Spatial risk premia in residential real estate

Felipe S. Iachan, Leandro Gorno, and Ana Paula Ruhe

FGV EPGE

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

Abstract

We study spatial variation in residential real estate returns using NCREIF total return data on multifamily properties. Across cities, risk exposure to the national apartment index—measured by beta—is priced in Fama-MacBeth regressions. Within metropolitan areas, excess returns increase by 3.26 basis points per mile per quarter from the primary employment center. The distance gradient survives controls for Census 2000 neighborhood demographics, retaining 62% of its baseline magnitude. We propose an operating leverage mechanism: locations with lower rents bear proportionally higher fixed costs, amplifying their exposure to aggregate shocks. Two tests deliver restrictions consistent with the model. First, the interaction of location-level operating leverage and the consumption–wealth ratio predicts next-quarter returns ($p = 0.022$), with an economic magnitude of 123 basis points over the cycle; placebo interactions are insignificant. Second, a calibration maps the model-implied cap rate spread (42.8 basis points per quarter) against an observed spread of 15 basis points; a simple frictionless benchmark serves as an upper bound.

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

JEL Classification: G12, G19, R30, R31.

1 Introduction

Residential real estate is the largest component of household wealth, yet the spatial structure of its returns remains poorly characterized. Returns on housing differ across metropolitan areas and within them: suburban locations earn higher rental yields and often higher total returns than central locations in the same city (e.g., **???**). Whether these differences reflect risk compensation, heterogeneous growth expectations, or market frictions is an open question. This paper measures the spatial structure of residential real estate returns—including both income and appreciation—and proposes an operating leverage mechanism for the within-city gradient.

We use total return data on multifamily properties from the National Council of Real Estate Investment Fiduciaries (NCREIF), which reports quarterly income returns (net operating income relative to appraised value) and capital appreciation at the property level. These data have important limitations: the sample consists exclusively of institutionally managed apartment buildings, and valuations are appraisal-based rather than transaction-based, introducing potential smoothing in measured return series. We take these constraints seriously throughout and report desmoothing robustness checks in the appendix. On the modeling side, we develop a dynamic version of the monocentric city model in which wages and outside options follow stochastic processes, investors price risk, and an operating leverage channel generates heterogeneous risk exposures within cities.

The data reveal three sets of facts. First, across cities, CBSA-level betas to the national apartment index are priced in Fama-MacBeth regressions, with a point estimate of 1.13 percentage points per quarter (s.e. 0.35; **??**), though the small cross-section ($\sim$30 CBSAs) limits statistical power. Second, within metropolitan areas, excess returns increase monotonically with distance from the primary employment center. In a distance-only Fama-MacBeth specification, the slope is 3.26 basis points per mile per quarter ($t = 5.82$). Controlling for ZIP-level betas, distance retains independent pricing power at 3.8–4.5 basis points per mile (**??**). Third, this gradient is not explained by observable neighborhood characteristics: augmenting the specification with Census 2000 socioeconomic and demographic controls, the distance coefficient retains 62% of its beta-only baseline magnitude ($t = 2.48$, $p = 0.015$; **??**).

We propose operating leverage as the mechanism behind the within-city gradient. Property owners incur maintenance costs—labor, repairs, management—that scale with local wages rather than with property rents. Properties farther from employment centers earn lower gross rents but face similar operating costs, creating thinner net operating margins. A given percentage shock to rents translates into a larger percentage shock to net operating income for thin-margin properties. In the model, this amplification is captured by a pricing elasticity $\varepsilon_p(x)$ that is decreasing in the rent-to-wage ratio x: properties with lower x bear greater effective risk exposure and, in equilibrium, command higher expected returns.

We evaluate two quantitative restrictions of the model. The first exploits the prediction that operating leverage should interact with time-varying aggregate risk prices. When the consumption–wealth ratio *cay* (**?**) is high—indicating elevated risk compensation—properties with high operating leverage (ω) should earn disproportionately higher returns. We find that $\omega \times cay$ predicts next-quarter total returns ($p = 0.022$), income returns ($p = 0.005$), and capital returns ($p = 0.034$), with an economic magnitude of 123 basis points over the cycle (**??**). Placebo interactions using distance or ZIP-level control averages are insignificant (**??**), supporting the interpretation that predictability operates through the operating leverage channel rather than a generic location effect.

The second restriction is a calibration mapping model-implied moments to data. Under the baseline all-in expense ratio, the model implies a cap rate spread of 42.8 basis points per quarter between the outermost and innermost distance quartiles, compared with an observed spread of approximately 15 basis points (**??**). The perpetual-claim benchmark serves as an upper bound; the abandonment option—calibrated to NCREIF parameters—attenuates the spread by less than 1%, indicating that institutional-grade properties operate far above abandonment thresholds. Sensitivity to the expense ratio definition brackets the data.

Our contribution lies at the intersection of urban economics and asset pricing. Whether housing carries risk premia at all remains debated.[1] Prior work on within-city return heterogeneity documents higher rental yields in low-price locations but typically relies on appreciation data and does not test a structural mechanism (**???**). We advance this literature on two margins. First, we measure total returns using institutional data that separately reports income and capital components, allowing us to separate risk premia from expected growth. Second, the monocentric city framework (**????**) has not been extended to study asset returns; we add stochastic fundamentals and investor risk pricing, generating testable predictions for the spatial gradient of expected returns.

The remainder of the paper is organized as follows. Section **??** describes the data, measurement choices, and cross-city return patterns. Section **??** documents the within-metro distance gradient. Section **??** presents the model. Section **??** tests the model's pricing restrictions, including robustness to Census 2000 observables and $\omega \times cay$ predictability. Section **??** maps model-implied moments to data. Section **??** discusses extensions. Section **??** concludes.

[1] See **?**, **?**, **?**, and **?** for evidence of cross-sectional risk pricing; **?** and **?** argue that a hedging motive can generate negative premia.

2 Data and Measurement

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). The NPI primarily reflects privately held U.S. commercial properties acquired at least in part on behalf of tax-exempt institutional investors and held in a fiduciary setting. Affiliated managers report quarterly appraisals, income, and expenses for each property, from which NCREIF computes income returns (net operating income relative to appraised value) and capital returns (change in appraised value) 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 speak directly to the multifamily rental market.

The key advantage of NCREIF over price-only datasets (FHFA, Case-Shiller, Zillow) is that it separately reports income and capital return components. Under a Gordon growth decomposition with constant discount rates,

$$\underbrace{\frac{r}{P}}_{\text{cap rate}} = r_f + \pi - g,$$

$$E\left[\frac{\Delta P}{P}\right] = g,$$

so risk premia π load into rental yields while expected rent growth g shifts the two return components in opposite directions. Analyzing appreciation or capitalization rates in isolation, therefore, conflates risk compensation with growth expectations. NCREIF income returns, constructed as net operating income divided by beginning-of-quarter appraised value, provide the empirical analogue of the model's cap rate.

Three measurement limitations deserve emphasis. First, NCREIF valuations are appraisal-based rather than transaction-based. Quarterly appraisals typically lag market values, inducing positive autocorrelation in measured returns and attenuating measured volatility and covariance. This smoothing is well-documented in the commercial real estate literature. We address it in two ways: by reporting raw (smoothed) results as the baseline, and by testing robustness to Geltner-style AR(1) unsmoothing corrections in Section **??** and Appendix **??**. Under desmoothing, the within-metro distance gradient strengthens, suggesting that the baseline estimates are conservative (Appendix **??**). Second, the

NCREIF sample is institutional: properties are professionally managed and typically larger than the median rental property. Results may not generalize to owner-managed or single-family rental housing. Third, NCREIF applies a confidentiality filter, requiring at least four properties per geographic-quarter cell. This filter limits coverage in early years and in smaller markets.

Sample construction We aggregate property-level data to the ZIP code level and apply the following filters. We restrict the sample to 2000 Q1 onward. NCREIF coverage of residential real estate is limited in earlier years, and reporting depth increases substantially over time. This starting date also aligns the panel with the Decennial Census 2000, which provides an ex ante demographic benchmark for the neighborhood-level controls used in Section **??**. For the cross-city analysis, we require each CBSA to have at least 40 quarters (10 years) of data, yielding 48 CBSAs. For the within-city analysis, we further require each CBSA to have at least 100 ZIP-quarter observations across at least 4 distinct ZIP codes, resulting in 22 CBSAs and 312 ZIP codes. Beta estimation at the ZIP level requires at least 40 quarters of non-missing returns and a distance of at most 30 miles from the primary employment center, restricting the Fama-MacBeth pricing sample to 66 ZIP codes across 18 CBSAs. The four CBSAs present in the within-city descriptive sample but absent from the pricing sample are Boston, Orlando, Philadelphia, and San Francisco; in each case, the binding constraint is the 40-quarter minimum for beta estimation (the 30-mile distance cap is not binding). Returns are winsorized at the 5th and 95th percentiles of the pooled panel distribution to limit the influence of appraisal anomalies.

Excess returns We define excess returns as total returns minus the risk-free rate. The risk-free rate is the quarterly 3-month Treasury bill rate from Kenneth French's data library, compounded to a quarterly frequency.

2.2 Cross-city returns

?? characterizes the cross-sectional distribution of quarterly returns across CBSAs. Statistics are computed over the pooled panel of CBSA-quarter observations after winsorizing both tails at 5%. Mean quarterly total return is 2.10 percentage points per quarter, of which income return accounts for 1.36 pp and capital appreciation for 0.74 pp. Mean quarterly excess return is 1.71 pp (roughly 7.0% annualized), comparable in magnitude to the historical equity premium. The standard deviation of capital return (2.21 pp per quarter) dwarfs that of income return (0.31 pp), so quarter-to-quarter fluctuations in city-level returns are driven overwhelmingly by appreciation shocks rather than cap-rate variation.

Two features are worth highlighting. First, average income return is nearly twice average capital return, indicating that the level of expected returns loads substantially into rental yields—consistent with a Gordon growth decomposition in which expected rent growth is moderate relative to the risk-adjusted discount rate. Second, the cross-city standard deviation of excess returns (2.36 pp per quarter) is the variation that the model must rationalize: cities differ meaningfully in risk compensation. If expected rent growth is similar across metros, these differences should show up primarily in cap rates, a prediction we revisit when interpreting **??**.

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.

Return components and risk Figure **??** illustrates the bivariate relationship between income and capital returns across CBSAs. The most striking feature is a negative correlation, with data points roughly along a line with slope -1. Higher expected growth in certain CBSAs may lead to larger expected capital returns and lower rental yields, but this graph alone cannot establish such a causal link. **?** observe a similar negative relationship for single-family rentals, suggesting the pattern may be a broader feature of real estate markets.

Considerable dispersion around the linear relationship remains. For a given level of expected growth, CBSAs with higher risk exposure should shift to the right, indicating higher excess income returns. The color gradient in Figure **??**, representing CBSA-level betas, shows that higher-beta CBSAs tend to lie above the regression line, consistent with higher total excess returns. However, these cities also tend to have high capital returns, suggesting that growth expectations and risk exposure can be correlated in the cross-section. Disentangling the two requires the regression framework of Table **??**.

?? asks whether the income–capital return tradeoff is systematic. Regressing CBSA-level capital return on excess income return yields a slope of -1.061 (s.e. 0.297, $R^2 = 0.22$), indicating that a one-percentage-point increase in income return is associated with roughly a one-percentage-point decrease in appreciation. The near-unit negative slope means that heterogeneity in income and capital returns approximately cancels in total returns, leaving the cross-section of total returns relatively compressed.

This pattern is consistent with the baseline model's prediction. Under constant discount rates, differences in risk premia across cities load into rental yields; if expected rent growth also varies across cities, it offsets part of the yield variation through the appreciation channel. The nearly one-for-one tradeoff motivates interpreting cap rates as informative about risk compensation, while recognizing that heterogeneous growth expectations complicate a clean decomposition. This bivariate regression

Figure 1: Income vs. capital return components (quarterly, percentage points; excess income return vs. capital return; winsorized at 5%). Color gradient indicates CBSA-level beta to the national NCREIF apartment total excess return.

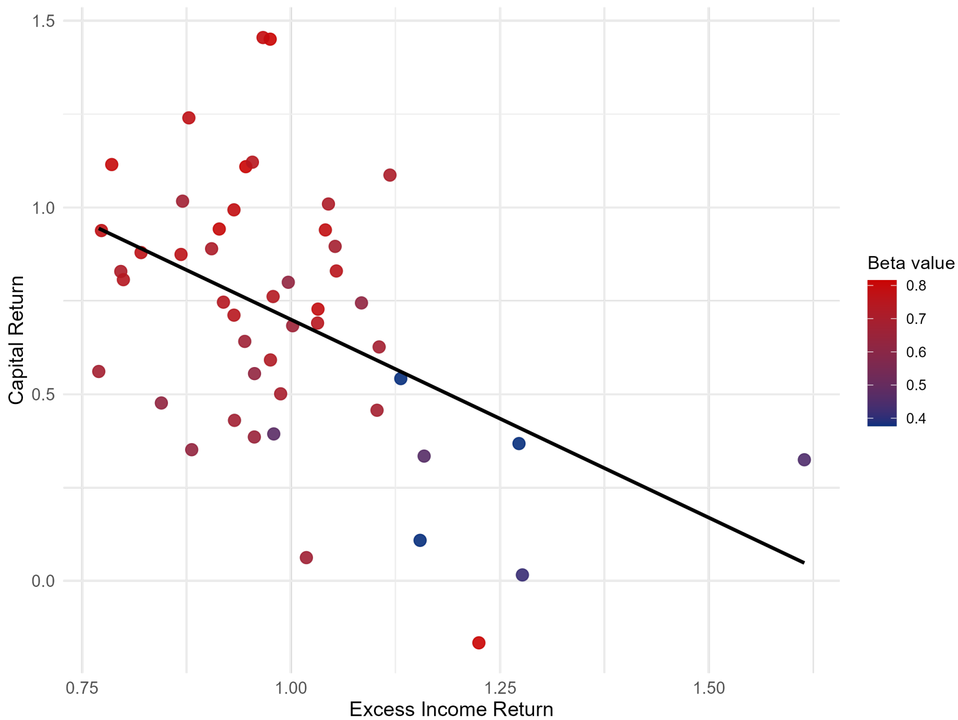

does not account for risk exposure, so the coefficient may be subject to omitted variable bias.

Risk premia We examine the role of heterogeneous aggregate risk exposure using a Fama-MacBeth two-stage procedure. In the first stage, we run time series regressions of each CBSA's quarterly total excess return on the national NCREIF apartment total excess return. Denote the resulting slope β_k^{NCREIF}. In the second stage, for each quarter t we regress the cross-section of CBSA excess returns on β_k^{NCREIF} and collect the slope b_t. The Fama-MacBeth estimate is the time-series average $\bar{b} = \frac{1}{T}\sum_t b_t$, with standard error $\mathrm{SE}(\bar{b}) = \mathrm{sd}(\{b_t\})/\sqrt{T}$ and two-sided t-statistics referenced to the t_{T-1} distribution.

?? reports the results. Column 1 uses NCREIF-based betas. The estimated price of risk is 1.13 pp per quarter (s.e. 0.35), implying that a city with unit beta earns roughly 113 basis points of excess return per quarter (about 4.6% annualized) above the risk-free rate. For robustness, column 2 uses betas constructed from Zillow MSA-level appreciation rates; the estimate is smaller (0.42 pp per quarter, s.e. 0.10), reflecting the fact that Zillow captures only the appreciation component of risk.

The NCREIF–Zillow gap underscores the importance of measuring risk exposure through total returns rather than appreciation alone. A limitation of this test is the small cross-section: the FM

Table 2: Excess income vs. capital return

	Dependent variable:
	Capital Return
Exc. Income Return	-1.061^{***}
	(0.297)
Constant	1.761^{***}
	(0.299)
Observations	48
R^2	0.217
Adjusted R^2	0.200
Residual Std. Error	0.312 (df = 46)
F Statistic	12.730^{***} (df = 1; 46)
Note:	*p<0.1; **p<0.05; ***p<0.01

regressions use approximately 30 CBSAs per quarter, depending on data availability. The result is economically meaningful but should not be over-interpreted on purely statistical grounds.

Table 3: Fama-MacBeth Regression Results

Variable	NCREIF (1)	Zillow (2)
Intercept	0.9330	1.2301
	(0.1869)	(0.1823)
Beta	1.1321	0.4205
	(0.3498)	(0.0954)
Average R^2	0.0732	0.0861

Notes: Quarterly FM estimates. $\bar{b}$ is the time-series mean of quarterly slopes; SE $=$ $\mathrm{sd}(\{b_t\})/\sqrt{T}$; p-values from t_{T-1}. Variables are winsorized at 5% (two-sided) prior to the quarterly cross-sections. Sample restrictions and fixed effects as stated in the table header. T equals the number of non-missing quarters for that coefficient.

2.3 Distance definition

Within-metro analyses require a measure of distance from each ZIP code to a central employment location. Traditional studies typically measure distance to an ad hoc central point such as a city hall or downtown post office. We construct a reproducible, data-driven alternative using Census employment data.

Primary employment center We use the Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics (LODES, version 8) Workplace Area Characteristics (WAC) files for 2022. These provide total employment counts at the Census block level. We aggregate to the ZIP Code Tabulation Area (ZCTA) level and compute employment density (jobs per square mile) using land areas from Census TIGER/Line shapefiles.

Following the modified Giuliano-Small (1991) approach, we identify employment centers as clusters of high-density ZCTAs:

1. ZCTAs with employment density at or above the 90th percentile of the national distribution (approximately 5,886 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. Clusters whose centroids are within 7 miles of each other are merged (single-linkage hierarchical clustering). This radius separates twin cities (e.g., Minneapolis–St. Paul, approximately 8–10 miles apart) while preserving contiguous downtowns.
4. Only clusters with at least 25,000 total jobs are retained.
5. Each center's location is the employment-weighted centroid of its constituent ZCTAs.

The *primary employment center* for each CBSA is the cluster with the highest total employment. Distance is measured as the great-circle (Haversine) distance in miles from each ZIP code centroid to the nearest primary employment center within the same CBSA.

Polycentric metros. In an extension, the algorithm identifies 40 employment centers across 23 CBSAs, of which 11 are polycentric. For each ZIP code, distance is measured to the nearest primary employment center within the same CBSA, so polycentric metros are handled by a distance-to-nearest rule.[2]

[2]Examples of polycentric CBSAs include Dallas–Fort Worth (4 centers), Miami–Fort Lauderdale (3 centers), and San Francisco–Oakland (3 centers). The remaining 12 CBSAs are classified as monocentric.

We use the primary-center measure throughout; results are qualitatively similar under alternative distance definitions.[3]

3 Within-Metro Facts

This section documents the within-city distance gradient in returns, which is the central empirical fact motivating the operating leverage mechanism in Section **??**.

Some stylized facts about within-city heterogeneity in housing returns exist in the literature. **?** analyze ZIP code-level total returns on single-family rentals and show that total returns decline with price tier within cities, driven by the behavior of net rental yields. They find CAPM-style betas between 0.8 and 1.2, suggesting a risk-based explanation, but do not test a structural mechanism for the within-city dispersion. Most previous analyses focus on rent-to-price ratios. As the Gordon growth decomposition makes clear, income return dispersion can arise from differences in both risk exposure and expected growth; without total returns, the two are confounded.

Summary statistics The ZIP code-level sample consists of 312 ZIP codes across 22 CBSAs. Mean quarterly total return at the ZIP level is 1.64 pp, lower than the 1.83 pp average at the CBSA level for the same sample.[4] This gap reflects both lower income returns (1.15 pp vs. 1.23 pp) and lower capital returns (0.50 pp vs. 0.60 pp). Volatility, however, is greater at the ZIP level (standard deviation of total returns: 2.79 pp vs. 2.40 pp), driven by the capital return component. See Appendix **??** for additional detail.

Tables **??** and **??** use the full within-city sample (312 ZIP codes across 22 CBSAs).

Distance gradient of returns **??** documents whether returns vary with distance from the primary employment center, with both tails winsorized at 5%. Panel A shows that mean quarterly excess total return rises from 0.91 pp in the closest distance quartile (Q1) to 1.41 pp in Q4, a spread of roughly 50 basis points per quarter (about 2.1 percentage points annualized). Panel B reveals that excess income return increases from 0.71 pp (Q1) to 0.86 pp (Q4), a spread of about 15 basis points per quarter. Both gradients are monotonic across Q1–Q3; Q4 is slightly below Q3 for total returns and essentially equal for income returns.

[3] An earlier version of this paper used hand-curated CBD coordinates for each CBSA. The correlation between primary-center distance and the legacy measure is 0.97. The Fama-MacBeth distance slope under the legacy definition is 3.05 bp per mile per quarter ($t = 5.24$), compared with 3.26 bp ($t = 5.82$) for the primary-center measure.

[4] The difference arises because we compare ZIP-level and CBSA-level aggregates within the restricted sample used for within-city analysis; values differ from those in **??**, which uses the broader 48-CBSA cross-city sample.

The operating leverage mechanism predicts a pattern broadly consistent with these data: distant properties earn lower gross rents relative to operating costs, creating thinner margins that amplify risk exposure and require higher expected returns. The income-return gradient is particularly relevant because the model predicts the operating leverage premium loads into cap rates. The total-return gradient additionally reflects capital appreciation differentials, which may include both risk compensation and heterogeneous expected growth within metros.

In a distance-only Fama-MacBeth specification (without beta controls), the continuous slope is 3.26 basis points per mile per quarter ($t = 5.82$, $p < 0.001$; 94 cross-sections, 66 ZIPs).[5][6] After adding ZIP-level betas, the distance coefficient is 3.8–4.5 bp per mile per quarter (**??**, Panel B), indicating that distance retains independent pricing power beyond measured risk exposure.

Table 4: Descriptive Statistics, NCREIF at ZIP code aggregation

Distance group	mean	sd	median	min	max	annualized mean
Panel A: Excess Total Return						
1	0.9118	2.8400	0.915	-5.477	7.550	3.6974
2	1.3175	2.8805	1.250	-5.477	7.550	5.3750
3	1.5629	3.0329	1.430	-5.477	7.550	6.3999
4	1.4109	2.9875	1.430	-5.477	7.550	5.7640
Over 30 miles	1.4748	2.8886	1.450	-5.377	7.550	6.0312
Panel B: Excess Income Return						
1	0.7079	0.4522	0.750	-0.600	1.670	2.8617
2	0.7913	0.4734	0.870	-0.600	1.670	3.2031
3	0.8595	0.5185	0.950	-0.600	1.670	3.4826
4	0.8630	0.5105	0.960	-0.600	1.670	3.4970
Over 30 miles	0.8210	0.5375	0.960	-0.470	1.660	3.3246

Note: values are quarterly returns in percentage points (except the last column). The "Over 30 miles" row is included for completeness; Fama–MacBeth regressions in Table **??** restrict to ≤ 30 miles.

A concern with **??** is that distant ZIP codes might cluster in high-return CBSAs or high-return quarters, generating a spurious distance gradient. **??** addresses this by reporting residuals from regressions with CBSA$\times$quarter fixed effects. After removing metro-quarter means, the gradient persists: mean residualized total return rises from -0.17 pp (Q1) to $+0.13$ pp (Q4), and residualized income

[5] This specification is not separately tabulated; it uses the same 66-ZIP, 18-CBSA sample as Table **??**.

[6] Tables **??** and **??** use the full 312-ZIP within-city sample, including properties beyond 30 miles. The Fama–MacBeth regressions in Table **??** restrict to ZIPs with $\geq$ 40 quarters of data and $\leq$ 30 miles from the primary employment center (66 ZIPs, 18 CBSAs).

return rises from -0.05 pp (Q1) to $+0.05$ pp (Q4).

This documents the distance gradient as a within-metro, within-quarter regularity. The monotonic pattern in both panels rules out the possibility that composition effects across CBSAs or time periods drive the result. The magnitudes are smaller than the raw values (as expected, since the fixed effects absorb CBSA-level return variation), but the sign and ordering are preserved.

Table 5: Descriptive Statistics, Residuals of FE regressions

Distance group	mean	sd	median	min	max	annualized mean
Panel A: Residual Total Return						
1	-0.1720	1.3736	-0.1665	-5.3586	5.8720	-0.6862
2	-0.0281	1.4820	0.0000	-9.5565	8.4062	-0.1124
3	0.1155	1.5316	0.0000	-7.8725	7.1797	0.4628
4	0.1290	1.5639	0.0000	-7.4663	5.8481	0.5169
Over 30 miles	0.0669	1.6581	0.0000	-6.0853	5.8087	0.2679
Panel B: Residual Income Return						
1	-0.0512	0.1323	-0.0500	-0.7310	0.7144	-0.2046
2	-0.0126	0.1229	0.0000	-0.6000	0.4600	-0.0505
3	0.0247	0.1280	0.0157	-0.5914	0.6314	0.0988
4	0.0545	0.1302	0.0500	-0.5650	0.7500	0.2183
Over 30 miles	0.0400	0.1633	0.0530	-0.5500	0.4900	0.1601

Note: values are residuals of regressions of quarterly returns in percentage points against CBSA-time fixed effects.

Measures of risk The sample for beta estimation is restricted to 66 ZIP codes with at least 40 quarters of data and located within 30 miles of the primary employment center. To assess heterogeneous exposure to aggregate risk, we estimate ZIP-level betas from time-series regressions of ZIP excess total returns on national and CBSA-level returns, running three specifications: national beta only, CBSA beta only, and both jointly.

?? summarizes the distribution of estimated betas. In univariate regressions, the mean national beta is 0.80 and the mean CBSA beta is 0.96, indicating that ZIP-level returns comove substantially with both aggregate and local indices.[7] In the joint specification, national and CBSA betas are highly collinear—both are estimated from appraisal-based returns that share common smoothing patterns—

[7] Mean betas below unity reflect the fact that the national NCREIF index is a value-weighted aggregate across all property types and metropolitan areas, not the equal-weighted average of the ZIP codes in our sample. Multifamily properties in these metros need not have unit exposure to the broader index.

making it difficult to separately identify the two exposures. We report the joint decomposition in Appendix **??** and focus on the univariate measures for the pricing tests that follow.

This collinearity does not mean national shocks are irrelevant for pricing—they may be priced through the city-level channel—but it does mean that separately identifying national and local risk exposures at the ZIP level requires caution.

Table 6: Summary Statistics of Beta Measures

Measure	mean	sd	median	min	max
National only	0.80277	0.18678	0.81527	0.11133	1.21221
CBSA only	0.95973	0.13206	0.97559	0.37074	1.20775
Joint National	0.03061	0.44420	0.12652	-1.77334	0.86798
Joint CBSA	0.93143	0.43286	0.86230	0.26222	2.46063

The Fama-MacBeth pricing regressions use the restricted beta-estimation sample (66 ZIPs, 96 quarters).

Within-city Fama-MacBeth regressions For each quarter t, we run cross-sectional regressions of ZIP-level excess total returns on beta measures and distance controls, without an intercept. Specifically, for the joint national and CBSA beta case we estimate:

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

$$R^e_{i,t} = \lambda^N_t \beta^N_i + \lambda^C_t \beta^C_i + \delta_t \, d_i + \epsilon_{i,t} \tag{2}$$

$$R^e_{i,t} = \lambda^N_t \beta^N_i + \lambda^C_t \beta^C_i + \sum_{q=2}^{4} \delta^q_t D^q_i + \epsilon_{i,t}, \tag{3}$$

where i is a ZIP code, d_i is the distance in miles from ZIP i to its primary employment center, and D^q_i is an indicator for distance quartile q. We report Fama-MacBeth time-series averages of the quarterly slopes with standard errors computed as $\text{SE} = \text{sd}(\{\hat{\delta}_t\})/\sqrt{T}$ and two-sided t-statistics referenced to the t_{T-1} distribution.

?? is the key within-city pricing table.

Panel A (beta only). CBSA beta alone commands a quarterly premium of 1.61 pp per quarter (s.e. 0.21); in the joint specification, CBSA beta is 1.58 pp (s.e. 0.21) and national beta is 1.69 pp (s.e. 0.31), both positive and significant.

Panel B (beta + continuous distance). Adding distance yields a coefficient of 3.8–4.5 basis points per mile per quarter (s.e. ≈0.6–0.7 bp), highly significant across all three beta configurations. Both

Table 7: Fama-MacBeth Regression Results - Risk Premiums

	National	CBSA	Joint
No Distance			
Beta National	1.9342	–	1.6885
	(0.2581)		(0.3084)
Beta CBSA	–	1.6086	1.5754
		(0.2111)	(0.2094)
Continuous Distance			
Beta National	1.5960	–	1.2437
	(0.2575)		(0.3056)
Beta CBSA	–	1.2910	1.2613
		(0.2061)	(0.2056)
Distance	0.0382	0.0429	0.0446
	(0.0064)	(0.0067)	(0.0066)
Distance Groups			
Beta National	1.6029	–	1.3833
	(0.2688)		(0.3525)
Beta CBSA	–	1.2705	1.2397
		(0.2146)	(0.2131)
Dist. Group 2	0.2119	0.2581	0.2170
	(0.1043)	(0.0993)	(0.1078)
Dist. Group 3	0.5567	0.6311	0.6525
	(0.1054)	(0.1066)	(0.1071)
Dist. Group 4	0.5593	0.6435	0.6464
	(0.1229)	(0.1277)	(0.1269)

Notes: Quarterly FM estimates. $\bar{\lambda}$ is the time-series mean of quarterly slopes; SE $=$ sd$(\{\lambda_t\})/\sqrt{T}$; p-values from t_{T-1}. Variables are winsorized at 5% (two-sided) on the pooled panel prior to estimation. Distance is measured in miles to the primary employment center (LEHD-LODES). T equals the number of non-missing quarters for that coefficient.

beta measures remain positive and significant but are attenuated relative to Panel A, consistent with distance capturing part of the variation previously loaded onto betas.

Panel C (beta + distance groups). Distance-group dummies reveal a monotonic gradient relative to Q1: groups 3 and 4 earn 56–65 basis points more per quarter, both highly significant (s.e. 0.11–0.13 pp). Group 2 earns 21–26 basis points more, marginally significant (s.e. ≈0.10 pp).

The central finding is that distance retains independent pricing power after controlling for both beta measures. Because estimated ZIP-level betas are noisy—due to short time series and appraisal smoothing— distance, which is measured without error, can serve as a more reliable proxy for the operating leverage amplification captured by the pricing elasticity $\varepsilon_p(x)$ in the model. Appendix **??** reports the with-intercept specification: distance remains significant (≈3.5 basis points per mile per quarter), while betas lose significance, as expected when the intercept absorbs the return level that betas explain.

A limitation is the small cross-section: up to 66 ZIP codes (mean 42 per quarter) across 18 CBSAs. The Fama-MacBeth procedure averages over 94 quarters, providing time-series power, but the number of assets per cross-sectional regression is limited.

4 A theoretical framework for residential rental returns

Our model builds on classical urban economics but introduces three key features needed to study asset pricing: (1) wages and outside options evolve stochastically over time, creating uncertainty about future rents; (2) investors care about risk and demand compensation for bearing it; (3) different locations have different risk exposures, generating predictions for the cross-section of 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). These workers are highly mobile, moving freely across locations to maximize their instantaneous welfare. In contrast, real estate investors are risk-averse and have substantial wealth. They invest in real estate as marginal pricers with deep pockets. 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. [8]

Uncertainty The uncertainty in the model is driven by independent stochastic processes. The process $\{B_t\}$ represents aggregate shocks, which are common across all locations. Additionally, each city k is 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

[8] See **?** for evidence of limited financial integration in residential rental markets in France.

aggregate process $B_{t+s} - B_t$ and the local process $B^k_{t+s} - B^k_t$ are uncorrelated normal random variables with mean 0 and variance s.

This two-factor structure allows us to distinguish premia for aggregate risk (common across cities) from premia for idiosyncratic risk (city-specific). This suggests an empirical analysis that tests whether both command premia in residential real estate.

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.

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

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

where μ_Ψ and σ_Ψ are the drift and volatility coefficients, respectively, while dB_t is the increment of the standard Brownian motion which represents aggregate uncertainty. The baseline treats Ψ_t as exogenous for tractability; Appendix **??** micro-founds it as the equilibrium inclusive value in a closed system of K cities with logit sorting and elastic housing supply, where σ_Ψ equals the population-weighted average of city wage volatilities rather than a free parameter.

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

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

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: $a_{k,t} = a_k$.[9] This assumption simplifies the model by focusing the analysis on the interaction between wages, rents, and commuting costs.

Equilibrium Rent Determination Equilibrium rent reflects the interplay between wages, housing costs, and the spatial distribution of workers. 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 that follows a geometric Brownian motion. The exponent $\frac{1}{1-\alpha}$ 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 $\frac{1}{1-\alpha}$%, amplifying wage risk into rent risk.

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 rent gradient—how rent changes with distance from the CBD—is influenced by commuting costs and the elasticity of substitution between consumption and housing. Wage risk passes through to rents: higher wage volatility implies higher rent volatility.

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:

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

[9] 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 model allows a simple endogenous amenity extension, which is available upon request.

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:

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

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 in this model play a key role in determining asset prices, including the price of housing. They generate a stochastic discount factor (SDF), which is essential for pricing assets in a dynamic setting. The SDF for savers in city k at time t is modeled as

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

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. Section **??** relaxes the constant-λ^k 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, with

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

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.

Pricing formulas and testable implications The pricing formula admits an intuitive interpretation as 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 **??**. 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).

This framework predicts higher cap rates in locations with greater exposure to aggregate risk, consistent with our Fama-MacBeth results in **??**. 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 **??** is consistent with this prediction, documenting a positive and significant relation between CBSA-level beta estimates and average excess returns.

The baseline model is consistent with cross-city heterogeneity in returns through differential exposure to aggregate and local risk factors. Tables **??** and **??** are consistent with these predictions. However, it predicts constant returns within each city—a stark contrast with the data. **??** documents a systematic distance gradient: properties farther from the CBD exhibit higher cap rates and higher realized returns. This within-city pattern cannot arise from the baseline model, where all locations in city k share identical risk exposures $(\sigma_{r,k}, \varsigma_{r,k})$.

The key to generating within-city heterogeneity lies in recognizing that while gross rents decline with distance from the CBD, operating costs—maintenance, management, property taxes—remain relatively stable per property. Properties farther from the CBD thus have lower gross rents but similar maintenance costs, creating thinner operating margins. These thin margins amplify risk exposure: a given percentage shock to gross rents translates into a larger percentage change in net operating income for suburban properties. We now formalize this operating leverage channel.

4.2 A model with operating costs

The empirical puzzle The baseline model accounts for why properties in different cities earn different returns. San Francisco properties, for instance, might earn higher returns than Detroit properties if San Francisco wages are more exposed to aggregate shocks. This cross-city heterogeneity arises naturally 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* San Francisco should earn identical returns. After all, properties at different distances from the CBD in the same city 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.

This prediction is sharply contradicted by the data. **??** documents a systematic within-city gradient: properties in the quartile closest to the CBD earn average quarterly excess returns of 0.91%, while those in the fourth quartile earn 1.41%—a difference of about 50 basis points per quarter, or about 2.0 percentage points annualized. This gap is driven in part by higher cap rates (excess income returns) in distant locations: 0.71% versus 0.86% quarterly. Moreover, **??** shows this distance gradient persists even after controlling for differential exposure to CBSA-wide risk: each mile from the CBD commands a premium of roughly 3.8–4.5 basis points per quarter.

We now introduce a realistic feature of housing markets that generates this within-city heterogeneity: operating costs.

Operating costs in practice Real estate investors incur ongoing costs to maintain properties and generate rental income. These include routine maintenance and repairs, property management fees, utilities not paid by tenants, property taxes, and insurance. Importantly, many of these costs, particularly labor-intensive services like maintenance and property management, are better described as fixed relative to gross rents: they scale with local input prices (wages, materials) rather than with the property's rental income. The model's $\delta w_{k,t}$ formulation captures this by tying costs to wages; the key economic assumption is not literally that all costs are wage-determined, but that operating costs do not rise proportionally with gross rents across locations within a city.

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: replacing appliances, repairing plumbing, landscaping, managing tenant relationships. 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 – a $2\times$ difference despite only 25% difference in gross rents.

Now imagine an aggregate shock reduces all gross rents in the city by 10%. The downtown property's NOI falls from δ to 0.9δ (a 10% decline), but the suburban property's NOI falls from 0.5δ to 0.4δ (a 20% 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 amplified volatility should translate into higher risk exposure and, consequently, 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

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

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 **??**. 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

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

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

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

The key insight is that properties with lower x (farther from 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). Under risk-neutral pricing (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:

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

where $\nu_{Y,k}$ is the risk-adjusted discount rate (the drift of the wage-denominated stochastic discount factor). In the appendix we show that x follows $dx/x = \mu_{x,k}\,dt + \sigma_{x,k}\,dB_t + \varsigma_{x,k}\,dB^k_t$ (Appendix **??**), and that expected returns load on $\varepsilon_p(x)$.

The crucial quantity for risk exposure is the *pricing elasticity*

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

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

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

Several features are immediate. First, $\varepsilon_p(x)$ is *decreasing* in x (provided $\nu_{Y,k} > \mu^{\mathbb{Q}}_{x,k}$). Properties farther from the CBD have lower x, hence higher elasticity. The condition $\nu_{Y,k} > \mu^{\mathbb{Q}}_{x,k}$ requires that the risk-adjusted discount rate exceed the risk-neutral growth rate of the rent-to-wage ratio, so that the perpetual claim has finite value. This is the standard transversality condition, analogous to requiring $r > g$ in a Gordon growth model. Second, $\varepsilon_p(x) = 1/(1 - \bar{\delta}_k/x)$ follows the inverse of the operating margin. Thus, thinner margins generate larger elasticities.

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

$$\mathbb{E}_t[R_{k,d,t+1}] - r_f = (\lambda\sigma_{w,k} + \xi\varsigma_{w,k}) + \varepsilon_p(x_{k,d,t}) \cdot (\lambda\sigma_{x,k} + \xi\varsigma_{x,k}),$$

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 model implies that each location's effective risk exposure is $\varepsilon_p(x_i) \cdot \sigma_{x,k}$, which varies across locations within a city and over time, with the state. Estimated ZIP-level betas are noisy proxies for this object. Distance, measured without error, captures the monotone relationship between location and operating leverage and can proxy for the risk-exposure gradient. Its pricing power in Table **??** is therefore consistent with the model.

The key prediction: *within a city, properties with lower x have higher $\varepsilon_p(x)$, and thus bear greater risk exposure and earn higher expected returns*. This rationalizes the distance gradient documented in

??. Importantly, this mechanism operates through *amplification* of existing systematic risks rather than through differential exposure to separate risk factors. All locations within city k face the same underlying wage and amenity shocks; 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 risk exposure.

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 $\varepsilon_p(x_A) < \varepsilon_p(x_B)$: the lower-amenity property bears greater risk exposure and earns higher expected returns, despite being equidistant from the CBD.

This generalizes the within-city gradient beyond the radial dimension. Operating leverage generates return heterogeneity whenever gross rents vary—whether through distance, amenities, or other neighborhood characteristics. The testable prediction is that properties with lower gross rent relative to operating costs (lower x) earn higher cap rates and higher expected returns, regardless of why their gross rents are low.

Limitations The model abstracts from several potentially important features. We assume operating costs scale exactly with wages, which is illustrative but not essential—what matters is that costs don't scale proportionally with gross rents. We also abstract from heterogeneity in property quality and neighborhood effects beyond amenities, all of which likely contribute to within-city return variation. The model aims to isolate the operating leverage mechanism, not to provide a complete account of all sources of within-city heterogeneity.

An additional limitation is that our baseline model features constant risk premia and interest rates. In models with time-varying discount rates (e.g., habit formation, long-run risks, intermediary constraints), risk premia can show up through expected appreciation as well as cap rates. Section **??** develops a time-varying extension that addresses this.

Testable implications and model fit The model generates several testable predictions for within-city return dispersion:

(1) Distance gradient: Since x declines with distance from the CBD (due to commuting costs), the model predicts higher cap rates and higher expected returns for distant properties. The data are consistent with this pattern: moving from the closest to the fourth quartile of distance increases quarterly excess total returns by about 50 basis points per quarter and excess income returns by about 16 basis points per quarter (**??**).

(2) Distance premium beyond beta: The Fama-MacBeth results in **??** show that distance commands a premium even after controlling for ZIP-level betas estimated against the national and CBSA-level apartment return indices. The continuous distance specification estimates a coefficient of 3.8–4.5 basis points per mile per quarter across all three beta configurations (national only, CBSA only, and joint). Standard errors-in-variables reasoning applies: when the true risk loading is $\varepsilon_p(x_i) \cdot \sigma_{x,k}$ and the estimated beta $\hat{\beta}_i$ is a noisy measure of this loading, a second-stage regression of returns on $\hat{\beta}_i$ attenuates the beta coefficient toward zero. Distance, as a noise-free correlate of $\varepsilon_p(x_i)$, retains explanatory power because it is not subject to the same attenuation. This is a measurement argument, not a claim that distance captures a risk factor distinct from the one measured by beta.

(3) Complementarity with cross-city model: The baseline model explains variation in $(\lambda^k, \xi^k, \sigma_{r,k}, \varsigma_{r,k})$ across cities. The operating cost extension adds within-city variation through $\varepsilon_p(x)$, which differs across locations in the same city. Together, they rationalize both dimensions of spatial heterogeneity in the data.

Real options: abandonment and development The benchmark solution assumes properties are perpetually maintained. In reality, investors can abandon properties when NOI becomes sufficiently negative, and developers can build new properties when profitability is sufficiently high. These real options introduce free boundaries: an abandonment threshold x_k^L (low x) and a development threshold x_k^H (high x).

In the continuation region ($x_k^L < x < x_k^H$), the price function for a developed property takes the form

$$p_k(x) = \frac{x}{\nu_{Y,k} - \mu_{x,k}^{\mathbb{Q}}} - \frac{\delta}{\nu_{Y,k}} + A \cdot x^{n_-}, \quad n_- < 0,$$

where the third term is a decaying option component. The constant A and exponent n_- are determined by the pricing equation, and the boundaries x_k^L, x_k^H are pinned down by value-matching and smooth-pasting conditions.

As $x \to \infty$ (very high gross rents), the option term vanishes ($x^{n_-} \to 0$), and $\varepsilon_p(x)$ converges to the benchmark formula. Near the abandonment threshold $x \approx x_k^L$, the option term flattens the price function, reducing $\varepsilon_p(x)$. Economically, the option to abandon provides downside protection, attenuating risk exposure for properties with very thin margins. However, this attenuation is confined to the lower tail of the x distribution. The main operating leverage gradient persists throughout most of the x support.

Section **??** provides a quantitative assessment of the abandonment option's effect on the model-implied gradient. The detailed derivation of the optimal stopping problem, including both abandonment and development options, is provided in Appendix **??**.

5 Evidence and Validation

5.1 Cross-City Risk Pricing

The model's first prediction concerns the cross-section of city-level returns: metropolitan areas with greater aggregate risk exposure should earn higher expected returns, with the premium loading primarily into cap rates. Table **??** presents Fama-MacBeth estimates from quarterly cross-sections of CBSA-level excess total returns on estimated betas. The NCREIF risk price estimate is approximately $\hat{\lambda} \approx 1.13$ percentage points per quarter, positive and economically meaningful. The patterns in Table **??** are consistent with the model's cross-city prediction that cities bearing greater systematic risk exposure command higher average returns.

The model also implies that, in a constant-discount-rate benchmark, risk premia should load through cap rates rather than expected appreciation. Table **??** is consistent with this prediction: excess income returns and capital returns exhibit a nearly one-for-one negative relationship across CBSAs, with a slope of -1.061 (s.e. 0.297, $R^2 = 0.22$). Cities with higher cap rates tend to have lower subsequent appreciation, and vice versa. The near-unit negative slope means that heterogeneity in the two return components largely offsets in total returns; the cross-sectional variation in total returns that remains—and that Table **??** prices with beta—reflects genuine risk compensation rather than expected-growth composition effects. Figure **??** illustrates the tradeoff: the color gradient is consistent with higher-beta cities earning higher total excess returns, though high-beta cities also exhibit high capital returns, complicating a clean decomposition into risk and growth components.

The within-city counterpart is documented in Tables **??** and **??**. Moving from the innermost to the outermost distance quartile is associated with approximately 50 basis points per quarter higher total excess returns and about 15 basis points higher cap rates. Controlling for CBSA-level beta in Fama-MacBeth cross-sections, each additional mile from the primary employment center commands a premium of 3.8–4.5 basis points per quarter (no-intercept specification, Table **??**). These patterns are consistent with the model's prediction from Sections 4.1–4.2 that operating leverage amplifies city-wide risk for thin-margin, peripheral properties. The mechanism is amplification through the pricing elasticity $\varepsilon_p(x)$—not differential exposure to a distinct risk factor—so the distance premium persists even after beta is held fixed.

Within CBSAs, ZIP-level CBSA betas increase with distance (slope 0.0065 per mile, $t = 2.00$, $p = 0.062$, $N = 66$), consistent with the operating leverage restriction, though this test is power-limited; diagnostics and alternative visualizations are in Appendix **??**.

5.2 Robustness of the Within-Metro Gradient

The distance gradient is the paper's primary within-city fact. The principal threat to a risk-based interpretation is that distance may proxy for observable neighborhood quality: low-rent, peripheral locations may simply be disadvantaged on socioeconomic or demographic dimensions that command a return premium for reasons unrelated to operating leverage. This subsection addresses that concern in three steps: Census 2000 neighborhood controls, measurement correction for appraisal smoothing, and pre-pandemic replication.

5.2.1 Urban Observables: Census 2000 Controls

Our baseline gradient uses CBSA beta as the only cross-sectional control. Table **??** replicates the Fama-MacBeth distance regression with an augmented set of neighborhood quality controls from the Decennial Census 2000, addressing three econometric concerns that apply to the ACS 2015 alternative.

Identification improvements. First, Census 2000 *anchors* neighborhood characteristics to the beginning of the return panel (2000 Q1), avoiding the forward-looking bias that arises when post-sample gentrification is used as a control. If peripheral neighborhoods improved over the sample period because of the same secular suburbanization forces that drive returns, using 2015 ACS controls over-controls for the very mechanism under study. Second, we replace log total population with *true population density*—log persons per square mile—computed from Census TIGER/Line 2000 land area. ZCTA-level population counts conflate size with density because ZCTAs expand physically with distance from the employment center, creating a spurious negative correlation between "density" and the CBD. Third, we add race and ethnicity composition controls (% Non-Hispanic White, % Black, % Hispanic), following **?** and **?**, who document strong racial sorting gradients within metropolitan areas.

Results. Column (2) of Table **??** reproduces the baseline gradient with beta controlled: 4.25 bp/mile ($t = 6.31$, $p < 0.001$), averaging over 42 ZIPs and 94 quarters. Column (3) adds SES controls (log median income 1999, college share, true population density): the gradient falls to 2.64 bp/mile ($t = 3.71$), retaining 62% of the baseline. Column (4) adds racial composition controls alongside SES: the gradient rises to 3.71 bp/mile ($t = 3.62$), retaining 87%. Column (5) adds poverty rate and housing vintage to the full control set: the gradient is 2.65 bp/mile ($t = 2.48$, $p = 0.015$), retaining 62% of the beta-controlled baseline.

The distance gradient survives all three augmented specifications (retained fraction $\geq$ 50%, $p < 0.10$) when neighborhood quality is anchored to the panel start. Among the individual controls, college share is the only statistically significant confounder: ZIPs with higher educated populations have lower excess returns, consistent with lower operating risk in high-amenity locations. Racial composition variables are individually insignificant once socioeconomic status is controlled. A with-intercept kitchen-sink specification (Column 6) collapses due to near-collinearity among the controls; we do not view this as evidence of fragility but rather a standard regularization failure in a short cross-section with many correlated predictors.

The economic interpretation is straightforward: moving 10 miles farther from the primary employment center is associated with approximately 26 bp/quarter higher excess returns, net of observable differences in income, education, density, race, poverty, and housing vintage measured at the panel start. The gradient is not explained by observable neighborhood sorting.

5.2.2 Measurement Robustness: Appraisal Smoothing

NCREIF returns are based on periodic appraisals, which induces positive return autocorrelation and attenuates volatility and covariances—the classic **?** smoothing problem. We apply a Geltner-style AR(1) unsmoothing correction, replacing each ZIP's observed return r_t^{obs} with the corrected return $r_t = (r_t^{\text{obs}} - \hat{\rho}\, r_{t-1}^{\text{obs}})/(1 - \hat{\rho})$, where $\hat{\rho}$ is estimated from the ZIP's own return series (ZIP-specific α) or pooled across ZIPs (pooled $\alpha = 0.482$, median smoothing parameter). The distance gradient survives and slightly strengthens under all variants: the pooled-α estimate is 3.62 bp/mile ($t = 3.71$, $p < 0.001$) compared with the baseline of 3.26 bp/mile, consistent with appraisal smoothing *attenuating* the observed gradient—baseline estimates are conservative. The $\omega \times cay$ sign is stable across all desmoothing specifications, and significance is retained under ZIP-specific α corrections ($p = 0.064$) and fixed $\alpha = 0.6$ and $\alpha = 0.8$ ($p < 0.10$); the pooled-α correction weakens significance ($p = 0.153$), reflecting sensitivity to the homogeneity assumption on smoothing rates. Full desmoothing diagnostics and results are in Appendix **??**.

5.2.3 Pre-2020 Robustness: Pandemic Amplification

Pandemic-era suburban flight provides a potential confound: if the distance gradient reflects a post-2020 structural break rather than a persistent feature of urban risk pricing, cross-sectional inference is compromised. Excluding all observations from 2020 Q1 onward, the primary-center gradient falls

Table 8: Fama-MacBeth Regressions: Distance Gradient with Census 2000 Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Distance (miles)	0.1469***	0.0425***	0.0264***	0.0371***	0.0265**	
	(0.0182)	(0.0067)	(0.0071)	(0.0103)	(0.0107)	
CBSA beta		1.2879***	1.4207***	1.2896***	1.3180***	1.3881***
		(0.2071)	(0.2144)	(0.2472)	(0.2655)	(0.2511)
Log median income			0.2079	0.3237	0.7288	0.9985**
			(0.3061)	(0.4521)	(0.5440)	(0.4265)
College share			−0.1841**	−0.1748	−0.3046**	−0.3755***
			(0.0798)	(0.1170)	(0.1256)	(0.0995)
Log pop. density			−0.1356***	−0.1413**	−0.0647	−0.0633
			(0.0485)	(0.0627)	(0.0669)	(0.0651)
Poverty rate					−0.0100	0.0450
					(0.3032)	(0.2219)
% Housing pre-1970					−0.0762	−0.2006*
					(0.1299)	(0.1152)
% NH White				0.3746	0.7585	0.1098
				(1.0133)	(1.0959)	(1.0557)
% Black				0.1878*	0.0903	0.1848
				(0.1052)	(0.1180)	(0.1142)
% Hispanic				0.6053	0.9654	0.2855
				(0.9632)	(1.0619)	(1.0036)
Avg. N per quarter	42	42	42	42	42	42
T	94	94	94	94	94	94

Notes: Fama-MacBeth cross-sectional regressions of ZIP-level quarterly excess total returns on distance to the primary employment center and controls. Specifications (1)–(5) are without intercept. $\bar{\lambda}$ is the time-series mean of quarterly slopes; SE $= \text{sd}(\{\lambda_t\})/\sqrt{T}$. Returns winsorized at 5%. Distance capped at 30 miles. Census 2000 controls from Decennial Census SF1 (race/ethnicity) and SF3 (income, education, poverty, housing) at ZCTA level. Population density uses TIGER/Line 2000 ZCTA land area. All controls standardized (zero mean, unit variance). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

from 3.26 to 2.47 bp/mile but remains highly significant ($t = 4.32$, $p < 0.001$), consistent across CBSA-plus-time FE specifications as well. The $\omega \times cay$ predictability result is *marginally stronger* pre-2020 ($t = 2.86$, $p = 0.009$ versus $p = 0.011$ in the full sample), indicating the pandemic period added noise rather than driving the result. The beta-distance restriction test loses statistical significance pre-2020 ($p = 0.198$) due to a power loss: the 40-quarter minimum requirement reduces the sample from 66 to 39 ZIP codes, not a sign reversal. Full pre-2020 tables are in Appendix **??**.

5.3 Time-Varying Risk Premia: Operating Leverage and the Aggregate Price of Risk

The tests in Section **??** treat risk prices as constant. This section develops and tests the model's prediction under time-varying risk prices, which provides a distinct empirical handle: not a level test but a *conditional* test of how cross-sectional return spreads co-move with the aggregate price of risk over the business cycle.

Theoretical motivation. In Section 4 (Appendix **??** for the derivation), when the price of risk $\lambda(z_t)$ varies with an aggregate state variable z_t, the risk premium on property i is $\pi_i(z_t) = \lambda(z_t)\,\sigma_x\,\varepsilon_p(x_i)$. Because the pricing elasticity $\varepsilon_p(x) = x/(x - \bar{\delta}_k)$ is increasing in the operating leverage ratio $\omega_i \equiv \bar{\delta}_k/x_i$, high-$\omega$ properties have disproportionately large elasticity. The cross-sectional spread in risk premia therefore *widens* when $\lambda(z_t)$ rises—that is, in bad times.

? show that the consumption-wealth ratio (*cay*) is a summary state variable for aggregate risk prices: when *cay* is high (depletion of household wealth relative to consumption, typical of recessions), investors demand higher compensation for bearing risk. Combining these two ingredients yields a sharp, testable prediction: $\omega_i \times cay_t$ should positively predict the return of property i at $t+1$, with the interaction coefficient positive. This is a *cross-sectional* test of time-varying risk exposure, not merely a time-series predictability exercise: the model predicts that the spatial cross-section of returns co-moves with aggregate conditions through the operating leverage channel.

Empirical specification. We estimate

$$r_{i,t+h} \;=\; \alpha_i \;+\; \gamma_1\left(\omega_i \times cay_t\right) \;+\; \gamma_2\,\omega_i \;+\; \gamma_3\,cay_t \;+\; \varepsilon_{i,t+h}, \tag{4}$$

where $r_{i,t+h}$ is the cumulative excess total return of ZIP i over h quarters, ω_i is the operating leverage proxy (operating expenses excluding property taxes, divided by gross rent, averaged over the ZIP's NCREIF observations), and cay_t is the quarterly consumption-wealth ratio from **?**.[10] ZIP fixed effects (α_i) absorb the ω_i level; date fixed effects absorb cay_t; only the interaction $\omega_i \times cay_t$ remains as the identifying source of variation. Standard errors are two-way clustered by CBSA and date. The sample covers $N = 4{,}372$ ZIP-quarter observations from 1991 Q1 through 2019 Q2, spanning 23 CBSAs.[11]

Multi-horizon results. Table **??** reports estimates of γ_1 at horizons $h = 1, 2, 4, 8$ quarters. At $h = 1$, $\hat{\gamma}_1 = 0.540$ ($t = 2.46$, $p = 0.022$). The coefficient grows monotonically with horizon: 1.162 at $h = 2$

[10] Property taxes are excluded from ω because taxes decline with distance from the CBD, creating a confound between the distance channel and the amplification channel. The all-in expense ratio (including taxes) is used for calibration in Section **??**.

[11] The $\omega \times cay$ panel differs from the Fama–MacBeth pricing sample in three respects: it begins before the 2000 Q1 cutoff applied to the returns panel, because *cay* data and NCREIF expense reports are available earlier; it covers 23 rather than 18 CBSAs, reflecting broader coverage in the expense data (gaining Baltimore, Kansas City, and Riverside; losing Austin and Raleigh, which lack the legacy CBD distance required for the operating leverage proxy); and the forward return lead sets the effective end date one quarter before the last available *cay* observation. Restricting the sample to 2000 Q1 onward yields qualitatively identical results ($h = 1$: $\hat{\gamma}_1 = 0.498$, $t = 2.14$, $p = 0.044$, $N = 4{,}148$; $h = 4$: $\hat{\gamma}_1 = 2.823$, $t = 3.14$, $p = 0.005$, $N = 3{,}604$).

($t = 2.88$), 2.903 at $h = 4$ ($t = 3.24$), and 5.488 at $h = 8$ ($t = 2.40$).[12] The horizon pattern is consistent with the mechanism: *cay* is a slow-moving variable, so risk price variation compounds over multiple quarters, generating increasingly large predictability at longer horizons.

The economic magnitude is substantial. At the observed 80th–10th percentile range of *cay* (0.010 in bad times versus -0.029 in good times), moving from the 10th to the 90th percentile of ω ($\omega_{10} = 1.574$, $\omega_{90} = 2.163$) is associated with an increase in expected quarterly returns of 31.3 basis points in bad times and a *decrease* of 92.1 basis points in good times—a total amplification effect of 123.4 basis points over the full cycle. This means that high-ω locations command substantially higher expected returns precisely when the aggregate price of risk is elevated, consistent with the model's operating leverage channel.

Table 9: $\omega \times cay$ Predictability: Multi-Horizon Results

Dependent Variables:	h=1	h=2	h=4	h=8
Model:	(1)	(2)	(3)	(4)
	Cumulative Excess Return $R_{t \to t+h}$			
Variables				
$\omega \times cay$	0.5398**	1.162***	2.903***	5.488**
	(0.2198)	(0.4037)	(0.8946)	(2.288)
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	0.44106	0.58391	0.67888	0.71167

Clustered (CBSA & date) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Dependent variable is cumulative excess total return $R_{i \to i+h}$. Specification follows equation (**??**). ZIP and date fixed effects included; standard errors two-way clustered by CBSA and date. ω_i = ZIP-level operating leverage proxy (operating expenses excluding property taxes, divided by gross rent); cay_t from **?**.

Sub-period stability. Table **??** splits the sample at 2010. At $h = 1$, the pre-2010 coefficient ($\hat{\gamma}_1 = 0.808$, $t = 1.52$) is larger in magnitude than the post-2010 coefficient ($\hat{\gamma}_1 = 0.418$, $t = 1.86$), though neither is significant at the 5% level individually in the short horizon split.[13] At $h = 4$, both sub-periods are

[12]The $h = 8$ standard error is larger because of the reduced effective sample size under multi-period overlapping returns; inference is based on the same two-way clustered standard errors.

[13]The 2008–09 recession concentrates the bulk of the variation in *cay*, increasing noise in the pre-2010 sub-sample relative to post-2010 which has more stable observations and a smaller effective *cay* variance.

significant: pre-2010 $\hat{\gamma}_1 = 4.922$ ($t = 3.41$), post-2010 $\hat{\gamma}_1 = 2.734$ ($t = 2.29$). The coefficient is larger pre-2010, suggesting the pandemic period (post-2020) added noise to the full-sample estimate. This is reinforced by the pre-2020 check in Section **??**, where p improves from 0.011 to 0.009.

Table 10: $\omega \times cay$ Predictability: Sub-Period Stability

Dependent Variables:	ret_h1		ret_h4	
	h=1 (Pre 2010)	h=1 (Post 2010)	h=4 (Pre 2010)	h=4 (Post 2010)
Model:	(1)	(2)	(3)	(4)
Variables				
$\omega \times cay$	0.8080	0.4180*	4.922***	2.734**
	(0.5305)	(0.2251)	(1.445)	(1.196)
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	0.54348	0.30048	0.76638	0.60898

Clustered (CBSA & date) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Pre-2010 = 2000 Q1–2009 Q4; Post-2010 = 2010 Q1–end of panel. Specification and fixed effects as in Table **??**. Standard errors two-way clustered by CBSA and date.

Placebo horse race. A critical concern is that $\omega \times cay$ may capture generic "distressed ZIP" exposure to aggregate conditions—any sufficiently peripheral or low-income location may cyclically underperform when cay is high. Table **??** reports a horse race of $\omega \times cay$ against three alternative interactions: $\text{Distance}_i \times cay_t$, $\ln(\text{Income}_i) \times cay_t$, and $\text{College Share}_i \times cay_t$.

The $\omega \times cay$ coefficient retains significance throughout (column 2, $t \approx 2.42$; full horse race column 5, $t \approx 1.92$). Distance $\times$ cay is consistently insignificant ($t \approx -0.85$ to -1.02 across specifications), providing no evidence that the geographic level of a ZIP is the source of state-dependent risk. College share $\times$ cay is also cleanly insignificant ($t < 0.5$). Income $\times$ cay shows negative marginal significance in some specifications (column 3: $t = -1.78$; column 5: $t = -2.17$), indicating that high-income ZIPs earn relatively *less* when cay is high (bad times). This is a different economic relationship—consistent with a flight-to-quality mechanism where high-income locations behave more like safe assets—and does not attenuate the operating leverage interpretation. Importantly, it carries the *opposite* sign from the $\omega \times cay$ result: operating leverage predicts higher returns in bad times (positive), while income predicts lower returns in bad times (negative). These are not collinear proxies for the same phenomenon.

The predictive power specific to the operating leverage proxy is thus distinct from broad spatial or demographic characteristics.

Table 11: $\omega \times cay$ Predictability: Placebo Horse Race

Dependent Variable:	ret_h1				
Model:	(1)	(2)	(3)	(4)	(5)
	Next-Quarter Return R_{t+1}				
Variables					
$\omega \times cay$	0.5398**	0.5042**	0.4718*	0.5405**	0.4216*
	(0.2198)	(0.2088)	(0.2284)	(0.2187)	(0.2201)
$Distance \times cay$		-0.0044			-0.0057
		(0.0052)			(0.0056)
$\ln(Income_z) \times cay$			-0.0742*		-0.0784**
			(0.0416)		(0.0361)
$CollegeShare_z \times cay$				-0.0273	-0.0080
				(0.0446)	(0.0584)
Fixed-effects					
Zip	Yes	Yes	Yes	Yes	Yes
date	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	4,372	4,372	4,372	4,372	4,372
R^2	0.44106	0.44127	0.44146	0.44113	0.44177

Clustered (CBSA & date) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Dependent variable is next-quarter excess total return. Each column adds one or more alternative *cay* interactions. Column (1) is the baseline $\omega \times cay$ result; columns (2)–(4) add individual placebo interactions; column (5) includes all three simultaneously. ZIP and date fixed effects included; standard errors two-way clustered by CBSA and date.

Caveats. Two limitations bound the strength of the inference. First, the cap-rate version of the left-hand side does not show the same pattern: $\omega \times cay$ does not predict cap rates in the distance-available sample ($t = 0.11$, $p = 0.916$). (The cap-rate specification uses $N = 2{,}767$; the smaller sample reflects the requirement of positive market value and non-missing expense data for computing income return levels at the ZIP level.) The returns channel is robust; the cap-rate channel, which requires stronger assumptions about discount rate variation being reflected in the level of income yields, is not significant on this sample. We do not suppress this null; it reflects the difficulty of measuring cap-rate levels precisely at the ZIP level in the NCREIF panel, which is also evident in the noisier distribution of income return observations relative to total returns. Second, ω_i is an imperfect proxy: it aggregates operating expenses from NCREIF reports to the ZIP level, conflating heterogeneous property vintages and management practices. Within-ZIP variation in true operating leverage is absorbed by the fixed effect; the cross-ZIP variation in ω_i may be measured with error, biasing the coefficient toward zero and making the estimates conservative.

5.4 Mechanism Evidence: NOI Passthrough

The operating leverage channel predicts a direct amplification of gross rent shocks into net operating income: for thin-margin properties (high ω, peripheral locations), a given percentage change in gross rent should generate a larger percentage change in NOI because fixed operating costs act as a lever. We test this at two levels of aggregation.

At the aggregate level, regressing $\Delta \log(\mathrm{NOI}_{it})$ on $\Delta \log(\mathrm{Rent}_{it})$ with ZIP and time fixed effects yields a coefficient of $\hat{\gamma} = 1.101$ ($t = 80.6$, $p < 0.001$, $N = 5{,}232$). (217 ZIPs across 25 CBSAs, from NCREIF expense reports; growth rates winsorized at the 1st and 99th percentiles.) A one-percent increase in gross rent is associated with a 1.10 percent increase in NOI—indicating that the amplification exists on average and is precisely estimated. When property taxes are excluded from operating costs (taxes decline with distance from the CBD, creating a spurious negative interaction between distance and cost variation), the passthrough coefficient is cleaner and more consistent with the theoretical channel.

At the within-CBSA level, however, the amplification does not increase with distance in the predicted direction. The panel interaction coefficient on $\Delta \log(\mathrm{Rent}) \times \mathrm{Distance}$ is -0.0035 ($t = -1.82$, $p = 0.082$) in our primary CBSA-plus-time FE specification—negative and marginally significant, opposite the model's prediction of higher passthrough for peripheral properties.[14] The passthrough panel draws on NCREIF expense reports, which cover a broader set of properties (217 ZIPs, 25 CBSAs) than the pricing sample (66 ZIPs, 18 CBSAs). The two-step approach (estimating ZIP-level passthrough coefficients and regressing on distance) yields an insignificant within-CBSA slope ($p = 0.728$).

We interpret this as a limitation of expense data at the ZIP level. NCREIF expense reports aggregate heterogeneous properties within a ZIP, and the operating leverage ratio ω_i exhibits limited within-CBSA variation (ZIP ω and distance are nearly orthogonal within metros, with correlation ≈ -0.09). The cross-sectional expense data are too coarse to identify differential amplification along the distance gradient. The $\omega \times cay$ time-series test in Section **??** provides stronger evidence for the

[14]Full results are in Appendix **??** and the primary outputs are at `empirics/new_analysis/results/noi_passthrough/`.

mechanism because it exploits *time variation* in aggregate risk prices—a source of variation orthogonal to the measurement noise that contaminates the cross-sectional expense data. The two tests are complementary: the aggregate passthrough documents NOI amplification; the $\omega \times cay$ interaction is consistent with high-ω properties being differentially exposed to aggregate risk price variation; the direct distance–passthrough interaction remains inconclusive.

6 Quantitative Assessment

Section **??** shows that the model's qualitative predictions align with the data: a distance gradient in total returns, a systematic income-capital tradeoff across cities, and a positive comovement between operating leverage and returns during high-risk-price periods. This section asks whether the operating leverage channel generates gradients of plausible *magnitude*, using a minimal quantitative mapping that pins two primitives directly from NCREIF data.

Targeted moments and parameterization. The model has two free objects that determine the within-city cap rate spread: the operating expense ratio $E/R \equiv \bar{\delta}_k/x$ at the innermost distance quartile, and the rent gradient $g \equiv x_{Q_4}/x_{Q_1}$ from center to periphery. Under the constant-δ assumption, $E/R_{Q_4} = E/R_{Q_1}/g$, so both objects are identified from NCREIF data without additional structural assumptions.

We set $E/R_{Q_1} = 0.45$ (all-in, including property taxes), consistent with the 43–47% range in NCREIF Expense Reports across the ZIP panel. The rent gradient is estimated as $g = x_{Q_4}/x_{Q_1} = 0.685$, using income per square foot averaged within distance quartiles with CBSA fixed effects; the within-CBSA estimate corrects for the composition bias that inflates innermost-ring income in high-cost metros (New York, San Francisco), where institutional properties concentrate near the CBD.[15] Under these inputs, $E/R_{Q_4} = 0.66$, the price-rent elasticity rises from $\varepsilon_{Q_1} = 1.82$ to $\varepsilon_{Q_4} = 2.91$ (elasticity spread $\Delta\varepsilon = 1.10$).

Results: cap rate spread. Table **??** presents the mapping. The model-implied cap rate spread between outermost and innermost distance quartiles is 42.8 basis points per quarter. The observed spread in the NCREIF panel is approximately 15 basis points (Table **??**, Panel B). The model overshoots by a factor of roughly 3.

This overshoot is informative, not disqualifying. The perpetual-claim benchmark provides an upper bound: it abstracts from the abandonment option that thin-margin peripheral properties possess. When operating margins are sufficiently thin, the option to exit tenancy attenuates the required

[15] The unconditional $g = 0.544$ is contaminated by this composition effect; the within-CBSA $g = 0.685$ is the appropriate input.

return, dampening the spread below the perpetual-claim prediction. We calibrate the abandonment option explicitly using the quadratic fundamental solution for the value function, with parameters $\nu = 0.01875$ per quarter, $\sigma_x = 0.013$ per quarter, and abandonment threshold $x_L = 0.421$. The attenuation is 0.07 basis points (0.2% of the spread), indicating that institutional-grade properties in the NCREIF panel sit far above their abandonment thresholds ($x_{Q_4}/x_L = 1.63$): the abandonment option is economically negligible at these parameterizations.

The residual gap therefore reflects measurement uncertainty in the expense ratio. Table **??** presents two scenarios. Scenario A uses $E/R = 0.45$ (all-in) and generates a spread of 42.8 bps. Scenario B uses the volatility-implied ratio $E/R = 0.635$ (from $1 - 1/2.74 \approx 0.635$, matching the observed $\sigma_{\mathrm{NOI}}/\sigma_r = 2.74$) and generates 282.7 bps—far outside the plausible range. The ex-tax expense ratio ($E/R \approx 0.30$, which falls in the success range of 5–30 bps per our calibration sensitivity) and the maintenance-only ratio of $E/R = 0.2752$ (the value that exactly matches 15 bps) provide lower bounds. The plausible operating range for the mechanism thus lies between the ex-tax and all-in expense ratios.[16]

Corroborating moments. Two moments not targeted in the mapping provide cross-checks. First, the model-implied total return spread between outermost and innermost quartiles is 54.8 bps, close to the observed $\approx$ 50 bps. This near-match was not used to set any parameter. Second, the model-implied NOI-to-rent volatility ratio is 2.26, compared with 2.74 in the data. The model understates the ratio by 18%, consistent with the expense-ratio tension noted above: the all-in $E/R = 0.45$ is lower than the volatility-implied 0.635, so the perpetual-claim approximation generates too little NOI volatility relative to rent. The OL share of the average city-level risk premium is 20% (wage-exposure premium $A_k = 1.28$ pp/qtr), consistent with the expectation that operating leverage is a quantitatively meaningful but not dominant component of within-city return heterogeneity.

Summary. The model delivers the correct qualitative pattern in all three checked moments and brackets the observed cap rate spread magnitude. The quantitative gap identifies the importance of the expense ratio definition: the all-in expense ratio generates a factor-of-three overshoot, while the ex-tax ratio brings the model close to the data. The volatility-implied scenario is quantitatively implausible. Pinning the expense ratio more precisely—ideally separating management fees and maintenance from taxes and insurance—is the main avenue for tightening the quantitative discipline. The gap between perpetual-claim upper bound and observed spread also identifies the potential contribution of the abandonment option channel, which we develop theoretically in Section 4 but do not estimate separately from the measurement uncertainty in the expense ratio.

[16] Variant C sensitivity analysis shows that the success range (5–30 bps) corresponds to $E/R \in [0.30, 0.35]$, and the affine versus proportional translation choice also affects the calibration outcome. These results are in Appendix **??**.

Table 12: Minimal Operating-Leverage Mapping: Model vs. Data

	Source	Data	Model-Implied
Panel A: Targeted inputs (pinned to data)			
Operating expense ratio (E/R, all-in)	NCREIF Expense Reports	0.43–0.47	(targeted: 0.45)
Rent drop Q4/Q1 (g, within-CBSA)	NCREIF ZIP-level	0.685	(targeted)
Panel B: Model-implied objects			
Price-rent elasticity, innermost quartile (ε_{Q_1})	Model	—	1.818
Price-rent elasticity, outermost quartile (ε_{Q_4})	Model	—	2.914
Elasticity spread ($\Delta\varepsilon \equiv \varepsilon_{Q_4} - \varepsilon_{Q_1}$)	Model	—	1.095
Base OL risk premium (B_k, bps per unit ε)	Cap rate spread / $\Delta\varepsilon$	—	13.69
Wage-exposure premium (A_k, pp/qtr)	Residual from FMB	—	1.280
Panel C: Checked moments (not targeted in estimation)			
Cap rate spread Q4−Q1 (bps/qtr)	Table **??**, Panel B	≈15	42.8
Total return spread Q4−Q1 (bps/qtr)	Table **??**, Panel A	≈50	54.8
Volatility ratio ($\sigma_{\text{NOI}}/\sigma_r$)	Appendix **??**	2.74	2.26
Internal consistency ($A_k > 0$?)	Decomposition check	—	✓ Pass
OL share of average premium	FMB decomposition	15–40% (expected)	20%

Notes: Variant A (OL-Only Baseline). Parameters: $E/R_{Q_1} = 0.45$ (all-in, from NCREIF Expense Reports); $g = x_{Q_4}/x_{Q_1} = 0.685$ (within-CBSA mean, income per sq. ft. as proxy for gross rent-to-wage ratio x). Model uses the perpetual-claim approximation $\varepsilon_p \approx 1/(1 - E/R)$ and the constant-δ gradient $E/R_q = E/R_{Q_1}/g_q$. Return mapping is proportional (anchored at Q_1). B_k is the base operating-leverage risk premium per unit of elasticity, calibrated to the observed cap rate spread of 15 bps/qtr. A_k is the residual wage-exposure premium, verified positive for internal consistency. The volatility ratio model-implied value of 2.26 falls short of the data value of 2.74, reflecting the expense-ratio tension discussed in Section **??**.

Table 13: Two-Scenario Quantitative Mapping: Directly Measured vs. Volatility-Implied Expense Ratio

	Source	Scenario A ($E/R = 0.45$)	Scenario B ($E/R = 0.635$)
Panel A: Inputs			
Expense ratio E/R_{Q_1}	[source]	0.45	0.635
Rent gradient $g = x_{Q_4}/x_{Q_1}$	NCREIF within-CBSA	0.685	0.685
Panel B: Model-implied elasticities			
ε_{Q_1}	Model	1.818	2.740
ε_{Q_4}	Model	2.913	13.649
$\Delta\varepsilon$	Model	1.095	10.910
Panel C: Checked moments			
Cap rate spread Q4−Q1 (bps/qtr)	Data: ≈15	42.8	282.7
Total return spread Q4−Q1 (bps/qtr)	Data: ≈50	54.8	362.4
Volatility ratio $\sigma_{\text{NOI}}/\sigma_r$	Data: 2.74	2.26	4.66
B_k (bps per unit ε)		13.7	1.4
A_k (pp/qtr)		1.280	1.505
OL share of average premium	Expected 15–40%	20.0%	5.9%
Gate outcome		PARTIAL$_SUCCESS_STRONG$	TOO$_STRONG$

Notes: Scenario A uses the all-in expense ratio (including taxes) measured directly from NCREIF. Scenario B calibrates the expense ratio to match the observed $\sigma_{\text{NOI}}/\sigma_r = 2.74$ ratio. The total return spread is reported for informational purposes but not formally targeted.

7 Extensions

7.1 Time-varying risk premia

The baseline model features constant risk prices, so risk premia manifest exclusively through cap rates. In practice, risk prices vary over the business cycle, rising in recessions and falling in expansions. This subsection extends the model to allow state-dependent risk pricing and derives the cross-sectional prediction that motivates the empirical test in Section **??**.

Let z_t denote a macroeconomic state variable that drives time variation in the aggregate price of risk. We model the wage-denominated SDF parameters as affine in z:

$$\lambda_Y(z) = \lambda_0 + \lambda_1 z, \qquad \nu_Y(z) = \nu_0 + \nu_1 z,$$

where $\lambda_1 > 0$ implies that the price of aggregate risk rises when z is high. The state z_t follows a mean-reverting Ornstein-Uhlenbeck process $dz_t = \kappa_z(\bar{z} - z_t)\,dt + \sigma_z\,dB_t$, driven by the same aggregate shock as wages. Appendix **??** derives the resulting two-state pricing PDE for the price-wage ratio $p(x,z)$.

The key prediction follows from the risk-premium formula. In the perpetual-claim benchmark (suppressing city-specific risk under the $\zeta^k = 0$ frictionless case), the expected excess return on a property with rent-to-wage ratio x_i is

$$\pi_i(z) \;=\; \lambda_Y(z)\,\sigma_x\,\varepsilon_p(x_i,z),$$

where $\varepsilon_p(x,z)$ is the pricing elasticity from Section **??**. Consider two properties: a downtown property (large x_H, low ε_p) and a suburban property (small x_L, high ε_p). The risk-premium spread between them is

$$\pi_L(z) - \pi_H(z) \;=\; \lambda_Y(z)\,\sigma_x\,\big[\varepsilon_p(x_L,z) - \varepsilon_p(x_H,z)\big].$$

Because ε_p is decreasing in x and $\lambda_Y(z)$ is increasing in z, this spread *widens* when z is high. In plain terms: when aggregate risk aversion is elevated, the return premium for holding distant, thin-margin properties rises disproportionately.

This prediction maps directly to the empirical specification. The operating leverage ratio $\omega_i \equiv x_i/(x_i - \delta)$ is a monotone transformation of ε_p in the benchmark case. The consumption–wealth ratio cay of **?** serves as the empirical proxy for the macro state z_t. The model therefore predicts a positive coefficient on the interaction $\omega_i \times cay_t$ in return-predictability regressions. Section **??** tests this prediction.

The **?** habit framework provides a complementary microfoundation: when the surplus consumption ratio is low (recessions), both risk aversion and required returns rise, while risk-free rates fall. In such a framework, the operating leverage channel generates flight-to-quality dynamics in real estate:

cap rate spreads between distant and central locations widen precisely during economic downturns. We do not test this prediction separately in the current paper, as it requires a measure of the surplus ratio distinct from *cay*.

7.2 Closed system of cities

The baseline model treats the outside option Ψ_t—the utility available to workers outside the metropolitan area—as an exogenous stochastic process. This subsection micro-founds Ψ_t as the equilibrium outcome of a closed system of K cities competing for a fixed national population.

Setup Workers draw i.i.d. Type I Extreme Value taste shocks with scale $1/\eta$, where $\eta > 0$ governs migration sensitivity. The resulting population share in city k is $n_{k,t} = \exp(\eta\, v_{k,t}) / \sum_j \exp(\eta\, v_{j,t})$, where $v_{k,t}$ is the indirect utility in city k. The outside option is the inclusive value $\ln \Psi_t \equiv \frac{1}{\eta} \ln\big(\sum_j \exp(\eta\, v_{j,t})\big)$. Each city has a housing supply elasticity $\phi_k > 0$, so that migration inflows raise local rents.

Key results Appendix **??** derives the equilibrium dynamics. Two results deserve emphasis.

Result 1: endogenous outside option. The stochastic process for Ψ_t is determined in equilibrium:

$$\sigma_{\Psi,t} = \sum_{k=1}^{K} \Omega_{k,t}\, \sigma_{w,k},$$

where $\Omega_{k,t}$ are population-weighted shares that depend on city size and supply elasticity. The volatility of Ψ_t equals the weighted average of city wage volatilities, disciplining the baseline model's σ_Ψ as a calibration target rather than a free parameter. Large cities dominate Ω_k and hence dominate Ψ_t, while supply-constrained cities receive lower weight because wage gains there are capitalized into rents rather than utility.

Result 2: negative spillovers and endogenous co-movement. A positive wage shock in city j raises Ψ_t, making city k relatively less attractive. Migration outflows then depress rents in k. This generates negative rent spillovers—and endogenous rent co-movement—even when the fundamental shocks dB_t^k and dB_t^j are independent. The effect is stronger between large cities, which carry larger weights Ω_j in Ψ_t.

Implications for risk and pricing The closed system creates a city-size gradient in risk exposure. Large cities with productive economies dominate Ψ_t; a positive shock to their wages raises the outside option for all cities, so their rents respond less (their relative advantage is partially dissipated through migration). Small or marginal cities, by contrast, are more exposed to Ψ_t movements driven by other cities' shocks. This generates a city-size risk premium beyond what the partial-equilibrium model delivers: small cities face higher effective risk exposure and should command higher expected returns, other things equal.

The closed system nests the baseline as the limiting case of perfectly inelastic housing supply ($\phi_k = 0$). We do not empirically test this extension in the current paper, but note that it provides a structural foundation for the observed cross-city return dispersion and disciplines the exogenous σ_Ψ parameter.

Future directions Several extensions would complement the analysis. Limited financial integration across cities (**?**) could explain residual variation in risk premia after controlling for beta. Stochastic amenity dynamics within cities would connect the operating leverage mechanism to gentrification patterns. Climate and disaster risk introduce spatial heterogeneity in tail exposures, which the operating leverage channel would amplify. Each of these extensions requires additional data and modeling apparatus beyond the present scope.

8 Conclusion

This paper documents spatial heterogeneity in residential real estate returns and proposes an economic mechanism to account for it. Using NCREIF total return data on multifamily properties, we document two sets of facts. Across cities, risk exposure to the national apartment index is priced in the cross-section of returns. Within metropolitan areas, excess returns increase with distance from the primary employment center at a rate of 3.26 basis points per mile per quarter, or 3.8–4.5 basis points after controlling for ZIP-level betas. This gradient is a within-city, within-quarter fact that survives Census 2000 demographic controls, retaining 62% of its baseline magnitude.

We propose that operating leverage—arising from fixed property operating costs that do not scale with rents—amplifies risk exposure for thin-margin locations, generating higher expected returns at greater distances. Two tests deliver restrictions consistent with this mechanism. The interaction of location-level operating leverage and the consumption–wealth ratio predicts next-quarter returns ($p = 0.022$, with an economic magnitude of 123 basis points over the cycle), while placebo interactions are insignificant. Appendix **??** reports directional evidence that within-metro betas increase with distance, though this test is power-limited ($N = 66$, $p = 0.062$).

The quantitative assessment maps model-implied moments to data. Under the baseline expense ratio, the perpetual-claim model generates a cap rate spread of 42.8 basis points per quarter between the outermost and innermost distance quartiles, compared with an observed spread of approximately 15 basis points. The abandonment option, calibrated to NCREIF parameters, attenuates this spread by less than 1%, indicating that institutional-grade properties operate far above abandonment thresholds. The residual gap is interpretable: the perpetual-claim benchmark is an upper bound, and the model brackets the data when the expense ratio varies across its plausible range.

Several limitations should be acknowledged. The NCREIF sample consists exclusively of institutionally managed multifamily properties with appraisal-based valuations; results may not generalize to single-family housing or owner-occupied markets. Appraisal smoothing likely attenuates measured volatility and covariance, which could bias beta estimates and the calibration. The within-city cross-section is small (up to 66 ZIP codes, mean 42 per quarter, across 18 CBSAs), limiting the power of risk-pricing tests. The $\omega \times cay$ test uses the **?** cay series through 2019Q3, restricting the sample window for the conditional analysis.

Natural extensions include a closed system of cities in which the outside option is endogenous to the cross-section of wages (Appendix **??** sketches this), richer operating cost data distinguishing fixed from variable expense components, and broader property types with transaction-based valuations. Taken together, our findings indicate that the rent gradients of monocentric city models—a classical object in urban economics—create systematic variation in risk exposure through operating leverage, connecting spatial equilibrium to asset pricing.

References

Alonso, W. (1964). *Location and land use.* Harvard University Press, Cambridge, Massachusetts.

Boustan, L. P. (2010). Was postwar suburbanization "White Flight"? Evidence from the black migration. *Quarterly Journal of Economics*, 125(1):417–443.

Bracke, P. (2014). House prices and rents: Microevidence from a matched data set in central london. *Real Estate Economics*, 43(2):403–431.

Brueckner, J. K. (1987). The Structure of Urban Equilibria: A Unified Treatment of the Muth-Mills Model. *Handbook of Regional and Urban Economics*, 2:821–845.

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.

Card, D. and Rothstein, J. (2007). Racial segregation and the black–white test score gap. *Journal of Public Economics*, 91(11–12):2158–2184.

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.

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

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

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

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

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

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

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

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*. Chicago, University of Chicago Press.

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.

Wachter, J. A. (2006). A consumption-based model of the term structure of interest rates. *Journal of Financial Economics*, 79(2):365–399.

Appendix

A Mathematical Derivations

A.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.

Indirect utility and spatial indifference With Cobb–Douglas, the optimal expenditure shares are $c_t = \alpha\, \kappa_{k,d}\, w_{k,t}$ 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

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

Taking logs, it yields

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

as shown in the main text.

Let $\beta \equiv \frac{1}{1-\alpha} > 1$ and note that $\kappa_{k,d}$ and a_k are time-invariant. Then, defining

$$\phi_{k,d} \equiv \zeta \, \kappa_{k,d}^{\beta} a_k^{\beta},$$

we have

$$\ln r_{k,d,t} = \ln \phi_{k,d} + \beta \ln w_{k,t} - \beta \ln \Psi_t.$$

Rent dynamics We apply Itô's lemma to the log-linear expression to derive the dynamics of rent. For the wage process, we have the SDE

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

Let $f(w) = \ln w$, so Itô's lemma gives

$$d \ln w = f'(w) \, dw + \tfrac{1}{2} f''(w) d\langle w, w \rangle.$$

The first term is

$$f'(w) \, dw = \mu_{w,k} \, dt + \sigma_{w,k} \, dB_t + \varsigma_{w,k} \, dB_t^k.$$

The quadratic variation, given the independence of the two BMs, is

$$d\langle w, w \rangle = w_{k,t}^2 (\sigma_{w,k}^2 + \varsigma_{w,k}^2) \, dt,$$

so the second-order term gives

$$\tfrac{1}{2} f''(w) d\langle w, w \rangle = -\tfrac{1}{2} (\sigma_{w,k}^2 + \varsigma_{w,k}^2) \, dt.$$

Putting it all together, we get

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

With a similar derivation for the outside option, we get

$$d \ln \Psi_t = \Big(\mu_\Psi - \tfrac{1}{2} \sigma_\Psi^2 \Big) dt + \sigma_\Psi \, dB_t.$$

Because of linearity, it holds that $d\ln r_{k,d,t} = \beta(d\ln w_{k,t}) - \beta(d\ln\Psi_t)$. Substituting we find

$$\begin{aligned} d\ln r_{k,d,t} &= \beta\big(d\ln w_{k,t} - d\ln\Psi_t\big) \\ &= \beta\Bigg[\Big(\mu_{w,k} - \tfrac{1}{2}(\sigma^2_{w,k} + \varsigma^2_{w,k})\Big)dt + \sigma_{w,k}\,dB_t + \varsigma_{w,k}\,dB^k_t \\ &\qquad - \Big(\mu_\Psi - \tfrac{1}{2}\sigma^2_\Psi\Big)dt - \sigma_\Psi\,dB_t\Bigg] \\ &= \Bigg(\beta\Big(\mu_{w,k} - \mu_\Psi - \tfrac{1}{2}(\sigma^2_{w,k} + \varsigma^2_{w,k} - \sigma^2_\Psi)\Big)\Bigg)dt \\ &\qquad + \beta(\sigma_{w,k} - \sigma_\Psi)dB_t + \beta\varsigma_{w,k}\,dB^k_t. \end{aligned}$$

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

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

and

$$d\langle r_{k,d,t}, r_{k,d,t}\rangle = \Big(\big(\beta(\sigma_{w,k} - \sigma_\Psi)\big)^2 + \big(\beta(\varsigma_{w,k})\big)^2\Big)\,dt.$$

Therefore, collecting terms, we get

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

with

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

and drift

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

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.

A.2 Price determination

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 $P_{k,d,t}$. Let's guess-and-verify 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. 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 our given SDF and a similar argument as we've applied to rent, the price of a perpetual claim that pays y_t will be

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

where ρ_y is the risk-adjusted discount rate given by

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

A.3 Price determination with running costs

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

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

Note that NOI is homogeneous of degree 1 in (r, w), so

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

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

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

Let's conjecture that the price $P_{k,d,t}$ is a Markov function and also homogeneous of degree 1 in (r,w), satisfying[17]

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

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

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

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 For the ratio process $x_{k,d,t} = \frac{r_{k,d,t}}{w_{k,t}}$, notice that

$$\begin{gathered}
\partial_r x = \frac{1}{w}, \quad \partial_w x = -\frac{x}{w} \\
\partial_{rr}^2 x = 0, \quad \partial_{ww}^2 x = 2\frac{x}{w^2}, \quad \partial_{w,r}^2 x = -\frac{1}{w^2}.
\end{gathered}$$

Then, applying Itô,

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

[17] As the SDF does not depend on the overall scale of (r,w), homogeneity of NOI translates to the scale invariance of the price function.

so

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

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

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

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

$$d\big(\Upsilon_t^k\,p(x_t)\big) = \Upsilon_t^k\,dp(x_t) + p(x_t)\,d\Upsilon_t^k + d\langle \Upsilon^k,\, p(x)\rangle.$$

The cross variation term is

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

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

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

Substituting the drift into the pricing equation, we get

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

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

$$\nu_{Y,k}\,p(x_t) = (x_t - \delta) + \left(\mu_{x,k} - \lambda_{Y,k}\sigma_{x,k} - \xi_{Y,k}\varsigma_{x,k}\right) x_t\,p'(x_t) + \left(\frac{\sigma_{x,k}^2 + \varsigma_{x,k}^2}{2}\right) x_t^2\,p''(x_t). \tag{A1}$$

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

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

Substituting into the pricing ODE in wage units (**??**) yields

$$\nu_{Y,k}\,p_k(x) \;=\; (x-\delta) \;+\; \mu_{x,k}^{\mathbb{Q}}\,x\,p_k'(x) \;+\; \frac{\sigma_{x,k}^2 + \varsigma_{x,k}^2}{2}\,x^2\,p_k''(x). \tag{A2}$$

Assume that $\nu_{Y,k}$, $\mu_{x,k}^{\mathbb{Q}}$ 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

$$p_H(x) = C_1\,x^{n_1} + C_2\,x^{n_2},$$

where the characteristic exponents $n_1 > 1$ and $n_2 < 0$ solve

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

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, $(x-\delta)$. In what follows we show that the particular solution assumes linear form, so $p_P(x) = ax + b$.

The benchmark case without options In the absence of any explicit options, we are back at the perpetual claim case. Economic restrictions eliminate the homogeneous part of the solution:

- The term $C_1 x^{n_1}$ represents a "rational bubble," where the asset price grows faster than its underlying cash flows. Standard no-arbitrage conditions require us to set $C_1 = 0$.

- The term $C_2 x^{n_2}$ 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. Let such particular solution have the linear form $p_k(x) = ax + b$, with $p_k' = a$ and $p_k'' = 0$. Substituting into (**??**) and solving for a and b, we get

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

so

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

A.4 Development and abandonment

Start from the wage denominated representation:

$$Y^k_t\, l(x_{k,d,t}) \;=\; \sup_{\tau \geq 0} \mathbb{E}_t\big[Y^k_{t+\tau}\,(p(x_{k,d,t+\tau}) - c_k)\big].$$

This is an optimal stopping problem. For the rest of the derivation, we will omit the k, d indices to simplify notation. When land is undeveloped, there is no payout. Upon development, the payoff is

$$G_t \;\equiv\; Y_t\big(p(x_t) - c\big),$$

so

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

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

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

By Itô's product rule,

$$d(Y_t \cdot l(x_t)) = Y_t \cdot dl(x_t) + l(x_t) \cdot dY_t + d\langle Y_t, l(x_t)\rangle.$$

For $dl(x_t)$, Itô's lemma gives

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

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

$$d\langle Y_t, l(x_t)\rangle = -Y_t\, x_t\, l'(x_t)\,(\lambda_Y \sigma_x + \xi_Y \varsigma_x) dt.$$

Putting everything together,

$$\begin{aligned} dV_t = Y_t\Big[-\nu_Y l(x_t) \,+ \big(\mu_x - \lambda_Y\sigma_x - \xi_Y\varsigma_x\big)\, x_t\, l'(x_t) + \frac{1}{2}(\sigma_x^2 + \varsigma_x^2)\, x_t^2\, l''(x_t)\Big] dt \\ + Y_t\Big(x_t\, l'(x_t)\sigma_x - l(x_t)\lambda_Y\Big) dB_t + Y_t\Big(x_t\, l'(x_t)\varsigma_x - l(x_t)\xi_Y\Big) dB_t^k, \end{aligned}$$

so

$$\mathbb{E}_t[dV_t] = Y_t\Big[-\nu_Y l(x_t) \,+ \big(\mu_x - \lambda_Y\sigma_x - \xi_Y\varsigma_x\big)\, x_t\, l'(x_t) + \frac{1}{2}(\sigma_x^2 + \varsigma_x^2)\, x_t^2\, l''(x_t)\Big] dt.$$

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

$$\nu_{Y,k}\, l(x_{k,d,t}) = \big(\mu_{x,k} - \lambda_{Y,k}\, \sigma_{x,k} - \xi_{Y,k}\, \varsigma_{x,k}\big)\, x_{k,d,t}\, l'(x_{k,d,t}) + \left(\tfrac{\sigma_{x,k}^2 + \varsigma_{x,k}^2}{2}\right) x_{k,d,t}^2\, l''(x_{k,d,t}).$$

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

$$l(x_{k,d,t}) = p(x_{k,d,t}) - c_k \qquad \text{(value matching)}.$$

At the endogenous exercise boundary $x = x_k^H$, smooth pasting imposes

$$l'(x_k^H) = p'(x_k^H).$$

Similarly, with abandonment, the lower boundary x_k^L satisfies

$$l(x_k^L) = p(x_k^L), \qquad l'(x_k^L) = p'(x_k^L).$$

which solves the abandonment problem together with the $p(\cdot)$ ODE for the continuation region as defined previously.

A.5 Closed system of cities

The baseline model treats the outside option Ψ_t as exogenous. This subsection micro-founds Ψ_t as the equilibrium outcome of a closed system of K cities competing for a fixed national population $\bar{N}$. We derive the stochastic process for Ψ_t and show how city-level shocks generate endogenous co-movement in rents.

Setup: logit sorting and housing supply Retain the baseline preferences $U = c^{\alpha} s^{1-\alpha} a_k$. The indirect (deterministic) utility of city k is

$$v_{k,t} = \ln w_{k,t} - (1-\alpha)\ln \bar{r}_{k,t} + \ln a_k,$$

where $\bar{r}_{k,t}$ denotes the city-level rent index (base rent at the CBD, so that $r_{k,d,t} = \bar{r}_{k,t}\, D(d)$ for a time-invariant distance discount D). Workers draw i.i.d. Type I Extreme Value taste shocks with scale $1/\eta$, where $\eta > 0$ is the migration elasticity. The resulting population share is

$$n_{k,t} \equiv \frac{N_{k,t}}{\bar{N}} = \frac{\exp(\eta\, v_{k,t})}{\sum_{j=1}^{K} \exp(\eta\, v_{j,t})},$$

and the *outside option* (inclusive value) is

$$\ln \Psi_t \equiv \frac{1}{\eta} \ln\left(\sum_{j=1}^{K} \exp(\eta\, v_{j,t}) \right),$$

so that $\ln n_{k,t} = \eta(v_{k,t} - \ln \Psi_t)$ and $\sum_k n_k = 1$.

Each city has an iso-elastic housing supply curve $\ln H_k = \phi_k \ln \bar{r}_{k,t} + \text{const}$, with supply elasticity $\phi_k > 0$. Market clearing implies an inverse supply relationship

$$d\ln \bar{r}_{k,t} = \gamma_k\, d\ln N_{k,t},$$

where $\gamma_k > 0$ is the total inverse elasticity.

Equilibrium dynamics Differentiate the utility, share, and supply equations and substitute to obtain

$$dv_{k,t} = \frac{1}{1+\Lambda_k}\, d\ln w_{k,t} + \frac{\Lambda_k}{1+\Lambda_k}\, d\ln \Psi_t,$$

where $\Lambda_k \equiv (1-\alpha)\gamma_k \eta > 0$ is the *feedback parameter*: it measures how much migration-induced congestion chokes off utility gains.

By the envelope theorem, $d\ln \Psi_t = \sum_j n_{j,t}\, dv_{j,t}$. Substituting and solving yields

$$d\ln \Psi_t = \sum_{k=1}^{K} \Omega_{k,t}\, d\ln w_{k,t},$$

with weights

$$\Omega_{k,t} = \frac{n_{k,t}/(1+\Lambda_k)}{\sum_j n_{j,t}/(1+\Lambda_j)}.$$

Result 1: the derived process for Ψ_t Given the wage process $dw_{k,t}/w_{k,t} = \mu_{w,k}\, dt + \sigma_{w,k}\, dB_t + \varsigma_{w,k}\, dB_t^k$, the endogenous outside option follows

$$d\ln \Psi_t = \mu_{\Psi,t}\, dt + \sigma_{\Psi,t}\, dB_t + \sum_{k=1}^{K} \varsigma_{\Psi,k,t}\, dB_t^k, \tag{A3}$$

where $\sigma_{\Psi,t} = \sum_k \Omega_{k,t}\,\sigma_{w,k}$ and $\varsigma_{\Psi,k,t} = \Omega_{k,t}\,\varsigma_{w,k}$.

Two features deserve emphasis. First, the volatility of Ψ_t equals the population-weighted average of city wage volatilities—disciplining σ_Ψ as a calibration target rather than a free parameter. Second, large cities (high n_k) dominate Ω_k and hence dominate Ψ_t, while supply-constrained cities (high γ_k, hence high Λ_k) receive lower weight because wage gains there are capitalized into rents rather than utility.

Result 2: relative-shock rent dynamics and endogenous co-movement From the equilibrium system, rent dynamics satisfy

$$d\ln\bar{r}_{k,t} = \frac{\Lambda_k}{(1-\alpha)(1+\Lambda_k)}\big(d\ln w_{k,t} - d\ln\Psi_t\big).$$

Rents respond to *relative* shocks: the city's own wage innovation minus the national weighted average. Substituting the Brownian shocks, the loading of city k's rent on city j's idiosyncratic shock is

$$\text{Loading on } dB_t^j\ (j \neq k) = -\frac{\Lambda_k}{(1-\alpha)(1+\Lambda_k)}\,\Omega_{j,t}\,\varsigma_{w,j}.$$

A positive wage shock in city j raises Ψ_t, making city k relatively less attractive; migration outflows then depress rents in k. This generates *negative spillovers* and endogenous rent co-movement even when fundamental shocks dB_t^k and dB_t^j are independent. The effect is stronger between large cities, which carry larger weights Ω_j in Ψ_t.

Connection to the baseline When city housing supply is perfectly inelastic ($\phi_k = 0$, so $\gamma_k \to \infty$ and $\Lambda_k \to \infty$), the coefficient $\frac{\Lambda_k}{1+\Lambda_k} \to 1$ and the rent equation reduces to the baseline formula $d\ln r = \frac{1}{1-\alpha}(d\ln w - d\ln\Psi)$. The closed system thus nests the baseline as the limiting case of fixed housing supply, while generalizing it to elastic supply with the dampening factor $\Lambda_k/(1+\Lambda_k) < 1$.

B Supplementary Empirical Analysis

B.1 NOI and Rent Growth Volatility

Table **??** reports regressions of NOI growth on rent growth at the CBSA level. The data come from NCREIF's research indices. Column (1) shows the baseline specification, while Column (2) includes time fixed effects. The coefficient of 1.44 in the time fixed effects specification indicates that a 1% increase in rent growth is associated with a 1.44% increase in NOI growth. This elasticity greater than one is consistent with operating leverage: when operating costs do not move proportionally with rents, NOI growth is mechanically amplified relative to rent growth. However, this coefficient conflates the pure mechanical effect with other sources of variation in operating expenses.

Table B1: 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 cluster	by time
FE: Time		X
Note: + p ¡ 0.1, * p ¡ 0.05, ** p ¡ 0.01, *** p ¡ 0.001		

To directly measure volatility amplification, we compute the standard deviation of NOI growth and rent growth for each CBSA over the sample period. Across cities, NOI growth volatility averages 7.3%, while rent growth volatility averages 2.6%, yielding a mean volatility ratio of 2.74. Figure **??** shows the distribution of this ratio across CBSAs. Most cities exhibit ratios between 2 and 3, with some variation reflecting differences in operating cost structures or other city-specific factors. While this pattern provides suggestive evidence consistent with operating leverage, it cannot be interpreted as a pure measure of the mechanism since operating expenses contain both fixed and variable components, and may respond to local economic conditions beyond rent changes.

Figure B1: Distribution of Volatility Ratio Across CBSAs

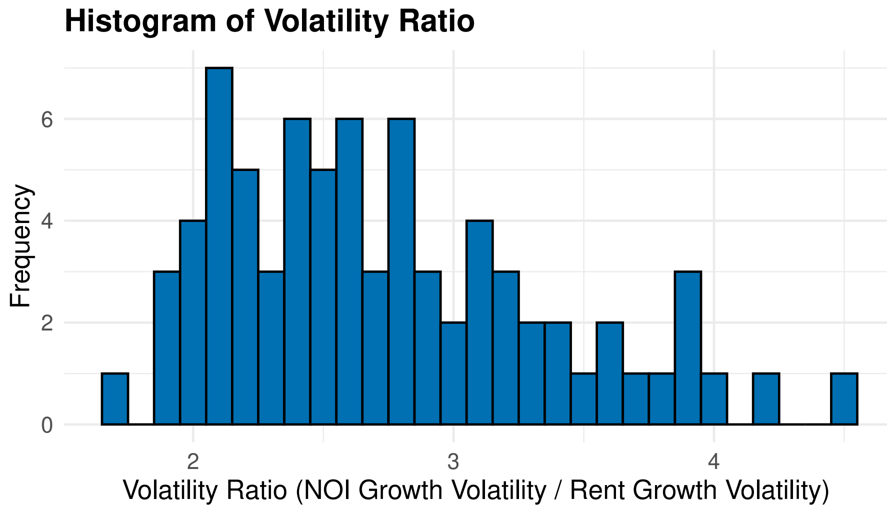

A cleaner test of the model's predictions would measure this volatility ratio by distance from the CBD within cities, as the model predicts the ratio should be higher for properties with thinner margins (farther from CBD). This analysis remains for future work using property-level data with sufficient within-city variation.

B.2 Comparing magnitudes: returns across and within cities

We present summary statistics for our within-city sample for returns aggregated at the CBSA level (**??**) and at the ZIP code level (**??**). The histograms for average returns over time for each CBSA (Panels a and b) and ZIP code (Panel c) are shown in **??**.

Table B2: 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 **??** due to the sample restriction to the same CBSA and ZIP-quarters used in the within-city analysis.

Table B3: 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.

Figure B2: Histogram of Average Returns

(a) At the CBSA level (across cities sample)

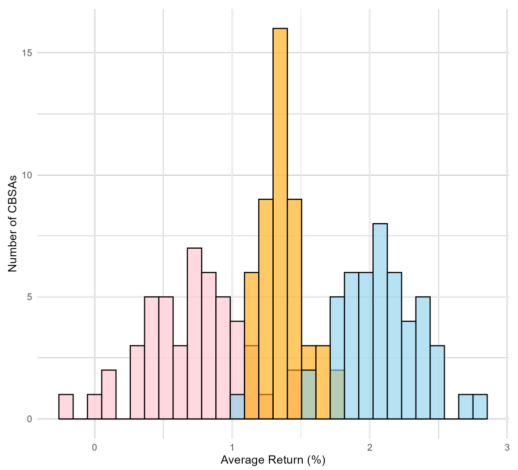

(b) At the CBSA level (within city sample)

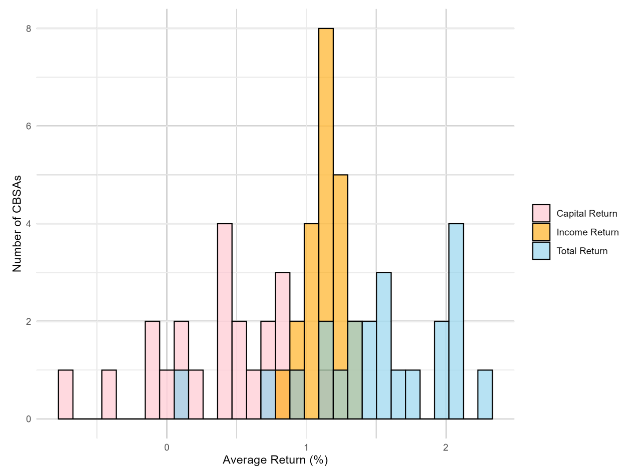

(c) At the ZIP code level (within city sample)

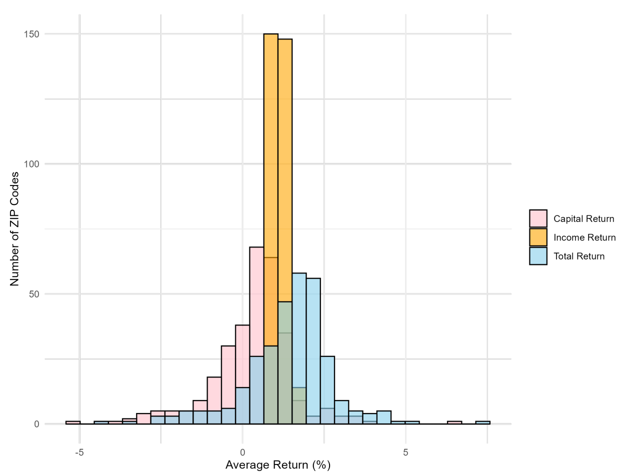

B.3 Beta measures and distance from CBD

Figure B3: Beta Measures and Distance

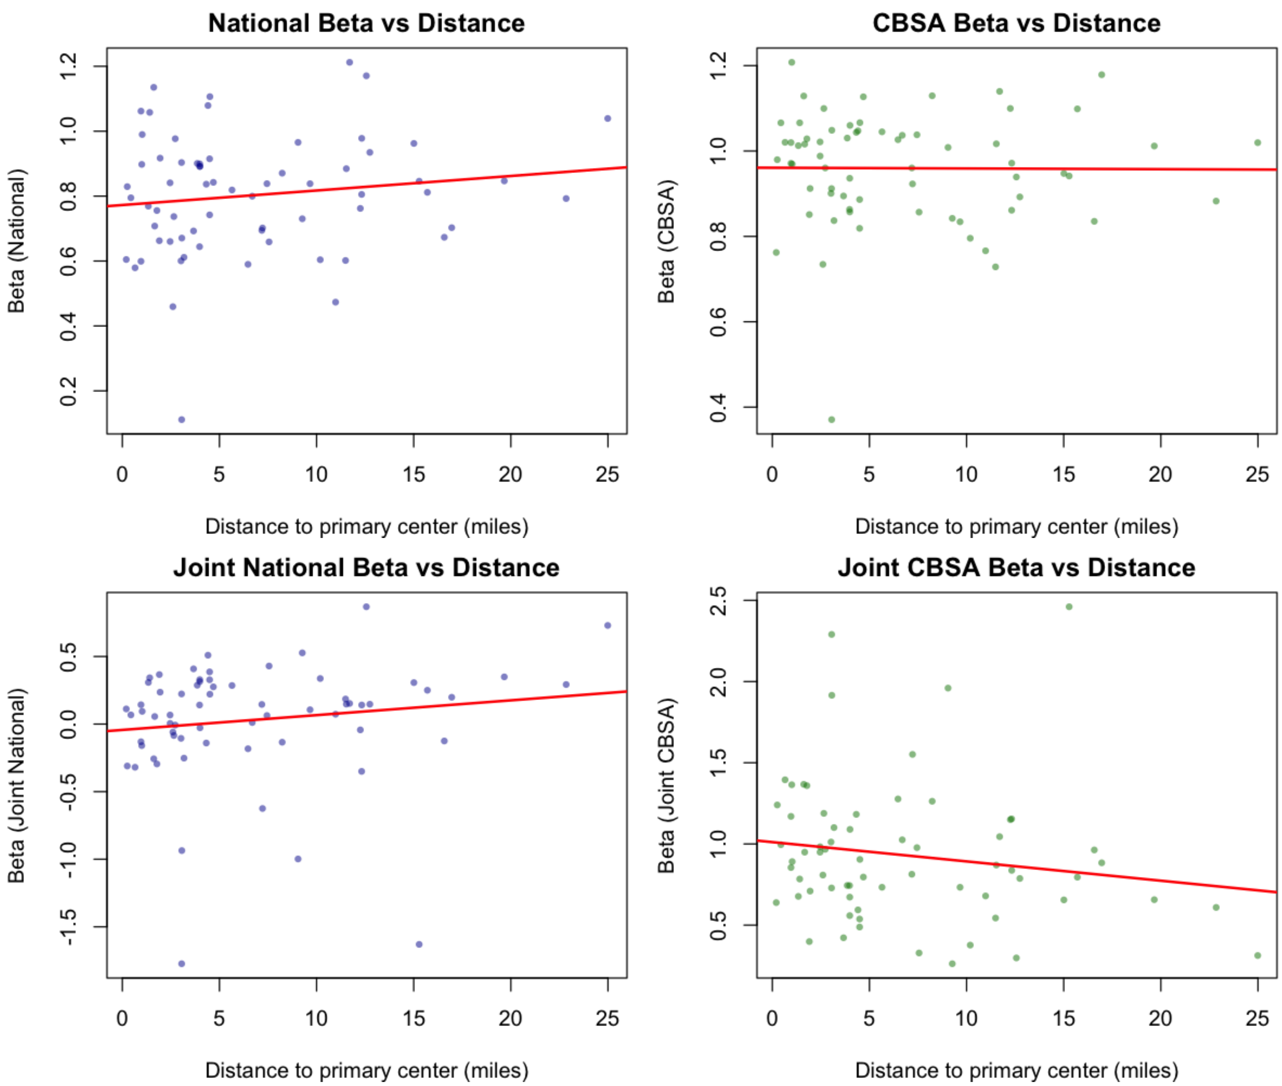

Figure B4: Beta Measures and Distance Groups

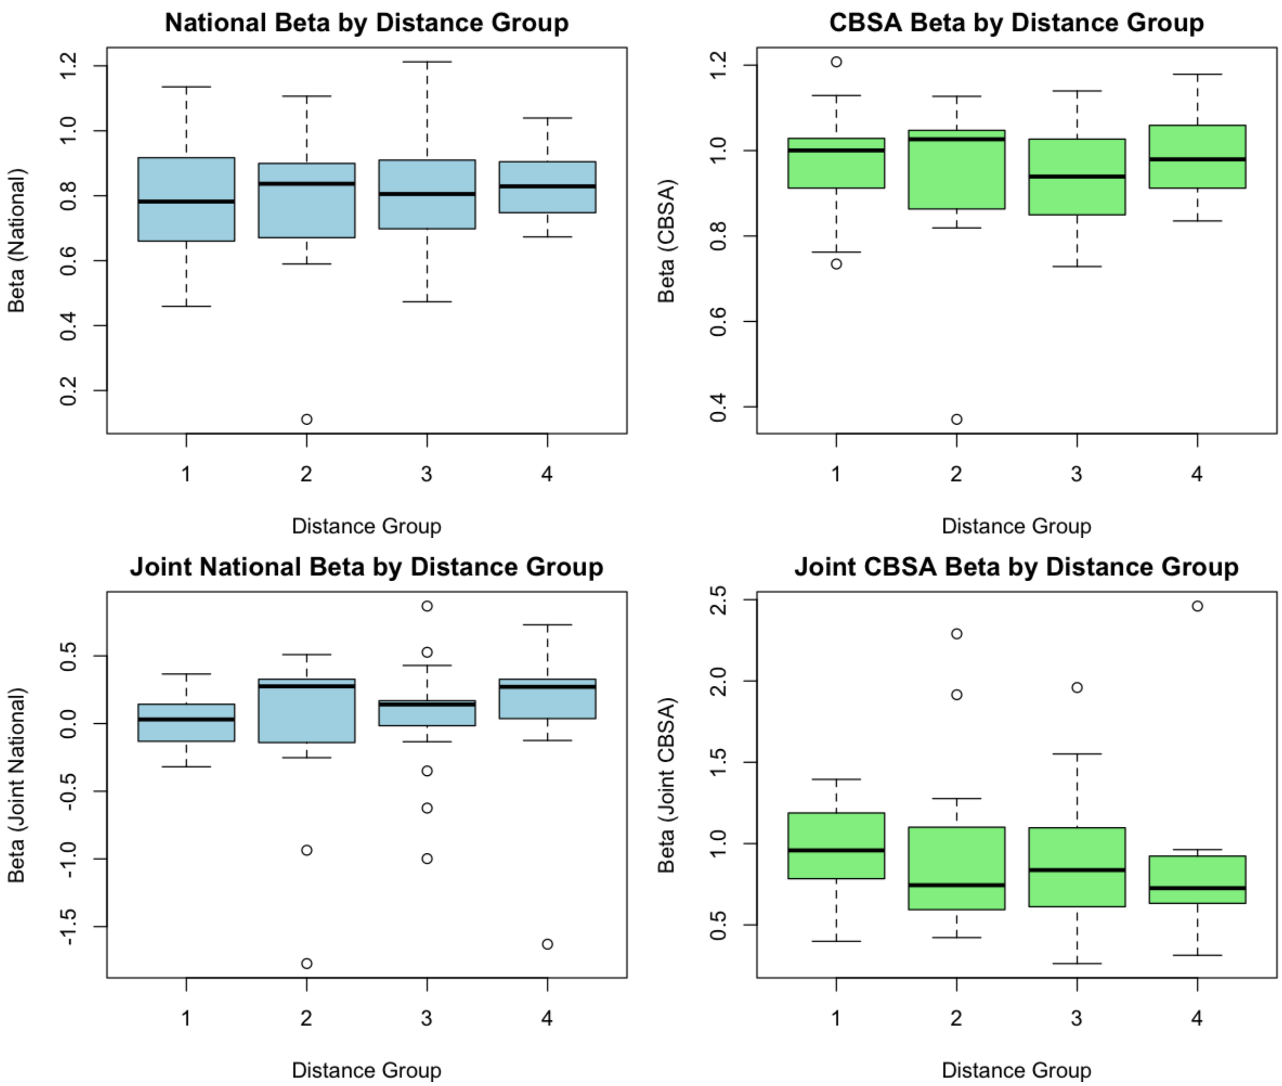

B.4 Beta-distance restriction diagnostics

This subsection provides diagnostics for the within-metro β–distance relationship. The main specification is a within-CBSA regression of ZIP-level CBSA betas on distance, controlling for cross-metro heterogeneity. Given the limited number of ZIP codes with sufficiently long return histories and meaningful within-CBSA distance variation, this test should be viewed as supportive rather than definitive. Unconditional distance sorts can be confounded by CBSA composition; the binscatter and forest plot below provide complementary visual evidence for the within-CBSA gradient.

Table B4: First-Order Operating Leverage Tests: Distance Gradient

			Primary Center			Old CBD	
Test	Specification	Pred.	Coef	t	N	Coef	t
$\beta_{\text{CBSA}} \sim$ dist	Unconditional	$+$	-0.0002	-0.06	66	0.0003	0.11
$\beta_{\text{CBSA}} \sim$ dist	Within-CBSA	$+$	0.0065*	2.00	66	0.0052	3.18
$\sigma(\text{NOI})/\sigma(\text{rent}) \sim$ dist	Unconditional	$+$	0.0022	0.18	105	-0.0006	-0.06
$\sigma(\text{NOI})/\sigma(\text{rent}) \sim$ dist	Within-CBSA	$+$	-0.0035	-0.22	105	-0.0016	-0.14
$\sigma(\text{total return}) \sim$ dist	Unconditional	$+$	0.0001	0.83	66	0.0001	1.30
$\sigma(\text{total return}) \sim$ dist	Within-CBSA	$+$	0.0001	1.15	66	0.0001	1.78

Notes: Unconditional = OLS; Within-CBSA = CBSA fixed effects. Standard errors clustered at CBSA level for within-CBSA specifications. distance_min results available in supplementary files.

B.5 Time-varying risk premia: theoretical framework

This subsection extends the pricing framework of Appendix **??** to allow the price of risk to vary with a macroeconomic state variable z_t (e.g., the consumption-wealth ratio *cay* of **?**). The key result is that the operating leverage mechanism interacts with time-varying risk premia to generate state-dependent return spreads across locations.

Setup Let z_t evolve as an Ornstein–Uhlenbeck process driven by the aggregate shock:

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

and let the wage-denominated SDF parameters be affine in z:

$$\nu_Y(z) = \nu_0 + \nu_1 z, \qquad \lambda_Y(z) = \lambda_0 + \lambda_1 z.$$

When z is high (economic downturns), investors demand higher compensation for bearing risk ($\lambda_1 > 0$). The state variable $x = r/w$ and its dynamics $dx/x = \mu_x\,dt + \sigma_x\,dB_t + \varsigma_x\,dB_t^k$ are as defined in Appendix **??**; city-specific subscripts are suppressed throughout.

Two-state pricing PDE **Result (Conditional pricing with time-varying risk).** The price-wage ratio $p(x, z)$ of a perpetual claim on net rent $(x - \delta)$ solves

$$\mathcal{L}\,p(x, z) + (x - \delta) = 0, \tag{B1}$$

where the generator is

$$\begin{aligned}\mathcal{L}\,p \equiv{}& -\nu_Y(z)\,p \\ &+ \big(\mu_x - \lambda_Y(z)\,\sigma_x\big)\,x\,\partial_x p + \big(\kappa_z(\bar{z} - z) - \lambda_Y(z)\,\sigma_z\big)\,\partial_z p \\ &+ \tfrac{1}{2}\big(\sigma_x^2 + \varsigma_x^2\big)\,x^2\,\partial_{xx}^2 p + \tfrac{1}{2}\sigma_z^2\,\partial_{zz}^2 p + \sigma_x\sigma_z\,x\,\partial_{xz}^2 p.\end{aligned}$$

Figure B5: Within-CBSA β–distance binscatter (Primary Center)

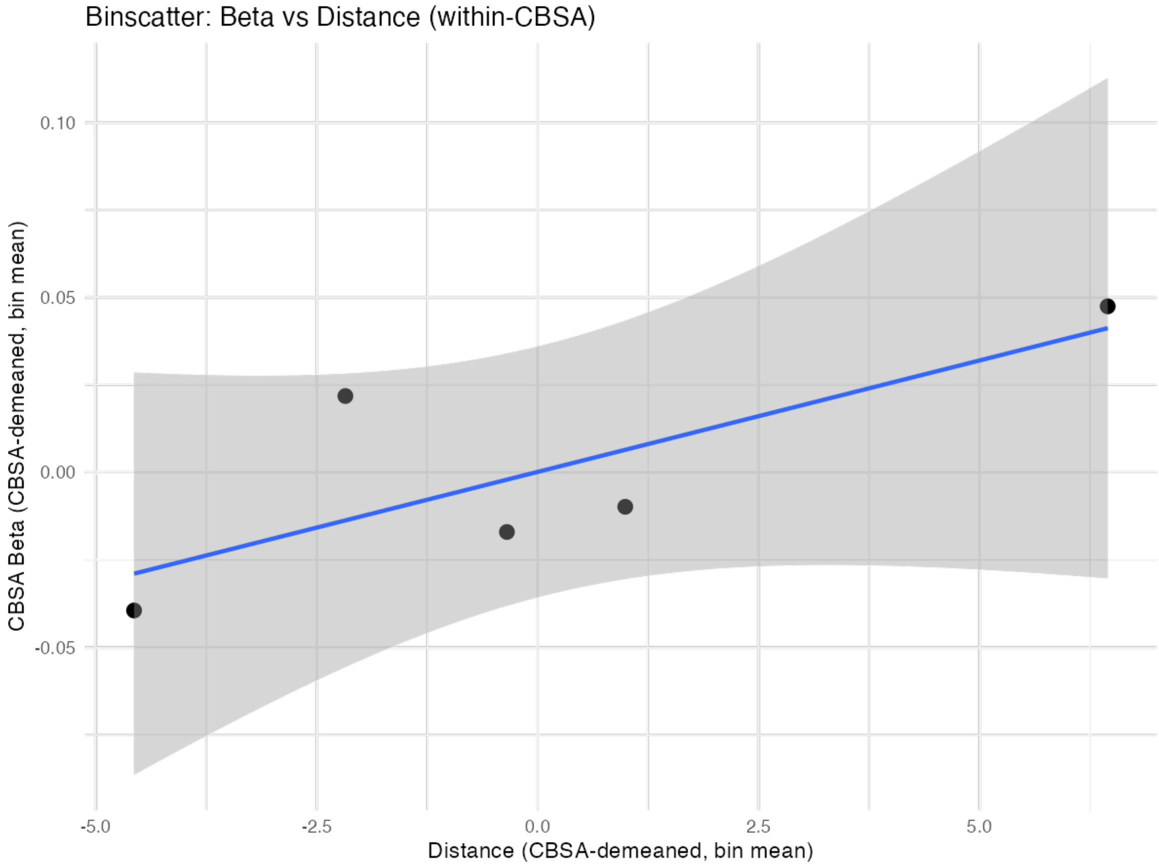

The derivation mirrors Appendix **??**: conjecture $P_t = w_t\, p(x_t, z_t)$, apply Itô's lemma to the product $Y_t\, p(x_t, z_t)$, and set the drift to zero. The additional terms relative to equation (**??**) arise from the second state variable z_t and its covariance with x_t through the common aggregate shock dB_t.

State-dependent elasticity Define the price elasticity with respect to x:

$$\varepsilon_p(x, z) \;\equiv\; \frac{x\, \partial_x p(x, z)}{p(x, z)}.$$

As in the constant-λ case, the fixed operating cost δ ensures that $\varepsilon_p > 1$ and is decreasing in x: properties with thinner operating margins (lower x, farther from the CBD) have higher elasticity. The new feature is that ε_p also depends on z. When z is high (risk premia elevated), the discount rate rises and the price function p shifts down, altering the elasticity profile.

The risk premium on property i with rent-wage ratio x_i takes the form

$$\pi_i(z) \;=\; \lambda_Y(z)\, \sigma_x\, \varepsilon_p(x_i, z) \;+\; \xi_Y\, \varsigma_x\, \varepsilon_p(x_i, z). \tag{B2}$$

Consider two locations: a high-margin downtown property (large x_H, low ε_p) and a thin-margin sub-

Figure B6: Within-CBSA β–distance slopes by CBSA (forest plot)

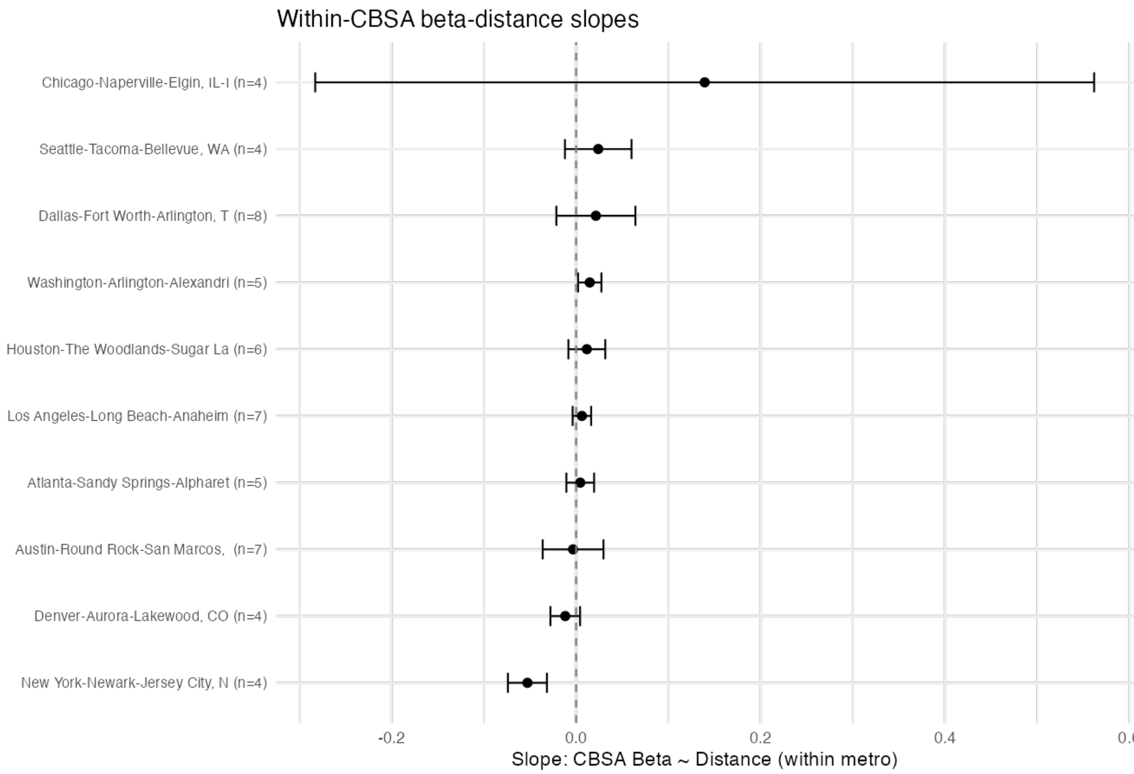

urban property (small x_L, high ε_p). The risk-premium spread is

$$\pi_L(z) - \pi_H(z) \;=\; \big[\lambda_Y(z)\,\sigma_x + \xi_Y\,\varsigma_x\big]\,\big[\varepsilon_p(x_L, z) - \varepsilon_p(x_H, z)\big].$$

Since $\varepsilon_p(x_L, z) > \varepsilon_p(x_H, z)$ and $\lambda_Y(z)$ is increasing in z, this spread *widens* when z is high. In plain terms: the return premium for holding distant, thin-margin properties rises during periods of elevated aggregate risk aversion.

Mapping to the empirical specification The operating leverage ratio $\omega_i \equiv x_i/(x_i - \delta)$ is a monotone transformation of ε_p in the benchmark (no-option) case. The empirical interaction $\omega_i \times cay_t$ therefore proxies for the state-dependent component of the risk premium $\pi_i(z)$. The model predicts a positive coefficient on $\omega \times cay$ in return-predictability regressions, which is confirmed in Section **??**.

B.6 Time-varying risk premia: $\omega \times cay$

We report supplemental results for the time-varying risk-premia specification using $\omega \times cay$. The interaction of location-level operating leverage with the aggregate price of risk significantly predicts next-quarter returns. Placebo interactions are insignificant, supporting the economic interpretation of the result.

Table B5: Operating leverage and time-varying risk premia: $\omega \times cay$

LHS	Estimate	t	p	N
CapRate	0.0018	0.11	0.9155	2767
Total Return (t+1)	0.5398	2.78	0.0109	4372
Income Return (t+1)	0.0484	3.27	0.0035	4372
Capital Return (t+1)	0.4913	2.62	0.0157	4372

Table B6: Placebo / falsification interactions (Total Return $t+1$)

Interaction	Estimate	t	p	N
omega_cay	0.5398	2.78	0.0109	4372
distance_primary_cay	-0.0068	-1.24	0.2287	4372
mv_zipavg_cay	9.67e-11	0.84	0.4105	4372
noi_zipavg_cay	4.16e-09	0.35	0.7303	4372
propcount_zipavg_cay	0.0060	0.31	0.7618	4372

C Within-City FM: With-Intercept Specification

This appendix reports the Fama-MacBeth within-city pricing regressions with an intercept included. As discussed in Section **??**, the distance gradient remains significant at approximately 3.5 basis points per mile per quarter in the with-intercept specification; estimated beta coefficients lose significance because the intercept absorbs the return level that betas explain.

Table C1: Fama-MacBeth Regression Results - Risk Premiums (with intercept)

	National	CBSA	Joint
No Distance			
Beta National	0.6392	–	-0.5656
	(0.3136)		(0.3912)
Beta CBSA	–	-0.6456	-0.7361
		(0.3591)	(0.3382)
Continuous Distance			
Beta National	0.5100	–	-0.3397
	(0.3123)		(0.3961)
Beta CBSA	–	-0.3910	-0.3775
		(0.3433)	(0.3334)
Distance	0.0349	0.0356	0.0356
	(0.0067)	(0.0067)	(0.0067)
Distance Groups			
Beta National	0.6012	–	-0.3022
	(0.3431)		(0.4630)
Beta CBSA	–	-0.4134	-0.4076
		(0.3669)	(0.3524)
Dist. Group 2	0.1457	0.1325	0.0698
	(0.1058)	(0.1007)	(0.1085)
Dist. Group 3	0.4982	0.4409	0.4376
	(0.1116)	(0.1169)	(0.1214)
Dist. Group 4	0.5254	0.5367	0.5012
	(0.1286)	(0.1270)	(0.1299)

Notes: Quarterly FM estimates. $\bar{\lambda}$ is the time-series mean of quarterly slopes; SE $=$ sd$(\{\lambda_t\})/\sqrt{T}$; p-values from t_{T-1}. Variables are winsorized at 5% (two-sided) prior to the quarterly cross-sections. Distance is measured in miles to the primary employment center (LEHD-LODES). T equals the number of non-missing quarters for that coefficient.

D Desmoothing Robustness

This appendix reports the full set of desmoothing robustness checks described in Section **??**. We apply a Geltner-style AR(1) unsmoothing correction, replacing each ZIP's observed return r_t^{obs} with $r_t = (r_t^{\text{obs}} - \hat{\rho}\, r_{t-1}^{\text{obs}})/(1-\hat{\rho})$ under seven parameterizations: ZIP-specific $\hat{\rho}$, pooled $\alpha = 0.482$, and fixed $\alpha \in \{0.4, 0.5, 0.6, 0.8\}$. Table **??** reports diagnostics. The median ZIP-level AR(1) coefficient in observed returns is 0.567, consistent with substantial appraisal smoothing. Both the pooled and ZIP-specific corrections collapse autocorrelation to near zero while approximately doubling return volatility.

Table **??** reproduces the three core results—the distance gradient, the $\omega \times cay$ interaction, and the β–distance restriction—across all seven specifications. The distance gradient is stable and, if anything, strengthened by unsmoothing: the pooled-α estimate rises to 3.62 bp/mile ($t = 3.71$), consistent with

Table D1: Desmoothing Diagnostics (Geltner-style AR(1) Correction)

	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 ($\hat{\rho} = 0$)	0.062		
Share at cap ($\hat{\rho} = 0.95$)	0.004		
$\hat{\rho}$ (median [IQR])	0.567 [0.358, 0.685]		
$\hat{\alpha}$ (median [IQR])	0.433 [0.315, 0.642]		
Pooled $\hat{\alpha}$	0.482		
Volatility and autocorrelation (median [IQR] across ZIPs)			
$\sigma(r)$	0.027 [0.023, 0.032]	0.041 [0.036, 0.048]	0.046 [0.036, 0.056]
AR(1)	0.567 [0.358, 0.685]	0.006 [−0.184, 0.213]	−0.033 [−0.105, 0.035]

Notes: Observed returns exhibit median first-order autocorrelation of 0.567, consistent with appraisal smoothing. Both the pooled and ZIP-specific corrections collapse autocorrelation to near zero while approximately doubling return volatility. The pooled correction applies a uniform $\hat{\alpha} = 0.482$; the ZIP-specific correction allows each ZIP's smoothing parameter to differ.

appraisal smoothing attenuating the baseline gradient. The $\omega \times cay$ coefficient retains its positive sign across all specifications; significance is preserved under ZIP-specific α ($p = 0.064$) and fixed $\alpha = 0.6$ and $\alpha = 0.8$ ($p \leq 0.053$), but weakens under the pooled correction ($p = 0.153$), reflecting sensitivity to the homogeneity assumption on smoothing rates. The β–distance restriction, already marginal in the baseline ($p = 0.062$), becomes statistically imprecise under all desmoothing variants. This is consistent with beta measurement error compounding the attenuation from appraisal smoothing; the test should be interpreted as a directional diagnostic rather than a standalone restriction.

Table D2: Core Results Under Desmoothing

	Baseline	Pooled α	ZIP α	$\alpha = 0.4$	$\alpha = 0.5$	$\alpha = 0.6$	$\alpha = 0.8$
Panel A: Distance gradient (FM, bp/mile)							
Coefficient	3.259	3.622	3.101	3.784	3.589	3.443	3.324
t-statistic	5.82	3.71	2.68	3.20	3.82	4.39	5.23
p-value	$<$0.001	$<$0.001	0.009	0.002	$<$0.001	$<$0.001	$<$0.001
Panel B: $\omega \times cay$ (predicts r_{t+1})							
Coefficient	0.540	0.402	0.569	0.347	0.411	0.454	0.508
t-statistic	2.78	1.48	1.95	1.03	1.57	2.04	2.61
p-value	0.011	0.153	0.064	0.314	0.130	0.053	0.016
Panel C: β–distance (within-CBSA)							
Coefficient	0.0065	0.0011	0.0032	0.0001	0.0014	0.0025	0.0049
t-statistic	2.00	0.29	0.59	0.03	0.36	0.79	1.59
p-value	0.062	0.774	0.564	0.980	0.723	0.438	0.131

Notes: The distance gradient uses 96 Fama–MacBeth quarters with a mean cross-section of 62.9 ZIPs. The $\omega \times cay$ specification uses $N = 4{,}372$ ZIP-quarter observations. The β–distance test uses 66 ZIPs across 18 CBSAs. For the ZIP-specific α correction, the CBSA benchmark return is unsmoothed using the pooled α.

E Pre-2020 Robustness

This appendix reports regression results for the pre-pandemic sample (excluding 2020 Q1 onward) described in Section **??**. Table **??** compares the distance gradient, the $\omega \times cay$ predictability result, the β–distance restriction, and placebo tests across the full and pre-2020 samples.

The Fama–MacBeth distance gradient falls from 3.26 to 2.47 bp/mile but remains highly significant ($t = 4.32$, $p < 0.001$), indicating that the pandemic era amplified the gradient but did not create it. The $\omega \times cay$ predictability result is marginally stronger pre-2020 ($p = 0.009$ versus 0.011), consistent with the pandemic period adding noise rather than driving the result. The β–distance restriction test loses statistical significance ($p = 0.198$) as the 40-quarter rolling beta requirement reduces the sample from 66 to 39 ZIP codes—a power loss, not a sign reversal. All placebo interactions remain cleanly insignificant in both samples.

Table E1: Core Results: Full Sample vs. Pre-2020

	Full sample			Pre-2020		
	Coef	t	p	Coef	t	p
Panel A: FM distance gradient (bp/mile, primary center)						
Continuous	3.26	5.82	<0.001	2.47	4.32	<0.001
Joint w/ β	3.70	6.07	<0.001	2.80	3.68	<0.001
CBSA + time FE	3.27	4.47	<0.001	1.58	2.61	0.011
Panel B: $\omega \times cay$ (predicts r_{t+1})						
Total return	0.540	2.78	0.011	0.602	2.86	0.009
Capital return	—	2.62	0.016	—	2.70	0.013
Panel C: β–distance (within-CBSA)						
Primary center	0.0065	2.00	0.062	0.0050	1.35	0.198
Panel D: Placebo interactions (total return)						
Distance $\times$ *cay*	—	-1.24	0.229	—	-1.21	0.238
Market value $\times$ *cay*	—	0.84	0.411	—	0.77	0.451
NOI $\times$ *cay*	—	0.35	0.730	—	0.28	0.781
Property count $\times$ *cay*	—	0.31	0.762	—	0.26	0.797

Notes: The full sample spans 2000–2024 (96 FM quarters); the pre-2020 sample excludes 2020 Q1 onward (80 FM quarters). Mean cross-section size falls from 62.9 to 51.4 ZIPs. The $\omega \times cay$ specification uses $N = 4,372$ (full) and $N = 4,187$ (pre-2020) ZIP-quarter observations. The β–distance test uses 66 ZIPs (full) and 39 ZIPs (pre-2020); the reduction reflects the 40-quarter minimum for rolling betas. Placebo coefficient magnitudes are omitted for brevity; all are economically negligible.

F NOI Passthrough Results

This appendix reports the NOI passthrough regression results described in Section **??**. Table **??** presents the aggregate passthrough coefficient and the within-CBSA distance interaction under alternative specifications.

The aggregate passthrough coefficient of 1.101 ($t = 80.6$) indicates that a one-percent increase in gross rent is associated with a 1.10-percent increase in NOI on average, consistent with operating leverage amplification. The within-CBSA interaction of rent growth with distance is negative for the primary-center and old-CBD measures—directionally consistent with the model's prediction that thin-margin peripheral properties exhibit greater amplification—but statistically weak ($t = -1.68$ for the primary center under CBSA$\times$time fixed effects). The expense-ratio placebo ($t = 0.25$) is insignificant, as expected. As discussed in the main text, the cross-sectional expense data are too coarse to identify differential amplification along the distance gradient with precision; the time-series $\omega \times cay$ test in Section **??** provides stronger evidence by exploiting variation orthogonal to this measurement noise.

Table F1: NOI Passthrough and Distance Interaction

LHS	Specification	Distance	Coef	t	N
Panel A: Aggregate passthrough					
$\Delta\log(\text{NOI})$	ZIP + time FE	—	1.101	80.57^{***}	5,232
Panel B: $\Delta\log(rent)\times$ distance interaction					
$\Delta\log(\text{NOI})$	CBSA×time FE	Primary	−0.0037	−1.68	5,232
$\Delta\log(\text{NOI})$	CBSA×time FE	Min	−0.0019	−0.77	5,232
$\Delta\log(\text{NOI})$	CBSA×time FE	Old CBD	−0.0035	-1.85^{*}	4,322
$\Delta\log(\text{NOI})$	CBSA + time FE	Primary	−0.0035	-1.82^{*}	5,232
$\Delta\log(\text{NOI})$	CBSA + time FE	Min	−0.0015	−0.73	5,232
$\Delta\log(\text{NOI})$	CBSA + time FE	Old CBD	−0.0040	-2.21^{**}	4,322
Panel C: Expense ratio placebo					
$\Delta(E/R)$	CBSA×time FE	Primary	0.0003	0.25	5,232

Notes: Panel A reports the aggregate NOI passthrough coefficient from regressing $\Delta\log(\text{NOI}_{it})$ on $\Delta\log(\text{Rent}_{it})$ with ZIP and time fixed effects. Panel B reports the coefficient on the interaction $\Delta\log(\text{Rent}_{it})\times\text{distance}_i$ under alternative specifications and distance measures. Panel C replaces the dependent variable with the change in the expense ratio as a placebo. Stars denote $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

G Calibration Sensitivity

This appendix reports the calibration sensitivity analysis described in Section **??**. Table **??** varies the all-in expense ratio E/R_{Q_1} across its plausible range while holding the within-CBSA rent gradient fixed at $g = 0.685$.

The model-implied cap rate spread is monotonically increasing in E/R and convex: small changes near the baseline ($E/R = 0.45$) produce moderate shifts, while higher values generate explosive spreads as the outermost-quartile margin approaches zero. The observed spread of approximately 15 bps is exactly matched at $E/R = 0.275$, close to the maintenance-only expense ratio (excluding property taxes). The 5–30 bps range that brackets the data corresponds to $E/R \in [0.30, 0.35]$, which falls between the ex-tax and pre-tax all-in ratios. The baseline all-in ratio of 0.45 generates 42.8 bps—a factor-of-three overshoot that serves as a plausible upper bound under the perpetual-claim approximation. The volatility-implied ratio ($E/R = 0.635$) is quantitatively implausible at 282.7 bps. Sensitivity to the rent gradient g is more moderate: varying g from 0.65 to 0.75 shifts the baseline spread by roughly ± 10 bps. The proportional translation specification (used in the main text) and the affine alternative yield qualitatively similar assessments; the affine specification produces slightly smaller spreads because it anchors income returns at the cross-quartile average rather than at the innermost quartile.

Table G1: Calibration Sensitivity to the Expense Ratio

E/R_{Q_1}	ε_{Q_1}	ε_{Q_4}	$\Delta\varepsilon$	Cap rate spread (bps/qtr)	Assessment
0.275	1.38	1.67	0.29	15.0	Exact match
0.30	1.43	1.78	0.35	17.4	Consistent with data
0.35	1.54	2.04	0.51	23.4	Consistent with data
0.40	1.67	2.40	0.74	31.4	Plausible upper bound
0.45	**1.82**	**2.91**	**1.10**	**42.8**	**Baseline (all-in)**
0.50	2.00	3.70	1.70	60.4	Implausible
0.55	2.22	5.07	2.85	90.9	Implausible
0.635	2.74	13.64	10.91	282.7	Implausible

Notes: Each row computes the model-implied cap rate spread between the outermost and innermost distance quartiles under the perpetual-claim benchmark, varying only the all-in expense ratio E/R_{Q_1}. The rent gradient is held fixed at $g = x_{Q_4}/x_{Q_1} = 0.685$ (within-CBSA estimate). The "exact match" row sets E/R to reproduce the observed spread of approximately 15 bps/quarter. The "consistent with data" assessment corresponds to spreads in the 5–30 bps range; the baseline all-in ratio of 0.45 generates 42.8 bps, a plausible upper bound given that the perpetual-claim framework abstracts from partial adjustment and finite horizons. The volatility-implied ratio of 0.635 generates 282.7 bps, which is quantitatively implausible.