

Asset-level risk and return in real estate investments*

Jacob S. Sagi[†]

This version: Feb. 28, 2015

Abstract

Relatively little is known in the academic literature about the idiosyncratic returns of individual real estate investments, though quite a few commercial properties command prices commensurate with the market values of small to medium publicly traded companies. I use purchase and sale data from the National Council of Real Estate Investment Fiduciaries (NCREIF) to compute holding period price-appreciation returns for commercial properties. In stark contrast with liquid asset returns, idiosyncratic drift and volatility estimates diverge as the holding period shrinks. This puzzling phenomenon survives a variety of controls for vintage effects, systematic risk heterogeneity, and sample selection biases. I derive an equilibrium search-based illiquid asset pricing model which, when calibrated, fits the data very well. Thus a structural model of illiquidity seems crucial to a descriptive theory of real estate investment returns. The model can be extended to other illiquid asset classes and used to price derivatives such as debt claims.

Keywords: Commercial real estate, illiquid assets, holding period returns, search, idiosyncratic risk.

*I am grateful to the National Council of Real Estate Investment Fiduciaries (NCREIF) for providing me with the data, and especially to Jeff Fisher and Jeff Havsy for helping me understand it. I also benefited from discussions with Lynn Fisher, Dave Hartzell, and Andra Ghent. All errors are attributable to me.

[†]Kenan-Flagler Business School, UNC Chapel Hill. Email: sagij@kenan-flagler.unc.edu

1 Introduction

Research on real estate investments primarily focuses on aggregated or portfolio-level attributes. Much is therefore known about the systematic returns of real estate assets. By contrast, very little seems to have been written about idiosyncratic or property-specific real estate risk. This is notwithstanding the importance of the subject. Whereas relatively few investors trade real estate indices (using derivative contracts such as total return swaps), a great number of commercial real estate investors hold concentrated portfolios. Geltner et. al (2013) suggest that in 2010 the stock of institutional quality commercial real estate was valued at around \$3 Trillion, with roughly a third being privately held. This suggests that many direct investors in commercial real estate do not hold well-diversified portfolios of properties and could benefit from a quantitative approach to understanding the sources and magnitudes of idiosyncratic risk to which they are exposed. Even ignoring that consideration, commercial real estate properties are often acquired individually and not as portfolios. The absence of a quantitative approach to assessing individual asset risk forces investors to rely on pro-forma analyses that ignore variability or at best subject assumptions to sensitivity tests based on rules of thumb. Finally, commercial real estate debt, totaling more than \$2 Trillion as of 2013, is mostly secured to individual properties and, in the absence of recourse reflects both the systematic and idiosyncratic risks to which the property is exposed.¹ A solid understanding of these risks is important for the efficient pricing of debt obligations and structured debt products such as commercial mortgage backed securities.

How much of a property's risk is diversifiable? Or in other words, how large is the property level idiosyncratic risk? To my knowledge, the only other work that attempts to address the same question is Downing, Stanton, and Wallace (2008) who back out property-level implied volatility from CMBS loan information (much as one might do using a Merton, 1974, model). Their annualized estimates are high and range from 22.7% for apartments to 27.2% for industrial properties. Given that annual volatilities for transaction based property indices (such as the NCREIF TBI) are about 12%, the estimates produced by Downing, Stanton, and Wallace (2008) suggest that the property-specific or idiosyncratic risk comprises between 70% and 80% of the total individual asset risk (or alternatively, the idiosyncratic volatilities range from 20% to 24%). This has important implications for property portfolio diversification, and the pricing and risk assessment of commercial mortgages (held by many regional banks). A potential criticism of the indirect approach taken by Downing, Stanton, and Wallace (2008) is that it is subject to influences outside the property market, such as movement in credit spreads due to changes in liquidity or lending standards. In addition, the loans used in their estimates are securitized and

¹The figure is taken from Table L.217 of the 2013 U.S. Federal Reserve Bank Balance sheet.

may not be representative of the universe of commercial real estate properties. Finally, the implied volatility may depend on the particular model of real estate price appreciation that is used. Model misspecification may conceivably lead to volatility estimates that are not descriptive. An important advantage to the methodology used by Downing, Stanton, and Wallace (2008) is that it provides a forward looking assessment of volatility.

In this paper I use a more direct, though complementary, approach to measuring commercial real estate asset-level return properties. Using data on purchases and sales of properties from the National Council of Real Estate Investment Fiduciaries (NCREIF), I construct holding period price appreciation returns at the individual property level.² While this means that each property bought and sold is typically associated with a single holding period observation, the cross-section of properties is large and, with adequate controls, the holding period return dispersion among properties can be viewed as a “Monte Carlo” simulation of idiosyncratic property risk.³

In Section 2 I apply this methodology and find that the idiosyncratic mean and variance of property log-price appreciation is a linear function of the holding period with a *positive* intercept (see Figure 1). In particular, the slope of the variance is about 0.01 per year held, while the intercept of the variance is 0.04, corresponding to an additional 20% return standard deviation *independent of the holding horizon*. Under the random walk hypothesis (e.g., geometric Brownian Motion as in the Black and Scholes, 1973, model), the slope coefficient corresponds to an annualized idiosyncratic volatility of 10%.⁴ The non-zero intercept, however, is not consistent with frictionless market pricing. Specifically, the data, when extrapolated to arbitrarily short holding horizons, suggests that the idiosyncratic asset drift and volatility are infinite. These findings (and the methodology) are qualitatively similar to what Goetzmann (1993) and Case and Shiller (1987) report in the context of residential properties. They model a property’s observed price as equal to some “true value” plus uncorrelated “noise” and identify the variance intercept with twice the variance of the noise component. The magnitude of the noise in the case under study, nearly four years’ worth of volatility, calls for a deeper investigation into the underlying economics.

I consider two types of explanations for these results: One is that they arise from the illiquidity of the underlying asset, while the other explanation is that the results are spurious. To explore the illiquidity-based explanation, I derive in Section 3 a search-based pricing model (see, for example, Duffie, Garleanu, and Pedersen, 2005, 2007) where investors vary in how they value the income

²Henceforth, I use the term “returns” and “price appreciation returns” interchangeably.

³To motivate this, consider a simple example: Suppose that one observes the purchase and sale of 1000 distinct but similar assets over a similar time horizon. Co-movement, or systematic risk, in the asset returns will contribute little to the dispersion in the holding period returns. Instead, the dispersion will reflect the “typical” idiosyncratic variance of an individual asset and measured with high precision because of the large cross section

⁴This is nearly half of the estimate in Downing, Stanton, and Wallace (2008).

stream from properties, thus leading to potential gains from trade. Market inefficiency arises from the presence of transaction costs and the inability of asset owners to entertain more than one bid per period.⁵ I demonstrate that a random matching and bargaining equilibrium, in which investor private valuations are dynamic but persistent, fits the holding period return data very well (see Figure 3). In the model, an investor that recently acquired a property is unlikely to change his or her valuation over a short period. Thus a short-horizon transaction results only when an investor that recently acquired the asset meets another with a higher valuation (net of the transaction cost). This explains why the mean observed holding period returns extrapolate to a positive value as the holding horizon is taken to zero. When assets are illiquid, an investor will only hold for a short period of time if a much better offer happens to come along. Moreover, the randomness in matching and bargaining impacts observed transaction prices regardless of the horizon, and this explains why the return variance does not vanish as the observed holding period horizon is extrapolated to zero. If transaction costs are eliminated and if all investors can bid for any given property at any time, the model holding period returns converge to the usual random walk result (asset drift and volatility are constant). If each investor's private valuation is serially uncorrelated, then one arrives at the Goetzmann (1993) and Case and Shiller (1987) reduced-form model of "noisy transactions" around a "true" fundamental price. With positive transaction costs, such a model predicts a *negative* intercept for the mean holding period returns and is therefore inconsistent with the data. This underscores the importance of employing a structural model of asset illiquidity in place of one that is reduced form.

When calibrated to the data, the model provides additional insights into the illiquidity of commercial real estate properties. The illiquidity discount, defined using the expected transaction price relative to the average owners' private value of the asset, is calculated to be 11.4%. Moreover, the model predicts that the quarterly turnover in properties would be about 3%. Finally, the model allows one to calculate the probability of sale as a function of price, which can be important for modeling bank "fire sales" and feeds into the pricing of derivative assets such as mortgage loans and mortgage backed securities. In the particular calibration I use, a bank would have to discount the expected transaction price by about 21% to sell the property with probability higher than two thirds.

While an illiquidity-based model fits the data well and can be useful for pricing, there is still the possibility that the holding period return data employed in the calibration suffer from challenges

⁵Real estate properties under contract for purchase are subject to a due diligence period, typically lasting several weeks or months, during which the price can be renegotiated by the prospective buyer and no other offer may be entertained by the seller. Professionals refer to this as "tying up the property". Between the due diligence period and contracted closing date, a period that can also last several weeks to several months, the buyer may back out by forfeiting a deposit of "earnest money" (usually a small percentage of the contract purchased price).

that render the short-horizon inferences (and therefore calibrated model parameters) spurious. For instance, time variation in asset volatility can cause vintage effects that bias the slope and intercept of holding-period return variance. Random exposures to systematic risk can contribute to spurious idiosyncratic variance. More significantly, selection bias can be important to consider in the case of the holding returns of illiquid assets because the decision to sell is endogenous and potentially costly, and may be linked to attributes of the return distribution.⁶ The latter raises concerns that measuring means and variances using holding period returns in the manner just described will yield results that are not representative. Section 4 demonstrates that the naive estimates of idiosyncratic holding period means and variances appear robust to controlling for vintage effects and random exposure to systematic risk. I also rule out several potential selection biases. For instance, it is unlikely that the effect arises because investors prefer to hold safer properties over longer periods of time. In addition, in order to control for the endogeneity of sale decisions, I develop a model of optimal asset disposition and demonstrate that it cannot explain away the effects without also making counter-factual predictions.

This paper contributes to the literature in several ways. First, I document that real estate investment holding period returns deviate from the frictionless market prediction. Second, I establish that this deviation is not spurious or a result of selection bias. Finally, I derive a tractable equilibrium model for illiquid asset pricing that, when calibrated, fits the data very well. Thus, in explaining the non-standard features of observed idiosyncratic real estate returns, illiquidity appears the most plausible driver. In principle, the model and analysis can also be applied to other highly illiquid assets, such as private equity transactions, rarely traded bonds, and complex financial assets.

2 Real Estate Holding Period Returns

Real estate is costly to trade and any given property is traded infrequently. In practice, short-horizon (e.g., quarterly) property “returns” are imputed from appraisals rather than market prices. Appraisals, however, are based on averages of historical prices and of comparable assets, and/or often employ simple rules of thumb (e.g., using a net income multiple). Like any kind of estimate, appraisal-based returns tend to be smoother than the returns they attempt to capture. Thus by their very nature, appraisal based returns can be expected, at best, to understate the volatility of the underlying asset. At worst, appraisals based on random guesses will add idiosyncratic noise unrelated to the actual property-level risk.

⁶See Fisher, Gatzlaff, Geltner, and Haurin (2003), Gatzlaff and Haurin (1997), Goetzmann and Peng (2006), and Hwang and Quigley (2004).

To properly measure attributes of typical property-level returns, one must therefore deploy a different strategy. The first challenge is to obtain actual returns based on prices of purchases and sales.⁷ With real estate assets, save for the rare instance, there is not enough information on repeat sales to calculate time series attributes for a single property. Thus one must rely on a different estimation strategy. The approach I employ is to assume that an observed transaction-based holding period return consists of two random components: One that is idiosyncratic to the property, and one that reflects systematic risk (e.g., a market index). Because one can in principle observe the holding return of the systematic risk benchmark, it is possible (e.g., using a random effects model) to separate the two components given a sufficiently large panel.

In this section, I first describe the dataset of transaction-based property returns. I then provide estimates of the means and variances of idiosyncratic price appreciation log-returns, and find that these exhibit the peculiar property of having a component that is independent of the holding horizon. Section 4 is devoted to confirming the robustness of these findings using either additional tests or more sophisticated approaches.

2.1 The Data

Data of sold and unsold properties comes from the National Council of Real Estate Investment Fiduciaries (NCREIF) and consists of financial and accounting information for sold properties reported by member firms between 1978Q1 and 2014Q1. The database contains, for each property, the acquisition date and price, the sale (or partial sale) date and net price if a sale took place, the property net operating income (NOI), and the total capital expenditures in each quarter, starting from the acquisition quarter (or 1987Q1 if acquisition is earlier) until 2014Q1 (or until the property is sold or otherwise exits the database).⁸ For each property, appraisal-based price appreciation and income returns are reported each quarter. Information about the property location, real estate category, leverage, and member firm is also available.⁹ The following initial filter is applied to properties in the data: The acquisition date must be documented along with a positive purchase price. For sold properties, a sale date and a sale price greater than \$1 must be

⁷Income flow must also be obtained to arrive at total returns. I focus on price appreciation returns for reasons to be stated shortly.

⁸According to the procedures manual, NCREIF members are instructed that partial sales “...may include items such as the sale of an easement, a parcel of land, or a single building in an industrial park” and are to report “the consideration received, less any selling expenses incurred.” When partial sales are reported in the quarter of final disposition, they should coincide with the net proceeds from disposition.

⁹Fields that are available but not used in this study include detailed appraisal valuations and breakdowns of income, operating expenses, and capital expenditures.

reported. In addition, the acquisition year must be equal to or precede the year in which the property first appears in the database by no more than two years. Properties are dropped if reported partial sales are negative or different from the net proceeds reported from the final disposition. I restrict the study to properties belonging to one of the four major type categories: Apartments, Office, Industrial, and Retail. Finally, I drop properties for which the time-series sum of capital expenditures is negative or for which the transaction-based total capital appreciation is negative.¹⁰ This results in 7010 properties that experienced a sale, 3785 properties that exited the database for a reason other than a sale, and 7463 that have not exited.¹¹

The holding period of property i is calculated as $(q_{is} - q_{ia} + 1)/4$ where q_{is} is the sale quarter and q_{ia} is the acquisition quarter. Let $r_{f,t}$ denote the continuously compounded 3-month Treasury Bill quarterly rate. For a property bought at date t and sold at date T , denote the purchase price by P_{it} , capital expenditures at quarter $s \geq t$ by C_{is} , partial net sales at $s \geq t$ by p_{is} , and the net final disposition price by P_T . The treatment of interim cash flows presents a complicating factor in calculating holding-period returns. In particular, the riskiness of the returns will depend on the reinvestment strategy for the interim cash flows. Real estate assets produce a significant amount of interim cash flows and the reinvestment strategy can matter a great deal to the risk-return profile. With liquid assets, it is customary to assume that income is reinvested in the same asset but this strategy is not implementable with real properties. For this reason, I choose to focus on price appreciation returns.¹² To that end, the excess log of price-appreciation return over the holding period is given by,

$$r_i^{\text{App},e} = \ln \left(\frac{\sum_{s=t}^{T-1} p_{is} e^{\sum_{s'=s+1}^T r_{I,s'}} + P_{iT}}{P_{it} e^{\sum_{s'=t}^T r_{f,s'}} + \sum_{s=t}^T C_{is} e^{\sum_{s'=s+1}^T r_{f,s'}}} \right), \quad (1)$$

where $r_{I,s}$ is the quarter s return on investing the proceeds from partial sales. While this expression also appears to depend on a discretionary investment strategy, in practice only 297 properties report partial sales and the reinvestment strategy chosen has negligible impact on the analysis. The default reinvestment return I employ with partial sales is the NCREIF Property Index (NPI) corresponding to the property's major asset type.

¹⁰Companies may report capital expenditures that are allocated but not actually spent. Funds that are not spent should be eventually recorded as negative capital expenditures.

¹¹According to the NCREIF procedures manual, exit other than a sale can also take place due to a transfer of ownership, disqualification as an institutional-quality stabilized asset, a split into multiple properties, the destruction or consolidation of a property, or its foreclosure by a lender.

¹²Another reason to focus on price appreciation returns is that mortgage debt — the most prevalent form of individual asset financing in real estate — is typically secured to the value of the property and not directly to the capitalized value of its past income. Thus a lender would be most interested in the risk characteristics of the asset's price appreciation which is independent of the income reinvestment strategy.

The data contains many outliers and entry errors. To filter these out I first calculate the total holding period return for each sold property assuming income is reinvested in the corresponding NPI. I then repeat the calculation using the reported appraisal-based quarterly returns and calculate the difference between the two returns. Absent errors in the data, this difference should not be excessively large. I therefore drop all properties for which this difference is below the first or above the 99th percentile (-86.15% and 126.53% , respectively).¹³

Although on occasion (as with the previous paragraph) I will employ appraisal-based returns, unless specified otherwise return analyses in this paper are done using the transaction-based holding period returns as calculated in (1). Thus, for the vast majority of properties, there is only one return observation per sold property.

2.2 The basic result

Under the random walk hypothesis, a liquid asset indexed by i should exhibit an excess log price appreciation return of

$$r_i^{\text{App},e} = -(\delta_i + \xi_i)\tau + \beta_i r_m^e + \sigma_i \sqrt{\tau} \varepsilon_i, \quad (2)$$


where τ is the holding period of the asset, r_m^e is the excess log-return on an appropriate benchmark index over the same holding period, δ_i is the dividend or income yield of the asset, σ_i is its volatility, ε_i is a standardized mean-zero random variable that is idiosyncratic and thus uncorrelated across assets, and ξ_i is a non-negative contribution from Jensen's inequality.¹⁴ The random walk hypothesis applied to r_m^e implies that its variance over the holding horizon τ is equal $\tau\sigma_m^2$, where σ_m denotes the market volatility.

For a property purchased at date t and sold at $t + \tau$, it is possible to observe τ and r_m^e . If the model in (2) applies to commercial real estate properties, then ε_i is uncorrelated across properties and, for a fixed τ , a regression of $r_i^{\text{App},e}$ against r_m^e should yield a residual variance and an intercept that are both proportional to τ . In particular, holding period residual variance should vanish with τ as should the holding period adjusted income, $(\delta_i + \xi_i)\tau$.

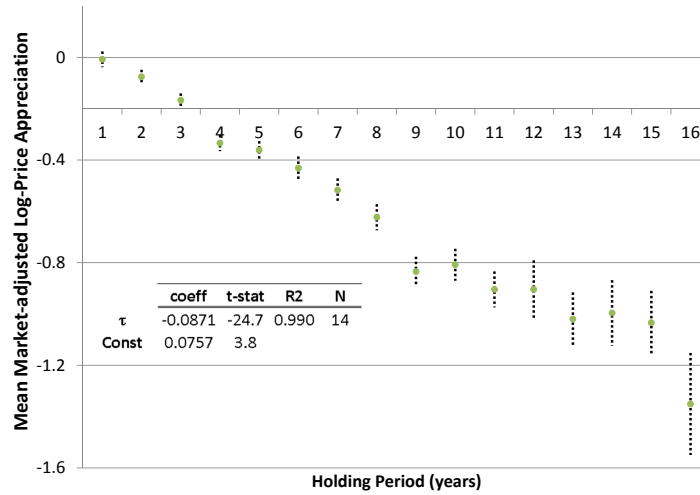
For various observed holding periods, Figures 1(a) and 1(b) plot the residual variances and

¹³Without any censorship, the correlation squared between the two return calculations is 0.49. If each return calculation is censored independently, the correlation squared is 0.69. Censoring the difference, as described in the text, results in fewer dropped observations and a correlation squared of 0.75.

¹⁴In principle, the parameters $\delta_i, \sigma_i, \beta_i$ and ξ_i may vary with time. Thus one can view the constants in (2) as reflecting holding period averages of these quantities. If ε_i and r_m^e are normally distributed, then $\xi_i = \frac{1}{2}(\sigma_i^2 + \beta_i(\beta_i - 1)\sigma_m^2)$.



(a)



(b)

Fig. 1: The top figure depicts point estimates and their 95% confidence intervals for the regression residual variance in Equation (2) (each holding period corresponds to an independent regression). The residual variance is an estimate of idiosyncratic property-level risk. The box in the figure reports a best-fit using weighted least squares. The bottom figure reports the analogous results for the constant regression coefficient term in Equation (2).

intercepts from regressions as in (2).¹⁵ For each property, the benchmark used is the NPI corresponding to the specific property type.¹⁶ In the plots the point coefficient estimate for each horizon is shown as a dot and the 95% confidence interval is depicted using a dashed line. To perform the regressions, the horizon of each property is rounded to the nearest integer.¹⁷ Each plot also reports a weighted least-squares fit to the point estimates, where the inverse of the square of the standard error in the point estimates is used as a weight. The slope of 0.0108 in Figure 1(a) can be associated with an idiosyncratic volatility of $\sigma_i \approx 10\%$ while the slope of -0.087 in Figure 1(b) is an estimate of $-(\delta_i + \xi_i)$. A striking feature of both graphs is that neither vanishes as the maturity approaches zero. The linear fit, in fact, extrapolates the graphs at $\tau = 0$ to positive and highly significant values both economically and statistically. If not an artifact, these would correspond to mean and variance components that accrue to investors independent of the holding horizon. I therefore term these the “time-independent” components of holding period return mean and variance, respectively.

One can view the intercepts of both graphs as representing diverging drift and volatility as the holding horizon tends to zero. These results suggest either that the holding period returns are subject to some bias or that they are distorted by some market frictions at short horizons. In the next section, I offer an illiquidity-based model to explain the presence of the time-independent components. Before doing so, it is worthwhile listing the various concerns that plague the preceding empirical analysis. First, the sample of properties used in the regression for holding period τ versus the sample with holding period $\tau' \neq \tau$ are in general taken from different populations. In particular, the average acquisition years for properties declines with horizon (e.g., 2002 for $\tau = 1$ versus 1983 $\tau = 16$). This raises the possibility that properties with shorter horizons might reflect a more volatile period thereby biasing the slope of the plot in Figure 1(a) down and correspondingly giving rise to the time-independent variance. A second concern is that parameter heterogeneity in (2) may lead to biased estimates of the residual variance.¹⁸ A further concern is that the observed holding period returns suffer from a selection bias because actual holding period returns are only observed when a property is sold, and a sale is an endogenous decision. Two such potential biases that can lead to effects documented in Figure 1 are that

¹⁵To further reduce the impact of outliers at the different horizons, in each regression based on (2) the sample of properties is censored at the first and 99th percentiles.

¹⁶For example, if the property type is “Apartments” then the Apartment NPI is used. The different property type indices are highly correlated, resulting in robust results regardless of the index used.

¹⁷The rounding is done to increase power. Half-integer holding horizons are randomly assigned to the integer above or below with equal probability. The data is quarterly and the actual horizon is uniformly distributed about the rounded horizons depicted, thus rounding should not introduce a bias into the coefficient estimate.

¹⁸When the regression is performed without using the market return as an explanatory variable, the time-independent components of the mean and variance *increase*. This suggests that the effect is not driven by statistical noise in measuring betas.

property risk and holding periods are negatively correlated, and/or that the option to sell a property is exercised contingent on property performance. These empirical issues are addressed in the Section 4.

3 A model of holding-period returns for illiquid assets

In the Black and Scholes (1973) model, the Sharpe Ratio for holding the asset over a period τ is proportional to $\sqrt{\tau}$, and diminishes to zero as τ approaches zero. In Figures 1(a) and 1(b), this ratio is $0.076/\sqrt{0.041} \approx 0.37$ for arbitrarily small τ . However, because commercial real estate assets are highly illiquid and take months to transact, there is no possibility of arbitraging off this fact. A natural conjecture is that illiquidity borne of market frictions may help explain the peculiar properties of observed holding period returns documented in the Figures. This section offers a liquidity-based explanation for the phenomenon, supported by a calibrated equilibrium model.

I consider an approach reminiscent of the search models employed in Duffie, Garleanu, and Pedersen (2005, 2007). There are N infinitely-lived properties, each producing a per-period income of \tilde{d}_t . Each property is held by an investor of type i . To calculate his or her private value for a property at date t , investor i discounts next period's expected income and property (private) value by a factor $\frac{1}{1+r_{i,t}}$. The discount rate $r_{i,t}$ is a Markov process and can take on one of a finite number of values indexed by $a \in \mathcal{A}$. The transition probability matrix for $r_{i,t}$ is $\Pi_{aa'} \equiv \text{Prob}(r_{i,t+1} = a | r_{i,t} = a')$, where $a, a' \in \mathcal{A}$. Denote by π_a^U the unconditional probability that an investor is in state a at date t . It is assumed that π_a^U is non-zero for all $a \in \mathcal{A}$.¹⁹ The process $r_{i,t}$ is assumed to be independent and identical across investors, and independent of \tilde{d}_t .

Each period, a property receives an offer from some, randomly chosen, investor and its owner must decide whether or not to sell at a cost of c_t . It is assumed that the number of investors is sufficiently large so that the probability that a property receives an offer from an investor of type a at date t is π_a^U . The motivation to trade in this model comes from the heterogeneity generated by different private values. The frictions in this model consist of the limited trading opportunities — each period the counter party is a single potential buyer rather than a market of potential buyers — and the cost of transacting a sale. There are no constraints on the number of properties that an investor may hold, and thus the ratio of investors to properties is immaterial.²⁰

¹⁹If there are no absorbing states, the row vector π^U is given by any row of the matrix defined by $\lim_{n \rightarrow \infty} (\Pi^T)^n$.

²⁰The ratio of investors to properties would become important if one wished to make endogenous the number of offers received each period by, say, including a cost to investors of embarking on a search. Here, this is finessed by assuming that each property receives a single offer each period.

A sale takes place if and only if the owner's certainty equivalent (i.e., private value) of the property value, $p_t(r_{i,t})$ is smaller than the investor's certainty equivalent, $p_t(r_j)$ less c_t .²¹ Assume that in the ensuing bargaining the seller receives a random fraction $\tilde{\lambda} \in [0, 1]$ of the gains from trade. Thus, when a sale takes place the transaction price net of costs is $p_t(r_{i,t}) + \tilde{\lambda} \left\{ p_t(r_{j,t}) - p_t(r_{i,t}) - c_t \right\}^+$ (where $\{x\}^+ = \max\{0, x\}$). I assume that $\tilde{\lambda}$ is identically and independently distributed across time and buyers/sellers, and unrelated to $d_t, r_{i,t}$ or c_t . Thus, the property certainty equivalent of a seller of type $r_{i,t}$ is determined by

$$p_t(r_{i,t}) = \frac{1}{1 + r_{i,t}} E \left[\tilde{d}_{t+1} + p_{t+1}(\tilde{r}_{i,t+1}) + \tilde{\lambda} \left\{ p_{t+1}(\tilde{r}') - p_{t+1}(\tilde{r}_{i,t+1}) - \tilde{c}_{t+1} \right\}^+ \right] \quad (3)$$

where $\tilde{r}_{i,t+1}$ and \tilde{r}' correspond, respectively, to the seller's and buyer's random discount rates at date $t + 1$ and are assumed to be independent. While $\tilde{r}_{i,t+1}$ depends on $r_{i,t}$ via the Markov transition process, the distribution of \tilde{r}' only depends on the unconditional probability vector π^U . The owner's valuation of $p_t(r_{i,t})$ depends on the continuation value of holding the income producing property as well as the alternative strategy of selling (at a cost) to a prospective investor. In turn, the continuation value depends on the long-run property value to investors in the market.

Definition. *An equilibrium is a positive and finite random variable $p_t(r_a)$ that solves (3) for every $a \in \mathcal{A}$.*

If \mathcal{A} is a singleton set, then $\tilde{r}_{i,t+1} = \tilde{r}' = r$ and $p_{t+1}(\tilde{r}') = p_{t+1}(\tilde{r}_{i,t+1})$, so that $p_t = \frac{E[\tilde{d}_t + p_{t+1}]}{1+r}$ defines the equilibrium. In other words, if all investors are identical then prices are set as if the market is frictionless and each investor discount all cash flows at the rate r . Liquidity has not role to play in such a market because there are no gains from trade.

It is instructive to consider a situation where the owner faces multiple bidders, each having different valuation and bargaining power. In this case, (3) becomes

$$p_t(r_{i,t}) = \frac{1}{1 + r_{i,t}} E \left[\tilde{d}_{t+1} + p_{t+1}(\tilde{r}_{i,t+1}) + \max \left\{ 0, \tilde{\lambda}' (p_{t+1}(\tilde{r}') - p_{t+1}(\tilde{r}_{i,t+1}) - c_{t+1}), \right. \right. \\ \tilde{\lambda}'' (p_{t+1}(\tilde{r}'') - p_{t+1}(\tilde{r}_{i,t+1}) - c_{t+1}), \\ \left. \left. \tilde{\lambda}''' (p_{t+1}(\tilde{r}''') - p_{t+1}(\tilde{r}_{i,t+1}) - c_{t+1}), \dots \right\}^+ \right].$$

If $c_t = 0$ for all t and the number of independent bidders is sufficiently large (e.g., dense in the support of \tilde{r} and $\tilde{\lambda}$), then the equilibrium will approach one where only the investors with lowest discount rates and least bargaining power will acquire the asset, and the property price will reflect

²¹It is assumed that investors face no funding constraints.

their valuation. This can be viewed as the frictionless limit in which the asset is held by those who derive the most utility from it. If $c_t > 0$, then even with a large number of bidders the equilibrium price will not collapse to a single value but, transactions will be limited to prices within a band determined by c_t .

To investigate the model implications, it is specialized to the case in which the logarithm of asset income is a geometric random walk with iid shocks. Thus $\tilde{d}_{t+1} = d_t e^{(\mu - \frac{\sigma^2}{2}) + \sigma \tilde{\epsilon}_{t+1}}$, with constant drift μ and volatility σ , and where $\tilde{\epsilon}_{t+1}$ is a standard normally distributed random variable. Assume further that the transaction cost is proportional to the level of income, $c_{t+1} = cd_{t+1}$. Conjecture an equilibrium private valuation of

$$p_t(r_a) = d_t Q_a.$$

Then from (3) and the model assumptions, Q_a must solve the linear system of equations:

$$\forall a \in \mathcal{A}, \quad (1 + r_a)e^{-\mu} Q_a = (1 + \sum_{a' \in \mathcal{A}} \Pi_{aa'}^T Q_{a'}) + \bar{\lambda} \sum_{a', b \in \mathcal{A}} \Pi_{aa'}^T \pi_b^U \{Q_b - Q_{a'} - c\}^+,$$

where Π^T denotes the transpose of Π . Assuming such an equilibrium exists, it is possible to refer to an investor of type $r_{i,t}$ via his or her per-dividend private valuation $Q_{i,t} \equiv \frac{p_t(r_{i,t})}{d_t}$.

A transaction takes place at date t if and only if the arriving buyer's private value less the transaction cost exceeds the private value of the seller. This is true if and only if $Q_b - Q_{a'} \geq c$, where Q_b corresponds to the valuation of the prospective investor while $Q_{a'}$ to that of the incumbent owner. Thus, the realization of a transaction is a random variable whose distribution depends on the incumbent owner type. If a trade occurs between an owner of type a and an investor of type b at date t , the observed net transaction price is

$$P_{t,ab} = d_t (Q_a + \tilde{\lambda} (Q_b - Q_a - c)).$$

To analyze holding-period returns, consider a property that at date t is purchased from some arbitrary owner of type \tilde{Q}_O by an investor i of type $Q_{i,t}$. If the property is subsequently sold at date $t + \tau$ to an arbitrary buyer of type \tilde{Q}_S , then the holding period return to investor i is:

$$\tilde{R}_{i,t,\tau} = \frac{Q_{i,t+\tau} + \tilde{\lambda}' (\tilde{Q}_S - Q_{i,t+\tau} - c)}{\tilde{Q}_O + c + \tilde{\lambda} (Q_{i,t} - \tilde{Q}_O - c)} e^{(\mu - \frac{\sigma^2}{2})\tau + \sigma\sqrt{\tau}\tilde{n}},$$

where \tilde{n} is a standard normally distributed random variable and where $\tilde{\lambda}$ and $\tilde{\lambda}'$ are iid. Note that the purchase price is gross of costs but the selling price is net of costs. The logarithm of the

holding period return separates into three sources of risk:

$$\ln \tilde{R}_{i,t,\tau} = \underbrace{\ln \left(Q_{i,t+\tau} + \tilde{\lambda}' (\tilde{Q}_S - Q_{i,t+\tau} - c) \right)}_{\text{Selling shock}} - \underbrace{\ln \left(\tilde{Q}_O + c + \tilde{\lambda} (Q_{i,t} - \tilde{Q}_O - c) \right)}_{\text{Purchasing shock}} \quad (4)$$

$$+ \underbrace{\sigma \sqrt{\tau} \tilde{n}}_{\text{Income shock}} + \left(\mu - \frac{\sigma^2}{2} \right) \tau.$$

The income shock, whose variance grows linearly with the holding period, is independent of the purchasing and selling shocks. The random variable, \tilde{Q}_O depends on the distribution of ownership at date t while \tilde{Q}_S depends on the unconditional distribution of investor types. The two variables are independent of each other and of $Q_{i,t}$ or $Q_{i,t+\tau}$. The latter two variables are related via the Markov process. If τ is sufficiently large, $Q_{i,t}$ will not be related to $Q_{i,t+\tau}$, meaning that the purchasing and selling shocks will be nearly independent. At that point, the contribution of the non-income shocks to the holding period variance will remain constant as the horizon is increased. Even at horizons where the distribution of $Q_{i,t+\tau}$ depends on $Q_{i,t}$ non-negligibly, the selling and purchasing shocks will not generally offset each other and will contribute to the overall variance.

3.1 A three-state example

Suppose that investors' discount rates can take on one of three values, $r_1 \leq r_2 \leq r_3$ and that the transition matrix across the discount rate states is given by

$$\Pi_{aa'}^T = \begin{pmatrix} 1 - x - y & x & y \\ x & 1 - 2x & x \\ y & x & 1 - x - y \end{pmatrix},$$

where $x, y > 0$ and $x + y < 1$. The unconditional distribution of investor types is $\pi^U = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})^T$, which also determines the distribution of \tilde{Q}_S in (4).

To characterize the shock coming from purchasing the asset in (4) one must pin down the distribution of \tilde{Q}_O . This is non-trivial when $Q_{i,t}$ is persistent. For instance, suppose that $y \leq x \ll 1$ so that the unconditional fraction of investors with high discount rates (or low $Q_{i,t}$) is always significant yet the transition into high discount rate states is slow. This means that investors with high discount rates are unlikely to be holding properties in the steady state because the properties will have been bought from them before they ever reached the high discount rate state (and to have purchase them in the first place the investors would have had to be in a higher discount rate state).

Let $\Pi_{ab,t}^S$ correspond to the probability that a property owned by an investor of type a , just before an offer is received at date t , is subsequently owned by an investor of type b just after sales are consummated (or declined) at date t . Based on the modeling assumptions, one can write,

$$\Pi_{ab,t}^S = \begin{cases} \pi_b^U & \text{if } Q_b - Q_a \geq c \\ \sum_{b' \text{ s.t. } Q_{b'} - Q_a < c} \pi_{b'}^U & \text{if } b = a \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where it should be recalled that investor b is drawn from the unconditional distribution of investors. Note also that $\Pi_{ab}^S \equiv \Pi_{ab,t}^S$ is time-independent.

A steady state is achieved once the distribution of ownership across properties is not expected to change. This is defined by the following condition on π_a^O , the unconditional probability that, before the arrival of offers, an arbitrary property is owned by an investor of type a :

$$\pi_a^O = \sum_{b,b'} \pi_b^O \Pi_{bb'}^S \Pi_{b'a}^T. \quad (6)$$

In words, the distribution of ownership changes because the types of current owners (π^O) change through sales (Π^S) and then through the Markov evolution of types (Π^T). Equation (6) asserts that distribution of ownership is stable in the steady state. It also pins down the distribution corresponding to the random variable \tilde{Q}_O in (4).

Suppose that c, r_1, r_2 and r_3 are such that an owner with a higher discount rate will always sell to an investor with a lower discount rate. Then

$$\pi^O = \left(\frac{4x - y + 2}{6x^2 + 2x(6y + 5) + 5y + 2}, \frac{3(2x + y)}{6x^2 + 2x(6y + 5) + 5y + 2}, \frac{3(2x^2 + 4xy + y)}{6x^2 + 2x(6y + 5) + 5y + 2} \right)^T.$$

To calculate the distribution of the random variables in (4), consider an arbitrary property at date t prior to the bid offered to its owner. The owner's type is drawn from π^O . Consider now the randomly drawn investor, of type $Q_{i,t}$ who bids on this property. The investor's type is drawn from π^U . The property will be sold to the investor making the bid only in one of three cases: If the owner is type 2 and the investor is type 1, or if the owner is type 3 and the investor is either type 1 or 2. The probability that the property is sold from a steady state owner of type b to a random investor of type a is $\pi_b^O \pi_a^U$, such that $Q_a > Q_b$.

A subsequent sale at date $t + \tau$ by the same investor that purchased at date t depends on the distribution of offers (determined by π^U) as well as the likelihood that the original investor did not sell at any date strictly between t and $t + \tau$. To pin down the latter, consider the following

modified transition probability matrix defined via.

$$\hat{\mathcal{P}}^T = \begin{array}{c} Q_1 \\ Q_2 \\ Q_3 \\ \text{Sold} \end{array} \begin{array}{cccc} Q_1 & Q_2 & Q_3 & \text{Sold} \\ \left(\begin{array}{cccc} \Pi_{11}^T & \Pi_{12}^T(1 - \pi_1^U) & \Pi_{13}^T(1 - \pi_1^U - \pi_2^U) & \Pi_{12}^T\pi_1^U + \Pi_{13}^T(\pi_1^U + \pi_2^U) \\ \Pi_{21}^T & \Pi_{22}^T(1 - \pi_1^U) & \Pi_{23}^T(1 - \pi_1^U - \pi_2^U) & \Pi_{22}^T\pi_1^U + \Pi_{23}^T(\pi_1^U + \pi_2^U) \\ \Pi_{31}^T & \Pi_{32}^T(1 - \pi_1^U) & \Pi_{33}^T(1 - \pi_1^U - \pi_2^U) & \Pi_{32}^T\pi_1^U + \Pi_{33}^T(\pi_1^U + \pi_2^U) \\ 0 & 0 & 0 & 1 \end{array} \right) \end{array} \quad (7)$$

The matrix $\hat{\mathcal{P}}$ corresponds to the probability that a current owner of type a sells or remains the owner next period and transitions into type a' . The ‘‘Sold’’ state is absorbing. Thus, the probability that investor i still owns the same property after $\tau - 1$ offers, and $Q_{i,t+\tau-1} = Q_{a'}$ conditional on $Q_{i,t} = Q_a$, is given by the aa' element of $(\hat{\mathcal{P}}^{\tau-1})^T$. Let $(\hat{\mathcal{P}}^{\tau-1})_{[3 \times 3]}^T$ be the upper 3×3 submatrix of $(\hat{\mathcal{P}}^{\tau-1})^T$.

Finally, define

$$\hat{\Pi}_\tau^T \equiv (\hat{\mathcal{P}}^{\tau-1})_{[3 \times 3]}^T \Pi^T.$$

The matrix $\hat{\Pi}_\tau^T$ corresponds to the type transition probability for a single investor between t (post-offer) and $t + \tau$ (pre-offer), for those paths in which the property was not sold before $t + \tau$.

Consider the following ‘‘path’’: An arbitrary property is sold to an investor of type $Q_{i,t} = Q_a$ at date t by an owner of type $\tilde{Q}_O = Q_b$, and is held without being successfully sold until the new owner transitions to type $Q_{i,t+\tau} = Q_{a'}$ and subsequently receives a satisfactory bid at date $t + \tau$ from an investor of type $\tilde{Q}_S = Q_{b'}$. From the preceding discussion, we can calculate the unconditional probability of such a path as

$$\pi(\tilde{Q}_O = Q_b, Q_{i,t} = Q_a; Q_{i,t+\tau} = Q_{a'}, \tilde{Q}_S = Q_{b'}) = \pi_a^U \pi_b^O (\hat{\Pi}_\tau^T)_{aa'} \pi_{b'}^U,$$

where $Q_a - c \geq Q_b$ and $Q_{b'} - c \geq Q_{a'}$. The observation of a holding period return is tantamount to observing one of these paths. Thus, the distribution of the purchasing and selling shocks in (4), holding $\tilde{\lambda}$ and $\tilde{\lambda}'$ constant, and conditioning on the observation of a holding period return, is

$$\pi(\tilde{Q}_O = Q_b, Q_{i,t} = Q_a; Q_{i,t+\tau} = Q_{a'}, \tilde{Q}_S = Q_{b'} \mid \tilde{\lambda}, \tilde{\lambda}') = \frac{\pi_b^O (\hat{\Pi}_\tau^T)_{aa'}}{\sum_{Q_a - c \geq Q_b, Q_{b'} - c \geq Q_{a'}} \pi_b^O (\hat{\Pi}_\tau^T)_{aa'}}, \quad (8)$$

where the fact that π^U is uniform was used to simplify the expression. For the specific example considered, equation (8) completes the specification of the probability law for the holding period log-return equation (4).

3.2 Can the model explain the data?

The purchasing and selling shocks in (4) contribute to the holding period return volatility, even if the holding period is short. An important question is whether this contribution can be as large as identified in the data. A more difficult issue to address is whether the observed average holding period return can have a positive intercept as suggested by the data. To see the problem, consider the case where there is no persistence in private values. I.e., each row of Π^T is identical, which can be achieved by setting $x = y = \frac{1}{3}$. In this case, all of $Q_{i,t}$, $Q_{i,t+\tau}$, \tilde{Q}_S and \tilde{Q}_O are independently and identically distributed, and every observed holding period return path in (8) is equally likely. The selling shock component in (4) can be denoted as $\ln \tilde{A}$ while the purchasing shock will contribute $\ln(\tilde{A}' + c)$, where \tilde{A} and \tilde{A}' are identically and independently distributed.²² Because $\tilde{A}' + c$ first-degree stochastically dominates \tilde{A} , $E[\ln \frac{\tilde{A}}{\tilde{A}'+c}] \leq 0$. In other words, without persistence, it is not possible to generate a positive average holding period return as τ approaches zero. Note that this includes the Goetzmann (1993) and Case and Shiller (1987) setting, suggesting that their model is not consistent with the data.

To see why persistence might counteract this result, consider a situation in which types change very slowly. In this case, when an asset is first purchased at date t by an investor with discount rate $r_{i,t} = r_a$, it is most likely that $r_{i,t+1} = r_a$ as well (i.e., $\text{Prob}(r_{i,t+1} \neq r_a) \ll 1$). Because the presence of transaction costs ensure that a sale will not be consummated between two investors with the same valuation, a short holding period is most likely to be characterized by new buyer with a higher valuation (i.e., $r_{j,t+1} < r_a$). Thus the most likely situation for a short holding period, when valuations are persistent, is that the selling price is higher than the recent purchase price. This may appear as a large premium for short holding periods, but the causality is reversed: The apparent “premium” is function of the fact that an investor will only hold for a short period if a much better offer comes along. Despite the intuitive appeal of such an argument, with realistic transaction costs, it is not clear that the effect can overcome the “negative premium” that would be observed in the absence of persistence.

To get a better sense of whether the model delivers plausible magnitudes beyond the qualitative effects outlined above, I undertake a crude calibration exercise where each period represents a quarter. There are nine parameters to consider when attempting to fit to the data, and specifically to Figures 1(a) and 1(b).²³ The transition parameters x and y are set equal to each other, and calibrated so that the observed percentage of properties sold within 6.25 years or less,

²²The random variables \tilde{A} and \tilde{A}' have the same distribution as $Q_{i,t+\tau} + \tilde{\lambda}'(\tilde{Q}_S - Q_{i,t+\tau} - c)$.

²³The empirical analysis in subsequent sections suggests the presence of vintage effects in Figures 1(a) and 1(b). By fitting to the data without explicitly controlling for such time-variation, I take the view that vintage effects have greatest impact on the standard errors but do not impact the average effect.

50.0% in the data, is matched by the model.²⁴ The parameters r_1, r_2 and r_3 are set so as to match the model holding period variance to the data and fixed at 0.0175, 0.075 and 0.125, respectively. The parameters c and μ are calibrated to the observed median proportional selling costs and the average annualized dividend-price ratio (the “cap rate”) in the data (respectively, 0.0212 and 0.0690).²⁵

The holding period excess return attributes depicted in Figures 1(a) and 1(b) are adjusted for market return and volatility. One can likewise state a risk-adjusted excess return version of Equation (4) as follows,

$$\ln \tilde{R}_{i,t,\tau}^{RA} = \underbrace{\ln \left(Q_{i,t+\tau} + \tilde{\lambda}'(\tilde{Q}_S - Q_{i,t+\tau} - c) \right)}_{\text{Selling shock}} - \underbrace{\ln \left(\tilde{Q}_O + c + \tilde{\lambda}(Q_{i,t} - \tilde{Q}_O - c) \right)}_{\text{Purchasing shock}} \quad (9)$$

$$+ \underbrace{\sigma_I \sqrt{\tau} \tilde{n}}_{\text{Idiosyncratic asset shock}} + \left(\mu - \frac{\sigma_I^2}{2} - r_m \right) \tau,$$

where, consistent with the modeling, it is assumed that selling and purchasing shocks are not systematic. The idiosyncratic component of asset volatility, σ_I , is set to $\sqrt{0.0108}$, the estimate in Figure 1(a). The property market rate of return, r_m , is assumed to be 10%, consistent with the time series returns of the NCREIF Property Index and the NCREIF Transaction Based Index.²⁶ Finally, $\tilde{\lambda}$ is assumed to be binomial with an equal likelihood of being 0 or 1.

Table 1 reports the results of the calibration as well as an exercise in which the transition parameter x (which is set equal to y) is increased (thereby reducing the persistence of investor valuation states). In varying these (see Models 2-4), the parameters c and μ are accordingly varied to so that the model proportional selling costs and cap rates match the data. All other variables are fixed. Thus, one should only view $x = y$ as the “free parameter” in the table. The first two columns provide the data-derived confidence interval for income growth, the one-year holding period adjusted excess return moments, and the fraction of properties sold within 6.25 years. The estimated statistics from the data compare very well with the calibrated model values.

²⁴Allowing y to vary freely does not make a significant impact on the calibration. The percentage of properties sold within τ years of purchase is calculated using only properties purchased τ years prior to the end of the sample. The denominator consists of both sold and unsold properties. To calculate this figure in the model, one multiples π^O by the vector created when the fourth entry is eliminated from the last column of $(\hat{P}^\tau)^T$ in (7).

²⁵The selling costs ratio is calculated in the data as the difference between the gross price and the net price, divided by net price. I ignore observations with non-positive costs. In addition, the median proportional selling cost is much more robust to outliers and is therefore employed in the calibration. The cap rate is calculated as the sum of the net operating income in the four quarters prior to the sale divided by the net selling price.

²⁶It is assumed in (9) that every property has a property market beta of 1, and thus the adjusted returns are net of r_m . The choice of $r_m = 10\%$ is made to fit to Figure 1(b) and is well within the 95% confidence interval of the observed index returns.

Table 1: **Calibration of the model to the raw data.**

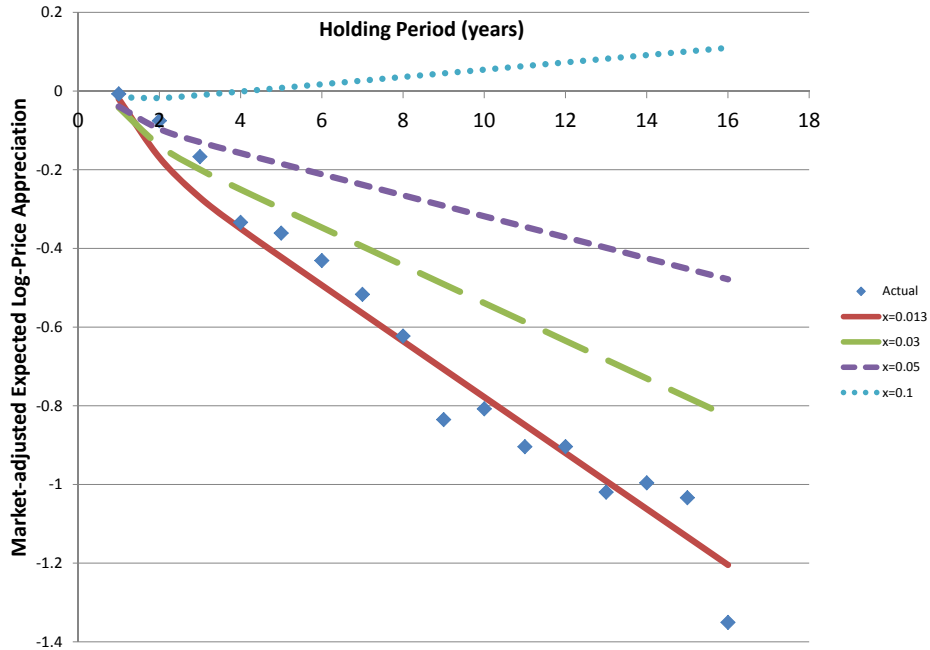
For different parameters, the table compares model-generated and actual one-year holding period adjusted excess return moments, and the fraction of properties sold 6.25 years after purchase. The parameters r_1, r_2 and r_3 are set to 0.0175, 0.075 and 0.125, respectively (each period is a quarter). The idiosyncratic property volatility is set via $\sigma_I^2 = 0.0108$, consistent with Figure 1(a), and the property market return r_M is set to 10% consistent with time-series estimates. The proportional cost parameter c and income growth rate μ are set so as to match observed property income to price ratios (cap rates) and observed proportional selling costs. The transition parameters, x and y are set equal to each other and represent the only “free” parameters varied in the table. The table reports a model liquidity premium calculated as the average observed net selling price divided by the average private valuation of the asset by investors that hold it in the steady state. The “fire sale discount” is one minus the ratio of the lowest private valuation of the asset to the average observed net selling price, and corresponds to the highest transaction price for which a sale is guaranteed.

Quantity	Estimated	Estimated	Calibrated			
	2.5 th percentile	97.5 th percentile	Model	Model 2	Model 3	Model 4
$x = y$ (Qtrly)	NA	NA	0.013	0.030	0.050	0.100
c	NA	NA	1.228	1.228	1.228	1.228
μ (Qtrly)	-0.0022	0.0141	0.0082	0.0140	0.0193	0.0282
1yr Holding Per.						
Adj. Exp. Return	-0.038	0.024	-0.021	-0.043	-0.041	-0.017
Adj. Variance	0.048	0.060	0.051	0.044	0.036	0.024
Fraction Sold						
After 6.25 Years	0.491	0.508	0.497	0.790	0.922	0.991
Quarterly Turnover	NA	NA	0.030	0.065	0.100	0.168
Liquidity Premium	NA	NA	0.114	0.084	0.059	0.021
Fire Sale Discount	NA	NA	0.210	0.188	0.167	0.131

The importance of persistence is demonstrated by the fact that data statistics are generally incompatible with the model once x is increased. Further evidence for the goodness of model fit is given in Figure 2 where average adjusted excess returns are calculated for various holding periods using the various parameterizations in Table 1, and plotted alongside the data (diamonds). The benchmark calibration appears to fit well, while the alternative models do not. Here too, the impact of reducing the model persistence is evident. Figure 3(a) provides a closer look at the goodness of fit of the calibration. While most of the model predictions (squares) lie within the 95-percent confidence interval (dashed vertical line) of the data (circles), there are a few significant departures. As mentioned earlier, the standard errors reflected in the confidence interval might be too tight given the presence of vintage effects. Figure 3(b) depicts the model fit to the variance data reported in Figure 1(a). Though the model is simple, it appears to capture most of the effects exhibited in the data and explain the time-independent variance component and non-standard dependence of average returns on the holding period.

The model allows one to infer important quantities from the calibration that are not readily observable in data. There are several measures of the illiquidity of the commercial property market that can be naturally quantified using the model calibration and are documented in Table 1. The first is the asset turnover. The model predicts a turnover of about 3% per quarter, calculated as the proportion of assets held by an owner that may gain from a sale, multiplied by the probability that he or she will be matched with a suitable buyer. A second measure is the departure of an asset from its “fundamental value”. While it is not always clear how one may objectively define the fundamental value of an asset in the presence of frictions and heterogeneous private values, for the sake of concreteness I define fundamental value to be the average private value of the asset by those who hold it in the steady state. Per unit income, this is $Q' \cdot \pi^O$. This value is high when investor types are persistent because most of the assets will be held by high-value types (92.8% in the case of the calibrated model) who will not sell because they receive no gains from trade. The average observed gross transaction, on the other hand, is likely to be relatively discounted because observed transactions are always with lower-valuation owners who are willing to sell to someone with a higher valuation. In the steady state of the calibrated model, this discount is 11.4%. A third measure, is the discount that must be applied to an asset to achieve a high probability of sale. In the three-state case, one can achieve a probability of sale greater than $\frac{2}{3}$ only by selling the asset at a discount corresponding to the ratio of the lowest private value by prospective investors to the expected net transaction price. In the steady state of the calibrated model, this “fire sale” discount amounts to 21%. The ability to quantify a fire sale discount can be useful in pricing debt contracts, in which the lending institution lacks the expertise to manage the asset while waiting for the right investor. Likewise, it may be possible to fine-tune the model calibration by backing out the implicit fire sale discount in loan spreads for

highly illiquid collateral.

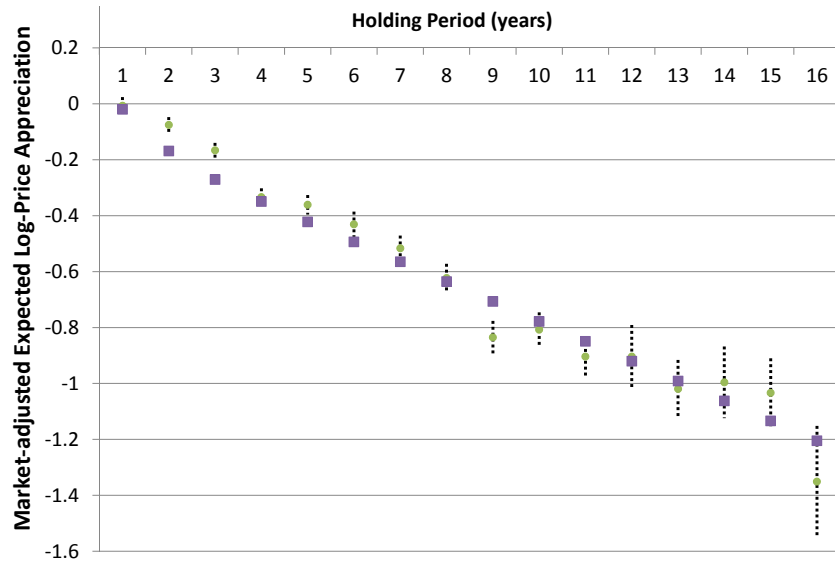


(a)


Fig. 2: Average adjusted excess returns are calculated for various holding periods using the various parameterizations in Table 1, and plotted alongside the data (diamonds).

4 Robustness of empirical findings

In this section, I argue that the economically and statistically significant time-independent return mean and variance components are not artifacts. First I show that the effects remain at roughly the same magnitude even after controlling for property vintage. Next, I develop an econometric model that accounts for possible parameter heterogeneity and demonstrate that the main effects persist. Then, I demonstrate that property-level measures of riskiness have little ability to predict holding horizons, thus dispensing with one potential form of selection bias. Finally, while I find some evidence for the endogeneity of sales and that this is linked to poor performance, there is also evidence that the impact of this is negligible in practice.




(a)



(b)

Fig. 3: The top figure plots the model predictions (squares) for expected adjusted excess returns using the $x = 0.013$ calibration in Table 1. The data (circles) along with 95-percent confidence intervals (dashed vertical line) are also depicted. The bottom figure depicts the model fit to the variance data reported in Figure 1(a).

Table 2: The histogram reports entries of properties into the database (i.e., “vintages”). The table reports the average vintage for various holding periods.



Holding Period	Average Vintage	Number of Props
1	2002	724
2	2001	770
3	2001	905
4	2000	731
5	2000	700
6	1997	593
7	1997	589
8	1994	477
9	1990	277
10	1991	236
11	1990	174
12	1990	125
13	1988	108
14	1987	75
15	1986	63
16	1983	34

4.1 Controlling for vintage

Table 2 lists the average vintage (roughly, the acquisition date) for sold properties in the data by holding period and presents a histogram of their initial year in the database. If property market volatility varies through time, because vintage is not uniform across holding horizons, it stands to reason that holding period variance will also vary with time.

To control for a vintage effect I employ a procedure motivated by the following intuition. Consider a property purchased in 1998 and held for eight years, and suppose one can randomly match it with two properties held for four years, one purchased in 1998 and another purchased in 2002. By “rolling” the investment from the first to the second of the four-year properties one effectively creates an eight-year property investment initiated in 1998. Assuming all three properties are chosen at random and that a property’s risk characteristics are not related to its holding horizon, the model in (2) predicts no difference between the actual and the synthetic eight year investment. For various values of k and τ , a linear program is used to match the largest possible number of properties having holding period $k \times \tau$ with k properties of (non-overlapping) holding periods τ . When the matching entails a strict subset of eligible properties, the subset is chosen randomly. To increase the accuracy of the estimates, and allow for all property returns to enter the estimation (through randomization), the procedure is repeated 100 times and the resulting statistics are averaged.²⁷

Table 3 documents the results of the vintage matching analysis with N corresponding to the number of “longer maturity” matched properties. The difference in regression residual variances and intercepts (or “ α ’s”) is normalized by $k - 1$ to account for the fact that each time a property is rolled over the time-independent variance and mean should be accrued. The t -statistics, which reflect pooled values of the coefficients and their standard errors, are likely to be conservative. In every case, regardless of vintage, a sequential multi-property investment strategy yields a higher residual variance than a matched single-property investment. In seven of ten cases this difference is statistically and economically significant, and comparable with the time-independent variance estimate of 0.041 in Figure 1. The sum of squared t -statistics under the null that the difference is zero is resolutely rejected by a Chi-squared test. Weighted least-squares yields an estimate of the time-independent variance of 0.0404 with a standard error of 0.006, suggesting that the estimate in Figure 1 does not suffer from bias because of a vintage effect. The results for the time-independent mean component in (2) are less conclusive. The null of zero alpha difference is only rejected at the 5% level by a Chi-squared test. The 0.0314 weighted least squares estimate of

²⁷I report the average coefficients and their average standard errors. The standard errors provide a conservative estimate of the error in the pooled coefficient estimates.

Table 3: To control for vintage, properties with holding period $k \times \tau$ are matched with k properties of non-overlapping holding period τ so that the total horizon of the latter matches that of the former. The Equation (2) regression residual variance of the $k \times \tau$ horizon properties are compared with those of the compounded returns of the matched properties. The normalized variance difference measures the time independent variance component while controlling for vintage (likewise for the normalized difference in regression intercepts). The GLS estimate is the error-weighted mean of the 10 variance (or alpha) differences.

k	τ	N	Avg Vintage	(Var Diff)/($k - 1$)	t -stat	(α Diff)/($k - 1$)	t -stat
2	1	321	2002	0.056	5.37	-0.004	0.09
2	2	315	2000	0.054	5.14	0.099	-2.27
2	3	240	1999	0.056	4.30	0.068	-1.27
2	4	237	1998	0.033	2.61	-0.023	0.37
3	3	126	1997	0.029	2.59	0.095	-2.08
2	5	140	1994	0.022	1.23	0.008	-0.11
2	6	76	1993	0.033	1.24	-0.055	0.58
2	7	58	1988	0.003	0.11	-0.183	2.23
3	5	39	1988	0.015	1.06	0.020	-0.43
2	8	31	1982	0.134	2.49	0.067	-0.59

Prob $\chi(\text{Var Diff} = 0)$	4.7E-17
Prob $\chi(\alpha = 0)$	0.0474

Statistic	GLS Estimate	se
Var Diff	0.0404	0.0060
α	0.0314	0.0226

the time-independent premium, while not statistically distinguishable from the 0.0757 estimate in Figure 1 (difference t -test of 1.47), is also itself not statistically different from zero.

It is worth noting that there is a correlation of 50% between the time-independent variances and the time-independent means as measures in Table 3. This supports the liquidity shock interpretation. In addition, the outlier in each case corresponds to the $\tau = 7$ and $k = 2$ matched properties. Removing this observation changes the least squares estimates of the time-independent variance and mean to 0.041 and 0.042, with the latter becoming statistically significant ($t = 2.46$). In summary, the time-independent variance is robust to controlling for vintage while the time-independent mean component may not be significant after controlling for vintage using matched holding periods.

4.2 Controlling for random effects

In practice, the appropriate market benchmark is not known and the factor loading β_i is time-varying as is the income yield, δ_i . Addressing all such realities is difficult enough when the underlying asset is relatively liquid and near impossible in the case at hand where only a single return (over an endogenously determined horizon) is available for each property. Some headway, however, can be made by making assumptions about the nature of the property-level return process. Specifically, assume that for each property β_i and δ_i in (2) are “sampled” from some distribution independently of the holding horizon, the market returns, and the property-level shock ε_i .²⁸ In addition, to account for the time-independent components I add to (2) a random variable that is also assumed to be idiosyncratic to the property. The resulting model for property price appreciation excess log-returns can be rewritten as

$$\tilde{r}^{\text{App},e} = \tilde{\alpha}_0 + \tilde{\alpha}_1\tau + \tilde{\beta}r_m^e + \sigma\sqrt{\tau}\tilde{\varepsilon}, \quad (10)$$

where $\tilde{\varepsilon}$ is a standard normal random variable. The new term, $\tilde{\alpha}_0$, captures the one-time (i.e., horizon independent) risk component. The $\tilde{\alpha}_1$ term captures heterogeneity in properties’ δ_i ’s and accrues with time. In assuming that $\tilde{\beta}$ is stochastic, one allows for different sensitivities to systematic risk across properties. Finally, $\tilde{\varepsilon}$ corresponds to a standardized shock to fundamental property value that is independent across properties.²⁹ An observation of a property’s holding period returns is assumed to be an independent draw from the distributions of $\tilde{\alpha}_0$, $\tilde{\alpha}_1$, $\tilde{\beta}$ and $\tilde{\varepsilon}$, and further assumed to be independent of τ and r_m^2 . Furthermore, I assume that while the

²⁸In the next subsection, I address the possibility the holding period is related to the risk characteristics of the property.

²⁹With essentially only one observation per property, heterogeneity in σ cannot be identified separately from $\tilde{\varepsilon}$. Likewise, one cannot separately identify heterogeneity in the mean of $\tilde{\alpha}_0$ (i.e., the time-independent mean component).

property's beta and mean can be known by the manager or investors when a property is purchased, the same is not true for the liquidity shock and return residual. Thus $\tilde{\beta}$ and $\tilde{\alpha}_1$ may be correlated with each other but are independent of $\tilde{\alpha}_0$ and $\tilde{\alpha}_1$ (the latter of which can also be correlated with each other).

The identification strategy for the joint distributional properties of $\tilde{\alpha}_0, \tilde{\alpha}_1, \tilde{\beta}$ and $\sigma\tilde{\varepsilon}$ loosely follows Beran and Hall (1992) who advocate a sequence of n regressions to estimate the first n moments of these distributions. The idea is as follows. Normalize (10) by $\sqrt{\tau}$ and write it as

$$\frac{\tilde{r}^{\text{APP},e}}{\sqrt{\tau}} = \alpha_0 \frac{1}{\sqrt{\tau}} + \alpha_1 \sqrt{\tau} + \beta \frac{r_m^e}{\sqrt{\tau}} + \left(\tilde{\varepsilon}_0 \frac{1}{\sqrt{\tau}} + \tilde{\varepsilon}_1 \sqrt{\tau} + \tilde{\varepsilon}_\beta \frac{r_m^e}{\sqrt{\tau}} + \sigma\tilde{\varepsilon} \right),$$

where $\tilde{\varepsilon}$ and the $\tilde{\varepsilon}$'s are independent of βr_m^e , $\frac{1}{\sqrt{\tau}}$ or $\sqrt{\tau}$.³⁰ Under the structural assumptions made earlier, one can therefore consistently estimate α_0, α_1 and β via an OLS regression of $\frac{\tilde{r}^{\text{APP},e}}{\sqrt{\tau}}$ against $\frac{1}{\sqrt{\tau}}, \sqrt{\tau}$ and $\frac{r_m^e}{\sqrt{\tau}}$. Denote the resulting residuals as \hat{z} . The latter are consistent (but not efficient) estimates of $\left(\tilde{\varepsilon}_0 \frac{1}{\sqrt{\tau}} + \tilde{\varepsilon}_1 \sqrt{\tau} + \tilde{\varepsilon}_\beta \frac{r_m^e}{\sqrt{\tau}} + \sigma\tilde{\varepsilon} \right)$. The squared residuals can be written as

$$\begin{aligned} \hat{z}^2 &= E[\hat{z}^2] + \tilde{U}, \\ &= \frac{1}{\tau} \sigma_0^2 + \tau \sigma_1^2 + \frac{(r_m^e)^2}{\tau} \sigma_\beta^2 + \sigma^2 + \frac{2}{\sqrt{\tau}} \sigma_{0r} + 2r_m^e \sigma_{1\beta} + O\left(\frac{1}{N}\right) + \tilde{U}, \end{aligned} \quad (11)$$

where $\sigma_x^2 = \text{Var}[\tilde{\varepsilon}_x]$, $\sigma_{0r} = \sigma \text{Covar}[\tilde{\varepsilon}_0 \tilde{\varepsilon}]$, and $\sigma_{1\beta} = \text{Covar}[\tilde{\varepsilon}_1 \tilde{\varepsilon}_\beta]$, while \tilde{U} is a zero-mean random variable uncorrelated with $\frac{1}{\tau}, \tau, \frac{(r_m^e)^2}{\tau}, \frac{1}{\sqrt{\tau}}$ or r_m^e . The $O\left(\frac{1}{N}\right)$ term signifies an expression that is inversely proportional to the number of observations and arising from the estimation error in the coefficients of first regression. For large N , the distributional parameters, $\sigma_0^2, \sigma_1^2, \sigma_\beta^2, \sigma^2, \sigma_{0r}$ and $\sigma_{1\beta}$ can again be estimated via OLS. In principle, one can continue this process and estimate higher joint moments of the error distributions. In practice, the power can rapidly dwindle after the second pass (i.e., beyond second moments). To avoid collinearity and identify the estimated moments, it is important for there to be sufficient dispersion in the explanatory variables. This is a particular problem in the data and motivates the choice of normalization by $\sqrt{\tau}$. Normalizing by $\sqrt{\tau}$ leads to substantially less collinearity in the second pass regression than either dispensing with a normalization or normalizing by τ . Despite that, there is still sufficient collinearity to make identifying σ_0^2, σ^2 and σ_{0r} difficult unless one further omits the σ_{0r} interaction terms. The actual estimation procedure attempts to correct for the inefficiency of the coefficient estimates from the first pass arising from the heteroskedasticity of the residual (see, for instance, Hildreth and Houck, 1968; Swamy, 1970; Raj, Srivastava, and Ullah, 1980), and proceeds as follows:

1. Regress $\frac{\tilde{r}^{\text{APP},e}}{\sqrt{\tau}}$ against $\frac{r_m^e}{\sqrt{\tau}}, \frac{1}{\sqrt{\tau}}$ and $\sqrt{\tau}$ (without a constant) and calculate the residuals, \hat{z} .

³⁰The reason for the choice of normalization has to do with reducing collinearity in the sequence of regressions to be described.

2. Regress \hat{z}^2 against $\frac{1}{\tau}$, τ , $\frac{(r_m^e)^2}{\tau}$, $\frac{1}{\sqrt{\tau}}$, r_m^e and a constant. Calculate \hat{z}_p^2 , the predicted value of \hat{z}^2 using the regression coefficient estimates. Thus \hat{z}_p^2 is an estimate of the property-specific residual variance from the first pass regression.³¹
3. Correct for heteroskedasticity in the first regression by regressing $\frac{\tilde{r}^{\text{APP},e}}{\sqrt{\tau}\hat{z}_p}$ against $\frac{\beta r_m^e}{\sqrt{\tau}\hat{z}_p}$, $\frac{1}{\sqrt{\tau}\hat{z}_p}$ and $\frac{\sqrt{\tau}}{\hat{z}_p}$ (again, without a constant).
4. Use the heteroskedasticity-adjusted first-moment coefficient estimates to recalculate \hat{z}^2 and regress against $\frac{1}{\tau}$, τ , $\frac{(r_m^e)^2}{\tau}$, $\frac{1}{\sqrt{\tau}}$, r_m^e and a constant to obtain estimates of the heteroskedasticity adjusted second-moments. Rerun this regression as needed to adjust for collinearity.

For each integer holding horizon, the return data is censored at the first and 99th percentiles, as was done in the regressions behind Figure 1.

Table 4 reports the results. Comparing the first two columns in each panel, adjusting for heteroskedasticity makes little impact on the coefficient estimates. The average beta of the properties, as might expected, is close to one. The α_1 component, corresponding to an average continuously compounded mean yield (roughly, $\delta - \xi_i$ in (2)) is statistically indistinguishable from the less sophisticated τ coefficient estimate from Figure 1b. The same is true for the α_0 estimate, which is the time-independent mean component. The second pass regression exhibits collinearity which is eliminated if the correlation between the shock $\tilde{\epsilon}_0$ and the residual return shock $\tilde{\epsilon}$ is set to zero (i.e., $\sigma_{0r} = 0$). Note that this is consistent with the statistically insignificant role that this interaction plays in the second pass regression. Once this is done, the time-independent variance, σ_0^2 and idiosyncratic return volatility, σ , are well identified, relatively stable, and also statistically indistinguishable from the estimates in Figure 1a. The third and fourth regressions in Table 4 panel B imply that there is little to no identifiable variation in betas and deltas across properties.³² In summary, controlling for heterogeneity makes little to no impact on detecting and measuring the presence of the time-independent return components in holding period returns.

A side benefit of the methodology used in the subsection is that it can be extrapolated to incorporate additional controls and factors influencing α_0 , σ_0^2 and σ^2 . This will be further pursued in Section 5

³¹Note that collinearity should not impact the estimate of \hat{z}_p^2 .

³²One can constrain the variance estimate σ_β^2 to be positive using non-linear least squares. All this achieves is to drive the estimate of σ_β^2 to zero.

Table 4: Panels A and B report the first and second stage regressions for estimating the model with random effects in Equation (10). The estimation procedure is described in the text. The time-independent variance component is σ_0^2 in Panel B while σ^2 corresponds to the volatility component of idiosyncratic risk.

Panel A.

	(1)	(2)
	$\frac{\hat{r}^{\text{APP},e}}{\sqrt{\tau}}$	Adj $\frac{\hat{r}^{\text{APP},e}}{\sqrt{\tau}}$
β	0.960 (48.80)	0.946 (52.34)
α_1	-0.0875 (-63.62)	-0.0864 (-65.20)
α_0	0.0787 (12.37)	0.0748 (10.54)
Observations	6287	6287

Panel B.

	(1)	(2)	(3)	(4)	(5)
	\hat{z}^2	Adj \hat{z}^2	Adj \hat{z}^2	Adj \hat{z}^2	Adj \hat{z}^2
σ_1^2	-0.0000736 (-0.13)	-0.0000900 (-0.16)	0.000242 (1.01)	0.000184 (0.79)	
$2\sigma_{1\beta}$	-0.00445 (-1.18)	-0.00443 (-1.18)	-0.00421 (-1.12)		
σ_β^2	-0.0749 (-1.62)	-0.0755 (-1.63)	-0.0791 (-1.72)	-0.118 (-3.89)	
σ_0^2	0.0618 (2.05)	0.0616 (2.04)	0.0419 (10.17)	0.0421 (10.21)	0.0404 (14.56)
$2\sigma_{0r}$	-0.0290 (-0.66)	-0.0289 (-0.66)			
σ^2	0.0227 (1.38)	0.0228 (1.38)	0.0121 (4.81)	0.0122 (4.85)	0.0113 (11.72)
Collinearity	Yes	Yes	No	No	No
Observations	6287	6287	6287	6287	6287

t statistics in parentheses

4.3 Endogeneity I: Investor risk preferences

If property investors prefer to hold riskier properties for shorter periods of time, then this might be confused for the time-independent components in Figures 1a and 1b. To explore whether there is a basis for this channel, I investigate whether determinants of property risk have any power to predict holding periods. To do this, various property-level attributes are compiled to create a time series panel where each quarter the dependent variable is set to one if a property is within four quarters of its final sale.³³ All properties in the data are used, including those that exit for reasons other than a sale. The following is a list of the cross-sectional variables used in the regression along with their expected relation (if any) to the property's risk. Each, in principle, would be available to the property investor before the decision to sell is made.

SqFt — Square footage proxies for size. A larger property is more likely to have a well diversified pool of tenants and therefore less idiosyncratic risk.

JV — A property owned as a joint venture may be more speculative.

Age when acquired — Older properties may have more idiosyncratic risk.

Percent Leased — Higher vacancy may signify higher risk .

Loan spread — If the property is mortgaged, then the spread of the initial interest rate over the average prevailing mortgage rate paid by other properties proxies for the property's risk.

Property type — Apartment/Office/Industrial/Retail.

Region — East/Midwest/South/West

Appraisal-based lagged return — The accumulate total return on the property (appraisal based) lagged four quarters.

³³The decision to sell a property typically takes place months before an actual sale is consummated.

Appraisal-based time series properties — The property’s appraisal-based time-series of returns, up to four quarters before its last quarter in the database, is regressed against its NPI benchmark to calculate the property’s appraisal-based market adjusted R^2 , market β , and idiosyncratic variance.

Mgr Type — A dummy variable equal to one if the property is owned by the largest firms. Firms are sorted by number of owned properties in the database and the cutoff defined at the 20th percentile (i.e., the smallest firms in aggregate own 20% of the properties).

Table 5 reports the results of a logistic regression (vintage year dummies are included and are significant, but are not reported in the table). For continuous variables, the “Marginal impact” column reports the change in probability of disposition if the explanatory variable is shifted from its 10th percentile in the database of sold properties to the 90th percentile. For dummy variables, this column reports the impact of changing the variable from zero to one.

Small size, JV ownership, property age, vacancy, and high loan spreads all seem to be associated with a higher probability of sale and therefore a propensity towards shorter holding horizons. Intuitively, these are all linked to higher risk. On the other hand, there is much evidence that apartment buildings are the least risky property type, and yet the regression associates it with significantly higher probability of turnover. In addition, arguably the most direct proxy linking idiosyncratic risk to holding horizon should be the imperfectly measured idiosyncratic volatility obtained from appraisal-based returns. This measure, however, is significantly linked to a *lower* propensity to sell and therefore a longer holding period. It is also important to note that the sum of the absolute impact of the the variables with the “wrong” sign (i.e., Apartments and Idiosyncratic variance) is about 5.3% which is 66% greater than the sum of the absolute impact of the effects with the “correct” sign. Finally, by far, the variables with greatest impact on the probability of asset disposition, with a sum of absolute impact of 16.0%, are not obviously linked to idiosyncratic risk: The percentage and magnitude of systematic risk in appraisal-based returns (adjusted R^2 and β), whether the property is owned by a large investor, and the property’s accumulated appraisal-based performance.

The single most impactful property characteristic linked to the decision to sell is, surprisingly, the diversifiability of its returns. The more idiosyncratic the return, as measured by the appraisal-based adjusted R^2 relative to the corresponding NPI benchmark, the more likely it is to be sold.³⁴ That said, the time-independent effects identified are *net* of a systematic risk exposure,

³⁴A natural concern is that the appraisal-based R^2 might be artificially low for properties with short holding horizons because of the shorter time series available. This motivates the use of properties’ adjusted R^2 in the logistic

Table 5: Results of a logistic regression for the probability of a property sale. The dependent variable is one if a property is within four quarters of a sale and zero otherwise. All properties (sold and unsold) and quarters in the data are used. The explanatory variables are property-specific attributes.

Variable	Coeff	<i>t</i> -stat	Marginal impact
SqFt	-1.08E-07	-4.47	-0.0034
JV	0.1060	5.43	0.0075
Age when acquired	0.0061	11.35	0.0119
Percent Leased	-0.5002	-8.75	-0.0085
Loan spread	2.2742	2.65	0.0005
Apartments	0.5007	20.49	0.0394
Industrial	-0.0430	-1.76	-0.0030
Office	0.1705	6.83	0.0123
East	0.0527	2.50	0.0037
Midwest	-0.0093	-0.39	-0.0006
South	0.1723	8.71	0.0124
Lagged return	-0.7748	-20.05	-0.0272
Idiosyncratic variance	-6.9031	-13.50	-0.0135
R_a^2	-1.7288	-38.69	-0.0950
β	0.0145	9.23	0.0043
Mgr Type	0.5835	20.09	0.0339
Const	-6.154	-14.61	
Observations			228,935

and to explain them away one must argue for a positive relationship between the magnitude of appraisal-based idiosyncratic time series risk and the likelihood of a sale — a relationship that appears to be solidly rejected by the data.

Summarizing, while some risk-related property-level variables do predict disposition, this is not consistent across measures and the economic magnitude does not appear decisive. In particular, there is no compelling evidence demonstrating that the magnitude of idiosyncratic risk is positively linked to shorter holding periods. Thus it does not appear that the time-independent effects are an artifact of a preference among managers for holding properties with less idiosyncratic risk over longer periods.

4.4 Endogeneity II: Strategic asset disposition and holding period returns

Property investors have the option but seldom the obligation to sell their investments. Consequently, the holding period of an asset, τ , is endogenously determined and potentially contingent on the property’s performance. This raises concerns that a property’s holding horizon may be related to observable aspects of the return distribution. A natural economic link between τ and idiosyncratic risk attributes is as follows: An investor will purchase a property only if it is deemed to provide risk-adjusted value (i.e., “alpha”) on a forward-looking basis and will sell the asset when this is no longer the case. Thus, to the extent that the previous statement is correct, the observation of a sale indicates that the investor has updated the property’s return distribution parameters over the holding period.³⁵

Suppose, for instance, that the investor purchases the property believing that its annualized alpha is 2%. If the property underperforms, the manager will update his or her beliefs. Sufficient underperformance will eventually cause the manager to conclude that the alpha is zero or negative, at which point the property will be sold. The logistic regression in Table 5 supports this view in that negative appraisal-based performance in the year prior to a sale is one of the most significant predictors of a sale. If properties are mostly sold because they underperform up to some threshold, then the risk-adjusted price paths associated with sold properties will only

regression. To further allay such concerns, the adjusted R^2 is only calculated when the number of observation available for the time-series regression equal or exceed 10. I substitute the median adjusted R^2 of the former for the remaining properties in the database. Whether followed or not, this procedure does not qualitatively alter the results of the logistic regression.

³⁵Properties will also be sold to finance the purchase of more lucrative properties. As long as a held property delivers risk-adjusted value, selling for this purposes is only necessary for financial constrained entities. The companies in the database tend to be large and seemingly well-capitalized institutional investors. In addition, real estate assets possess high collateral value suggesting that selling over-performing assets to fund other purchases will be exceptional.

portray a strict subset of possible price paths and thus exhibit a lower variance. One concern is that this will be more pronounced for properties held over a longer period of time.

To explore this and the empirical implications, consider a model of a single property whose risk-adjusted, cumulative and continuously compounded price appreciation returns evolve as:

$$dr_t = dr_t^O + dr_t^U,$$

where r_t^O corresponds to the part of the return process observable by the investor (e.g., the income stream from the property, its appraised value, etc.), and r_t^U is not observed. Suppose that the investment's abnormal returns or $\alpha(t)$ is a constant that is not known to the manager. When the property is first purchased, the investor's prior over $\alpha(t)$ is normally distributed with mean α_0 and variance η^2 . From observing r_t^O , the investor updates $a_t = E_t[\alpha(t)]$ accordingly. Here, $E_t[\cdot]$ corresponds to a conditional expectations based on all observable information. Assuming that the observed part of the return process is Brownian motion with instantaneous innovation $\sigma_W dW_t$, the manager observes

$$dr_t^O = \alpha(t)dt + \sigma_W dW_t$$

with $r_0^O = 0$, and must use this to update his or her estimate of the expected value of $\alpha(t)$. Because r_t^O and the manager's prior are jointly normally distributed, a standard result (see Theorem 12.1 in Liptser and Shiriyayev, 1978) is that their joint dynamics is given by

$$dr_t^O = a_t dt + \sigma_W dW_t \tag{12}$$

$$da_t = \frac{\sigma_W}{\kappa + t} dW_t, \quad \kappa = \frac{\sigma_W^2}{\eta^2}, \tag{13}$$

where $a_0 = \alpha_0$. Assume that the unobserved (until the sale at $t = \tau$) component of the property return evolves as

$$dr_t^U = \sigma_Z dZ_t,$$

where Z_t is standard Brownian motion that is uncorrelated with W_t such that $W_0 = 0$, and with $r_0^U = 0$.³⁶ Thus, the evolution of the total returns realized when the property is sold is

$$dr_t = dr_t^O + dr_t^U = a_t dt + \sigma_W dW_t + \sigma_Z dZ_t, \tag{14}$$

where $r_0 = 0$ and $Z_0 = W_0 = 0$, and where a_t evolves according to (13). Suppose that the manager will sell the property at the first instance, τ , that the expected value of $\alpha(\tau)$ falls below some threshold $\alpha_L < \alpha_0$. I.e., $\tau = \inf_t E_t[\alpha(t+s)] \leq \alpha_L$ for all $s > 0$. Given the martingale property of the updating rule (which is an optimal forecast), $E_t[\alpha(t+s)] = a_t$ is constant for all s . Thus a sale takes place at holding period of τ if and only if τ is the first passage time for $a_t = \alpha_L$. An important result follows from this observation.

³⁶One can also consider r_t^U to be observed by the manager but containing no information about $\alpha(t)$.

Proposition 1. Let $\hat{\alpha} \equiv \alpha_0 - \alpha_L$. At the first passage time, $\tau = \inf_t \{a_t \leq \alpha_L\}$,

$$r_\tau = \sigma_Z Z_\tau - \kappa \hat{\alpha} + \alpha_L \tau. \quad (15)$$

Thus $E[r_\tau] = -\kappa \hat{\alpha} + \alpha_L \tau$ and $VAR[r_\tau] = \sigma_Z^2 \tau$.

Proof. The result follows from the observation that $r_t^O = (\kappa + t)(a_t - \alpha_0) + \alpha_0 t$. To see this, first note that this equation holds at $t = 0$. Next, consider that

$$d[(\kappa + t)(a_t - \alpha_0) + \alpha_0 t] = a_t dt + (\kappa + t) da_t = a_t dt + \sigma_W dW_t = dr_t^O.$$

Thus equality follows from the fact that r_t^O and $(\kappa + t)(a_t - \alpha_0) + \alpha_0 t$ are processes with identical evolutions that coincide at $t = 0$. One can therefore write,

$$r_t = \sigma_Z Z_t + (\kappa + t)(a_t - \alpha_0) + \alpha_0 t.$$

Evaluating this at $t = \tau$, the first passage time defined by $a_\tau = \alpha_L$ yields (15). The expected value and variance of r_τ follow from the fact that Z_t is independent of τ . ■

If $\kappa \hat{\alpha}$ is fixed, then Proposition 1 establishes that, despite the endogeneity of the decision to sell, the variance of holding period returns should still be proportional to the holding period. While it may be possible that variability in $\kappa \hat{\alpha}$ across investors leads to the time-independent variability in holding period returns, the model makes a counter-factual prediction.³⁷ Specifically, the additional variation should also appear in the observable cumulative returns for holding periods near the first passage time. In other words, the additional variation should also appear in appraisal-based cumulative returns with holding periods close to the actual holding period of the property. Figure 4 depicts the appraisal-based cumulative returns for holding the property until two quarters before the sale.³⁸ The dots and dashed lines portray the point estimate and associated 95% confidence interval for each holding period variance. It is evident that the time-independent component is missing. In fact, the variance is in line with what one would predict for a random walk process using time-series (i.e., quarter by quarter) estimates of volatility. Indeed, the square symbols correspond to the average quarterly time-series variation in the appraisal-based returns multiplied by the number of periods.³⁹

³⁷Assume κ is on average equal to ten, corresponding to an investor with ten years of experience for forming a prior and an average $\hat{\alpha}$ of 5%, and that $\kappa \hat{\alpha}$ is uniformly distributed from 25% to 75%. This is sufficient to deliver a time-independent variability of 0.042, in line with the observed quantity.

³⁸In practice, it makes little qualitative difference in the results whether one uses a lag of one or two quarters. In principle, the sales price may be known a few months before the sale is recorded and may therefore be incorporated into the appraised value one quarter before the sale.

³⁹To employ an analogy, the endogenous sale model in Proposition 1 predicts that appraisal-based cumulative returns would behave like a translated Brownian bridge rather than Brownian motion. This implies that the time-series variation and holding period variation should not line up as they do in Figure 4.

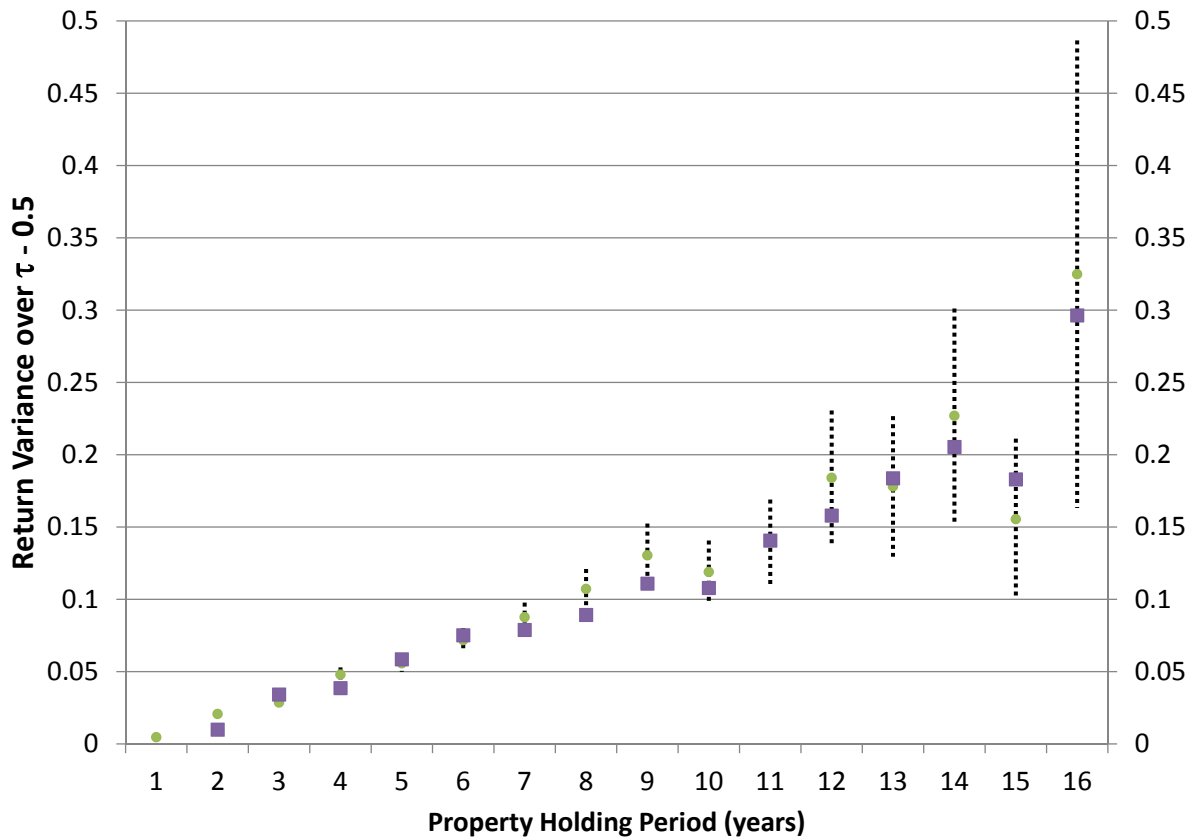


Fig. 4: Point estimates and their 95% confidence intervals for the regression residual variance in Equation (2) (each holding period corresponds to an independent regression). Appraisal-based holding period returns are used with holding horizons lagged two quarters before a sale occurs. The square boxes correspond to estimates of the quarterly appraisal-based time-series volatility multiplied by the holding horizon.

A related alternative model is that properties are sold if they underperform to a low threshold α_L (as in the model above) or overperform to high threshold α_H (e.g., the property is sold to exercise an option for redevelopment). In this case, the distribution of holding period returns should be bimodal, and the results in Proposition 1 will be modified to

$$\text{VAR}[r_\tau] = \sigma_Z^2 \tau + p(\tau)(1 - p(\tau))(\tau + \kappa)^2(\alpha_H - \alpha_L)^2, \quad (16)$$

$$E[r_\tau] = -\kappa\alpha_0 + (\kappa + \tau)\left(p(\tau)\alpha_L + (1 - p(\tau))\alpha_H\right), \quad (17)$$

where $p(\tau)$ is the probability that a sale occurs because of underperformance rather than because of overperformance.⁴⁰ This expression is non-zero at $\tau = 0$, thus mimicking the time-independent variance effect. Given sufficiently high precision it may even be made to fit both Figure 1a and 1b, as long as $\alpha_0 - \alpha_L > \alpha_H - \alpha_0$. The problem, however, is the prediction of bi-modality: The holding period return distributions of underperforming and overperforming properties should separate with time. Table 6 reports the results of the Hartigan and Hartigan (1985) “dip” test for unimodality at each integer holding horizon and vintage year combination for which at least 40 observations are available. All 43 tests, save for one, are unable to reject the null of a unimodal distribution at any conventional level.⁴¹

In summary, while models of strategic disposition can account for the time-independent effects, they also make counter-factual predictions about the appraisal-based and actual return distributions, suggesting that this type of endogeneity is not a main driver behind the observed time-independent holding period return components.

5 Determinants of individual return variance

The analysis in Section 4.2 can be expanded to investigate the determinants of heterogeneity among property holding period returns. The covariance between $\tilde{\varepsilon}$ and $\tilde{\alpha}_0$ in (10) is shown to be statistically insignificant in Table 4 and the corresponding interaction term subsequently dropped because it is collinear with the other regressors in (11). This leaves open the possibility of further identifying heterogeneity in both the volatility component, corresponding to σ^2 and the time-independent variance component corresponding to σ_0^2 . To this end, rewrite (11), leaving out the terms that were found to be insignificant, as

$$\hat{z}^2 = \frac{1}{\tau}\sigma_0^2 + \sigma^2 + \tilde{U}.$$

⁴⁰It is straight forward to derive an expression for $p(\tau)$.

⁴¹Visual inspection and kernel density estimates of the number of modes confirm the unimodal distribution of the holding period returns.

Table 6: Results of a Hartigan and Hartigan (1985) “dip” test for deviations from unimodality of the holding period returns. Tests are done separately for different vintages to control for time variations in property return distributions. A low p -value rejects the null of a unimodal distribution.

Holding Period	Vintage	n	dip	p-value	Holding Period	Vintage	n	dip	p-value
1	2001	43	0.034	0.96	4	2000	83	0.031	0.89
1	2002	53	0.045	0.67	4	2001	54	0.063	0.22
1	2003	88	0.027	0.94	4	2002	55	0.052	0.41
1	2004	65	0.042	0.60	4	2003	44	0.049	0.69
1	2005	98	0.050	0.17	4	2005	42	0.057	0.48
1	2006	78	0.053	0.24	4	2006	47	0.052	0.53
					4	2007	55	0.030	0.97
2	1996	43	0.071	0.20					
2	2001	51	0.055	0.40	5	2000	53	0.035	0.92
2	2002	50	0.046	0.68	5	2001	50	0.073	0.15
2	2003	76	0.049	0.32	5	2002	46	0.049	0.63
2	2004	82	0.038	0.66	5	2006	56	0.032	0.96
2	2005	65	0.034	0.90	5	2007	78	0.033	0.84
2	2011	52	0.036	0.93	5	2008	65	0.030	0.96
3	2000	41	0.044	0.83	6	1999	50	0.043	0.79
3	2001	78	0.033	0.84	6	2000	45	0.045	0.76
3	2002	127	0.113	0.00	6	2006	65	0.033	0.91
3	2003	123	0.043	0.22	6	2007	59	0.053	0.35
3	2004	81	0.021	0.99					
3	2007	58	0.025	0.98	7	2000	41	0.063	0.38
3	2010	62	0.064	0.16	7	2004	43	0.036	0.96
					7	2005	40	0.035	0.95
					7	2006	122	0.028	0.83
					8	1999	41	0.046	0.80
					8	2005	43	0.042	0.85

Both σ^2 and σ_0^2 can be viewed as average effects of a more complex variance structure that can be captured via

$$\hat{z}^2 = \frac{1}{\tau} X_1' \cdot b_1 + X_2' \cdot b_2 + \tilde{U}. \quad (18)$$

Where the X_i 's are vectors of property-specific characteristics and the b_i 's can be interpreted as their marginal contribution to the idiosyncratic variance of the property. The vector b_1 corresponds to contributions to the time-independent variance (i.e., σ_0^2), while the vector b_2 corresponds to contribution to time series volatility (i.e., σ^2). Here too one has to be watchful of collinearity when including the same characteristics in the sets corresponding to X_1 and X_2 .

I begin by adding vintage effects to the original specification. This is done by identifying X_2 with a constant and a set of year dummies, where each dummy variable takes the value of one if it corresponds to the property's entry year into the database and zero otherwise. The X_1 is only identified with a constant and therefore measures σ_0^2 "directly". By recalling that each property appears once in the regression, it can be seen that these dummies control for time variation in property market volatility that may give rise to an artificial time-independent variance component. Table 7 reports estimations of the model in (18) both unconditionally as well as when the set of properties is restricted by region, property type, investor type and ownership structure. The unconditional regression demonstrates that the time-independent variance is robust to vintage effects. Moreover, the idiosyncratic volatility component (the σ^2) is higher than estimated without controlling for vintage (see Table 4B, Model 5). The coefficients in the regressions restricted by geography, while suggestive of some heterogeneity, are not statistically distinguishable from their unconditional counterparts. The evidence for heterogeneity among property types is a bit stronger. The idiosyncratic variance of Apartments resides almost exclusively in the time-independent component, meaning that when held over long horizons the idiosyncratic risk of apartments is amortized and investors face mainly systematic holding period risk. By contrast, office properties carry substantially higher time-series volatility and are thus more likely to separate from their benchmark index over long holding horizons. The table also reveals some difference between properties held by large versus small investors: Large investors hold properties that are significantly less susceptible to a time-independent shock, consistent with the notion that they hold more liquid properties. While there is an economically significant difference between the volatility of separately owned properties versus those held through joint ventures, this difference is not statistically significant.

Table 8 reports the results from a different set of regressions. This time, the full set of properties is employed and dummies are used, as part of X_2 , to isolate the effects of geography and property types. In addition, the property's log of age at acquisition, average vacancy, and log of square footage are also incorporated into X_2 along with the vintage effects. Only a constant (i.e., $\frac{1}{\tau}$) and

Table 7: Structural investigation of the second-stage regression, Equation (18). Year dummies are included to control for vintage. Each row reports the regression coefficients for restricted property subset. The σ_0^2 column reports the estimated time-independent variance component and the σ^2 column reports estimates of the idiosyncratic volatility component.

Subset	σ_0^2	t -stat	σ^2	t -stat	N	Vintage effects
Unconditional	0.0428	13.89	0.0189	4.82	6287	Yes
Region						
W	0.0524	9.32	0.0206	3.58	2040	Yes
S	0.0387	7.25	0.0094	1.00	1890	Yes
M	0.0367	4.09	0.0214	2.82	981	Yes
E	0.0417	6.43	0.0166	1.25	1376	Yes
Property type						
A	0.0604	8.93	-0.0067	-0.14	1526	Yes
I	0.0411	8.99	0.0163	3.92	2191	Yes
O	0.0351	5.55	0.0325	3.05	1685	Yes
R	0.0428	4.63	0.0192	1.92	885	Yes
Investor Type						
Small	0.0547	6.95	0.0087	0.84	1239	Yes
Large	0.0378	11.21	0.0217	5.23	5048	Yes
JV						
No	0.0450	13.46	0.0187	4.87	5099	Yes
Yes	0.0355	4.50	-0.0208	-0.41	1188	Yes

average vacancy are used for X_1 because of collinearity considerations. From this exercise it is once again evident that geography does not significantly contribute to the variance estimates. Property type, however, does appear to be important for the volatility component, with office properties being the riskiest, followed by retail. Size is not a key determinant of the volatility but property log-age and vacancy are positively and significantly related to it. By contrast, vacancy appears to reduce the time-independent risk (corresponding to the $\frac{1}{\tau}$ interaction term), although the regressions suggest the possibility that this is due to collinearity with the separate vacancy and $\frac{1}{\tau}$ terms.

Overall, these estimates together with the analysis in the previous sections suggest that a “typical” property will exhibit an idiosyncratic σ^2 between 0.01 and 0.02 (i.e., $\sigma \sim 10\% - 14\%$), and a time-independent variance resembling a one-time illiquidity shock (assessed when the property is sold) of roughly 0.04 ($\sigma_0 \sim 20\%$), with adjustments made for property type and property risk characteristics such as vacancy, size, and age.

6 Conclusions

Real estate risk is different from the risk of liquid traded assets. Though this may seem self-evident, quantifying the risk of individual real estate assets has been left relatively unexplored in the literature. I use purchase and sale data from the National Council of Real Estate Investment Fiduciaries (NCREIF) to compute holding period returns for commercial properties and estimate the typical idiosyncratic risk associated with individual assets. Analysis of this data suggests that idiosyncratic risk comes in two forms: a volatility component similar to that exhibited by liquid assets measuring between 10% and 14% per year (annualized); and a time-independent variance component that persists at all horizons with standard deviation of roughly 20%. Average holding period returns also appear to exhibit a positive time-independent component. The estimates appear robust to various specifications and controls. The most plausible explanation for the presence of the time-independent components in holding period returns is the underlying asset illiquidity. A search-based illiquid asset pricing model calibrated to model parameters does a good job of capturing all of the effects. The model can also be applied to the pricing real estate derivative instruments such as debt or mortgage backed securities. It can also be extended to other illiquid assets, such as rarely-traded bonds, private equity, or complex financial assets.

Table 8: Structural investigation of the second-stage regression, Equation (18). Year dummies are included to control for vintage. Explanatory variables are listed in the second column and contribute to the risk component in the first column.

		(1)	(2)	(3)	(4)	(5)
σ^2	size	-0.000224 (-0.26)	-0.000299 (-0.35)		-0.000580 (-0.69)	-0.000672 (-0.80)
σ^2	Avg. Vacancy	0.0705 (5.67)	0.0271 (3.98)		0.0766 (6.27)	0.0315 (4.76)
σ^2	Age at acquisition	0.00134 (1.97)	0.00137 (2.02)		0.00149 (2.24)	0.00154 (2.31)
σ_0^2	$\frac{1}{\tau}$	0.0490 (10.93)	0.0385 (10.37)	0.0427 (13.84)	0.0496 (11.10)	0.0386 (10.42)
σ_0^2	Avg. Vacancy $\times \frac{1}{\tau}$	-0.110 (-4.16)			-0.116 (-4.39)	
σ^2	Industrial	-0.0000698 (-0.04)	0.000274 (0.14)	0.000832 (0.51)		
σ^2	Office	0.00539 (2.62)	0.00605 (2.94)	0.00684 (4.07)		
σ^2	Retail	0.00393 (1.65)	0.00412 (1.72)	0.00377 (1.89)		
σ^2	Midwest	0.00143 (0.60)	0.00152 (0.64)	0.00155 (0.79)		
σ^2	South	0.00267 (1.35)	0.00301 (1.51)	0.00174 (1.05)		
σ^2	West	0.00200 (1.01)	0.00209 (1.06)	0.00143 (0.87)		
σ^2	Const.	0.00552 (0.32)	0.00963 (0.55)	0.0153 (3.50)	0.0106 (0.62)	0.0155 (0.91)
Observations		4033	4033	6287	4033	4033
Vintage effects		Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

References

- Beran, R., and P. Hall, 1992, "Estimating coefficient distributions in random coefficient regressions," *The annals of Statistics*, pp. 1970–1984.
- Black, F., and M. S. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81(3), 637–54.
- Case, K., and R. Shiller, 1987, "Prices of Single Family Homes Since 1970: New Indexes for Four Cities," *National Bureau of Economic Research Working Paper Series*, 2393.
- Downing, C., R. Stanton, and N. Wallace, 2008, "Volatility, mortgage default, and CMBS subordination," *University of California, Berkeley*.
- Duffie, D., N. Garleanu, and L. H. Pedersen, 2005, "Over-the-Counter Markets," *Econometrica*, 73(6), 1815–1847.
- , 2007, "Valuation in Over-the-Counter Markets," *Review of Financial Studies*, 20(6), 1865–1900.
- Fisher, J., D. Gatzlaff, D. Geltner, and D. Haurin, 2003, "Controlling for the Impact of Variable Liquidity in Commercial Real Estate Price Indices," *Real Estate Economics*, 31(2).
- Gatzlaff, D. H., and D. R. Haurin, 1997, "Sample Selection Bias and Repeat-Sales Index Estimates," *The Journal of Real Estate Finance and Economics*, 14(1-2), 33–50.
- Geltner, D., N. Miller, J. Clayton, and P. Eichholtz, 2013, *Commercial Real Estate Analysis and Investments*. OnCourse Learning, 3rd edn.
- Goetzmann, W., and L. Peng, 2006, "Estimating House Price Indexes in the Presence of Seller Reservation Prices," *The Review of Economics and Statistics*, 88(1), 100–112.
- Goetzmann, W. N., 1993, "The single family home in the investment portfolio," *The Journal of Real Estate Finance and Economics*, 6(3), 201–222.
- Hartigan, J. A., and P. Hartigan, 1985, "The dip test of unimodality," *The Annals of Statistics*, pp. 70–84.
- Hildreth, C., and J. P. Houck, 1968, "Some estimators for a linear model with random coefficients," *Journal of the American Statistical Association*, pp. 584–595.
- Hwang, M., and J. Quigley, 2004, "Selectivity, Quality Adjustment and Mean Reversion in the Measurement of House Values," *The Journal of Real Estate Finance and Economics*, 28(2-3), 161–178.

- Liptser, R. S., and A. N. Shiryaev, 1978, *Statistics of Random Processes II: Applications*. Springer-Verlag, New York, NY.
- Merton, R. C., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29(2), 449–70.
- Raj, B., V. Srivastava, and A. Ullah, 1980, "Generalized Two Stage Least Squares Estimators for a Structural Equation with Both Fixed and Random Coefficients," *International Economic Review*, pp. 171–183.
- Swamy, P. A., 1970, "Efficient inference in a random coefficient regression model," *Econometrica: Journal of the Econometric Society*, pp. 311–323.