The Commercial Real Estate Ecosystem

Ralph S.J. Koijen* Neel Shah† Stijn Van Nieuwerburgh‡

May 18, 2026

Abstract

We develop a new framework to study the joint dynamics of prices and trading volume in commercial real estate markets. We start from a micro-founded model where buyers and sellers differ in private valuations of building characteristics. We model listing and meeting probabilities based on investor and property traits. The model implies that transaction prices depend on both building features and participant identities—unlike traditional hedonic models. Using detailed transaction data and machine-learning methods, we find buyer and seller identities have a first-order effect on prices and transaction likelihood. We apply the framework to out-of-sample price prediction and counterfactual analyses on investment flows and prices. Our methodology extends to other private markets, including private equity, private credit, and infrastructure.

JEL: G11, G12, R30, R33
keywords: private and real assets, demand system asset pricing, commercial real estate

*Booth School of Business, University of Chicago, ralph.koijen@chicagobooth.edu.
†Columbia Business School, ns3481@columbia.edu.
‡Columbia Business School, svnieuwe@gsb.columbia.edu.

The authors would like to thank Dayou Xi, Luofeng Zhou, Zi Wang, Laurenz De Rosa, and especially Michael Leahy for excellent research assistance. We thank the Fama-Miller Center at Chicago Booth and the Paul Milstein Center (PMC) at Columbia Business School for financial support to acquire the MSCI Real Capital Analytics data. The PMC also provided the Compstak data. We thank Jeff Fisher for providing the NCREIF data. The authors thank Neroli Austin, Simon Buechler, Lu Liu, Sophie Calder-Wang, and the participants at the Inquire Europe conference in Valencia, the NBER Long Term Asset Management conference, the ABFER conference, the European Finance Association conference in Paris, the CREDA conference at UNC, The FSU Truist conference, and seminar participants at the University of Michigan, Columbia Business School, NYU Stern, UPenn Wharton, and Stanford GSB for helpful comments. Koijen thanks the Center for Research in Security Prices at the University of Chicago and the Fama Research Fund at the University of Chicago Booth School of Business.

1 Introduction

Financial markets have shifted markedly from public to private over the past 25 years. The number of U.S. publicly traded stocks fell from about 8,000 in 1997 to roughly 4,000 in 2024. Over the same period, assets under management in private equity broadly defined grew from a small base to $7.1 trillion in the U.S. and $13 trillion globally.[1] Many large institutions have reduced public equity and fixed-income allocations, increasing exposure to private-asset alternatives.[2] This "great rotation" has major implications for finance research and practice. Private and real assets differ from public ones: they trade infrequently, often in bilateral search-and-matching markets; prices are rarely observed; assets are lumpy and heterogeneous; location and other asset idiosyncrasies matter; there is an ecosystem of specialized investors. This paper develops a new methodology to value such assets.

One way to microfound standard hedonic models is to assume that investors have the same valuation for asset characteristics. Real-world markets, especially private and real assets, often violate these assumptions. We focus on commercial real estate (CRE)—a large, private asset segment with rich data on buyer and seller characteristics, spanning a diverse and evolving participant mix. When different investors value asset characteristics differently, it matters for the transaction price which investors are in the market at a given point in time.

Two CRE examples illustrate investor composition effects. First, a suburban mall loses anchor tenant Sears, hurting rents. Owned by a publicly-listed mall REIT with little capacity to repurpose it, the property may be worth more to an opportunistic REPE fund like Blackstone, with flexible capital and turnaround expertise. Second, large foreign institutional investors—sovereign wealth and pension funds—favor high-end "superstar city" CRE. The surge in such foreign capital after 2013 raises questions: How much did they push up prices? What assets did displaced domestic buyers acquire instead? How similar were those assets? The answers shed light on how foreign purchases spill over to other locations, CRE sectors, and asset features. Such questions are timely as cross-border investors became net sellers more recently.

We systematically document the CRE buyer–seller composition over 2001–2023 using MSCI Real Capital Analytics (RCA) data, which cover about 470,000 transactions worth $10

[1] McKinsey Global Private Markets Review 2024. Private asset markets include buyout, venture capital, growth, and other private equity, private debt, real estate, infrastructure, and natural resources. AUM as of June 2023; expected to double by 2029.

[2] For example, public pension funds raised their private and real asset share from below 10% in the late 1990s to over 25% by 2020 (Mittal, 2024). Alternatives include private equity, hedge funds, real estate, infrastructure, commodities, and miscellaneous alternatives. Using a broad definition including privately held businesses and collectibles, Goetzmann et al. (2021) estimate $84 trillion in private and real assets in 2020.

trillion across apartments, offices, retail, and industrial CRE. RCA's buyer–seller mapping reveals 325,000 unique investors holding $8.6 trillion in CRE by end-2023. We classify investors into eight types and construct their holdings from past transactions. We observe segmentation along multiple dimensions. For example, smaller investors focus on smaller, local assets; larger investors trade nationally in bigger deals.

We build a model reflecting decentralized private-asset trade consisting of valuation and transaction components. Each investor's valuation depends on asset characteristics and an idiosyncratic component (e.g., liquidity shocks). Valuations vary across investors for features like size, location, and quality. Allowing for investor-specific feature valuation greatly improves the fit of the model for transaction prices: adding investor traits to a rich Light Gradient Boosting Machine (LGBM) hedonic model raises R^2 by 10–20 points across sectors, with final in-sample R^2 of 84–88% versus 48–59% for a linear hedonic model. Our main model reduces the unexplained variation of the best non-linear model without investor characteristics between 42-55% when adding investor features.[3] We follow best practices in choosing hyper-parameters, mindful of overfitting concerns. Out-of-sample R^2 values reach 81–85%.

The transaction model conceptualizes bilateral search and matching between buyer and seller. We model the probabilities of listing, meeting, and transacting. The likelihood that a buyer meets a listed building is shaped by buyer size, similarity of buyer and seller size, prior relationships, and asset-type fit. Estimating the meeting model by maximum likelihood is computationally infeasible. To overcome this hurdle, we leverage the concept of negative sampling from the NLP literature (Mikolov et al., 2013), reducing computational burden while preserving the consistency of the estimator (Ma and Collins, 2018). Our estimated matching model with characteristics-based consideration sets is about 200 times more accurate in identifying the correct buyer than the standard random matching model.

Having estimated the model parameters, we turn to the main counterfactuals. A central object is the *potential price distribution* of a listed property. The seller could have met many other buyers besides the one who completed the deal. By sampling (with replacement) potential buyers in proportion to their meeting probabilities, and discarding cases with no gains from trade, we obtain the distribution of potential transaction prices. This distribution is valuable in its own right—for assessing the seller's luck or broker's skill, setting future price targets, and managing risk—and serves as a key input in four applications.

[3]These gains occur notwithstanding the possibility that a rich hedonic LGBM model without investor features may infer investor-related traits from property, market, and fixed-effect variables through *triangulation*, much like Fuster et al. (2022) show a random forest model can infer race from other borrower data in mortgage default prediction.

The first application addresses building owners planning to sell in the next period, who seek to predict the transaction price. Our valuation model implies that building, market, and investor characteristics all influence observed prices. Forecasting prices thus requires predicting these inputs. Here, the transaction model is crucial, drawing potential buyers in proportion to their likelihood of engaging and completing a transaction. Our benchmark model outperforms the linear hedonic model in out-of-sample tests, producing prices closer to actual transactions much more often and with significantly smaller pricing errors. The strength and stability of the out-of-sample evidence lends credibility to the counterfactual analyses that we explore.

The second application examines heterogeneity in valuations through an episode just before the Global Financial Crisis, involving several large transactions between REIT sellers and REPE fund buyers. Operating in a strong fundraising environment, REPE funds may have been compelled to deploy capital quickly. We ask what transaction probabilities and prices would have been if these REPE funds had valued the properties like all other investors. To measure the impact of a buyer group, we remove it from the potential matching pool and recompute the potential price distribution. We find a large impact of eager REPE buyers on both the transaction volume and average transaction prices.

The third application uses the model to analyze substitution patterns in deal flows across building size, sector, and location. Concretely, if many foreign investors purchase large, high-end Manhattan office buildings, what happens to smaller low-end offices in Manhattan, local apartments, or offices in other cities? The model traces transaction flows along these dimensions, with the counterfactual price distribution quantifying the impact of foreign investors on prices. We find that average transaction prices in the top-quartile of Manhattan office assets would have been 4.3% lower over the full sample and 5.5% lower over the post-2013 period in the absence of foreign buyers. These effects become larger for the top-10% of office properties, illustrating the large impact of foreign buyers on trophy office properties in gateway markets. Buyers displaced by foreign purchases predominantly substituted towards smaller Manhattan offices.

The last application studies the potential deployment of significant additional real estate private equity capital. REPE funds raised a lot of funds in 2021-22 but faced an unfavorable investment environment in 2023-24. They are now running out of time to deploy this “dry powder." We study what assumptions on REPE’s valuations and matching intensity are required to reach a target $150 billion investment in 2025, and what the implications are for the various CRE segments in which they are active. We find heterogeneous price increases by sector and market, highlighting heterogeneity in markets’ capacity to absorb this capital inflow.

Related Literature Our paper contributes to the literature on valuing private assets. Practitioner approaches often ignore risk, simply comparing total cash outflows from operations and disposition to the initial investment or adjusting for timing and the time value of money (IRR). The finance literature underscores the need to recognize the systematic risk in cash flows, measured in increasingly sophisticated ways. The public market equivalent (PME) approach Kaplan and Schoar (2005) benchmarks against a stock index with identical cash-flow timing. The generalized PME of Korteweg and Nagel (2016) allows for equity market exposures that differ from unity. The strip-by-strip approach of Gupta and Van Nieuwerburgh (2021) incorporates a broader set of public-market risk factors, including cross-sectional equity and interest rate risks, and allows exposures of operating and disposition cash flows to differ.[4] While Gupta et al. (2025) value CRE assets with richer cash-flow models, they also assume that the market prices of cash-flow risk can be inferred from public-market asset pricing data. We depart from this literature by not assuming that public markets set the prices of risk for private market exposures. Instead, our valuation model allows investor composition to influence prices. Specifically, who is in the deal affects the values of the coefficients on the hedonics.

We also contribute to the literature on asset prices using asset demand systems (Koijen and Yogo, 2019; Koijen et al., 2024), sharing the goal of jointly explaining valuations, asset characteristics, holdings, and trades. This paper contains the first demand system for private assets. Because such assets trade in bilateral markets, our approach differs sharply: instead of modeling portfolio shares of divisible public securities, we model the probability that a whole private asset changes hands in a buyer–seller meeting. We provide a micro-foundation for the valuation equation suited to this market structure that connects private valuations to heterogeneity in beliefs about first and second moments of cash flows, including the ability to transform the cash flows through asset repositioning, risk aversion, and pre-existing portfolios. Our framework extends naturally to other private or over-the-counter markets, including corporate and municipal bonds, private credit, private equity, infrastructure, natural resources, residential real estate, and collectibles.

The standard hedonic regression framework—valuing assets as a linear function of fixed-priced physical characteristics—has a long tradition (Lancaster, 1966; Griliches, 1971; Rosen, 1974; Witte et al., 1979; Wallace, 1996). Our model is richer: investor composition shapes how characteristics map to prices. Transaction prices reflect both buyer and seller valuations; changes in investor mix thus affect residual variance across time and markets. Time-varying coefficients in a standard hedonic model may inadvertently capture investor composition effects. For example, "number of floors" may appear to gain value in Man-

[4]Other important contributions to measuring risk and return in private equity include (Korteweg and Sørensen, 2010; Driessen et al., 2012; Ang et al., 2018).

hattan office data simply because foreign investors prefer tall buildings. Controlling for investor mix can "cleanse" hedonic price indexes of such effects.

We also enhance the set of characteristics to represent buildings, such as neighborhood quality based on granular rent data. We explore the importance of nonlinearities and interactions using modern machine learning (ML) techniques, specifically LGBM. See Gu et al. (2020); Francke and van de Minne (2024) for prominent examples of ML applied to stock and CRE markets, respectively. Even relative to the richer hedonic LGBM model, investor characteristics substantially improve explanatory power as they uncover crucial interaction effects between asset and investor characteristics, consistent with our theory of heterogeneous valuations for asset features.

The REIT literature uses standard empirical asset pricing methods to explain returns of publicly listed CRE firms (Van Nieuwerburgh, 2019), but REITs account for only 12% of transaction value in our data. As CRE trading shifts from public to private markets, understanding risk and return in private CRE becomes both harder and more important.[5] Peng (2016), Sagi (2021), and Giacoletti (2021) analyze risk and return at the property level, emphasizing search and matching frictions.[6] Plazzi et al. (2008, 2010) examine cross-sectional variation in expected returns across U.S. cities. Our potential price distribution complements these papers by providing a new liquidity risk measure at property and city levels.

Recent studies highlight the role of investor heterogeneity in CRE valuation. Ghent (2021) and Cvijanović et al. (2022) show that institutional CRE investors cluster in markets with similar institutions, raising valuations through added liquidity. Badrinza et al. (2022) stress same-nationality buyer–seller matches as a mechanism for resolving the gravity puzzle. Badarinza and Ramadorai (2018) use differences in foreign concentration across London neighborhoods to identify cross-border purchase effects on home prices. Damen et al. (2025) shows landlord types segment rental markets by quality tier.

Finally, we contribute a new estimation technique to the matching literature. We adapt negative sampling from AI and large language model training (Mikolov et al., 2013; Ma and Collins, 2018) to handle the computational burden of many assets potentially trading each period to many potential buyers. Similar ideas appear in the earlier discrete-choice literature in economics (McFadden, 1977; Fox, 2007). Fox (2018) proposes a maximum score estimator to address dimensionality in many-to-many matching. Gabaix et al. (2024) applies LLM-inspired techniques to define stock similarity using portfolio data. Badarinza

[5] Van Nieuwerburgh et al. (2015) discuss challenges in measuring returns for infrequently traded CRE and approaches to address them.

[6] A long literature captures frictional trade in real estate markets using search and matching models (e.g., Wheaton, 1990; Piazzesi and Schneider, 2009; Genesove and Han, 2012; Han and Strange, 2015; Piazzesi et al., 2013).

et al. (2024) estimates a housing-market matching function using unique search and broker request data.

The rest of the paper is organized as follows. Section 2 develops the theoretical framework, which contains a valuation and a matching model, bridging key methodological gaps in the current literature. Section 3 details the estimation procedure for the models. Section 4 describes the data and provides summary statistics. Section 5 shows results for the parameter estimates of the valuation model and Section 6 for the transaction model. Section 7 explores the four applications outlined above. Section 8 concludes. The online appendix contains additional details on the data and empirical methods.

2 Model

2.1 Notation

There are four building types, which we call sectors: office, retail, industrial, and apartments. We estimate separate models for each sector. Time is indexed by $t, t = 1, \ldots, T$. Buildings are indexed by n, $n = 1, \ldots, N$, and investors by $i, i = 1, \ldots, I$. If a building is sold, then we refer to the seller as investor s and the buyer as investor b. The transaction price is $P_t(n)$. Buildings are differentiated by a vector of asset characteristics, x_{nt}, to be spelled out in the empirical section. They include deal features and local macro-economic variables. Investor heterogeneity is captured by a vector of six characteristics, z_{it}. This vector includes information on the investor's type (e.g., REIT, foreign investor, etc.) and portfolio composition.

2.2 Investors' Private Valuations

Section 2.2.1 discusses the main empirical specification of an investor's private valuation while Section 2.2.2 provides a micro foundation for this specification.

2.2.1 A characteristics-based model of investors' private valuations

We assume that investor i's valuation of building n is given by

$$v_{it}(n) \equiv \ln V_{it}(n) = h\left(x_{nt}, z_{it}; \gamma_t\right) + \epsilon_{it}(n), \tag{1}$$

where γ_t is a vector of parameters that may vary over time. The error term, $\epsilon_{it}(n)$, captures liquidity shocks, funding constraints, or idiosyncratic preference shocks. We assume that the valuation shocks are uncorrelated across investors and buildings, conditional on all

characteristics, and normally distributed, $\epsilon_{it}(n) \sim N\left(0, \sigma^2_{m,t}\right)$, where $\sigma^2_{m,t}$ is allowed to vary by market m and time t.

In our empirical specification, we use a non-linear ML model for $h(\cdot)$ that flexibly captures interactions between investor and asset characteristics without imposing a specific functional form. Market fixed effects are included to capture local heterogeneity. Time fixed effects capture aggregate economic conditions.

If the investor's valuation depends linearly on asset characteristics $x_{n,t}$, as in the standard hedonic model, $h_{it}(n) = \beta'_{x,i} x_{n,t}$, but with "hedonic prices" $\beta_{x,i}$ that are investor-specific, and those hedonic prices are a linear function of investor characteristics, $\beta_{x,i} = \beta_x z_{i,t}$, then we obtain a private valuation given by

$$h_{it}(n) = \beta'_x \, (x_{n,t} \otimes z_{i,t}), \tag{2}$$

as a special case of our richer non-linear model. Equation (2) illustrates that when different investors have heterogeneous valuations for asset features, interactions of asset and investor features are central in valuation. Before proceeding, we provide a micro foundation of the valuation equation in (2).

2.2.2 A micro foundation of the private valuation model

We consider a two-period model, $t = 0, 1$. In period $t = 0$, the investor considers buying a building. In period $t = 1$, the investor receives the net cash flow and the resale value of the building. In mapping the model to the data, we think of $t = 1$ as the payoff of the building over a longer holding period.[7] The building in question may be part of a broader property portfolio, and in case of a pension fund or insurer that portfolio may include equities and fixed income assets. Without the new building, the broader portfolio generates a payoff D_{1i}. If the investor has no other assets, $D_{1i} = 0$.

At $t = 0$, the investor has access to cash C_{0i} that can be used to purchase the building. Any remaining cash earn a gross risk-free rate that is normalized to one. This implies that, without purchasing the new building, the investor's wealth in $t = 1$ is $A_{1i} = D_{1i} + C_{0i}$.

The new building generates a time $t = 1$ payoff of N_{1i}. A potentially important source of heterogeneity in CRE comes from differences in investors' ability to improve the cash flows of a building by managing it more efficiently or by renovating or repurposing it. We therefore endogenize the cash flows. Without any intervention, the cash flows are given by $\overline{N}_1$. The investor can improve the cash flows at cost $\frac{1}{2\phi_i} C_i^2$ to $N_{1i} = \overline{N}_1 + C_i$. The optimal investment in cash flow improvements is then $C_i^\star = \phi_i$, leading to investor-specific cash

[7] At the expense of some additional notation, this micro foundation can be extended to a multi-period model.

flows $N_{1i} = \overline{N}_1 + \frac{1}{2}\phi_i$, where the latter term is net of improvement costs.

We assume that $(D_{1i}, N_{1i}) \sim N(\mu_i, \Sigma_i)$. We allow investors to have heterogeneous beliefs about future payoffs and index all moments by i. For the new building, heterogeneity comes from differences in beliefs about $\overline{N}_1$ and from differences in investors' ability to improve cash flows, ϕ_i.

If the investor adds the building to her portfolio, period $t = 1$ wealth equals $A_{1i}^P = D_{1i} + C_{0i} - P_0 + N_{1i}$, where P_0 is the purchase price of the property.[8] To complete the model, we assume that the investor has mean-variance preferences over terminal wealth:

$$\mathbb{E}_i[A_{1i}] - \gamma_i \text{Var}_i(A_{1i}),$$

where γ_i is risk aversion.

As the building is a discrete purchase (as opposed to buying shares in a company), we compute the investor's valuation as the $t = 0$ price $P_0 = V_{0i}$ that makes the investor indifferent between purchasing the building or not. This valuation then solves the following equation:

$$\mathbb{E}_i[D_{1i} + C_{0i}] - \gamma_i \text{Var}_i(D_{1i}) = \mathbb{E}_i[D_{1i} + C_{0i} - V_{0i} + N_{1i}] - \gamma_i \text{Var}_i(D_{1i} + N_{1i}),$$

where the left-hand side is the investor's utility if she does not buy the building and the right-hand side is the utility when she does buy the building. We use this equation to solve for the investor's private valuation:

$$V_{0i} = \mathbb{E}_i[N_{1i}] - \gamma_i \text{Var}_i(N_{1i}) - 2\gamma_i \text{Cov}_i(D_{1i}, N_{1i}), \tag{3}$$

which depends on the expected payoff of the property, $\mathbb{E}_i[N_{1i}]$, a discount for its variance, $\text{Var}_i(N_{1i})$, and a further discount or premium that depends on the property's covariance with other assets in the investor's portfolio, $\text{Cov}_i(D_{1i}, N_{1i})$.

The main valuation equation is (3). To obtain a characteristics-based model of investors' private valuations, we follow Koijen and Yogo (2019) and model first and second moments as functions of characteristics with investor-specific coefficients,

$$\begin{aligned} \mathbb{E}_i[N_{1i}] &= \beta_{i0}' x_n, \\ \gamma_i \text{Var}_i(N_{1i}) &= \beta_{i1}' x_n, \\ \gamma_i \text{Cov}_i(D_{1i}, N_{1i}) &= \beta_{2i}' x_n, \end{aligned}$$

[8] We assume that the investor has deep pockets and $C_{0i} > P_0$. The model can be extended to allow the investor to finance a fraction of the property with a mortgage, or to cap the investor's valuation at C_{0i}.

Beliefs about first and second moments, risk aversion γ_i, ability to improve cash-flows ϕ_i, and time-1 payoffs D_{1i} from pre-existing portfolios are all allowed to be heterogeneous across investors. Hence the model accommodates five different sources of heterogeneity. We model this heterogeneity as a function of the investor characteristics (including investor type, size of the portfolio, etc.):

$$\beta_{ki} = \beta_k' z_i, \quad k = \{0, 1, 2\}.$$

This yields the valuation equation in (2).[9]

2.3 Transaction Prices

There is a proportional cost of $cP_t(n)$ if investor i and j transact. The bargaining power of the seller is measured by $\beta \in [0, 1]$, which is used to split the proportional transaction cost and in setting the price. Investor b is willing to buy the building when

$$V_{bt}(n) > P_t(n)(1 + \beta c),$$

or, equivalently,

$$v_{bt}(n) - \chi_b > p_t(n),$$

where $\chi_b \equiv \ln(1 + \beta c) \simeq \beta c$. Analogously, investor s is willing to sell the building when

$$V_{st}(n) < P_t(n) - (1 - \beta)cP_t(n),$$

or, equivalently,

$$v_{st}(n) - \chi_s < p_t(n),$$

where $\chi_s \equiv \ln(1 - (1 - \beta)c) \simeq -(1 - \beta)c$.

As is standard in search and matching models, we assume that the price is set by Nash bargaining,[10] implying

$$p_t(n) = \beta v_{bt}(n) + (1 - \beta)v_{st}(n) - c.$$

This implies that trade takes place when

$$v_{bt}(n) - v_{st}(n) > c.$$

In our empirical analysis, we set $c = 0$ and $\beta = 0.5$. Appendix C relaxes these assumptions,

[9]The difference between a valuation in levels versus logs arises because we assume mean-variance preferences. Alternatively, we can assume CRRA preferences and log-linearize the model.

[10]Formally, the price is the solution to $p^* = \arg\max \left((p - c) - v_s\right)^{\beta} \left(v_b - (p + c)\right)^{1-\beta}$.

derives the corresponding selection correction for $\beta \neq 0.5$ and $c > 0$, and shows that the valuation and meeting model estimations are robust to alternative bargaining weights $\beta \in \{0.4, 0.5, 0.6\}$ and realistic transaction costs of $c = 0.05$.

2.4 Transaction Probabilities

We now discuss the listing model, the meeting probability conditional on a listing, and the probability of transacting conditional on a meeting. At this stage, we assume that we know the parameters of the valuation model. We therefore have the component of valuations explained by characteristics, h_b and h_s, for the buyer and seller.

We model the probability of transaction as:

$$\pi(b, s) = \pi_l(s)\pi_m(b, s)\pi_\tau(b, s), \tag{4}$$

where the first term, π_l, is the probability of listing, π_m the probability of meeting the listed building, and π_τ the probability that the transaction goes through (conditional on a meeting). For ease of notation, we omit the time subscript on these probabilities.

We model listing at the property-year level. For a building s and year t, we specify the listing probability as

$$\pi_{l,t}(s) = \Lambda\Big(\eta_t + X'_{s,t}\theta\Big), \tag{5}$$

where $\Lambda(z)$ is the logistic cdf, η_t is a sector-year intercept, and θ is a sector-specific vector of slope coefficients. The vector $X_{s,t}$ contains information known at the beginning of year t: the seller-side predicted log total valuation of the property, defined as the predicted seller-side log unit valuation from the first-step valuation model plus log property size; an indicator for whether the owner sits on a paper gain; log market size; and seller-type and market fixed effects. The paper-gain indicator is motivated by models of reference point dependence (Genesove and Mayer, 2001; Andersen et al., 2022), under which owners may be more willing to list assets with unrealized gains than otherwise similar assets with unrealized losses. More broadly, by allowing listing probabilities to vary across properties and investors, the model accommodates the possibility that listing behavior reflects market conditions and changes in the composition of market participants. The intercept η_t is chosen so that, given the meeting and transaction probabilities, the model-implied number of transactions in a sector-year matches the number observed in the data exactly. Thus, η_t captures common time-varying listing conditions within a sector, while θ governs which properties are more likely to be listed than others within the same sector-year.

The second term in (4) is the meeting probability, the key object in our transaction

model. We assume that the likelihood of a buyer b meeting a building listed for sale depends on buyer and seller characteristics, past interactions, and the degree to which the property characteristics align with the buyer's pre-existing property portfolio:

$$\pi_m(b,s) = \begin{cases} 0 & \text{if } b = s, \\ \frac{\exp\left(\lambda_1 S_b + \lambda_2 N_b + \lambda_3 \Delta S_{b,s}^{-1} + \lambda_4 T_{b,s} + \lambda_5' \delta_{b,s}\right)}{\sum_{c \neq s} \exp\left(\lambda_1 S_c + \lambda_2 N_c + \lambda_3 \Delta S_{c,s}^{-1} + \lambda_4 T_{c,s} + \lambda_5' \delta_{c,s}\right)} & \text{otherwise.} \end{cases} \tag{6}$$

When $\lambda_1 > 0$, larger buyers (high S_b) are more likely to meet the asset, consistent with their higher trading activity and greater market presence.[11] When $\lambda_2 > 0$, buyers that own multiple assets ($N_b = 1$) have a higher likelihood of being matched. When $\lambda_3 > 0$, matching probabilities increase when buyers and sellers have similar size (high $\Delta S_{b,s}^{-1}$). Investors may prefer to deal with investors of similar size to reduce deal execution risk. When $\lambda_4 > 0$, buyers are more likely to match with sellers with whom they have a prior transactional relationship ($T_{b,s} = 1$), reflecting the role of established counterparties in reducing search and execution frictions.

The vector $\delta_{b,s}$ captures the buyer's consideration set. It measures the similarity between the characteristics of seller s's building and the portfolio characteristics of buyer b. Similarity has four dimensions. First, asset size alignment, the similarity between the focal building's log dollar value and the average log dollar value of the assets in the buyer's portfolio. Second, market alignment, the similarity between the building's geography and the share of the buyer's portfolio in that geography. Third, sector alignment, the similarity between the building's sector and the share of the buyer's portfolio in that sector. Fourth, quality alignment, proxied by the log NER of the building and the average log NER in the buyer's portfolio. Hence, λ_5 is a 4×1 vector.[12]

When $\lambda_i = 0$, $i = 1, \ldots, 5$, the matching probability is the inverse of the number of buyers and the model specializes to have random matching. When $\lambda_i \neq 0$, $i = 1, \ldots, 5$, (6) tilts the model away from random matching, towards search in ways that will be estimated from the transaction data that identifies properties that are more likely to be considered by buyers.

The third term in (4) measures whether a transaction takes place conditional on a meeting: $\pi_\tau(b,s) = P(V_s < V_b)$. We assume that this is the case when the valuation of the

[11] In reality, a seller may have bilateral meetings with multiple buyers, and choose among those making a bid for the asset based on multiple criteria ranging from offered price to certainty of execution. The buyer size as well as some of the other variables in (6) capture in reduced-form that larger buyers are more likely to emerge as winners from this process.

[12] For a given characteristic x_s of the building and x_b, the average characteristic of the buyer's portfolio, similarity is derived from the distance $|x_b - x_s|$ by converting smaller distances into higher similarities. Specifically, we use 1 - $|x_b - x_s|$ for geography and sector alignment and $\frac{1}{1+|x_b - x_s|}$ for size and NER alignment.

buyer exceeds that of the seller:

$$P\left(V_s < V_b\right) = P\left(\epsilon_s - \epsilon_b < h_b - h_s\right) = \Phi\left(\frac{h_b - h_s}{\sqrt{2}\sigma_{m,t}}\right).$$

under our assumption that $\epsilon_b, \epsilon_s \sim N(0, \sigma^2_{m,t})$.

The probability that a building does not trade is given by:

$$\pi_n(s) = (1 - \pi_l(s)) + \pi_l(s) \sum_b \pi_m(b,s)(1 - \pi_\tau(b,s)). \tag{7}$$

The first term, $1 - \pi_l(s)$, is the probability that the seller does not list the building. The second term is the probability that the building lists, and thus a match will be formed, but the valuation of the buyer does not exceed the valuation of the seller, as captured by $1 - \pi_\tau(b,s)$. As $\sum_b \pi_m(b,s) = 1$, $\sum_b \pi_m(b,s)(1 - \pi_\tau(b,s)) \leq 1$ and $\pi_n(s) \leq 1$. The probabilities that the building transacts with some buyer and that it does not transact sum to one:

$$\pi_l(s) \sum_b \pi_m(b,s)\pi_\tau(b,s) + \pi_n(s) = 1.$$

3 Estimation Procedure

The model is estimated sequentially, in two steps. In Section 3.1, we discuss the estimation of the valuation model. We discuss the estimation of the transaction model in Section 3.2.

3.1 Estimating the Valuation Model

We begin by estimating the valuation model, which is done separately for each property sector. We use observed transaction prices to estimate the parameters of h_b and h_s. The log price is given by

$$p_t(n) \quad = \quad \frac{1}{2}(h_{bt}(n) + h_{st}(n)) + \frac{1}{2}(\epsilon_{bt}(n) + \epsilon_{st}(n)). \tag{8}$$

We note that prices are observed only when $v_b > v_s$, which is when $h_b - h_s > \epsilon_b - \epsilon_s$. The difference in shocks is uninformative about the sum of the shocks because

$$\mathbb{E}\left[\epsilon_b + \epsilon_s \mid \epsilon_b - \epsilon_s\right] = 0,$$

when $\epsilon_b, \epsilon_s \sim N(0, \sigma^2_{m,t})$. This implies that estimating the valuation model based solely on the observed transactions does not create a bias.[13]

The investor valuations h_{bt} and h_{st} contain $\dim(x) + \dim(z) + 2$ predictors, where the last two variables refer to the market and year categorical variables. To flexibly capture the complex relationships between asset and investor characteristics, we use LightGBM (Ke et al., 2017), a gradient-boosted decision tree model. LightGBM is well-suited for this task as it efficiently handles high-dimensional data, categorical variables, captures non-linear interactions without requiring pre-specified terms, and incorporates built-in regularization to mitigate overfitting.

We use a two-step "coordinate descent" procedure with separate models for the buyer and the seller. In each iteration, we first hold the seller-side model fixed and update the buyer-side function by setting the training target to $2p_t(n) - h_s(x_n, z_s; \gamma_t)$. Next, we hold the buyer-side fixed and update the seller-side function using $2p_t(n) - h_b(x_n, z_b; \gamma_t)$ as the target. Alternating these steps ensures that, at convergence, the final predictions satisfy (8). The algorithm converges in a few iterations. We discuss the empirical implementation of the valuation model in Section 5. Appendix C.4 provides the detailed step-by-step implementation of the iterative estimation procedure.

3.2 Estimating the Transaction Model

Having estimated the valuation equation and obtained estimates of $h_{bt}(n)$ and $h_{st}(n)$, we turn to the estimation of the listing and meeting probabilities. In theory, we can use maximum likelihood. The contribution of the building of seller s to the likelihood is:

$$\ell_s = \sum_{b=1}^{B} y_{b,s} \ln \pi(b,s) + (1 - \sum_{b=1}^{B} y_{b,s}) \ln \pi_n(s),$$

where $y_{b,s}$ equals one when a transaction takes place and zero otherwise.

In practice, computing this likelihood is computationally infeasible due to the large number of buildings (and thus sellers) and buyers. We instead estimate the model using negative sampling, inspired by applications in the natural language processing literature (Mikolov et al., 2013). Intuitively, instead of using the full likelihood, we consider a classification problem. For each transaction, we sample another K potential buyers c that did not purchase the building, $c \in \mathcal{N}_s$. The cardinality of the negative set is $|\mathcal{N}_s| = K$. We then compute the probability that the actual buyer b purchases the building, conditional

[13] Appendix C extends the estimation procedure to allow for unequal bargaining weights and positive transaction costs. In that case, the conditional expectation is no longer zero, and we derive the corresponding selection correction.

on one of the $K+1$ potential buyers (including the actual buyer) purchasing the property:

$$\pi_r(b,s) = \frac{\xi_{b,s}}{\xi_{b,s} + \sum_{k \in \mathcal{N}_s} \xi_{k,s}}, \tag{9}$$

where $\xi_{b,s} = \exp\left(\lambda_1 S_b + \lambda_2 N_b + \lambda_3 \Delta S_{b,s}^{-1} + \lambda_4 T_{b,s} + \lambda_5' \delta_{b,s}\right) \pi_\tau(b,s)$.[14] We then minimize the loss function over the observed transactions, i.e., (b,s) pairs: $-\sum_s \ln \pi_r(b,s)$. This estimator is the ranking estimator in Ma and Collins (2018), who also develop the asymptotic theory. They show that the estimator is consistent for $K > 1$, asymptotically normal, and converges to the maximum-likelihood estimator as $K \to \infty$. Similar ideas appear in the discrete-choice literature.[15]

Lastly, we estimate the listing probability model separately for each sector. For property n in year t, let

$$q_{nt} = \sum_b \pi_{mt}(b,n)\pi_{\tau t}(b,n)$$

denote the model-implied probability that the property transacts conditional on being listed, where $\pi_{mt}(b,n)$ and $\pi_{\tau t}(b,n)$ are obtained from the previously estimated meeting and transaction blocks. We combine it with the property-level listing specification in equation (5). Thus, conditional on observables, the probability that property n transacts in year t is

$$P(T_{nt} = 1 \mid X_{nt}, q_{nt}) = q_{nt}\pi_{l,t}(n) = q_{nt}\Lambda(\eta_t + X_{nt}'\theta)\,.$$

where X_{nt} includes the seller-side predicted log total valuation of the property, an indicator for whether the seller sits on a paper gain, log market size, seller type and market fixed effects. The slope coefficients θ are estimated pooling all property-years within a sector using a penalized Bernoulli criterion, with a Ridge penalty to stabilize the estimates; the penalty parameter is selected by cross-validation. The intercept η_t is allowed to vary by year. For any candidate value of θ, we choose η_t so that the model-implied number of transactions in year t matches the number observed in the data exactly:

$$\sum_{n \in t} q_{nt}\Lambda(\eta_t + X_{nt}'\theta) = T_t,$$

where $T_t = \sum_{i \in t} T_{it}$ is the number of observed transactions in that sector-year.

In summary, transaction prices, observed characteristics of the buildings and the buyers

[14] The common denominator, $\sum_c \exp\left(\lambda_1 S_c + \lambda_2 N_c + \lambda_3 \Delta S_{c,s}^{-1} + \lambda_4 T_{c,s} + \lambda_5' \delta_{c,s}\right) \pi_\tau(c,s)$, which creates the computational bottleneck, drops out, radically simplifying the objective function.

[15] Specifically, section 7 of the working paper by McFadden (1977) discusses limiting the number of alternative choices in the context of a multinomial logit model. He shows that consistent estimates can be obtained from random samples of alternatives (Theorem 2). Fox (2007) considers a semi-parametric random utility model that produces a similar result.

and the sellers identify the coefficients of the valuation model in h_b and h_s. The coefficients of the listing and meeting model are identified based on which transaction do and do not take place. Transaction prices therefore inform the coefficients of the valuation model whereas transaction volumes identify the coefficients of the meeting and listing model.

3.3 Potential Price Distribution and Counterfactuals

For several applications discussed below, we are interested in computing the distribution of prices at which a trade could have occurred for an asset that is listed for sale. To do this systematically, we introduce a new tool, the potential price distribution, the distribution of prices that could have arisen if the asset had met different buyers and if both sides' idiosyncratic shocks had realized differently. Our algorithm delivers two objects: the trade probability and the distribution of prices conditional on trade. We summarize the latter by its mean and interquartile range, a measure of price risk (liquidity).

Once the parameters of the valuation and matching models have been estimated, we implement the following algorithm to get the potential price distribution. For each asset n at time t, we first draw a listing indicator $l_{n,t} \sim \text{Bernoulli}(\pi_{l,t}(s))$ using the sector-year listing probability $\pi_{l,t}(s)$ defined in (5). We simulate only the listed assets with $l_{n,t} = 1$. For each listed asset, we know the asset features $x_{n,t}$ and the systematic component of the sellers' valuation $h_{s,t}$. We draw a seller shock $\epsilon_s \sim N(0, \sigma^2_{m,t})$, resulting in a seller valuation $v_{s,t} = h_{s,t} + \widehat{\epsilon}_{s,t}$. Next, we compute the probability $\pi_m(c, s)$, given by (6), that this seller meets a potential buyer c from the pool of all investors. We randomly draw one potential buyer from the distribution $\{\pi_m(c, s)\}_{c=1}^B$. For that potential buyer, we draw a latent demand shock $\epsilon_{c,t} \sim N(0, \sigma^2_{m,t})$ and construct the valuation $v_{c,t} = h_{c,t} + \epsilon_{c,t}$. Using the trading rule from Section 3.1, the property transacts if $v_b > v_s$, in which case, the price is $\hat{p}_t(n) = \frac{1}{2}(v_{c,t} + v_{s,t})$. If the valuation of the buyer is too low, $v_{c,t} < v_{s,t}$, we set the price to missing. We can now compute the fraction of assets for which a trade takes place, $\tau = \frac{1}{N}\sum_n 1_{\hat{p}(n)>0}$, and the average price conditional on a trade $Ep = \frac{1}{\sum \mathbb{1}_{\hat{p}(n)>0}} \sum_n \hat{p}(n) 1_{\hat{p}(n)>0}$. We repeat this procedure $C = 1,000$ times, drawing different potential buyers for each property with replacement from the model-implied distribution $\{\pi_m(c, s)\}_{c=1}^B$. The $C = 1,000$ repetitions deliver a distribution over the statistics τ and Ep. For each property, we can also compute the transaction probability and the potential price distribution. Implementation details are in Appendix A.4.

As the procedure draws both buyer and seller shocks and simulates matches without conditioning on which buildings traded or which buyers met which sellers, it is not subject to the selection bias inherent in observed transactions. The procedure replays the same core ingredients –valuations, meeting probabilities, and the trading rule that generate actual trades, but over the full set of feasible counterfactual matches.

The potential price distribution is a useful tool in private markets with frictional trade, with applications to price setting strategy, assessment of broker skill, and risk management. The probability of trade as well as the interquartile range of the potential price distribution are useful measures of the liquidity of an asset or market (segment).

This algorithm is also how we compute counterfactuals. To understand the importance of a subset G of investors, we recompute the potential price distribution removing that group from the buyer pool and replacing them by potential buyers drawn from all other types, the set $B \backslash G$. Comparing the baseline and counterfactual means of potential price distributions Ep, help us see the effect of each group of buyers on average prices. Comparing IQRs provides a metric of how the buyer pool affects the price risk of an asset. Comparing trading probabilities τ allows us to assess the importance of group G for the liquidity of an asset or market.

Beyond examining the price impact of excluding certain buyer types from specific markets, we are also interested in studying spillover effects. When a particular investor type purchases a property, other potential buyers who might have acquired that asset are displaced. By identifying these displaced buyers, we can track which properties they end up buying. This analysis provides insight into the substitution patterns along dimensions such as as asset sector, location, and quality. Comparing these patterns across different investor types and time periods helps us understand how competition for assets in one segment influences liquidity, pricing, and allocation dynamics in other segments.

4 Data

Our data consists of the universe of commercial real estate transactions in the United States over $0.5 Million in the four main commercial property sectors (apartments, office, retail, industrial) between 2001 and 2023, about 370,000 transactions representing $10 trillion in transaction value. The data are from MSCI Real Capital Analytics (RCA). We define 60 geographies, which roughly correspond to metropolitan areas, which we refer to as markets.[16]

4.1 Property and Investor Characteristics

We include fifteen building features in x_{nt}. The first eight are measured at the property level: log asset size (measured in square feet for office, retail, and industrial, and in number of units for apartments), an indicator for whether the building is located in a superstar

[16] Appendix A.2 shows the mapping from RCA's markets to our 60 markets.

city,[17] an indicator for whether the property is located in the central business district (CBD), the number of floors, log building age, log renovation-adjusted age, building-level occupancy rate, and a property sub-type indicator.

The next five characteristics in x_{nt} are local macro-economic variables. Two vary at the market-level: log population (for apartments, from the BEA) or log employment (for the other sectors, BEA) and log real personal income per capita (BEA). Two vary at the market-sector level: the occupancy rate (NCREIF) and the net operating income (NOI) growth rate (NCREIF).[18] The fifth one is the net effective rent (NER) per square foot for office, retail, and industrial buildings and the NOI per unit for apartments. For simplicity, we will refer to this variable as NER for all four sectors. The NER is measured at the property level, if possible, or at the block level (small geographic area) where the property is located otherwise. Appendix A.3 describes the details. We interpret the NER as a (time-varying) measure of asset or neighborhood quality. This NER variable, which is the most important of the five local macroeconomic variables, is new to the hedonic literature.

The final two characteristics that enter into x_{nt} are two transaction type categorical variables: whether the asset belongs to a distressed sale and whether the asset belongs to an entity sale–where the entire property portfolio of a legal entity transacts.

We include six investor characteristics in z_{it}. RCA has spent enormous effort unraveling the true buyer(s) and seller(s) involved in each transaction, as well as determining each investor's type classification.[19] The first element of z_{it} is the investor's type (e.g., REIT, REPE fund or foreign investor). Next, we build portfolios for each investor based on transactions data using the perpetual inventory method. We include in z_{it} the log dollar value of the investor portfolio across all sectors, the share of the overall portfolio in superstar markets, the share of the portfolio in the same market as the focal asset, the share of the portfolio in the same sector as the focal asset, and whether the investor transacted as a part of a Joint Venture (JV). These portfolio variables are computed just prior to the transaction for the sellers and just after the transaction for the buyers, so that the transacted asset is always included for both the seller and the buyer.

[17]Superstar cities are defined as those markets with office prices above $350 per square foot in 2023. The superstar dummy is one for the following markets (for all sectors): Washington DC, DC VA burbs, DC MD burbs, Chicago, Los Angeles, Orange Co, Manhattan, NYC Boroughs, San Jose, San Francisco, East Bay, Seattle, Boston, Austin, Miami/Dade Co, and San Diego.

[18]The National Council for Real Estate Investment Fiduciaries (NCREIF) is a consortium of large institutional real estate investors that share data on their institutional-quality portfolios allowing for the construction of market-level cash flow and occupancy statistics. We thank NCREIF for making these data available to us. We build a bridge between the NCREIF markets and our 60 markets. Appendix A.2 provides the details. Net operating income is a standard cash-flow metric in CRE markets that subtracts operating expenses from rent revenues.

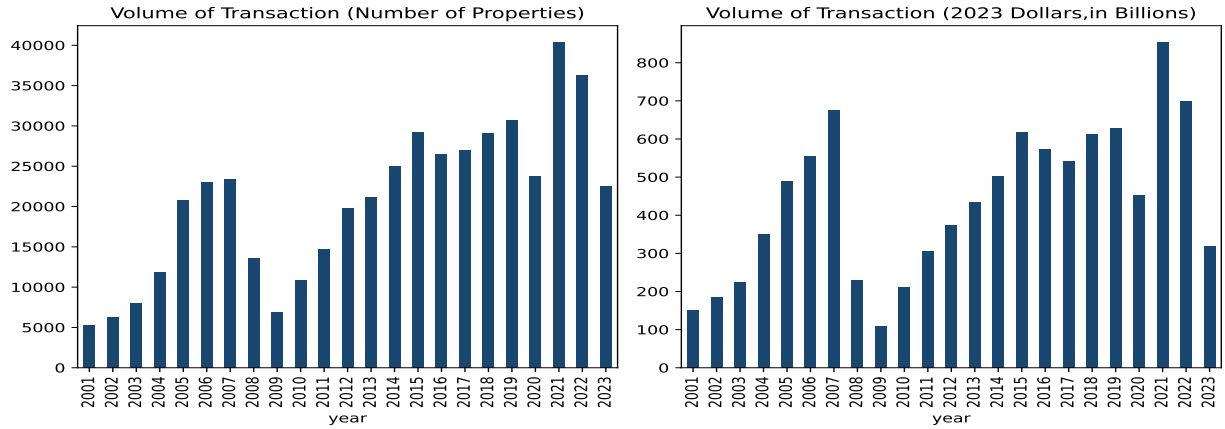
[19]This is a difficult task since CRE properties are usually bought using a separate LLC for each building, typically with a non-descriptive name. The RCA team has spent years unraveling who is behind each LLC, often cross-checking each transaction using multiple sources.

Appendix A.1 shows our data cleaning algorithm that leads to the final data sample used in the estimation.

4.2 Summary Statistics of Transactions

Figure 1 shows the number of transactions (left) and the dollar volume of transactions (right) in our final sample. All dollar values in the paper are expressed in 2023 real dollars. Henceforth, we use the word volume to describe the real dollar value of transactions. The graph shows three boom periods, from 2001 until 2007, from 2009 until 2019, and in 2021-22. It also shows three busts, the Great Financial Crisis in 2008-09, the Covid-19 recession in 2020, and the interest rate hiking cycle in 2022-23. Appendix Table A.4 shows transaction volume by year in table format. There are additional transaction dynamics that are sector specific, such as the decline in retail due to the rise in e-commerce and the decline in office due to the rise in working-from-home in the post-2020 period (Gupta et al., 2025). Figures A.1 and A.2 show transaction volume by sector and property subtype, respectively.

Figure 1: Commercial Real Estate Transactions



Notes: The figure shows the aggregate number of transactions across all property sectors, i.e. Apartments, Office, Industrial and Retail, by year between 2001 and 2023. The left panel shows the transaction volume in terms of the number of properties traded while the right panel shows the total dollar value of properties traded, expressed in billions of 2023 U.S. dollars.

Property transactions are about evenly split between six size categories: >$250M, $100-250M, $50-100M, $25-50M, $10-25M, < $10M. Table A.3 shows the breakdown. Most investors stay within one of these size buckets.

Table A.5 reports transaction volume for the 16 largest markets, as well as all other 44 markets combined. Not surprisingly, Manhattan is the largest property market in the sample, followed by Los Angeles. There are large differences in the importance of each sector across markets. Office accounts for over 60% of volume in Manhattan, Washington

DC, and San Francisco, but only 34% of volume in Los Angeles and 24% in Dallas. About half of total transaction volume (and 40% of the number of transactions) is concentrated in the largest 16 markets.

About 75% of transactions (accounting for 70% of volume) are conventional sales, where one property changes hands. About 19% of transactions (21% of volume) are portfolio sales, where a group of properties is sold together. The remainder of the transactions are entity transactions, where an entire corporate entity changes hand, or distressed sales (foreclosures or debtor/trustee sales). In the case of portfolio and entity sales, RCA imputes the price of each property that is part of the transaction. We further improve on RCA's imputation as explained in Appendix A.1. Table A.6 shows the transaction breakdown by transaction type.

4.3 Summary Statistics By Investor Characteristics

Next, we turn to investor characteristics. Table A.7 shows how we map RCA's investor types to a smaller group of eight investor types. Our breakdown creates meaningful heterogeneity in terms of constraints, investment objectives, geographic reach, etc. Table 1 shows the transaction breakdown by investor type. About 600 Real Estate Private Equity funds (REPE) account for 10-12% of volume. About 3,400 Institutional investors, made up of investment managers, pension funds, insurance, banks and finance companies, and open-ended funds, make up about 14% of volume. The two largest investment categories are local and national Owner-Operator-Developers (OOD_L and OOD_N, respectively), accounting for about 15% and 30% of volume, respectively. These are typically smaller, more geographically-concentrated, private firms that are fully dedicated to commercial real estate. They tend to have substantial discretion in terms of investments, but are subject to tighter financial constraints. We define local OODs as those we see transact in only one market, and national OODs as those we see transact in more than one market. REITS are publicly-listed commercial real estate vehicles and account for 12% of volume. We include real estate operating companies and listed funds in this category as well. As listed entities, REITs are subject to much stricter oversight from the regulatory and financial analyst communities, and face difficulties raising additional equity when their share price trades below net asset value. Foreign investors include sovereign wealth funds and foreign OODs, but also the foreign investors of all other investor types. They represent about 6-8% of volume.[20] Users consist of corporations, governments, non-profits, educational, and religious institutions and charities that own their own real estate. They account for

[20]Table A.7 shows that there are some foreign entities among most investor types. For example, about 33% of the pension fund buys and 21% of pension fund sells are by foreign pension funds. We reclassify all foreign investors into the foreign category irrespective of their original type classification.

about 5% of transactions. Individuals are made up of high-net worth buyers as well as non-traded REITs who serve high-net worth households, and account for about 3-4% of volume. For about 1-2% of volume we do not know the investor type, and create an Unknown category.

Table 1: Transactions By Investor Type

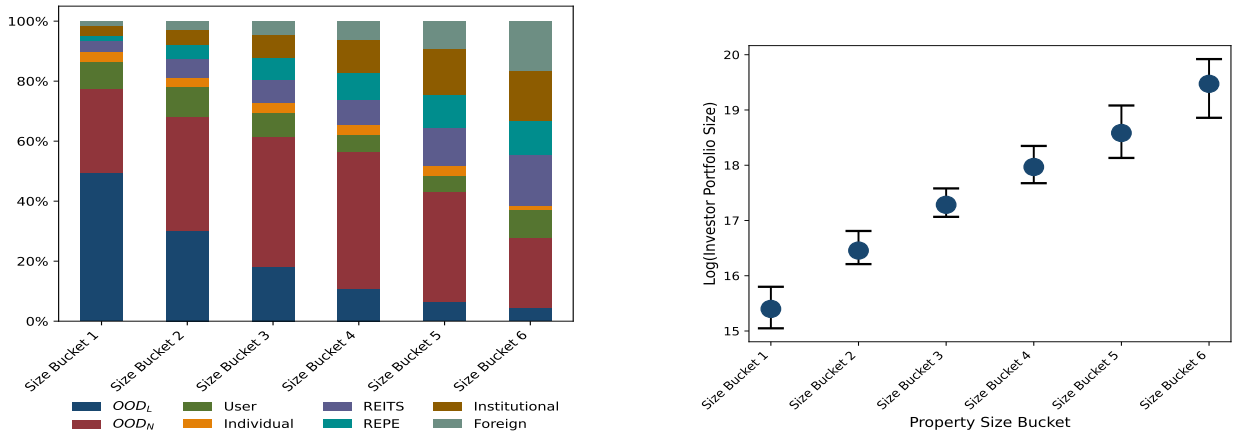
	Buyer (#Trans)	Buyer ($ Vol)	Buyer (% Vol)	Seller (#Trans)	Seller ($ Vol)	Seller (% Vol)	Unique Investors
REPE	28,853	1241	12.30	22,058	1031	10.22	596
Institutional	38,066	1479	14.66	38,148	1371	13.59	3,435
OOD_L	147,030	1316	13.05	161,967	1656	16.41	238,140
OOD_N	150,678	3083	30.56	131,966	3081	30.54	25,385
Individual	19,453	406	4.02	19,399	285	2.82	15,811
REITS	33,518	1182	11.72	35,135	1254	12.43	389
Foreign	17,606	844	8.37	13,055	616	6.11	2,782
User	28,771	418	4.14	33,663	554	5.49	29,845
Unknown	12,044	119	1.18	20,627	241	2.39	7,802
Total	476,018	10,088	100	476,018	10,088	100	324,185

Notes: The table shows the number and dollar volume of buy and sell transactions for various investor types, i.e REPE (Real Estate Private Equity), Institutional Investors, OOD_L (Local Developers), OOD_N (National Developers), Individual investors, REITs, Foreign Investors and Users from 2001 to 2023. Volume is expressed in billions of 2023 U.S. dollars. The table also reports the share of buy and sell volume, which sums to 100% across investor types.

Table A.8 shows that 87% of transactions by count and 73% by volume involve a single buyer and a single seller. A further 10% by transactions and 20% by volume involve one buyer and two sellers or two buyers and one seller. The remainder has multiple buyers and multiple sellers where we observe up to four buyers and up to four sellers. Table A.9 provides additional detail on these Joint Ventures, in terms of which investor types tend to collaborate with each other.

Investors are also heterogeneous in terms of their portfolio size, the size of the typical asset they trade, and their geographic, sector, and asset quality concentration. Figure A.3 shows the distribution of investor portfolio sizes at the end of 2023. The left panel of Figure 2 shows that OOD_L predominantly hold smaller properties. The share of OOD_L transactions diminishes for larger property sizes, giving way to a more diverse mix of investor types in higher size buckets. In contrast, Foreign investors, Institutions, and REITs become increasingly prominent in the larger size buckets. This could indicate a preference or ability to invest in larger properties. The right panel of Figure 2 illustrates the average log of investor portfolio size by property size bucket, with error bars showing the inter-quartile range. A clear upward trend in log portfolio size is evident, suggesting that investors transacting larger buildings (by dollar value) tend to have larger portfolio sizes. This could reflect economies of scale or better access to capital for larger investors.

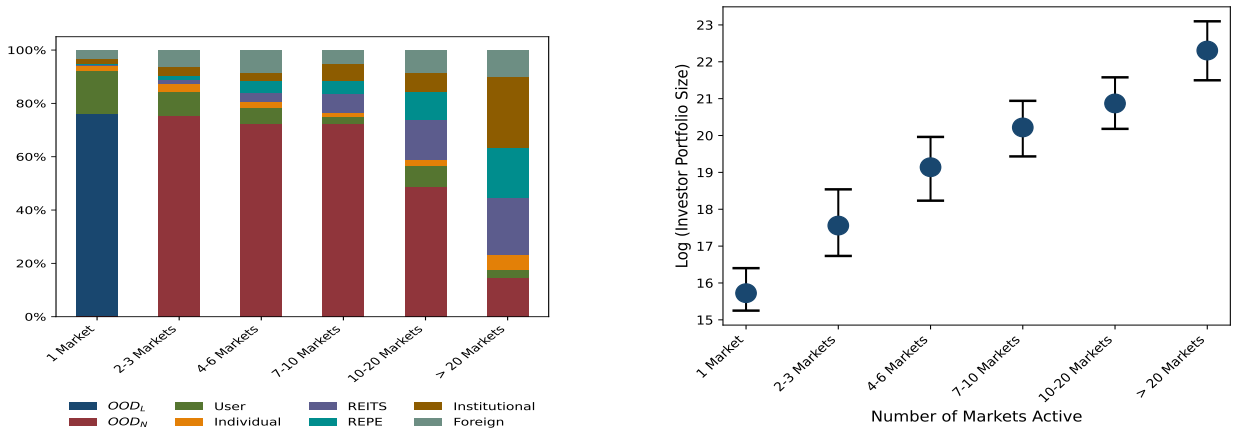
Figure 2: Size Concentration



Notes: This figure shows investor concentration by the aforementioned property size buckets at the end of 2023 for investors that own at least one property. The left figure shows the distribution of investor types by the size buckets while the right figure shows average log investor size (and inter-quartile range) of average log(investor portfolio size) for each property size bucket.

The left panel of Figure 3 summarizes the distribution of investor types by the number of markets the investor is active in. Investors of type OOD_L are active in a single market by definition. Many Users also only show up in a single market. REPE, REITs, and institutional investors, on the other hand, are more likely to hold geographically-diversified portfolios. The right panel shows a strong positive correlation between investor size and geographic diversification. Overall, these facts suggest that there is a substantial amount of market segmentation.

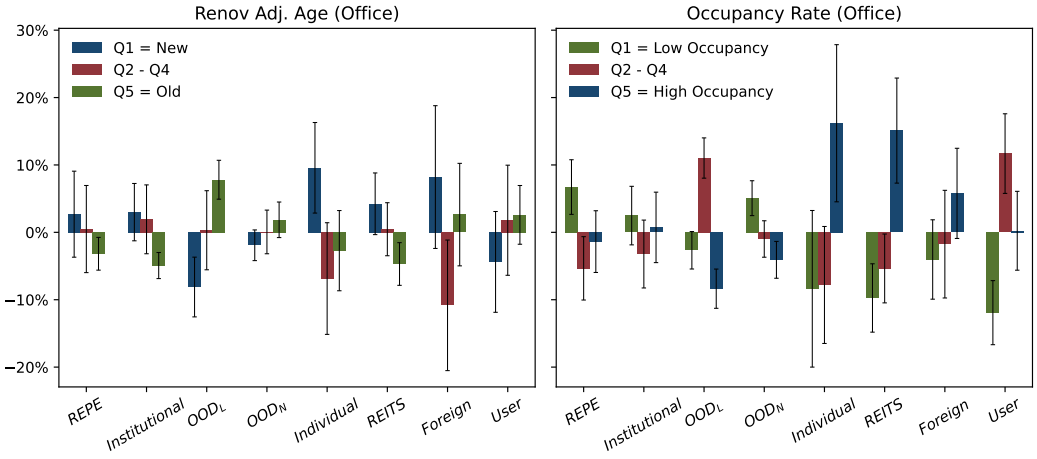
Figure 3: Geographic Concentration



Notes: This figure shows investor concentration by geography at the end of 2023 for investors that own at least one property[17]. The left figure shows the distribution of investor types by the number of aggregated RCA markets they are active in while the right figure shows average log investor size (and inter-quartile range) of average log(investor portfolio size) by the number of markets the investor is active in.

Figure 4 provides an interesting example of preference heterogeneity across investor types for two property characteristics: renovation-adjusted age and occupancy rate. Offices that are older and have lower occupancy are considered to have higher cash-flow risk. We see that REITs and foreign investors disproportionately buy newer, higher-occupancy properties, while REPE funds tend to invest in properties that are older with occupancy issues, usually with a view towards improving them.[21] We find similar results for other sectors beside office in Figures A.6 and A.7.

Figure 4: Investor Heterogeneity and Property Cash-flow Risk



Notes: For each year, we sort all office properties into quintiles of the relevant characteristic: renovation-adjusted age in the left panel and occupancy rate in the right panel. For each investor type, we compute the share of purchases across quintiles and subtract the market-wide shares for the same year. Bars show the time-series average of this difference, with vertical lines indicating the standard deviation across years.

4.4 Flow Dynamics

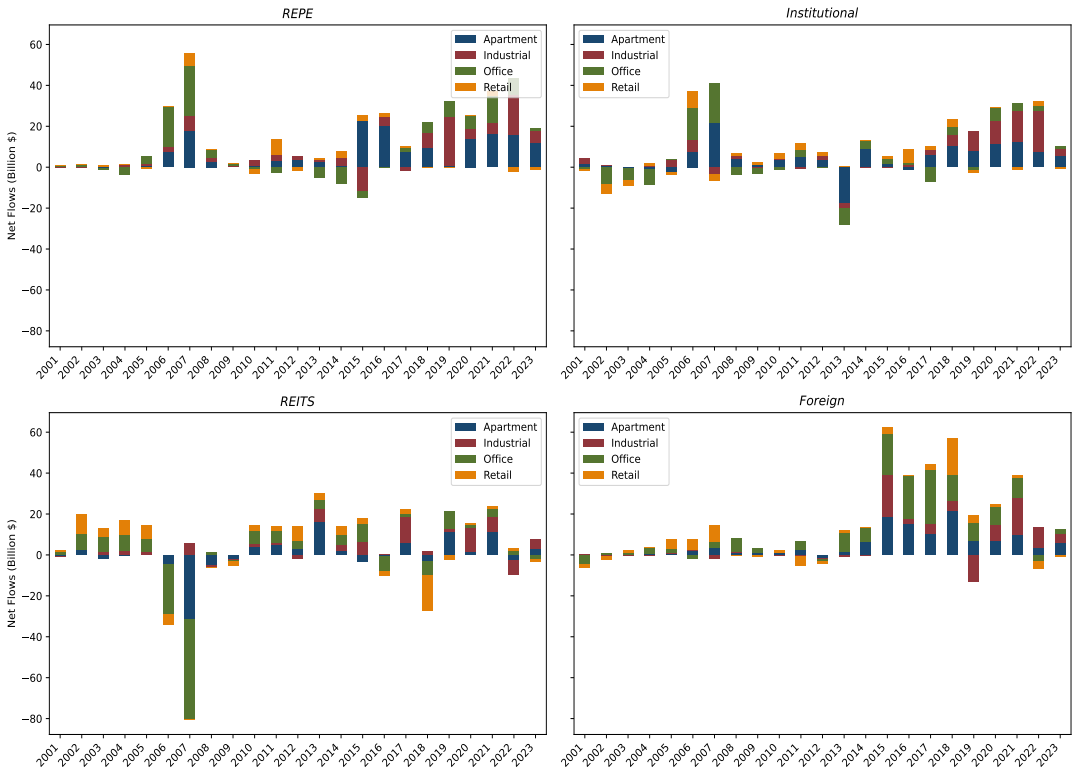
Identifying buyers and sellers at the transaction level, and aggregating their buy and sell trades, allows us to analyze aggregate capital flows by investor type. Figure 5 shows net flows for select investor types by sector and year. It offers insights into how capital was allocated across CRE sectors over time.

One interesting trend is the large inflow of foreign investors from 2013 onward. These inflows totaled $317 billion, with a significant concentration in Office ($123 billion, 39%) and Apartments ($106 billion, 33%).[22] Figure A.4 further reveals that most of this foreign

[21]The typical 10-year lifespan of commingled REPE funds gives them the flexibility to execute on such "value-add" or "opportunistic" investment strategies. By contrast, REITS are subject to quarterly reporting to their investor base and tend to focus on "core" or "core plus" investment strategies.

[22]This surge included landmark deals such as Canada's Brookfield Property Partners' $18 billion acquisition of the REIT GGP and its $11.4 billion acquisition of the REIT Forest City in 2018. Singapore's sovereign

Figure 5: Net Flows by Investor Types



Notes: The figure plots the net purchases of REPE funds, REITs, Institutional, and Foreign investors, calculated as the dollar value of purchases minus sales in 2023 dollars for each sector-year.

investment in Offices and Apartments was concentrated in superstar cities, primarily Manhattan, Los Angeles, San Francisco, and Boston. These patterns motivate us to study the price and quantity effects of such foreign capital inflows in Section 7.

Another notable pattern is the high volatility in REIT flows, particularly in the lead-up to the 2008 financial crisis. Several REITs aggressively offloaded properties in 2006 and 2007. Notable examples include Sam Zell's $39 billion entity sale of his Equity Office Properties REIT to Blackstone in 2007 and the $22 billion purchase of Archstone-Smith REIT by Tishman Speyer and Lehman Brothers. These major transactions, occurring on the eve of the Great Financial Crisis, provide an interesting setting for counterfactual analyses, helping us examine how market outcomes might have differed had these REITs not liquidated their holdings.

Finally, we document trading patterns among investor types in Table A.10. We see significant trade both within and across types. One quarter of all transaction volume involves buyers and sellers of the same type, reflecting differences in valuation even when

wealth fund GIC made a major push into the US industrial market by acquiring IndCor Properties from the REPE fund Blackstone for $8.1 billion in 2015.

investment strategies are similar.

5 Estimation of the Valuation Model

We estimate the valuation model (8) from observed transaction prices. The left-hand side variable is the log real price per square foot for sectors O, I, and R and the log price per unit for sector A.

5.1 The Linear Hedonic Valuation Model

To establish a reference point, we first consider the standard valuation model in the literature: the linear hedonic model (LHM) which assumes that log price per square foot or per apartment unit is a linear function of property characteristics. This model ignores the role of investor characteristics in valuation.

Panel A of Table 2 reports the goodness of fit (R^2) of the LHM. The first row of the table shows the results of a standard LHM that only includes features of the building (age, size, etc.) and of the deal (whether the transaction is part of an entity sale and whether it is a distressed sale), for a total of ten features. The R^2 ranges from 34.4% for office to 51.7% for retail. Table B.1 reports the coefficient estimates and their standard errors in columns (1), (3), (5), and (7) for the four sectors, respectively. Point estimates for the "hedonic prices" are precisely estimated and have the expected sign.

In the second row, we add the five local macro variables, such as the occupancy rate and the cash-flow (NOI) growth rate in that market-sector-year or the per-capita income of the residents in that market-year. This includes the NER variable, a cash-flow measure that is either available at the property level or at a fine geographic level, as explained above. The NER variable enters highly statistically significantly in all sectors, as does personal income. Combined, these captures key dynamics in the investment opportunity set. They add between 5 and 10% points to the R^2.

In the third row, we add year fixed effects to capture all other, national macro dynamics such as the interest rate environment or the overall liquidity of the market. Finally, in the fourth row, we add market fixed effects to capture all other sources of constant geographic heterogeneity. This model explains 59% of the variation in prices in the apartment sector, 59% in industrial, 48% in office, and 58% in retail. Columns (2), (4), (6), and (8) of Table B.1 report the point estimates for this specification. This model is the hedonic model against which we wish to assess the role of non-linearities and investor covariates. We note that this is already a much richer hedonic model in terms of covariates than the traditional hedonic model, which corresponds to the first row rather than the fourth row of Table

2. In addition, market and time fixed effects could stand in for the effects of investor composition. By including them in the hedonic model, we make it harder for investor characteristics to matter.

5.2 LGBM model Without Investor Characteristics

Next, we allow for the full set of property characteristics x_{nt}, including time and market fixed effects, to affect valuations in a non-linear way. We also allow for all building features to interact with one another. Importantly, we still leave out investor characteristics.

We estimate a Light Gradient Boosted Machine (LGBM) model, which provides a flexible, data-driven and scalable approach compared to traditional parametric models. LGBM is well-suited for this estimation as it efficiently handles large datasets, captures complex non-linear relationships without requiring manual specification, and is robust to overfitting through built-in regularization.[23] We estimate the valuation model for each sector and model specification separately, calibrating hyper-parameters in each case. We standardize continuous variables to facilitate comparison across sectors.[24,25]

The first four rows of the panel B of Table 2 show the results. The first row shows the full-sample R^2 from the LGBM model estimated on just the 10 asset and deal characteristics, the same features as in the first row of Panel A. The second row adds the 5 local macro variables. The third row adds the year-fixed effects, and the fourth row the market-fixed effects. We refer to the model in row 4 as the hedonic LGBM model. The hedonic LGBM model produces a R^2 that ranges from 65% for Office to 75% for Apartments.

This R^2 in row 4 of Panel B is directly comparable to the R^2 in row 4 of Panel A. The difference between the two sets of R^2 statistics quantifies the importance of non-linearities and interaction effects in asset characteristics. The added flexibility in the functional form between hedonics and valuations boosts the R^2 by 16% points for Apartments, 12% points for Industrial, 17% points for Office, and 14% points for Retail. The substantial improvement in fit from non-linearities and interaction effects is a first main empirical result, and is new in the context of CRE valuation. Below, we analyze which interaction effects and non-linearities matter most.

[23]The hyperparameters are optimized using 10-fold cross-validation to minimize out-of-sample mean squared error. The optimization procedure searches over key parameters, including tree depth, the number of leaves, feature and bagging fractions, and regularization terms to ensure that the model effectively balances flexibility and generalization. Details on the tuning process and parameter search space are provided in Appendix B.2.

[24]Standardizing variables is not a requirement for tree-based models such as LGBM which make splits based on the relative ordering of feature values rather than their absolute magnitudes.

[25]We make use of the flexibility of LightGBM to incorporate categorical variables as predictors directly.

5.3 The Main Valuation Model

We are now ready to add investor characteristics. Again, we proceed in steps. In the first row of Panel C of Table 2, we add investor type-fixed effects, leaving out the indicator for the Unknown category: REPE, Institutional, OOD_L, OOD_N, Individual, REIT, Foreign, User. In the final row, we add the other five investor characteristics in z_{it}: the log portfolio value of the investor, the share of the total portfolio value in superstar markets, in the same market as the focal building, in the same sector as the focal building, as well as the JV dummy.

Table 2: Valuation Model, Goodness of Fit

Sector	Apartment		Industrial		Office		Retail	
	R^2	R2C	R^2	R2C	R^2	R2C	R^2	R2C
Panel A: Linear Hedonic Model								
Asset & deal features	41.1		44.7		34.4		51.7	
+ Macro Vars	51.3		52.6		44.5		56.0	
+ Year FE	52.8		53.6		45.3		57.0	
+ Market FE	**59.2**		**58.5**		**48.0**		**58.4**	
Panel B: Light Gradient Boosting machine Model, No Investor Features								
Asset & deal features	54.1		53.9		50.1		63.1	
+ Macro Vars	71.3		67.2		62.3		68.7	
+ Year Dummies	72.3		68.1		61.9		70.0	
+ Market Dummies	**75.3**		**70.1**		**64.9**		**72.0**	
Panel C: Light Gradient Boosting machine Model, With Investor Features								
+ Investor Types	78.9	14.6***	74.7	15.2***	69.9	14.1***	76.6	16.3***
+ Portfolio Vars	**85.6**	41.6***	**86.0**	53.9***	**84.5**	55.7***	**87.5**	55.4***

Notes: This table summarizes the performance R^2 from valuation model estimations in equation (8) using the linear hedonic model in Panel A, the hedonic LGBM model in Panel B, and the main LGBM model with investor characteristics in Panel C. In each panel, we successively add features, as described in the text.

We compute the comparative R-squared (R2C) for the models in Panel C (the alternative models), relative to the LGBM model without investor characteristics in row 4 of Panel B (the null model). The R2C is the ratio of the difference in R-squared values between the alternative and null models to the unexplained variation in the dependent variable under the null model:

$$R2C = \frac{R^2_{alt} - R^2_{null}}{1 - R^2_{null}}$$

Kim (2024) develops the asymptotic distribution of R2C and a test of the null hypothesis that R2C equals zero. The stars on the R2C statistic reflect the p-values of this test.

The first row of Panel C shows that the investor type provides useful additional information for valuations. Relative to the LGBM model without investor characteristics, the

R^2 goes up between 3.6 and 4.9% points, reducing the unexplained variation in the LGBM model without investor features by between 14 and 16% (R2C). The gains from adding the other investor characteristics in the second row of Panel C are substantially larger. This model is our main valuation model. The R^2 reach values of 84.5% points for Office, 85.6% for Apartments, 86.0% for Industrial, and 87.5% points for Retail. Comparing the last row of Panel C to row 4 of Panel B, investor variables increase the R^2 by 10.3% points (A), 15.8% points (I), 19.5% points (O), and 15.5% points (R). They reduce the unexplained variation in model 4 of Panel B by 41.6% for A, 53.0% for I, 55.7% for O, and 55.4% for R. These R2C gains are highly statistically significant in all four sectors and economically large.

In sum, investor characteristics are a crucial "new hedonic" in commercial real estate valuation models that strongly add to the valuation model's explanatory power. When different investors value asset characteristics differently, it matters for the transaction price which investors are in the market at a given point in time.

Table B.5 shows the out-of-sample (OOS) performance of the LGBM model. We use a 10-fold cross-validation procedure where the R^2 is calculated as the average of the hold-out samples across all folds. Naturally, the main model has a lower OOS than in-sample R^2. However, both the absolute level of the R^2 and the gain in explanatory power relative to the hedonic model without investor characteristics remain large. Investor characteristics explain an additional 12.4% (A), 17.5% (I), 20.4% (O), and 17.6% (R) points of the variation in property transaction prices, reducing unexplained variation between 42.8% and 54.4%. These OOS gains are similar to the in-sample gains and remain highly statistically significant in all four sectors. We return to OOS predictability in Section 7.1.

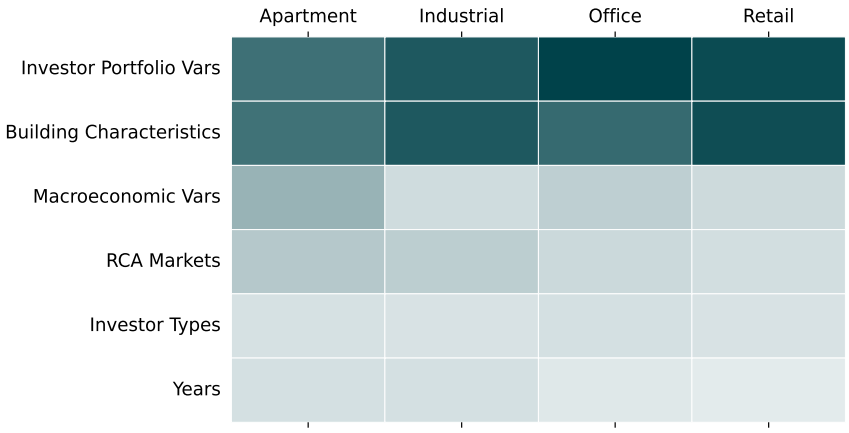
The main valuation model also allows us to compute the residual variance for buyers and sellers by market and year. We allow for market-time variation but impose: $\sigma^2_{b,m,t} = \sigma^2_{s,m,t}$. We estimate this from the residual variation in prices $\sigma^2_{b,m,t} = 2\sigma^2_{p,m,t}$.

5.4 Feature Importance in the Valuation Model

The previous analysis showed that (i) investor characteristics matter and (ii) that allowing for non-linearities and interaction effects matters. Our main model allows for traditional hedonics to be valued differently by different types or sizes of investors, for example. We systematically investigate the importance of each feature using SHAP values (Lundberg and Lee, 2017), defined in Appendix B.3.1.

Figure 6 highlights the most important features selected by the model, organized by category. Investor portfolio characteristics are the most important feature category, ahead of property characteristics and macro variables. Figure B.2 breaks out the features further. Among all variables, investor size and property size are the two most important features in all sectors. The share of properties the investor has in the same market and in the

Figure 6: Valuation Model Feature Importance by Category



Notes: Figure shows the importance of features selected by the LightGBM Model. Feature importance is defined as the absolute SHAP values for each feature summed across all observations. We aggregate the importance by feature types. To facilitate comparability, feature importance is normalized so that the sum across feature groups equals 1 in each sector.

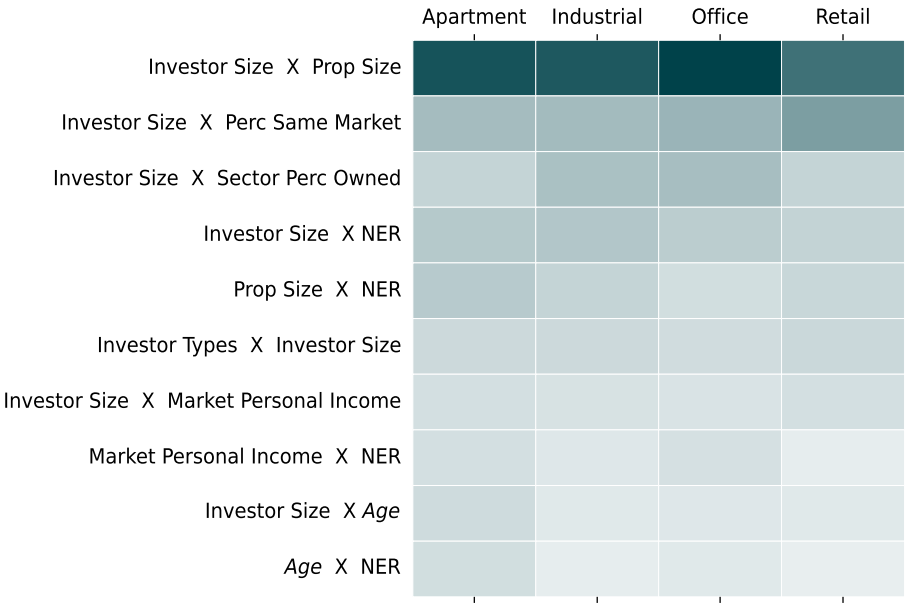
same sector as the focal asset, and the investor's type are all among the ten most important features. NER is the most important macro variable and market fixed effects are important as well.

To analyze how transaction prices vary with predictors, we use SHAP dependence plots. SHAP values measure how the observed level of a feature shifts the predicted price, with positive (negative) SHAP values indicating that the feature at its realized level raises (lowers) predicted prices. Figure B.3 shows the results for the most important continuous features. Valuations are concave in investor size but convex in property size. Investors who specialize in the office sector have higher valuations for office, while sector focus matters much less for valuations in other sectors. A high NER (expensive neighborhood with high-quality assets) affects office and apartment valuations disproportionately. Many of the effects are nonlinear in ways that would be difficult to uncover without the help of a ML model like LGBM.

To study time trends in valuations, we plot the SHAP values for the year fixed effects in Figure B.4. Valuations decline sharply in the GFC and recover steadily after 2012 with industrial and apartments showing the strongest recovery. Retail valuations begin falling post 2016, coinciding with the retail apocalypse and rise of e-commerce. All sectors saw their valuations fall in 2023 following a sharp rise in interest rates.

SHAP values capture both direct and interaction effects, making them ideal for analyzing feature interactions. Beyond attributing predictions to individual features, SHAP can be extended to decompose contributions into pairwise interactions. We compute SHAP interaction values for all pairwise interactions as detailed in Appendix B.3.1. Figure 7 shows the 10 most important interaction effects for each sector. Interactions between

Figure 7: Most Important Interaction effects



Notes: This figure presents the ten most important interaction effects, measured by averaging the absolute SHAP interaction values across all observations for each feature pair within a sector, averaged across buyer and seller valuation models. To facilitate comparability, values are normalized so that the sum across all feature pairs equals 1 in each sector.

investor characteristics, notably investor portfolio size, and building features such as property size, NER, and building age are among the most influential. There are also several interactions of two investor characteristics in the top ten. This result reinforces the critical role of investor characteristics in the valuation model, as well as an important mechanism: different investors value different hedonics differentially so it matters which investors are in the market at a given point in time. These interaction effects of asset and investor characteristics are at the core of valuation model (2).

Figure B.5 studies the interaction between investor size and property size, the most important interaction identified in Figure 7. Large investors tend to value small properties less, while small investors assign lower valuations to large properties across all sectors.

5.5 Addressing Overfitting Concerns

To address concerns that the gains in explanatory power from the LGBM model reflect overfitting rather than a genuine signal, we select our hyperparameters conservatively, including a maximum tree depth of 5, as discussed in Appendix B.2. In addition, Figure B.1 and Table B.4 show that adding the most important non-linear and interaction effects, identified by the LGBM model, to the standard linear hedonic model significantly improves predictive accuracy the LHM.

6 Estimation of the Transaction Model

6.1 Point Estimates in the Meeting Model

Next, we estimate the parameters of the meeting model in (6). The variables that determine whether a potential buyer b meets a building that has been listed for sale by a seller s are: the buyer's log portfolio size S_b, an indicator N_b which is one if the buyer owns more than two buildings, the log difference between the buyer and the seller's portfolio sizes $\Delta S_{b,s}^{-1}$, an indicator $T_{b,s}$ which is one if the buyer and the seller have traded with each other before, and four consideration set variables in $\delta_{b,s}$ which measure how similar the building for sale is to the assets in the buyer's portfolio in terms of asset size, sector, location, and NER. We estimate one model for each sector-year using all transactions from that sector-year.[26] Table 3 reports the coefficient estimates for $(\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5)$ that minimize the loss function $-\sum_s \ln \pi_r(b, s)$, with π_r as defined in (9). The parameter estimates $\hat{\lambda}_t$ are averaged over time: $\hat{\lambda} = (T)^{-1} \sum_{t=1}^{T} \hat{\lambda}_t$. The standard errors for $\hat{\lambda}_t$ are obtained using bootstrap and then also averaged over time.

Since most investors in our sample transact infrequently, our baseline specification restricts the negative set to investors that have transacted at least once in the preceding five years. Appendix D.3 shows that the results are robust to alternative definitions of the negative sample. For potential buyers that enter the transaction year with zero assets, we assume that they own the median number of properties (which is 1) with characteristics equal to the average portfolio in that year. Since the asset listed for sale is included in the portfolio of the actual buyer, we also add it to the portfolio of all other potential buyers to avoid a bias. Not doing so would make it artificially easy for the algorithm to select the true buyer from the set of all potential buyers.[27]

As explained above, rather than including all potential buyers in the likelihood estimation, we use an important insight from the NLP literature showing it is sufficient to consider a negative set with small cardinality K. We explore values of K between 50 and 1,500 and choose $K = 1,000$ for our main results, since parameter estimates stabilize around that value. We discuss robustness of our results to alternative values for K in Appendix D.2.

The first row of Table 3 shows that the size of the buyer has a large positive effect $\hat{\lambda}_1$ on the meeting probability. A potential buyer which is 1% larger has about a 1.5% higher

[26] While the valuation model is estimated over 2001-2023, we begin estimating the meeting model in 2002 to ensure that all covariates, including the prior buyer–seller transaction variable $T_{b,s}$, are available.

[27] For investors with no existing holdings in year t, consideration set variables are constructed using a benchmark portfolio that holds the cross-sectional median number of properties (which is one) with characteristics equal to the year-t average portfolio. For symmetry in computing consideration set variables, the focal property is appended to every potential buyer's portfolio, including the realized buyer.

Table 3: Meeting Model Parameter Estimates

Parameter	Description	Apartment	Industrial	Office	Retail
λ_1	Buyer Size	1.48	1.70	1.43	1.39
		(0.03)	(0.04)	(0.04)	(0.04)
λ_2	Buyer has > 2 properties	2.43	2.51	2.43	2.26
		(0.08)	(0.09)	(0.08)	(0.09)
λ_3	Buyer-Seller Size Ratio	2.25	2.89	2.23	2.39
		(0.17)	(0.19)	(0.19)	(0.18)
λ_4	Prior transactional relationship	2.79	2.42	2.68	2.76
		(0.29)	(0.40)	(0.39)	(0.38)
$\lambda_{5,1}$	Asset Value Similarity	7.52	8.25	7.15	7.29
		(0.24)	(0.31)	(0.29)	(0.30)
$\lambda_{5,2}$	Market similarity	5.63	6.36	5.65	5.34
		(0.15)	(0.17)	(0.17)	(0.17)
$\lambda_{5,3}$	Sector Similarity	4.30	4.18	3.29	4.08
		(0.15)	(0.15)	(0.16)	(0.14)
$\lambda_{5,4}$	NER Similarity	7.25	10.14	10.27	7.86
		(0.32)	(0.50)	(0.54)	(0.39)

Notes: This table reports the average estimated values of the meeting model parameters along with their standard errors in bracket. The point estimates $\hat{\lambda}_t$ are obtained by optimizing the objective function (9) each year, and averaged across years. Standard errors are calculated by bootstrap from 1,000 samples for each year, and then averaged across years. See Figure D.1 for parameter estimates by year. λ_1 is the coefficient on log investor size, λ_2 is the coefficient on the indicator that equals one when the buyer owns more than 2 buildings, λ_3 is the coefficient on the buyer-seller size ratio, λ_4 is the coefficient of the indicator that equals one when the buyer and seller have a prior transactional relationship, and $\lambda_{5,1}$ through $\lambda_{5,4}$ are the coefficients of asset, market, sector and NER similarity to the focal asset.

likelihood of meeting the seller. This parameter is identified from the fact that large investors, for example large REITS or REPE funds, are engaged in many more transactions than small investors. The positive and significant $\hat{\lambda}_2$ suggests that owning more than two assets makes a buyer about 2.4% points more likely to meet the property for sale, a useful investor size non-linearity.

The positive and significant $\hat{\lambda}_3$ estimate in the third row shows that buyers are more likely to meet sellers with a similar portfolio size. This is consistent with the positive assortative matching between investor size and asset size shown in the right panel of Figure 2. A positive $\hat{\lambda}_4$ shows that a buyer who has transacted with the focal seller before is about 3% more likely to meet that seller than a buyer who has not traded with the seller before. Related investor familiarity findings have been found in other contexts (Badrinza et al., 2022).

The consideration set variables (the vector $\hat{\lambda}_5$ in rows 5–8) all have the expected positive sign, are large in magnitude, and are estimated precisely. Buyers are much more likely to meet assets for sale that have similar size, location, sector, and quality (NER) to the

assets they already own. Property size alignment ($\lambda_{5,1}$) and asset quality alignment ($\lambda_{5,4}$) are the most important similarity measures. A potential buyer of apartments that is 1% more invested in the same market is 5.63% more likely to meet the apartment property for sale. A 1% increase in the alignment of NER increases the chance of a buyer-seller meeting by 10.27% for offices. The consideration set variables suggest that specialization is an important feature of CRE portfolios.

6.2 Feature Importance in the Meeting Model

To better understand the importance of each of the variables in the meeting model, Table D.1 successively leaves out each of the variables, reporting the resulting coefficient estimates and log likelihood. The last column reports the percentage reduction in the probability that the true buyer meets the seller when a specific variable is excluded from the estimation.

We find that buyer size is the most important determinant of the meeting probability. Excluding it from the estimation results in a 56-64% drop in the probability that the algorithm identifies the correct buyer across sectors. The likelihood of the restricted model is about 100 log points lower than that of the baseline model across all four sectors. Excluding buyer size leads to large changes (biases) in the remaining coefficient estimates. The second size measure, owning at least two properties, increases the probability of identifying the right buyer by 9–19% across sectors.

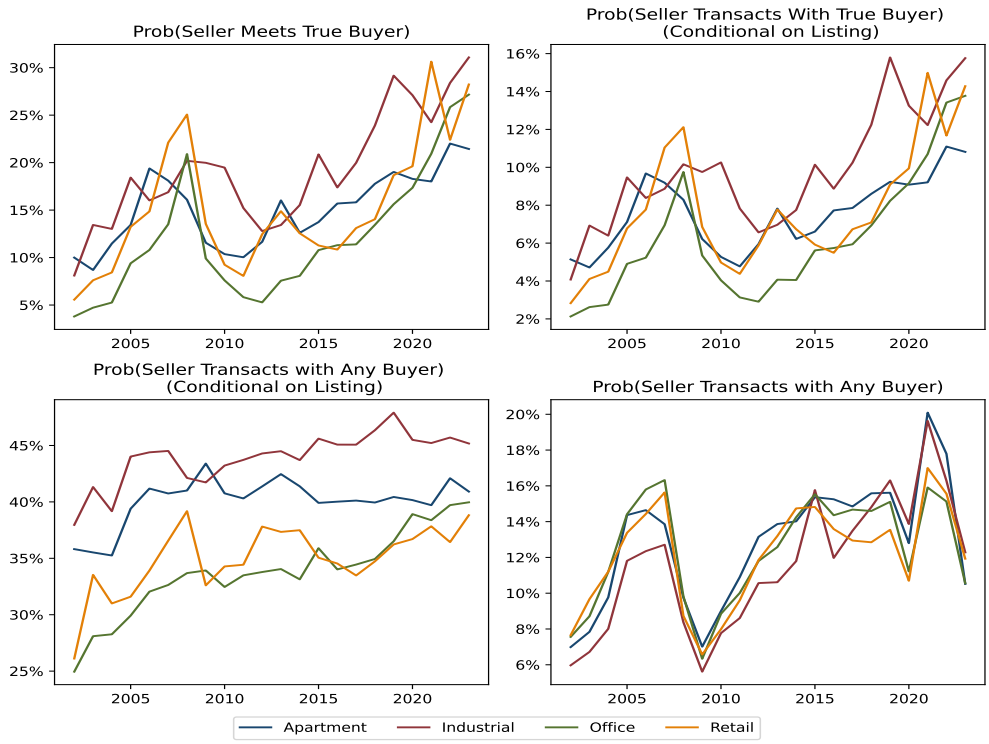
Size similarity between buyers and sellers has modest importance. Excluding it lowers the probability of identifying the correct buyer by about 14–23%. Prior transactional relationships between buyer and seller also matter: excluding this variable reduces the probability of identifying the true buyer by 5–10%. Its overall impact is muted by the fact that such prior transactions are only present for about 8% of buyer–seller pairs.

The consideration set variables are all quantitatively important, with some variation across sectors. The most important one in every sector is asset location similarity ($\lambda_{5,2}$). Excluding it lowers the probability of identifying the correct buyer by 42–56% and the model's likelihood by 60–75%. Asset value similarity is the second most important consideration set variable in every section, with a 35-44% reduction in likelihood from excluding it.

The top left panel of Figure 8 displays the likelihood that the model identifies the true buyer from among the negative set of 1,000 potential buyers plus the true buyer. That likelihood is consistently high, averaging 20% at the end of the sample. For comparison, random search would result in a 0.1% probability of identifying the true buyer. Our meeting model therefore delivers a 200x gain over random search. This underscores the model's ability to capture the key determinants of buyer-seller matches in CRE markets. Our model therefore provides a natural benchmark for other matching models (including

structural models).

Figure 8: Meeting and Transaction Probabilities



Notes: The top left panel shows the probability of meeting the actual buyer from a set of $K = 1,000$ potential buyers (negative set) and the actual buyer, as implied by the benchmark model estimation. The top right panel displays the probability that the seller transacts with the actual buyer, combining the meeting and transaction probabilities. The bottom left panel shows the probability that the seller transacts with any buyer, conditional on listing, reflecting aggregate transaction likelihoods across all potential buyers. Bottom right panel shows the probability that the seller transacts at all, combining the meeting, transaction, and listing probabilities (Section 6.4).

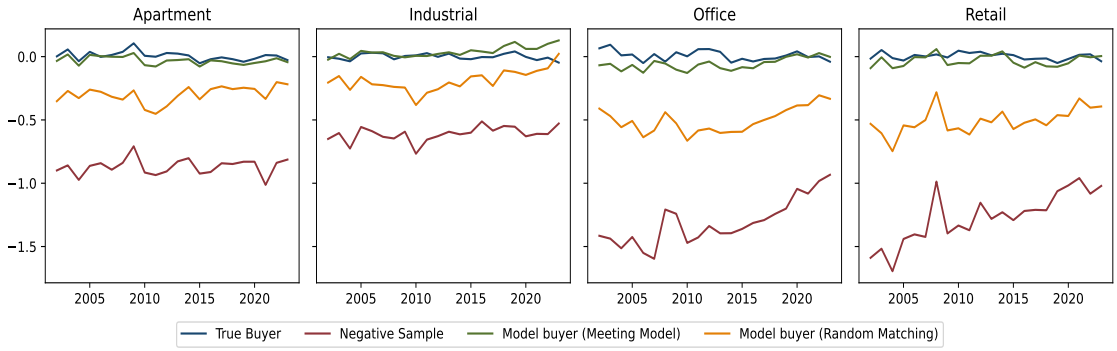
A drop in meeting probabilities during the financial crisis is evident across all sectors. This may reflect the heightened uncertainty and reduced liquidity in the market, disrupting the typical alignment between buyer preferences and assets available for sale. The graph also shows that the model's ability to identify the correct buyer rises from about 8% at the start of the sample to 25% at the end of the sample. While the estimation is done year-by-year, the number of transactions in the data set is growing over time. Also, with every year of additional transaction data, we obtain a clearer picture of the investors' portfolios, as holdings have to be built from past transactions.

Identifying the correct buyer is a high hurdle. After all, in a pool of 1,000 potential buyers, there are usually several buyers with characteristics that are very similar to the true buyer, making exact identification difficult. What matters in practice is whether the model ranks the true buyer near the top. Figure D.2 shows that the true buyer appears among the top 1%, 2.5%, or 5% of predicted matches with high probability (65–85%), highlighting the model's ability to zoom in on the relevant subset of potential buyers.

6.3 Transaction Probability

Having estimated the probability of a meeting, we next turn to the probability that a transaction takes place conditional on a meeting, π_τ. The latter depends on the valuation gap between the buyer and the seller, and hence combines the estimates from the meeting model and the valuation model. Figure 9 plots the valuation gap $h_b - h_s$ for the positive set (the actual buyer) and the negative set (the average among the K potential buyers).

Figure 9: Valuation Gap



Notes: This figure shows the valuation gaps, i.e. difference between the log valuation of the buyer and the log valuation of the seller $h - b - h_s$ for the positive set (the actual buyer), the negative set (the K = 1,000 potential buyers), the model-matched buyer using search with characteristcs-based consideration sets, and the buyer selected in a random-matching model.

Across all sectors, the true buyers consistently exhibit much higher valuation gaps than the other potential buyers. This reflects not only the model's ability to select the correct buyer, but also reinforces the model's assumption that matches occur when buyers perceive gains from trade. True buyers have about a 3–8% higher valuation for the asset than the seller. For the negative set (the non-buyers), the valuation gap is large and negative (-75% to -150% on a log scale).[28] Figure 9 also compares outcomes under the matching of our main model to an alternative model with random matching. We sample 1,000 potential buyer using either random matching or matching with characteristics-based consideration sets, then draw buyer and seller shocks, and verify whether $v_b > v_s$ (following Section 3.3). The valuation gap is $\mathbb{E}[h_b - h_s \mid v_b > v_s]$. Under matching with characteristics-based consideration sets, the valuation gap is positive, about 9% for apartments and 21% for retail, while under random matching it is consistently negative, ranging from –10% to -23%. This highlights the effectiveness of the matching model with characteristics-based consideration sets. Figure D.3 shows that the true buyer typically ranks among the very top predicted matches, with larger positive valuation gaps even

[28]Since it is randomly drawn, the negative set includes some investors who are highly unlikely buyers of a given property. For example, a small investor with a single property in a single market and sector is an unlikely buyer of another property in a different market and sector.

compared to close alternatives.

The top right panel of Figure 8 shows the model-implied probability that the actual buyer transacts with the seller. This probability is high, reflecting the model's ability to accurately capture both which buyer matches with the property and the relative valuations of the buyer and the seller for the asset, i.e., the valuation model delivers a positive probability of sale conditional on a meeting. The bottom left panel shows the model-implied probability of a transaction occurring with any buyer. These probabilities are around 35%. This shows that the model has a high likelihood of predicting a trade for assets that actually traded in the data.

6.4 Listing Probability

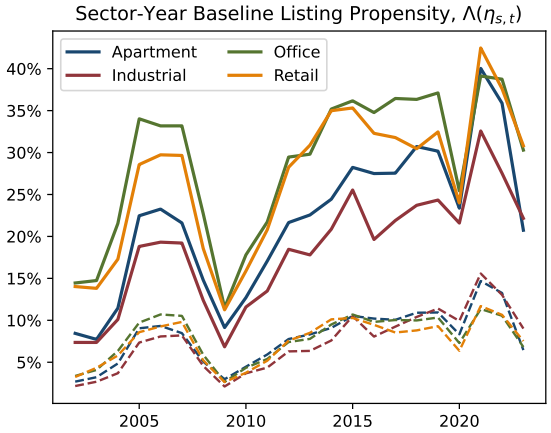
The final component of the transaction model is the listing probability. The left panel of Figure 10 reports the estimated slope coefficients (θ) listing model in (5). Since the seller-side predicted log total valuation and log market size are standardized, their coefficients are the log-odds effects of a one-standard-deviation increase. Converting them to odds ratios, a one-standard-deviation increase in the seller-side predicted log total valuation raises the odds of listing by about 89% in Apartments, 71% in Industrial, 186% in Office, and 136% in Retail. Listing probabilities are increasing in market size for Industrial and Retail but declining for Apartment and Office, once market fixed effects are considered. Paper gains also matter positively for listings in every sector; roughly half the sellers sit on paper gains (0.57 for apt, 0.48, 0.47, 0.45 for retail). Properties sitting on gains have 88% higher listing odds in Apartments, 40% in Industrial, 40% in Office, and 49% in Retail than otherwise similar properties without gains. The last row reports the difference between the model-implied listing probability for properties that actually transact (π_l^{Trans}) and the corresponding probability for properties that do not transact ($\pi_l^{No\ Trans}$). The uniformly positive gap shows that the model assigns systematically higher listing propensities to properties that ultimately trade. The listing specification improves the within-sector-year cross-sectional ranking of likely transactors relative to a listing model that only contains sector-year fixed effects. These differences average 5 percentage points, which is economically meaningful given average listing probabilities of 20-25%.

The right panel in Figure 10 plots the transformed sector-year intercepts, $\Lambda(\eta_{s,t})$, from the listing model in (5). They capture the baseline listing propensity in each sector-year for a property with average covariates. Listing odds fall sharply in 2009 by 10-20% points across sectors, then recover steadily, reaching about 40% in Apartments, Office, and Retail, and 20% in Industrial at the 2021 peak. The Covid crisis of 2020 and the monetary policy hiking cycle of 2022-23 result in reductions in listing probabilities.

Figure 10: Listing Model

Listing Model parameters

	Apartment	Industrial	Office	Retail
Seller-side predicted log total val	0.639	0.538	1.051	0.858
log(Market Size)	-0.408	0.103	-0.106	0.183
Paper Gain Indicator	0.312	0.167	0.168	0.198
Market FE	✓	✓	✓	✓
Seller Type FE	✓	✓	✓	✓
$\pi_l^{Trans} - \pi_l^{No\ Trans}$	0.074	0.045	0.072	0.054



Notes: The left panel shows the estimated slope coefficients (θ) listing model in (5). This right panel shows the transformed sector-year intercept, $\Lambda(\eta_{s,t})$, from the listing model, over time in each sector, as defined in (5). Dotted lines show the fraction of properties that transacted in that year. The gap between the two lines reflects the model-implied likelihood that the seller listed the property but failed to find a suitable buyer.

7 Applications: Prediction and Counterfactuals

Having estimated the valuation and matching models, we now turn to four applications of the model. Section 7.1 uses the model to predict the transaction price out-of-sample. Section 7.2 shows how the model can be used to understand the dynamics of CRE prices, focusing in particular on the period leading up to the 2008-09 financial crisis. Section 7.3 uses the model to uncover investors' substitution patterns across building size, sector, and location. This analysis is an important input to understanding the impact on the CRE ecosystem of, for instance, an inflow of foreign buyers in the high-quality segment of the Manhattan office market. Section 7.4 studies the potential deployment of significant additional real estate private equity capital, focusing on the assumptions about REPE valuations and matching intensity required to reach a target of $150 billion of investment in 2025 and the implications for the CRE segments in which REPE funds are active. All applications make use of the potential price distribution concept explained in Section 3.3, which relies on both the valuation and the transaction model.

7.1 Predicting Transaction Prices Out-of-sample

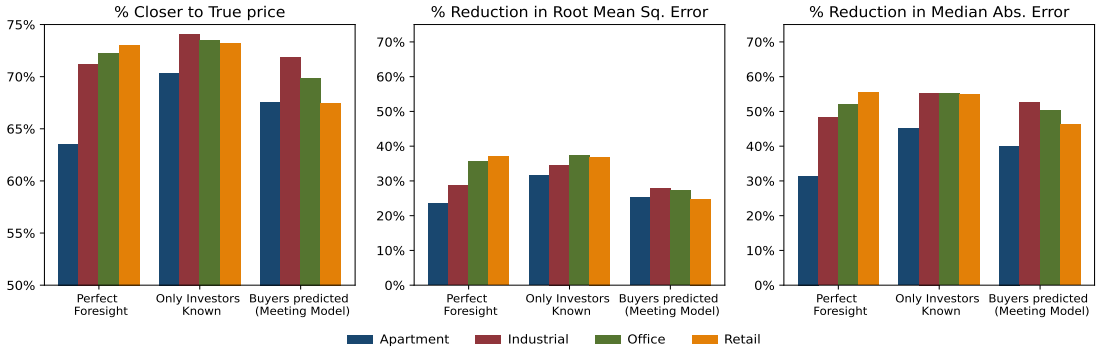
For an owner looking to sell their property in the next period, a key issue is to predict the transaction price and to assess the price risk. Our CRE model is well positioned to answer both questions. The two main insights from the valuation model are that building, market, and investor characteristics all matter for observed transaction prices, and that it is important for accurate valuations to capture the nonlinear and interaction effects of those variables on prices. We now explore the model's ability to predict transaction prices

out-of-sample for buildings that we know transacted.

Standing at time t, predicting prices p_{t+1} for buildings that trade at time $t+1$ in the data amounts to (i) estimating the model based on data available until time t, and (ii) predicting the inputs in the valuation model $(x_{n,t+1}, z_{t+1})$. To understand which inputs are most important, we proceed in three steps.

In the first step, we assume that the building, market, and investor characteristics can be predicted with perfect accuracy. We estimate the valuation and transaction model parameters using information up to time t and predict the prices in $t+1$ using the observed inputs (x_{t+1}, z_{t+1}). Figure 11 reports gains in OOS predictability of our baseline LGBM model relative to the linear hedonic model (LHM) reported in Table B.1, a natural benchmark in CRE valuation. The left panel shows the fraction of transactions for which the LGBM is closer to the true transaction price than the LHM. The middle panel shows the percent gains in terms of root mean squared error (RMSE), and the right panel reports the gains in median absolute pricing error (MAE). The first set of bars in each plot (labeled "perfect foresight") refer to the case where inputs $(x_{n,t+1}, z_{t+1})$ are assumed to be known. The figure shows that the LGBM strongly outperforms the LHM out-of-sample. The LGBM model's forecast is closer to the true observed price than the LHM in about 60-75% of transactions across sectors. In terms of the magnitude of the pricing error, the LGBM model's predicted price is typically 20-40% closer to the true price than the LHM's prediction (RMSE), while MAE gains indicate a 30-55% error reduction rate across sectors.

Figure 11: Comparing Out of Sample performance



Notes: The left panel reports the fraction of transactions for which the LGBM out-of-sample prediction is closer to the true transaction price than the LHM out-of-sample prediction. The middle panel reports the percent reduction in root mean squared prediction error of the LGBM relative to the LHM model. The right panel reports the percent reduction in median absolute prediction error of the LGBM relative to the LHM model. While the prediction model is always estimated with information up until time t in order to predict the price at $t+1$, the rows differ in how we determine the covariates (x_{t+1}, z_{t+1}) that enter in the prediction of the time-$t+1$ price. The first set of bars in each plot assumes that (x_{t+1}, z_{t+1}) are known. The middle set of bars assumes that z_{t+1} is known and that x_{t+1} must be predicted using a random walk. The right groups of bars assumes that (x_{t+1}, z_{t+1}) must both be predicted, and we use our matching model with characteristics-based based consideration sets to identify the potential buyers.

In the second set of bars in Figure 11, we assume that the time-$(t+1)$ building and market characteristics and time fixed effects are unknown. We use their time-t values as estimates; i.e., we assume that the macro variables and building occupancy rate follow a random walk: $\mathbb{E}_t[x_{n,t+1}] = x_{n,t}$.[29] We continue to assume that the buyer characteristics z_{t+1} are known. The LGBM model continues to produce predictions that are closer to the true transaction price than the LHM's predictions in about 70% of transactions. The gains in RMSE and MAE are larger for Apartments and Industrial and similar for Office and Retail when hedonics have to be predicted than when they are known.

In the third set of bars, we assume that all of the building and investor characteristics need to be predicted. For the building characteristics, we continue to use the time-t values. For the buyer information, we use our matching model with characteristics-based consideration sets. Specifically, for each property that transacts, the forecaster who stands at time t knows the identity of the seller, but not that of the buyer who materializes at time $t+1$, $z_{b,t+1}$. We draw $C = 1,000$ potential buyers in proportion with their likelihood according to our meeting model, estimated only on data available as of time t. We also do not know the ex-post characteristics of the potential buyers and so we use their time-t characteristics: $\mathbb{E}_t[z_{b,t+1}] = z_{b,t}$. For each potential buyer, we compute the fundamental valuation h_b plus a shock ϵ_b, and check whether the buyer's value exceeds that of the seller ($v_b > v_s$). If it is, the price is then the average of the buyer and seller valuation. We average all such prices across the potential buyers. This is similar to the process we use to obtain the potential price distribution (see Section 3.3). Here, we obtain potential prices for a given building n and compute statistics across the C potential buyer draws. The implementation details are in Appendix A.4.

The third row of Figure 11 shows that the LGBM generates better out-of-sample predictions than the LHM in about 65% of transactions. The average reduction in pricing errors relative to the linear hedonic model is about 25% in terms of RMSE and around 40% in terms of MAE. Figures E.1 and E.2 decompose the total out-of-sample performance gains into three components: (i) the improvement from moving from the LHM to the non-linear LGBM without investor characteristics, (ii) the gain from adding investor characteristics to the LGBM valuation model but with transactions that are based on random matching, and (iii) the gain from using our matching model with characteristics-based consideration sets for predicting buyers compared to random matching. About 50% of the gains in RMSE come from the first step, 20% from the second step, and 30% from using our matching model.

Even when investor identities and characteristics must be predicted, our joint valuation

[29] Of course, building subtype, CBD indicator, market location indicator all remain unchanged from one year to the next. The building's age and renovation-adjusted age grow by one year between t and $t+1$.

and transaction model delivers significant improvements in pricing accuracy fully out-of-sample. Figure E.3 shows that the gains in prediction power are consistent across years and sectors. These robust results should lend credibility to the counterfactual exercises that follow.

7.2 Understanding CRE price dynamics

We now explore the first counterfactual application of the CRE model in which we quantify the role of heterogeneous valuations. Figure 5 showed that a large number of offices and apartments changed hands from REITS to REPE funds and Institutions in the lead-up to the Great Financial Crisis. As our model indicates, some of these transactions may have reflected differences in beliefs between these types of investors about expected returns or differences in risk tolerance. Some may have arisen because REPE funds had strong fundraising years in 2005–2007 and may have been compelled to execute deals.[30] To understand the impact of this heterogeneity across investors on Office and Apartment prices in 2006 and 2007, we explore what would have happened had REPE and Institutional investors not been among the potential buyers for these REIT-owned properties.

We follow the procedure outlined in Section 3.3 to compute potential price distributions. For every office and apartment building owned by a REIT in 2006 and 2007, we simulate $C = 1,000$ potential buyers from the model's meeting probability distribution.[31] If the buyer's valuation exceeds the seller's, the building transacts, and we record the price. We compute the fraction of assets that trade (τ) and the size-weighted average price among the assets that transact (ppsf for offices and ppu for apartments).

Table 4 presents the results. In the first row of each panel, we show the fraction of REIT-owned properties that trade and the average price per unit (square feet for office, units for apartments) conditional on trade in the data, as well as the empirical composition of buyers for the properties that were sold by REITS. In the second row, we show the same statistics from a model simulation where all investors are present in the potential buyer pool. This baseline model predicts similar transaction probabilities to those in the data for office and somewhat higher transaction probabilities for apartments. Prices conditional on sale are similar in model and data for 2006 offices and apartments, and lower in the model than in the data in 2007.[32] This benchmark simulation sets a baseline for studying

[30] In typical REPE fee structures, fund managers only earn asset management fees on funds actually invested after an initial period.

[31] We define the universe as the properties owned by REITs rather than those sold by REITS to avoid selection bias in which assets were transacted.

[32] One possibility for the discrepancy between model and data is that the selection in the model, in terms of the average idiosyncratic shocks $\epsilon_b + \epsilon_s$ of the properties that actually trade, is smaller than in the data. Another possibility is that the model under-prices the specific set of office and apartment assets held by REITs in 2007.

the impact of replacing specific groups of buyers.

Table 4: CRE Price Dynamics in the REIT sell-off of 2006-2007

2006 Offices (3,889 Buildings)

Model	τ (%)	ppsf	Major Buyers (% Value of properties bought)
Truth	22.0	261.7	[REPE: 31.5, Institutional: 22.6, Foreign: 15.0, REITS: 14.1]
Benchmark (No Exclusions)	18.2	257.7	[Institutional: 29.0, REITS: 24.0, REPE: 18.4, OOD_N: 7.6]
Exclude REPE and Institutional	13.3	230.9	[REITS: 39.9, OOD_N: 23.8, Foreign: 10.9, OOD_L: 10.3]
Exclude REPE, Institutional and REITS	9.3	218.6	[OOD_N: 41.1, OOD_L: 18.0, User: 14.4, Foreign: 14.2]

2007 Offices (3,997 Buildings)

Model	τ (%)	ppsf	Major Buyers (% Value of properties bought)
Truth	28.7	382.5	[REPE: 60.9, Institutional: 14.9, OOD_N: 10.5, REITS: 10.3]
Benchmark (No Exclusions)	18.6	327.4	[Institutional: 27.8, OOD_N: 23.6, REPE: 18.2, REITS: 10.7]
Exclude REPE and Institutional	12.6	269.2	[REITS: 33.3, OOD_N: 31.4, Foreign: 11.3, OOD_L: 8.7]
Exclude REPE, Institutional and REITS	9.3	249.5	[OOD_N: 47.5, Foreign: 14.9, OOD_L: 13.4, Individual: 13.0]

2006 Apartments (3,560 Buildings)

Model	τ (%)	ppu	Major Buyers (% Value of properties bought)
Truth	16.4	103,386	[OOD_N: 39.2, Institutional: 32.6, REITS: 8.5, Foreign: 8.3]
Benchmark (No Exclusions)	27.6	110,013	[REITS: 36.3, OOD_N: 30.1, Institutional: 20.7, REPE: 6.1]
Exclude REPE and Institutional	24.3	113,358	[REITS: 49.6, OOD_N: 41.2, OOD_L: 5.2]
Exclude REPE, Institutional and REITS	17.9	100,374	[OOD_N: 82.1, OOD_L: 10.1]

2007 Apartments (3,375 Buildings)

Model	τ (%)	ppu	Major Buyers (% Value of properties bought)
Truth	18.8	200,547	[REPE: 38.4, Institutional: 38.3, OOD_N: 15.0, Foreign: 5.1]
Benchmark (No Exclusions)	29.5	164,113	[REITS: 34.7, OOD_N: 26.9, Institutional: 25.4, REPE: 6.2]
Exclude REPE and Institutional	24.7	126,142	[REITS: 53.4, OOD_N: 37.0, OOD_L: 5.1]
Exclude REPE, Institutional and REITS	17.1	109,922	[OOD_N: 80.7, OOD_L: 9.5, Foreign: 6.3]

Notes: τ is average fraction of assets that trade in every draw as defined in 3.3; ppu/ppsf shows the price per unit for Apartments and price per square foot for Offices, conditional on a sale. The first row in each panel refers to the observed RCA data.

The third row shows transaction probabilities and average prices conditional on a transaction when we exclude REPE and Institutional investors from the pool of potential buyers. We find that there would have been 27% (31%) fewer office and 12% (9%) fewer apartment transaction in 2006 (2007) without REPE and Institutional buyers. Average sales prices for offices sold by REITs in 2006 (2007) would have been 10.4% (17.8%) lower, while apartment properties would have sold for 2.7% (23.2%) less in 2006 (2007). The reason for these price drops is that the model identifies REPE and Institutional investors as having particularly high valuations for these properties relative to other investors in

those years (high $h_{b,t}$). Specifically, the last column shows that national owner-operator-developers (OOD_N) assume a much higher share of transactions in the counterfactual. These OOD_N investors can be considered to be the "value investors" in this market, looking for attractively priced assets. There are many of them; Table 1 showed that OOD_N is the largest investor category by dollar volume, accounting for 30% of buys and sells in our data set.

We note that in this counterfactual, REITs also emerge as prime candidate buyers of the assets sold by (other) REITs. As an additional counterfactual, we ask what would have happened if all REITs had stayed away from the office and apartment properties considered in this counterfactual. The fourth row in each panel of Table 4 shows that REITs would have sold a lot fewer assets than in the benchmark (between 35-49% drop in transactions across the four groups of assets) and that transaction prices would have fallen significantly more. The average price for these offices would have been 15.1% (23.8%) lower in 2006 (2007), while the average price for apartments would have been 11.5% (32.9%) lower in 2006 (2007). In this counterfactual, local owner-operator-developers and end-users emerge as alternative buyers of offices and some foreigners step in to buy the apartments. Local OOD and Users have even lower valuations for these properties, possibly because of a size, location, or quality mismatch with their existing holdings or because of tighter financial constraints. This counterfactual illustrates that the presence of REITs buyers lifts office values by 6% and apartment values by 13.5%. The model produces substantial heterogeneity in valuations, even within investor types such as REITs.

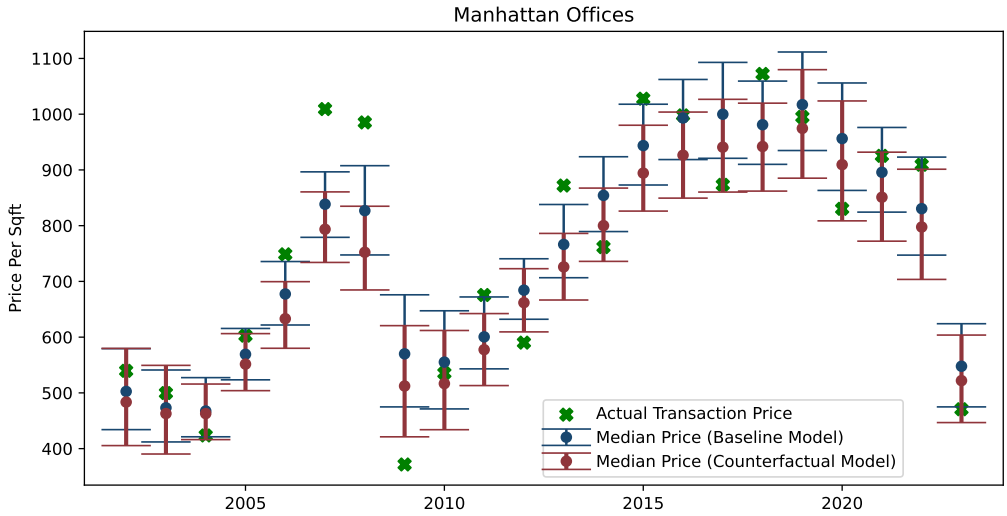
This exercise suggests that REPE funds and Institutions propped up prices just prior to the GFC. Their high sale prices, relative to the counterfactual valuation, may have blurred an early-warning signal of the impending financial crisis.

7.3 Substitution Patterns across Size, Sector, and Location

In our third application, we show how the model can be used to analyze substitution patterns in deal flows across building size, sector, and location. This is of interest to CRE investors and policymakers alike. Concretely, we focus on the role of foreign buyers who became large net buyers of U.S CRE in 2013–2018 and again in 2021, as shown in Figures 5 and A.5. Prominent foreign buyers, such as Middle Eastern sovereign wealth funds and large Canadian pension plans, had a strong preference for large, high-end properties in superstar cities. In addition to understanding the price impact on those types of office assets, we also use the model to study substitution patterns. If more foreign investors purchase large, high-end Manhattan office buildings, what is the impact on transaction volume and prices of smaller, lower-quality office buildings in Manhattan, on retail and apartment buildings in Manhattan, or on office buildings in other cities such as Chicago,

Los Angeles, and San Francisco?

Figure 12: Manhattan Office Market: Counterfactuals

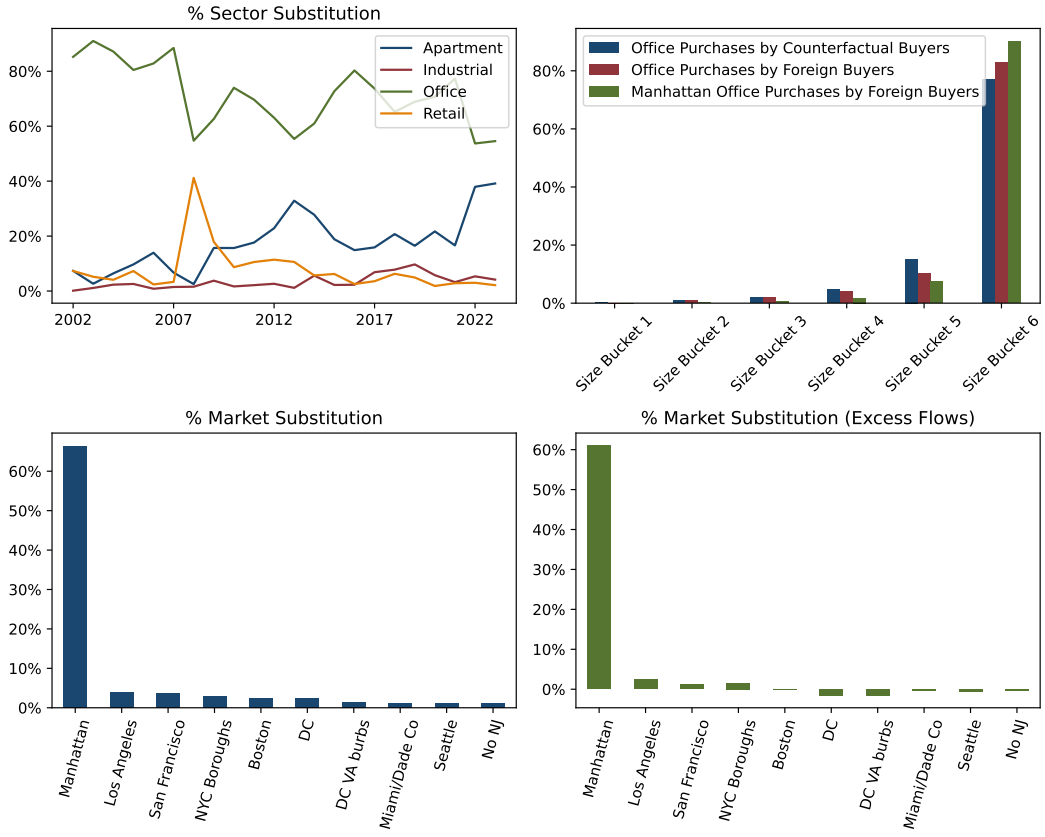


Notes: This figure shows the potential price distribution for the top-quartile of Manhattan office properties in the baseline model (in blue) and in a counterfactual model without foreign investors (in red) for each year between 2001 and 2023. The top (bottom) line segment indicates the 75th (25th) percentile of potential price distribution, and the dot, the median. The actual transaction price is given by the green cross. All prices are in 2023 dollars.

We start from the universe of high-end office buildings in Manhattan, defined as properties in the top quartile of the valuation distribution. Following the algorithm described in Section 3.3, we assess the impact of foreign buyers by comparing a benchmark model where all buyers are present to a counterfactual economy in which foreign buyers are replaced by buyers of all other types. We draw $C = 1,000$ potential buyers for each of the high-end Manhattan office properties in proportion to their model-implied likelihood of buying the specific asset under consideration.

The blue line segments in Figure 12 show the potential price distribution for top-quartile Manhattan office properties in each year, according to the baseline model that includes foreign buyers in the potential buyer pool. The top and bottom whiskers indicate the 75th and 25th percentiles and the blue dot the median of the potential price distribution. The red whiskers show the potential price distribution for the same assets in a counterfactual world where we replace all foreign buyers with buyers from all other types, sampled in proportion to their model-implied transaction probabilities. The assets that in reality were bought by foreigners are now bought by other types of investors. We see that the counterfactual price distribution (red) shifts down meaningfully from the benchmark one (blue). Removing foreign buyers lowers high-end Manhattan office prices by 4.3% over the full sample period 2001–2023, and by 5.5% in the 2013-2023 subsample.

Figure 13: Manhattan Office Market: Substitution Patterns



Notes: The top left panel plots the share of transaction volume invested by these potential buyers across the four sectors. The top right panel shows the property size distribution of office purchases across all markets by the foreigners (red bars) and by the counterfactual buyers (blue bars), and compares them to the size distribution of Foreign purchases of Manhattan office properties (green bars). The bottom left panel plots the top ten markets by average percentage of purchase volume by the counterfactual buyers. The bottom right panel subtracts from the bottom left panel the average percentage of purchase volume by all Manhattan office buyers.

If we define the high-end of the Manhattan office market as the top decile instead of the top quartile of properties, the valuation impact of foreign buyers grows to 6.3% for the full sample period and to 7.5% post-2013. Put differently, the foreign purchase activity meaningfully lifted the prices of Manhattan office assets, and especially at the top end of the market.

The counterfactual exercise gives us a list of potential buyers who would have bought the properties in question in the absence of foreign investors. With this list in hand, we explore whether those buyers bought another property in the focal year or in year thereafter. By analyzing what kinds of properties these plausibly displaced investors bought instead, we learn about the substitution patterns across property sizes, sectors, and locations.

The results are in Figure 13. The top left panel shows the sectoral composition of

purchase volume by the counterfactual buyers. The majority of their purchases are in the office sector, about 60% on average. There is substitution to apartments, which represent about 25% of purchases by the counterfactual buyers post-2013. The top right panel studies substitution across property sizes. For reference, the red (green) bars plot the size distribution of the actual foreign purchases of office properties in all markets (in Manhattan). The blue bars show the size distribution of actual office purchases made by the counterfactual buyers in those same years in all markets. We find that foreigners predominantly purchased very large offices and pushed other office buyers into buying somewhat smaller properties. The bottom left panel shows the locations of the purchases made by the counterfactual buyers across all sectors. There is limited substitution across space. Over 60% of actual purchase activity by this group took place in Manhattan. Foreign buyers may have crowded them out of specific Manhattan office assets, but it mostly did not crowd them out of Manhattan. To the extent that there is geographic substitution, it is fairly evenly directed towards other superstar markets such as San Francisco, Boston, Los Angeles, Chicago and DC. The bottom right panel of Figure 13 subtracts from the geographic portfolio shares of the counterfactual buyers, shown in the bottom left panel, the geographic portfolio shares of all investors who bought Manhattan offices in the same period. It shows that the counterfactual buyer pool is much more heavily concentrated in Manhattan, relative to the average Manhattan office buyer. Put simply, foreigners crowded out Manhattan office specialists.

7.4 Real Estate Private Equity Dry Powder

This section studies a final counterfactual, motivated by the large amount of capital raised by REPE funds in 2020–2021 and the subsequent CRE market slowdown in 2022–2024. This situation results in large amounts of "dry powder," waiting to be unleashed into the CRE markets.[33] We study under what conditions REPE funds can win sufficiently many trades to put a significant amount of dry powder to work.

Concretely, we target a combined CRE investment of $150 billion by REPE funds in 2025, three times the historical average annual REPE purchase level, shown in the left panel of Figure 15 for each sector. Because the transaction data end in 2023, we freeze the investor ecosystem, portfolios, and relationships, at their 2023 year-end values. We impose a 2025 macro scenario, by updating the local macroeconomic variables and the NER.

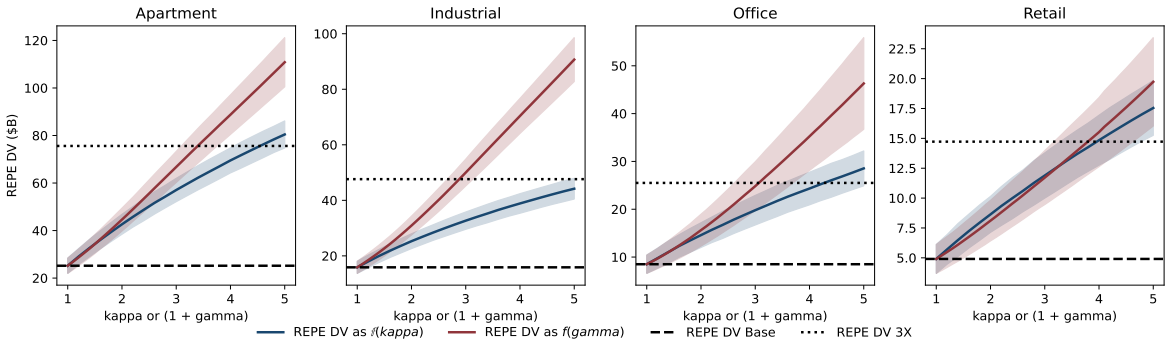
Seen through the lens of the model, the increase in REPE purchase volume must arise from a higher meeting intensity and/or a higher valuation for properties. We set

[33] As noted above, REPE funds lose asset management fees on capital they raise but do not deploy (typically) within 3 years.

$\pi_m^{cf}(b,s) \propto \pi_m(b,s) \times \kappa$ for b $\in$ REPE, with renormalization across all investor types, where $\pi_m(b,s)$ is the baseline model's meeting probability between seller s and buyer b. A value of $\kappa > 1$ increases the likelihood that a REPE fund is the matched counter-party, holding fixed property and other investor characteristics. For valuations, we shift the deterministic component of REPE buyers' private valuations $h_b^{cf} = h_b + \gamma\sigma_b$ for $b \in REPE$, while holding fixed all non-REPE buyers' valuations. A value of $\gamma > 0$ raises REPE bids by $\gamma\sigma_b$ making it more likely that a REPE buyer's valuation exceeds the seller's valuation and increasing the average price of transactions involving REPE.

Figure 14 varies either the search intensity κ (blue line) or the valuation parameter γ (red line) and traces out the model-implied REPE purchase volume, holding the other parameter at its baseline value. The dashed horizontal line is the baseline ($\kappa = 1, \gamma = 0$) where the horizontal dotted line is the target where REPE funds buy three times the baseline volume. The figure shows that a value of $\gamma = 2$ delivers the target purchase volume. This means that REPE valuations experience a latent demand shock that is two standard deviations above normal. Alternatively, REPE funds must have a matching probability that is about four times higher than in the baseline ($\kappa = 4$) to generate the targeted buying activity. The concavity of the blue lines show that the matching-intensity channel has diminishing returns because higher κ primarily reallocates meeting probability mass toward REPE within a fixed opportunity set.

Figure 14: Model-implied REPE purchase dollar volume as a function of κ and γ



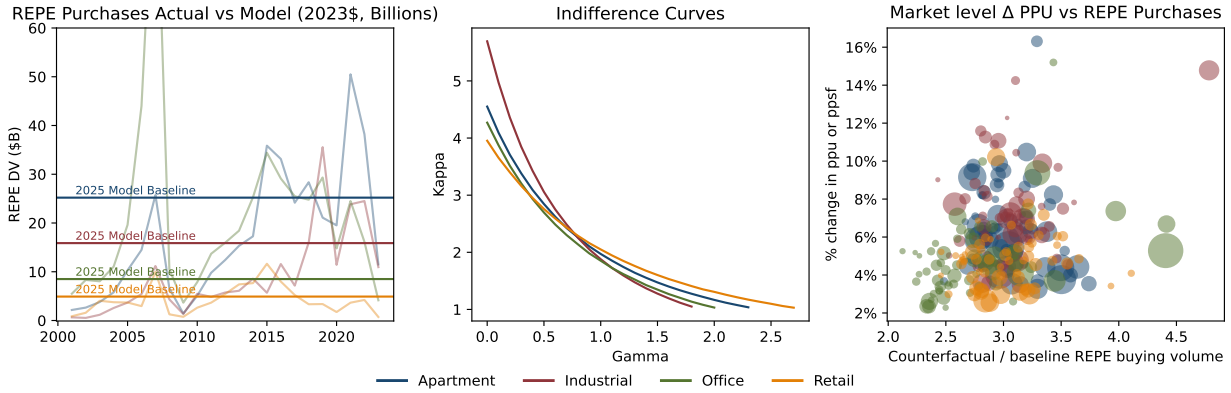
Notes: The figure shows the model-implied REPE purchase dollar volume (DV) as a function of the search-intensity multiplier κ or the shock multiplier γ. The blue curve plots the DV when κ is varied from 1 to 5, holding $\gamma = 0$. The red curve plots DV when γ is varied from 0 to 4, holding $\kappa = 1$. The dashed horizontal line denotes the baseline model ($\kappa = 1, \gamma = 0$), and the dotted horizontal line denotes the 3 times that number, the dry powder deployment target. Shaded regions report dispersion across Monte Carlo draws.

The middle panel of Figure 15 summarizes combinations of (γ, κ) that attain the 3× deployment target in each sector. The curves are downward sloping and convex, implying that higher REPE valuation shocks substitute for higher matching intensity, but with diminishing marginal substitutability. For example, a combination of twice higher meeting intensities ($\kappa = 2$) and a one-standard-deviation valuation increase ($\gamma = 1$) is sufficient to

meet the capital deployment target. So is a combination of $\kappa = 1.25 - 1.5$ and $\gamma = 1.5$. These calculations reflect the macroeconomic state of local CRE markets, the existing asset stock, and the composition of the buyer pool with whom the REPE funds compete.

What are the implications of the REPE buying spree on CRE markets? The right panel of Figure 15 relates the increase in REPE purchase volume in the counterfactual relative to the baseline to percentage changes in CRE transaction prices.[34] The size of the circle denotes market size and the color the sector. Prices increase between 2% and 16%, highlighting heterogeneity in market absorption capacity.

Figure 15: REPE Deployment: historical purchases and (κ, γ) counterfactuals



Notes: The left panel shows REPE purchase volume in the data by sector, together with horizontal lines showing the baseline ($\kappa = 1, \gamma = 0$) model-implied REPE purchase in the 2025 counterfactual environment. The middle panel shows the indifference curves in (γ, κ)-space that achieve the 3× deployment target, plotted separately by sector, where κ tilts meeting probabilities toward REPE and γ multiplies the idiosyncratic shock standard deviation σ_b for REPE potential buyers. The right panel shows a market-sector scatter plot of counterfactual excess REPE buying intensity (x-axis), measured as the ratio of counterfactual REPE buying to REPE buying in that same market in the baseline model, versus the % change in average PPU/PPSF (y-axis), where the counterfactual uses $\gamma = 1$ and κ set to the sector indifference-curve value, marker color denotes sector, and bubble size is proportional to market size.

8 Conclusion

We provide a novel approach to determine the joint evolution of prices and transactions in decentralized markets for private and real assets. We apply our approach to a comprehensive data set of commercial real estate (CRE) transactions that contain detailed information on property and investor characteristics. We find that using the investor composition in the CRE ecosystem is critical to understanding asset prices. The meeting model demonstrates that buyer-seller interactions are shaped by factors such as investor size, portfolio similarity between buyers and seller, and asset feature alignment in terms of property size,

[34] The baseline model has $(\kappa, \gamma) = (1, 0)$ while the counterfactual sets $\gamma = 1$ and κ equal to the sector-specific indifference-curve value that hits the deployment target, as indicated by the middle panel.

location, sector, and quality. These insights underscore the role of matching frictions and heterogeneous preferences in shaping transaction dynamics. The counterfactual analyses highlight the important role of investor composition in determining price outcomes. For example, foreign buyers increased Manhattan office prices by more than 5% in the post-2013 period. REPE fund purchases of large REIT sales in the office and apartment sectors in 2006 and 2007 increased property valuations by more than 10%, emphasizing how specific investor groups influence market outcomes.

Our paper offers a new methodology for understanding price formation and transaction dynamics in other private asset classes, such as private equity, private credit, infrastructure, and residential real estate, where buyer-seller interactions are similarly crucial. Furthermore, our new estimation technique for matching functions, which leverages insights from the computer-science literature, should have broad applicability to other search and matching settings.

Bibliography

Andersen, Steffen, Badarinza, Cristian, Liu, Lu, Marx, Julie and Ramadorai, Tarun. (2022). 'Reference Dependence in the Housing Market', *American Economic Review* 112(10), 3398–3440.

Ang, Andrew, Chen, Bingxu, Goetzmann, William N. and Phalippou, Ludovic. (2018). 'Estimating Private Equity Returns from Limited Partner Cash Flows', *Journal of Finance* 73(4), 1751–1783.

Arora, Abhishek and Dell, Melissa. (2023). 'LinkTransformer: A Unified Package for Record Linkage with Transformer Language Models'.

Badarinza, Cristian, Balasubramaniam, Vimal and Ramadorai, Tarun. (2024), In Search of the Matching Function in the Housing Market, Technical report, SSRN.

Badarinza, Cristian and Ramadorai, Tarun. (2018). 'Home Away from Home? Foreign Demand and London House Prices', *Journal of Financial Economics* 130(3), 532–555.

Badrinza, Cristian, Ramadorai, Tarun and Shimizu, Chihiro. (2022). 'Gravity, Counterparties and Foreign Investment', *Journal of Financial Economics* 145(2), 132–152.

Cvijanović, Dragana, Milcheva, Stanimira and van de Minne, Alex. (2022). 'Preferences of Institutional Investors in Commercial Real Estate', *The Journal of Real Estate Finance and Economics* 65, 321–359.

Damen, Sven, Korevaar, Matthijs and Van Nieuwerburgh, Stijn. (2025), An Alpha in Affordable Housing?, Working Paper 33470, National Bureau of Economic Research, Cambridge, MA.

Driessen, Joost, Lin, Tse-Chun and Phalippou, Ludovic. (2012). 'A new method to estimate risk and return of nontraded assets from cash flows: The case of private equity funds', *Journal of Financial and Quantitative Analysis* 47(3), 511–535.

Fox, Jeremy T. (2007). 'Semiparametric estimation of multinomial discrete-choice models using a subset of choices', *RAND Journal of Economics* 38(4), 1002–1019.

Fox, Jeremy T. (2018). 'Estimating Matching Games with Transfers', *Quantitative Economics* 9(1), 1–38.

Francke, Marc and van de Minne, Alex. (2024). 'Combining Machine Learning and Econometrics: Application to Commercial Real Estate Prices', *Real Estate Economics* 52(5), 1308–1339.

Fuster, Andreas, Goldsmith-Pinkham, Paul, Ramadorai, Tarun and Walther, Ansgar. (2022). 'Predictably Unequal? The Effects of Machine Learning on Credit Markets', *Journal of Finance* 77(1), 5–47.

Gabaix, Xavier, Koijen, Ralph S.J., Richmond, Robert and Yogo, Motohiro. (2024), Asset Embeddings, Technical report, SSRN Working Paper Series.

Genesove, David and Han, Lu. (2012). 'Search and matching in the housing market', *Journal of Urban Economics* 72, 31–45.

Genesove, David and Mayer, Christopher. (2001). 'Loss Aversion and Seller Behavior: Evidence from the Housing Market', *Quarterly Journal of Economics* 116, 1233–1260.

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

Giacoletti, Marco. (2021). 'Idiosyncratic Risk in Housing Markets', *The Review of Financial Studies* 34(8), 3695–3741.

Goetzmann, William N, Spaenjers, Christophe and Van Nieuwerburgh, Stijn. (2021). 'Real and private-value assets', *The Review of Financial Studies* 34(8), 3497–3526.

Griliches, Zvi. (1971), Introduction: Hedonic Price Indexes Revisited, *in* **Zvi Griliches**., ed., 'Price Indexes and Quality Change: Studies in New Methods of Measurement', Harvard University Press, Cambridge, MA, pp. 3–15.

Gupta, Arpit, Mittal, Vrinda and Van Nieuwerburgh, Stijn. (2025). 'Work From Home and the Office Real Estate Apocalypse', *American economic Review* forthc.

Gupta, Arpit and Van Nieuwerburgh, Stijn. (2021). 'Valuing private equity investments strip by strip', *The Journal of Finance* 76(6), 3255–3307.

Gu, Shihao, Kelly, Bryan and Xiu, Dacheng. (2020). 'Empirical Asset Pricing via Machine Learning', *The Review of Financial Studies* 33(5), 2223–2273.

Han, Lu and Strange, William. (2015), The microstructure of housing markets: Search, bargaining, and brokerage, *in* 'Handbook of regional and urban economics', Elsevier.

Kaplan, Steven N. and Schoar, Antoinette. (2005). 'Private equity performance: Returns,

persistence, and capital flows', *Journal of Finance* 60(4), 1791–1823.

Ke, Guolin, Meng, Qi, Finley, Thomas, Wang, Taifeng, Chen, Wei, Ma, Weidong and Liu, Qiwei Yeand Tie-Yan. (2017). 'LightGBM: A Highly Efficient Gradient Boosting Decision Tree', *Advances in Neural Information Processing Systems* 30, 3149–3157.

Kim, Min. (2024), Comparative R-Squared. SSRN Working Paper No. 3790200.

Koijen, Ralph S.J., Richmond, Robert J. and Yogo, Motohiro. (2024). 'Which Investors Matter for Equity Valuations and Expected Returns', *Review of Economic Studies* 91(4), 2387–2424.

Koijen, Ralph S.J. and Yogo, Motohiro. (2019). 'A Demand System Approach to Asset Pricing', *Journal of Political Economy* 127(4), 1475–1515.

Korteweg, Arthur and Nagel, Stefan. (2016). 'Risk-Adjusting the Returns to Venture Capital', *Journal of Finance* 71(3), 1437–1470.

Korteweg, Arthur and Sørensen, Morten. (2010). 'Risk and Return Characteristics of Venture Capital-Backed Entrepreneurial Companies', *Review of Financial Studies* 23, 3738–3772.

Lancaster, Kevin. (1966). 'A New Approach to Consumer Theory', *Journal of Political Economy* 74, 132–157.

Lundberg, Scott M, Erion, Gabriel, Chen, Hugh, DeGrave, Avery, Prutkin, Jordan M, Nair, Bharat, Katz, Ronit, Himmelfarb, Jonathan, Bansal, Nisha and Lee, Su-In. (2020). 'From local explanations to global understanding with explainable AI for trees', *Nature Machine Intelligence* 2(1), 252–259.

Lundberg, Scott M and Lee, Su-In. (2017), A unified approach to interpreting model predictions, *in* 'Advances in Neural Information Processing Systems', Vol. 30.

Ma, Zhuang and Collins, Michael. (2018). 'Noise Contrastive Estimation and Negative Sampling for Conditional Models: Consistency and Statistical Efficiency', *CoRR* abs/1809.01812.

McFadden, Daniel L. (1977), Modelling the Choice of Residential Location, Discussion Paper 477, Cowles Foundation for Research in Economics, Yale University, New Haven, CT.

Mikolov, Tomás, Sutskever, Ilya, Chen, Kai, Corrado, Greg and Dean, Jeffrey. (2013). 'Distributed Representations of Words and Phrases and their Compositionality', *CoRR* abs/1310.4546.

Mittal, Vrinda. (2024), Desperate Capital Breeds Productivity Loss: Evidence from Public Pension Investments in Private Equity. SSRN Working Paper No. 4283853.

Peng, Liang. (2016). 'The risk and return of commercial real estate: A property level analysis', *Real Estate Economics* 44(3), 555–583.

Piazzesi, Monika and Schneider, Martin. (2009). 'Momentum Traders in the Housing

Market: Sruvey Evidence and a Search Model', *American Economic Review, Papers and Proceedings,* forthcoming .

Piazzesi, Monika, Schneider, Martin and Stroebel, Johannes. (2013), Segmented Housing Search. Working paper, Stanford University.

Plazzi, Alberto, Torous, Walter and Valkanov, Ross. (2010). 'Expected Returns and the Expected Growth in Rents of Commercial Real Estate', *Review of Financial Studies* 23, 3469–3519.

Plazzi, Alberto, Torous, Walter and Valkanov, Rossen. (2008). 'The Cross-Sectional Dispersion of Commercial Real Estate Returns and Rent Growth: Time Variation and Economic Fluctuations', *Real Estate Economics* 36(3), 403–439.

Rosen, Sherwin. (1974). 'Hedonic Prices and Implicit Markets: Production Differentiation in Pure Competition', *Journal of Political Economy* 82, 34–55.

Sagi, Jacob S. (2021). 'Asset-Level Risk and Return in Real Estate Investments', *Review of Financial Studies* 34(8), 3647–3686.

Takuya Akiba, Shotaro Sano, Toshihiko Yanase Takeru Ohta and Koyam, Masanori. (2019). 'Optuna: A Next-generation Hyperparameter Optimization Framework', *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery Data Mining* pp. 2623–2631.

Van Nieuwerburgh, Stijn. (2019). 'Why Are REITs Currently So Expensive?', *Real Estate Economics* 47(1), 18–65.

Van Nieuwerburgh, Stijn, Stanton, Richard and de Bever, Leo. (2015), 'A Review of Real Estate and Infrastructure Investments by the Norwegian Government Pension Fund Global (GPFG)'.

Wallace, Nancy E. (1996). 'Hedonic-based price indexes for housing: Theory, estimation, and index construction', *Economic Review* pp. 34–48.

Wheaton, William. (1990). 'Vacancy, Search, and Prices in a Housing Market Matching Model'. Journal of Political Economy.

Witte, Ann D., Sumka, Howard J. and Erekson, Homer. (1979). 'An Estimate of a Structural Hedonic Price Model of the Housing Market: An Application of Rosen's Theory of Implicit Markets', *Econometrica* 47(5), 1151–1173.

Online Appendix

A Data Appendix

A.1 Data Cleaning Process

This appendix describes the process used to clean and filter the raw transaction data from RCA. The dataset includes transaction-level data for four core sectors—apartments, industrial, office, and retail—from 2001 to 2023.

Data Filters Using the transaction type variable, we retain only "Sale," "Construction," "Entity Level," and "Transfer" transactions, excluding land sales, air rights, and leases. Additionally, we drop incomplete transactions (e.g., contracts not completed) and those transactions valued below $500,000. These filters eliminate approximately 5% of all transactions. We also exclude a handful of transactions with implausibly high prices, defined as price per square foot greater than $10,000 for non-apartment transactions.

Property Characteristics Next, we address missing property characteristics. Property age is measured as the years since construction. Renovation-adjusted age is years since the last renovation and uses the most recent renovation year when available. Properties explicitly flagged as vacant in the notes section are assigned an occupancy rate of 0%. For the property's occupancy rate, about 30% of values are missing even after manually filling some cases using supplementary information from the RCA portal. We first forward-fill for up to one year using the property's latest observation on occupancy. Remaining gaps are imputed with a linear model based on market-level occupancy, renovation-adjusted age, neighborhood net effective rent, and the property's most recent occupancy when available. If no estimate can be produced, we use the sector–market average.

Investor Names For each transaction, RCA provides the names of up to four buyers and sellers. Fewer than 1% of transactions involve four names, suggesting that truncation at four buyers or sellers is not an important limitation. We consolidate investor names with inconsistent spellings or formats using the method in Arora and Dell (2023). Specifically, we use two independent models to generate the 20 closest matches per name, apply a 90% similarity threshold, manually discard false positives, and replace names with their most frequent version. This standardization allows us to better track investors over time. This helps in two key tasks: (i) identifying property owners for each year, critical for the seller role in our meeting model, and (ii) constructing investor-level portfolios to study asset evolution. The latter supports variables like investor transaction activity, size, and the consideration set characteristics such as the size, geographic, sectoral, and quality distribution of the asset portfolio.

RCA's investor names typically reflect the parent entity. If the data references a subsidiary, we replace it with its parent entity.

We assign unique investor IDs based on the cleaned investor names. When an investor name is missing, which happens in approximately 5% of cases, we assign a distinct ID to that record. These investor IDs form the basis for tracking investor portfolios.

Investor Types We map RCA's raw investor types (around 20 groups) into eight types, as shown in Table A.7, to better capture differences in investment objectives, size, and geographic origin. We manually correct certain investor attributes, such as incorrect country assignments or misclassified types.

Transactions Cleaning and Ownership Stakes Transactions suspected of not being at arms' length (e.g., multiple deals involving the same property within a 30-day span) are consolidated. We merge such transactions into a single record, using the price and characteristics of the final transaction and assigning buyers and sellers from the first and last transactions, respectively.

For properties that change sector during the sample period due to conversions, for example from office to apartments, we assign separate property IDs for periods before and after such conversions, to ensure we include the property in the right sector at each point in time.

Transactions where the buyer and seller are the same entity are removed. For cases where one party transacts on behalf of another (indicated by OBO), we retain the primary investor and drop the intermediary.

If there are multiple buyers or sellers, we allocate ownership shares equally unless the transaction history of the property indicates otherwise.

We track property ownership over time based on the observed transactions. Challenges such as missing or incorrect transactions are addressed by inferring missing links from prior and subsequent transactions. For example, if investor A sells to B, and C later sells the same property to D, we know that there is a missing transaction in the interim period where B sells to C. We adjust buyer and seller shares to ensure total ownership sums to 100% without negative values. Approximately 20% of properties require such adjustments during the sample period. These corrections are used solely for portfolio construction and are excluded from valuation and meeting model calibrations, which rely strictly on actual transaction data.

Before computing portfolio variables, we first establish the current owner of each property. We assume that the most recent known owner held the property continuously until another transaction is observed. When a transaction occurs, ownership is updated accordingly. This step is critical for identifying the seller in each transaction, as the seller is simply the current owner immediately prior to the sale.

Portfolio Holdings Portfolio variables are then computed sequentially from the transaction history. For each investor, we calculate:

1. Portfolio size: Total market value of properties owned, adjusted for fractional ownership. The investor size is in 2023 dollars and is obtained by using the annual property price panels described below.
2. Number of properties: Count of unique holdings, adjusted for fractional stakes.
3. Sector and geographic composition: share of portfolio value in each RCA sector and aggregated market (see Section A.2).
4. Average property value: The average market value of properties in the investor's portfolio.

Entity Sales For entity sales, where RCA equally distributes transaction value across properties, we adjust property-level prices by assigning per-square-foot values proportional to RCA's sector-market price indices. This provides a more accurate estimate than assuming uniform prices across locations.

RCA Price Indices We construct a panel of real property prices by market-sector by combining RCA's nominal price indices with the seasonally-adjusted CPI index (CPIAUCSL from FRED), expressing all indices in 2023 dollars. We use these property price estimates to calculate the investor portfolio sizes and other portfolio holding variables discussed above.

Market Occupancy and NOI Growth Market occupancy data is primarily sourced from NCREIF at the CBSA-sector level. For missing values, we calculate volume-weighted occupancy rates using RCA data, provided that at least 15 properties in the sector-year contain valid occupancy information. Net

Operating Income (NOI) data, also from NCREIF, is CPI-adjusted and annualized. NOI growth is calculated logarithmically if both the current and previous NOI values are positive; otherwise, we use a smoothed growth rate formula. Extreme outliers (top/bottom 1%) are replaced with median values.

Geography Geographic variables (e.g., zip codes, FIPS codes) are filled where possible. For markets with incomplete coverage (but at least 50% available), we impute missing values using the average from available observations within the same market. Gaps are more common in remote areas and for retail and industrial properties, which tend to be farther from city centers. In the rare cases where no market-level information exists, we use sector-level national averages.

Defining the Estimation Sample We next define the set of investors and properties used in our analyses. Properties with incomplete geographic information that cannot be recovered (e.g., missing RCA market or geo codes) are excluded. For each year, we include all properties that exist at that time: properties are added once constructed and reassigned if converted to another sector. Constructions are excluded from the main analysis. The number of properties rises from roughly 220,000 in 2001 to over 300,000 in 2023. As of 2023, the dataset contains around 87,000 apartments (65,000 garden and 22,000 mid/high-rise), 78,000 industrial (60,000 warehouses and 18,000 flex), 59,000 offices (50,000 suburban and 9,000 CBD), and 77,000 retail properties (43,000 shops and 34,000 centers). These counts vary over time, and about one-third of properties in the dataset transact more than once.

For investors, our base universe includes all entities that ever buy or sell a property in the dataset. We exclude investors that only ever transact in 'construction' properties.

When a group of investors forms a joint venture (JV) to transact, we retain the individual members while also assigning a unique JV identifier. Most investors, and particularly JVs, appear only once as JVs are typically created for a specific investment. Investors are removed from the sample once they cease to exist (as flagged by entity sales) or five years after they have fully liquidated their holdings. This limits noise from inactive entities. Symmetrically, investors begin to appear in the data five years before their first observed transaction, ensuring they are included as potential buyers. JVs are included from the year of their formation until their final joint transaction. This procedure yields an average of about 1,000 JVs per year (rising over time) and roughly 350,000 investors overall.

In the meeting model, potential buyers include all active investors and JVs. To keep the negative set tractable, we restrict it to active investors, defined as those that have transacted at least once in the previous five years. Some investors enter with no holdings in the transaction year. In such cases, we assign them the median number of properties (which is one) with characteristics equal to the average portfolio in that year. Robustness checks confirm that results are not sensitive to this treatment.

Next, we construct the buyer–seller relationship variable for the meeting model, defined as an indicator equal to one if the buyer and seller have had a prior transaction. On average, about 8% of buyers have such a relationship. For joint ventures, the indicator is set to one if any member of the buying group has previously transacted with the seller. Prior relationships are measured from the start of the sample up to the year prior to the transaction.

Since the dataset records only actual transactions, property and investor characteristics are directly observed only at the time of trade. For analyses that require the full stock of properties, such as listing probabilities and counterfactuals, we extend these characteristics to all properties and investors, including those not transacting in a given year. Static attributes such as location, market, and size are carried forward, while age is updated annually using construction and renovation years. Occupancy for non-transacting properties is imputed using the model described above. Transaction types (regular sale, transfer, entity sale) are assigned probabilistically to match the annual distribution of observed types, ensuring consistency with

market conditions. Investor types are assumed to be fixed over time. Portfolios are constructed as described above, and macroeconomic characteristics are observed directly for all years without adjustment.

Variable Transformations Finally, we transform variables to reduce skewness and improve comparability across sectors. Property size is measured in apartment units for multifamily properties and in thousands of square feet for office, retail, and industrial assets. We enter property size in logarithmic form, along with property age, renovation-adjusted age, number of floors, net effective rent (NER), market personal income, and market population.[35] Investor portfolio size for both buyers and sellers is logged. Transaction prices are expressed as the logarithm of price per unit (for apartments) or per square foot (for other sectors), ensuring comparability across sectors and mitigating the influence of extreme values.

A.2 Mapping Geographies

The RCA dataset provides geo-codes for every property, which we leverage to map properties to various geographic indicators. This is essential because the macroeconomic time series data we use comes from multiple sources, each often organized by different geographic units. Using the 3-digit rounded latitude and longitude data provided by RCA, we generate ZIP codes for each property. This is done using standard geolocation algorithms that match coordinates to corresponding ZIP codes.

We use the county-level Federal Information Processing Standards (FIPS) codes to map properties to Core-Based Statistical Areas (CBSAs). This mapping is based on a crosswalk file obtained from the National Bureau of Economic Research (NBER).

RCA defines 165 individual markets in its classification system. To align with other datasets and guarantee sufficient coverage in each market, we aggregate these into 60 broader geographic categories, which we refer to as our "markets." The aggregation is guided by logical groupings and shared economic characteristics. These markets serve as the level at which we map properties to RCA price indices, ensuring consistency when linking property-level information to market-level indicators.

For most large markets, we retain RCA's original definitions, while smaller markets are combined into broader groups. Table A.1 reports the mapping between RCA's markets and our aggregated markets.

A.3 Net Effective Rent

This appendix describes how we construct net effective rent statistics of the property or the neighborhood of each property that we see transact in the RCA data. Net effective rents are observed for office, retail, and industrial properties at the lease level in the Compstak database.

We first merge the RCA and the CompStak databases at the property level using the following procedure: for each RCA property, we consider CompStak buildings within a 1 km radius of its geo point. We assign each CompStak building a matching score based on the geographical distance, the difference between the years the properties were built, the difference in log building size, and the fuzzy matching score on property names. Specifically, the matching score equals the sum of: (1) $0.3 \times \max\{1 - \frac{\text{geo distance (km)}}{0.3}, 0\}$, (2) $0.3 \times \max\{1 - \frac{|\Delta \text{year built}|}{5}, 0\}$, (3) $0.3 \times \max\{1 - \frac{|\Delta \log \text{building size}|}{5\%}, -1\}$, and (4) $0.4 \times \max\{\frac{\text{name similarity} - 70}{30}, 0\}$, where name similarity is the highest partial Levenshtein similarity between elements in the RCA property name list and the CompStak name list (including all street addresses and building names of that property). Additionally, if the similarity score is exactly 100, we assign this part of the score as 0.4×2 instead of $0.4 \times$

[35] For Apartments, where we use NOI as discussed in Section A.3 below, we set the negative values which occur in a handful of cases to zero prior to transformation.

Table A.1: Mapping RCA Markets to Our Markets

Our Market	RCA Markets
Tertiary Mid-Atlantic	All Others - DE,MD,VA,WV; Ann Arbor, Albany, All Others - NJ, All Others - NY, All Others - PA, Allentown, Buffalo, Camden, NJ, Eastern PA, Harrisburg, Rochester, Syracuse, Trenton, NJ
Tertiary Midwest	Lexington, Louisville, Akron, All Others - IL, IN, OH, Dayton, Kankakee, IL, Lansing Toledo, All Others - Plains States, Des Moines, Omaha, Wichita, All Others - MI,WI, Flint, MI, Grand Rapids, Madison, Milwaukee, Monroe, Racine, All Others - MN
Tertiary Northeast	All Others - CT, All Others - New England, New Haven, Portland, ME, Providence, Southern NH, Worcester
Tertiary Southeast	Polk Co, All Others - FL, Brevard Co, Daytona Beach, Florida Panhandle, Fort Myers, Gainesville, Martin/Saint Lucie, Sarasota, Tallahassee, All Others - GA,NC,SC,TN,KY, Charleston, Chattanooga, Columbia, Greensboro, Greenville, Knoxville, Myrtle Beach, Savannah, Winston-Salem, All Others - AR,LA, All Others - MS,AL, Baton Rouge, Huntsville, AL, Jackson, MS, Little Rock, New Orleans
Tertiary Southwest	Boulder, Greeley, Albuquerque, All Others - Southwest, Tucson, All Others - OK,TX, Corpus Christi, El Paso, McAllen, Oklahoma City, Tulsa, Co Springs, Fort Collins, Provo
Tertiary West	Santa Barbara, Ventura Co, All Others - Northwest, Spokane, Bakersfield All Others - Southern California, Reno, Fresno, Madera, CA, Modesto, Salinas, Vallejo-Fairfield, All Others - Northern California, Santa Rosa, Boise, Napa, All Others - ID,MT,WY; Hawaii, Honolulu

Same name as RCA Markets

Atlanta	Austin	Baltimore	Birmingham	Boston	Broward
Charlotte	Chicago	Cincinnati	Cleveland	Columbus	DC
DC MD burbs	DC VA burbs	Dallas	Denver	Detroit	East Bay
Hartford	Houston	Indianapolis	Inland Empire	Jacksonville	Kansas City
Las Vegas	Long Island	Los Angeles	Manhattan	Memphis	Miami/Dade Co
Minneapolis	NYC Boroughs	Nashville	No NJ	Orange Co	Orlando
Palm Beach Co	Philadelphia	Phoenix	Pittsburgh	Portland	Raleigh/Durham
Richmond/Norfolk	Sacramento	Salt Lake City	San Antonio	San Diego	San Francisco
San Jose	Seattle	St Louis	Stamford	Tampa	Westchester

Notes: This table provides the mapping we use from the raw RCA markets to the aggregated RCA markets we use where we performed the aggregation.

1. We keep the CompStak building with the highest matching score. If the maximum score ≤ 0.3, we deem that there is no match at the property level.

We manually check the matching result for the top 100 RCA office buildings transacted after 2015 in Manhattan: the correct matching rate is 92%, wrong matching rate is 3%, and the missing rate is 5%. Wrong matches occur when very similar and adjacent buildings have characteristics closer than the true match, while the missing matches occur because CompStak does not provide coverage for all RCA properties. The full-sample building-level matching rates are 55.1%, 39.3%, and 48.3% for office, retail, and industrial properties, respectively. Matching rates at the building-year level are naturally lower as we do not necessarily observe new leases in the matched buildings in every year.

For RCA properties that are not matched with a CompStak building, or matched with a CompStak building lacking sufficient recent lease data to compute measures, we create variables from the lease-level

data at a fine level of geography (call it a “block"), and match each property we see transact in RCA to its block-level variables.

The finest level of geography we consider uses the first 2 digits of the latitude and longitude. For example, for a property with latitude and longitude (40.754932,-73.984016), we define a 2d area as [40.75,40.76;-73.99,-73.98]. This corresponds to an area of roughly 0.6 miles by 0.6 miles (1km by 1km). We collect all the leases in this block to calculate the average rent measures.

Given that leasing activity is sparse, we use all leases signed in the past three years to calculate the property-year or block-year average rent. We weigh leases by their transaction sqft, and assign a lower weight for leases signed in the more distant past. In the baseline, we assign weights of 1, 1/2, and 1/3 to leases signed one, two, and three years before, and then renormalize to make the weights sum to one. As a robustness check, we also compute a variable using the past 5 years of leasing data, with weights on past leases of 1, 1/2, 1/3, 1/4, and 1/5, renormalized to sum to one.

When there are fewer than 5 buildings in the block, or when there are no leases in the past three/five years, we use the average rent calculated for a larger block. The next-larger block size we consider is defined based on the first decimal place of the latitude and longitude. For example, a property at latitude and longitude (40.754932,-73.984016) belongs to block [40.7,40.8;-74.0,-73.9]. The typical size of this 1-digit block is 6 miles by 6 miles (10km by 10km).

When the average rent is still missing in this 1-digit block, we fill it with the value in the block defined based on the unit digit of the latitude and longitude (0-digit block). In our example, this would be the area given by [40,41;-73,-74], an area of roughly 60 miles by 60 miles.

For the remaining missing values, we sequentially fill them with the average rent at the RCA market, average rent at our own market definition (which combines several RCA markets, see appendix A.2), and if necessary, average rent measured at the national level.

For apartment properties, we observe the net operating income per apartment unit at the loan level in the Fannie Mae’s multifamily mortgage origination data. We follow a similar approach to calculate the NOI (referred to as NER below for simplicity). Given that loan origination activity is also sparse, we use all originations in the past three years to calculate the average NER. We first use the same matching algorithm to conduct a property-level matching, obtaining a matching rate of 45.8% at the building level, and then use the block-year-level NER to fill in missing values. We start with the 2-digit blocks, use the same weighting scheme explained above, and fill missing values sequentially with values computed on increasingly coarse geographic levels (1-digit, 0-digit, RCA market, our market, national).

Table A.2 reports the proportion of the final NER observations for each sector (columns) that come from each level of geography (rows). Most of the NER are measured at the 1-decimal level (6 miles by 6 miles). For office NER of the last three years, we have 24.9% of observations measured at the even finer 2-decimal level (0.6 miles by 0.6 miles) and 20.8% at the property level.

To understand how well the block-level NER measures captures property-level variation in NER, we estimate a linear regression of property-level NER in the Compstak/Fannie Mae data on the block-level NER variable, where we use the most spatially granular variable available for that property. The R^2 of those property-level regressions is reported in the last row of Table A.2.

We have explored several more variables available in the Compstak data, such as occupancy, quality grade (class A, B, C), remaining lease duration, tenant size (number of employees of the tenant), tenant concentration (Herfindahl index computed from the leases in a building), but none of these variables add explanatory power to the valuation model, once the NER variable is included.

Table A.2: NER Information Source

Panel A: NER (Last 3 Years)

Level	Office Sector	Industrial Sector	Retail Sector	Apartment Sector
Property level	20.8%	11.7%	6.5%	8.4%
2-decimal lat/lon	24.9%	25.3%	15.4%	31.2%
1-decimal lat/lon	38.9%	40.2%	49.5%	50.1%
unit-digit lat/lon	10.8%	15.2%	21.3%	9.5%
RCA Market	2.5%	3.3%	4.4%	0.8%
Market	1.7%	2.9%	2.4%	0.0%
National	0.4%	1.5%	0.6%	0.0%
Regression R^2	69.4%	59.7%	44.3%	45.2%

Panel B: NER (Last 5 Years)

Level	Office Sector	Industrial Sector	Retail Sector	Apartment Sector
Property level	25.9%	15.9%	9.3%	12.5%
2-decimal lat/lon	23.2%	24.7%	15.9%	35.7%
1-decimal lat/lon	37.4%	38.8%	49.6%	44.3%
unit-digit lat/lon	9.7%	14.1%	19.0%	7.0%
RCA Market	2.2%	2.9%	4.0%	0.5%
Market	1.4%	2.5%	1.9%	0.0%
National	0.2%	1.0%	0.4%	0.0%
Regression R^2	66.9%	54.1%	43.2%	43.2%

A.4 Potential Price Distribution: Process

This appendix describes the practical steps for constructing the potential price distribution (PPD) described in Section 3.3. For the counterfactuals in Sections 7.2 and 7.3, we study how removing specific investor groups affects (i) the probability that a asset trades and (ii) the expected price conditional on trade. For the out-of-sample prediction in Section 7.1, we use the same machinery to forecast next-period transaction prices, conditional on transacting.

We consider the market defined by a calendar year t and a property sector. To reduce clutter, t subscripts are omitted in what follows. For each asset n with known seller s (current owner), we observe asset characteristics x_n, seller characteristics z_s, and the set of active buyers B (investors that transacted at least once in the preceding five years) with characteristics z_b. This allows us to estimate systematic valuations $h_s(n)$ and $h_b(n)$ from the valuation model, meeting probabilities for every potential buyer b': $\{\pi_m(b' \mid s, n)\}_{b' \in B}$, and a common idiosyncratic shock scale σ with $\epsilon_s, \epsilon_b \sim N(0, \sigma^2_{m,t})$ (see Section 5.3). We also have the listing probability $\pi_{l,t}(s)$, at asset-year level as defined in (5). We then follow the steps below:

1. For each asset n at time t, draw a listing indicator $l_{n,t} \sim \text{Bernoulli}(\pi_{l,t}(s))$ using the asset-year listing probability $\pi_{l,t}(s)$. We simulate only the asset with $l_{n,t} = 1$.
2. For each listed asset, draw a seller shock $\epsilon_s \sim N(0, \sigma^2_{m,t})$. Then set $v_s = h_s(n) + \epsilon_s$.
3. For each listed asset n, draw a candidate buyer from B with replacement using weights $\{\pi_m(b' \mid s, n)\}_{b' \in B}$.
4. For each potential buyer, draw a latent demand shock $\epsilon_{b'} \sim N(0, \sigma^2_{m,t})$ and set $v_{b'} = h_{b'}(n) + \epsilon_{b'}$.
5. If $v_{b'} > v_s$, record the price $\widehat{p}(n) = \frac{1}{2}\left(v_{b'} + v_s\right)$; otherwise set $\widehat{p}(n)$ to missing (no trade).
6. Compute:
   (i) the mean price conditional on trade, $Ep = \frac{1}{\sum \mathbb{1}_{\hat{p}(n)>0}} \sum_n \hat{p}(n) 1_{\hat{p}(n)>0}$

(ii) the fraction of assets that trade, $\tau = \frac{1}{N}\sum_n 1_{\hat{p}(n)>0}$.

7. Repeat steps 1–6 for $C = 1000$ replications, obtaining $\mathbb{E}[\widehat{p}]$ and τ for each run. The empirical distribution of these statistics is the market-level PPD. Asset-level PPDs are obtained by retaining outcomes for each n across replications.

8. To understand the importance of a subset G of investors (e.g., foreign buyers in Manhattan), compute the PPD after removing that group from the buyer pool: set $\pi_m(b' \mid s, n) = 0$ for $b' \in G$, renormalize over $B \setminus G$, and repeat steps 1–6. Compare the distributions of $\mathbb{E}[\widehat{p}]$ and τ to the baseline.

For market-level counterfactuals, we compare baseline and alternative buyer-composition outcomes by reporting, across replications, the distributions of the mean simulated price conditional on trade ($\mathbb{E}[\widehat{p}]$) and the fraction of assets that trade τ. We interpret the differences as effects on prices and liquidity.

For out-of-sample prediction (predicting prices in t+1 from the vantage point of t), we do not observe next year's asset characteristics $x_{n,t+1}$, investor characteristics z_{t+1}, or macroeconomic conditions. We therefore use their time-t values to construct a time-t prediction of the seller's next-period deterministic valuation $h'_{s,t+1}(n)$. The price-shock variance for t+1 is also unknown; we set $\sigma^2_{p,t+1}$ to the historical estimate $\sigma^2_{p|t} = \mathrm{Var}_t(\epsilon_p)$ computed from ϵ_p over 2001–t. Next, we use $\sigma^2 = 2\,\sigma^2_{p|t}$ since $\mathrm{Var}(\epsilon_p) = \frac{1}{2}\sigma^2$. We then get the asset-level PPD and set $\widehat{p}^{\,\mathrm{PPD}}_{t+1}(n)$ to the mean of non-missing simulated prices for asset n (conditional on trade). Forecast improvement over the linear hedonic model (LHM) is measured on the set $\mathcal{T}_{t+1}$ of assets that actually trade in t+1. For instance, the reduction in mean-squared-error is:

$$1 - \frac{\sum_{n\in\mathcal{T}_{t+1}} \left(\widehat{p}^{\,\mathrm{PPD}}_{t+1}(n) - p_{t+1}(n)\right)^2}{\sum_{n\in\mathcal{T}_{t+1}} \left(\widehat{p}^{\,\mathrm{LHM}}_{t+1}(n) - p_{t+1}(n)\right)^2},$$

where $\widehat{p}^{\,\mathrm{LHM}}_{t+1}(n)$ is the linear hedonic benchmark, and $p_{t+1}(n)$ is the realized price.

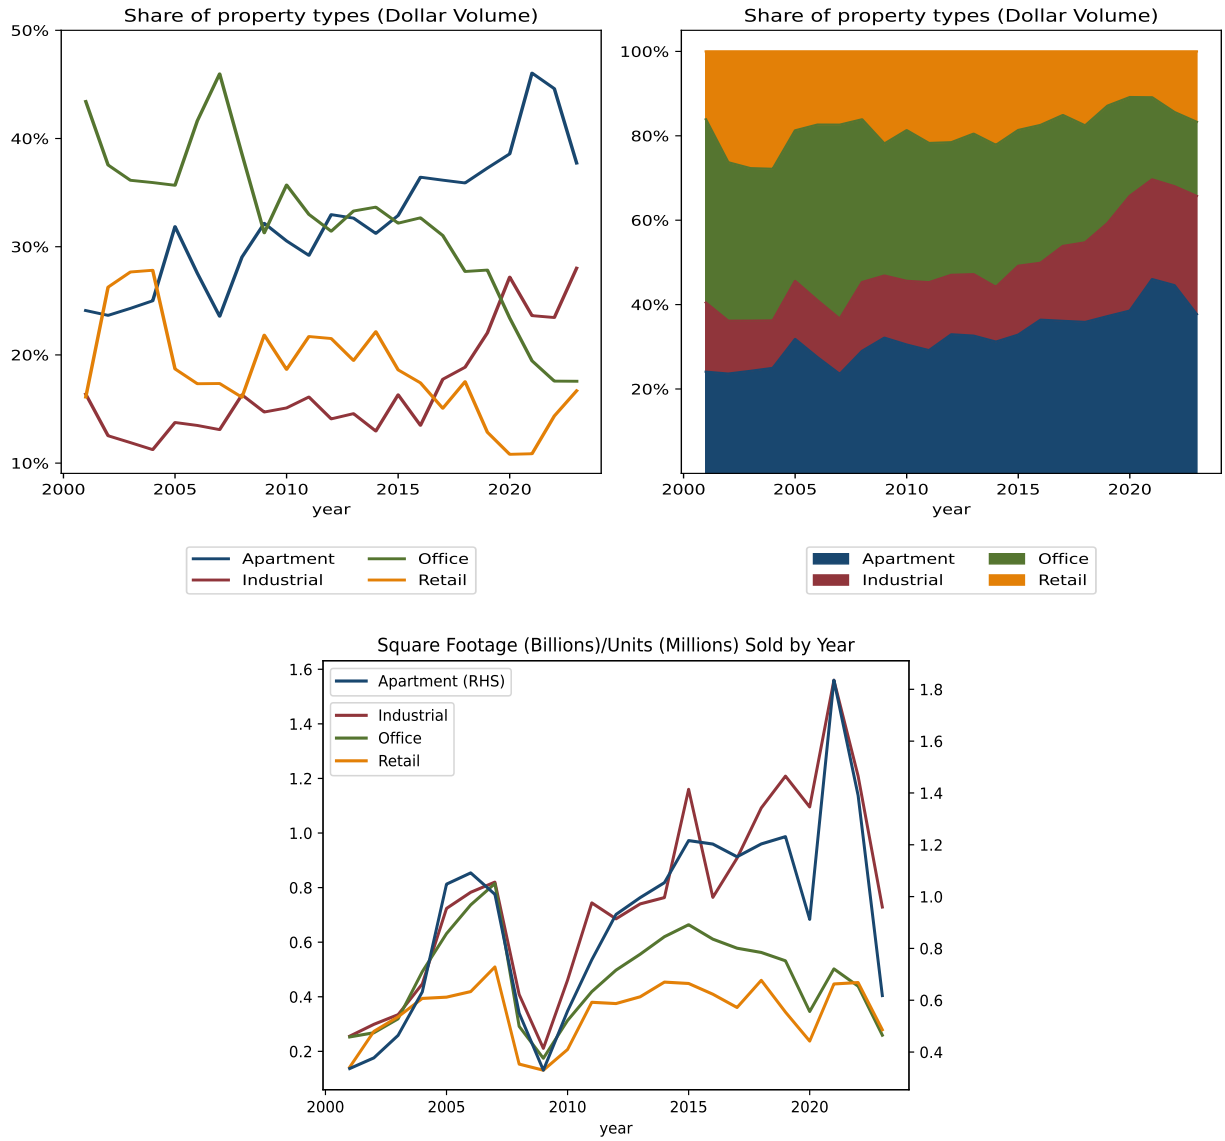
A.5 Transaction Summary Statistics

Table A.3: Transactions By Transaction Size

	# Trans	% Trans	Cum. % Trans	$ Vol	% Vol	Cum. % Vol
Above 1 Bil	269	0.06	0.06	327	3.24	3.24
500 Mil - 1 Bil	701	0.15	0.20	374	3.71	6.95
250-500 Mil	2,368	0.50	0.70	704	6.97	13.92
100-250 Mil	12,525	2.63	3.33	1,726	17.10	31.02
75-100 Mil	9,301	1.95	5.29	772	7.65	38.68
50-75 Mil	19,926	4.19	9.47	1,181	11.71	50.38
25-50 Mil	52,693	11.07	20.54	1,814	17.97	68.35
20-25 Mil	22,517	4.73	25.27	496	4.91	73.26
15-20 Mil	33,779	7.10	32.37	578	5.72	78.99
10-15 Mil	57,414	12.06	44.43	695	6.89	85.87
5-10 Mil	135,100	28.38	72.81	951	9.42	95.30
Below 5 Mil	129,429	27.19	100.00	474	4.70	100.00

Notes: This table shows the number of transactions, the % of transactions, Value and % of Value for transactions by various transaction size buckets. Dollar volume is defined in terms of Billions of 2023 US Dollars.

Figure A.1: Transaction Volume by Property Type



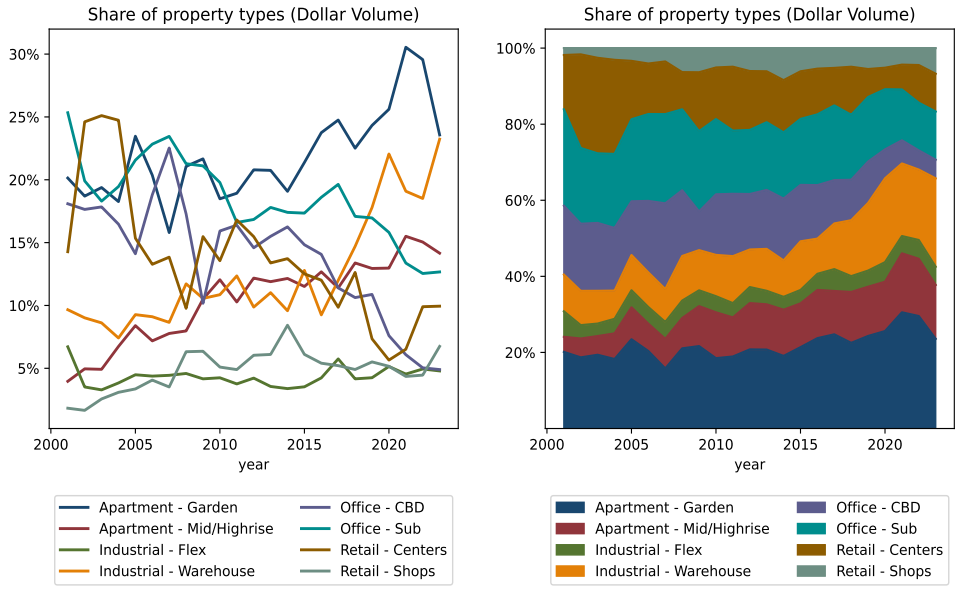
Notes: The top panel shows share of properties transacted by property types, i.e. Apartments, Office, Industrial and Retail, from 2001 to 2023. Value is defined in terms of 2023 USD. The bottom panel shows the total square feet transacted in billions of Industrial, Office and Retail space, and the number of Apartment units bought in millions.

Table A.4: Statistics by Transaction Year

Year	# Trans	% Trans	Cum. % Trans	Volume	% Vol	Cum. % Vol	% Foreclosures	% EntitySales
2001	5,272	1.11	1.11	150.71	1.49	1.49	0.14	6.82
2002	6,287	1.32	2.43	186.11	1.84	3.34	0.06	2.19
2003	7,977	1.68	4.10	223.00	2.21	5.55	0.21	2.96
2004	11,821	2.48	6.59	348.93	3.46	9.00	0.18	7.44
2005	20,725	4.35	10.94	490.05	4.86	13.86	0.15	3.14
2006	23,025	4.84	15.78	555.20	5.50	19.36	0.13	9.39
2007	23,385	4.91	20.69	675.31	6.69	26.05	0.34	21.01
2008	13,531	2.84	23.53	230.60	2.28	28.34	2.70	2.83
2009	6,884	1.45	24.98	109.09	1.08	29.42	17.25	0.38
2010	10,873	2.28	27.26	211.82	2.10	31.52	11.88	0.33
2011	14,749	3.10	30.36	304.41	3.02	34.53	6.64	8.58
2012	19,785	4.16	34.52	374.61	3.71	38.24	4.94	5.74
2013	21,200	4.45	38.97	434.17	4.30	42.54	2.30	5.71
2014	24,974	5.25	44.22	503.24	4.99	47.53	2.27	5.90
2015	29,252	6.15	50.36	616.83	6.11	53.64	0.48	10.14
2016	26,540	5.58	55.94	574.22	5.69	59.33	0.71	4.87
2017	26,996	5.67	61.61	542.27	5.37	64.70	0.73	4.33
2018	29,044	6.10	67.71	611.31	6.06	70.76	0.66	12.52
2019	30,679	6.44	74.16	627.41	6.22	76.98	0.36	1.26
2020	23,793	5.00	79.15	452.95	4.49	81.46	0.48	4.37
2021	40,410	8.49	87.64	854.42	8.47	89.93	0.55	4.27
2022	36,266	7.62	95.26	697.79	6.91	96.84	0.86	8.10
2023	22,554	4.74	100.00	318.54	3.16	100.00	1.68	2.95

Notes: This table shows the number of transactions, the share and cumulative share of total transactions, the volume (dollar value), share and cumulative share of volume by transaction year. The last two columns displays the share of transactions that are foreclosures and the share of transactions that are entity sales. Volume is defined in billions of 2023 U.S. dollars.

Figure A.2: Transaction Volume by Property Subtype



Notes: Figure shows share of properties transacted by property subtypes., i.e. Apartments, Office, Industrial and Retail, from 2001 to 2023. Value is defined in terms of 2023 USD.

Table A.5: Transactions By Major Market

	# Trans	% Trans	$ Vol	% Vol	%A	%I	%O	%R
Manhattan	12,617	2.65	733.41	7.27	26.15	0.85	63.82	9.18
Los Angeles	30,892	6.49	578.30	5.73	29.03	20.25	33.95	16.77
Dallas	18,720	3.93	448.09	4.44	44.79	18.30	24.20	12.72
Chicago	19,060	4.00	405.98	4.02	21.55	24.14	36.34	17.97
Atlanta	15,828	3.33	372.71	3.69	43.72	17.27	24.37	14.64
Houston	12,937	2.72	303.41	3.01	42.57	14.15	28.35	14.92
Boston	8,268	1.74	303.20	3.00	20.36	12.95	57.92	8.78
Seattle	10,744	2.26	279.32	2.77	34.78	14.64	39.27	11.30
Phoenix	13,512	2.84	277.81	2.75	46.14	16.21	22.14	15.51
San Francisco	7,561	1.59	242.48	2.40	21.18	8.48	60.49	9.85
DC VA burbs	5,051	1.06	236.12	2.34	36.26	10.96	42.38	10.40
Northern NJ	10,114	2.12	205.36	2.03	24.83	28.42	32.81	13.94
San Diego	9,332	1.96	199.01	1.97	31.34	19.69	33.75	15.22
San Jose	6,280	1.32	197.36	1.96	17.02	23.26	50.56	9.15
Washington DC	2,395	0.50	147.88	1.47	16.20	1.18	78.14	4.48
Miami	7,239	1.52	142.94	1.42	30.99	19.78	27.49	21.74
All Others	285,472	59.97	5,019.61	49.73	36.85	20.13	21.08	21.94

Notes: This table shows the number of transactions, the % of transactions,volume and share of volume for transactions for the largest 16 markets, ranked by volume, as well as in the remaining 44 markets (last row). The last four columns shows the percentage breakdown of volume by sector (Apartments, Industrial, Office, and Retail) within each market. Volume is expressed in billions of 2023 U.S. dollars. The cities in bold are 11 of the 16 superstar cities. The other five superstar cities that are not explicitly listed in the top of this table are: East Bay, DC MD burbs, Orange Co, NYC Boroughs, and Austin.

Table A.6: Transactions By Transaction Type

Type	Subtype	# Trans	% Trans	Vol	% Vol
Conventional Sale	Ind. Asset	354,712	74.53	7068.52	70.06
	Portfolio	89,256	18.75	2087.41	20.69
Entity Sale	Private Merger	1,852	0.39	61.10	0.61
	Privatization	5,832	1.23	289.44	2.87
	Public Merger	7,549	1.59	274.93	2.73
	Publicization	2,098	0.44	61.28	0.61
Foreclosure	Ind. Asset	7,895	1.66	128.55	1.27
	Portfolio	1,867	0.39	22.48	0.22
Debtor/Trustee Sale	Ind. Asset	3,486	0.73	74.64	0.74
	Portfolio	1,415	0.30	20.50	0.20

Notes: This table shows the number of transactions, the share of transactions,volume and share of volume by transaction type and subtype. Volume is expressed in billions of 2023 U.S. dollars.

Table A.7: Mapping Investor Types

RCA Group	RCA Type	Investor Type	%Trans Buy	% Vol Buy	%For Buy	%Trans Sell	% Vol Sell	%For Sell
CMBS	CMBS	Institutional	0.90	0.71	0.68	0.93	0.62	0.43
Equity Fund	Equity Fund	REPE	6.79	12.54	2.03	5.34	10.43	1.89
Institutional	Bank	Institutional	1.25	1.83	26.65	1.75	2.15	25.37
	Endowment	Institutional	0.09	0.12	0.00	0.09	0.12	0.00
	Finance	Institutional	0.33	0.32	22.93	0.32	0.30	13.16
	Insurance	Institutional	0.84	1.88	25.77	0.97	2.13	14.98
	Investment Manager	Institutional	5.90	11.36	19.58	5.00	8.85	17.31
	Open-Ended Fund	Institutional	0.01	0.08	98.74	0.00	0.02	94.23
	Pension Fund	Institutional	1.12	2.73	33.09	1.01	2.29	21.31
	Sovereign Wealth Fund	Foreign	0.57	0.89	100	0.22	0.30	100
Private	Developer/Owner/Operator	Foreign	0.67	0.98	100	0.45	0.68	100
		DOO_L	29.40	12.96	0.02	32.62	16.33	0.02
		DOO_N	32.07	30.43	0.00	28.24	30.44	0.00
	High Net Worth	Individual	1.81	1.20	22.71	2.97	1.41	7.14
	Non Traded REIT	Individual	2.37	3.19	0.06	1.22	1.64	5.09
Public	Listed Funds	REITS	0.01	0.02	89.01	0.01	0.03	78.36
	REIT	REITS	7.09	12.27	6.66	7.34	12.93	5.60
	REOC	REITS	0.61	0.94	72.49	0.49	1.05	64.98
Unknown	Unknown	Unknown	2.39	1.22	0.83	4.12	2.44	0.27
User/Other	Cooperative	User	0.01	0.01	3.68	0.03	0.04	0.20
	Corporate	User	4.17	3.04	6.49	5.84	4.85	6.59
	Educational	User	0.31	0.34	0.10	0.12	0.16	0.13
	Government	User	0.62	0.49	1.12	0.38	0.34	0.73
	Non-Profit	User	0.48	0.31	1.69	0.39	0.30	7.96
	Other	User	0.06	0.06	4.33	0.03	0.03	0.00
	Religious	User	0.14	0.07	2.10	0.11	0.10	2.54

Notes: The top table shows the mapping from RCA investor group (column 1) and RCA investor type (column 2) into our investor type (column 3). We show the share of buy transactions and buy volume for each type in columns 4 and 5 and the share of sell transactions and sell volume for each type in columns 7 and 8. Columns 6 and 9 show the share of buy and sell transactions that include a foreign entity. Volume is expressed in billions of 2023 U.S. dollars.

Table A.8: Transactions by Number of Investors

	1 Seller	2 Sellers	3 Sellers	4 Sellers
1 Buyer	87.30	5.723	0.584	0.154
2 Buyers	4.73	0.654	0.099	0.020
3 Buyers	0.50	0.076	0.004	0.003
4 Buyers	0.14	0.022	0.001	0.001

	1 Seller	2 Sellers	3 Sellers	4 Sellers
1 Buyer	73.34	10.033	1.270	0.367
2 Buyers	10.31	1.914	0.422	0.118
3 Buyers	1.38	0.424	0.026	0.014
4 Buyers	0.29	0.071	0.003	0.018

Notes: The left panel splits the number of transactions into sixteen groups, by the number of buyers and the number of sellers involved in each transactions. The numbers in the table add up to 100%. The right panel does the same based on transaction volume, expressed in billions of 2023 U.S. dollars.

Table A.9: Joint Ventures

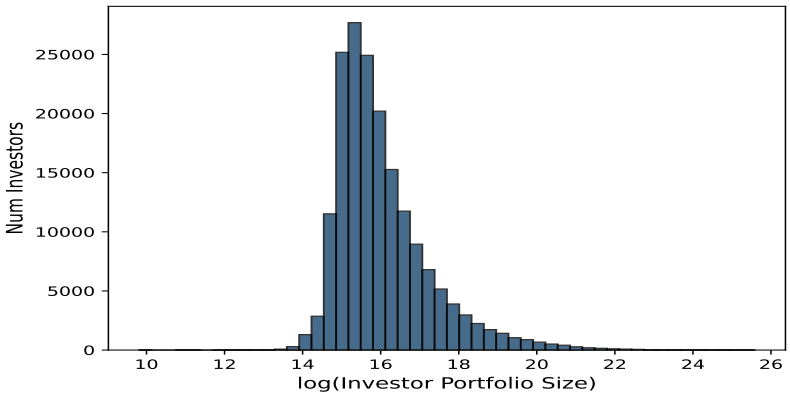
	REPE	Institutional	OOD_L	OOD_N	Individual	REITS	Foreign	User
REPE	2.35							
Institutional	5.03	2.51						
OOD_L	1.22	1.12	5.2					
OOD_N	13.31	14.06	7.45	13.63				
Individual	0.66	1.07	1.0	1.96	0.46			
REITS	1.46	2.75	0.23	2.37	1.09	0.07		
Foreign	1.62	3.41	0.64	6.26	0.72	1.83	4.49	
User	0.09	0.28	0.35	1.01	0.04	0.04	0.04	0.15

	REPE	Institutional	OOD_L	OOD_N	Individual	REITS	Foreign	User
REPE	9.12	16.63	7.10	22.16	9.40	14.84	8.54	4.68
Institutional	19.53	8.29	6.52	23.41	15.31	27.94	17.96	14.16
OOD_L	4.74	3.71	30.24	12.41	14.32	2.28	3.36	17.36
OOD_N	51.70	46.50	43.29	22.70	28.07	24.10	32.94	50.25
Individual	2.55	3.54	5.82	3.27	6.54	11.07	3.77	1.85
REITS	5.68	9.10	1.31	3.95	15.59	0.73	9.62	2.22
Foreign	6.31	11.29	3.71	10.43	10.25	18.58	23.59	1.97
User	0.36	0.94	2.03	1.68	0.53	0.45	0.21	7.51

	REPE	Institutional	OOD_L	OOD_N	Individual	REITS	Foreign	User
% Total Buys	27.34	24.77	4.43	13.99	11.66	10.73	34.16	2.42

Notes: This first panel breaks down the total number of buy transactions involving multiple buyers into who forms a JV with whom. The second panel shows this number as a share of JV property purchases by each investor type; every column sums to 100%. The third panel table shows the share of properties bought in JV with every other investor type, as a share of each investor's total purchases.

Figure A.3: Investor Size Distribution



Notes: Figure shows a histogram of log(investor portfolio size) at the end of 2023 for investors that own at least one property. Log investor size varies from 11.8 (~ $135K) to 25.5 (~ $119 bi). Note that 16.2 = $10 mi, 17.7 = $50 mi, 20.7 = $1 bi.

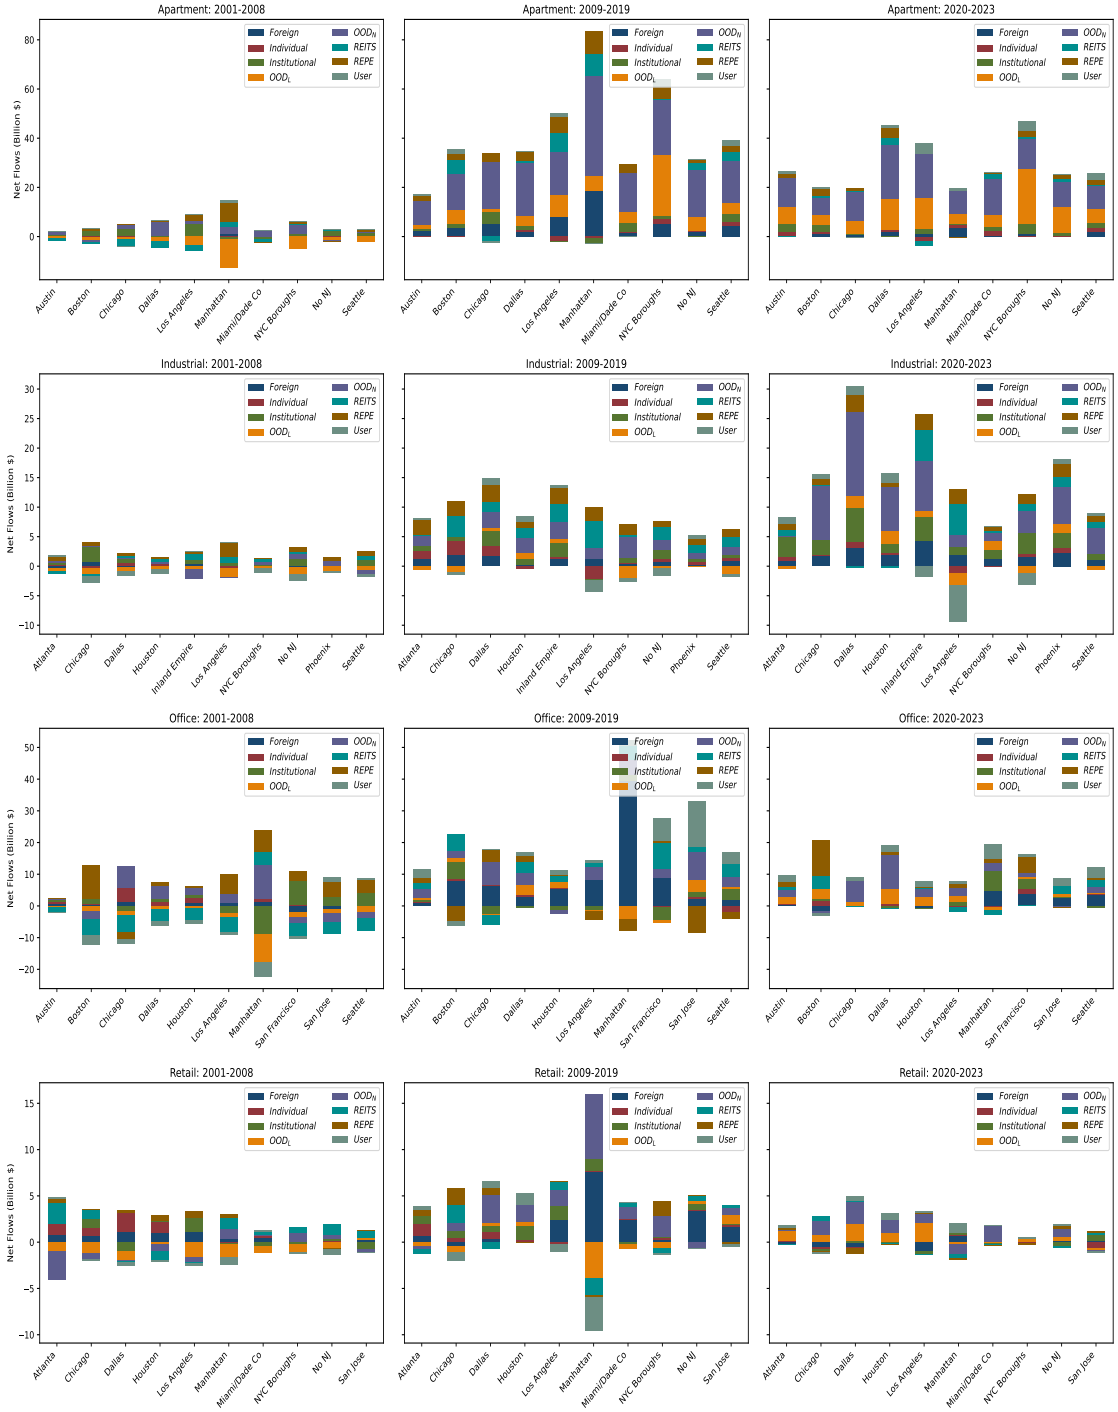
Table A.10: Trade Networks

	REPE	Institutional	OOD_L	OOD_N	Individual	REITS	Foreign	User	Unknown
REPE	205.87	241.76	55.18	334.05	44.70	132.03	159.68	28.33	4.71
Institutional	243.07	321.74	82.39	461.55	56.27	192.81	176.05	42.65	7.44
OOD_L	102.32	137.32	590.45	610.10	52.38	83.67	55.94	96.05	27.51
OOD_N	368.63	536.94	395.71	1263.38	152.55	309.66	235.54	111.02	29.29
Individual	27.06	36.34	52.15	82.97	17.98	62.32	31.71	9.40	2.23
REITS	318.93	243.37	39.15	278.62	51.98	324.63	186.06	30.75	8.51
Foreign	145.44	131.80	29.09	164.04	32.83	71.97	123.05	25.37	1.93
User	68.04	81.48	88.76	182.84	22.79	80.48	34.79	74.60	3.90
Unknown	9.67	23.89	48.21	76.17	9.65	27.93	8.43	12.82	37.86

	REPE	Institutional	OOD_L	OOD_N	Individual	REITS	Foreign	User	Unknown
REPE	1.81	2.13	0.49	2.94	0.39	1.16	1.40	0.25	0.04
Institutional	2.14	2.83	0.72	4.06	0.49	1.70	1.55	0.38	0.07
OOD_L	0.90	1.21	5.19	5.37	0.46	0.74	0.49	0.84	0.24
OOD_N	3.24	4.72	3.48	11.11	1.34	2.72	2.07	0.98	0.26
Individual	0.24	0.32	0.46	0.73	0.16	0.55	0.28	0.08	0.02
REITS	2.80	2.14	0.34	2.45	0.46	2.85	1.64	0.27	0.07
Foreign	1.28	1.16	0.26	1.44	0.29	0.63	1.08	0.22	0.02
User	0.60	0.72	0.78	1.61	0.20	0.71	0.31	0.66	0.03
Unknown	0.09	0.21	0.42	0.67	0.08	0.25	0.07	0.11	0.33

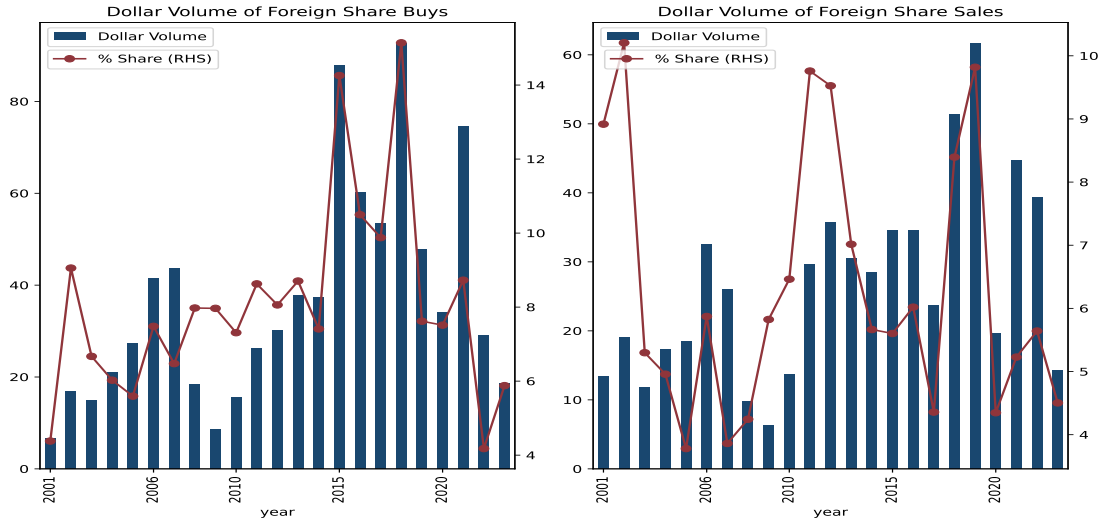
Notes: This table summarizes the trade networks among various investor types. We show buyers along the columns and sellers along the rows. Top Panel reports the $ Volume of transaction between different investor types, expressed in billions of 2023 U.S. dollars. If a transaction involves multiple buyers (sellers) of different types, transaction value is split among them proportionately. The bottom table reports the share of overall transaction $ Volume. Values in the entire table sum up to 100%

Figure A.4: Net Flows by Investor Types



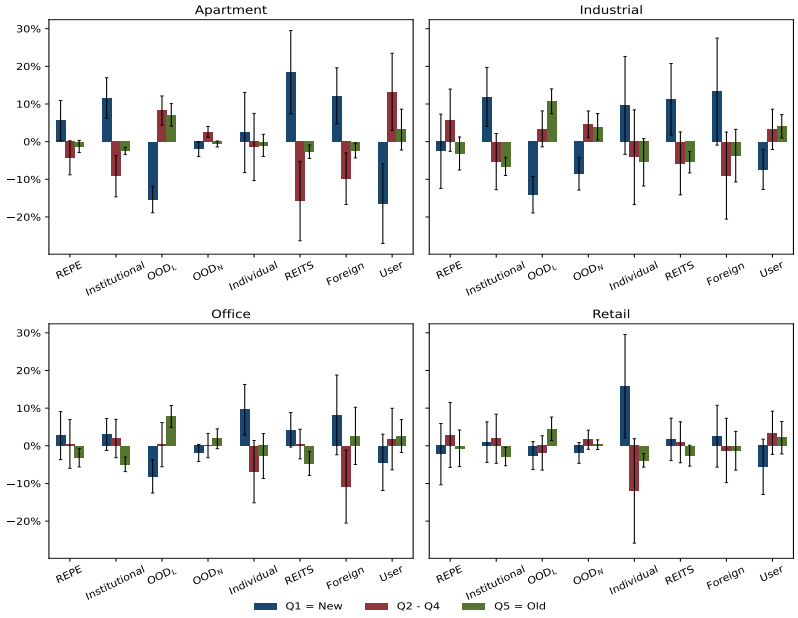
Notes: This figure shows the net acquisitions of properties by investor types and sector during 2001-2008, 2009-2019 and 2020-2023, calculated as the dollar value of purchases minus sales (in 2023 dollars) for each sector.

Figure A.5: Foreign Investor Share



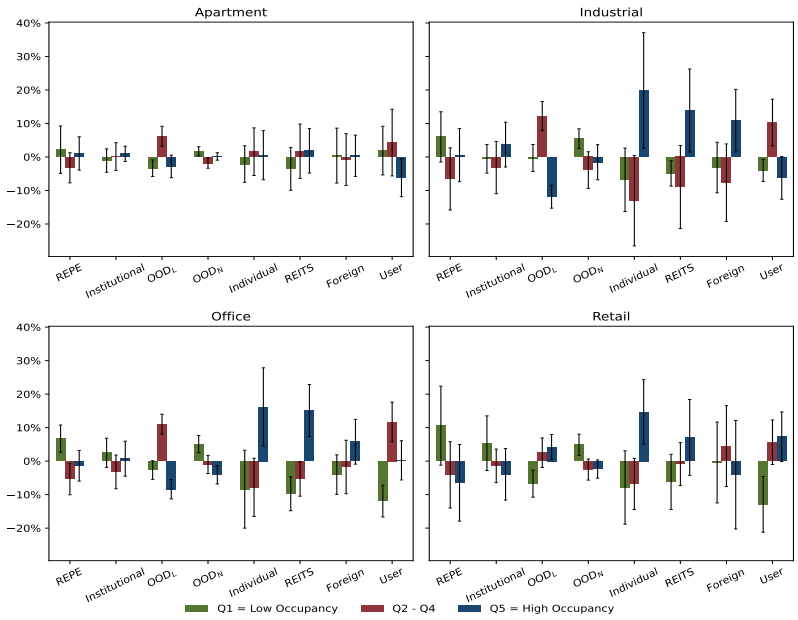
Notes: Figures show foreign investor activity in US Commercial Real Estate Ecosystem during 2001-2023. The blue bars show the dollar value (in Billions, 2023 USD) traded, aggregated by property type while the red lines show the share of foreign investors among all investors. The left panel summarizes foreign investor buys while the right panel summarizes foreign investor sales.

Figure A.6: Investor Type Heterogeneity: Renovation Adj. Age



Notes: Figures show Investor heterogeneity in preferences for renovation adjusted age. Bars show the average excess share of transaction volume by investor type in each quintile of renovation-adjusted age relative to the market-wide share. Vertical lines indicate time-series standard deviations.

Figure A.7: Investor Type Heterogeneity: Occupancy Rate



Notes: Figures show Investor heterogeneity in preferences for occupancy. Bars show the average excess share of transaction volume by investor type in each quintile of occupancy relative to the market-wide share. Vertical lines indicate time-series standard deviations.

B Valuation Model Appendix

B.1 Linear Hedonic Model

Table B.1 shows the results for the LHM estimation where the log price per square foot (I, O, R) or per unit (A) is a linear function of property characteristics. Columns (1), (3), (5), and (7) include only asset characteristics: log asset size, occupancy rate, log age, log renovation-adjusted age, subtype indicator, number of floors, CBD indicator, superstar city indicator, entity sale indicator, and distressed sale indicator. Since the office subtypes are CBD and suburban office, the CBD indicator is co-linear with the subtype category CBD, and we drop the CBD indicator from the model for Office. For Apartments, the subtype indicator is 1 for Mid/Highrise and 0 for Garden, for Industrial, it is 1 for Flex and 0 for warehouse and for Retail properties, it is 1 for shops and 0 for centers.

All property characteristics enter with the expected sign and most "hedonic prices" are precisely estimated. The deal variables show that transactions that are part of an entity sale trade at a premium whereas properties that are part of a distressed deal trade at a large discount.

In columns (2), (4), (6), and (8) of Table B.1 we consider the full LHM which adds the five local macroeconomic variables, market fixed effects, and year fixed effects. The market fixed effects absorb the superstar indicator variable, which is consequently dropped form the regression model. NER and personal income per capita are strongly statistically significant in each sector. Inclusion of local macro variables cuts the coefficient on the distressed deal (the fire sale discount) indicator in half.

B.2 LGBM Hyperparameter tuning and Model Complexity

Light Gradient Boosting Machine (Ke et al., 2017), henceforth LGBM, is a gradient-boosted decision tree algorithm designed for efficiency and scalability. Unlike traditional tree-based models, it uses histogram-based learning to speed up training and leaf-wise tree growth to capture complex patterns in the data. It can capture complex, non-linear relationships without requiring explicit interaction terms, automatically select relevant features and interactions, improving model interpretability and regularize effectively to prevent overfitting.

We fine-tune LGBM hyperparameters for each model specification using Optuna (Takuya Akiba and Koyam, 2019), a Bayesian optimization framework. Hyperparameters are selected via 10-fold cross-validation, minimizing mean squared error (MSE) on the validation folds. Each model undergoes 50 Optuna trials, and the configuration with the lowest average validation MSE is retained for out-of-sample evaluation. The search space is deliberately bounded to avoid overly complex models that may overfit. We vary the number of estimators between 500 and 1,000 in increments of 10, and the maximum tree depth between 2 and 5 for our baseline models. The number of leaves is constrained by the chosen depth, ensuring trees remain interpretable and regularized. Minimum child samples vary from 1% to 10% of the number of observations, in increments of 50, to control leaf-level granularity. The feature and bagging fractions range from 0.5 to 1.0 in steps of 0.01, encouraging subsampling for robustness. L1 and L2 regularization coefficients are searched between 0.0 and 5.0 in increments of 0.1, while the minimum split gain is restricted to 0.2–1.0 to discourage splits with marginal predictive value, thus keeping the maximum depth under control. The learning rate is tuned between 0.01 and 0.1 on a log scale.

Restricting tree depth to 2–5 in the baseline search produces trees that capture main effects and low-order interactions without memorizing idiosyncratic patterns in the training data. In the "deep" LGBM variant, we relax this bound and allow max depth to be selected by cross-validation without restriction, resulting in depths of 7–11. While deeper trees can model more complex nonlinearities and higher-order

Table B.1: Reference Point: Linear Hedonic Model

	Apartment		Industrial		Office		Retail	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CBD Indicator	0.061	0.165**	0.427***	0.288***	0.272***	0.105**	0.425***	0.269***
	(0.129)	(0.060)	(0.068)	(0.060)	(0.093)	(0.049)	(0.095)	(0.052)
Age	-0.064***	-0.078***	0.015*	0.001	-0.013	-0.030***	0.010	-0.004
	(0.010)	(0.009)	(0.009)	(0.007)	(0.008)	(0.007)	(0.007)	(0.006)
Renovation Adj. Age	-0.054***	-0.033***	-0.104***	-0.092***	-0.095***	-0.078***	-0.108***	-0.104***
	(0.013)	(0.007)	(0.012)	(0.010)	(0.012)	(0.011)	(0.009)	(0.009)
Property Size	-0.197***	-0.091***	-0.370***	-0.266***	-0.268***	-0.221***	-0.395***	-0.369***
	(0.024)	(0.021)	(0.016)	(0.015)	(0.021)	(0.021)	(0.015)	(0.014)
Property Subtype	0.138***	0.145***	0.160***	0.129***			-0.015	0.027
	(0.048)	(0.027)	(0.030)	(0.020)			(0.024)	(0.025)
No. of Floors	0.203***	0.120***	0.064**	0.023	0.139***	0.097***	0.198***	0.059***
	(0.024)	(0.017)	(0.030)	(0.021)	(0.016)	(0.010)	(0.024)	(0.020)
Occupancy Rate	1.031***	0.709***	0.174**	0.162***	0.974***	0.835***	0.591***	0.568***
	(0.193)	(0.140)	(0.067)	(0.038)	(0.073)	(0.052)	(0.074)	(0.055)
Superstarcity	0.561***		0.400***		0.520***		0.416***	
	(0.092)		(0.111)		(0.077)		(0.075)	
Entity Sale	0.270**	0.220**	0.142	0.118	0.161	0.140	-0.020	0.044
	(0.106)	(0.091)	(0.098)	(0.094)	(0.109)	(0.090)	(0.119)	(0.109)
Distressed	-0.498***	-0.229***	-0.466***	-0.226***	-0.447***	-0.292***	-0.350***	-0.276***
	(0.048)	(0.029)	(0.048)	(0.032)	(0.035)	(0.033)	(0.043)	(0.040)
Market Occupancy		0.094		-0.140		0.217*		0.006
		(0.306)		(0.091)		(0.117)		(0.091)
NOI Growth		0.073		0.013		0.027		-0.069
		(0.093)		(0.046)		(0.033)		(0.043)
Personal Income		0.603***		0.288***		0.359***		0.431***
		(0.071)		(0.066)		(0.052)		(0.031)
Population/Employment		0.025*		0.019		-0.027**		0.038***
		(0.013)		(0.013)		(0.012)		(0.006)
NER		0.043***		0.224***		0.446***		0.173***
		(0.007)		(0.033)		(0.062)		(0.033)
Year FE	×	✓	×	✓	×	✓	×	✓
Market FE	×	✓	×	✓	×	✓	×	✓
Observations	141,135	141,135	116,737	116,737	96,139	96,139	114,223	114,223
Adj. R^2	41.1	59.2	44.7	58.5	34.4	48.0	51.7	58.4

Notes: This table shows the full-sample coefficient estimates and goodness of fit statistics for a linear hedonic valuation model. Standard errors in bracket are clustered by RCA Market and Year.

interactions, they carry a greater risk of overfitting, especially in settings with persistent residual correlation. Our design mitigates this by combining depth constraints (in the baseline), strong regularization, feature and row subsampling, and a relatively high min child samples threshold, all of which reduce variance and improve stability across folds. Together, these choices ensure that any gains from increased flexibility in the deep model reflect genuine predictive signal rather than noise-fitting.

Tables B.2 and B.3 present the hyperparameter search spaces and optimized values for the final LGBM models, i.e. those with investor characteristics.

We evaluate the effect of increasing model complexity on valuation performance by gradually expanding the set of predictors and allowing for progressively richer functional forms. Figure B.1 shows the in-sample R^2 as we move from a linear baseline hedonic model with only asset and transaction characteristics to specifications that incorporate macroeconomic variables, time and market fixed effects, investor characteristics, non-linear transformations of continuous predictors, and select interaction terms. This staged build-up

Table B.2: LGBM Hyperparameters: Buyer Model

	Search Space	Apartment	Industrial	Office	Retail
Number of Estimators	500 - 1000	750	860	750	600
Maximum Tree Depth	2 - 5	5	5	4	4
Number of Leaves	3 - $2^{MaxDepth}$	27	27	13	11
Minimum Child Samples	1% Obs - 10% Obs	1410	1170	1360	1140
Feature Fraction	0.5 - 1.0	0.76	0.74	0.87	0.56
Bagging Fraction	0.5 - 1.0	0.89	0.62	0.95	0.54
L1 Regularization	0.0 - 5.0	3.60	0.80	0.30	0.80
L2 Regularization	0.0 - 5.0	2.10	0.50	2.80	1.80
Minimum Split Gain	0.2 - 1.0	0.44	0.37	0.29	0.45
Learning Rate	0.01 - 0.1	0.08	0.06	0.06	0.06
Number of Obs		141,135	116,737	96,139	114,223

Table B.3: LGBM Hyperparameters: Seller Model

	Search Space	Apartment	Industrial	Office	Retail
Number of Estimators	500 - 1000	930	670	600	870
Maximum Tree Depth	2 - 5	5	5	4	5
Number of Leaves	3 - $2^{MaxDepth}$	20	29	15	23
Minimum Child Samples	1% Obs - 10% Obs	2360	1170	1280	1140
Feature Fraction	0.5 - 1.0	0.73	0.64	0.89	0.85
Bagging Fraction	0.5 - 1.0	0.66	0.65	0.72	0.73
L1 Regularization	0.0 - 5.0	1.90	1.70	0.20	3.50
L2 Regularization	0.0 - 5.0	1.70	1.60	1.40	0.80
Minimum Split Gain	0.2 - 1.0	0.26	0.88	0.34	0.27
Learning Rate	0.01 - 0.1	0.05	0.03	0.07	0.06
Number of Obs		141,135	116,737	96,139	114,223

illustrates how much explanatory power is attributable to each successive modeling enhancement.

Similarly, Table B.4 reports the results of a simple linear model that adds investor type fixed effects and the other five investor portfolio characteristics in z_{it} in columns (1), (3), (5), and (7) for the various sectors. Columns (2), (4), (6), and (8) add the top-5 interaction effects, identified by the LGBM model, to the linear model. The investor characteristics and the key interaction effects are typically estimated precisely. Even simple linear inclusion of investor variables and some basic interaction effects significantly boost the fit for transaction prices.

This shows that moving from a linear model to then allowing for more flexible functional forms—first through targeted transformations and selected interactions, then via a fully non-linear LGBM specification, substantially improves the fit. The final step in our complexity ladder is the "deep LGBM" model, in which the maximum tree depth is determined by cross-validation without restriction.

Increasing model complexity comes with important trade-offs. More complex models, particularly those with deeper trees, tend to achieve higher in-sample R^2 but at the risk of fitting patterns that do not generalize out-of-sample. In our max-depth analysis (unreported), the marginal gains in in-sample R^2 beyond depths of 4–5 are modest, while the gap between in-sample and out-of-sample R^2 widens. This widening gap reflects the classic bias-variance trade-off: deeper models reduce bias by capturing more intricate patterns

Table B.4: Linear Model with Investor Characteristics

	Apartment		Industrial		Office		Retail	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CBD Indicator	0.120*	0.098*	0.248***	0.185***	0.126***	0.094***	0.204***	0.170***
	(0.058)	(0.053)	(0.054)	(0.044)	(0.038)	(0.033)	(0.042)	(0.039)
Age	-0.065***	-0.059***	0.004	0.007	-0.020***	-0.011**	-0.005	-0.005
	(0.008)	(0.007)	(0.006)	(0.004)	(0.006)	(0.005)	(0.006)	(0.004)
Renovation Adj. Age	-0.026***	-0.021***	-0.085***	-0.068***	-0.067***	-0.058***	-0.087***	-0.061***
	(0.008)	(0.007)	(0.008)	(0.005)	(0.008)	(0.006)	(0.008)	(0.005)
Property Size	-0.233***	-1.690***	-0.369***	-1.898***	-0.386***	-2.112***	-0.456***	-2.244***
	(0.018)	(0.094)	(0.017)	(0.080)	(0.024)	(0.120)	(0.020)	(0.058)
Property Subtype	0.099***	0.135***	0.106***	0.099***			0.031	0.004
	(0.022)	(0.020)	(0.018)	(0.016)			(0.022)	(0.021)
No. of Floors	0.122***	0.110***	0.047**	0.064***	0.087***	0.048***	0.050**	0.035*
	(0.014)	(0.012)	(0.019)	(0.013)	(0.008)	(0.006)	(0.020)	(0.017)
Occupancy Rate	0.542***	0.449***	0.241***	0.240***	0.699***	0.607***	0.459***	0.409***
	(0.102)	(0.094)	(0.034)	(0.028)	(0.044)	(0.034)	(0.054)	(0.036)
Entity Sale	-0.071	-0.088	-0.172*	-0.161*	-0.237**	-0.204**	-0.277**	-0.189**
	(0.079)	(0.078)	(0.095)	(0.086)	(0.088)	(0.072)	(0.107)	(0.070)
Distressed	-0.218***	-0.175***	-0.136***	-0.112***	-0.164***	-0.144***	-0.142***	-0.100***
	(0.031)	(0.024)	(0.032)	(0.027)	(0.026)	(0.026)	(0.037)	(0.031)
Market Occupancy	0.043	0.025	-0.149*	-0.112*	0.126	0.081	-0.066	-0.045
	(0.281)	(0.247)	(0.079)	(0.061)	(0.091)	(0.082)	(0.099)	(0.084)
NOI Growth	0.047	0.026	-0.011	0.019	0.044	0.056*	-0.057	-0.033
	(0.079)	(0.069)	(0.041)	(0.022)	(0.027)	(0.028)	(0.038)	(0.031)
Personal Income	0.501***	0.450***	0.226***	0.175***	0.287***	0.217***	0.387***	0.285***
	(0.067)	(0.059)	(0.054)	(0.043)	(0.049)	(0.042)	(0.037)	(0.029)
Population/Employment	0.011	-0.001	0.004	0.004	-0.019*	-0.021**	0.036***	0.025***
	(0.012)	(0.011)	(0.010)	(0.008)	(0.010)	(0.008)	(0.007)	(0.006)
NER	0.038***	0.001	0.204***	-0.628***	0.363***	-0.323**	0.144***	-0.403***
	(0.005)	(0.023)	(0.027)	(0.125)	(0.049)	(0.150)	(0.024)	(0.086)
Investor Size	0.146***	-0.367***	0.145***	-0.394***	0.212***	-0.365***	0.187***	-0.304***
	(0.012)	(0.024)	(0.012)	(0.030)	(0.011)	(0.038)	(0.015)	(0.026)
Same Market Share	0.345***	-2.159***	0.465***	-4.463***	0.506***	-4.021***	0.759***	-5.297***
	(0.063)	(0.324)	(0.044)	(0.521)	(0.030)	(0.490)	(0.052)	(0.494)
Same Sector Share	-0.234***	-1.644***	-0.211***	0.791*	0.428***	-0.286	0.056	1.336***
	(0.075)	(0.287)	(0.038)	(0.415)	(0.040)	(0.390)	(0.058)	(0.370)
Superstarcity Share	0.117**	-0.045	0.048	-0.035	-0.033	-0.041*	0.127***	-0.005
	(0.047)	(0.046)	(0.032)	(0.022)	(0.039)	(0.023)	(0.041)	(0.030)
JV	0.085**	0.113***	-0.049	-0.030	0.166***	0.154***	0.148**	0.094*
	(0.031)	(0.023)	(0.054)	(0.044)	(0.025)	(0.023)	(0.068)	(0.046)
Investor Size × Property Size		0.082***		0.084***		0.087***		0.093***
		(0.005)		(0.004)		(0.004)		(0.003)
Investor Size × Same Market Share		0.141***		0.273***		0.244***		0.333***
		(0.018)		(0.031)		(0.028)		(0.028)
Investor Size × Same Sector Share		0.071***		-0.043**		0.036*		-0.061***
		(0.013)		(0.020)		(0.020)		(0.020)
Investor Size × NER		0.004**		0.052***		0.030***		0.030***
		(0.002)		(0.008)		(0.009)		(0.006)
Property Size × NER		-0.008***		-0.035***		0.013		-0.012
		(0.002)		(0.012)		(0.018)		(0.010)
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Market FE	✓	✓	✓	✓	✓	✓	✓	✓
Investor Type FE	✓	✓	✓	✓	✓	✓	✓	✓
Adj. R^2	65.1	68.9	64.1	71.3	57.7	66.0	64.3	73.0

Notes: This table shows the full-sample coefficient estimates and goodness of fit statistics for a linear valuation model with investor characteristics and select interactions. Standard errors in bracket are clustered by RCA Market and Year.

but increase variance by fitting idiosyncrasies in the training data.

In our setting, where both asset characteristics and investor attributes can exhibit persistence and cross-

sectional correlation, these safeguards are important. A model that is too shallow may fail to capture economically important interaction effects, but a model that is too deep may overfit to transient, investor-specific shocks embedded in the training sample. The optimal balance depends on the objective of the model, forecasting horizon and the stability of the underlying relationships: for short-horizon prediction, some residual persistence can legitimately be exploited; for structural interpretation, more conservative complexity controls are required.

Having tuned the model, we can estimate the model out of sample (OOS) using 10-fold cross-validation. Table B.5 reports the goodness of fit (R^2) statistics, computed as the averages of the to hold-out samples from the 10 folds.

Table B.5: Valuation Model Goodness of Fit, Out of Sample (OOS)

Sector	Apartment		Industrial		Office		Retail	
	R^2	R2C	R^2	R2C	R^2	R2C	R^2	R2C
Hedonic Model	49.56		50.39		47.11		60.0	
+ Macro Vars	67.14		62.93		58.37		65.09	
+ Year Dummies	68.21		63.62		58.53		66.13	
+ Market Dummies	70.91		66.09		60.57		67.65	
+ Investor Types	76.34	18.67***	72.3	18.31***	67.08	16.51***	73.83	19.1***
+ Portfolio Vars	83.35	42.76***	83.63	51.73***	80.94	51.66***	85.24	54.37***

Notes: This table summarizes the performance R^2 from valuation model estimations in equation (8) using LGBM. The two panels summarize the performances of our non-linear valuation model by successively adding variables as described in 5. Each row chooses the optimal LGBM hyperparameters using cross validation. The R^2 are out-of-sample, computed as averages of 10 hold-out samples in a 10-fold cross-validation.

B.3 Valuation Model: Variable Importance

B.3.1 SHAP Values

SHAP (Shapley Additive Explanations) values provide a game-theoretic approach to interpreting machine learning models by quantifying the marginal contribution of each predictor to the model's predictions. For a given observation, SHAP values measure how much each feature pushes the model's output away from the mean prediction. Mathematically, for a prediction function $f(x)$ and a set of features N, the SHAP value for feature j is calculated as:

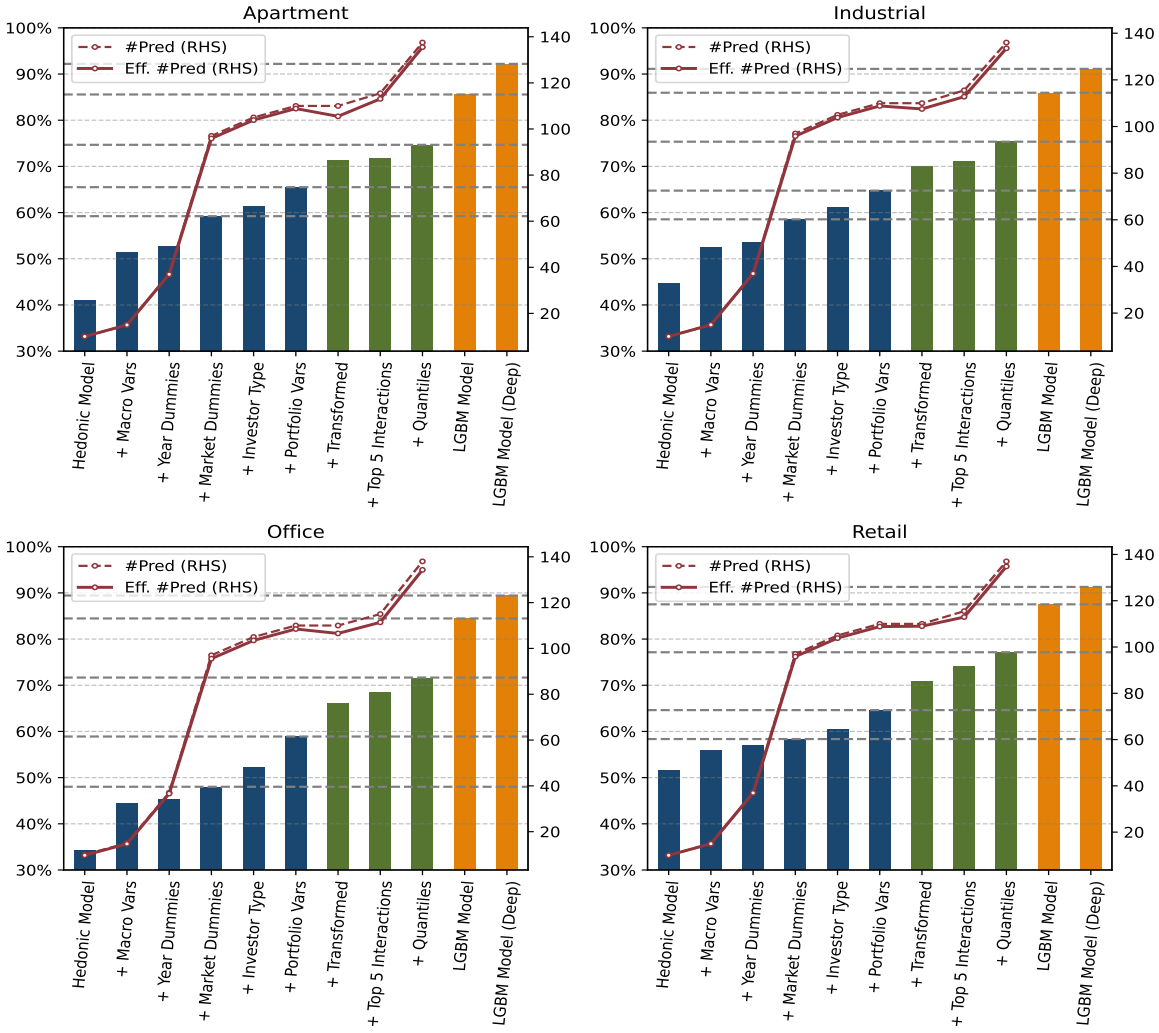
$$\phi_j = \sum_{S \subseteq N \setminus \{j\}} \frac{|S|!(M - |S| - 1)!}{M!} (f(S \cup \{j\}) - f(S)),$$

where M is the total number of features and $f(S)$ represents the model output when only the features in subset S are known. Feature Importance (I_j) is obtained by averaging its absolute SHAP value over all n observations.

$$Imp_j = \frac{1}{n} \sum_{i=1}^{n} \left| \phi_j^{(i)} \right|$$

SHAP values provide a more comprehensive importance measure by accounting for both direct and interaction effects, unlike traditional importance measures such as split frequency.

Figure B.1: Incremental R^2 from Sequential Model Enhancements



Notes: Figure shows the In-sample R^2 for all sectors as we sequentially expand the set of predictors, starting from the baseline hedonic model containing only asset and transaction characteristics (the first blue bar). The next five blue bars add variables step-by-step: macro variables, year and market fixed effects, investor type variables, and investor portfolio variables (see sections 5.2–5.3), but retain the linear setting. The "+ Transformed" bar replaces continuous predictors with their best-fit polynomial transformations, as implied by the LGBM SHAP dependence plots in Figure B.3. The next green bar ("+ Top 5 Interactions") adds the five most important interaction terms identified in Figure 7. Given the prominence of investor size, property size, and their interaction, the following green bar ("+ Quantiles") includes quintile dummies for investor size and property size and their interaction, while dropping terms eliminated by a simple LASSO selection. The final two yellow bars report R^2 for our full LGBM non-linear model and for the "deep LGBM" specification as described in section B.2 where the maximum tree depth is selected via cross-validation without restriction. All R^2 values are obtained from ridge regressions to ensure stable coefficient estimates in the presence of multi-collinearity. The solid red line plots the effective number of predictors-calculated from the singular values of the standardized design matrix as $\sum s^2/\left(s^2+\alpha\right)$-while the dotted red line shows the actual number of predictors included at each stage.

Following Lundberg et al. (2020), we compute SHAP interaction values for all pairwise interactions. The SHAP interaction values quantify how a feature pair jointly influences predictions beyond their individual

effects. For a pair of features i and j these interaction importance values are defined as:

$$Imp_{ij} = \sum_{S \subseteq N \setminus \{i,j\}} \frac{|S|!(M - |S| - 2)!}{2(M-1)!} \left[f(S \cup \{i,j\}) - f(S \cup \{i\}) - f(S \cup \{j\}) + f(S) \right]$$

where the term inside the square brackets captures the marginal effect of including both features i and j together compared to including them individually.

B.3.2 Additional Variable Importance Results

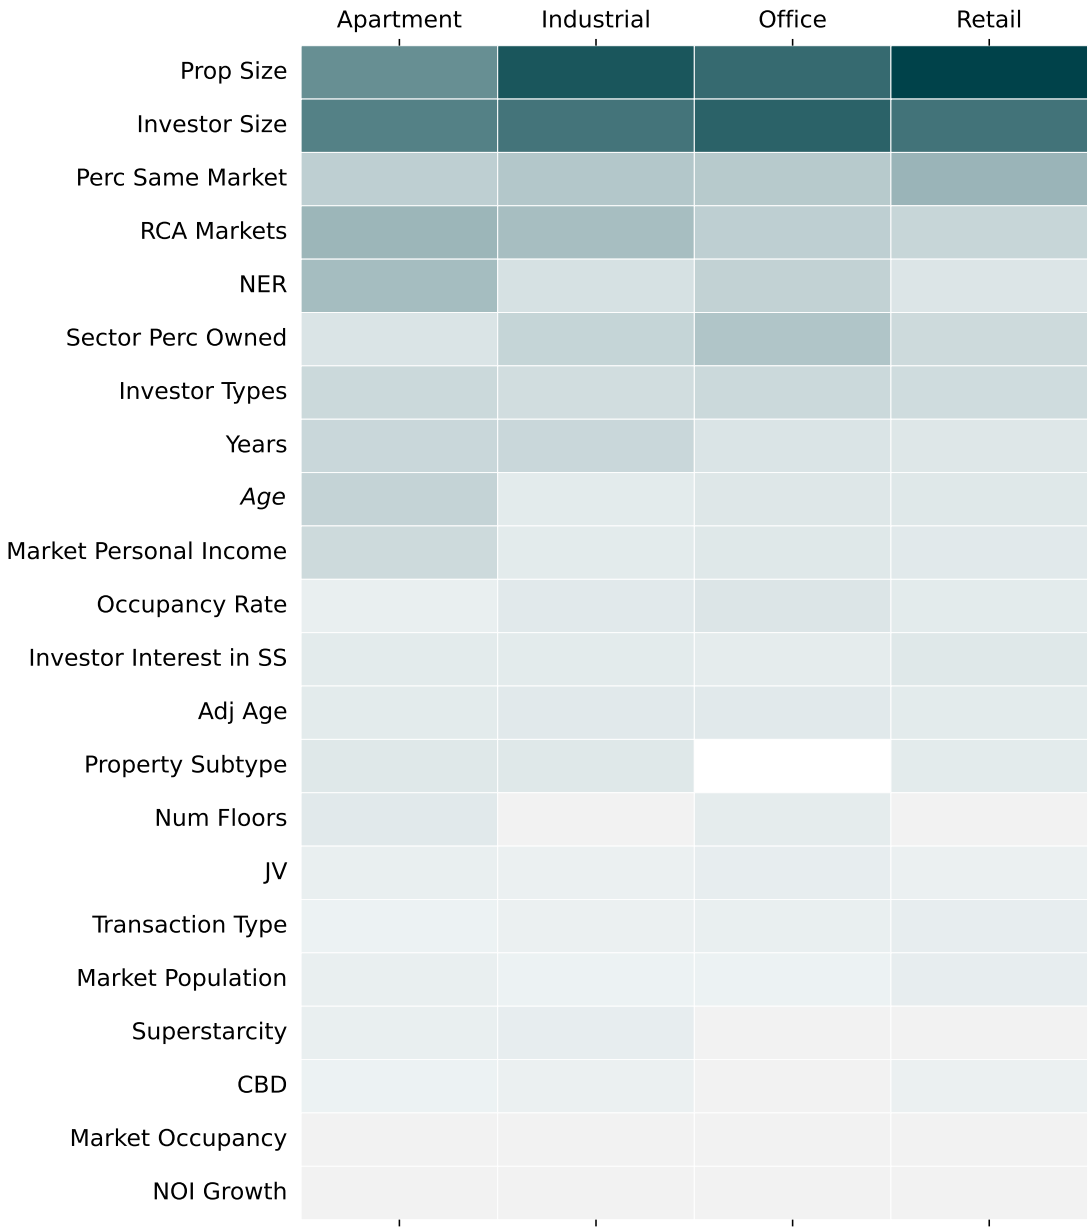
Figure B.2 shows the most important features, ranked by SHAP value I_j selected by the LightGBM Model.

Figure B.3 shows the results for the nine most important continuous features, reported in the title of each panel. For each feature, we compute the average SHAP value at different percentiles of the feature (from the 5th to the 95th percentile) and plot these against the corresponding feature values using a binned scatter plot. The X-axis represents the standardized feature value, while the Y-axis represents the standardized SHAP value (averaged across the buyer and seller valuation models). The first panel shows that small investor size has negative SHAP contributions, pulling predicted prices downward, whereas large investor size raises predicted prices.

Figure B.4 plots the average SHAP value for the year categorical variable.

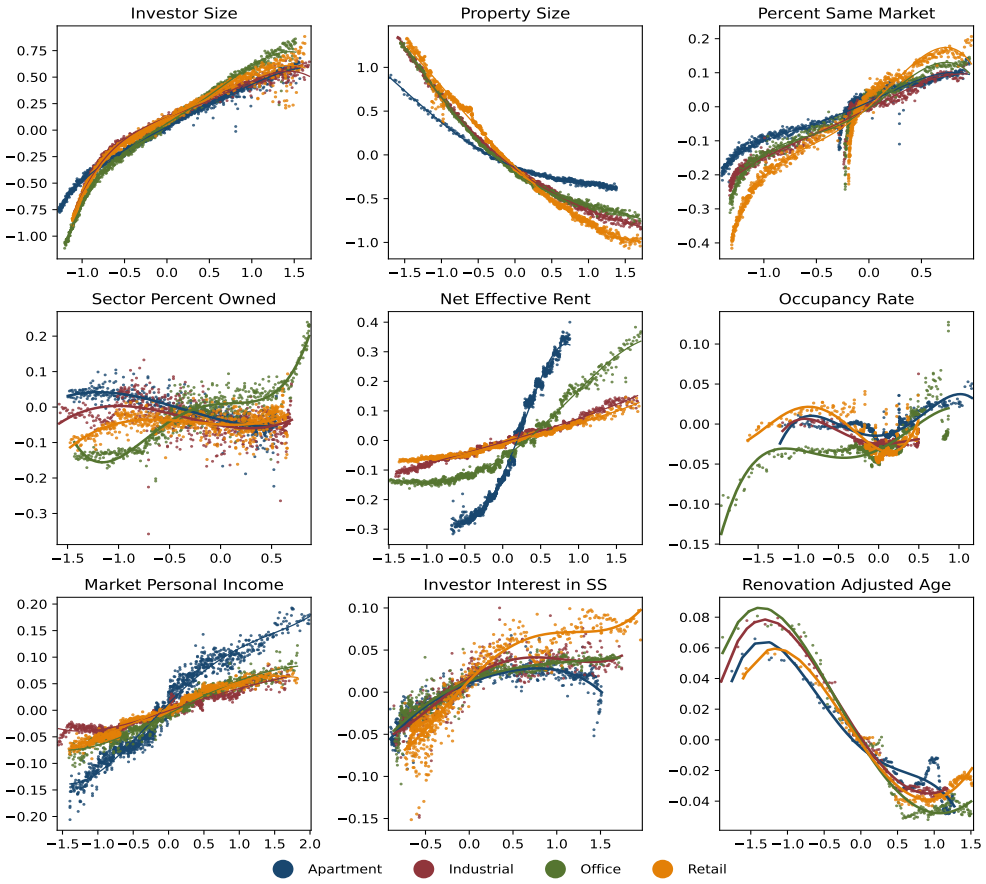
Figure B.5 studies the interaction effects between investor size and property size, the most important interaction effect identified in Figure 7. For the pair, we compute the average interaction SHAP value at different percentiles of the feature and plot these against the corresponding feature values using a binned scatter plot. The X-axis shows the standardized investor size, while the Y-axis represents the standardized property size. The color intensity reflects the magnitude and sign of the SHAP values, where red indicates a negative interaction effect and blue represents a positive interaction effect. Unlike the main SHAP values, which include both direct and interaction effects, this plot isolates only the interaction component, which measures how the effect of one variable changes depending on the level of the other variable. A key observation is that large investors have lower valuations for small properties and small investors have dampened valuations for large properties in every sector.

Figure B.2: LightGBM Feature Importance



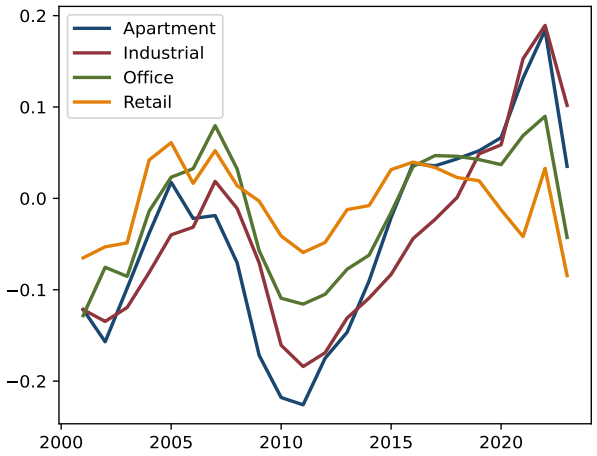
Notes: The figure reports the importance of features selected by the LightGBM Model. Feature importance is defined as the absolute SHAP (Shapley Additive Explanations) value for each feature summed across all observations and is normalized such that the importance of all features sums to 1 in each sector. Features are sorted by the average importance within each sector. Investor type is treated as a single feature.

Figure B.3: SHAP Dependence: Non-linearities



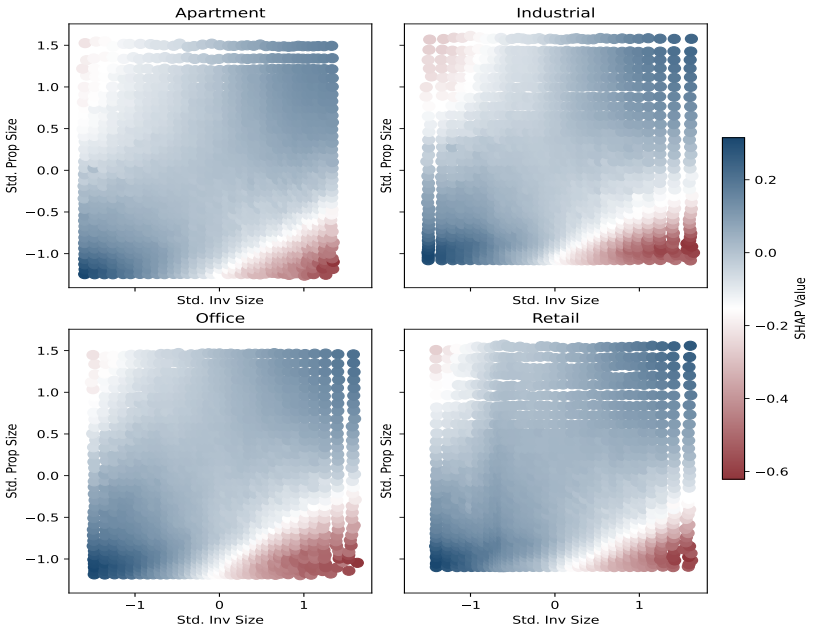
Notes: This figure shows SHAP dependence plots for the nine most important continuous features in our valuation models. Each plot is a binned scatter plot with a best-fit polynomial trend line. The X-axis represents the standardized feature value, while the Y-axis represents the standardized SHAP value, averaged across the buyer and seller valuation models, to allow comparability across sectors.

Figure B.4: Year Fixed Effects in the Valuation Model



Notes: This figure plots the average SHAP value for the year categorical variable. SHAP values are averaged across buyer and seller models and normalized to ensure comparability across sectors.

Figure B.5: SHAP Interaction Effects



Notes: This figure shows a binscatter plot to understand the interaction effects between investor size and property size, capturing how their combined influence deviates from the sum of their individual contributions. The color intensity reflects the magnitude and direction of the interaction SHAP values, where deep red indicates a negative interaction effect, while blue represents a positive interaction effect.

C Extensions: Bargaining Power and Transaction Costs

In the baseline specification of Section 2.3, we impose equal bargaining weights, $\beta = 0.5$, and zero proportional transaction costs, $c = 0$. These assumptions simplify the valuation equation estimated on observed transactions. This appendix shows how the analysis changes when either assumption is relaxed.

C.1 Selection Bias When $\beta \neq 0.5$

Let

$$v_b = h_b + \epsilon_b, \qquad v_s = h_s + \epsilon_s,$$

where $\epsilon_b, \epsilon_s \sim \mathcal{N}(0, \sigma_t^2)$ are independent. The log transaction price is

$$p = \beta v_b + (1-\beta)v_s - c = \beta h_b + (1-\beta)h_s - c + \beta\epsilon_b + (1-\beta)\epsilon_s,$$

and trade occurs when $v_b - v_s > c$. Define

$$u \equiv \epsilon_b - \epsilon_s, \qquad w \equiv \epsilon_b + \epsilon_s.$$

Since ϵ_b and ϵ_s are independent Normal shocks with equal variance, u and w are independent. The price residual can be written as

$$r(\beta) \equiv \beta\epsilon_b + (1-\beta)\epsilon_s = \frac{1}{2}w + \left(\beta - \frac{1}{2}\right)u.$$

When $\beta = 1/2$, the residual depends only on w, while trade depends only on u, so conditioning on trade does not induce selection bias. When $\beta \neq 1/2$, the residual contains the additional term $(\beta - \frac{1}{2})u$, and selection on trade generates a non-zero mean residual. Let

$$\Delta \equiv h_b - h_s - c, \qquad z \equiv \frac{\Delta}{\sqrt{2}\sigma_t}.$$

Since trade occurs when $u > -\Delta$ and $u \sim \mathcal{N}(0, 2\sigma_t^2)$,

$$\mathbb{E}[u \mid \text{trade}] = \sqrt{2}\sigma_t \frac{\phi(z)}{\Phi(z)}.$$

Hence,

$$\mathbb{E}[r(\beta) \mid \text{trade}] = \left(\beta - \frac{1}{2}\right)\sqrt{2}\sigma_t \frac{\phi(z)}{\Phi(z)},$$

and the conditional mean of observed transaction prices is

$$\mathbb{E}[p \mid \text{trade}] = \beta h_b + (1-\beta)h_s - c + \left(\beta - \frac{1}{2}\right)\sqrt{2}\sigma_t \frac{\phi(z)}{\Phi(z)}.$$

Thus, unequal bargaining weights generate a selection term in the valuation equation estimated on traded properties. Base case bargaining power $\beta = 0.5$ is a special case in which this correction is zero.

C.2 Transaction Costs

A positive transaction cost shifts the gains-from-trade threshold:

$$\pi_\tau(b,s) = \mathbb{P}(V_b - V_s > c) = \Phi\left(\frac{h_b - h_s - c}{\sqrt{2}\sigma_t}\right).$$

Thus, $c > 0$ lowers transaction probabilities and shifts transaction prices downward by c in logs. Importantly, when $\beta = 0.5$, a non-zero c does not by itself generate selection bias; it affects trade probabilities and price levels, but not the orthogonality argument.

C.3 Modified Valuation Equation and Estimation Procedure

For candidate values of β and c, define

$$\text{bias}(\Delta) \equiv \left(\beta - \frac{1}{2}\right)\sqrt{2}\sigma_t \frac{\phi(z)}{\Phi(z)},$$

and the adjusted transaction price

$$p^\star \equiv p + c - \text{bias}(\Delta). \tag{10}$$

Then

$$\mathbb{E}\left[p^\star - \beta h_b - (1-\beta)h_s \mid \text{trade}\right] = 0.$$

Therefore, even when $\beta \neq 1/2$, the valuation equation can still be estimated on the traded sample after correcting for selection.

C.4 Valuation Model Estimation Procedure

Initialize $\hat{h}_b^{(0)}, \hat{h}_s^{(0)}$. Now, for iterations $m = 0, 1, 2, \ldots$:

1. Compute $\Delta_i^{(m)} = \hat{h}_{b,i}^{(m)} - \hat{h}_{s,i}^{(m)} - c$ and $z_i^{(m)} = \frac{\Delta_i^{(m)}}{\sqrt{2}\sigma_{t(i)}}$.
2. Compute correction term: $\text{bias}_i^{(m)} = \left(\beta - \frac{1}{2}\right)\sqrt{2}\sigma_{t(i)} \cdot \frac{\phi\left(z_i^{(m)}\right)}{\Phi\left(z_i^{(m)}\right)}$.
3. Form adjusted target price $p_i^{\star(m)} = p_i + c - \text{bias}_i^{(m)}$.
4. Buyer update (hold seller fixed): $y_i^{(b,m)} = \frac{p_i^{\star(m)} - (1-\beta)\hat{h}_{s,i}^{(m)}}{\beta}$, fit $\hat{h}_b^{(m+1)}$ on buyer covariates to predict $y^{(b,m)}$.
5. Recompute correction term and form adjusted target price
6. Seller update (hold buyer fixed): $y_i^{(s,m)} = \frac{p_i^{\star(m)} - \beta\hat{h}_{b,i}^{(m+1)}}{1-\beta}$, fit $\hat{h}_s^{(m+1)}$ on seller covariates to predict $y^{(s,m)}$.
7. Iterate until changes in fitted values are small.

C.5 Empirical Robustness

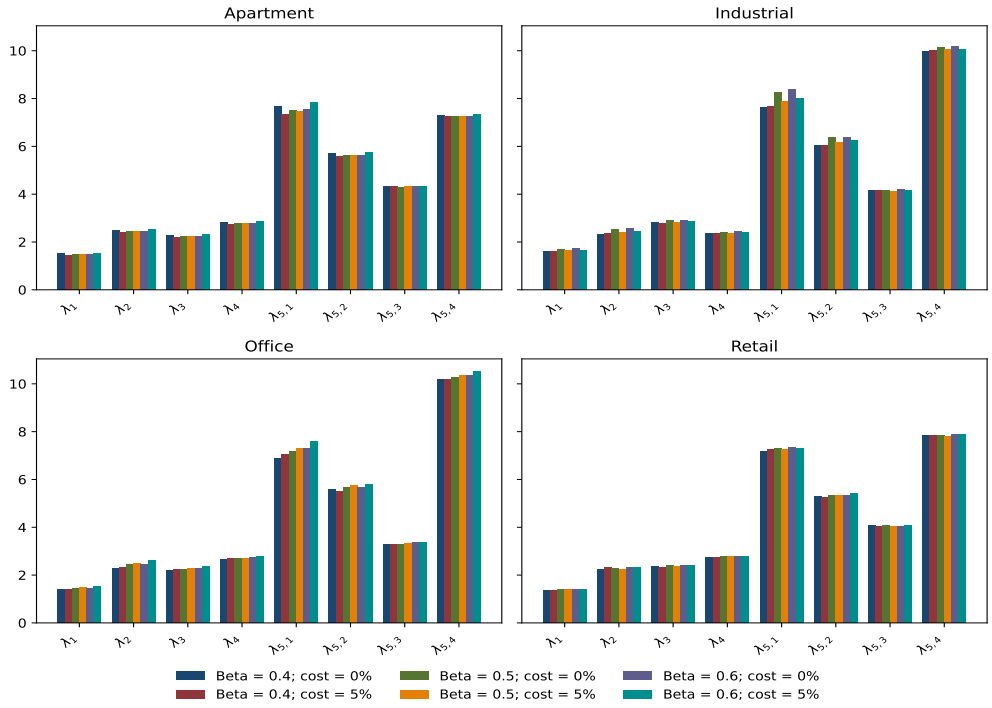
We assess the robustness of the empirical results to alternative values of bargaining power and transaction costs. Specifically, we consider $\beta \in \{0.4, 0.5, 0.6\}$ and $c \in \{0, \log(1+\tau)\}$ with $\tau = 5\%$ [36]. Thus, the robustness exercise considers six specifications in total.

[36]The transaction cost is therefore expressed in the same log-price units as the dependent variable in the valuation model.

In a first step, we re-estimate the valuation model under each of the six specifications using the selection correction described in Section C.4. The resulting buyer- and seller-side valuations are nearly unchanged relative to the baseline case with $\beta = 0.5$ and $c = 0$: across specifications, the fitted valuations are more than 99% correlated with the baseline. This result indicates that the main empirical variation in valuations is not sensitive to moderate departures from equal bargaining power or to the introduction of a modest transaction cost.

In a second step, we re-estimate the meeting model using the corrected valuations from each specification. Figure C.1 reports the average estimated meeting-model parameters. The parameter estimates are very similar across all six cases, indicating that the search component of the model is highly stable. In particular, the ranking of the most important determinants of meetings is unchanged: buyer size and the consideration-set variables continue to play the dominant role, while variation induced by alternative values of β and c is modest.

Figure C.1: Meeting model parameters under alt. bargaining weights and transaction costs



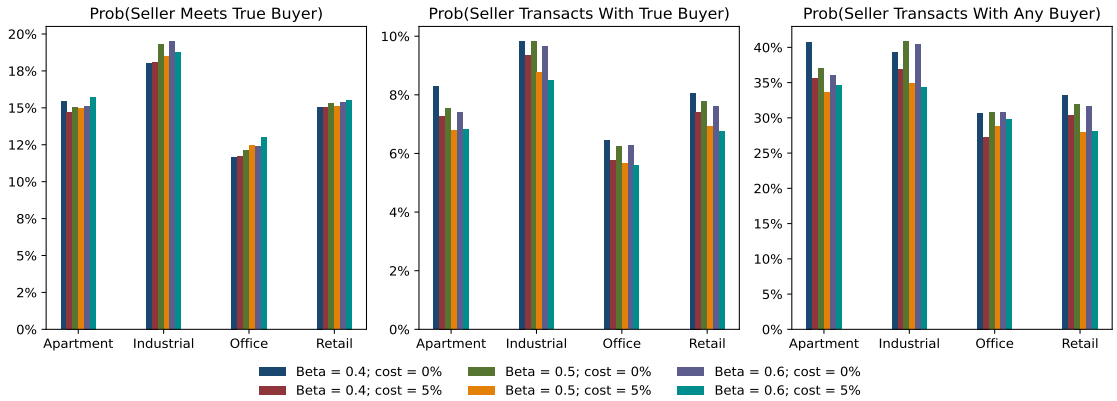
Notes: This figure reports the average estimated meeting-model parameters under six alternative specifications: $\beta \in \{0.4, 0.5, 0.6\}$ and $c \in \{0, \log(1.05)\}$. The parameters are estimated using the meeting model in Section 6. Parameter definitions and interpretations are the same as in Table 3.

Figure C.2 reports three performance measures for the same six specifications: the probability that the seller meets the realized buyer, the probability that the seller transacts with the realized buyer, and the probability that the seller transacts with any buyer. The first is largely unchanged across specifications, again confirming that the matching structure is robust to alternative assumptions on bargaining power and costs.

The main effect of introducing transaction costs can be seen in the transaction probabilities. When $c = \log(1.05)$, the probability that the seller transacts with any buyer falls by about 4 percentage points

in every sector. This is mechanical: a positive transaction cost raises the gains-from-trade threshold from $v_b - v_s > 0$ to $v_b - v_s > c$, so some marginal matches that would have generated trade in the baseline no longer do so. Transaction costs have only a limited effect on who meets whom, but they reduce the fraction of meetings that result in a sale. The decline is economically meaningful but not large, which suggests that the model's implications for matching with characteristics-based consideration sets are robust, while liquidity is naturally more sensitive to trading frictions.

Figure C.2: Matching performance under alt. bargaining weights and transaction costs



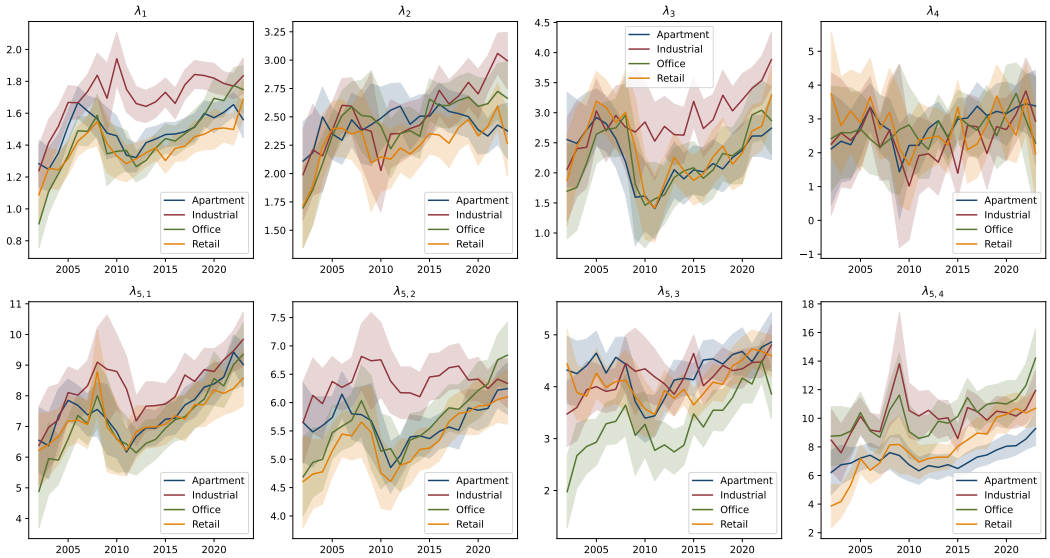
Notes: The figure reports three performance metrics under the same six specifications as in Figure C.1. The left panel shows the probability that the seller meets the realized buyer from among the realized buyer and a negative set of $K = 1{,}000$ potential buyers. The middle panel shows the probability that the seller transacts with the realized buyer. The right panel shows the probability that the seller transacts with any buyer. All probabilities are computed from the meeting model in Section 6.

Overall, the robustness exercise shows that the empirical results are not materially affected by moderate variation in bargaining power or by the introduction of a modest transaction cost. Valuation estimates remain essentially unchanged, meeting model parameters are stable, and the main effect of a positive transaction cost is a lower transaction probability, as expected.

D Meeting Model Appendix

D.1 Meeting Model: Additional Figures and Tables

Figure D.1: Meeting Model Parameters by year



Notes: This figure shows, the estimated values of the annual meeting model parameters (shaded area shows their confidence intervals). The point estimates $\hat{\lambda}_t$ are obtained by maximum likelihood for each year and standard errors are calculated using bootstrap (with 1,000 samples) for each year. The maximum likelihood estimation minimizes the loss function in (9). λ_1 is the coefficient of investor size, λ_2 is the coefficient of the indicator that equals one when the buyer owns more than 2 assets, λ_3 is the coefficient on buyer-seller size difference. λ_4 is the coefficient of the indicator that equals one when the buyer and seller have a prior transactional relationship. $\lambda_{5,1}$ through $\lambda_{5,4}$ are the coefficients of asset, market, sector and NER alignment.

Table D.1: Meeting Model Analysis

Apartments

	λ_1	λ_2	λ_3	λ_4	$\lambda_{5,1}$	$\lambda_{5,2}$	$\lambda_{5,3}$	$\lambda_{5,4}$	1 - π_{m0}′ / π_{m0}	Avg LL
All λs	1.48	2.43	2.25	2.79	7.52	5.63	4.30	7.25		-3.15
λs Excluding λ_1		3.78	1.86	4.57	4.76	3.20	3.97	6.76	55.7%	-4.05
λs Excluding λ_2	1.88		2.12	2.92	7.96	5.45	4.64	6.58	14.3%	-3.42
λs Excluding λ_3	1.41	2.30		3.18	7.90	5.52	4.21	7.33	13.9%	-3.24
λs Excluding λ_4	1.52	2.43	2.39		7.58	5.63	4.33	7.16	9.7%	-3.20
λs Excluding $\lambda_{5,1}$	1.04	2.32	3.15	3.25		5.09	4.26	7.69	40.4%	-3.66
λs Excluding $\lambda_{5,2}$	0.91	1.65	1.98	2.88	6.42		4.70	8.57	55.8%	-3.91
λs Excluding $\lambda_{5,3}$	1.27	2.36	2.16	3.27	7.12	6.09		7.74	33.4%	-3.66
λs Excluding $\lambda_{5,4}$	1.47	2.15	2.42	2.77	7.81	6.09	4.50		18.7%	-3.41

Industrial

	λ_1	λ_2	λ_3	λ_4	$\lambda_{5,1}$	$\lambda_{5,2}$	$\lambda_{5,3}$	$\lambda_{5,4}$	1 - π_{m0}′ / π_{m0}	Avg LL
All λs	1.70	2.51	2.89	2.42	8.25	6.36	4.18	10.14		-3.03
λs Excluding λ_1		3.76	2.84	5.10	2.62	3.30	3.61	9.72	63.8%	-4.38
λs Excluding λ_2	2.03		2.68	2.46	8.57	6.13	3.78	9.78	9.1%	-3.26
λs Excluding λ_3	1.64	2.24		3.09	8.75	6.28	4.10	10.38	23.1%	-3.19
λs Excluding λ_4	1.73	2.50	3.01		8.29	6.37	4.21	9.96	4.9%	-3.07
λs Excluding $\lambda_{5,1}$	1.20	2.32	3.72	2.83		5.83	4.53	10.48	34.9%	-3.51
λs Excluding $\lambda_{5,2}$	1.20	1.46	2.95	2.59	7.31		4.77	12.75	47.8%	-3.80
λs Excluding $\lambda_{5,3}$	1.64	1.70	2.88	2.75	8.54	6.74		10.87	31.0%	-3.49
λs Excluding $\lambda_{5,4}$	1.71	2.24	3.22	2.29	8.55	7.15	4.49		24.0%	-3.39

Office

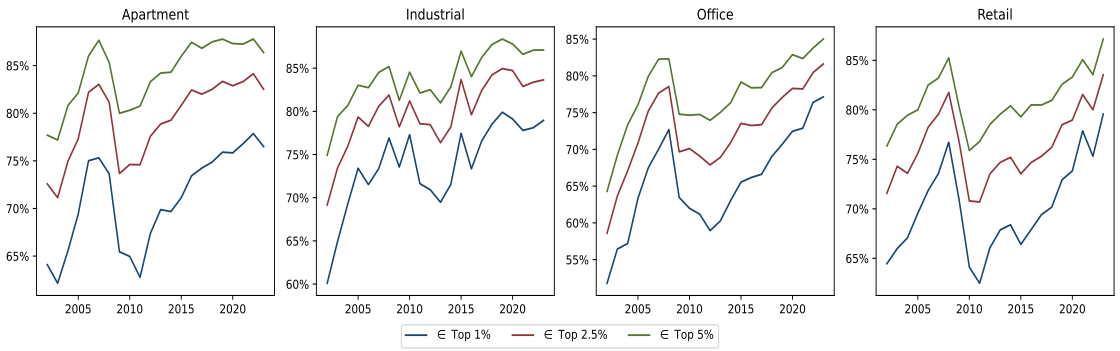
	λ_1	λ_2	λ_3	λ_4	$\lambda_{5,1}$	$\lambda_{5,2}$	$\lambda_{5,3}$	$\lambda_{5,4}$	1 - π_{m0}′ / π_{m0}	Avg LL
All λs	1.43	2.43	2.23	2.68	7.15	5.65	3.29	10.27		-3.26
λs Excluding λ_1		3.52	2.10	4.39	3.64	3.22	2.29	10.17	59.4%	-4.13
λs Excluding λ_2	1.73		2.10	2.77	7.11	5.41	2.86	9.88	19.0%	-3.52
λs Excluding λ_3	1.38	2.28		3.12	7.60	5.55	3.27	10.38	20.4%	-3.36
λs Excluding λ_4	1.46	2.43	2.35		7.20	5.64	3.32	10.11	8.4%	-3.31
λs Excluding $\lambda_{5,1}$	0.96	2.13	2.98	3.06		5.15	2.89	11.05	43.9%	-3.72
λs Excluding $\lambda_{5,2}$	0.92	1.58	2.10	2.71	6.29		3.73	12.36	54.9%	-3.95
λs Excluding $\lambda_{5,3}$	1.26	1.93	2.24	2.90	6.71	6.00		10.89	28.7%	-3.54
λs Excluding $\lambda_{5,4}$	1.42	2.21	2.43	2.59	7.60	6.21	3.57		31.4%	-3.61

Retail

	λ_1	λ_2	λ_3	λ_4	$\lambda_{5,1}$	$\lambda_{5,2}$	$\lambda_{5,3}$	$\lambda_{5,4}$	1 - π_{m0}′ / π_{m0}	Avg LL
All λs	1.39	2.26	2.39	2.76	7.29	5.34	4.08	7.86		-2.97
λs Excluding λ_1		3.36	2.41	4.61	3.20	2.81	3.95	7.28	60.1%	-3.86
λs Excluding λ_2	1.67		2.23	2.85	7.36	4.92	3.85	7.46	10.4%	-3.18
λs Excluding λ_3	1.35	2.07		3.19	7.66	5.27	4.00	7.97	20.9%	-3.09
λs Excluding λ_4	1.42	2.25	2.50		7.34	5.34	4.13	7.73	7.6%	-3.01
λs Excluding $\lambda_{5,1}$	0.95	2.01	2.92	3.14		4.91	4.35	8.12	38.8%	-3.38
λs Excluding $\lambda_{5,2}$	0.97	1.33	2.41	2.76	6.80		4.59	9.79	42.1%	-3.55
λs Excluding $\lambda_{5,3}$	1.35	1.67	2.33	3.38	7.98	5.71		8.23	33.9%	-3.48
λs Excluding $\lambda_{5,4}$	1.38	2.00	2.65	2.65	7.56	5.99	4.30		24.0%	-3.29

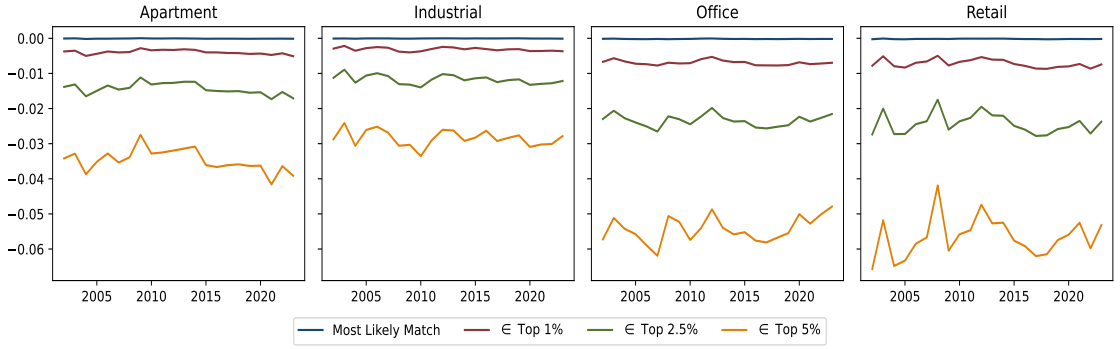
Notes: The table shows the point estimates of the meeting model for each of the four sectors. The last column reports the log-likelihood. The first row in each panel is the baseline model. The successive rows leave out one of the variables in the meeting model, each of which is a special case of the baseline model.

Figure D.2: Meeting Probabilities: True Buyer in Top Predicted Matches



Notes: This figure reports the probability that the true buyer ranks among the top 1%, 2.5%, or 5% of the 1,000 negative sample buyers—i.e., among the top 10, 25, or 50 buyers respectively—based on model-implied meeting probabilities.

Figure D.3: Valuation Gaps: Among Top Predicted Matches

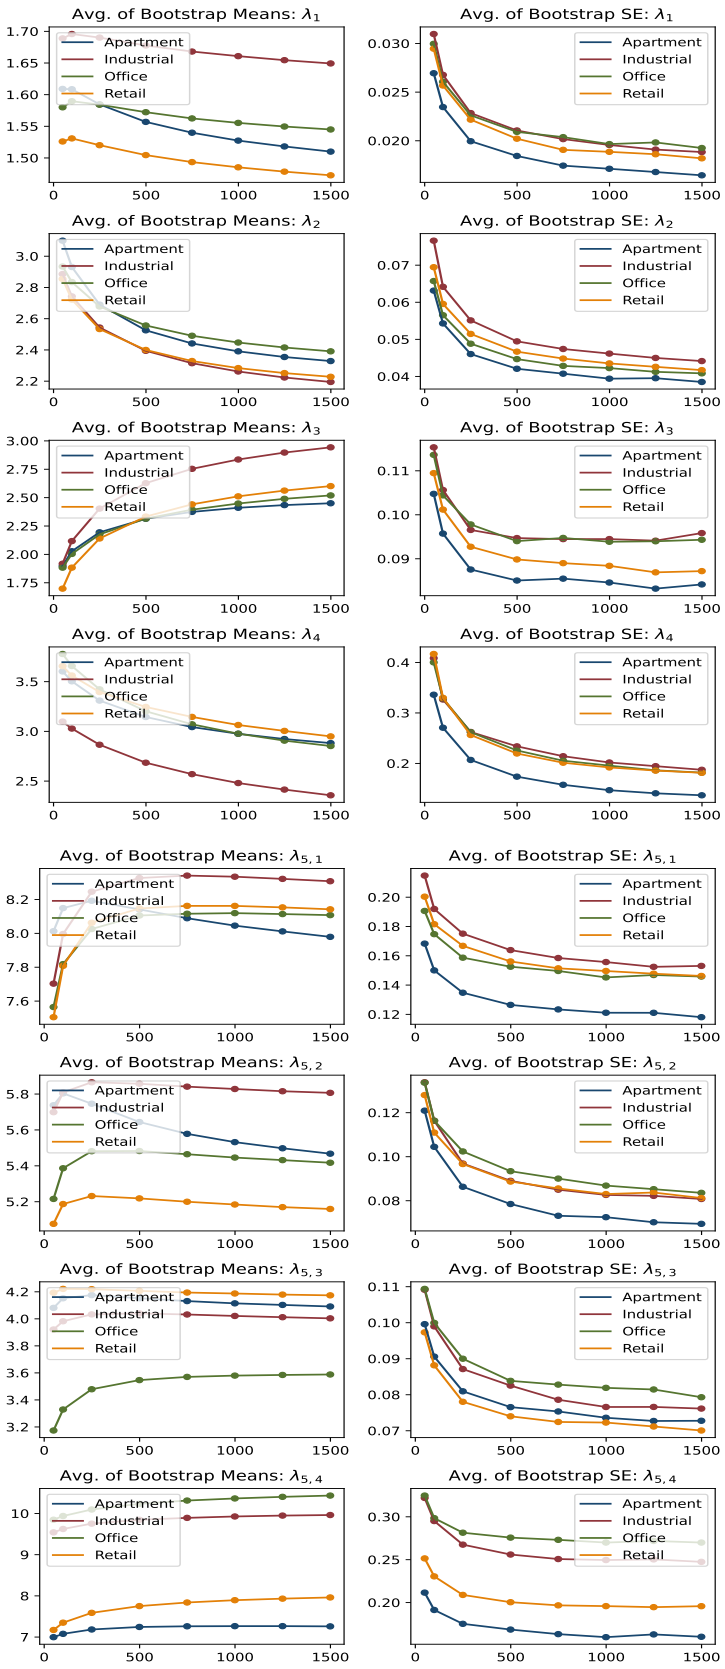


Notes: This figure reports the valuation gap i.e. difference between the log valuation of the buyer and the log valuation of the seller, for the potential buyer with the highest model-implied meeting probability, and the average valuation gaps for potential buyers in the top 1%, 2.5%, or 5% of all the 1000 sampled buyers ranked by modeled meeting probabilities.

D.2 Sensitivity to the Size of the Negative Sample

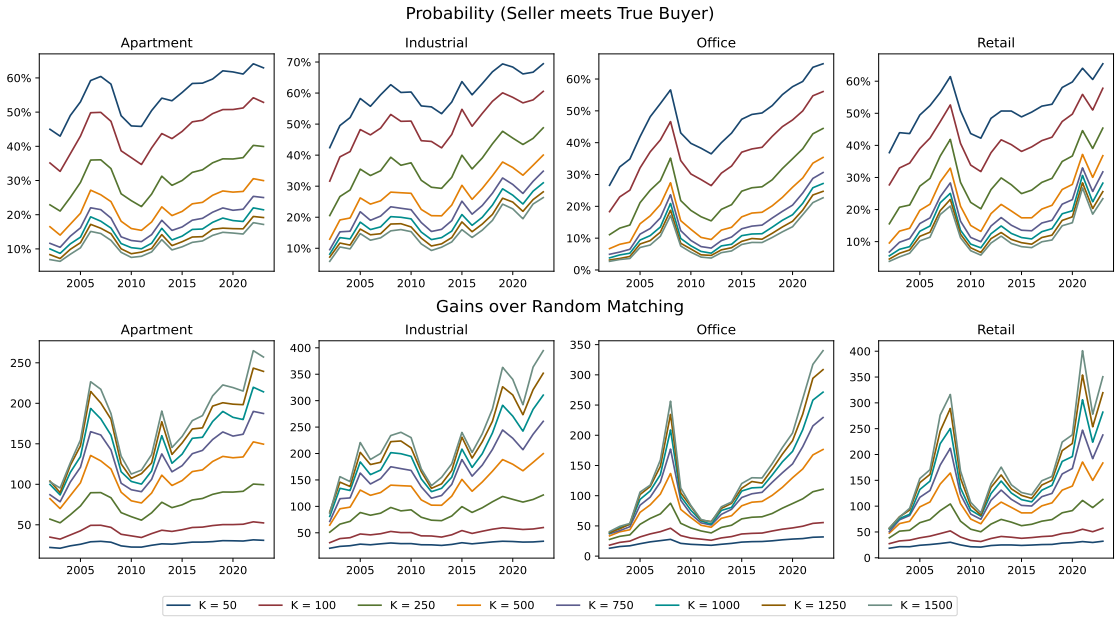
As discussed in Section 3.2, rather than including all potential buyers in the likelihood estimation, we use an important insight from the NLP literature showing it is sufficient to consider a negative set with small cardinality K. In this section, we verify that the parameter values and standard errors stabilize for K = 1,000. We test this by using bootstrap with 1,000 samples for various values of K between 50 and 1,500.

Figure D.4: Bootstrap Parameter Means and Standard Errors



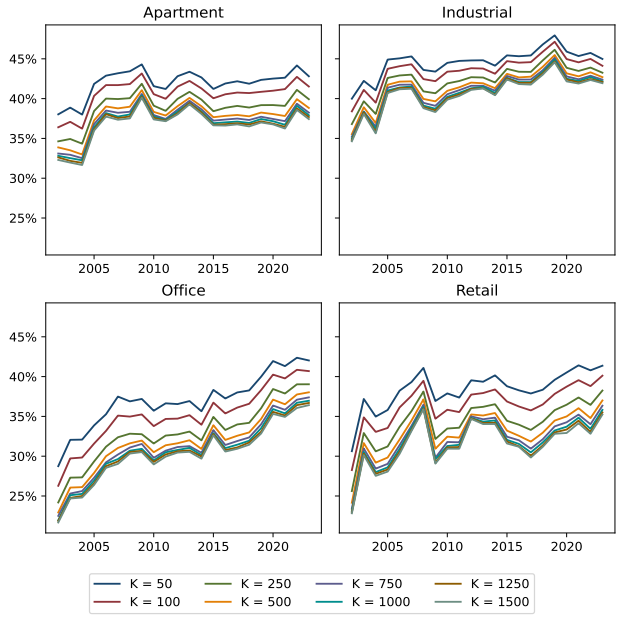
Notes: This figure shows, for each parameter in the matching model, the point estimates (left column) and the standard errors (right column) for varying values of K, the size of the negative sample.

Figure D.5: Probability of Meeting the True Buyer by K



Notes: This figure shows, the likelihood that the model identifies the true buyer from among the set of K potential buyers and the true buyer. The raw probabilities for each sector and K are shown in the top Panel. Since the probability of meeting the true buyer under random search varies with the size of K, in the bottom panel, we normalize the raw probability by the probability under random matching. The corresponding ratio thus gives the gains from search with characteristics-based consideration sets for various sizes of the negative sample. The incremental gains from increasing K beyond 1,000 are small.

Figure D.6: Probability of Transaction (Conditional on Listing) by K



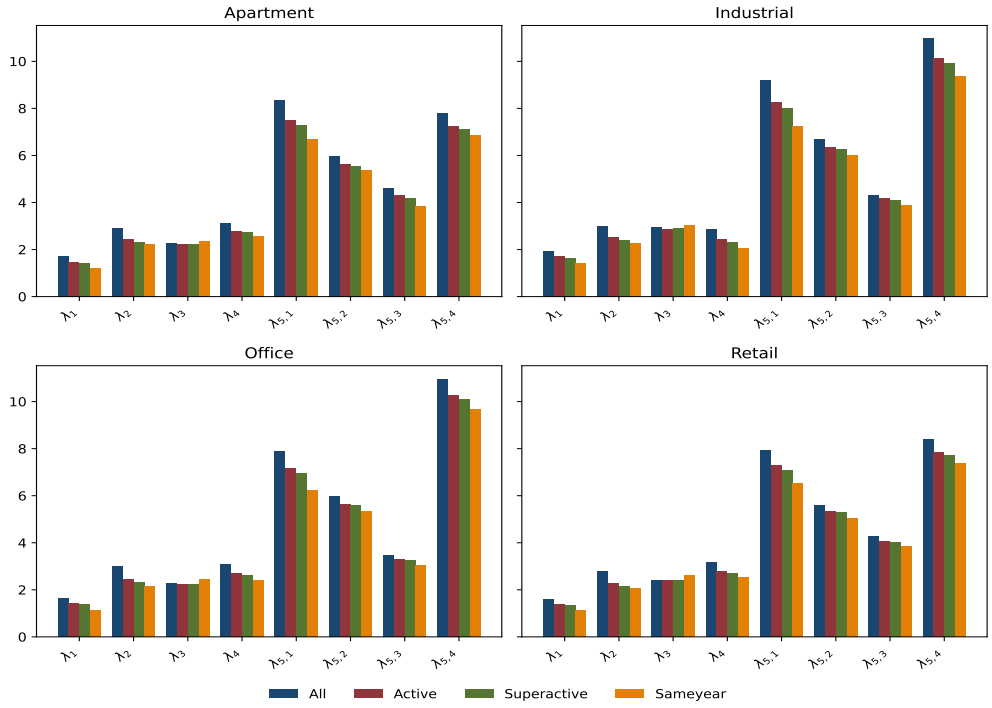
Notes: This figure shows, the model-implied probability of a transaction occurring with any buyer as we vary the size of the negative sample. This stabilizes around K = 1,000.

D.3 Robustness to Alternative Negative-Sample Definitions

In the baseline calibration of the meeting model in Section 6, we restrict the negative sample to investors that have transacted at least once in the previous five years in order to limit the number of inactive investors. We assess the robustness of the results to alternative definitions of this negative sample. We consider four definitions of the negative sample: all potential buyers ("All"); buyers that have transacted at least once in the previous five years ("Active"), which is our baseline specification; buyers that have transacted at least once in the previous three years ("Superactive"); and buyers that transact in the transaction year ("Same Year").

Figure D.7 reports the resulting parameter estimates. The coefficients retain the same signs and are of similar magnitude to those in Table 3, indicating that the main results are not sensitive to the choice of negative sample.

Figure D.7: Meeting Model Parameters under Alternative Negative Samples



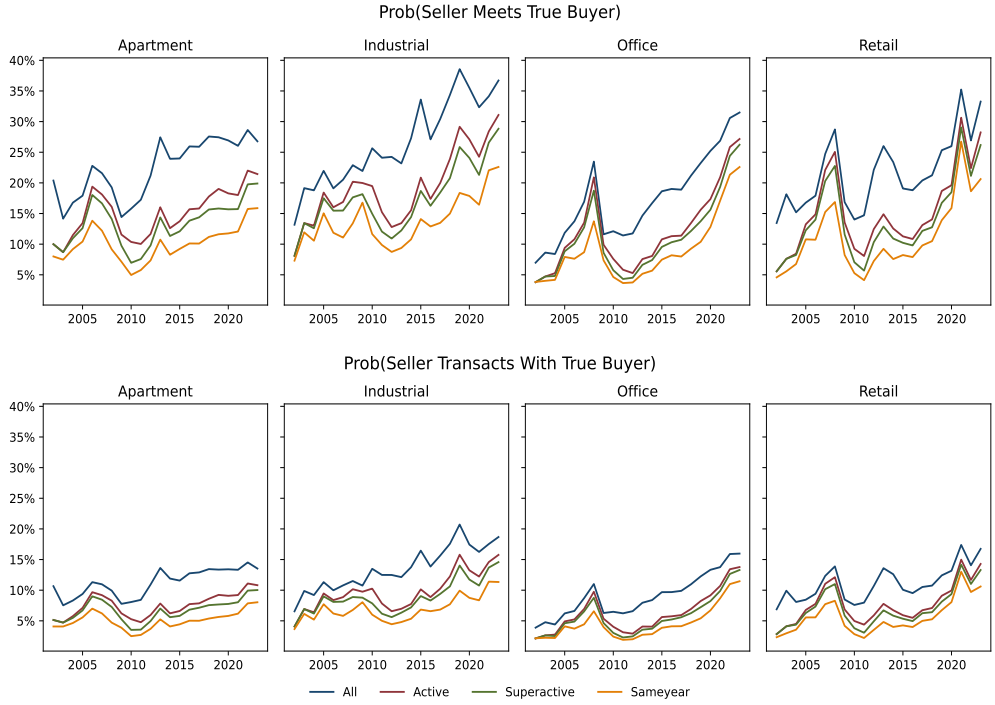
Notes: This figure reports the average estimated meeting-model parameters under alternative negative-sample definitions: all includes all potential buyers; active includes buyers that transacted in the previous five years, as in the baseline specification; superactive includes buyers that transacted in the previous three years; and same year includes buyers that transacted in the year under consideration. The parameters are estimated by maximum likelihood using the loss function in (9). Parameter interpretations are the same as in Table 3.

Figure D.8 illustrates the model's ability to identify the true buyer from the set of potential buyers under this alternative specification. In case of the all buyer sample, the relaxed negative sample includes a larger share of infrequent traders—who are less likely to transact—the model performs better at matching sellers with actual buyers. The probability of correctly identifying the buyer rises to 20-25%, compared to a little over 15% in the base model, as shown in figure 8.

Figure D.8 reports both the probability that the seller meets the realized buyer and the probability that the seller transacts with the realized buyer under the alternative negative-sample definitions. Across specifications, the patterns are very similar. The specification that includes all potential buyers delivers somewhat

higher meeting and transaction probabilities, consistent with the larger negative sample containing more infrequent investors that are unlikely to transact.

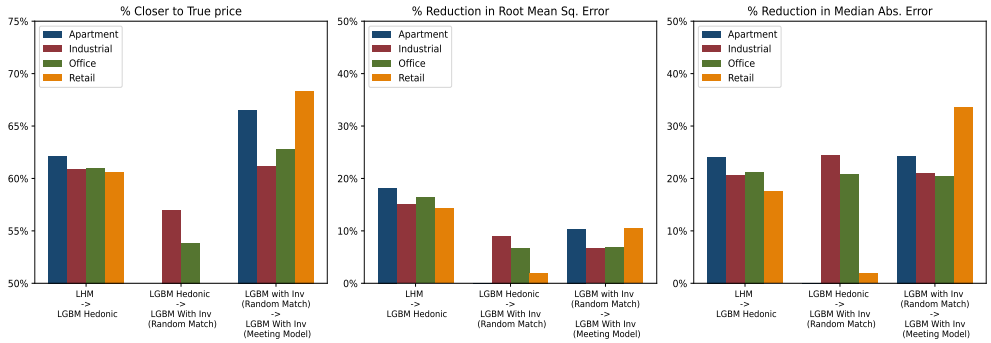
Figure D.8: Matching Performance under Alternative Negative Samples



Notes: The top panel in the figure shows the probability of meeting the actual buyer from among a set of $K = 1,000$ potential buyers (negative set) and the actual buyer, as implied by the meeting model described in Section 6, under alternative negative-sample definitions. The bottom panel shows the probability that the seller transacts with the realized buyer under the same specifications. The four specifications are: all potential buyers, buyers active in the previous five years, buyers active in the previous three years, and buyers active in the transaction year.

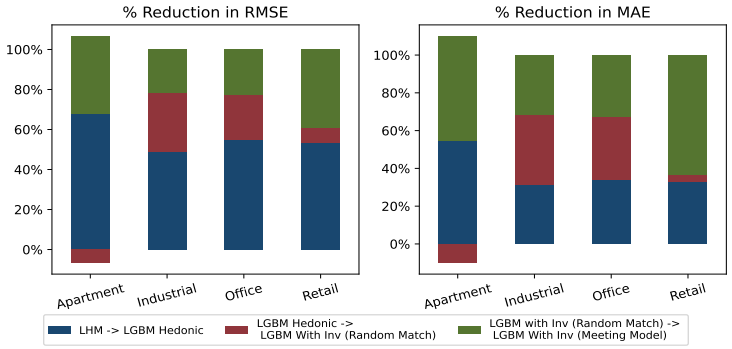
E Out of Sample Prediction Appendix

Figure E.1: LGBM Predictive performance improvement breakdown



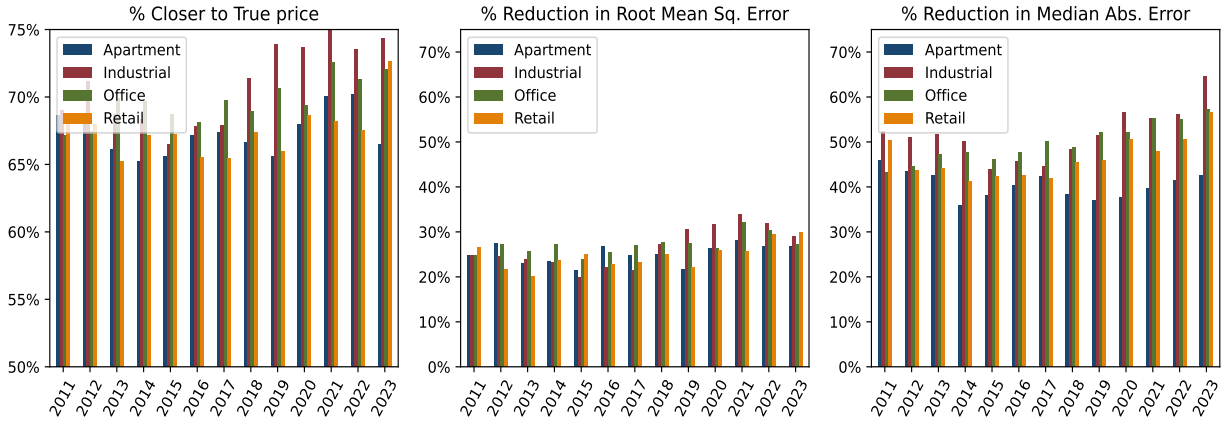
Notes: This figure compares the out-of-sample performance of the LGBM model to the linear hedonic model (LHM), breaking the gains into three components. The left bars show the improvement from moving from the LHM to a non-linear hedonic LGBM. The middle bars capture the additional gain from adding investor characteristics to the LGBM, with potential buyers drawn via random matching. The right bars show the further improvement when buyers are identified using our matching model with characteristics-based search. All models are estimated with information available at time t to predict prices at $t+1$. Missing middle bars indicate cases where random matching used to identify buyers performs worse than the LGBM hedonic model.

Figure E.2: LGBM Predictive performance improvement breakdown



Notes: This figure decomposes the gains in out-of-sample performance of the LGBM model relative to the linear hedonic model (LHM) into three components: (i) the improvement from the LHM to a non-linear hedonic LGBM, (ii) the additional gain from identifying buyers using random matching, and (iii) the further improvement from identifying buyers using matching with characteristics-based consideration sets. All models are estimated using information available at time t to predict prices at $t+1$.

Figure E.3: Out of Sample performance



Notes: This figure shows the relative out-of-sample performance of the LGBM model relative to the linear hedonic model (LHM). The left chart reports the fraction of transactions for which the LGBM prediction is closer to the true transaction price than the LHM prediction. The middle chart reports the percent reduction in root mean squared prediction error of the LGBM relative to the LHM model. The right chart reports the percent reduction in median absolute prediction error of the LGBM relative to the LHM model. In each case, the prediction model is always estimated with information up until time t in order to predict the price at $t+1$. Further, it assumes that (x_{t+1}, z_{t+1}) must both be predicted, and we use our matching model to identify the potential buyers.