d\ln x = \varepsilon_p(x) = 1/(1 - \theta_P e)$, which exceeds one under operating leverage. Tax payments are proportional to value, so they inherit this elasticity. Table 8 estimates the tax–rent elasticity across locations at 1.19, and Section 5.2 relates it to ε_p when costs also vary across locations. Because the non-tax share of rental income rises as rents fall while the tax share falls, value-based taxes offset part of the operating-leverage effect in all-in expenses.

Limitations The model abstracts from differences in property quality and from neighborhood effects other than amenities. Both likely contribute to within-city return variation, and a complete account of that variation would include them alongside operating leverage.

Extensions Appendix B.4 adds options to abandon, reuse, and build. Their contribution to the value of an operating property is concentrated where margins are thin and vanishes as rents rise relative to wages, so properties far from the abandonment threshold are priced close to the perpetual claim above. In the calibration, zero-salvage abandonment provides downside protection and narrows the income spread across distance quartiles (Appendix C.2). The options also create hysteresis: an idled building restarts only after rents recover above the level at which it stopped, and idle buildings return to use before new ones are built. The Online Appendix closes the system of cities. The outside option becomes an index of city wages, a wage gain in one city lowers rents elsewhere through migration, and cities with more elastic housing supply weigh more in the index because less of their wage gains is absorbed by rents.

4.3 Time-varying risk premia

When the price of aggregate risk is high, the premium for bearing operating leverage should be high too, and the return spread between thin-margin and thick-margin properties should widen. This section introduces time-varying risk premia and derives this prediction, which we later test in Section 5.4. We let the dollar price of aggregate risk $\lambda(z)$ depend on an aggregate state z_t, ordered so that the price of risk rises with z. The state reverts to its mean and is driven by the same aggregate shock as wages:

$$dz_t = \kappa_z(\bar{z} - z_t)\,dt + \sigma_z\,dB_t.$$

A moving price of risk also means that property values move when it changes, so part of the risk premium can be earned through expected appreciation as well as through the cap rate. A property's value now depends on two states, the rent-wage ratio and the aggregate state. Appendix B.5 derives the pricing equation for $p(x,z)$, which has the following exposures to x and z:

$$\varepsilon(x,z) = \frac{x\,p_x(x,z)}{p(x,z)}, \qquad \chi(x,z) = \frac{p_z(x,z)}{p(x,z)}.$$

The first is the pricing elasticity of Section 4.2. The second, χ, is the percentage change in the property's value per unit increase in the aggregate state, that is, how much the property revalues when the price of risk changes. Prices are in wage units, but returns are in dollars, so the expected dollar excess return carries the wage exposure σ_w and uses the dollar prices of risk:

$$\pi(x,z) = \lambda(z)\big[\sigma_w + \sigma_x\varepsilon(x,z) + \sigma_z\chi(x,z)\big] + \xi\big[\varsigma_w + \varsigma_x\varepsilon(x,z)\big].$$

As in Section 4.2, the direct term through which operating leverage is priced is $\lambda(z)\sigma_x\varepsilon$, the price of risk times the pricing elasticity, now at the current price of risk. The revaluation term $\lambda(z)\sigma_z\chi$ is also priced at $\lambda(z)$, because the aggregate state is driven by the same shock as wages.

Without real options, the price is the difference between a claim to rents and a claim to operating costs, $p(x,z) = a(z)x - \delta c(z)$. In properties where that difference is thin, any gap between the revaluations of the two claims is magnified into a large percentage change in property value, in the same way that this thin margin magnifies rent shocks. The part of χ that differs across locations is therefore proportional to ε (see equation (B19) in Appendix B.5).

Collecting terms, the expected excess return at location i has a part common to the city and a part proportional to its pricing elasticity:

$$\pi_i(z) = \pi_0(z) + \Gamma(z)\,\varepsilon_i(z), \qquad \pi_i(z) - \pi_j(z) = \Gamma(z)\big[\varepsilon_i(z) - \varepsilon_j(z)\big].$$

$\Gamma(z)$ is the dollar premium per unit of pricing elasticity. It equals the dollar price of aggregate risk times the aggregate loading of the rent claim in excess of the cost claim, plus the local premium $\xi\varsigma_x$ (see equation (B21)). A higher price of risk widens the spread between thin-margin and thick-margin properties when $\Gamma(z)$ is positive and rising and the gap between their elasticities does not narrow (Appendix B.5).

Prediction 5 (Leverage premium over time)**.** *Within a city, the expected return spread between high- and low-leverage properties widens when the price of aggregate risk is high.*

The prediction needs both ingredients of the model. Without operating costs, ε equals one at every location and the spread is zero. Without a time-varying price of risk, Γ is constant and, at given leverage, the spread does not move with the state.

Section 5.4 takes Prediction 5 to the data. At any date, a location's pricing elasticity is an increasing function of its non-tax operating leverage, $\omega_i^{NT} = x_i/(x_i - \delta)$, which we measure from expense accounts, and the consumption-aggregate wealth ratio *cay* serves as a proxy for the aggregate state. Up to a first-order approximation, and with coefficients common across cities, the location-specific premium $\Gamma(z)\varepsilon_i(z)$ becomes a regression of returns on leverage interacted with *cay*, with location and date fixed effects (Appendix B.5). The interaction coefficient measures how much the leverage premium moves with the state. It mixes the change in the price of risk with changes in the loadings and in the elasticity, so it does not recover the price of risk by itself. The model also says where the interaction should appear in returns. When the price of risk mean-reverts quickly relative to the rates at which rents and costs are discounted, a change in the state barely moves yields, and the interaction is earned mostly through expected capital gains (Appendix B.5).

5 Quantitative Assessment

This section takes the model of Section 4 to the data. Section 5.1 presents direct evidence for operating leverage: non-tax operating costs rise less than proportionally with rents across locations within a city. Section 5.2 calibrates the model to the within-city gradients in income and total returns. Section 5.3 asks whether the implied price of risk generates a cross-city spread consistent with the data, and Section 5.4 tests whether the leverage premium moves with the aggregate price of risk.

5.1 Cash flows: the cost structure behind operating leverage

The operating-leverage mechanism requires that non-tax operating costs fall less than rents across locations within a city. It implies that NOI responds more than one-for-one to shocks to rents and that property taxes, which are levied on value, rise more than one-for-one with rents across locations.

Costs fall less than rents across locations. In the model, the non-tax cost of operating a property is $\delta w_{k,t}$, which follows the city wage. As before, let $x = \frac{r}{w}$ be the ratio of gross rent to the wage. Where gross rent is lower, x is lower, the non-tax cost share $e = \frac{\delta}{x}$ is higher, and the margin is thinner. Operating leverage does not need costs to be literally fixed across locations, only that they rise less than proportionally with rents. We measure this proportionality in the NCREIF expense accounts, which report rental income and expenses per square foot, by expense category, for each reporting property. For each ZIP code i we average rental income per square foot, r_i, and expenses per square foot in category j, C_i^j, over 2000 to 2025, and estimate

$$\ln C_i^j \;=\; \alpha_{c(i)} \;+\; \eta_j \ln r_i \;+\; \beta_j' Z_i \;+\; u_i, \tag{4}$$

where $\alpha_{c(i)}$ is a fixed effect for the CBSA of ZIP code i and Z_i contains the log average property size and the log number of reporting properties. The CBSA fixed effect absorbs the city wage and the city-wide level of rents and costs, so η_j is estimated from differences across ZIP codes within a city. Property size controls for scale economies, and the number of reporting properties for differences in coverage. Standard errors are clustered by CBSA because ZIP codes in the same city share cost and rent shocks. An elasticity of one means that costs scale with rents and margins are the same everywhere. An elasticity of zero means that costs are fixed, so a ZIP code with ten percent lower rents has the same costs and a thinner margin.

NCREIF expense accounts cover fewer ZIP codes than the return panel, 258 against 312. Table 8 reports the estimates for the 112 of these, in 22 CBSAs, with at least 20 quarters of expense data and a measured distance to the primary employment center, and Panel A of Table 7 summarizes this sample: in the average ZIP code, non-tax operating expenses absorb 31 percent of rental income. Non-tax

operating costs have an elasticity of 0.80 (s.e. 0.04): where rental income is ten percent lower, these costs are just eight percent lower and thus absorb a larger share of income.[15]

Table 7: Summary statistics for the expense panel and the NOI and rent growth indices

	Mean	SD	P25	Median	P75	*N*
Panel A. NCREIF expense panel: ZIP-code averages						
Rental income ($ per sq. ft. per quarter)	6.28	2.66	4.39	5.63	8.41	112
Non-tax operating expenses ($ per sq. ft. per quarter)	1.86	0.76	1.22	1.64	2.38	112
Property taxes ($ per sq. ft. per quarter)	0.89	0.45	0.51	0.86	1.17	112
Total operating expenses ($ per sq. ft. per quarter)	2.77	1.16	1.84	2.55	3.55	112
Non-tax expenses / rental income	0.305	0.045	0.276	0.298	0.330	112
Market value ($ per sq. ft.)	346	191	199	302	498	112
Average property size (thousand sq. ft.)	297	133	231	284	337	112
Reporting properties per quarter	4.6	1.8	3.5	4.2	5.0	112
Quarters observed	40	18	26	37	52	112
Housing built before 1970 (share)	0.39	0.28	0.10	0.41	0.62	111
Distance to employment center (miles)	7.5	7.5	2.0	4.5	11.8	112
Panel B. NCREIF research indices: metro-quarters (first two rows) and metros						
Rent growth (% per quarter)	1.09	2.51	−0.47	1.04	2.60	4,839
NOI growth (% per quarter)	1.15	6.64	−2.91	0.98	4.80	4,839
SD of rent growth (% per quarter)	2.63	0.55	2.22	2.60	2.88	76
SD of NOI growth (% per quarter)	7.28	2.67	5.67	6.56	8.21	76
SD of NOI growth / SD of rent growth	2.74	0.63	2.24	2.59	3.11	76
Quarters observed	61	33	34	72	92	79

Notes: Panel A describes the sample of Table 8: 112 ZIP codes in 22 CBSAs with at least 20 quarters of NCREIF expense data between 2000Q1 and 2025Q2 and a measured distance to the primary employment center. Per-square-foot amounts and the expense share are winsorized at the 5th and 95th percentiles across ZIP-quarters and then averaged within ZIP code. Non-tax operating expenses are total operating expenses minus property taxes. The pre-1970 housing share is from Census 2000 (111 ZIP codes). Distance is to the primary employment center (Section 2.2). Panel B uses the NCREIF research indices of apartment NOI and rent growth for 79 metropolitan areas, 2001Q1 to 2023Q4. Quarterly growth rates are winsorized at the 5th and 95th percentiles, each series separately. The first two rows pool metro-quarters; the next three describe per-metro standard deviations and their ratio, for the 76 metros with at least two quarters of data.

NOI amplifies changes in rents. When rents rise and costs do not, the whole increase in rent passes to NOI, but NOI is a smaller base. Owners therefore see net income swing by more, in percentage terms, than rents. If costs instead scaled with rents, NOI would be a fixed fraction of rent and the

[15]Better buildings may command higher rents and also cost more to run. Operating leverage depends on the share of rental income that operating costs absorb, which the expense accounts measure directly, whatever the reason those costs are high. The estimates in Table 8 already control for property size, and adding neighborhood vintage lowers the elasticity to 0.75.

Table 8: Elasticity of operating costs with respect to rents across locations.

Expense measure	Elasticity with respect to rental income	
	Size controls	+ neighborhood vintage
Operating expenses excluding property taxes	0.80 (0.04)	0.75 (0.05)
Property taxes	1.19 (0.07)	1.23 (0.09)
Total operating expenses	0.94 (0.03)	0.90 (0.03)

Notes: Estimates of equation (4). ZIP-level regressions of log expenses per square foot on log rental income per square foot, both averaged within ZIP code over 2000 to 2025. NCREIF expense accounts, 112 ZIP codes in 22 CBSAs. All specifications include CBSA fixed effects. Standard errors, in parentheses, are clustered by CBSA. Size controls are log average property size and log number of reporting properties. Neighborhood vintage is the share of housing built before 1970 (Census 2000, 111 ZIP codes).

two would be equally volatile. The NCREIF research indices of NOI and rent growth, which cover 79 metropolitan areas from 2001Q1 to 2023Q4 and which we winsorize at the 5th and 95th percentiles, show the amplification (Panel B of Table 7). In the average metropolitan area, NOI growth is 2.7 times as volatile as rent growth, and the standard deviations average 7.3 and 2.6 percent per quarter (Appendix C.1).

Volatility alone does not isolate operating leverage. NOI also moves for reasons unrelated to rents, such as a roof replacement, a tax reassessment, or a change in insurance premiums, and these raise its volatility even where costs scale with rents. The more direct test asks how NOI responds when rents move. If costs do not respond to a change in rents, the elasticity of NOI with respect to gross rent is

$$\frac{d \ln \mathrm{NOI}}{d \ln r} = \frac{r}{r - C} = \frac{1}{1 - e}, \tag{5}$$

where $e = C/r$ is the cost share. The thinner the margin, the larger the percentage response of NOI to the same change in rents. A regression of quarterly NOI growth on rent growth estimates this elasticity, and expense shocks unrelated to rents fall into the residual rather than the slope. With quarter fixed effects, which absorb shocks common to all metros, the slope is 1.44 (s.e. 0.06), with standard errors clustered by quarter.[16] A one percentage point rise in rent growth thus comes with a 1.44 percentage point rise in NOI growth. This is close to what equation (5) implies for a property whose non-tax costs absorb 30 percent of rent, the cost share of the closest distance quartile in the calibration below: $1/(1 - 0.297) = 1.42$. In the ZIP-level expense accounts, a regression of four-quarter growth in NOI before property taxes on growth in rental income, with ZIP-code and quarter fixed effects, gives a slope of 1.45 (s.e. 0.05, clustered by CBSA). If each ZIP code's non-tax costs had stayed at their initial level, the observed changes in rental income would imply a slope of 1.47 in the same regression.

[16]The pooled slope without the quarter fixed effects is 1.34 (s.e. 0.07). Appendix Table C1 reports both.

Property taxes rise more than one-for-one with rents. Unlike other items of operating costs, property taxes are levied on value. Under operating leverage, value moves more than one-for-one with gross rent across locations within a city (Section 4.2), so taxes levied on value should do the same. Regression (4) with property taxes on the left gives 1.19 (s.e. 0.07), and 1.23 (s.e. 0.09) with the vintage control. All-in expenses, which combine non-tax costs that rise less than rents with taxes that rise more, move nearly one-for-one with rents, at 0.94 (s.e. 0.03). Section 5.2 uses these two elasticities to compare cost profiles.

5.2 Calibrating the within-city return gradients

Both return components rise with distance from the employment center. In Table 3, excess income returns rise from 0.72 pp/qtr in the closest distance quartile to 0.89 in the farthest, and excess total returns rise from 0.94 to 1.55. We now map the operating-cost model of Section 4.2 into the distance quartiles used in the calibration. In the model, both gradients run through the pricing elasticity ε_q, the elasticity ε_p of Section 4.2 evaluated at quartile q. We calibrate the elasticities to the income gradient and then use the total-return gradient to measure the price of exposure.

Income returns in the model. Section 4.2 prices a property as a perpetual claim on NOI, with property taxes a flow proportional to value. Here we evaluate the same pricing equation for quartile q, with rent-wage ratio x_q and non-tax cost share $e_q = \delta_q / x_q$. The expected excess income return in quartile q is then (Appendix C.2)

$$m_q \;=\; \underbrace{(\nu_Y - r_f)}_{\text{common to the city}} \quad \underbrace{-\,\mu_x^{\mathbb{Q}}\,\varepsilon_q}_{\text{operating leverage}}, \qquad \varepsilon_q \;=\; \frac{1}{1-\theta_P\, e_q}, \qquad \theta_P \;=\; 1 - \frac{\mu_x^{\mathbb{Q}}}{\nu_Y + \tau_P}. \tag{6}$$

Here m_q is NOI divided by value, minus the risk-free rate, which is the excess income return that NCREIF measures. The first term is common to the city. Its main ingredient, $\nu_Y = r_f + \lambda\sigma_w - \mu_w$, is the yield on a claim to the city wage: the risk-free rate plus the premium on aggregate wage risk, minus expected wage growth. The second term varies across locations through the pricing elasticity ε_q, the percentage response of value to gross rent. Its coefficient, $\mu_x^{\mathbb{Q}}$, is the risk-adjusted drift of the rent-to-wage ratio: its expected growth minus the risk adjustment. When $\mu_x^{\mathbb{Q}} < 0$, investors value rents as if they will fall relative to costs. The expected decline lowers value relative to current NOI, and by more where value is more sensitive to rents, so the income return rises with ε_q.

To calibrate equation (6), we need to convert the quartile cost shares e_q into the model's pricing elasticities ε_q derived in Section 4.2. A property is a claim on future rents net of future costs, so its value is the difference between two present values. Holding wages fixed, a one percent increase in gross rent raises the value of the gross-rent claim by one percent, while leaving the wage-indexed cost

claim unchanged. The percentage increase in the net property value is therefore the value of gross rents divided by the value of rents net of costs, which is the pricing elasticity ε_q:

$$p_q = \underbrace{\frac{x_q}{\nu_Y + \tau_P - \mu_x^Q}}_{\text{PV of gross rents}} - \underbrace{\frac{\delta_q}{\nu_Y + \tau_P}}_{\text{PV of operating costs}}, \qquad \varepsilon_q = \frac{\text{PV of gross rents}}{p_q} = \frac{1}{1 - \theta_P e_q}. \tag{7}$$

Here δ_q is the non-tax cost level that generates quartile q's cost share, e_q. In the fixed-cost version below, $\delta_q = \delta$ for all quartiles. This is the present-value analogue of the flow formula $1/(1-e)$ in equation (5). The flow formula uses the current cost share e_q, while the pricing formula uses the present-value cost share $\theta_P e_q$. The adjustment θ_P appears because rents and costs are capitalized at different rates. Costs follow wages and are capitalized at $\nu_Y + \tau_P$. Rents drift relative to wages at the risk-adjusted rate μ_x^Q and are capitalized at $\nu_Y + \tau_P - \mu_x^Q$. The wedge θ_P is the ratio of the two rates. When $\mu_x^Q < 0$, rents are capitalized at the higher rate, so costs weigh more in value than in current income and $\theta_P > 1$. Property taxes add τ_P to both rates and pull their ratio toward one, attenuating operating leverage.

Two cost profiles. The calibration needs the non-tax cost share e_q in each distance quartile. The model of Section 4.2 holds the non-tax cost $\delta w_{k,t}$ fixed across locations within a city, which keeps it tractable: the cost share is then inversely proportional to rent. The cost regressions point instead to partial fixedness, with an elasticity of 0.80 (Table 8). We calibrate the model under both cost profiles. With η the elasticity of non-tax costs with respect to rents across locations, the cost share in quartile q is

$$e_q = e_1 \left(\frac{r_1}{r_q} \right)^{1-\eta}, \tag{8}$$

where $e_1 = 0.297$ is the share in the closest quartile of the expense panel and r_q/r_1 is the ratio of gross rent in quartile q to gross rent in the closest quartile, measured from log rents net of CBSA means. Table 9 reports the rent ratios and both cost-share profiles. With a fixed cost, the cost share rises from 0.297 in the closest quartile to 0.453 in the farthest, and with the estimated elasticity only to 0.323. The shares measured directly in each quartile stay between 0.297 and 0.324, and a calibration to them does not reproduce the income gradient (Appendix C.3).

The estimated profile requires only that the cost level differ across locations. With a non-tax cost of $\delta_d w_{k,t}$ at location d, constant over time in wage units, the pricing formulas of Section 4.2 hold location by location with $e = \delta_d / x$.[17] In both profiles, costs follow the wage over time, which is what produces the amplification in equation (5).

[17] The pricing equation is linear in the cost term, so with a location-specific cost $\delta_d w_{k,t}$ the price is $p_d(x) = x/(\nu_Y + \tau_P - \mu_x^Q) - \delta_d/(\nu_Y + \tau_P)$, and equations (6) and (7) hold with $e_d = \delta_d / x_d$. In this specification η is an elasticity across locations only: over time, costs follow the wage and do not respond to a location's own rents. One interpretation is building

Calibration procedure. The remaining inputs are the average Treasury bill rate over the ZIP-quarters of the return panel, $r_f = 0.35$ pp/qtr, and the effective property-tax rate on value, $\tau_P = 0.29$ pp of value per quarter. Given these inputs, equation (6) has two unknowns, ν_Y and μ_x^Q, which together set the common level, the wedge, and the slope on ε_q. There is no free intercept or slope. For each cost profile, we choose the pair that minimizes the equally weighted sum of squared differences between m_q and the four observed excess income returns, solving numerically from multiple starting values because the mapping is nonlinear in μ_x^Q. With two parameters and four moments, each fit leaves two overidentifying restrictions on the shape of the income profile. The search is restricted to the region where the model is well defined: $\nu_Y > 0$, so that a claim to the wage has finite value, $\nu_Y + \tau_P > \mu_x^Q$, the analogue of $r > g$ in a riskless Gordon model, and $\theta_P e_q < 1$ in every quartile, so that property values are positive. Both minimizers lie inside this region.

Table 9 reports the fitted profiles and Table 10 the parameters. Both calibrations fit the income gradient closely: the largest residual is 1.26 bp with a fixed cost and 0.79 bp with the estimated elasticity, against a 17 bp spread between the farthest and closest quartiles. Both require negative risk-adjusted growth of rents relative to wages, with $\mu_x^Q = -0.233$ and -0.387 pp/qtr. In the model, income returns rise with distance because $\mu_x^Q < 0$: the risk adjustment on rents relative to wages exceeds their expected relative growth. We return to this condition below, where total returns give it a simpler form.

The two cost profiles differ mainly in how a current cost share translates into value leverage. The estimated profile fits the income gradient by making properties more levered in value than in income. In both calibrations, non-tax costs absorb 30 percent of rent in the closest quartile. With a fixed cost, this current cost share becomes a present-value cost share of 37 percent, so a one percent rise in rent raises value by $1/(1-0.37) = 1.6$ percent. With the estimated cost profile, the same current cost share becomes a present-value cost share of 61 percent, so the same one percent rise in rent then raises value by $1/(1-0.61) = 2.6$ percent. The estimated profile needs this stronger value leverage because cost shares vary little across quartiles, so small differences in cost shares have to produce the same 17 bp spread in income returns.

Costs weigh more in value than in income because investors value rents as if they will fall relative to wages, so the cost share of rent is expected to rise. The wedge $\theta_P = 1 + (-\mu_x^Q)/(\nu_Y + \tau_P)$ compares that rise with the wage-unit discount rate. With a fixed cost, the cost share rises at a quarter of the discount rate, and $\theta_P = 1.23$. With the estimated profile, it rises slightly faster than the discount rate, and $\theta_P = 2.06$. Most of the change comes from the discount rate: ν_Y falls from 0.71 to 0.08 pp/qtr, so future wage-indexed costs are discounted lightly and carry more weight in value. The leverage term $-\mu_x^Q \varepsilon_1$ then accounts for 93 percent of the closest quartile's income return, against about a third with

quality. Higher-rent locations hold buildings with more amenities and services, whose cost level is fixed in wage units. The alternative in which costs also respond to rents over time, $C = \delta r^\eta w^{1-\eta}$, keeps $x = r/w$ as the state variable but caps the response of NOI to rents at $(1-\eta e)/(1-e) \approx 1.08$, well below the estimated 1.44 in Section 5.1.

Table 9: Calibration of the within-city return gradients.

	Q_1	Q_2	Q_3	Q_4
Panel A: Data				
Gross rent relative to Q_1, r_q/r_1	1	0.893	0.699	0.656
Excess income return (pp/qtr)	0.72	0.77	0.85	0.89
Excess total return (pp/qtr)	0.94	1.31	1.47	1.55
Panel B: Fixed non-tax cost ($\eta = 0$)				
Non-tax cost share e_q	0.297	0.332	0.425	0.453
Fitted excess income return (pp/qtr)	0.73	0.76	0.85	0.89
Pricing elasticity ε_q	1.58	1.69	2.10	2.27
Panel C: Estimated cost elasticity ($\eta = 0.80$)				
Non-tax cost share e_q	0.297	0.304	0.319	0.323
Fitted excess income return (pp/qtr)	0.72	0.76	0.86	0.89
Pricing elasticity ε_q	2.58	2.68	2.92	3.00

Notes: Equally weighted least-squares fits of equation (6) to the four distance-quartile mean excess income returns, restricted to $\nu_Y > 0$, $\nu_Y + \tau_P > \mu_x^Q$, and $\theta_P e_q < 1$. Cost shares follow equation (8) from the closest-quartile share of the expense panel and the within-CBSA rent gradient. Rent ratios are computed from mean log rents per square foot in each quartile, net of CBSA means. Both fits use $r_f = 0.35$ pp/qtr and $\tau_P = 0.29$ pp of value per quarter. Return panel: 312 ZIP codes in 22 CBSAs, 2000Q1 to 2023Q4. Table 10 reports the parameters.

a fixed cost.

The property-tax elasticity, which the calibrations do not fit, favors the estimated profile over the fixed cost. Property taxes are levied on value, so their elasticity with respect to rent across locations measures how value varies with rent within a city. That elasticity is not ε_q, which holds costs fixed: across locations, a property with higher rent also has higher costs when $\eta > 0$. Because the wage is common within a city,

$$\frac{d\ln(\tau_P P)}{d\ln r} \;=\; 1 + (1-\eta)\,(\varepsilon_q - 1), \tag{9}$$

where $\varepsilon_q - 1$ is the present value of costs relative to value and $1-\eta$ is the part of costs that does not rise with rent across locations. Consider two properties in the same city whose rents differ by 10 percent. With a fixed cost, their values, and hence their tax bills, differ by about 20 percent. With the estimated profile, their costs differ by 8 percent and their tax bills by about 14 percent. In the data, tax bills differ by 12 percent (Table 8).

The income gradient pins down the sign of risk-adjusted rent growth, but not how a property's yield splits between the base yield ν_Y and the operating-leverage term. Different assumptions on the cost elasticity between $\eta = 0$ and $\eta = 0.80$ give calibrations that fit the four income returns about equally well, with a largest residual between 0.8 and 1.3 bp, and $\mu_x^{\mathbb{Q}}$ stays between -0.23 and -0.40 pp/qtr throughout. The split, however, depends on the assumption. As η rises from 0 to 0.80, ν_Y falls from 0.71 to 0.08 pp/qtr, θ_P rises from 1.23 to 2.06, and the leverage share of the closest quartile's income return rises from about a third to 93 percent. Within this family of cost profiles, the tax elasticity, which none of these calibrations fits, favors the scenarios near the estimated cost profile, where the base yield ν_Y is low. Separating the two components further would require different data, such as operating statements for comparable buildings in different locations, which would measure how costs vary with rent when building quality is held fixed.

Total returns and the price of exposure. The same elasticity ε_q governs expected total returns. Applying Itô's lemma to the dollar price $P = w\,p(x)$ and pricing with the stochastic discount factor gives

$$\mathbb{E}[R_q] - r_f \;=\; \underbrace{(\lambda\sigma_w + \xi\varsigma_w)}_{\text{wage exposure}} + B\,\varepsilon_q, \qquad B \;=\; \lambda\sigma_x + \xi\varsigma_x \tag{10}$$

(see Appendix C.2). Every property in a city carries the same exposure to wage shocks, which sets the common level. Operating leverage scales the exposure to rent-to-wage shocks, and B is the premium per unit of that scaled exposure. Holding priced exposure fixed, a difference in expected growth moves yields and expected appreciation by offsetting amounts and leaves expected total returns unchanged.

We hold ε_q at its calibrated values and regress the four quartile excess total returns on ε_q with a free intercept, which absorbs the common wage-exposure term. The slope is $B = 0.75$ pp/qtr per unit

Table 10: Parameters and implied expense elasticities under the two cost profiles.

	Fixed non-tax cost ($\eta = 0$)	Estimated cost elasticity ($\eta = 0.80$)	Data
Panel A: Parameters			
ν_Y (pp/qtr)	0.71	0.08	
$\mu_x^{\mathbb{Q}}$ (pp/qtr)	−0.233	−0.387	
Wedge θ_P	1.23	2.06	
Panel B: Fit and implications			
Largest income residual (bp)	1.26	0.79	
Leverage share of Q_1 income return	0.34	0.93	
Total-return slope B (pp/qtr per unit of ε_q)	0.75	1.27	
Implied property-tax elasticity	1.96	1.36	1.19 (0.07)

Notes: Parameters of the fits in Table 9. The leverage share is $-\mu_x^{\mathbb{Q}}\varepsilon_1$ divided by the Q_1 income return (NOI over value). B is the slope of a regression of the four quartile excess total returns on ε_q with a free intercept (equation (10)). The implied tax elasticity is equation (9), averaged across quartiles with cost-share weights. Data column from Table 8, with standard errors in parentheses.

of elasticity with a fixed cost and 1.27 with the estimated elasticity, which spreads the same gradient over a narrower range of elasticities. The implied outer-minus-inner spread is about 51 and 53 bp, against 61 bp in the data. Because ε_q rises with distance, the model implies total returns that rise monotonically across quartiles, as they do in the data. Returns in the data rise most between the first two quartiles, while the elasticities rise most between the second and third.

The gap between the two gradients. Total returns rise with distance about three times as much as income returns, and capital gains account for the difference. In the model, the slope of expected appreciation on ε_q is the expected growth of dollar rents relative to wages,[18]

$$B + \mu_x^{\mathbb{Q}} \;=\; g \;\equiv\; \mu_x + \sigma_w\sigma_x + \varsigma_w\varsigma_x. \tag{11}$$

The slope of income returns, $-\mu_x^{\mathbb{Q}} = B - g$, is therefore positive when the price of exposure exceeds the expected relative growth of rents. When that growth is modest, the condition is close to $B > 0$, which the total-return gradient requires.

The total-return gradient requires a large capital-gain component. In the model, the gap between the total-return and income-return slopes is the expected growth of rents relative to wages. That gap is 0.51 pp/qtr with a fixed cost and 0.88 with the estimated elasticity, implying rents would have to grow 2.1 to 3.5 percentage points a year faster than wages throughout the city, a large drift to sustain

[18]The covariance terms are not risk premia. They arise because x is measured in wage units: when wages and x move together, dollar rents grow faster than the drift of x alone.

over long horizons.[19] Nationally, rents outgrew wages by less than half a percentage point a year from 2000 to 2023, and by about 0.8 in the 2010s.[20] The calibration reads sample-average capital gains as expected returns. If part of the capital-gain gradient is realized in the sample in excess of expectations, the calibrated $\widehat{B}$ overstates the price of exposure.

5.3 One price of risk across and within cities

In our baseline model, a single stochastic discount factor prices both locations within a city and cities in the cross section, and there is no premium on city-specific risk ($\xi = 0$). The price of aggregate risk recovered from the within-city return gradient should therefore also price cross-city differences in exposure to aggregate wage shocks. With $\xi = 0$, equation (10) delivers that the within-city premium per unit of operating-leverage exposure is $B = \lambda\sigma_x$, so

$$\lambda = B/\sigma_x.$$

Using B from the fixed-cost calibration and a proxy of 2.6 percent per quarter for σ_x, the average within-CBSA volatility of quarterly rent growth (Appendix C.4),[21] gives $\lambda \approx 0.29$ per square-root quarter. This is the aggregate price of risk implied by the within-city gradient.

We then ask what this price of risk implies across cities. Cities more exposed to aggregate wage shocks earn higher returns. We measure this exposure by the slope of their apartment returns on national payroll growth from 2001 to 2019, for the 28 CBSAs with payroll data and long enough return series.[22] The third of cities with the highest exposure earns 0.26 pp/qtr (s.e. 0.22) more than the third with the lowest. Their difference in exposure is 0.59 percent per square-root quarter. Multiplying this exposure gap by the within-city price of risk, $\lambda = 0.29$, gives an implied spread of 0.17 pp/qtr.

The implied spread is proportional to B and inversely proportional to the volatility proxy used for σ_x. Under the estimated cost profile, $B = 1.27$ raises it to 0.29 pp/qtr. With the directly measured volatility of the rent-to-wage ratio, 4.36 percent per quarter, it is 0.10 pp/qtr with a fixed cost and 0.17

[19] In the model, rents can outgrow wages indefinitely: with a constant expenditure share, households rent less space as rents rise relative to wages.

[20] Bureau of Labor Statistics consumer price index for rent of primary residence, against average hourly earnings of production and nonsupervisory workers and against median usual weekly earnings of full-time workers; annual averages. The rent index tracks sitting tenants and lags market rents on new leases. See also Quigley and Raphael (2004) and Albouy et al. (2016).

[21] Strictly, σ_x is the volatility of the rent-to-wage ratio. We use the volatility of rent growth because quarterly local earnings are noisy, which inflates the measured volatility of the ratio. The implied price of risk, about 0.6 a year, is of the same order as the Sharpe ratio of the aggregate stock market. Appendix C.4 and Table C4 report the results with the directly measured ratio.

[22] Payroll combines earnings and employment, so the payroll beta is an empirical proxy for the model's wage loading. Across CBSAs it has a correlation of 0.40 with the apartment-index beta priced in Section 3.1, and the relation is tighter over 2001 to mid-2010 (Appendix C.5).

with the estimated profile. Across the two cost profiles and volatility inputs, it ranges from 0.10 to 0.29 pp/qtr, and it reaches 0.45 under a calibration to directly measured cost shares, which does not reproduce the distance gradient (Table C4). Therefore, the price of risk recovered from the within-city gradient calibration implies a cross-city spread of the same order as the observed one, which the data measure imprecisely.

5.4 The leverage premium over time

Prediction. The spread between high- and low-leverage properties should widen when aggregate expected returns are high. In the model, the expected return spread between two locations in a city is $\Gamma(z)\,[\varepsilon_i(z) - \varepsilon_j(z)]$, where z is the aggregate state and $\Gamma(z)$ is the premium per unit of exposure (Prediction 5). We test this prediction using operating leverage measured directly from each location's expense accounts.

Measurement. The measure is each ZIP code's non-tax operating leverage,

$$\omega_i^{NT} = \frac{\text{gross rental income}_i}{\text{gross rental income}_i - \text{non-tax operating expenses}_i},$$

the accounting counterpart of $1/(1-e)$ in equation (5), which is monotone in ε_i at each date. It excludes property taxes because these are a claim on value. We compute the ratio for each ZIP quarter and average it over the ZIP code's expense history. For the aggregate state we use cay_t, the consumption-wealth ratio of Lettau and Ludvigson (2001a), a standard predictor of aggregate expected returns, following the scaled-factor approach of Cochrane (1996), Lettau and Ludvigson (2001a), and Lettau and Ludvigson (2001b). We estimate

$$R_{i,t+h} \;=\; \alpha_i \;+\; \delta_t \;+\; \gamma\left(\omega_i^{NT} \times cay_t\right) \;+\; u_{i,t+h}, \tag{12}$$

where $R_{i,t+h}$ is the compounded total return of ZIP code i over its next h observed ZIP-quarter records, in decimal units. A ZIP code does not report in every quarter, so the next h records can span more than h calendar quarters. The notes to Table 11 report how often. ZIP fixed effects absorb each location's average return, including its average leverage premium. Date fixed effects absorb everything common to all ZIP codes at a date, including the direct effect of cay_t. The coefficient γ is therefore identified from whether high-leverage ZIP codes outperform low-leverage ZIP codes by more when cay_t is high. Standard errors are two-way clustered by CBSA, which allows for correlation within a city over time, including the overlap of compounded returns when $h > 1$, and by date, which allows for shocks common to all ZIP codes in a quarter. The sample has 4,372 ZIP-quarter observations from 1991Q1 to 2019Q2, covering 188 ZIP codes in 23 CBSAs. The *cay* series ends in 2019, so the test extends the return sample back before 2000. It does not use the Census 2000 controls that set the 2000 start of the

main panel. Like other scaled-factor regressions, equation (12) is a reduced form. In the model, γ mixes the state dependence of the price of risk with the aggregate loading and the valuation elasticity (Appendix B.5), so the prediction concerns its sign.

Estimates. The leverage premium rises with *cay*. For the next record, which is almost always the next quarter, $\gamma = 1.140$ (s.e. 0.442), and Table 11 shows that the estimates are positive at every horizon and grow with it, as one expects from a persistent conditioning variable.

This next-quarter estimate is economically large. Compare ZIP codes at the 10th and 90th percentiles of ω^{NT} (1.331 and 1.614). When *cay* moves from its 10th to its 90th percentile, the predicted return spread between them widens by 125 bp per quarter. The fixed effects absorb the unconditional spread, so only this conditional change is identified. The coefficient is positive in both halves of the sample split at 2010 (Appendix C.7), and it remains positive when the *cay* slope is also allowed to vary with distance, neighborhood income, and college share (Appendix C.6.1).

Table 11: $\omega^{NT} \times cay$ Predictability over Observed ZIP-Quarter Records

Observed records included:	1	2	4	8
Model:	(1)	(2)	(3)	(4)
	Compounded Total Return $R_{t\to t+h}$			
Variables				
$\omega^{NT} \times cay$	1.140**	2.028***	4.772***	7.520**
	(0.4420)	(0.6877)	(1.309)	(3.063)
Fixed-effects				
Zip	Yes	Yes	Yes	Yes
date	Yes	Yes	Yes	Yes
Fit statistics				
Observations	4,372	4,184	3,826	3,200
R^2	0.44112	0.58302	0.67693	0.70854

Notes: Estimates of equation (12). The dependent variable is the compounded total return over the next h observed ZIP-quarter records, in decimal units (0.01 = 1%), unlike the percentage-point convention used for the tables in Section 3. The shares of usable leads that span more than h calendar quarters are 2.29%, 4.52%, 8.52%, and 14.88% for $h = 1, 2, 4, 8$, respectively. Date fixed effects are indexed by the formation quarter t. ZIP and date fixed effects included. Parentheses contain standard errors two-way clustered by CBSA and date. ω_i^{NT} is the ZIP-level ratio of gross rental income to gross rental income less non-tax operating expenses, and cay_t is from Lettau and Ludvigson (2001a). Significance codes: *** at 1%, ** at 5%, and * at 10%.

Almost all of the effect runs through capital returns. Estimating equation (12) separately for the two components of the next-quarter return gives 1.074 (s.e. 0.440) for capital returns and 0.066 (s.e. 0.050) for income returns. This is the split the model predicts when the price of risk mean-reverts quickly relative to the rates at which property cash flows are discounted (Section 4.3).[23]

6 Conclusion

We began with measurement. Housing returns have an income component and a capital component, and either one studied alone misleads. A high yield may mean a location with a high expected return or one where rents are expected to grow slowly, and the yield alone does not distinguish the two. Price appreciation alone leaves out the income return, which is most of the total return in our data and, in the simplest pricing models, where the risk premium appears. The excess total return combines both components and measures the premium.

With both components, we document two facts. Across metropolitan areas, cities more exposed to the national apartment market earn higher average excess returns. Within metropolitan areas, excess returns rise with distance from the main job center. In the model, both facts reflect compensation for exposure to aggregate risk. A location's exposure is its city's wage exposure to aggregate shocks, amplified by operating leverage. Costs follow the wage but rents fall with distance, so peripheral margins are thin and net operating income swings more. Other sources of risk exposure, such as industry mix, population composition shifts, and climate, could be studied the same way.

The main open question is which location characteristics carry a premium. Households buy a bundle of job access, schools, safety, retail and transit. Each piece has its own value process and is capitalized into rents and prices, so the bundle is hard to unpack. The model compresses it into a wage, a commute and one amenity, and Census 2000 neighborhood characteristics account for part of the within-city gradient. Identification work using changes in transit, schools or amenities would show which pieces carry premia and separate operating leverage from neighborhood effects.

The model is frictionless, and natural extensions add moral hazard in management, investment and use, illiquidity, and heterogeneity among properties and owners. Our contribution is standard measurement paired with a tractable structural model that can incorporate them. Its restrictions can be taken to single-family rentals and to owner-occupied housing, where part of the return is an implicit rent, the amount an owner would pay for the bundle.

[23] A property's yield moves by its yield times the percentage change in its value, so a temporary rise in the price of risk lowers the value by much more than it raises the yield. Appendix B.5 shows that the higher expected return of high-leverage properties then appears mainly as expected appreciation while the price of risk reverts.

References

Ahlfeldt, G. M., Heblich, S., and Seidel, T. (2023). Micro-geographic property price and rent indices. *Regional Science and Urban Economics*, 98:103836.

Albouy, D., Ehrlich, G., and Liu, Y. (2016). Housing demand, cost-of-living inequality, and the affordability crisis. Working Paper 22816, National Bureau of Economic Research.

Alonso, W. (1964). *Location and Land Use*. Harvard University Press, Cambridge, Massachusetts.

Amaral, F., Dohmen, M., Kohl, S., and Schularick, M. (2025). Superstar returns? spatial heterogeneity in returns to housing. *The Journal of Finance*, 80(5):3057–3094.

Ang, A. and Liu, J. (2007). Risk, return, and dividends. *Journal of Financial Economics*, 85(1):1–38.

Bond, S. A., Hwang, S., and Marcato, G. (2012). Commercial real estate returns: An anatomy of smoothing in asset and index returns. *Real Estate Economics*, 40(4):637–661.

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

Campbell, S. D., Davis, M. A., Gallin, J., and Martin, R. F. (2009). What moves housing markets: A variance decomposition of the rent–price ratio. *Journal of Urban Economics*, 66(2):90–102.

Cannon, S., Miller, N. G., and Pandher, G. S. (2006). Risk and return in the us housing market: A cross-sectional asset-pricing approach. *Real Estate Economics*, 34(4):519–552.

Case, K., Cotter, J., and Gabriel, S. (2011). Housing risk and return: Evidence from a housing asset-pricing model. *The Journal of Portfolio Management*, 37(5):89–109.

Clark, S. and Lomax, N. (2020). Rent/price ratio for english housing sub-markets using matched sales and rental data. *Area*, 52(1):136–147.

Cochrane, J. H. (1996). A cross-sectional test of an investment-based asset pricing model. *Journal of Political Economy*, 104(3):572–621.

Couts, S. J., Gonçalves, A. S., and Rossi, A. (2024). Unsmoothing returns of illiquid funds. *The Review of Financial Studies*, 37(7):2110–2155.

Crone, T. M. and Voith, R. P. (1999). Risk and return within the single-family housing market. *Real Estate Economics*, 27(1):63–78.

Damen, S., Korevaar, M., and Van Nieuwerburgh, S. (2026). An alpha in affordable housing? Working paper; SSRN 5121139; earlier version circulated as NBER Working Paper 33470.

Davis, M. A. and Ortalo-Magné, F. (2011). Household expenditures, wages, rents. *Review of Economic Dynamics*, 14(2):248–261.

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

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

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

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

Fama, E. F. and MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636.

Favilukis, J., Ludvigson, S. C., and Van Nieuwerburgh, S. (2017). The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy*, 125(1):140–223.

French, K. R. (2026). Data library. Kenneth R. French Data Library. Available at https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Geltner, D. (1991). Smoothing in appraisal-based returns. *Journal of Real Estate Finance and Economics*, 4(3):327–345.

Geltner, D. M. (1993). Estimating market values from appraised values without assuming an efficient market. *Journal of Real Estate Research*, 8(3):325–345.

Giacoletti, M. (2021). Idiosyncratic risk in housing markets. *Review of Financial Studies*, 34(8):3695–3741.

Giuliano, G. and Small, K. A. (1991). Subcenters in the Los Angeles region. *Regional Science and Urban Economics*, 21(2):163–182.

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

Halket, J., Loewenstein, L., and Willen, P. S. (2023). The cross-section of housing returns. Working paper, Texas A&M University.

Han, L. (2013). Understanding the puzzling risk-return relationship for housing. *Review of Financial Studies*, 26(4):877–928.

Himmelberg, C., Mayer, C., and Sinai, T. (2005). Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives*, 19(4):67–92.

Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., and Taylor, A. M. (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics*, 134(3):1225–1298.

Landvoigt, T., Piazzesi, M., and Schneider, M. (2015). The housing market(s) of San Diego. *American Economic Review*, 105(4):1371–1407.

Lettau, M. and Ludvigson, S. C. (2001a). Consumption, aggregate wealth, and expected stock returns. *The Journal of Finance*, 56(3):815–849.

Lettau, M. and Ludvigson, S. C. (2001b). Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying. *Journal of Political Economy*, 109(6):1238–1287.

Levy, A. (2021). Housing policy with home-biased landlords: Evidence from french rental markets. Working paper.

Mills, E. (1967). An aggregative model of resource allocation in a metropolitan area. *American Economic Review*, 57(2):197–210.

Muth, R. (1969). *Cities and Housing*. University of Chicago Press, Chicago.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.

Novy-Marx, R. (2011). Operating leverage. *Review of Finance*, 15(1):103–134.

Piazzesi, M. and Schneider, M. (2016). Housing and macroeconomics. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 1547–1640. Elsevier.

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

Quigley, J. M. and Raphael, S. (2004). Is housing unaffordable? why isn't it more affordable? *Journal of Economic Perspectives*, 18(1):191–214.

Roback, J. (1982). Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6):1257–1278.

Rosen, S. (1979). Wage-based indexes of urban quality of life. In Mieszkowski, P. and Straszheim, M., editors, *Current Issues in Urban Economics*, pages 74–104. Johns Hopkins University Press, Baltimore.

Sagi, J. S. (2021). Asset-level risk and return in real estate investments. *Review of Financial Studies*, 34(8):3647–3694.

Sinai, T. and Souleles, N. S. (2005). Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics*, 120(2):763–789.

Van Nieuwerburgh, S. and Weill, P.-O. (2010). Why has house price dispersion gone up? *Review of Economic Studies*, 77(4):1567–1606.

Vasicek, O. A. (1973). A note on using cross-sectional information in bayesian estimation of security betas. *The Journal of Finance*, 28(5):1233–1239.

Appendix

Appendix A reports supplementary results on the data and the empirical facts of Sections 2 and 3. Appendix B gives the derivations behind the model of Section 4, including the real-options extension. Appendix C reports material behind the quantitative assessment of Section 5. A separate Online Appendix develops the closed-city equilibrium.

A Data and Empirical Facts

A.1 Distance measure construction

The baseline distance measure uses employment concentrations rather than a hand-curated downtown ZIP code. We use the 2022 LEHD Origin-Destination Employment Statistics Workplace Area Characteristics files, aggregate workplace employment to ZIP Code Tabulation Areas (ZCTAs), and compute employment density using Census TIGER/Line land areas. We then identify centers with a modified Giuliano-Small procedure:

1. ZCTAs with employment density at or above the 90th percentile across the ZCTAs of the 25 metropolitan areas for which we build centers, approximately 5,795 jobs per square mile, are flagged as high-density.
2. High-density ZCTAs within 3 miles of each other are grouped into initial clusters using DBSCAN.
3. Initial clusters whose centroids are within 7 miles of each other are merged using single-linkage hierarchical clustering. This step prevents contiguous downtown areas from being split into small fragments while still allowing separate centers in twin-city or polycentric regions.
4. Clusters with at least 25,000 jobs are retained. If no cluster in a CBSA satisfies this employment threshold, the largest cluster in that CBSA is retained, so the filter never drops a metropolitan area mechanically.
5. Each surviving center is assigned the employment-weighted centroid of its member ZCTAs.

The primary employment center is the surviving cluster with the largest total employment in the CBSA. The baseline distance variable, `distance_primary`, is the Haversine distance from each ZIP centroid to that primary center. On the main sample of 296 ZIPs in 22 CBSAs, its Fama–MacBeth slope is 2.31 basis points per mile per quarter (s.e. 0.80). For comparison, we also measure distance to a hand-curated downtown ZIP code in each metropolitan area. Comparing the measures requires the

193 ZIPs in 20 CBSAs for which all three distances are available and no greater than 30 miles. On this common sample, the primary-center and hand-curated downtown ZIP code slopes are 3.16 (s.e. 0.74) and 3.05 (s.e. 0.70), respectively, and the two distance measures have a correlation of 0.96.

The same construction also produces a multi-center distance measure for polycentric metros. The construction identifies 64 employment centers across 25 CBSAs, with 16 CBSAs having more than one center. Examples include Dallas–Fort Worth, Miami–Fort Lauderdale, San Francisco–Oakland, Seattle–Tacoma, and Tampa–St. Petersburg. The alternative variable, `distance_min`, measures distance from each ZIP to the nearest surviving center in its CBSA, whether primary or secondary. On the common comparison sample, its Fama–MacBeth slope is 4.12 basis points per mile per quarter (s.e. 1.11), against 3.16 (s.e. 0.74) for primary-center distance. Nearest-center distances are shorter and vary less across ZIP codes, so the per-mile slopes are not directly comparable. Per standard deviation of distance, the slopes are 20.9 and 22.3 basis points per quarter. We use distance from the primary center as the baseline because distance to a single center is the standard measure in urban theory and in empirical work. The nearest-center measure is a natural and data-driven alternative, but it depends more on the radii and employment thresholds we choose, which determine how many secondary centers each metropolitan area has.

A.2 Summary statistics for the within-city sample

Tables A1 and A2 report summary statistics for the within-city sample, aggregated to the CBSA and ZIP levels.

Table A1: Descriptive Statistics – CBSA aggregation

Variable	mean	sd	median	min	max
Total Return	1.83291	2.39982	1.81000	-4.1470	7.5500
Income Return	1.22744	0.24124	1.20667	0.7300	1.6700
Capital Return	0.60400	2.34711	0.53933	-5.1440	6.4000
Risk Free Rate	0.37551	0.44464	0.17000	0.0000	1.5700
Excess Return	1.45740	2.49477	1.42333	-5.2070	7.5500

Note: values are quarterly returns in percentage points. This differs from Table 1 due to the sample restriction to the same CBSA and ZIP-quarters used in the within-city analysis.

Table A2: Descriptive Statistics – ZIP code aggregation

Variable	mean	sd	median	min	max
Total Return	1.6378	2.7908	1.5500	-4.1470	7.5500
Income Return	1.1455	0.2542	1.1100	0.7300	1.6700
Capital Return	0.4967	2.7301	0.3500	-5.1440	6.4000
Excess Return	1.2917	2.9380	1.2500	-5.4770	7.5500

Note: values are quarterly returns in percentage points.

A.3 Cross-city Fama–MacBeth regressions: additional results

Table A3 reports Fama–MacBeth regressions on finer beta-sorted portfolios and on individual CBSAs with winsorized betas or without Baton Rouge. The price of risk falls as CBSAs are grouped more finely, from 2.15 with three portfolios to 1.80, 1.12 and 1.01 with four, six and eight, while the intercept rises. Winsorizing the betas at the 5th and 95th percentiles across CBSAs gives 1.13 (s.e. 0.57). Without Baton Rouge, whose beta is the least precise, a one-standard-deviation difference in beta is worth 0.18 pp per quarter under unadjusted, winsorized and shrunk betas.

Table A3: Cross-city Fama–MacBeth regressions: additional results

	Price of risk	Intercept	One-s.d. effect	Test assets
Tercile portfolios (Table 2)	2.15 (0.76)	0.22 (0.40)		3
Four portfolios	1.80 (0.71)	0.46 (0.36)		4
Six portfolios	1.12 (0.71)	0.94 (0.37)		6
Eight portfolios	1.01 (0.61)	1.03 (0.30)		8
CBSAs, winsorized betas	1.13 (0.57)	0.93 (0.28)	0.14 (0.07)	48
CBSAs without Baton Rouge:				
unadjusted betas	1.36 (0.55)	0.79 (0.27)	0.18 (0.07)	47
winsorized betas	1.44 (0.59)	0.73 (0.29)	0.18 (0.07)	47
shrunk betas	1.97 (0.77)	0.36 (0.37)	0.18 (0.07)	47

Notes: Fama–MacBeth regressions with an intercept, in percentage points per quarter, over 96 quarters from 2000Q1 to 2023Q4. Portfolios sort the CBSAs on full-sample unadjusted beta and use post-ranking betas, as in Table 2. Winsorized betas cap the unadjusted betas at the 5th and 95th percentiles across CBSAs. Without Baton Rouge, all three beta measures are recomputed on the remaining 47 CBSAs. The one-s.d. effect multiplies the price of risk by the cross-sectional standard deviation of the beta measure. Parentheses contain Newey–West standard errors with four lags.

The Zillow betas in Panel B of Table 2 are loadings of monthly metropolitan Zillow Home Value

Index growth on national Zillow growth. Zillow reports appreciation rather than total returns, and it is aggregated at the MSA level. Missing months are left missing. Before the regression, monthly MSA growth is capped between $\min\{4g_t^{US}, -2g_t^{US}\}$ and $\max\{4g_t^{US}, -2g_t^{US}\}$, where g_t^{US} is national monthly growth. The MSA-to-CBSA crosswalk covers all 48 CBSAs in the pricing sample. The second-stage regression still uses NCREIF excess total returns.

A.4 ZIP-beta pricing

Table A4 reports the full Fama–MacBeth estimates for the 66 ZIP codes with at least 40 quarters of returns, the sample of the last four columns of Table 6. Distance enters either continuously or through indicators for distance quartiles. The gradient is about 4.1 basis points per mile whichever beta enters, and the quartile coefficients rise from 0.15–0.20 pp per quarter in the second quartile to 0.70–0.76 pp in the fourth. Without distance, the national beta has a positive coefficient, 0.64 (s.e. 0.43), and the local beta a negative one, -0.65 (s.e. 0.43). Once distance enters, the beta coefficients are imprecise.[24]

[24] In beta form, the model implies that expected excess returns are proportional to betas, with no intercept. Imposing this restriction leaves the distance gradient at 4.6 to 5.0 basis points per mile. Measured ZIP betas cluster near one, so the restricted beta coefficients mostly absorb the average excess return, and we report the unrestricted estimates.

Table A4: ZIP-beta pricing and distance

	Beta controls			
	None	National	Local	Both
No distance				
National beta		0.6392	–	-0.5656
		(0.4331)		(0.4984)
Local beta		–	-0.6456	-0.7361
			(0.4290)	(0.4263)
Continuous distance				
National beta	–	0.5929	–	-0.3986
		(0.4428)		(0.5110)
Local beta	–	–	-0.3884	-0.3638
			(0.4190)	(0.4205)
Distance (miles)	0.0406	0.0409	0.0409	0.0406
	(0.0107)	(0.0105)	(0.0107)	(0.0102)
Distance quartiles				
National beta	–	0.6120	–	-0.4071
		(0.4642)		(0.5424)
Local beta	–	–	-0.5239	-0.4584
			(0.4662)	(0.4578)
Quartile 2	0.2018	0.1545	0.1925	0.1451
	(0.1240)	(0.1337)	(0.1263)	(0.1283)
Quartile 3	0.4541	0.4491	0.4174	0.3913
	(0.1466)	(0.1419)	(0.1530)	(0.1601)
Quartile 4	0.7505	0.7003	0.7581	0.7041
	(0.1707)	(0.1611)	(0.1660)	(0.1525)

Notes: Fama–MacBeth averages of quarterly cross-sectional slopes, in percentage points per quarter (distance per mile), with Newey–West standard errors with four lags in parentheses. Every quarterly cross-section includes an intercept, whose estimates are not reported. The sample is the 66 ZIP codes in 18 CBSAs within 30 miles of the primary employment center with at least 40 quarters of returns, over 94 quarters. The national beta is the ZIP code's loading on the national NCREIF apartment index and the local beta its loading on its own CBSA apartment index, and in the "Both" column the two betas come from a single first-stage regression of the ZIP return on both indices. Quartile indicators are relative to the closest distance quartile. The quartile specifications drop 2001Q1, which has no ZIP code in the closest quartile ($T = 93$). Returns are winsorized at 5/95 after the CBSA coverage screen.

A.5 Distance gradient with a separate intercept for each CBSA

The distance gradient remains positive when each quarterly regression includes a separate intercept for every CBSA, so that the slope uses only differences in distance within a metropolitan area. Table A5 reports 1.69 basis points per mile per quarter (s.e. 0.75) for the 296 ZIP codes of the main sample and 1.66 (s.e. 0.76) for the 288 of them with complete Census 2000 data.

Table A5: Distance gradient with a separate intercept for each CBSA

	Main sample	Complete Census 2000 data
Distance (bp/mile/qtr)	1.69	1.66
	(0.75)	(0.76)
ZIPs / CBSAs	296 / 22	288 / 22
ZIP-quarter observations	7,705	7,582
Quarterly cross-sections	96	96

Notes: Fama–MacBeth averages of quarterly cross-sectional slopes of excess total returns on distance, with ZIP codes weighted equally. Every quarterly regression includes a separate intercept for each CBSA present that quarter. The main sample has all ZIP codes within 30 miles of the primary employment center, with returns winsorized at 5/95 after the CBSA coverage screen. The second column keeps the ZIP codes with complete Census 2000 data and does not add Census controls. Parentheses contain Newey–West standard errors with four lags. Conventional standard errors are 0.76 and 0.77.

A.6 Local betas and distance

Table A6 regresses each ZIP code's local beta on its distance from the employment center, with CBSA fixed effects, for the 66 ZIP codes with at least 40 quarters of returns. Only 14 CBSAs contribute more than one ZIP code, so the table reports wild cluster bootstrap p-values alongside the clustered standard errors. The slope is positive with distance to the primary center and to the hand-curated downtown ZIP code, and negative but imprecise with distance to the nearest center. It is significant at the 5 percent level only with the hand-curated downtown ZIP code.

Table A6: Local betas and distance within CBSAs

Distance measure	Slope per 10 miles	s.e.	Wild bootstrap p	ZIP codes	CBSAs
Primary employment center	0.045	(0.022)	0.13	66	18
Hand-curated downtown ZIP code	0.052	(0.016)	0.05	57	16
Nearest employment center	-0.042	(0.050)	0.42	66	18

Notes: Cross-sectional regressions of each ZIP code's local beta, its loading on the return of its own CBSA apartment index, on distance, with CBSA fixed effects. ZIP codes within 30 miles of the primary employment center with at least 40 quarters of returns. Standard errors are clustered by CBSA. Wild bootstrap p-values impose the null and enumerate all Rademacher sign vectors across CBSAs. With primary-center distance, dropping the most influential ZIP code gives 0.018 (0.029), betas estimated from desmoothed returns give 0.010 (0.035), and betas on the national index give 0.011 (0.034).

A.7 Census 2000 neighborhood controls

Table A7 augments the beta-controlled within-city regression with neighborhood characteristics measured in Census 2000, at the start of the return panel: median income, college share, population density on land area, the non-Hispanic White share, and the poverty rate. Of the 66 ZIP codes with a beta, one (60654, in Chicago) lacks college and poverty data, leaving 65. Each quarter's cross-section includes an intercept, and the estimates average the 74 quarters in which at least 16 of these ZIP codes report.

Table A7: Fama-MacBeth Regressions: Distance Gradient with Census 2000 Controls

	(1)	(2)	(3)	(4)	(5)
Distance (miles)	0.0432***	0.0252***	0.0254***	0.0272***	
	(0.0121)	(0.0082)	(0.0083)	(0.0088)	
Local beta	−0.1837	−0.2481	−0.3216	−0.6570	−0.7619
	(0.4712)	(0.5713)	(0.5983)	(0.5851)	(0.5882)
Log median income		0.0573	0.0622	0.3106***	0.3803***
		(0.0484)	(0.0481)	(0.0865)	(0.0961)
College share		−0.0698	−0.0707	−0.1405	−0.2344**
		(0.0652)	(0.0807)	(0.0880)	(0.1090)
Log pop. density		−0.1676*	−0.1615*	−0.1514*	−0.1787**
		(0.0854)	(0.0853)	(0.0867)	(0.0895)
% NH White			−0.0010	0.0398	0.0330
			(0.0592)	(0.0667)	(0.0647)
Poverty rate				0.2468***	0.1968**
				(0.0893)	(0.0835)
Intercept	1.4890***	1.7030***	1.7657***	2.0734***	2.3588***
	(0.2025)	(0.3131)	(0.3271)	(0.3144)	(0.3356)
Avg. N per quarter	49	49	49	49	49
T	74	74	74	74	74

Notes: Fama-MacBeth cross-sectional regressions of ZIP-level quarterly excess total returns (pp) on distance to the primary employment center, the local beta (the ZIP code's loading on its own CBSA apartment index, as in Table 6), and Census 2000 controls. Every quarterly cross-section includes an intercept. The sample is the 65 ZIP codes with at least 40 quarters of returns and complete Census data; quarters with fewer than 16 ZIPs are excluded ($T = 74$). Parentheses contain Newey–West standard errors with four lags; conventional standard errors of the distance coefficient are 0.0073, 0.0075, 0.0075 and 0.0074. Returns winsorized at 5%. Distance capped at 30 miles. Census 2000 controls from Decennial Census SF1 (race/ethnicity) and SF3 (income, education, poverty) at ZCTA level, standardized over the estimation sample. Population density uses Census 2000 ZCTA land area (Gazetteer). Race and ethnicity enter as the non-Hispanic White share. The share of housing built before 1970 is excluded because, in cities that grow outward, it is largely determined by distance. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$, from the Newey–West t.

A.8 Desmoothing Robustness

We apply the one-step AR(1) unsmoothing filter used in the real estate literature (Geltner, 1991, 1993; Couts et al., 2024). For consecutive observations, the filter replaces each ZIP's observed return r_t^{obs} with

$$r_t = \frac{r_t^{\text{obs}} - \hat{\rho} r_{t-1}^{\text{obs}}}{1 - \hat{\rho}}.$$

Here, ρ is the lag weight and $\alpha = 1 - \rho$ is the weight on the current return. We estimate ρ separately by ZIP, and the pooled specification takes a pair-count-weighted average of these estimates. We also report fixed $\alpha \in \{0.4, 0.5, 0.6, 0.8\}$. Together with observed returns, these choices produce the seven columns reported below. Table A8 reports diagnostics. The median ZIP-level AR(1) coefficient is 0.564. Both the pooled and ZIP-specific filters reduce autocorrelation sharply and increase return volatility.

More general unsmoothing methods allow for several lags and distinguish smoothing in the common and local components of returns (Bond et al., 2012; Couts et al., 2024). Our ZIP series are short and unbalanced, and each aggregates a changing set of properties. We therefore use the pooled AR(1) filter as our main desmoothing specification.

Table A8: One-Step AR(1) Unsmoothing Diagnostics

	Observed	Pooled α	ZIP-specific α
$\hat{\rho}$ distribution (ZIP AR(1), capped to [0, 0.95])			
N ZIPs with $\hat{\rho}$	241		
Share at floor ($\hat{\rho} = 0$)	0.066		
Share at cap ($\hat{\rho} = 0.95$)	0.004		
$\hat{\rho}$ (median [IQR])	0.564 [0.354, 0.683]		
$\hat{\alpha}$ (median [IQR])	0.436 [0.317, 0.646]		
Pooled $\hat{\alpha}$	0.484		
Volatility and autocorrelation (median [IQR] across ZIPs)			
$\sigma(r)$	0.027 [0.023, 0.032]	0.041 [0.036, 0.047]	0.045 [0.035, 0.056]
AR(1)	0.563 [0.354, 0.683]	0.006 [-0.192, 0.205]	-0.033 [-0.108, 0.035]

Table A9 reproduces two results across all seven specifications: the distance gradient and the $\omega^{NT} \times cay$ interaction. The distance gradient remains positive under every filter. The pooled-α estimate is 2.90 bp/mile (s.e. 0.88), while the ZIP-specific estimate is 2.32 (s.e. 0.97), against the 2.31 (s.e. 0.80) baseline. The $\omega^{NT} \times cay$ coefficient remains positive across the reported filters, but the pooled and ZIP-specific estimates are imprecise. The fixed $\alpha = 0.8$ estimate is 1.046 (s.e. 0.537).

Table A9: Main Results under One-Step AR(1) Unsmoothing

Specification	Baseline	Pooled α	ZIP α	$\alpha = 0.4$	$\alpha = 0.5$	$\alpha = 0.6$	$\alpha = 0.8$
Distance gradient (FM; bp/mi)	2.31	2.90	2.32	3.14	2.86	2.65	2.43
	(0.80)	(0.88)	(0.97)	(0.95)	(0.87)	(0.82)	(0.80)
$\omega^{NT} \times cay$ (predict r_{t+1})	1.140	0.737	1.205	0.574	0.763	0.888	1.046
	(0.442)	(1.042)	(0.907)	(1.348)	(0.996)	(0.775)	(0.537)

Notes: Coefficients with standard errors in parentheses. The distance regression uses the 296-ZIP primary-center sample over 96 quarters; the mean quarterly cross-section contains 80.3 ZIPs. Its standard errors are Newey–West with four lags. The $\omega^{NT} \times cay$ specification uses 4,372 ZIP-quarter observations, with standard errors two-way clustered by CBSA and date.

A.9 Pre-2020 Robustness

This subsection reports regression results for the pre-pandemic sample (excluding 2020 Q1 onward) described in Section 3.3. Table A10 compares the distance gradient, the $\omega^{NT} \times cay$ predictability result, and placebo tests across the full and pre-2020 samples.

The Fama–MacBeth distance gradient falls from 2.31 to 1.62 bp/mile (s.e. 0.78), so the gradient predates the pandemic and was larger during it. Rebuilding the expense characteristic on the pre-2020 extract gives a total-return coefficient of 1.038 (s.e. 0.404), compared with 1.140 (s.e. 0.442) in the full construction, and the capital-return coefficients are 0.972 and 1.074. Both $\omega^{NT} \times cay$ columns use returns through 2019Q2, so this comparison changes only the expense data used to build ω^{NT} and, with them, the ZIP codes that have it. The placebo interactions remain weak in both samples.

Table A10: Core Results: Full Sample vs. Pre-2020

	Full sample	Pre-2020
Panel A: Fama–MacBeth distance gradient (bp per mile per quarter)		
Continuous	2.31	1.62
	(0.80)	(0.78)
Separate intercept for each CBSA	1.69	1.08
	(0.75)	(0.74)
With both betas	4.06	2.42
	(1.02)	(0.99)
Panel B: $\omega^{NT} \times cay$ (predicts r_{t+1})		
Total return	1.140	1.038
	(0.442)	(0.404)
Capital return	1.074	0.972
	(0.440)	(0.403)
Panel C: Placebo interactions (total return)		
Distance $\times$ *cay*	−0.0068	−0.0066
	(0.0056)	(0.0056)
Market value $\times$ *cay*	9.67e-11	8.78e-11
	(1.31e-10)	(1.31e-10)
NOI $\times$ *cay*	4.16e-09	3.34e-09
	(1.35e-08)	(1.35e-08)
Property count $\times$ *cay*	0.0060	0.0051
	(0.0188)	(0.0186)

Notes: Coefficients with standard errors in parentheses. The full sample spans 2000Q1–2023Q4 (1991Q1–2019Q2 for the $\omega^{NT} \times cay$ specification, as in Section 5.4), and the pre-2020 sample ends in 2019Q4. Panel A runs every step of each specification on the shorter sample. Panel A uses Fama–MacBeth regressions with a quarterly intercept and Newey–West standard errors with four lags. The continuous row uses the 296-ZIP, 22-CBSA primary-center sample (252 ZIPs and 80 quarters pre-2020; mean cross-section 80.3 and 64.0 ZIPs). The second row adds a separate intercept for each CBSA in every quarter, as in Appendix A.5. The third row adds the national and local betas of Table 6, estimated jointly for the ZIP codes with at least 40 quarters of returns: 66 ZIPs over 94 quarters in the full sample, and 39 ZIPs over 76 quarters pre-2020. Panels B and C report the specifications of Section 5.4 with their own standard errors. The $\omega^{NT} \times cay$ specification uses $N = 4,372$ (full) and $N = 4,187$ (pre-2020) ZIP-quarter observations. Both columns of Panels B and C use returns through 2019Q2. The pre-2020 column rebuilds ω^{NT} from expense data before 2020, which also changes the ZIP codes that have it.

B Model Derivations

In this appendix, we present the derivations behind the model expressions used in the main text.

B.1 Rent determination

We derive the equilibrium rent and its stochastic dynamics from the worker's utility maximization and spatial indifference.

Setup and worker problem Workers living in city k at distance d from the CBD at time t have utility given by

$$U(c, s, a_k) = c^{\alpha} s^{1-\alpha} a_k, \qquad \alpha \in (0, 1),$$

and face the following budget constraint with commuting-adjusted income

$$c + r_{k,d,t}\, s = \kappa_{k,d}\, w_{k,t}, \qquad \kappa_{k,d} = \max\{1 - \tau_k d, 0\}.$$

They can relocate freely across space, including to an outside option location k_0 with associated utility Ψ_t. Wages and the outside option value follow, respectively,

$$\begin{aligned} \frac{dw_{k,t}}{w_{k,t}} &= \mu_{w,k}\, dt + \sigma_{w,k}\, dB_t + \varsigma_{w,k}\, dB_t^k, \\ \frac{d\Psi_t}{\Psi_t} &= \mu_{\Psi}\, dt + \sigma_{\Psi}\, dB_t, \end{aligned}$$

where B_t is an aggregate Brownian motion and B_t^k is an idiosyncratic Brownian motion for city k, independent of B_t. Thus $dB_t\, dB_t^k = 0$ and $(dB_t)^2 = (dB_t^k)^2 = dt$.

Indirect utility and spatial indifference With Cobb–Douglas preferences, the worker spends shares α and $1 - \alpha$ of commuting-adjusted income on consumption and housing: $c_t = \alpha\, \kappa_{k,d}\, w_{k,t}$ and $r_{k,d,t} s_t = (1 - \alpha)\, \kappa_{k,d}\, w_{k,t}$. The worker's indirect utility is

$$\begin{aligned} U_{k,d,t}^{*} &= (\alpha\, \kappa_{k,d}\, w_{k,t})^{\alpha} \left((1 - \alpha) \frac{\kappa_{k,d}\, w_{k,t}}{r_{k,d,t}} \right)^{1-\alpha} a_k \\ &= \alpha^{\alpha} (1 - \alpha)^{1-\alpha}\, \kappa_{k,d}\, w_{k,t}\, r_{k,d,t}^{-(1-\alpha)}\, a_k. \end{aligned}$$

Free mobility and spatial indifference equate $U_{k,d,t}^{*}$ to the outside option Ψ_t:

$$\alpha^{\alpha} (1 - \alpha)^{1-\alpha}\, \kappa_{k,d}\, w_{k,t}\, r_{k,d,t}^{-(1-\alpha)}\, a_k = \Psi_t.$$

Hence, solving for $r_{k,d,t}$, we get

$$r_{k,d,t} = \zeta\, (\kappa_{k,d} w_{k,t})^{\frac{1}{1-\alpha}}\, a_k^{\frac{1}{1-\alpha}}\, \Psi_t^{-\frac{1}{1-\alpha}}, \qquad \zeta \equiv \left[\alpha^{\alpha} (1 - \alpha)^{1-\alpha} \right]^{\frac{1}{1-\alpha}}.$$

Taking logs gives

$$\ln r_{k,d,t} = \ln \zeta + \frac{1}{1 - \alpha} \ln (\kappa_{k,d} w_{k,t}) + \frac{1}{1 - \alpha} \ln a_k - \frac{1}{1 - \alpha} \ln \Psi_t,$$

as shown in the main text.

Rent dynamics By Itô's lemma, the logs of the wage and the outside option follow

$$d\ln w_{k,t} = \left(\mu_{w,k} - \tfrac{1}{2}(\sigma_{w,k}^2 + \varsigma_{w,k}^2)\right)dt + \sigma_{w,k}\, dB_t + \varsigma_{w,k}\, dB_t^k,$$
$$d\ln \Psi_t = \left(\mu_\Psi - \tfrac{1}{2}\sigma_\Psi^2\right)dt + \sigma_\Psi\, dB_t.$$

Because $\kappa_{k,d}$ and a_k are constant, the log rent equation gives $d\ln r_{k,d,t} = \frac{1}{1-\alpha}\big(d\ln w_{k,t} - d\ln \Psi_t\big)$. Substituting, we find

$$\begin{aligned} d\ln r_{k,d,t} &= \tfrac{1}{1-\alpha}\Bigg[\left(\mu_{w,k} - \tfrac{1}{2}(\sigma_{w,k}^2 + \varsigma_{w,k}^2)\right)dt + \sigma_{w,k}\, dB_t + \varsigma_{w,k}\, dB_t^k \\ &\qquad - \left(\mu_\Psi - \tfrac{1}{2}\sigma_\Psi^2\right)dt - \sigma_\Psi\, dB_t\Bigg] \\ &= \left(\tfrac{1}{1-\alpha}\left(\mu_{w,k} - \mu_\Psi - \tfrac{1}{2}(\sigma_{w,k}^2 + \varsigma_{w,k}^2 - \sigma_\Psi^2)\right)\right)dt \\ &\qquad + \tfrac{1}{1-\alpha}(\sigma_{w,k} - \sigma_\Psi)dB_t + \tfrac{1}{1-\alpha}\varsigma_{w,k}\, dB_t^k. \end{aligned}$$

Since $r_{k,d,t} = \exp(\ln r_{k,d,t})$, we have

$$dr_{k,d,t} = r_{k,d,t}(d\ln r_{k,d,t}) + \tfrac{1}{2}\, r_{k,d,t}\, d\langle \ln r_{k,d,t}, \ln r_{k,d,t}\rangle$$
$$\frac{dr_{k,d,t}}{r_{k,d,t}} = d\ln r_{k,d,t} + \tfrac{1}{2}\, d\langle \ln r_{k,d,t}, \ln r_{k,d,t}\rangle,$$

and

$$d\langle \ln r_{k,d,t}, \ln r_{k,d,t}\rangle = \left(\left(\tfrac{1}{1-\alpha}(\sigma_{w,k} - \sigma_\Psi)\right)^2 + \left(\tfrac{1}{1-\alpha}(\varsigma_{w,k})\right)^2\right)dt.$$

Therefore, collecting terms, we get

$$\frac{dr_{k,d,t}}{r_{k,d,t}} = \mu_{r,k}\, dt + \sigma_{r,k}\, dB_t + \varsigma_{r,k}\, dB_t^k,$$

with

$$\sigma_{r,k} = \frac{\sigma_{w,k} - \sigma_\Psi}{1-\alpha}, \qquad \varsigma_{r,k} = \frac{\varsigma_{w,k}}{1-\alpha},$$

and drift

$$\begin{aligned} \mu_{r,k} = &\left(\frac{1}{1-\alpha}\right)\mu_{w,k} - \left(\frac{1}{1-\alpha}\right)\mu_\Psi \\ &+ \left(\frac{1}{2}\frac{\alpha}{(1-\alpha)^2}\right)\sigma_{w,k}^2 + \left(\frac{1}{2}\frac{\alpha}{(1-\alpha)^2}\right)\varsigma_{w,k}^2 \\ &+ \left(\frac{1}{2}\frac{2-\alpha}{(1-\alpha)^2}\right)\sigma_\Psi^2 - \left(\frac{1}{(1-\alpha)^2}\right)\sigma_{w,k}\sigma_\Psi. \end{aligned}$$

Hence, rent inherits aggregate and local (city) risk from wages and the outside option, with elasticities scaled by $\frac{1}{1-\alpha}$, and a drift that includes the usual Itô correction terms and the covariance between $w_{k,t}$ and Ψ_t through the common aggregate shock.

B.2 Price determination

This subsection prices a claim to gross rent under the stochastic discount factor used in the main text. The proportional-price guess delivers the risk-adjusted Gordon formula.

For any asset with price Q_t and dividend (cash flow) y_t, it holds that

$$0 \;=\; \mathbb{E}_t\big[d(\Lambda_t^k Q_t)\big] + \Lambda_t^k\, y_t\, dt,$$

where we use the (potentially city-specific) Stochastic Discount Factor (SDF) as given in the main text:

$$\frac{d\Lambda_t^k}{\Lambda_t^k} \;=\; -r_f\, dt \;-\; \lambda^k\, dB_t \;-\; \xi^k\, dB_t^k.$$

A claim on rent is an asset that pays the instantaneous cash flow $y_t = r_{k,d,t}$ and has price $P_{k,d,t}$. Conjecture a price proportional to rent,

$$P_{k,d,t} \;=\; \theta\, r_{k,d,t}, \qquad \theta \text{ constant},$$

and show it verifies the fundamental pricing equation.

Starting with the inside of the expectation, using the product rule we have

$$d(\Lambda_t^k P_{k,d,t}) \;=\; \Lambda_t^k\, dP_{k,d,t} \;+\; P_{k,d,t}\, d\Lambda_t^k \;+\; d\Lambda_t^k\, dP_{k,d,t}.$$

With $P_{k,d,t} = \theta r_{k,d,t}$, it holds that $dP_{k,d,t} = \theta\, d\, r_{k,d,t}$, so

$$d(\Lambda_t^k P_{k,d,t}) \;=\; \theta\Big(\Lambda_t^k\, d\, r_{k,d,t} \;+\; r_{k,d,t}\, d\Lambda_t^k \;+\; d\langle \Lambda_t^k\,, r_{k,d,t}\rangle\Big).$$

Given the independence of the aggregate and the city-specific Brownian motions, the cross term is simply

$$d\langle \Lambda_t^k\,, r_{k,d,t}\rangle = \Lambda_t^k r_{k,d,t}\big(-\lambda^k\sigma_{r,k} - \xi^k\varsigma_{r,k}\big)\, dt.$$

Combining these terms with the original SDEs for $d\, r_{k,d,t}$ and $d\Lambda_t^k$, we get

$$d(\Lambda_t^k P_{k,d,t}) \;=\; \theta\Big(\mu_{\Lambda,r}\, dt \;+\; \Lambda_t^k r_{k,d,t}(\sigma_{r,k} - \lambda^k) dB_t \;+\; \Lambda_t^k r_{k,d,t}(\varsigma_{r,k} - \xi^k) dB_t^k\Big),$$

where the drift term $\mu_{\Lambda,r}$ is given by

$$\mu_{\Lambda,r} = \Lambda_t^k r_{k,d,t}\Big(\mu_{r,k} - r_f - \lambda^k\sigma_{r,k} - \xi^k\varsigma_{r,k}\Big).$$

Because B_t and B_t^k are standard Brownian motions, only the drift term survives when we take the conditional expectation:

$$\mathbb{E}_t\big[d(\Lambda_t^k P_{k,d,t})\big] \;=\; \theta\Lambda_t^k r_{k,d,t}\Big(\mu_{r,k} - r_f - \lambda^k\sigma_{r,k} - \xi^k\varsigma_{r,k}\Big)\, dt.$$

Substituting into the fundamental pricing equation, we get

$$0 \;=\; \mathbb{E}_t\big[d(\Lambda_t^k P_{k,d,t})\big] + \Lambda_t^k r_{k,d,t}\,dt \;=\; \Lambda_t^k r_{k,d,t}\left[\theta\big(\mu_{r,k} - r_f - \lambda^k\sigma_{r,k} - \xi^k\varsigma_{r,k}\big) + 1\right]dt.$$

Canceling $\Lambda_t r_{k,d,t} dt \neq 0$,

$$\theta\big(\mu_{r,k} - r_f - \lambda^k\sigma_{r,k} - \xi^k\varsigma_{r,k}\big) + 1 \;=\; 0,$$

hence

$$\theta \;=\; \frac{1}{r_f + \lambda^k\sigma_{r,k} + \xi^k\varsigma_{r,k} - \mu_{r,k}}.$$

Thus, the guessed functional form $P_{k,d,t} = \theta\, r_{k,d,t}$ satisfies the fundamental pricing equation as claimed. The formula requires a positive risk-adjusted discount rate, $r_f + \lambda^k\sigma_{r,k} + \xi^k\varsigma_{r,k} - \mu_{r,k} > 0$, and selects the fundamental solution with no bubble component, as in Appendix B.3. Substituting θ back, we find the price formula shown in the text:

$$P_{k,d,t} \;=\; \frac{r_{k,d,t}}{r_f + \lambda^k\sigma_{r,k} + \xi^k\varsigma_{r,k} - \mu_{r,k}}.$$

More generally, consider any cash flow y_t with dynamics

$$\frac{d\,y_t}{y_t} = \mu_y\,dt + \sigma_y\,dB_t + \varsigma_y\,dB_t^k.$$

Under the same SDF and the same argument applied above, the price of a perpetual claim that pays y_t is

$$P_t = \frac{y_t}{\rho_y},$$

where ρ_y is the risk-adjusted discount rate given by

$$\rho_y = r_f + \lambda^k\sigma_y + \xi^k\varsigma_y - \mu_y.$$

B.3 Price determination with running costs

This subsection derives the wage-normalized pricing ODE for a claim to net operating income. The normalization works because net operating income, $r - \delta w$, is homogeneous of degree one in rent and the wage, which leaves the rent-to-wage ratio as the only state variable.

Set up the environment Denote the rent-to-wage ratio as

$$x_{k,d,t} = \frac{r_{k,d,t}}{w_{k,t}}.$$

Net operating income can then be written as

$$\text{NOI}_{k,d,t} = w_{k,t}(x_{k,d,t} - \delta).$$

The price of a developed piece of land, $P_{k,d,t}$, satisfies

$$0 = E_t\left[d\left(\Lambda_t^k P_{k,d,t}\right)\right] + \Lambda_t^k \,\text{NOI}_{k,d,t}\, dt.$$

Conjecture that the price $P_{k,d,t}$ is a Markov function and is homogeneous of degree 1 in (r, w), satisfying[25]

$$P_k\left(r_{k,d,t}, w_{k,t}\right) = w_{k,t} p_k\left(x_{k,d,t}\right),$$

where $p_k(z) = P_k(z, 1)$.

Wage-denominated SDF Let $\Upsilon_t^k = \Lambda_t^k w_{k,t}$ be a wage denominated stochastic discount factor. The Itô product rule gives

$$\begin{aligned} d\Upsilon_t^k &= \Lambda_t^k dw_{k,t} + w_{k,t} d\Lambda_t^k + d\langle w_k, \Lambda^k\rangle_t \\ &= \Lambda_t^k w_{k,t}\left[\mu_{w,k}dt + \sigma_{w,k}dB_t + \varsigma_{w,k}dB_t^k\right] + w_{k,t}\Lambda_t^k\left[-r_f dt - \lambda^k dB_t - \xi^k dB_t^k\right] \\ &\quad + \Lambda_t^k w_{k,t}\left[-\lambda^k\sigma_{w,k} - \xi^k\varsigma_{w,k}\right]dt. \end{aligned}$$

Rearranging terms, we get

$$\begin{aligned} \nu_{\Upsilon,k} &= r_f - \mu_{w,k} + \lambda^k\sigma_{w,k} + \xi^k\varsigma_{w,k}, \\ \lambda_{\Upsilon,k} &= \lambda^k - \sigma_{w,k}, \\ \xi_{\Upsilon,k} &= \xi^k - \varsigma_{w,k}, \end{aligned}$$

and

$$\frac{d\Upsilon_t^k}{\Upsilon_t^k} = -\nu_{\Upsilon,k}\,dt - \lambda_{\Upsilon,k}\,dB_t - \xi_{\Upsilon,k}\,dB_t^k.$$

Substituting, the pricing equation becomes

$$0 = E_t\left[d\left(\Upsilon_t^k p_k\left(x_{k,d,t}\right)\right)\right] + \Upsilon_t^k\left(x_{k,d,t} - \delta\right)dt.$$

Deriving the price ODE The next step is to derive the dynamics of the rent-to-wage ratio and substitute them into the pricing equation written in wage units. For the ratio process $x_{k,d,t} = \frac{r_{k,d,t}}{w_{k,t}}$, notice

[25] Because the SDF does not depend on the scale of (r, w), scaling rent and the wage scales NOI, and therefore the price, by the same factor.

that

$$\partial_r x = \frac{1}{w}, \quad \partial_w x = -\frac{x}{w}$$
$$\partial^2_{rr} x = 0, \quad \partial^2_{ww} x = 2\frac{x}{w^2}, \quad \partial^2_{w,r} x = -\frac{1}{w^2}.$$

Then, applying Itô,

$$\begin{aligned} dx &= \frac{1}{w} dr - \frac{x}{w} dw + \frac{x}{w^2}\, d\langle w, w\rangle - \frac{1}{w^2}\, d\langle w, r\rangle \\ &= \frac{r}{w}\left(\mu_{r,k} dt + \sigma_{r,k} dB_t + \varsigma_{r,k} dB^k_t\right) - x\left(\mu_{w,k} dt + \sigma_{w,k} dB_t + \varsigma_{w,k} dB^k_t\right) \\ &\quad + x\left(\sigma^2_{w,k} + \varsigma^2_{w,k}\right) dt - x\left(\sigma_{r,k}\sigma_{w,k} + \varsigma_{r,k}\varsigma_{w,k}\right) dt, \end{aligned}$$

so

$$\frac{dx}{x} = \mu_{x,k}\, dt + \sigma_{x,k}\, dB_t + \varsigma_{x,k}\, dB^k_t, \tag{B1}$$

where

$$\sigma_{x,k} = \sigma_{r,k} - \sigma_{w,k}, \quad \varsigma_{x,k} = \varsigma_{r,k} - \varsigma_{w,k},$$
$$\mu_{x,k} = \mu_{r,k} - \mu_{w,k} + \sigma_{w,k}(\sigma_{w,k} - \sigma_{r,k}) + \varsigma_{w,k}(\varsigma_{w,k} - \varsigma_{r,k}).$$

For the pricing function, Itô gives

$$\begin{aligned} dp_k\left(x_{k,d,t}\right) &= p'(x)\, dx + \frac{1}{2} p''(x)\, d\langle x, x\rangle \\ &= p'(x)\, x\left(\mu_{x,k}\, dt + \sigma_{x,k}\, dB_t + \varsigma_{x,k}\, dB^k_t\right) \\ &\quad + \frac{1}{2} p''(x) x^2 \left(\sigma^2_{x,k} + \varsigma^2_{x,k}\right)\, dt \\ &= \left[p'(x_{k,d,t}) x_{k,d,t} \mu_{x,k} \;+\; p''(x_{k,d,t}) x^2_{k,d,t} \left(\frac{\sigma^2_{x,k} + \varsigma^2_{x,k}}{2}\right)\right] dt \\ &\quad +\; p'(x_{k,d,t})\, x_{k,d,t}\, \sigma_{x,k}\, dB_t \;+\; p'(x_{k,d,t})\, x_{k,d,t}\, \varsigma_{x,k}\, dB^k_t. \end{aligned}$$

For the inside of the expectation of the fundamental pricing equation, we use the product rule:

$$d\left(\Upsilon^k_t\, p(x_t)\right) = \Upsilon^k_t\, dp(x_t) + p(x_t)\, d\Upsilon^k_t + d\langle \Upsilon^k,\, p(x)\rangle.$$

The cross variation term is

$$\begin{aligned} d\langle \Upsilon^k,\, p(x)\rangle &= \Upsilon^k_t \Big[(-\lambda_{\Upsilon,k}) \cdot \sigma_{x,k}\, x_t\, p'(x_t) + (-\xi_{\Upsilon,k}) \cdot \varsigma_{x,k}\, x_t\, p'(x_t)\Big] dt \\ &= -\Upsilon^k_t\, x_t\, p'(x_t)\left(\lambda_{\Upsilon,k}\sigma_{x,k} + \xi_{\Upsilon,k}\varsigma_{x,k}\right) dt. \end{aligned}$$

Collect only the dt terms, as the martingale parts have zero conditional mean:

$$\begin{aligned}\mathbb{E}_t\big[d(\Upsilon_t^k\, p(x_t))\big] &= \Upsilon_t^k\Big[\underbrace{-\nu_{\Upsilon,k}\, p(x_t)}_{\text{from } p\, d\Upsilon} + \underbrace{\mu_{x,k}\, x_t\, p'(x_t)}_{\text{from } \Upsilon\, dp} \\ &\qquad \underbrace{- x_t\, p'(x_t)\,(\lambda_{\Upsilon,k}\sigma_{x,k} + \xi_{\Upsilon,k}\varsigma_{x,k})}_{\text{from } d\langle \Upsilon, p\rangle} + \underbrace{\tfrac{1}{2}(\sigma_{x,k}^2 + \varsigma_{x,k}^2)\, x_t^2\, p''(x_t)}_{\text{from } \Upsilon\, dp}\Big]dt \\ &= \Upsilon_t^k\Big[-\nu_{\Upsilon,k}\, p(x_t) \\ &\qquad + \Big(\mu_{x,k} - \lambda_{\Upsilon,k}\sigma_{x,k} - \xi_{\Upsilon,k}\varsigma_{x,k}\Big)\, x_t\, p'(x_t) \\ &\qquad + \Big(\frac{\sigma_{x,k}^2 + \varsigma_{x,k}^2}{2}\Big)\, x_t^2\, p''(x_t)\Big]dt.\end{aligned}$$

Substituting the drift into the pricing equation, we get

$$0 = \Upsilon_t^k\Big[-\nu_{\Upsilon,k}\, p(x_t) + (\mu_{x,k} - \lambda_{\Upsilon,k}\sigma_{x,k} - \xi_{\Upsilon,k}\varsigma_{x,k})\, x_t\, p'(x_t) + \tfrac{1}{2}(\sigma_{x,k}^2 + \varsigma_{x,k}^2)\, x_t^2\, p''(x_t) + (x_t - \delta)\Big]dt.$$

As $\Upsilon_t^k\, dt \neq 0$, the interior of the brackets must be zero. Isolating $\nu_{\Upsilon,k}\, p(x_t)$ yields

$$\nu_{\Upsilon,k}\, p(x_t) = (x_t - \delta) + \Big(\mu_{x,k} - \lambda_{\Upsilon,k}\sigma_{x,k} - \xi_{\Upsilon,k}\varsigma_{x,k}\Big)\, x_t\, p'(x_t) + \left(\frac{\sigma_{x,k}^2 + \varsigma_{x,k}^2}{2}\right) x_t^2\, p''(x_t). \tag{B2}$$

Let $\mu_{x,k}^{\mathbb{Q}}$ be the Υ-risk–neutral drift of x, given by

$$\mu_{x,k}^{\mathbb{Q}} \;=\; \mu_{x,k} \;-\; \lambda_{\Upsilon,k}\,\sigma_{x,k} \;-\; \xi_{\Upsilon,k}\,\varsigma_{x,k}.$$

Substituting into the pricing ODE in wage units (B2) yields

$$\nu_{\Upsilon,k}\, p_k(x) \;=\; (x - \delta) \;+\; \mu_{x,k}^{\mathbb{Q}}\, x\, p_k'(x) \;+\; \frac{\sigma_{x,k}^2 + \varsigma_{x,k}^2}{2}\, x^2\, p_k''(x). \tag{B3}$$

Assume that $\nu_{\Upsilon,k}$, $\mu_{x,k}^{\mathbb{Q}}$ and $\sigma_{x,k}^2 + \varsigma_{x,k}^2$ are constant. The general solution to this ODE is the sum of a particular solution, $p_P(x)$, and a homogeneous solution, $p_H(x)$:

$$p_{general}(x) = p_P(x) + p_H(x)$$

The homogeneous solution $p_H(x)$ solves the ODE when the dividend term $(x - \delta)$ is set to zero. It takes the form

$$p_H(x) = C_1\, x^{n_+} + C_2\, x^{n_-},$$

where, under $\nu_{\Upsilon,k} > \max\{\mu_{x,k}^{\mathbb{Q}}, 0\}$, the characteristic exponents $n_+ > 1$ and $n_- < 0$ solve

$$\left(\frac{\sigma_{x,k}^2 + \varsigma_{x,k}^2}{2}\right) n^2 \;+\; \left(\mu_{x,k}^{\mathbb{Q}} - \frac{\sigma_{x,k}^2 + \varsigma_{x,k}^2}{2}\right) n \;-\; \nu_{\Upsilon,k} \;=\; 0,$$

and C_1 and C_2 are constants to be determined by boundary conditions imposed by any options.

The particular solution $p_P(x)$ represents the value of the asset as a perpetual claim on its continuous stream of dividends, $(x - \delta)$. A linear function, $p_P(x) = ax + b$, solves the ODE, as we verify below.

The benchmark case without options In the absence of explicit abandonment or development options, the asset is a perpetual claim to NOI. Economic restrictions eliminate the homogeneous part of the solution:

- The term $C_1 x^{n_+}$, with $n_+ > 1$, would add a superlinear homogeneous component. Fundamental valuation adds a no-bubble (transversality) condition to the local pricing equation. For the perpetual claim, the admissible value has at most linear growth in x, which sets $C_1 = 0$.
- The term $C_2 x^{n_-}$ explodes in value as the state variable x approaches zero. Without an option to abandon the asset (i.e., a boundary condition at low x), this component has no economic basis, requiring us to set $C_2 = 0$.

With both constants of integration set to zero, only the particular solution remains for a simple, perpetual claim. We look for a particular solution of the linear form $p_k(x) = ax + b$, with $p_k' = a$ and $p_k'' = 0$. Substituting into (B3) and solving for a and b, we get

$$\begin{aligned} \nu_{Y,k}\,(ax + b) &= (x - \delta) + \mu^{Q}_{x,k}\, x\, a \;+\; 0 \\ (\nu_{Y,k}\, a)x + \nu_{Y,k} b &= (1 + \mu^{Q}_{x,k}\, a)x - \delta, \end{aligned}$$

so

$$\begin{aligned} \nu_{Y,k} a = 1 + \mu^{Q}_{x,k} a \implies a &= \frac{1}{\nu_{Y,k} - \mu^{Q}_{x,k}} \\ \nu_{Y,k} b = -\delta \implies b &= \frac{-\delta}{\nu_{Y,k}}. \end{aligned}$$

Therefore, in the benchmark case without options,

$$p_k(x) = \frac{x}{\nu_{Y,k} - \mu^{Q}_{x,k}} - \frac{\delta}{\nu_{Y,k}}.$$

Pricing elasticity The main text uses the pricing elasticity of the wage-normalized claim with respect to the state variable x:

$$\varepsilon_p(x) \equiv \frac{x\, p_k'(x)}{p_k(x)}.$$

From the benchmark price formula,

$$p_k'(x) = \frac{1}{\nu_{Y,k} - \mu^{Q}_{x,k}}.$$

Substituting $p_k(x)$ and $p_k'(x)$ into the elasticity definition gives

$$\begin{aligned}\varepsilon_p(x) &= \frac{x/(\nu_{Y,k} - \mu_{x,k}^{\mathbb{Q}})}{x/(\nu_{Y,k} - \mu_{x,k}^{\mathbb{Q}}) - \delta/\nu_{Y,k}} \\ &= \frac{x}{x - \delta\left(\frac{\nu_{Y,k} - \mu_{x,k}^{\mathbb{Q}}}{\nu_{Y,k}}\right)} \\ &= \frac{x}{x - \bar{\delta}_k},\end{aligned}$$

where

$$\bar{\delta}_k \equiv \delta\left(1 - \frac{\mu_{x,k}^{\mathbb{Q}}}{\nu_{Y,k}}\right).$$

This is the elasticity expression used in the operating-cost discussion in the main text.

Expected excess return In dollar terms, the property pays the net operating income flow $\text{NOI}_{k,d,t}\,dt$ and the ex-dividend property price is $P_{k,d,t} = w_{k,t}p_k(x_{k,d,t})$. Over an infinitesimal interval dt, the investor receives the price change plus the cash-flow payout. Thus, the instantaneous total return on the property claim is

$$dR_{k,d,t} \equiv \frac{dP_{k,d,t} + \text{NOI}_{k,d,t}\,dt}{P_{k,d,t}}.$$

To simplify notation, suppress the k, d subscripts inside the derivation and write $x_t = x_{k,d,t}$, $w_t = w_{k,t}$, $P_t = P_{k,d,t}$, and $y_t = \text{NOI}_{k,d,t}$. Taking conditional expectations, we get

$$\mathbb{E}_t[dR_t] = \frac{\mathbb{E}_t[dP_t]}{P_t} + \frac{y_t}{P_t}\,dt. \tag{B4}$$

The first step is to derive dP_t/P_t in dollars. Start from $P_t = w_t p_t$. Itô's product rule gives

$$dP_t = w_t\,dp_t + p_t\,dw_t + d\langle w, p\rangle_t.$$

Dividing by $P_t = w_t p_t$,

$$\frac{dP_t}{P_t} = \frac{dp_t}{p_t} + \frac{dw_t}{w_t} + \frac{d\langle w, p\rangle_t}{w_t p_t}.$$

We have dw from (3) in the main text and dx from (B1) above. We have also shown that

$$\begin{aligned}dp_t &= p'(x_t)\,dx_t + \frac{1}{2}p''(x_t)\,d\langle x, x\rangle_t \\ &= \left[p'(x_t)x_t\mu_{x,k} + \frac{1}{2}p''(x_t)x_t^2\left(\sigma_{x,k}^2 + \varsigma_{x,k}^2\right)\right]dt \\ &\quad + p'(x_t)x_t\sigma_{x,k}\,dB_t + p'(x_t)x_t\varsigma_{x,k}\,dB_t^k.\end{aligned}$$

Divide by $p_t = p(x_t)$ and use the definition of the pricing elasticity $\varepsilon_p(x_t) = \frac{x_t p'(x_t)}{p(x_t)}$ to get

$$\frac{dp_t}{p(x_t)} = \left[\varepsilon_p(x_t)\mu_{x,k} + \frac{1}{2}\frac{x_t^2 p''(x_t)}{p(x_t)}\left(\sigma_{x,k}^2 + \varsigma_{x,k}^2\right)\right] dt + \varepsilon_p(x_t)\sigma_{x,k}\, dB_t + \varepsilon_p(x_t)\varsigma_{x,k}\, dB_t^k.$$

The cross-variation term in dP_t is

$$d\langle w, p\rangle_t = w_t p_t\, \varepsilon_p(x_t)\left(\sigma_{w,k}\sigma_{x,k} + \varsigma_{w,k}\varsigma_{x,k}\right) dt,$$

so

$$\frac{d\langle w, p\rangle_t}{w_t p_t} = \varepsilon_p(x_t)\left(\sigma_{w,k}\sigma_{x,k} + \varsigma_{w,k}\varsigma_{x,k}\right) dt.$$

Substituting into dP_t/P_t yields

$$\frac{dP_t}{P_t} = \mu_{P,t}\, dt + \sigma_{P,t}\, dB_t + \varsigma_{P,t}\, dB_t^k,$$

where the drift is

$$\mu_{P,t} = \mu_{w,k} + \varepsilon_p(x_t)\mu_{x,k} + \frac{1}{2}\frac{x_t^2 p''(x_t)}{p(x_t)}\left(\sigma_{x,k}^2 + \varsigma_{x,k}^2\right) + \varepsilon_p(x_t)\left(\sigma_{w,k}\sigma_{x,k} + \varsigma_{w,k}\varsigma_{x,k}\right),$$

and the diffusion exposures are

$$\sigma_{P,t} = \sigma_{w,k} + \varepsilon_p(x_t)\sigma_{x,k}, \qquad \varsigma_{P,t} = \varsigma_{w,k} + \varepsilon_p(x_t)\varsigma_{x,k}.$$

Because

$$\mathbb{E}_t\left[\frac{dP_t}{P_t}\right] = \mu_{P,t}\, dt,$$

it holds that

$$\mathbb{E}_t[dR_t] = \left(\mu_{P,t} + \frac{y_t}{P_t}\right) dt.$$

Now use the dollar-denominated stochastic discount factor

$$\frac{d\Lambda_t^k}{\Lambda_t^k} = -r_f\, dt - \lambda^k\, dB_t - \xi^k\, dB_t^k$$

and the pricing equation

$$0 = \mathbb{E}_t[d(\Lambda_t^k P_t)] + \Lambda_t^k y_t\, dt.$$

Expand the product:

$$d(\Lambda_t^k P_t) = \Lambda_t^k\, dP_t + P_t\, d\Lambda_t^k + d\langle \Lambda^k, P\rangle_t.$$

Dividing the no-arbitrage condition by $\Lambda_t^k P_t\, dt$ and collecting drift terms gives the following pricing restriction. The first term gives

$$\frac{\mathbb{E}_t[\Lambda_t^k\, dP_t]}{\Lambda_t^k P_t\, dt} = \mu_{P,t}$$

and the second term gives

$$\frac{\mathbb{E}_t[P_t\, d\Lambda_t^k]}{\Lambda_t^k P_t\, dt} = -r_f.$$

The cross-variation term gives

$$\frac{d\langle \Lambda^k, P\rangle_t}{\Lambda_t^k P_t\, dt} = -\lambda^k \sigma_{P,t} - \xi^k \varsigma_{P,t}$$

and the dividend term gives

$$\frac{\Lambda_t^k y_t\, dt}{\Lambda_t^k P_t\, dt} = \frac{y_t}{P_t}.$$

Substituting these terms into the pricing equation, we get

$$0 = \mu_{P,t} - r_f - \lambda^k \sigma_{P,t} - \xi^k \varsigma_{P,t} + \frac{y_t}{P_t}.$$

Rearranging leads us to

$$\mu_{P,t} + \frac{y_t}{P_t} - r_f = \lambda^k \sigma_{P,t} + \xi^k \varsigma_{P,t}.$$

Finally, multiplying the previous display by dt and using $\mathbb{E}_t[dR_t] = \left(\mu_{P,t} + y_t/P_t\right) dt$, we get

$$\begin{aligned}\mathbb{E}_t[dR_t] - r_f\, dt &= \left[\lambda^k \left(\sigma_{w,k} + \varepsilon_p(x_t)\sigma_{x,k}\right) + \xi^k \left(\varsigma_{w,k} + \varepsilon_p(x_t)\varsigma_{x,k}\right)\right] dt \\ &= \left[\left(\lambda^k \sigma_{w,k} + \xi^k \varsigma_{w,k}\right) + \varepsilon_p(x_{k,d,t}) \left(\lambda^k \sigma_{x,k} + \xi^k \varsigma_{x,k}\right)\right] dt.\end{aligned}$$

This is the continuous-time expected excess return increment used in the main text.

B.4 Development, abandonment and reuse

We add development, abandonment, and reuse options to the perpetual-claim benchmark. Their contribution to the value of an operating property is concentrated where margins are thin, at low values of the rent-to-wage ratio x, and vanishes at high x. Where the perpetual-claim price is above zero, an option term that is positive and decays with x lowers the operating pricing elasticity relative to that benchmark. Near the abandonment threshold, the elasticity need not be monotone.

We consider three cases. In the two-regime model, a developed property can revert to undeveloped land, which keeps the option to develop again at cost c_k. In the reuse model, an idle structure is preserved and can restart at cost $0 < c_R < c_k$. In the zero-salvage case, an abandoned property has no

further use. The calibration uses only this last case, in the tax-adjusted, fixed-parameter sensitivity of Appendix C.2.

We use the constant-coefficient pricing operator of Appendix B.3 without property taxes, and suppress city and distance indices:

$$\mathcal{L}f(x) = \mu x f'(x) + \frac{1}{2}S_x^2 x^2 f''(x) - \nu f(x), \qquad \nu = \nu_{Y,k}, \quad \mu = \mu_{x,k}^{Q}, \quad S_x^2 = \sigma_{x,k}^2 + \varsigma_{x,k}^2.$$

We assume $\delta > 0$, $S_x > 0$, $\nu > \max\{\mu, 0\}$, and a positive construction cost c_k. The characteristic exponents of Appendix B.3, defined by $\mathcal{L}x^n = 0$, satisfy $n_- < 0$ and $n_+ > 1$.

An operating property has flow $x - \delta$, and idle structures and undeveloped land have zero flow. The switching costs c_k and c_R are lump sums in wage units, so a switch at time t is a one-time payment of $w_{k,t}c_k$ or $w_{k,t}c_R$ dollars. Under the wage-denominated pricing measure, x follows a geometric Brownian motion on $(0, \infty)$ and reaches neither endpoint in finite time. For simplicity, we omit property taxes. A value-proportional tax at a common rate τ_P in every regime would replace ν by $\nu + \tau_P$ throughout, which is the discount rate of the zero-salvage sensitivity.

An operating property continues to operate for all $x > x_L$, with value

$$q(x) = \frac{x}{\nu - \mu} - \frac{\delta}{\nu} + Ax^{n_-}, \qquad x > x_L,$$

where the switching conditions determine A. The x^{n_+} term is absent because fundamental values cannot exceed the present value of gross rents, $x/(\nu - \mu)$. The development threshold x_H concerns only undeveloped land and does not bound the operating region.

In the reuse model, an operating owner continues for $x > x_L$, an idle owner waits for $x < x_R$, and an owner of undeveloped land waits for $x < x_H$. Where the owner waits, the value of an idle structure is $p_I(x) = Cx^{n_+}$ and the value of land is $l(x) = Dx^{n_+}$, with no x^{n_-} terms because values vanish as $x \downarrow 0$. Elsewhere, $p(x) = p_I(x)$ below x_L, $p_I(x) = p(x) - c_R$ above x_R, and $l(x) = p(x) - c_k$ above x_H. Three value-matching and smooth-pasting pairs give six equations for the coefficients (A, C, D) and the thresholds (x_L, x_R, x_H). Below, we state the regime inequalities and verify that ordered solutions exist.

Owners can switch between operation and idleness repeatedly, paying the sunk cost c_R at each restart, and an idle structure remains developed. Because the operating and idle regions overlap on (x_L, x_R), the model has hysteresis: a property idled at x_L restarts only when x recovers to x_R. Locally, a higher restart cost lowers x_L and raises x_R, holding the pricing operator fixed. With $c_R < c_k$, idle structures restart before undeveloped land develops, so $x_R < x_H$. We can solve the operating and idle problem first, and then use c_k to determine the land value and the development threshold without changing the operating value.

In the reuse model, smooth pasting gives a pricing elasticity of $\varepsilon_p(x_L+) = n_+$ at the abandonment threshold. In the zero-salvage case, the price and its derivative both vanish at the threshold, so the elasticity can be large even where the price is flat.

We now state the optimal stopping problem that determines the development and abandonment thresholds.

Start from the wage-denominated representation:

$$\mathrm{Y}_t\, l(x_t) \;=\; \sup_{\tau \geq 0} \mathbb{E}_t\big[\mathrm{Y}_{t+\tau}\,(p(x_{t+\tau}) - c_k)\big].$$

This is an optimal stopping problem. When land is undeveloped, there is no payout. Upon development, the payoff is

$$G_t \;\equiv\; \mathrm{Y}_t\big(p(x_t) - c_k\big),$$

so

$$\mathrm{Y}_t\, l(x_t) \;=\; \sup_{\tau \geq 0} \mathbb{E}_t\big[\, G_{t+\tau}\,\big].$$

The Snell envelope $V_t \;\equiv\; \mathrm{Y}_t\, l(x_t)$ is the smallest supermartingale that dominates G_t. On the continuation region, $\mathbb{E}_t[dV_t] = 0$, so

$$\mathbb{E}_t\big[d(\mathrm{Y}_t\, l(x_t))\big] \;=\; 0.$$

By Itô's product rule,

$$d(\mathrm{Y}_t \cdot l(x_t)) = \mathrm{Y}_t \cdot dl(x_t) + l(x_t) \cdot d\mathrm{Y}_t + d\langle \mathrm{Y}_t, l(x_t)\rangle.$$

For $dl(x_t)$, Itô's lemma gives

$$dl(x_t) = \left[\mu_x\, x_t\, l'(x_t) + \frac{1}{2}(\sigma_x^2 + \varsigma_x^2)\, x_t^2\, l''(x_t)\right] dt + x_t\, l'(x_t)\left(\sigma_x\, dB_t + \varsigma_x\, dB_t^k\right).$$

The cross-variation between Y_t and $l(x_t)$ is

$$d\langle \mathrm{Y}_t, l(x_t)\rangle = -\mathrm{Y}_t\, x_t\, l'(x_t)\,(\lambda_{\mathrm{Y}}\sigma_x + \xi_{\mathrm{Y}}\varsigma_x)dt.$$

Putting everything together,

$$\begin{aligned} dV_t = \mathrm{Y}_t\Big[-\nu_{\mathrm{Y}} l(x_t) \;+ \big(\mu_x - \lambda_{\mathrm{Y}}\sigma_x - \xi_{\mathrm{Y}}\varsigma_x\big)\, x_t\, l'(x_t) + \frac{1}{2}(\sigma_x^2 + \varsigma_x^2)\, x_t^2\, l''(x_t)\Big] dt \\ + \mathrm{Y}_t\Big(x_t\, l'(x_t)\sigma_x - l(x_t)\lambda_{\mathrm{Y}}\Big) dB_t + \mathrm{Y}_t\Big(x_t\, l'(x_t)\varsigma_x - l(x_t)\xi_{\mathrm{Y}}\Big) dB_t^k, \end{aligned}$$

so

$$\mathbb{E}_t[dV_t] = \mathrm{Y}_t\Big[-\nu_{\mathrm{Y}} l(x_t) \;+ \big(\mu_x - \lambda_{\mathrm{Y}}\sigma_x - \xi_{\mathrm{Y}}\varsigma_x\big)\, x_t\, l'(x_t) + \frac{1}{2}(\sigma_x^2 + \varsigma_x^2)\, x_t^2\, l''(x_t)\Big] dt.$$

Imposing $\mathbb{E}_t[dV_t] = 0$ on the continuation region yields the ODE for $l(x)$:

$$\nu_{\mathrm{Y}}\, l(x) = \big(\mu_x - \lambda_{\mathrm{Y}}\,\sigma_x - \xi_{\mathrm{Y}}\,\varsigma_x\big)\, x\, l'(x) + \left(\tfrac{\sigma_x^2 + \varsigma_x^2}{2}\right) x^2\, l''(x).$$

Stopping region and boundary conditions On the stopping set (optimal to develop), $V_t = G_t$:

$$l(x) = p(x) - c_k \qquad \text{(value matching)}.$$

At the endogenous exercise boundary $x = x_H$, smooth pasting imposes

$$l'(x_H) = p'(x_H).$$

Similarly, with abandonment, the lower boundary x_L satisfies

$$l(x_L) = p(x_L), \qquad l'(x_L) = p'(x_L).$$

Together with the ODE for $p(\cdot)$ on the continuation region, these conditions solve the abandonment problem.

The lower-boundary pair describes abandonment *to land*. The owner keeps the option to redevelop, so $l(x_L)$ is positive. For this two-regime formulation, operating continuation is (x_L, ∞) and undeveloped-land continuation is $(0, x_H)$. Their values obey $p = l$ below x_L and $l = p - c_k$ above x_H.

The economic boundary conditions follow from fundamental valuation. Since operation yields at most the gross rent flow and all costs are nonnegative, while the owner can remain out of operation, each option value satisfies

$$0 \le V(x) \le \mathbb{E}_x^{\mathbb{Q}}\left[\int_0^\infty e^{-\nu t} X_t \, dt\right] = \frac{x}{\nu - \mu}.$$

Consequently $V(0+) = 0$, $V(x) = O(x)$ at infinity, and $\lim_{T\to\infty} \mathbb{E}_x^{\mathbb{Q}}[e^{-\nu T} V(X_T)] = 0$. The roots solve

$$\tfrac{1}{2} S_x^2 n(n-1) + \mu n - \nu = 0, \qquad n_- < 0, \quad n_+ > 1.$$

Thus the operating continuation branch retains only Ax^{n_-} in addition to the particular solution, and the land waiting branch retains only Dx^{n_+}. The four matching and pasting conditions are then equations for (A, D, x_L, x_H). The condition at x_H concerns only land, and the operating branch extends past it. A solution of these equations must also satisfy the continuation and switching inequalities. Setting $c_R = c_k$ in the reuse model below recovers this case.

Extension: reuse option. An operating property may be idled at zero cost, preserving its structure. An idle property may restart at cost c_R, whereas undeveloped land enters operation at cost c_k. Assume $0 < c_R < c_k$. Idling takes the place of the abandonment-to-land condition above. Operation and idleness can alternate, and each round trip costs $c_R > 0$. Structures do not deteriorate, idle properties carry no holding cost, and the state does not reset at a switch.

Let p, p_I, l denote the values when the property is currently operating, idle, or undeveloped. Define the operating continuation branch

$$q(x) = ax - b + Ax^{n_-}, \qquad a = \frac{1}{\nu - \mu}, \quad b = \frac{\delta}{\nu}.$$

The growth and zero-state conditions above exclude the operating positive-root coefficient and the idle/land negative-root coefficients. For an ordered threshold solution the values are

$$p(x) = \begin{cases} Cx^{n_+}, & 0 < x \le x_L, \\ q(x), & x > x_L, \end{cases} \qquad p_I(x) = \begin{cases} Cx^{n_+}, & 0 < x < x_R, \\ q(x) - c_R, & x \ge x_R, \end{cases}$$

$$l(x) = \begin{cases} Dx^{n_+}, & 0 < x < x_H, \\ q(x) - c_k, & x \ge x_H. \end{cases}$$

Operating continuation extends through x_H to infinity. The operating and idle continuation regions overlap on (x_L, x_R), where the owner's current status, operating or idle, determines whether the property operates.

The boundary equations are

$$\begin{aligned} q(x_L) &= Cx_L^{n_+}, & q'(x_L) &= n_+ Cx_L^{n_+-1}, \\ Cx_R^{n_+} &= q(x_R) - c_R, & n_+ Cx_R^{n_+-1} &= q'(x_R), \\ Dx_H^{n_+} &= q(x_H) - c_k, & n_+ Dx_H^{n_+-1} &= q'(x_H). \end{aligned}$$

These are six equations for three coefficients (A, C, D) and three thresholds (x_L, x_R, x_H). The first four solve the operating/idle subsystem, and the last two then determine development. Reuse leaves the roots unchanged because the pricing operator is held fixed.

Verification and threshold ordering. Optimality requires more than these equations. The values must be nonnegative, have admissible growth, and satisfy the following variational inequalities away from the thresholds, with continuous values and first derivatives at the thresholds:

$$\begin{aligned} \max\{\mathcal{L}p + x - \delta,\ p_I - p\} &= 0, \\ \max\{\mathcal{L}p_I,\ p - c_R - p_I\} &= 0, \\ \max\{\mathcal{L}l,\ p - c_k - l\} &= 0. \end{aligned}$$

In particular, $0 \le p - p_I \le c_R$ and $l \ge p - c_k$ on the entire state space. The continuation operator must also be nonpositive throughout the switching regions, not only at the exercise points.

Here a positive solution $A, C > 0$ with $0 < x_L < x_R$ to the first four equations admits a direct verification. Write $m = n_-$, $n = n_+$, and $\Delta(x) = q(x) - Cx^n$ on $[x_L, x_R]$. Its second derivative,

$$\Delta''(x) = m(m-1)Ax^{m-2} - n(n-1)Cx^{n-2},$$

changes sign at most once, from positive to negative. Since $\Delta'(x_L) = \Delta'(x_R) = 0$, it follows that $\Delta' > 0$ inside the band. Thus $0 < \Delta < c_R$ there, and the operator equation implies

$$x_L < \delta, \qquad x_R > \delta + \nu c_R.$$

These inequalities give $\mathcal{L}p + x - \delta = x - \delta \leq 0$ for $x \leq x_L$ and $\mathcal{L}p_I = \delta + \nu c_R - x \leq 0$ for $x \geq x_R$.

Holding (ν, μ, S_x, δ) fixed, implicit differentiation of the four operating/idle boundary conditions along a regular solution gives $A_{c_R} < 0$, $C_{c_R} < 0$, and, because $\Delta''(x_L) > 0 > \Delta''(x_R)$,

$$\frac{dx_L}{dc_R} < 0 < \frac{dx_R}{dc_R}.$$

This is a local comparative static in c_R. It does not extend to changes in volatility.

The admissible class is nonempty for every $c_R > 0$. For any $s > 1$, define

$$u_s = \frac{s^{n-1} - 1}{m(s^{m-1} - s^{n-1})}, \qquad v_s = \frac{1 - s^{m-1}}{n(s^{n-1} - s^{m-1})},$$

and set

$$L_s = \frac{b}{a(1 + u_s - v_s)}, \quad R_s = sL_s, \quad A_s = au_s L_s^{1-m}, \quad C_s = av_s L_s^{1-n}.$$

These quantities satisfy value matching at L_s and smooth pasting at both L_s and R_s. The implied restart cost is

$$c_R(s) = aL_s[s + u_s s^m - v_s s^n] - b.$$

It is continuous, tends to zero as $s \downarrow 1$, and tends to infinity as $s \to \infty$. Hence every positive restart cost admits at least one positive ordered operating/idle solution. This establishes existence, not uniqueness.

For development define, on the operating continuation branch,

$$h_c(x) = xq'(x) - n[q(x) - c] = (1 - n)ax + n(b + c) + (m - n)Ax^m.$$

This function is strictly concave. At x_R, $h_{c_k}(x_R) = n(c_k - c_R) > 0$, whereas $h_{c_k}(x) \to -\infty$ as $x \to \infty$. Its unique root above x_R determines x_H and obeys

$$x_H > x_R, \qquad x_H > \delta + \nu c_k, \qquad D = \frac{q(x_H) - c_k}{x_H^n} > 0.$$

The ratio $(p(x) - c_k)/x^n$ increases up to x_H: below x_L it equals $C - c_k/x^n$, and above x_L its derivative has the sign of h_{c_k}. Therefore $Dx^n \geq p(x) - c_k$ throughout the waiting region. On the development region, $\mathcal{L}l = \delta + \nu c_k - x \leq 0$.

The resulting functions are nonnegative, continuously differentiable, piecewise twice differentiable, and of at most linear growth. Applying Itô's formula to the discounted value plus accumulated net cash flow gives an upper bound on the payoff of every admissible switching policy. Equality holds for the threshold policy, because the terminal term vanishes when $\nu > \max\{\mu, 0\}$. Positive round-trip cost and distinct thresholds exclude costless switching cycles and finite-time accumulation of switches. The candidate is therefore optimal.

Because $c_R < c_k$ and the waiting regimes otherwise have the same dynamics and zero flow, $p_I \geq l$, since an idle owner can imitate any development timing and pay the smaller first-entry cost. Free disposal of an idle structure to land would therefore not improve value. If such disposal is allowed and $c_R > c_k$, the dominance argument reverses, and reuse need not precede development.

When $c_R = c_k$, idle property and undeveloped land have the same value, $C = D$ and $x_R = x_H$, recovering the preceding abandonment-to-land model. Zero-salvage abandonment is a different restriction, which rules out any use after exit. In the operating/idle subsystem it is the limit as the restart cost grows without bound. The tax-adjusted sensitivity in Appendix C.2 uses this zero-salvage case.

B.5 Time-varying risk premia

This subsection derives prices, expected returns, and return volatility from the same state system, as in Ang and Liu (2007). City subscripts are suppressed.

Wages, the rent-to-wage ratio, the aggregate state, and the dollar stochastic discount factor follow

$$\frac{dw_t}{w_t} = \mu_w\,dt + \sigma_w\,dB_t + \varsigma_w\,dB_t^k, \tag{B5}$$

$$\frac{dx_t}{x_t} = \mu_x\,dt + \sigma_x\,dB_t + \varsigma_x\,dB_t^k, \tag{B6}$$

$$dz_t = \kappa_z(\bar{z} - z_t)\,dt + \sigma_z\,dB_t, \tag{B7}$$

$$\frac{d\Lambda_t}{\Lambda_t} = -r_f\,dt - \lambda(z_t)\,dB_t - \xi\,dB_t^k. \tag{B8}$$

The wage-denominated SDF is $\mathrm{Y}_t = \Lambda_t w_t$. Itô's product rule gives

$$\frac{d\mathrm{Y}_t}{\mathrm{Y}_t} = -\nu_{\mathrm{Y}}(z_t)\,dt - \lambda_{\mathrm{Y}}(z_t)\,dB_t - \xi_{\mathrm{Y}}\,dB_t^k, \tag{B9}$$

where

$$\lambda_{\mathrm{Y}}(z) = \lambda(z) - \sigma_w, \qquad \xi_{\mathrm{Y}} = \xi - \varsigma_w, \qquad \nu_{\mathrm{Y}}(z) = r_f - \mu_w + \lambda(z)\sigma_w + \xi\varsigma_w. \tag{B10}$$

These are identities. If only $\lambda(z) = \ell_0 + \ell_1 z$ varies with the state, the slope of $\nu_{\mathrm{Y}}(z)$ is $\sigma_w\ell_1$, so the wage-unit objects cannot be parameterized independently under these primitives.

Tax-adjusted two-state pricing equation. The dollar property price is $P_t = w_t p(x_t, z_t)$. Non-tax operating cash flow is $w_t(x_t - \delta)$, while the value-based property-tax payment is $\tau_P P_t$. Define

$$\mu_x^{\mathrm{Q}}(z) = \mu_x - \lambda_{\mathrm{Y}}(z)\sigma_x - \xi_{\mathrm{Y}}\varsigma_x, \qquad \mu_z^{\mathrm{Q}}(z) = \kappa_z(\bar{z} - z) - \lambda_{\mathrm{Y}}(z)\sigma_z. \tag{B11}$$

The pricing condition $0 = E_t[d(\Lambda_t P_t)] + \Lambda_t[w_t(x_t - \delta) - \tau_P P_t]dt$ gives

$$0 = (x - \delta) - [\nu_Y(z) + \tau_P]p + \mu_x^{\mathbb{Q}}(z)xp_x + \mu_z^{\mathbb{Q}}(z)p_z + \frac{1}{2}(\sigma_x^2 + \varsigma_x^2)x^2 p_{xx} + \frac{1}{2}\sigma_z^2 p_{zz} + \sigma_x\sigma_z x p_{xz}. \tag{B12}$$

The mixed derivative has coefficient $\sigma_x\sigma_z x$ because x and z share the aggregate shock. Property taxes enter through $\nu_Y + \tau_P$ and remain separate from the non-tax cost δ.

Affine reduction. The fundamental no-option value is affine in x:

$$p(x, z) = a(z)x - \delta c(z). \tag{B13}$$

Substituting into equation (B12) and collecting the coefficient on x and the constant term gives

$$0 = 1 - [\nu_Y(z) + \tau_P - \mu_x^{\mathbb{Q}}(z)]a(z) + [\mu_z^{\mathbb{Q}}(z) + \sigma_x\sigma_z]a'(z) + \frac{1}{2}\sigma_z^2 a''(z), \tag{B14}$$

$$0 = 1 - [\nu_Y(z) + \tau_P]c(z) + \mu_z^{\mathbb{Q}}(z)c'(z) + \frac{1}{2}\sigma_z^2 c''(z). \tag{B15}$$

The function $a(z)$ is the wage-unit value of the claim to rents per unit of x, and $c(z)$ is the wage-unit value of the claim to operating costs per unit of δ. No-bubble, transversality, and admissibility conditions select these fundamental solutions and exclude homogeneous components. The affine representation is not imposed on the real-option extension.

Exact dollar-return exposure. Define

$$\varepsilon(x, z) = \frac{xp_x(x, z)}{p(x, z)}, \qquad \chi(x, z) = \frac{p_z(x, z)}{p(x, z)}. \tag{B16}$$

Applying Itô's lemma to $P = wp(x, z)$ gives

$$\sigma_P(x, z) = \sigma_w + \sigma_x\varepsilon(x, z) + \sigma_z\chi(x, z), \qquad \varsigma_P(x, z) = \varsigma_w + \varsigma_x\varepsilon(x, z). \tag{B17}$$

The exact expected dollar excess return is

$$\pi(x, z) = \lambda(z)[\sigma_w + \sigma_x\varepsilon(x, z) + \sigma_z\chi(x, z)] + \xi[\varsigma_w + \varsigma_x\varepsilon(x, z)]. \tag{B18}$$

Equation (B18) uses the dollar prices of risk, whereas equation (B11) uses their wage-unit counterparts.

Valuation-state exposure. For the affine solution, define $A(z) = a'(z)/a(z)$ and $C(z) = c'(z)/c(z)$, unrelated to the option coefficients A and C of Appendix B.4. Direct algebra yields

$$\chi(x, z) = C(z) + \varepsilon(x, z)[A(z) - C(z)]. \tag{B19}$$

Thus

$$\sigma_P(x,z) = [\sigma_w + \sigma_z C(z)] + \varepsilon(x,z)[\sigma_x + \sigma_z(A(z) - C(z))]. \tag{B20}$$

Define

$$\Gamma(z) = \lambda(z)[\sigma_x + \sigma_z(A(z) - C(z))] + \xi\varsigma_x. \tag{B21}$$

The premium is

$$\pi_i(z) = \pi_0(z) + \Gamma(z)\varepsilon_i(z), \qquad \pi_i(z) - \pi_j(z) = \Gamma(z)[\varepsilon_i(z) - \varepsilon_j(z)], \tag{B22}$$

where $\pi_0(z) = \lambda(z)[\sigma_w + \sigma_z C(z)] + \xi\varsigma_w$ is common across locations within the city. The valuation-state exposure also has a cross-sectional component, which is collinear with the same elasticity that carries operating leverage.

Local mapping to the empirical specification. For a common state, the non-tax accounting leverage $\omega_i^{NT} = x_i/(x_i - \delta)$ is strictly increasing in $\varepsilon_i(z)$ for $x_i > \delta$ and $p > 0$, where non-tax operating income and value are positive. Demean z around the expansion point. The location-varying component has the first-order expansion

$$\Gamma(z)\varepsilon_i(z) = \Gamma(0)\varepsilon_i(0) + z\{\Gamma'(0)\varepsilon_i(0) + \Gamma(0)\varepsilon_{i,z}(0)\} + O(z^2). \tag{B23}$$

A further local linearization in ω^{NT} gives

$$r^e_{i,t+1} = \alpha_i + \delta_t + \gamma\,\omega_i^{NT} z_t + u_{i,t+1}, \qquad \gamma \approx \left.\frac{\partial^2[\Gamma(z)\varepsilon(\omega^{NT},z)]}{\partial\omega^{NT}\partial z}\right|_{(\bar{\omega},0)}. \tag{B24}$$

ZIP fixed effects absorb the time-average location component, and date fixed effects absorb the common state component under the common-coefficient approximation used for this pooled test. If the state mapping differs freely by city, exact absorption would require CBSA-by-date effects or equivalent city-state controls. The derivation holds each location's current leverage fixed, and the regression uses the ZIP code's average leverage over its expense history as a proxy for it.

The coefficient in equation (B24) combines variation in the dollar price of aggregate risk, variation in the effective aggregate loading, and state dependence of the pricing elasticity. It does not separately identify those components. This scaled-factor interpretation follows Cochrane (1996) and Lettau and Ludvigson (2001b). A positive interaction requires a positive cross derivative of the composite $\Gamma(z)\varepsilon_i(z)$ and does not follow from $\lambda'(z) > 0$ alone.

Income and capital components of the premium. The dollar total return is the sum of an income return and a dollar capital gain,

$$dR_t = \frac{w_t(x_t - \delta) - \tau_P P_t}{P_t}\,dt + \frac{dP_t}{P_t}.$$

Under the affine solution both expected components are affine in the pricing elasticity. The income yield is a ratio, so it is the same in dollar and wage units. Using $x/p = \varepsilon/a$ and $\delta/p = (\varepsilon-1)/c$,

$$y(x,z) = \frac{x-\delta}{p(x,z)} - \tau_P = \Big[\frac{1}{c(z)} - \tau_P\Big] + \Gamma^y(z)\,\varepsilon(x,z), \qquad \Gamma^y(z) = \frac{1}{a(z)} - \frac{1}{c(z)}. \tag{B25}$$

When the price of risk does not vary with the state ($\ell_1 = 0$ below), $a = 1/(\nu_Y + \tau_P - \mu_x^{\mathbb{Q}})$ and $c = 1/(\nu_Y + \tau_P)$. Then $\Gamma^y = -\mu_x^{\mathbb{Q}}$, and equation (B25) is the after-tax cap rate $\nu_Y - \mu_x^{\mathbb{Q}}\varepsilon$ of the main text.

Itô's lemma applied to $P = wp(x,z)$, together with $\chi = C + \varepsilon(A-C)$ and $p_{zz}/p = \varepsilon\,a''/a + (1-\varepsilon)\,c''/c$, gives the expected dollar capital gain

$$q(x,z) = \frac{E_t[dP_t/P_t]}{dt} = q_0(z) + \Gamma^q(z)\,\varepsilon(x,z), \tag{B26}$$

where

$$q_0(z) = \mu_w + \big[\kappa_z(\bar z - z) + \sigma_w\sigma_z\big]C(z) + \tfrac{1}{2}\sigma_z^2\,\frac{c''(z)}{c(z)},$$
$$\Gamma^q(z) = \mu_x + \sigma_w\sigma_x + \varsigma_w\varsigma_x + \sigma_x\sigma_z A(z) + \big[\kappa_z(\bar z - z) + \sigma_w\sigma_z\big]\big[A(z) - C(z)\big] + \tfrac{1}{2}\sigma_z^2\Big[\frac{a''(z)}{a(z)} - \frac{c''(z)}{c(z)}\Big].$$

The two slopes add up to the premium slope of equation (B21),

$$\Gamma^y(z) + \Gamma^q(z) = \Gamma(z). \tag{B27}$$

To see this, divide equations (B14) and (B15) by a and c, which gives $1/a = \nu_Y + \tau_P - \mu_x^{\mathbb{Q}} - (\mu_z^{\mathbb{Q}} + \sigma_x\sigma_z)A - \tfrac{1}{2}\sigma_z^2 a''/a$ and $1/c = \nu_Y + \tau_P - \mu_z^{\mathbb{Q}}C - \tfrac{1}{2}\sigma_z^2 c''/c$. Subtracting,

$$\Gamma^y(z) = -\mu_x^{\mathbb{Q}}(z) - \mu_z^{\mathbb{Q}}(z)\big[A(z) - C(z)\big] - \sigma_x\sigma_z A(z) - \tfrac{1}{2}\sigma_z^2\Big[\frac{a''}{a} - \frac{c''}{c}\Big].$$

Adding Γ^q and using $\mu_x - \mu_x^{\mathbb{Q}} = \lambda_Y\sigma_x + \xi_Y\varsigma_x$ and $\kappa_z(\bar z - z) - \mu_z^{\mathbb{Q}} = \lambda_Y\sigma_z$ from equation (B11) leaves $\lambda(z)[\sigma_x + \sigma_z(A-C)] + \xi\varsigma_x = \Gamma(z)$. The intercepts, $y_0(z) = 1/c(z) - \tau_P$ and $q_0(z)$, satisfy $y_0 + q_0 = r_f + \pi_0$. Within a city, the leverage premium $\Gamma(z)[\varepsilon_i - \varepsilon_j]$ therefore divides into an income part $\Gamma^y(z)[\varepsilon_i - \varepsilon_j]$ and a capital part $\Gamma^q(z)[\varepsilon_i - \varepsilon_j]$.

How the interaction divides. The empirical interaction holds each location's leverage fixed while the state moves. Write ω for ω^{NT}, $e = \delta/x = 1 - 1/\omega$, and $\theta(z) = c(z)/a(z)$. The elasticity is $\varepsilon(\omega,z) = 1/[1-\theta(z)e]$, so $\varepsilon_\omega = \theta\varepsilon^2/\omega^2$ and $\varepsilon_{\omega z}/\varepsilon_\omega = (C-A)(2\varepsilon-1)$. For $j \in \{y,q\}$,

$$\frac{\partial^2[\Gamma^j(z)\varepsilon]}{\partial\omega\,\partial z} = \varepsilon_\omega\Big\{\Gamma^{j\prime}(z) + \Gamma^j(z)\big[C(z) - A(z)\big](2\varepsilon-1)\Big\}, \tag{B28}$$

and the two components add up to the cross derivative in equation (B24). The first term is the change in the premium per unit of elasticity. The second arises because a change in the state also changes the elasticity of a location with given accounting leverage, and it divides in proportion to the levels Γ^y

and Γ^q. These are signed contributions, and either can have the opposite sign from their sum. For two locations with given leverage, the same identities give

$$\frac{\partial[\pi_i(z)-\pi_j(z)]}{\partial z}=(\varepsilon_i-\varepsilon_j)\Big\{\Gamma'(z)+\Gamma(z)\big[C(z)-A(z)\big](\varepsilon_i+\varepsilon_j-1)\Big\}.$$

The spread widens with the state when the braced term is positive. A sufficient condition is that Γ is positive and rises with the state and that the elasticity gap does not narrow, which holds when $C \geq A$ because $\varepsilon \geq 1$.

Persistence of the state. Let the dollar price of risk be $\lambda(z) = \ell_0 + \ell_1 z$. The wage-unit price of risk $\lambda_Y(z) = \lambda(z) - \sigma_w$ has the same slope ℓ_1, so each claim's wage-unit discount rate is linear in the state. The rent claim is discounted at $R_a(z) = \nu_Y(z) + \tau_P - \mu_x^{\mathbb{Q}}(z)$, with slope $R_a' = \ell_1(\sigma_w + \sigma_x)$. The cost claim is discounted at $R_c(z) = \nu_Y(z) + \tau_P$, with slope $R_c' = \ell_1\sigma_w$.

The two claims are priced with the state drifts of equations (B14) and (B15), $\mu_z^{\mathbb{Q}} + \sigma_x\sigma_z$ for the rent claim and $\mu_z^{\mathbb{Q}}$ for the cost claim. Both drifts take the form $b_f - Kz$, with intercepts $b_c = \kappa_z\bar{z} - (\ell_0 - \sigma_w)\sigma_z$ and $b_a = b_c + \sigma_x\sigma_z$, and with the same mean-reversion coefficient

$$K \equiv \kappa_z + \ell_1\sigma_z.$$

Physical persistence is governed by κ_z. Pricing persistence is governed by K, which differs from κ_z because the price of risk varies with z, so the risk adjustment to the drift of z does too. By Feynman–Kac, each claim is a wage-unit perpetuity priced like a bond in a Gaussian affine term-structure model. For $f \in \{a, c\}$,

$$f(z)=\int_0^\infty \exp\{\alpha_f(s)+\beta_f(s)z\}\,ds, \qquad \beta_f(s)=-\frac{R_f'}{K}\big(1-e^{-Ks}\big), \qquad \dot{\alpha}_f=-R_f(0)+b_f\beta_f+\tfrac{1}{2}\sigma_z^2\beta_f^2,$$

with $\alpha_f(0) = 0$. For $K > 0$, the value of claim f is finite if and only if

$$R_f(m_f)-\tfrac{1}{2}\big(R_f'\sigma_z/K\big)^2>0, \qquad m_f=b_f/K.$$

The long-run discount rate, net of the convexity that a random discount rate creates, must be positive. Holding nonzero R_f' and σ_z fixed, with bounded $R_f(0)$ and b_f, it fails as $K \downarrow 0$. The restriction concerns pricing persistence, so a state with no physical mean reversion, $\kappa_z = 0$, can still have finite values when $K > 0$.

The income and capital-gain split. The semi-elasticity f'/f is the value-weighted average of $\beta_f(s)$. At $z = z_0$, fix $K > 0$ and positive benchmark rates $\bar{R}_f = R_f(z_0)$, and expand to first order in the discount-rate slopes around zero. The weights are then those of a perpetuity discounted at $\bar{R}_f$, which gives

$$A \approx -\frac{R_a'}{\bar{R}_a+K}, \qquad C \approx -\frac{R_c'}{\bar{R}_c+K}, \qquad \Gamma^{y\prime} \approx \ell_1\Big[(\sigma_w+\sigma_x)\frac{\bar{R}_a}{\bar{R}_a+K}-\sigma_w\frac{\bar{R}_c}{\bar{R}_c+K}\Big].$$

When the two claims' discount rates are close, $\bar{R}_a \approx \bar{R}_c \approx \bar{R}$, the income slope responds by the fraction $\bar{R}/(\bar{R}+K)$ of its static response $\ell_1\sigma_x$. The fraction falls as K rises relative to $\bar{R}$.

Consider two benchmarks. In the deterministic benchmark $\kappa_z = \sigma_z = 0$, the state is a fixed parameter, and $\Gamma^{y\prime} = \ell_1\sigma_x$ and $\Gamma^{q\prime} = 0$ hold exactly. Changes in the premium slope then pass entirely through Γ^y. The full interaction still has the capital component $\varepsilon_\omega\Gamma^q(C-A)(2\varepsilon-1)$, generated by the change in ε at given leverage.

In the opposite benchmark, $\kappa_z \to \infty$ with the other primitives fixed, suppose $R_a(\bar{z}), R_c(\bar{z}) > 0$ and $x/R_a(\bar{z}) - \delta/R_c(\bar{z}) > 0$, so that the net price stays away from zero and the elasticities remain bounded. Then $\Gamma^{y\prime} \to 0$ and $A - C \to 0$, the income cross derivative tends to zero, and the capital-gain cross derivative tends to $\ell_1\sigma_x\varepsilon_{\omega,\infty}$. When $\ell_1\sigma_x \neq 0$, the whole interaction is earned through expected dollar capital gains. A rise in z that is expected to reverse quickly moves prices and yields only slightly, since $y_z = -(y+\tau_P)\chi$ and χ is small. Physical mean reversion adds $\kappa_z(\bar{z}-z)\chi$ to the expected capital gain, and this product stays finite as χ shrinks. The attribution concerns the interaction, not the level of expected capital gains, which also reflect wage growth, rent growth, and covariance terms and can be negative.

C Quantitative Assessment

This appendix collects quantitative checks that support the calibration and the operating-cost interpretation.

C.1 NOI and rent growth volatility

We compute the standard deviation of NOI growth and of rent growth for the 76 metropolitan areas in the NCREIF research indices with at least two quarters of data, from 2001Q1 to 2023Q4, after winsorizing each series at the 5th and 95th percentiles. NOI growth volatility averages 7.3 percent per quarter and rent growth volatility 2.6 percent, for a mean ratio of 2.74. Figure C1 shows the distribution across metropolitan areas, and most ratios lie between 2 and 3. The ratio is consistent with operating leverage, but it is not a pure measure of it, because operating expenses also respond to local conditions other than rents.

Figure C1: Distribution of Volatility Ratio Across Metropolitan Areas

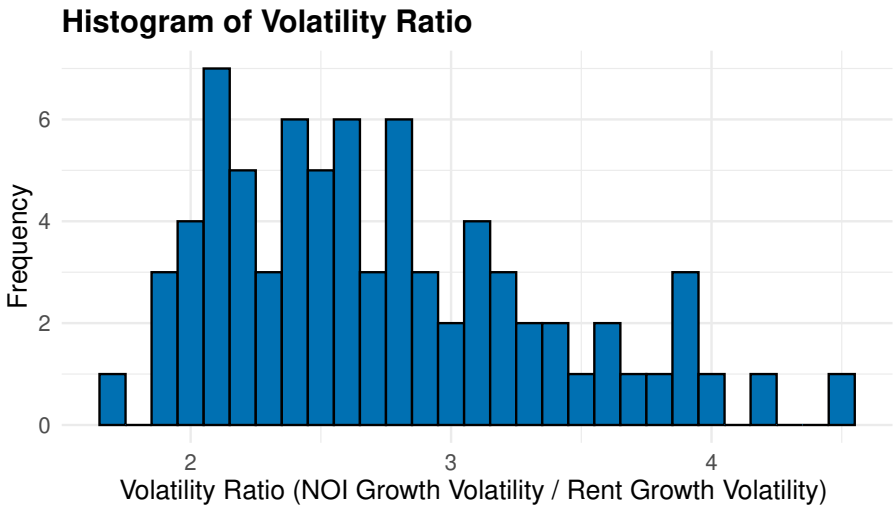


Table C1 reports regressions of NOI growth on rent growth for the same sample. Column (1) is pooled, column (2) adds quarter fixed effects, and both cluster standard errors by quarter. An elasticity above one is consistent with operating leverage, since costs that do not move one-for-one with rents amplify NOI growth.

Within CBSAs, the ratio of NOI-growth to rent-growth volatility and the volatility of total returns both rise with distance from the primary employment center, but neither slope is distinguishable from zero (Table C2).

Table C1: NOI Growth Regressed on Rent Growth (CBSA Level)

	(1)	(2)
(Intercept)	−0.003	
	(0.002)	
Rent Growth	1.341***	1.443***
	(0.066)	(0.060)
Num.Obs.	4839	4839
R^2 Adj.	0.256	0.334
R^2 Within Adj.		0.272
Std. Errors	by quarter	by quarter
FE: Time		X

Notes: NCREIF research NOI and rent growth indices for 79 CBSAs, 2001Q1–2023Q4. NOI growth and rent growth are each winsorized at 5/95 before the series are merged. Column (2) includes quarter fixed effects. Standard errors are clustered by quarter. Significance codes: *** at 1%, ** at 5%, and * at 10%.

C.2 Model-implied cap-rate mapping: derivation and estimation

This subsection derives the calibration equation of Section 5.2 and records the estimation conventions. Write $S_x^2 = \sigma_x^2 + \varsigma_x^2$ for the total instantaneous variance of the rent-to-wage state.

Price and cap rate. With non-tax operating costs δw and the stylized property-tax flow $\tau_P P$, the wage-normalized price solves

$$(\nu_Y + \tau_P)p(x) \;=\; (x - \delta) + \mu_x^{\mathbb{Q}} x p'(x) + \frac{1}{2} S_x^2 x^2 p''(x). \tag{C1}$$

The tax is proportional to the price in this benchmark, so it shifts the discount rate. The no-abandonment solution is linear:

$$p(x) \;=\; \frac{x}{\nu_Y + \tau_P - \mu_x^{\mathbb{Q}}} - \frac{\delta}{\nu_Y + \tau_P}, \qquad \varepsilon_p(x) \;=\; \frac{1}{1 - \theta_P e}, \quad \theta_P \;=\; 1 - \frac{\mu_x^{\mathbb{Q}}}{\nu_Y + \tau_P}, \quad e = \delta/x. \tag{C2}$$

Dividing equation (C1) by the linear price and subtracting the tax yield gives the after-tax income return

$$c(x) \;=\; \frac{x - \delta}{p(x)} - \tau_P \;=\; \nu_Y - \mu_x^{\mathbb{Q}} \varepsilon_p(x). \tag{C3}$$

The tax cancels as an additive yield because it is already capitalized into price, but it still changes the pricing elasticity through θ_P. Subtracting r_f from (C3) gives equation (6) in the main text. Setting $\tau_P = 0$ recovers the no-separate-tax formula.

Table C2: Volatility and distance within CBSAs

	Slope per 10 miles	ZIP codes
σ(NOI growth)/σ(rent growth)	0.040	105
	(0.116)	
σ(total return), pp/qtr	0.093	66
	(0.062)	

Notes: Cross-sectional regressions of ZIP-level volatility on distance from the primary employment center, with CBSA fixed effects. Standard errors clustered by CBSA in parentheses. Volatilities are standard deviations of quarterly log NOI and rent growth for ZIP codes within 30 miles with at least 20 quarters of positive NOI and rent, and of quarterly total returns for the 66 ZIP codes with at least 40 quarters of returns.

Risk–growth decomposition. This calibration holds the aggregate state fixed, so $p_z = 0$ in this one-state calculation. Appendix B.5 gives the full two-state exposure. Applying Itô's lemma to the dollar price $P = w\,p(x)$ and pricing with the SDF gives $m(x) = \pi(x) - \mu_P(x)$, with

$$\pi(x) = \big(\lambda\sigma_w + \xi\varsigma_w\big) \;+\; \varepsilon_p(x)\,\big(\lambda\sigma_x + \xi\varsigma_x\big), \tag{C4}$$

$$\mu_P(x) = \mu_w \;+\; \varepsilon_p(x)\,\big(\mu_x + \sigma_w\sigma_x + \varsigma_w\varsigma_x\big). \tag{C5}$$

The expected excess total return is the risk premium π, and expected excess income returns subtract expected price growth from it. Matching coefficients against (C3) identifies the intercept and slope of the income-return mapping in terms of primitives: $\nu_Y - r_f = \lambda\sigma_w + \xi\varsigma_w - \mu_w$ and $-\mu_x^{\mathbb{Q}} = \lambda_Y\sigma_x + \xi_Y\varsigma_x - \mu_x$, where $\lambda_Y = \lambda - \sigma_w$ and $\xi_Y = \xi - \varsigma_w$ are the wage-unit prices of risk. Thus $\mu_x^{\mathbb{Q}}$ equals the physical drift of the rent-to-wage ratio minus its risk adjustment. In dollar terms this is $-\mu_x^{\mathbb{Q}} = B - g$, with $B = \lambda\sigma_x + \xi\varsigma_x$ and $g = \mu_x + \sigma_w\sigma_x + \varsigma_w\varsigma_x$, as in equation (11). Total returns are affine in ε_p with slope $B = \lambda\sigma_x + \xi\varsigma_x$, the price of exposure in equation (10). This is the slope reported in Section 5.2. The diversified-owner benchmark $\xi = 0$ leaves $\xi_Y = -\varsigma_w$ in the income-return slope. Recovering the primitive λ from $\mu_x^{\mathbb{Q}}$ therefore requires the covariance term $\varsigma_w\varsigma_x$, whereas the total-return slope identifies $\lambda\sigma_x$ directly.

Functional form. As a function of the true elasticity ε_p, expected income returns are exactly affine with both coefficients pinned by $(\nu_Y, \mu_x^{\mathbb{Q}})$. As a function of the measured non-tax expense share e, the mapping (6) is strictly convex when $\mu_x^{\mathbb{Q}} < 0$. It is linear in the measured elasticity $1/(1-e)$ only when $\theta_P = 1$, that is, when $\mu_x^{\mathbb{Q}} = 0$, which also sets the slope to zero. The calibration therefore estimates $(\nu_Y, \mu_x^{\mathbb{Q}})$ on the exact mapping.

Estimation conventions. The least-squares estimator minimizes the sum of squared residuals over all four quartiles. Admissibility requires $\nu_Y > 0$, $\nu_Y + \tau_P > \mu_x^Q$, and $\theta_P e_q < 1$ for every quartile. All rates are quarterly (pp/qtr). Quarterly simple rates are treated as the discrete counterpart of the model's instantaneous rates.

Abandonment. Holding the parameters of the fixed-cost calibration fixed, imposing zero salvage value, and setting $S_x = 2.6\%$ per quarter, we add the decaying option term $A^- x^{n_-}$ and compute income yields directly as $(x - \delta)/p(x) - \tau_P$. Here $n_- < 0$ is the negative root of

$$\frac{1}{2} S_x^2 n(n-1) + \mu_x^Q n - (\nu_Y + \tau_P) = 0.$$

The perpetual-claim $Q_4 - Q_1$ spread is 16.03 basis points, and the option-adjusted spread is 12.35 basis points, an attenuation of 23 percent. This is a sensitivity exercise, in which we do not recalibrate these parameters jointly with the abandonment option. Sagi (2021) describes apartments in the NCREIF Property Index as institutional commercial real estate assets held by tax-exempt institutions and operating above minimum occupancy standards, so the exercise quantifies downside protection rather than observed abandonment in the sample.

C.3 Calibration sensitivity

Table C3 adds a third cost profile to the two of Section 5.2: the non-tax expense share measured directly in each distance quartile. Like the estimated profile, measured shares imply a location-specific cost level δ_d, so the pricing formulas apply location by location. The measured shares rise from 0.297 to 0.324 in the third quartile and fall to 0.306 in the farthest, so this calibration does not reproduce the distance gradient. Its largest income residual is 9.62 basis points, against 1.26 and 0.79 under the two profiles of the main text, and its fitted outer-minus-inner spreads are 3 basis points in income returns and 13 in total returns, against 17 and 61 in the data. Because measured costs rise almost one for one with rents across locations, its implied elasticity of value with respect to rent, 1.11, is close to the measured property-tax elasticity of 1.19 (s.e. 0.07). Table C4 reports the cross-city spread implied by each profile under each volatility input of Section 5.3. It ranges from 0.10 to 0.29 pp/qtr under the two profiles of the main text and reaches 0.45 under measured shares, against an observed 0.26 pp/qtr (s.e. 0.22).

C.4 The volatility parameter σ_x

The cross-city check of Section 5.3 and the abandonment exercise of Appendix C.2 use 2.6% per quarter as a proxy for the volatility of the rent-to-wage ratio $x = r/w$. This is the average per-CBSA standard

Table C3: Calibration sensitivity across cost profiles.

Cost profile	ν_Y (pp/qtr)	μ_x^Q (pp/qtr)	Wedge θ_P	Max. resid. (bp)	B (pp/qtr)	Tax elast.
Fixed non-tax cost ($\eta = 0$)	0.71	−0.233	1.23	1.26	0.75	1.96
Estimated cost elasticity ($\eta = 0.80$)	0.08	−0.387	2.06	0.79	1.27	1.36
Directly measured non-tax shares	0.29	−0.408	1.71	9.62	1.98	1.11

Notes: Least-squares fits of equation (6) to the four quartile excess income returns. The first two rows are the calibrations of Table 10. The third uses the non-tax expense share measured in each distance quartile, 0.297, 0.313, 0.324 and 0.306. B is the slope of the four quartile excess total returns on ε_q with a free intercept. The tax elasticity is the model's elasticity of value with respect to rent across locations: equation (9), averaged across quartiles with cost-share weights, in the first two rows, and the slope of log value on log rent across the four quartiles in the third, where measured shares imply no single cost elasticity. The same slope gives 1.86 and 1.36 for the first two profiles. The measured property-tax elasticity is 1.19 (s.e. 0.07; Table 8).

Table C4: Cross-scale sensitivity to cost profiles and the volatility proxy.

Cost profile	Volatility input	σ_x (%/qtr)	B (pp/qtr)	Model spread (pp/qtr)	Ratio
Fixed non-tax cost	Rent proxy	2.60	0.75	0.17	0.64
Fixed non-tax cost	Rent-to-wage sensitivity	4.36	0.75	0.10	0.38
Estimated cost elasticity	Rent proxy	2.60	1.27	0.29	1.09
Estimated cost elasticity	Rent-to-wage sensitivity	4.36	1.27	0.17	0.65
Measured shares	Rent proxy	2.60	1.98	0.45	1.70
Measured shares	Rent-to-wage sensitivity	4.36	1.98	0.27	1.02

Notes: The model spread is the high-minus-low difference in exposure to aggregate wage shocks, 0.59 percent per square-root quarter, times the price of risk $\lambda = B/\sigma_x$. The observed high-minus-low spread is 0.26 pp/qtr (Newey–West s.e. 0.22). CBSAs are sorted on their return loading on national payroll growth cumulated over the lagged, current, and lead quarters. Cost profiles and B are those of Table C3. The rent proxy is the 2.6% quarterly NCREIF rent-growth volatility. The rent-to-wage sensitivity uses the mean within-CBSA volatility of calendar-quarter-adjusted $\Delta \log(r/w)$, winsorized at 5/95. Ratios divide the model spread by the observed spread.

deviation of quarterly rent growth in the NCREIF research series (Appendix C.1). The underlying growth rates are quarter-over-quarter, and the series is winsorized at the 5th and 95th percentiles before the per-CBSA standard deviations are averaged. The proxy uses total CBSA rent volatility rather than a separately estimated aggregate component of r/w.

As a direct check, we measure the volatility of the rent-to-wage ratio itself. In the merged NCREIF–earnings panel, we construct $\Delta \log x = \log(1 + \text{rent growth}) - \Delta \log(\text{earnings})$, remove calendar-quarter means, winsorize the resulting growth rate at 5/95, compute its time-series standard deviation within each CBSA, and average across CBSAs. This gives 4.36% per quarter. Table C4 reports the cross-scale calculation under both the 2.6% rent proxy and this directly measured volatility of the rent-to-wage ratio.

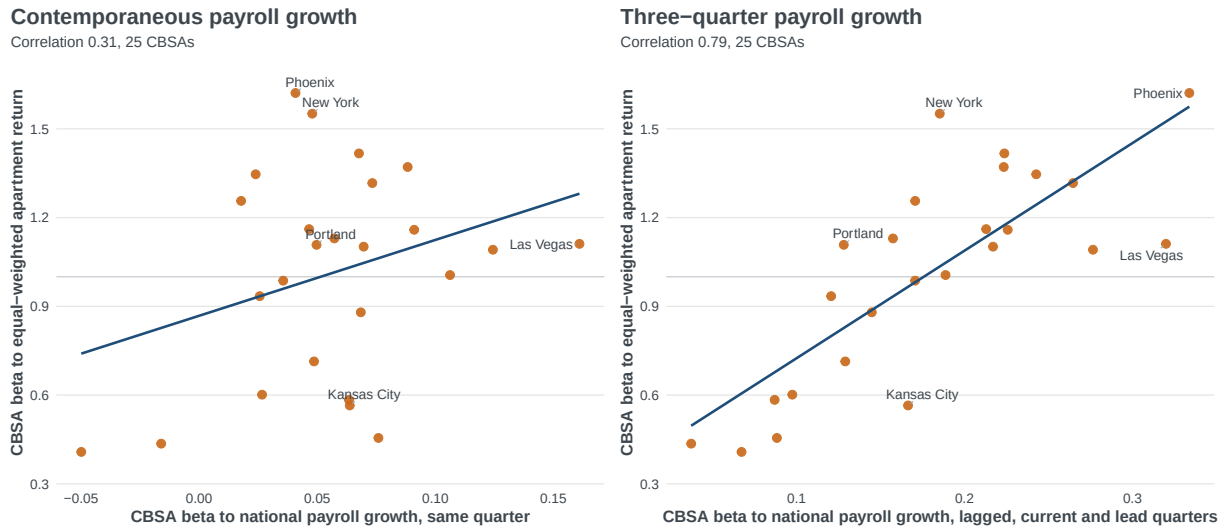
C.5 Cross-city payroll betas: construction and attenuation

This subsection documents the payroll betas used in the cross-city check of Section 5.3. For each CBSA, the payroll-return beta is the slope from a time-series regression of NCREIF apartment total returns on national payroll growth cumulated over the lagged, current, and lead quarters. The labor panel comes from the Quarterly Workforce Indicators (QWI) of the U.S. Census Bureau. Of the 30 CBSAs in the labor panel, Pittsburgh and Cleveland have too few quarters of apartment returns for a payroll beta, which leaves the 28 CBSAs of Section 5.3. Their correlation with the beta to the national NCREIF apartment index is 0.40 in the full sample, and 0.41 when we use the equal-weighted return of the 30 labor-panel CBSAs as the apartment factor, which keeps both beta regressions on the same panel. Figure C2 compares contemporaneous and cumulative payroll growth. The cumulative-beta correlation is 0.79 in the 2001Q1–2010Q2 subsample and 0.41 in the full sample. The exercise relates payroll exposure to the priced apartment factor and leaves open whether payroll risk is priced on its own.

C.6 Time-varying risk premia: $\omega^{NT} \times cay$

This subsection reports robustness checks on equation (12). Table C5 lets the *cay* slope also vary with distance, neighborhood income, and college share. Table C6 interacts *cay* with other ZIP characteristics, and Table C7 counts property taxes as operating costs. We also estimate the interaction for cap rates. Appendix C.7 splits the sample at 2010.

Figure C2: Alternative constructions of the cross-city payroll-return beta. The left panel uses contemporaneous national payroll growth, and the right panel uses the three-quarter cumulative construction adopted in the main text. The vertical axis is the beta to the equal-weighted apartment return of the labor-panel CBSAs. Betas are estimated over 2001Q1–2010Q2 for the 25 CBSAs with at least 20 quarters of returns in that window.



C.6.1 Alternative interactions with *cay*

Table C5 adds interactions of *cay* with distance, log income, and college share. With all three interactions included, the coefficient on $\omega^{NT} \times cay$ is 0.982 (s.e. 0.421), compared with 1.145 (s.e. 0.442) for the baseline on the same 4,316 ZIP-quarters. The estimate remains positive and falls by about 15 percent. In column (5), the income interaction is -0.504 (s.e. 0.235). Neighborhood income can reflect amenity and rent differences that also affect operating leverage. We include its interaction with *cay* as an additional control.

C.6.2 Cap rates

The cap-rate regression uses ZIP-quarter NOI divided by market value as the dependent variable, with ZIP-quarter ω^{NT}, its interaction with *cay*, and CBSA and date fixed effects. On 2,767 ZIP-quarter observations, the interaction coefficient is 0.025 (s.e. 0.037), with standard errors two-way clustered by CBSA and date.

Table C5: Alternative interactions with *cay*

	(1)	(2)	(3)	(4)	(5)
$\omega^{NT} \times cay$	1.145**	1.088**	1.004**	1.123**	0.982**
	(0.442)	(0.418)	(0.435)	(0.443)	(0.421)
Distance $\times$ *cay*		-0.005			-0.002
		(0.006)			(0.005)
ln(Income) $\times$ *cay*			-0.423**		-0.504**
			(0.199)		(0.235)
College Share $\times$ *cay*				-0.020	0.028
				(0.041)	(0.043)
ZIP fixed effects	Yes	Yes	Yes	Yes	Yes
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	4,316	4,316	4,316	4,316	4,316
R^2	0.44108	0.44141	0.44177	0.44114	0.44197

Notes: The dependent variable is total return in the next observed ZIP-quarter. Column (1) is the matched-sample $\omega^{NT} \times cay$ specification; columns (2)–(4) add one alternative interaction at a time; column (5) includes all three. Income is log Census 2000 median household income and college share is from Census 2000; both are standardized across all ZCTAs in the source file. All specifications include ZIP and quarter fixed effects. Standard errors are two-way clustered by CBSA and quarter. Significance codes: *** at 1%, ** at 5%, and * at 10%.

C.6.3 Placebo interactions

Table C6 interacts *cay*, one at a time, with four other ZIP characteristics: distance, market value, NOI, and the number of reporting properties. None of these interactions is distinguishable from zero.

C.6.4 Leverage including property taxes

Table C7 measures leverage with property taxes counted as operating costs. The interaction is 0.540 (s.e. 0.220) for the next observed quarter and positive at every horizon.

C.7 Subperiod stability of the $\omega^{NT} \times cay$ interaction

Table C8 splits the $\omega^{NT} \times cay$ sample at 2010, by formation quarter t. At $h = 1$, the coefficient is 1.331 (s.e. 0.819) before 2010 and 0.860 (s.e. 0.492) after, and each estimate is imprecise. At $h = 4$, the coefficients are positive and economically large, 6.906 (s.e. 1.689) before 2010 and 2.871 (s.e. 1.691) after.

Table C6: Placebo interactions with *cay* (total return, next observed ZIP-quarter)

Interaction	Estimate	N
$\omega^{NT} \times cay$	1.140	4,372
	(0.442)	
Distance $\times$ *cay*	−0.0068	4,372
	(0.0056)	
Market value $\times$ *cay*	9.67e-11	4,372
	(1.31e-10)	
NOI $\times$ *cay*	4.16e-09	4,372
	(1.35e-08)	
Property count $\times$ *cay*	0.0060	4,372
	(0.0188)	

Notes: Each row is a separate regression of equation (12) at $h = 1$ with the listed interaction, ZIP and date fixed effects. Parentheses contain standard errors two-way clustered by CBSA and date.

Table C7: Leverage including property taxes: $\omega \times cay$

	$h = 1$	$h = 2$	$h = 4$	$h = 8$
$\omega \times cay$	0.540	1.162	2.903	5.488
	(0.220)	(0.404)	(0.895)	(2.288)
Observations	4,372	4,184	3,826	3,200

Notes: Estimates of equation (12) with all-in operating leverage ω, which counts property taxes as operating costs, in place of ω^{NT}. ZIP and date fixed effects included. Parentheses contain standard errors two-way clustered by CBSA and date.

Table C8: $\omega^{NT} \times cay$ Predictability: Sub-Period Stability

Dependent variable:	$R_{t\to t+1}$		$R_{t\to t+4}$	
	h=1 (Pre 2010)	h=1 (Post 2010)	h=4 (Pre 2010)	h=4 (Post 2010)
Model:	(1)	(2)	(3)	(4)
Variables				
$\omega^{NT} \times cay$	1.331	0.8595*	6.906***	2.871
	(0.8193)	(0.4920)	(1.689)	(1.691)
Fixed-effects				
ZIP	Yes	Yes	Yes	Yes
Date	Yes	Yes	Yes	Yes
Fit statistics				
Observations	1,490	2,882	1,425	2,401
R^2	0.54373	0.30005	0.76532	0.60259

Notes: Pre-2010 = 1991Q1–2009Q4; Post-2010 = 2010Q1–2019Q2, by formation quarter. The two sub-periods partition the full estimation sample of Table 11. Specification and fixed effects as in Table 11. Parentheses contain standard errors two-way clustered by CBSA and date. Significance codes: *** at 1%, ** at 5%, and * at 10%.