

Asset-level risk and return in real estate investments*

Jacob S. Sagi[†]

Abstract

In stark contrast with liquid asset returns, commercial real estate idiosyncratic return means and variances do not scale with the holding period. This puzzling phenomenon survives a variety of controls for vintage effects, systematic risk heterogeneity, and sample selection biases. To explain these findings, I derive an equilibrium search-based illiquid asset-pricing model which, when calibrated, provides an excellent fit to the data. Thus a structural model of transaction risk seems crucial to understanding real estate investment returns. Recent empirical literature suggests that these insights can be extended to other illiquid asset classes such as private equity and residential real estate.

Keywords: Real estate, illiquid assets, holding period returns, search, idiosyncratic risk, transaction 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 feedback from Shaun Bond, Bob Connolly, Lynn Fisher, Andra Ghent, Adam Gurren, Dave Hartzell, Allen Head, Robert Kimmel, Crocker Liu, Ludo Phalippou, Tim Riddiough, Kanis Saengchote, Morten Sørensen, Chester Spatt, Richard Stanton, Toni Whited, seminar participants at Baruch College, the University of Calgary, and the University of Cincinnati, and conference participants at the NUS-IRES 2015 conference, the 2015 Summer Real Estate Symposium, the 2015 UBC-ULE Symposium, the 2015 Fall HULM Conference, Inquire Europe Spring 2016 Conference, and the 2017 Utah Winter Finance Conference. 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 primarily focuses on aggregated or portfolio-level attributes. Much is therefore known about the systematic returns of real estate assets held as investments.¹ 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 real estate investors hold concentrated portfolios. Geltner et. al (2013) estimate that in 2010 the stock of institutional quality commercial real estate was valued at around \$3 trillion, with more than a third held by non-fund entities they identify as “...private investors, relatively small, typically locally oriented often family-based enterprises.” This suggests that many direct real estate investors 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, 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. As further motivation, commercial real estate debt, totaling roughly \$2.5 trillion as of 2015, is mostly secured to individual properties and, in the presence of limited recourse, reflects both the systematic and idiosyncratic risks to which the property is exposed.² A solid understanding of asset-level price dynamics is important for the efficient pricing, monitoring, and regulation of debt obligations and structured debt products such as commercial mortgage backed securities. Finally, commercial real estate investments share much in common with other “alternative investments” that are infrequently traded. Lessons drawn from understanding property-level risk can be extrapolated to other highly illiquid assets such as other real assets, private equity deals, large whole loans, and mergers and acquisitions of entire firms.

Infrequent trading of commercial real estate (CRE) assets renders impractical the direct estimation of time-series attributes of individual assets. Instead, inference of “typical” property-level

¹In this paper I focus on real estate properties held by institutional investors rather than owner-occupants.

²The figure is taken from Table L.217 of the 2015 U.S. Federal Reserve Bank Balance sheet.

risk and return characteristics from transaction prices has to leverage the cross-section. A purchase and sale decision, or equivalently an observed holding period return, can be viewed as a sampling from the return distribution of the underlying property’s stochastic price process. If such decisions are unrelated to the underlying process, then the sampling is random and, with adequate controls, a large cross-section of properties can be viewed as a “Monte Carlo” simulation of idiosyncratic or asset-level risk and return.³ If decisions to purchase or sell properties depend on attributes of the stochastic price process, then the sampling is not random and the panel of transaction data is subject to selection bias. In the latter case, one must account for the selection bias to measure characteristics of the underlying price process.

To illustrate the central issue, suppose property price evolves as a standard geometric random walk, as in the standard pricing model of Black and Scholes (1973), and jointly assume that transaction decisions are unrelated to the characteristics of the price process. Theoretically, the idiosyncratic variance and mean of log price appreciation returns should scale linearly with the holding horizon.⁴ In Section 2 I show that this prediction is solidly rejected (see Figure 1). In the data, idiosyncratic return variance is approximately linear in the holding horizon but the intercept is three times larger than the slope and highly significant. For the mean, even a linear relationship is statistically untenable and one year holding periods are associated with 4% expected “abnormal” returns. Failure of the predictions leads to a rejection of the joint hypothesis of a random walk log-price process and independent transaction decisions. The conclusions are robust to controlling for cross-sectional heterogeneity and time-variation in property return parameters. Moreover, the anomalous effects largely disappear if returns from repeat transactions in the panel are replaced with REIT returns or appraisal-based individual property returns. Further validation comes from similar empirical findings in related asset classes: residential real estate (Case and Shiller, 1987; Abraham and Schauman, 1991; Goetzmann, 1993; Goetzmann and Spiegel, 1995; Calhoun, 1996;

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 can be measured with high precision because of the large cross section

⁴Henceforth, I use the term “returns” and “price appreciation returns” interchangeably. The reason for focusing on price appreciation returns rather than total returns is discussed in Section 2.

Giacoletti, 2016) and individual private equity deals (Axelson, Sorensen, and Stromberg, 2015; Lopez-de-Silanes, Phalippou, and Gottschalg, 2016).

Guided by the intuition that illiquidity can lead to pronounced differences between short- and long-term holding strategies, I develop in Section 3 a search-based equilibrium pricing model where investors vary in how they value the income 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 Because not all bids are accepted, transactions (and therefore holding periods) are endogenous. In a random matching and bargaining equilibrium the model simultaneously delivers a price process and a distribution of repeat transaction returns. When investor private valuations are dynamic but persistent, the model can be made to fit 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, when assets are illiquid, an investor will only hold for a short period of time if a much better offer happens to come along. The corresponding high short-term “alpha” is earned by luck and not by design. Less fortunate investors do not sell after a short period of time and do not show up in a panel of short holding period transactions. Selection bias is why the mean of observed holding period returns extrapolates to a positive value as the holding horizon shrinks.

In the model, 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 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. If each investor’s private valuation is serially uncorrelated (i.e., not persistent), then the model becomes one of “jump transactions” around a “true” fundamental price. Reduced-form jump transaction models have been proposed by Case and Shiller (1987),

⁵ Real 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 purchase price).

Goetzmann (1993), and Axelson, Sorensen, and Stromberg (2015). Such models, however, are only economically consistent if the expected returns are non-positive in the limit of a vanishing holding period (otherwise they may admit arbitrage opportunities).⁶ The data, however, suggests positive “alphas” for short holding periods, underscoring the limitation of reduced-form approaches.

In addition to capturing the holding period risk and return patterns, the model also delivers a good fit to the observed quarterly turnover, the fraction of properties sold after five years, and the average transacted income-to-price ratio (or “cap rate”). In contrast with a standard price diffusion model where the price at date t is known given information at date t , the market price of a property in the model additionally features *transaction risk*. This is represented by an equilibrium probability distribution of transacting as a function of the seller’s reservation cap rate. In the absence of illiquidity there is no transaction risk and the associated distribution is degenerate. In fitting to the holding period return data, I find that the implied transaction risk distribution exhibits significant positive skewness. This leads to a long expected “time on market” for owners who will not accept any offers below average prevailing market transaction. The model also provides a one-to-one mapping between a reservation selling price, the expected transaction price, and the expected (or median) time on the market. Finally, given a strict time horizon for disposition and a strategy for a seller to accept/reject offers, one can calculate the price of immediacy, corresponding to an expected discount relative to prevailing transaction prices. Such a measure can be useful in better understanding the private equity closed-end asset management structure which operates under a stated liquidation target horizon. More generally, the various model features can be important for modeling bank sales of foreclosed assets and feeds into the pricing of derivative assets such as mortgage loans and mortgage backed securities.

The few papers that attempt to quantify attributes of property level risk and return all assume that, after discounting for risk and time, an individual property’s log-value evolves as a random walk — a hypothesis on which this study casts serious doubts. The closest work is that of 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

⁶This is discussed further in Section 4.

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 figure, significantly higher than my 10% estimate for the idiosyncratic random walk volatility component, may be the result of neglecting the impact of the large transaction risk component I identify. In another study, Plazzi, Torous, and Valkanov (2008) employ average transaction prices in various U.S. metropolitan areas to conclude that geographic dispersion contributes from 4% to 7% per year to a property's volatility. Peng (2015) uses holding period returns, similar to that used in this study, to quantify the systematic risk exposure of individual assets.⁷ The latter two papers provide insights into individual property dispersion but, unlike the first, do not provide a full assessment of property-level risk. Because its inference approach is indirect, Downing, Stanton, and Wallace (2008) is subject to other criticisms.⁸

This paper contributes to the literature in several ways. First, I document that real estate investment holding period returns are not generated by a random walk process. This result is neither spurious nor, as I argue in Appendix C, consistent with selection bias when the underlying property price process is random walk. Second, I derive a tractable equilibrium model for illiquid asset prices that, when calibrated, fits the data very well. Thus, in explaining the sizable non-standard features

⁷Peng (2015) notes in his study that the variance of idiosyncratic holding period property returns does not strongly depend on the holding period. Although he does not delve into the possible reasons or the implications for his assumption of a random walk process, this finding is consistent with mine.

⁸An important advantage to the methodology used by Downing, Stanton, and Wallace (2008) is that it provides a forward looking assessment of volatility. However, 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 may not be representative of the universe of commercial real estate properties (see Ghent and Valkanov, 2014, for a comparison between the loan types). 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. Indeed, my results suggest that asset-level jump-like shocks at disposition should play an important role in how a bank might price securitized debt — something not captured by the model used in Downing, Stanton, and Wallace (2008).

of observed idiosyncratic real estate returns, illiquidity appears to be the most plausible driver. In principle, the model and analysis can also be applied to other search-based (e.g., broker-mediated) markets for highly illiquid assets in which repeat transactions are uncommon and the main source of asset-level data comes from a cross-section of holding period returns. Some examples include private equity transactions, rarely traded bonds or large whole loans, complex financial assets, mergers and acquisitions, and other real assets (ships, oil rigs, mines, etc.). Moreover, this study has important implications for refining the financial econometrics use to analyze such markets (see, for instance, the approach in Ang and Sorensen, 2012, and references therein).

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 poor methodology will add idiosyncratic noise unrelated to the actual property-level risk.

To properly measure attributes of typical property-level returns, one must therefore deploy a different strategy. The first challenge is to obtain transaction-based returns from purchases and sales. 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 holding period returns and make inferences by leveraging the cross-section.

2.1 The Data

Data of sold and unsold properties comes from the National Council of Real Estate Investment Fiduciaries (NCREIF) and consists of quarterly financial and accounting information for properties reported by member firms between 1978Q1 and 2013Q4. The dataset contains property acquisition dates and prices, as well as capital expenditures in each quarter starting from the acquisition quarter (or 1978Q1 if acquisition is earlier) until 2013Q4 (or until the property is sold or otherwise exits the database). If a sale (or partial sale) took place, the sale date and net/gross prices (excluding/including selling costs) are reported. For each property, appraisal-based price appreciation and income returns are reported each quarter. Details about the property location, age, real estate category, leverage, owning member firm, and the type of fund in which the property is held are also available. Fields that are available but not used in this study include detailed appraisal valuations and breakdowns of income, operating expenses, and capital expenditures. The dataset also reports whether a given property qualifies for inclusion in the NCREIF Price Index (NPI). Qualifying properties have a minimum 60% occupancy requirement, correspond to one of the major property types (Apartments, Industrial, Retail, or Office), and are owned by tax-exempt institutions. Such properties tend to be well-maintained, are located in high-demand markets, and are frequently referred to as “core assets” by professionals.

The data contains properties with missing or inconsistent entries that I attempt to filter using a set of criteria detailed in Appendix A. To summarize, I require that the property history in the dataset exhibits no gap or appraisal return below -99.9%. The acquisition date must be unique, documented to be no later than the first quarter the property appears in the dataset, and the corresponding purchase price must be positive. Properties that exit the dataset must either be sold or foreclosed. For sold properties, a sale date and a sale price greater than \$1 must be reported. Partial sales cannot be negative or different from the net proceeds if reported during a quarter of final disposition. The time-series sum of capital expenditures and the transaction-based total capital appreciation cannot be negative.

To help address concerns about tax motives for transactions, I restrict attention to properties

that qualified for the NPI when first acquired. Finally, to help mitigate bias from unreported capital expenditures, I require there to be at most a lag of one quarter from the acquisition date of a property to the first time it appears in the dataset. The intersection of criteria results in 3,976 properties that experienced a sale (16 of which were part of the sale of a portfolio), 30 properties that were surrendered to the bank, and 3,148 that had not exited as of 2013Q4. Appendix B reports the results of applying the main empirical tests of this section to various looser/tighter alternative filters.

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. One could employ some feasible income investment strategy (e.g., bonds, REITs, etc.) but because real estate investments tend to exhibit high payouts, the total holding period returns would eventually be dominated by the returns from the reinvestment alternative rather than the property. For this reason, I choose to focus on price appreciation returns. To that end, the excess log of price-appreciation return over property i 's holding period is given by,

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

where $r_{I,s}$ is the quarter s return on investing the proceeds from partial sales. The expression in (1) corresponds to an excess return because the denominator is capitalized to date T using a risk-free return. While this expression also appears to depend on a discretionary investment strategy, in practice only 157 of the final set of filtered 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 NPI corresponding to the property’s major asset type. *Henceforth, to avoid tedious expressions, unless otherwise indicated I employ the term “return” to refer to the “log of total price-appreciation return net of a risk-free alternative”*, as in equation (1).

The return calculated in (1) contains outliers. I deal with these in two ways. To help identify outliers, I compare each property’s transaction-based return to a similar quantity calculated using the sum of reported quarterly appraisal-based returns. Absent errors in the data, this difference should not be too large. I then group repeat transaction returns from (1) into integer holding period “buckets”. The assignment is made by rounding holding periods to the nearest integer with random assignment for half-integer holding periods. In each bucket, I then drop all properties for which the transaction and appraisal difference is greater than or equal to the 99th percentile. The percentiles are calculated separately for integer holding period buckets because it would be natural for discrepancies between appraisal and transaction returns to accumulate with horizon. Without dropping outliers, the squared correlation between the two return calculations is 73.4%. After dropping the outliers, the squared correlation increases to 83.0%. Whenever running return regressions, a second way in which I deal with outliers is to drop extreme return observations (below the first and above the ninety-ninth percentiles) from each integer holding period bucket.

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

2.2 A Simple Null Hypothesis

In this section I investigate the following:

Null Hypothesis. *Suppose that property price dynamics correspond to a stochastic diffusion process and that transaction decisions are unrelated to that process.*

It is worth emphasizing that the null is of considerable practical value. If true, it can be

exploited to design empirical strategies for measuring the diffusion parameters from repeat transactions (i.e., holding period returns) as, in fact, is done below. Key to this is the observation that, given the null, the expected value and variance of returns should be proportional to the holding period.

Under a standard random walk hypothesis, the growth in property i 's price over a short time interval, Δ , takes the form

$$r_{i,\Delta} = \alpha_i \Delta + \beta_i r_{m,\Delta} + \sigma_i \sqrt{\Delta} \varepsilon_{i,\Delta},$$

where α_i, β_i and σ_i are constants, ε_{Δ} represents a standardized idiosyncratic shock to the property value, and $r_{m,\Delta}$ captures the influence of systematic shifts in property market values over Δ . If the idiosyncratic shocks are serially independent, then when accumulated over a holding period of duration τ , the sequence of incremental property returns sums to

$$r_{i,\tau} = \alpha_i \tau + \beta_i r_{m,\tau} + \sigma_i \sqrt{\tau} \varepsilon_{i,\tau}, \quad (2)$$

where $r_{m,\tau}$ is the change in the systematic risk factor over the same holding period, and $\varepsilon_{i,\tau}$ is an idiosyncratic mean-zero shock with variance one. In equation (2), the idiosyncratic holding period return is $r_{i,\tau} - \beta_i r_{m,\tau} = \alpha_i \tau + \sigma_i \sqrt{\tau} \varepsilon_{i,\tau}$. Because the decision to sell, and therefore τ , is assumed independent of the parameters of the return process, the expected value and variance of the idiosyncratic return are both proportional to τ . Moreover, the random walk hypothesis applied to r_m implies that its expected value and variance are proportional to τ . Thus even without netting out the market component, the expected return and return variance of $r_{i,\tau}$ are proportional to the holding period. Finally, if the constant coefficients in the equation for the incremental return, $r_{i,\Delta}$, vary with time then equation (2) still holds but α_i, β_i , and σ_i^2 in (2) would be replaced by their time-series averages over the holding period.

For a property purchased at date t and sold at $t + \tau$, it is possible to observe τ and $r_{m,\tau}$. If the model in (2) applies to commercial real estate properties, then $\varepsilon_{i,\tau}$ is unrelated across properties. Moreover, because transactions are assumed independent of the price process, $\varepsilon_{i,\tau}$ is unrelated to τ . Along with the preceding discussion this suggests an empirical strategy for testing the null and, under the null, identifying the coefficients in (2). Specifically, one can view each property's

holding period return from a repeat transaction as a random sampling from a distribution of possible coefficients and from the distribution of $\varepsilon_{i,\tau}$ (which is standardized, by definition). The corresponding empirical model for the property return over a holding period of τ can be rewritten as

$$\tilde{r} = \tilde{\alpha}\tau + \tilde{\beta}r_{m,\tau} + \sigma\sqrt{\tau}\tilde{\varepsilon},$$

where the index i from equation (2) corresponds to a single realization from the distribution of $\tilde{\alpha}$, $\tilde{\beta}$, and $\sigma\tilde{\varepsilon}$. To the last equation, I add an intercept term, $\tilde{\alpha}_0$, which I allow to be heterogeneous across properties. Under the null $\tilde{\alpha}_0$ should be zero. To minimize collinearity in the regression specification to follow, I normalize the resulting equation by $\sqrt{\tau}$ and then add an intercept term to the normalized formulation. This results in

$$\frac{\tilde{r}}{\sqrt{\tau}} = \frac{\tilde{\alpha}_0}{\sqrt{\tau}} + \alpha_1 + \tilde{\alpha}\sqrt{\tau} + \tilde{\beta}\frac{r_{m,\tau}}{\sqrt{\tau}} + \sigma\tilde{\varepsilon},$$

where it bears repeating that $\tilde{\alpha}_0$ and α_1 are zero under the null. Further separating the expected values of the random coefficients above by setting $\tilde{\alpha}_0 \equiv \alpha_0 + \tilde{\varepsilon}_0$, $\tilde{\alpha} \equiv \alpha + \tilde{\varepsilon}$, and $\tilde{\beta} \equiv \beta + \tilde{\varepsilon}_\beta$, yields

$$\frac{\tilde{r}}{\sqrt{\tau}} = \alpha_0\frac{1}{\sqrt{\tau}} + \alpha_1 + \alpha\sqrt{\tau} + \beta\frac{r_{m,\tau}}{\sqrt{\tau}} + \left(\tilde{\varepsilon}_0\frac{1}{\sqrt{\tau}} + \tilde{\varepsilon}\sqrt{\tau} + \tilde{\varepsilon}_\beta\frac{r_{m,\tau}}{\sqrt{\tau}} + \sigma\tilde{\varepsilon}\right). \quad (3)$$

If $\tilde{\varepsilon}$ and the $\tilde{\varepsilon}$'s are uncorrelated with $r_{m,\tau}$ and τ then equation (3) is a regression equation with heteroskedastic residuals. While this assumption seems plausible for $r_{m,\tau}$, it may not be so for τ . This is because NCREIF properties have not been added to the dataset at a uniform rate (see Table 1).

[Table 1 about here.]

In particular, properties with holding periods shorter than five years have been typically added at or after 2000 while those with holding periods longer than nine years were typically added at or before 1992. Consequently, time variation in asset-specific characteristics may lead to a mechanical relationship between the distribution of holding periods and the distributions of $\sigma\tilde{\varepsilon}$ and the $\tilde{\varepsilon}$'s. For instance, if σ across properties was lower before 1996 than after, then properties held over shorter horizons (mostly dating from after 1996) would spuriously appear to exhibit more idiosyncratic variance per unit time than properties held over longer horizons. To control for such “vintage

effects” I decompose each of α and σ^2 into a sum of heterogeneous year-specific contributions as follows:

$$\frac{\tilde{r}}{\sqrt{\tau}} = \frac{\alpha_0}{\sqrt{\tau}} + \alpha_1 + \sum_y v_{\alpha,y} \frac{e_{\tau,y}}{\sqrt{\tau}} + \beta \frac{r_{m,\tau}}{\sqrt{\tau}} + \left(\tilde{\epsilon}_0 \frac{1}{\sqrt{\tau}} + \tilde{\epsilon} \sqrt{\tau} + \tilde{\epsilon}_\beta \frac{r_{m,\tau}}{\sqrt{\tau}} + \left(\sum_y v_{\sigma,y} \frac{e_{\tau,y}}{\tau} \right)^{\frac{1}{2}} \tilde{\epsilon} \right), \quad (4)$$

where $e_{\tau,y}$ is the overlap (in years) between the property’s holding period and year y , while $v_{\alpha,y}$ and $v_{\sigma,y}$ are, respectively, the year y contributions to α and σ^2 . By definition $\sum_y e_{\tau,y} = \tau$, thus if $v_{\alpha,y} = \alpha$ and $v_{\sigma,y} = \sigma^2$ for every year y , then equation (4) reduces to (3).

Because α_0 , α_1 , and $\tilde{\epsilon}_0$ are zero under the null, equation (4) is simultaneously a test of the null and, assuming the null holds, a model for estimating key property price diffusion parameters. Despite the fact that the residual term (in parentheses) is heteroskedastic, the coefficients α_0 , α_1 , the $v_{\alpha,y}$ ’s, and β can be consistently (but not efficiently) estimated via OLS. The estimation strategy for the distributional properties of the residual term loosely follows Beran and Hall (1992) who advocate a sequence of n regressions to estimate the first n moments. The idea is as follows. Denote the OLS residual from the regression equation (4) as \hat{z} . The squared residuals can be written as

$$\begin{aligned} \hat{z}^2 &= E[\hat{z}^2] + \tilde{U}, \\ &= \frac{1}{\tau} \sigma_0^2 + \tau \sigma_\alpha^2 + \frac{(r_{m,\tau})^2}{\tau} \sigma_\beta^2 + \sum_y v_{\sigma,y} \frac{e_{\tau,y}}{\tau} + 2r_m^e \sigma_{\alpha\beta} + O\left(\frac{1}{N}\right) + \tilde{U}, \end{aligned} \quad (5)$$

where $\sigma_0^2 = \text{Var}[\tilde{\epsilon}_0]$, $\sigma_\alpha^2 = \text{Var}[\tilde{\epsilon}]$, $\sigma_\beta^2 = \text{Var}[\tilde{\epsilon}_\beta]$, and $\sigma_{\alpha\beta} = \text{Covar}[\tilde{\epsilon}_\beta, \tilde{\epsilon}]$, while \tilde{U} is a zero-mean random variable uncorrelated with $\frac{1}{\tau}$, τ , $\frac{(r_m^e)^2}{\tau}$, $\frac{e_{\tau,y}}{\tau}$ or r_m^e . The $O(\frac{1}{N})$ 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. Missing from equation (5) are covariance terms between $\tilde{\epsilon}$ and $\tilde{\epsilon}_\beta$ or $\tilde{\epsilon}$, and between $\tilde{\epsilon}_0$ and the other variables. The former case is justified by considering that while an owner may be able to control a property’s α_i and β_i (and thus the two could be related), under the null, ε_i is unrelated to these quantities. In addition, under the null, $\tilde{\epsilon}_0 = 0$ and setting its covariance with other variables to zero does not bias the test against the null.

For large N , the coefficients, $\sigma_0^2, \sigma_\alpha^2, \sigma_\beta^2, \sigma_{\alpha\beta}$, and the $v_{\sigma,y}$ ’s in (5) 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 earlier 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 . 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 the normalized holding period returns $\frac{\tilde{r}}{\sqrt{\tau}}$ against $\frac{r_{m,\tau}}{\sqrt{\tau}}$, $\frac{1}{\sqrt{\tau}}$ and the $\frac{e_{\tau,y}}{\sqrt{\tau}}$'s, and calculate the residuals, \hat{z} .
2. Regress \hat{z}^2 against $\frac{1}{\tau}$, τ , $\frac{(r_{m,\tau})^2}{\tau}$, r_m^e and the $\frac{e_{\tau,y}}{\tau}$'s without a constant. Calculate \hat{z}_p^2 , the predicted value of \hat{z}^2 using the regression coefficient estimates, representing an estimate of the property-specific residual variance from the first pass regression.
3. Correct for heteroskedasticity in the first regression by regressing $\frac{\tilde{r}}{\sqrt{\tau\hat{z}_p}}$ against $\frac{r_{m,\tau}}{\sqrt{\tau\hat{z}_p}}$, $\frac{1}{\sqrt{\tau\hat{z}_p}}$ and the $\frac{e_{\tau,y}}{\sqrt{\tau\hat{z}_p}}$'s.
4. Use the heteroskedasticity-adjusted first-moment coefficient estimates to recalculate \hat{z}^2 and regress against $\frac{1}{\tau}$, τ , $\frac{(r_{m,\tau})^2}{\tau}$, r_m^e and the $\frac{e_{\tau,y}}{\tau}$'s without a constant to obtain estimates of the heteroskedasticity adjusted second-moments.

Before implementing this estimation approach I drop holding period return outliers: observations above or below the 1st and 99th percentiles within each integer holding period “bucket” (rounding holding periods to the nearest integer to determine their censorship bucket, as discussed earlier in this section). I also exclude observations with holding periods less than one year. This is consistent with how the NCREIF calculates its NPI index and is done out of concern that very short holding periods often correspond to portfolio acquisitions in which the acquirer quickly sells “undesirable” properties out of the portfolio. In such cases, there may not be an objective purchase price for the undesirable property and the acquirer may potentially allocate it an arbitrary purchase price (and

therefore an arbitrary price appreciation). The results are qualitatively unchanged if one includes these properties.

[Table 2 about here.]

For each property I use the corresponding property type-specific NPI price appreciation returns in excess of the three-month treasury bill to compute a proxy for the systematic factor, $\tilde{r}_{m,\tau}$, over the corresponding holding period. Table 2 reports the regression parameter estimates. Comparing the first two columns in each panel, it seems that adjusting for heteroskedasticity does not markedly impact the coefficient estimates. The average beta of the properties, as might expected, is close to one. The sum, $\frac{1}{36} \sum_y v_{\alpha,y}$, is an estimate of the annualized expected return diffusion drift in the property market averaged across all years. If the $v_{\alpha,y}$ were equal across years this would be an estimate of α in (3). The average diffusion drift is only marginally significant, suggesting that at least on average the NPI price appreciation index return accounts for most of the return drift. This, of course, is not surprising given that the analysis is restricted to NPI properties. The probability that the alpha is constant across years, denoted in the table as “ $P(v_{\alpha,y})$ const”, is negligible. Thus vintage effects are highly significant and controlling for them is important. Under the null, α_0 and α_1 should be zero. This is solidly rejected in the data. The variance analysis in Panel B suggests that dispersion in β_i 's and α_i 's does not contribute much to holding period return variance in the data.⁹ The sum $\frac{1}{36} \sum_y v_{\sigma,y}$ corresponds to the average annualized diffusion variance, and its most precise estimate (fifth column) is 0.0101 per year, corresponding to a diffusion volatility of roughly 10% per year. This would be an estimate of σ in (3) if the $v_{\sigma,y}$'s were equal across all years. The latter hypothesis is solidly rejected (“ $P(v_{\sigma,y})$ const” is smaller than 0.0001 across all specifications). The other (highly) significant variance component, σ_0^2 should be zero under the null. Instead, it represents nearly three times the annual diffusion variance.

A main takeaway from the estimation summarized in Table 2 is that the null is solidly rejected. Idiosyncratic expected returns are not proportional to the holding period, instead featuring both

⁹One can constrain the variance estimate of σ_β^2 and σ_α^2 to be positive using non-linear least squares. All this achieves is to drive their estimates to zero.

a non-zero intercept and a non-linear component. Likewise, idiosyncratic return variance is not proportional to the holding period, instead featuring a large intercept component. To examine the robustness of these findings using a simpler approach, I first group properties by integer year holding periods, rounding non-integer holding periods, and randomly rounding up or down with half-integer holding periods. For each such grouping, indexed by the integer holding period τ , I estimate the following regression after dropping returns below the first and above the 99th percentiles:

$$\tilde{r}_\tau = \alpha_\tau + \beta_\tau \tilde{r}_{m,\tau} + \sigma_\tau \tilde{\varepsilon}_\tau. \quad (6)$$

[Fig. 1 about here.]

Figure 1 plots the α_τ 's and σ_τ^2 's from the sequence of regressions in (6) against the holding period (point estimates in green triangles, 95%-confidence interval in thin dotted lines) alongside the predictions from the more elaborate specification in (4) (point estimates in black circles, 95%-confidence interval in thick gray lines). The two estimation approaches roughly agree both visually as well as statistically (chi-squared tests for the difference in the predictions between the procedures are highly insignificant). While, of the two approaches, the multi-stage estimation with vintage effects is clearly the more econometrically sophisticated, the agreement is reassuring. It is clear from these plots that expected return and variance tend toward significantly positive values as the holding period vanishes. The non-linear dependence of expected returns on holding period is also visually evident.

It is worth noting that each of Plazzi, Torous, and Valkanov (2008), Downing, Stanton, and Wallace (2008), and Peng (2015), employs a model consistent with the vanishing of return expected values and variances as $\tau \rightarrow 0$. In the language of continuous-time finance, Table 2 and Figure 1 suggest that the idiosyncratic drift and volatility of the return process, corresponding to the expected return and variance per unit time, are divergent as $\tau \rightarrow 0$. These results suggest either that the holding period returns are subject to some bias or that the underlying price process is distorted by some market frictions at short horizons. In the next section, I offer an illiquidity-based model to explain why holding period expected returns and variances deviate markedly from their prediction under the null. Before doing so, I summarize other evidence that reinforces the empirical

conclusions of this section.

2.3 Robustness

Appendix B provides additional evidence that the results of the previous section are not spurious. For instance, when one randomly matches the return of a single property held over consecutive periods Δ_1 and Δ_2 with the compounded return from two properties, one of which was held over Δ_1 and the other over Δ_2 , then the return variance of the single property is lower than the variance of the compound returns from the two matched properties by essentially the estimate of σ_0^2 from Table 2. Because the holding periods of the two strategies coincide, the only material difference between the two returns is that the latter involves an additional two-way transaction (i.e., purchase and sale). This hints that the anomalous return behavior observed in Figure 1 stems from transactional frictions.

Indeed, evidence similar in spirit to that presented in Section 2.2 has been documented for other highly illiquid asset classes, such as residential real estate and individual private equity deals. Case and Shiller (1987) and Goetzmann (1993) document an anomalously large variance for short-term residential housing repeat transaction returns, while Abraham and Schauman (1991) additionally detects nonlinearity in the variance growth against holding periods. Goetzmann and Spiegel (1995) suggest that such findings might be attributable to the fact that calculated repeat transaction returns for residential housing ignore renovations (i.e., capital expenditures), and model this by adding a jump component with positive mean to explain holding period returns. Giacoletti (2016) is able to include home improvement expenditures in his return calculations, yet continues to find that holding period returns deviate from the random walk prediction. In the context of private equity buyout deals, and similarly to this study, Axelson, Sorensen, and Stromberg (2015) document a significant intercept to both variance and expected return components as functions of a deal's holding period. Likewise, Lopez-de-Silanes, Phalippou, and Gottschalg (2016) document that private equity average holding-period returns are not proportional to the holding period. A common feature of these different asset markets is the presence of (sometimes severe) illiquidity.¹⁰

¹⁰Peng (2015) notes that the variance of idiosyncratic holding period commercial property returns does not strongly

The literature just surveyed provides some external validity to the notion that the results here are not spurious. Correspondingly, it is shown in Appendix B that the null hypothesis is *not* rejected when one repeats the analysis substituting publicly traded real estate investment trust (REIT) returns for individual property holding period returns. This too suggests that the findings are driven by the relative illiquidity of physical properties. Finally, the analysis reported in Table 2 and depicted in Figure 1 is qualitatively robust to various alternative specifications, such as: Excluding properties purchased before 1997, including non-NPI properties, and removing the constraint that the purchase date must not lag the initial reporting date for the property by more than a quarter.

The rejection of the null from the previous section implies that either property fundamentals deviate substantially from a random walk with time-varying coefficients or that transaction decisions, underlying observed holding period returns, are not independent from the underlying process. Each of these possibilities presents a challenge to understanding asset-level risk and return. The first suggests market inefficiency in that it may be possible to earn abnormally large risk-adjusted returns in the commercial property market. The second suggests that one cannot model and estimate property-level fundamentals without jointly modeling endogenous repeat transactions. The presence of similar evidence from other asset classes makes the market inefficiency story less persuasive for it would have to apply across a much larger set of markets. In addition, the set of properties forming the basis of analysis here tend to be competitively owned and managed by industry professionals in an institutional context. Such a market setting does not intuitively furnish fertile ground for easily exploitable inefficiencies.

This leaves the possibility that observed holding period returns suffer from a selection bias because actual holding period returns are only observed when a property is sold, and the selling decision may not be independent of the original purchase or subsequent sale price. In Appendix C I consider various potential biases that are unrelated to liquidity, but that can lead to effects documented in Figure 1. One simple possibility is that institutions prefer to hold riskier assets for

depend on the holding period, although he does not delve into the possible reasons or the implications for his assumption of a random walk process.

shorter periods of time, thus mechanically creating a negative correlation between holding period and property risk. Another is that the option to sell a property is exercised contingent on individual property performance. For instance, if properties are only sold if they underperform relative to some benchmark, then one might find a relationship between risk adjusted return characteristics and the holding period. To summarize the analysis in Appendix C, I find no empirical support for the predictions coming from the liquid market alternative hypotheses considered.

3 A model of holding-period returns for illiquid assets

Under the random walk model of equation (2), the idiosyncratic Sharpe ratio for holding the asset over a period τ is $\frac{\alpha_i \tau}{\sigma_i \sqrt{\tau}} = \frac{\alpha_i}{\sigma_i} \sqrt{\tau}$, which vanishes as τ approaches zero. In particular, in a frictionless market the absence of arbitrage rules out the possibility of a finite and positive Sharpe ratio for a vanishingly small holding period. In Figure 1, the extrapolated Sharpe ratio for arbitrarily small τ is $\frac{0.216}{\sqrt{0.0107}} \approx 2$, which is both finite, positive, and large. However, because commercial real estate assets are highly illiquid and take months to transact, there is no possibility of exploiting such lucrative extrapolated short-term returns. 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 such an explanation for the phenomenon, supported by a calibrated equilibrium model.

I consider an approach reminiscent of the search models employed in Duffie, Gârleanu, and Pedersen (2005, 2007). There are N infinitely-lived income-producing properties. Time is discrete and each period corresponds to a quarter, matching the structure of the data. Each property is held and managed by some investor of a type indexed by a . At date $t + 1$, a property owned by an investor of type a will pay income $\tilde{d}_{t+1} = d_t e^{(\mu - \frac{\sigma^2}{2}) + \sigma \tilde{\epsilon}_{t+1}}$, where d_t is the property's income in the previous period, $\tilde{\epsilon}_{t+1}$ is a standard normally distributed and serially uncorrelated random variable, while the volatility σ and drift μ are constant. I assume that $\frac{\tilde{d}_{t+1}}{d_t}$ is identically distributed across properties. Later, I will also assume that $\tilde{\epsilon}_{t+1}$ can be decomposed into an idiosyncratic and a common shock for each property. For now, to avoid burdensome notation, I suppress reference

to any specific property.

To calculate his or her private value for a property at date t , investor a discounts next period's expected income and private property value by a factor e^{-r_a} . If investors' private values do not change then eventually all properties would be owned by the investors with the lowest discount rate. To derive an equilibrium in which a steady state of ownership exists it seems sensible for investors' private discount factors to mean-revert. To that end, I assume that the growth-adjusted capitalization factor, $\eta_a \equiv e^{r_a - \mu}$, is a regular Markov process taking one of $2S + 1$ values, where S is a positive integer, and where each state is indexed by $a \in \mathcal{A} \equiv \{1, \dots, 2S + 1\}$. Without loss of generality assume that $\eta_1 \geq \eta_2 \geq \dots \geq \eta_{2S+1}$. Denote by $\eta(t)$ some investor's type at date t . The transition probability matrix for this investor's type is $\Pi_{\hat{a}a} \equiv \text{Prob}(\eta(t+1) = \eta_{\hat{a}} | \eta(t) = \eta_a)$, where $a, \hat{a} \in \mathcal{A}$. Let $x \in (0, \frac{1}{2})$ and assume that the transpose of the transition matrix is given by $\Pi^T \equiv P^n$ where n is a positive integer exponent and where

$$P = \begin{pmatrix} 1-2x & 2x & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ x & 1-2x & x & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & x & 1-2x & x & \dots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & x & 1-2x & x & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & x & 1-2x & x \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 2x & 1-2x \end{pmatrix} \quad (7)$$

Π^T describes a discretized reflected random walk on a finite interval. The exponent n determines the number of states accessible from the central state, $S + 1$.¹¹ This is a flexible modeling environment in that, other than the assumed monotonicity, it does not constrain the values taken by the η_a 's. Denote by π_a^U the unconditional probability that an investor is in state a at date t .¹² An advantage of modeling Π^T as above is that, as is easy to check, the unconditional distribution of $\eta(t)$ is simply

¹¹For example, if $n = 1$ then only the nearest states, S and $S + 2$ are accessible from the central state in one period's transition. If $n \geq S$ then all states are accessible within one period's transition.

¹²Regularity of the Markov process implies that the row vector $(\pi^U)^T$ is given by any row of the matrix defined by $\lim_{n \rightarrow \infty} (\Pi^T)^n$. The superscript T refers to the transpose of the vector or matrix in question.

given by

$$\pi^U = \left(\frac{1}{4S}, \frac{1}{2S}, \frac{1}{2S}, \dots, \frac{1}{2S}, \frac{1}{2S}, \frac{1}{4S} \right)^T.$$

The process $\eta(t)$ is assumed to be independent and identical across investors, and independent of $\tilde{\epsilon}_{t'} \forall t'$. One can view, η_a , the time-varying growth-adjusted capitalization factor of a property owner/manager as a reduced-form proxy for the effects of institutional liquidity and capital constraints, individual managerial beliefs and/or preferences, portfolio effects, managerial constraints, and agency concerns. Ultimately, this modeling approach is chosen for its tractability.

At date t , immediately after income is distributed, every investor and owner is subject to a type transition according to the probability law described by Π . Post-transition, each property receives an offer from some randomly chosen investor and its owner must decide whether or not to sell at a cost, $c \times d_t$, assumed for analytic convenience to be proportional to the property's income. It is also 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 $\pi_{a'}^U$. The motivation to trade in this model comes from the heterogeneity generated by different private values. The frictions in this model consist of the cost of transacting a sale and, more importantly, the limited trading opportunities — each period the counterparty is a single potential buyer rather than a market of potential buyers. As discussed in Footnote 5, the assumption of limited trading opportunities is particularly fitting in the context of real estate, but may be applicable to other broker-mediated (rather than dealer-mediated) markets. It is further assumed that investors face no explicit funding constraints and that there are no constraints on the number of properties that an investor may hold. Thus the ratio of investors to properties is immaterial.¹³

A sale between dates t and $t + 1$ takes place between an owner of type \hat{a} and investor of type a' if and only if the owner's certainty equivalent (i.e., private value) of the property value, $p_{t,\hat{a}}$, is smaller than the investor's certainty equivalent, $p_{t,a'}$, less cd_t . If the difference between bidder's and owner's private values exceeds the transaction costs, assume that in the ensuing bargaining

¹³The 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.

the seller receives a random fraction, $\tilde{\lambda} \in [0, 1]$ with mean $\bar{\lambda}$, of the gains from trade. Thus, when a sale takes place the transaction price net of costs is $p_{t,\hat{a}} + \tilde{\lambda} \left\{ p_{t,a'} - p_{t,\hat{a}} - cd_t \right\}^+$ (where $\{x\}^+ = \max\{0, x\}$). I assume that $\tilde{\lambda}$ is identically and independently distributed across time and buyers/sellers, and unrelated to $\tilde{\epsilon}_t, \eta_{\hat{a}}$ or $\eta_{a'}$.

An owner or prospective buyer between dates t and $t+1$ values the property (ex-dividend and post-transition) by taking into account the sequence of events described earlier, but projected one period ahead. This projection, from the point of view of a prospective owner of type a , is depicted as a time-line schematic in Figure 2.

[Fig. 2 about here.]

Thus, the ex-dividend and post-transition property certainty equivalent of an owner of type a is determined by

$$p_{t,a} = e^{-ra} E \left[\tilde{d}_{t+1} + \tilde{p}_{t+1,\hat{a}} + \tilde{\lambda} \left\{ \tilde{p}_{t+1,a'} - \tilde{p}_{t+1,\hat{a}} - c\tilde{d}_{t+1} \right\}^+ \right] \quad (8)$$

where $\tilde{p}_{t+1,\hat{a}}$ and $\tilde{p}_{t+1,a'}$ correspond, respectively, to the owner's and bidder's private values (ex-dividend and post transition) at date $t+1$. While \hat{a} depends on a via the Markov transition process, the distribution of a' only depends on the unconditional probability vector π^U . In words, the owner's valuation, $p_{t,a}$, equals the discounted continuation value of holding the income producing property plus the option value of selling (at a cost) to a prospective buyer.

Definition. *An equilibrium is a positive and finite random variable $p_{t,a}$ that solves (8) for every $a \in \mathcal{A}$.*

If \mathcal{A} is a singleton set, then $a = \hat{a} = a'$ and $p_{t+1,\hat{a}} = p_{t+1,a}$, so that $p_t = e^{-r} E[\tilde{d}_{t+1} + p_{t+1}]$ 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 some constant rate r . Liquidity has no 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 a

different valuation and bargaining power. In this case, (8) becomes

$$p_{t,a} = e^{-r_a} E \left[\tilde{d}_{t+1} + \tilde{p}_{t+1,\hat{a}} + \max \left\{ 0, \tilde{\lambda}(\tilde{p}_{t+1,a'} - \tilde{p}_{t+1,\hat{a}} - c\tilde{d}_{t+1}) \right. \right. \\ \left. \left. \tilde{\lambda}'(\tilde{p}_{t+1,a''} - \tilde{p}_{t+1,\hat{a}} - c\tilde{d}_{t+1}) \right. \right. \\ \left. \left. \tilde{\lambda}''(\tilde{p}_{t+1,a'''} - \tilde{p}_{t+1,\hat{a}} - c\tilde{d}_{t+1}), \dots \right\} \right].$$

If $c = 0$, then as the number of independent bidders grows the equilibrium will approach one where only the investors with highest private values and least bargaining power will acquire the asset, and the property price will reflect their valuation. This can be viewed as the frictionless limit in which the asset is always held by those who derive the most utility from it. If $c > 0$, then even with a large number of bidders the equilibrium price will not converge to a single value, but transactions will be limited to prices within a band determined by c .

To proceed with the analysis, conjecture an equilibrium private valuation (ex-dividend and post transition) of

$$p_{t,a} = d_t Q_a.$$

Then from (8) and the model assumptions, Q_a must solve the linear system of equations:

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

where Π^T denotes the transpose of Π .¹⁴ The dependence on types is completely captured by the vector of η_a 's. If η_a is constant across types then, assuming $r - \mu > 0$, an equilibrium solution is given by $Q_a = (\eta - 1)^{-1} \approx (r - \mu)^{-1}$ — a simple growing annuity factor. Henceforth, I assume model parameters are chosen so that an equilibrium exists.

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 $Q_{a'} - Q_a \geq c$, where $Q_{a'}$ 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

¹⁴The $\{\cdot\}^+$ appears to make the set of equations non-linear. However, if $\{Q_a\}_{a \in \mathcal{A}}$ solves this set of equations, then it also solves a corresponding linear system of the form $T \cdot \vec{Q} = \vec{b}$ for some matrix T , where \vec{Q} is the vector of Q_a 's and \vec{b} is a constant vector.

the incumbent owner type. If a trade occurs between an owner of type a and an investor of type a' at date t , the observed net transaction price is

$$p_{t,aa'} = d_t \left(Q_a + \tilde{\lambda}(Q_{a'} - Q_a - c) \right) \text{ s.t. } Q_{a'} - Q_a \geq c. \quad (9)$$

This expression is a function of property market characteristics as well as the identities of the seller and bidder, and their relative bargaining power. In other words, as a function of property market information alone, the property transaction price at date t is not a number but a distribution (i.e., it can take on multiple values). In this model, it is meaningful to speak of the probability of sale (by an owner) as a function of some reservation price, whereas with liquid assets such a probability is trivially set to be one below the market clearing price and zero above it. This additional stochastic feature of prices can be important for modeling the price of immediacy of execution and feeds into the pricing of derivative assets such as mortgage loans and mortgage backed securities.

3.1 Holding Period Returns

To analyze holding-period returns, consider a property that at date t is purchased from some owner of type $o \in \mathcal{A}$ by an investor of type $a \in \mathcal{A}$. Suppose the property is held until date $t + \tau$, at which point the current owner has transitioned to type \hat{a} and sells to a buyer of type $b \in \mathcal{A}$. Then the observed holding period price appreciation total return corresponding to the repeat transaction is:

$$\tilde{R}_{t,\tau}(o, a, \hat{a}, b) = \frac{Q_{\hat{a}} + \tilde{\lambda}'(Q_b - Q_{\hat{a}} - c)}{Q_o + c + \tilde{\lambda}(Q_a - Q_o - c)} \frac{\tilde{d}_{t+\tau}}{d_t} = \frac{Q_{\hat{a}} + \tilde{\lambda}'(Q_b - Q_{\hat{a}} - c)}{Q_o + c + \tilde{\lambda}(Q_a - Q_o - c)} e^{(\mu - \frac{\sigma^2}{2})\tau + \sigma\sqrt{\tau}\tilde{n}},$$

where $\tilde{n} = \frac{1}{\sqrt{\tau}} \sum_{j=1}^{\tau} \tilde{\epsilon}_{t+j}$ is a standard normally distributed random variable and where $\tilde{\lambda}$ and $\tilde{\lambda}'$ are iid. Note that the purchase price paid is gross of costs but the selling price received is net of costs (i.e., in the presence of transaction costs a buyer will pay more than the seller receives). In a frictionless setting (e.g., $c = 0$ and there is only one type of owner), $\tilde{R}_{t,\tau}(o, a, \hat{a}, b) = e^{(\mu - \frac{\sigma^2}{2})\tau + \sigma\sqrt{\tau}\tilde{n}}$. This is the standard geometric random walk result in which the identities and private values of the transactors are immaterial in that there is no dependence on o, a, \hat{a} or b . In the presence of limited trading opportunities, $\tilde{R}_{t,\tau}(o, a, \hat{a}, b)$ depends not only on property-specific characteristics between t and $t + \tau$ but also on the attributes of the investors involved in the repeat transaction.

This feature is what drives the joint hypothesis problem: One cannot properly infer the property parameters, μ and σ , from holding period returns without also modeling (implicitly or explicitly) transaction dynamics.

The logarithm of the holding period return separates into three sources of risk:

$$\begin{aligned} \ln \tilde{R}_{t,\tau}(o, a, \hat{a}, b) = & \underbrace{\ln \left(Q_{\hat{a}} + \tilde{\lambda}'(Q_b - Q_{\hat{a}} - c) \right)}_{\text{Selling shock}} - \underbrace{\ln \left(Q_o + c + \tilde{\lambda}(Q_a - Q_o - c) \right)}_{\text{Purchasing shock}} \quad (10) \\ & + \underbrace{\sigma \sqrt{\tau} \tilde{n}}_{\text{Income shock}} + \left(\mu - \frac{\sigma^2}{2} \right) \tau. \end{aligned}$$

In equation (10), both the purchasing and selling shocks are idiosyncratic to the property. Using the subscript i to refer to a specific property, the more conventional income shock component can be further decomposed into an idiosyncratic and systematic part as

$$\sigma \tilde{n}_i = \sigma_M \tilde{n}_M + \sigma_I \tilde{n}_{I,i},$$

where \tilde{n}_M is a systematic shock common to all properties, $\tilde{n}_{I,i}$ is specific to property i , and $\sigma_M^2 + \sigma_I^2 = \sigma^2$. The risk-adjusted, or idiosyncratic, distribution of property holding returns can therefore be written as

$$\begin{aligned} \ln \tilde{R}_{t,\tau}^I(o, a, \hat{a}, b) = & \ln \left(Q_{\hat{a}} + \tilde{\lambda}'(Q_b - Q_{\hat{a}} - c) \right) - \ln \left(Q_o + c + \tilde{\lambda}(Q_a - Q_o - c) \right) \quad (11) \\ & + \sigma_I \sqrt{\tau} \tilde{n}_{I,i} - \frac{\sigma_I^2}{2} \tau, \end{aligned}$$

where it is assumed that a well-diversified portfolio or index of properties will appreciate at an expected rate of $\mu - \frac{\sigma_M^2}{2}$ (i.e., risk-adjustment eliminates $(\mu - \frac{\sigma_M^2}{2})\tau$ and $\sigma_M \sqrt{\tau} \tilde{n}_M$ from (10)).

Equation (11) describes the main object of interest in Section 2. The idiosyncratic shock component, $\sigma_I \sqrt{\tau} \tilde{n}_{I,i}$, whose variance grows linearly with the holding period, is independent of the purchasing and selling shocks. The purchase and sale shocks, however are correlated because \hat{a} and a are not independent. As τ increases, because an owner's type is stationary, the dependence of \hat{a} on a grows weaker and the purchase and sale shocks become asymptotically independent. In this limit, the contribution of the purchase and sale shocks to the holding period variance will remain constant as the horizon is increased. In addition, because the purchase price is gross of costs while

the selling proceeds are net of costs, one expects that the purchase shock will (asymptotically) stochastically dominate the selling shock. Their expected difference in (11) will therefore approach a negative constant as τ increases.

3.1.1 Distribution of Holding Period Returns

To calculate the distribution of holding period (i.e., repeat transaction) returns from (11) and compare its predictions with the data, one can consider the following:

1. The probability that the initial transaction involved an owner of type o depends on the distribution of ownership at date t , post-transition and just before an offer is received (see Figure 2). In principle, this distribution could be nearly arbitrary and depends on the initial ownership conditions. To constrain this degree of freedom, I only consider a steady state equilibrium (one in which the distribution of ownership is not expected to change).
2. The model assumes that the distribution of offers coincides with the unconditional distribution of types, π^U . Thus the likelihood that some fixed type o is matched with and sells to type a , or that some fixed type \hat{a} is matched with and sells to type b , can be derived from π^U and the additional constraints that $Q_a - Q_o \geq c$ and $Q_b - Q_{\hat{a}} \geq c$.
3. The probability of observing a holding period of τ depends on the likelihood that $\tau - 1$ previous offers were rejected.

Steady State Distribution of Ownership. A steady state is achieved once the distribution of ownership across properties is not expected to change. If the number of properties and investors is sufficiently large, then the expected distribution of property ownership would correspond with the actual distribution of property ownership. The distribution of ownership from one period to the next changes because of sales and type transitions. Denote the initial distribution just before an offer is received as π^O . Once an offer is received, ownership can change as described via the matrix Π^{Sale} to be shortly derived. Finally, before arriving at an offer in the next period, ownership type transitions based on the Markov evolution described by $\Pi^T = P^n$ in Equation (7). A steady state

is achieved if

$$\pi_o^O = \sum_{i,j \in \mathcal{A}} \pi_i^O \Pi_{ij}^{\text{Sale}} \Pi_{jo}^T. \quad (12)$$

Equation (12) asserts that the pre-offer distribution of ownership is stable in the steady state. It also pins down the distribution corresponding to the random variable \tilde{Q}_o in (11).

To calculate Π_{ij}^{Sale} , let $\Pi_{ij,t}^{\text{Sale}}$ correspond to the probability that a property owned by an investor of type i just before an offer is received at date t , is subsequently owned by an investor of type j just after the offer is accepted/rejected. Because bids arrive from investors drawn from the unconditional distribution, π^U , and a transfer of ownership requires $Q_j - Q_i \geq c$, one can write,

$$\Pi_{ij,t}^{\text{Sale}} = \begin{cases} \pi_j^U & \text{if } Q_j - Q_i \geq c \text{ and } i \neq j \\ \sum_{Q_k - Q_j < c} \pi_k^U & \text{if } i = j \text{ and } c > 0 \\ \pi_i^U + \sum_{Q_k - Q_i < c} \pi_k^U & \text{if } i = j \text{ and } c = 0 \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

The first case corresponds to a transfer of ownership.¹⁵ The second and third cases correspond to rejected offers. The last case corresponds to the zero probability case in which an offer with *lower* valuation is accepted by the owner. Note also that $\Pi_{ij,t}^{\text{Sale}}$ is time-independent, justifying the absence of the time index in (12).

Probability of rejecting sequential offers. 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, first define

$$\pi_{\mathcal{S}a} = \sum_{Q_i - Q_a \geq c} \pi_i^U,$$

¹⁵It is assumed that if buyer and seller are indifferent to the transaction, i.e., $Q_i - Q_j = c$, the property is sold.

the probability that a property of type a is sold to any other type subsequent to an offer.¹⁶ and consider the $(2S + 2) \times (2S + 2)$ modified transition probability matrix defined by

$$\hat{\mathcal{P}}_{aa'} = \begin{cases} \Pi_{aa'}^T (1 - \pi_{S_{a'}}) & a, a' \in \mathcal{A} \\ \sum_{i \in \mathcal{A}} \Pi_{ai}^T \pi_{S_i} & i \in \mathcal{A}, a' = 2S + 2 \\ 0 & a = 2S + 2, a' \in \mathcal{A} \\ 1 & a = a' = 2S + 2 \end{cases} \quad (14)$$

To clarify intuition, consider the case $S = 1$, $c > 0$, and $Q_1 < Q_2 < Q_3$, in which case $\pi_{S_3} = 0$ and

$$\hat{\mathcal{P}} = \begin{array}{c} \begin{array}{cccc} & Q_1 & Q_2 & Q_3 & \text{Sold} \\ Q_1 & \left(\begin{array}{cccc} \Pi_{11}^T(1 - \pi_{S_1}) & \Pi_{12}^T(1 - \pi_{S_2}) & \Pi_{13}^T & \Pi_{11}^T \pi_{S_1} + \Pi_{12}^T \pi_{S_2} \\ \Pi_{21}^T(1 - \pi_{S_1}) & \Pi_{22}^T(1 - \pi_{S_2}) & \Pi_{23}^T & \Pi_{21}^T \pi_{S_1} + \Pi_{22}^T \pi_{S_2} \\ \Pi_{31}^T(1 - \pi_{S_1}) & \Pi_{32}^T(1 - \pi_{S_2}) & \Pi_{33}^T & \Pi_{31}^T \pi_{S_1} + \Pi_{32}^T \pi_{S_2} \\ 0 & 0 & 0 & 1 \end{array} \right) \\ Q_2 \\ Q_3 \\ \text{Sold} \end{array} \end{array} \quad (15)$$

Notice that $\hat{\mathcal{P}}$ is non-negative and that the elements in each row sum to one. Thus $\hat{\mathcal{P}}$ is (the transpose of) a Markov transition matrix. The matrix $\hat{\mathcal{P}}$ corresponds to the probability that a current owner of type a transitions into type a' next period, and then sells *or* remains the owner next period. The “Sold” state is absorbing. Correspondingly, for $a, a' \in \mathcal{A}$, the aa' element of $\hat{\mathcal{P}}^{\tau-1}$ is the probability that investor a remains the property owner after $\tau - 1$ periods (and $\tau - 1$ offers) but has transitioned to type a' by the end of the $(\tau - 1)^{\text{th}}$ period. This motivates defining $\hat{\mathcal{P}}_{\mathcal{A}}^{\tau-1}$ as the upper $(2S + 1) \times (2S + 1)$ submatrix of $\hat{\mathcal{P}}^{\tau-1}$.

Probability of observing $\tilde{R}_{t,\tau}(o, a, \hat{a}, b)$. Consider the following “path”: A property is sold to an investor of type a at date t by an owner of type o drawn from the steady state distribution of owners. The property is then held without being sold (despite the arrival of offers) until the new owner transitions to type \hat{a} at date $t + \tau$. Following this last type transition the owner receives a satisfactory bid from an investor of type b . From the preceding discussion, we can calculate the

¹⁶As with Equation (12), it is implicitly assumed that if buyer and seller are indifferent to the transaction, i.e., $Q_i - Q_a = c$, the property is sold.

unconditional probability of such a path as

$$\pi(o, a, \hat{a}, b, \tau) = \begin{cases} \pi_o^O \pi_a^U \left(\sum_{a' \in \mathcal{A}} (\hat{\mathcal{P}}_{\mathcal{A}}^{\tau-1})_{aa'} \Pi_{a'\hat{a}}^T \right) \pi_b^U & \text{if } Q_a - Q_o \geq c \text{ and } Q_b - Q_{\hat{a}} \geq c \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

The two terms in the parentheses appear because one must account for the path from t until $t+\tau-1$, during which the property receives offers and isn't sold, as well as the last transition at date $t+\tau$ (captured by $\Pi_{a'\hat{a}}^T$).

Observing a holding period return is tantamount to observing one of these paths. Thus, the distribution of the purchasing and selling shocks in (11), holding $\tilde{\lambda}$ and $\tilde{\lambda}'$ constant, and conditioning on the observation that a repeat transaction took place, is

$$\pi(o, a, \hat{a}, b, \tau \mid \tilde{\lambda}, \tilde{\lambda}', Q_a - Q_o \geq c, Q_b - Q_{\hat{a}} \geq c) = \frac{\pi(o, a, \hat{a}, b, \tau)}{\sum_{o, a, \hat{a}, b \in \mathcal{A}} \pi(o, a, \hat{a}, b, \tau)}, \quad (17)$$

Equation (17) completes the specification of the probability law for the holding period log-return equation (11).

3.2 Other transaction statistics

The NCREIF data from which holding period price-appreciation returns are calculated can be used to estimate other transaction statistics. Among these are: (1) the acquiring transaction's capitalization rate or "cap rate" (the ratio of the forecast of the next twelve months' property income to the price), (2) the transaction costs reported by owners when a property is sold, (3) the quarterly rate of property turnover, and (4) the probability that a property will be sold within τ periods.

Turnover rate. The turnover rate can be calculated from the distribution of property ownership as follows. For each type of owner, o , first calculate the probability that o will receive a satisfactory offer. Because offer arrivals correspond to the distribution of types, π^U , this probability is

$\sum_{\substack{a \in \mathcal{A} \\ Q_a - Q_o \geq c}} \pi_a^U$. One then multiplies this quantity by the share of properties owned by o and sums

across all owner types. If the property market is at a steady state, then the distribution of owner types is π^O as calculated in (12). The steady state per-period turnover rate is therefore

$$\text{Turnover rate} = \sum_{\substack{a,o \in \mathcal{A} \\ Q_a - Q_o \geq c}} \pi_o^O \pi_a^U. \quad (18)$$

Mean transaction costs. When they sell their properties, NCREIF member firms often report both a gross and a net transaction cost. From this, one can estimate the average transaction cost as a proportion of the property value. To do this in the model, define the proportional transaction cost to equal to one minus the ratio of the net sale price to the gross sale price. In the model, when the sale takes place between an owner of type o and a buyer of type a , this is $\frac{c}{Q_o + c + \lambda(Q_a - Q_o - c)}$. Assuming once more that the distribution of owners corresponds to the steady state, then averaging over all possible owner and buyer types, and over all realizations of λ , yields

$$\text{Steady state mean transaction costs} = \frac{\sum_{\substack{a,o \in \mathcal{A}, \lambda \\ Q_a - Q_o \geq c}} \pi_o^O \pi_a^U \pi_\lambda \frac{c}{Q_o + c + \lambda(Q_a - Q_o - c)}}{\sum_{\substack{a,o \in \mathcal{A} \\ Q_a - Q_o \geq c}} \pi_o^O \pi_a^U}. \quad (19)$$

The turnover rate in the denominator reflects the fact that the mean is calculated conditional on a sale taking place.

Acquisition cap rates. In the data, it is not clear whether the acquisition price of a property purchased by a NCREIF member is reported net or gross of costs. If one assumes that it is equally likely to be reported net and gross of costs, then one can calculate,

$$\text{Mean acquisition cap rate} = \frac{\sum_{\substack{a,o \in \mathcal{A}, \lambda \\ Q_a - Q_o \geq c}} \pi_o^O \pi_a^U \pi_\lambda \frac{1}{2} \left(\frac{4}{Q_o + \lambda(Q_a - Q_o - c)} + \frac{4}{Q_o + c + \lambda(Q_a - Q_o - c)} \right)}{\sum_{\substack{a,o \in \mathcal{A} \\ Q_a - Q_o \geq c}} \pi_o^O \pi_a^U}. \quad (20)$$

One half times the term in parentheses averages over the net and gross acquisition cap rates. The factor of four arises because the Q_a 's are quarterly. As with the previous two calculations, the independent probabilities π^O , π^U and π_λ correspond respectively to the steady state distribution of property ownership, the unconditional distribution of offer types, and the distribution of $\lambda \in [0, 1]$. Here too the denominator reflects the fact that the mean is calculated conditional on observing a transaction.

Distribution of holding periods. The probability that a property, purchased at date t by an investor of type $a \in \mathcal{A}$, will be sold by date $t + \tau$ can be gleaned from the a^{th} element of the last column of the matrix \hat{P}^τ defined in (14). To obtain the probability that a property, bought by any type at date t , will be sold within τ periods, one calculates (similar to the calculations above)

$$\text{Steady state mean transaction costs} = \frac{\sum_{\substack{a,o \in \mathcal{A} \\ Q_a - Q_o \geq c}} \pi_o^O \pi_a^U \hat{P}_{a,2S+2}^\tau}{\sum_{\substack{a,o \in \mathcal{A} \\ Q_a - Q_o \geq c}} \pi_o^O \pi_a^U}. \quad (21)$$

4 Fitting the model to the data

The purchasing and selling shocks in (11) 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 be positive for short holding periods 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.¹⁷ In this case, the unconditional distributions of Q_o , Q_a , $Q_{\hat{a}}$ and Q_b in (11) are independently distributed as π^U , thus rendering the purchase and sale shocks independent. The selling shock component in (11) 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.¹⁸ Because $\tilde{A}' + c$ first-degree stochastically dominates \tilde{A} , it must be that $E[\ln \frac{\tilde{A}}{\tilde{A}' + c}] \leq 0$. In other words, without persistence, it is not possible to generate a positive average holding period return for short holding periods. If $c = 0$ then one obtains the Goetzmann (1993) and Case and Shiller (1987) setting in which holding period returns exhibit two iid shocks (when the property is bought and subsequently sold). For any finite $c > 0$, however, expected returns from the purchase and sale shocks are negative, suggesting that their reduced-form model is not only inconsistent with the data but may lack a microeconomic foundation.

Axelsson, Sorensen, and Stromberg (2015) add an “alpha” to the independent jump processes at purchase and disposition considered by Goetzmann (1993) and Case and Shiller (1987). This

¹⁷Recall that $\Pi^T = P^n$, where P is define by Equation (7). As $n \rightarrow \infty$, all rows of Π^T approach the unconditional distribution of types, π^U .

¹⁸The random variables \tilde{A} and \tilde{A}' have the same distribution as $Q_o + \tilde{\lambda}'(Q_a - Q_o - c)$ conditional on $(Q_a - Q_o - c) \geq 0$.

however, may not be consistent with the absence of arbitrage because the implied Sharpe ratio does not vanish with the holding period, meaning that infinite profit can be obtained by a sequence of purchase and sales of arbitrarily small holding periods.¹⁹ In the arbitrage-free model offered here, the “jumps” at purchase and disposition are generally not independent because the observed holding periods are endogenous. Moreover, as explained next, the positive alpha in *observed* transactions arises endogenously as well and a reduced-form description can miss important structure that is informative about the underlying asset market.

Beside persistence, a necessary condition for modeling a positive alpha for short-hold transactions is the existence of more than two types of investors/owners. To see this, consider that if there are only two types and $c > 0$, a purchase can only take place if the high valuation type purchases from a low valuation type. A subsequent sale can only take place if the high valuation type transitions to a low type and sells to a high type. The net selling prices in the repeat transaction just described will necessarily have the same distribution because in each case a low type is selling to a high type. Similar to the argument above when there is no persistence, transaction costs will lead to negative expected alphas in the observed transaction *regardless of the horizon*. Thus, positive short-hold alphas require that the cardinality of \mathcal{A} is greater than two, which rationalizes setting the number of types to be $2S + 1$ with S a positive integer.²⁰

To see why both persistence and the presence of intermediate types might lead to positive alphas in observed transactions, consider a situation with three highly persistent types, $Q_H > Q_M > Q_L$. In this case, when an asset is first purchased at date t by a type a investor, it is most likely that the type remains a at $t + 1$ just before offers are received. Because the presence of transaction costs ensures that a sale will not be consummated between two investors with the same valuation,

¹⁹One could avoid arbitrage opportunities in the reduced-form model of Axelson, Sorensen, and Stromberg (2015) by restricting holding periods to be no less than some finite value τ^* . Additional requirements would then have to hold to ensure that it is never optimal to hold the asset for less than τ^* (which also seems to be inconsistent with the data). Such exogenous requirements limit one’s ability to shed light on the economic fundamentals (i.e., illiquidity) driving the jumps and the positive zero holding period alpha. I note that the model in this paper can be extended to continuous time (without allowing for arbitrage opportunities) by using a Poisson arrival process for the bids.

²⁰Heterogeneity across investor types also plays a key role in the search and bargaining model of Hugonnier, Lester, and Weill (2014).

a short holding period is most likely to be characterized by the arrival of a new buyer with an even higher valuation. Thus, an intuitively 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 — i.e., it represents the upside for having recently purchased the asset. The downside is not observed because the property will be otherwise held longer. This example demonstrates how, in equilibrium, transactions data in the model exhibit a selection bias. Moreover, it highlights how intermediate investors act as intermediaries in the property market (or, alternatively, property “flippers”). The presence of such intermediaries and the profits that they make in equilibrium contain information about structural illiquidity in the market, which is most pronounced in short-term transactions.²¹

4.1 Model parameter calibration

To get a better sense of whether the model delivers a plausible quantitative description beyond the qualitative effects outlined above, I undertake a fitting/calibration exercise where each period represents a quarter. I set $S = 5$ so that there are 11 type states, and fix $n = 4$ in $\Pi^T = P^n$, with P defined in (7). To simplify the calculation of model statistics, I assume that $\tilde{\lambda}$ is one or zero with 50%-50% probability. Three parameters are used to characterized the η_a 's:

$$\eta_a = (\eta_H - \eta_L) \left(\frac{2S + 1 - a}{2S} \right)^v + \eta_L, \quad a \in \{1, \dots, 2S + 1\}$$

where $v > 0$, η_H corresponds to the largest growth-adjusted capitalization factor, $e^{r_H - \mu}$, while η_L corresponds to the smallest. When $v = 1$, the η_a 's are distributed linearly across the $2S + 1$ states. When $v > 1$ or $v < 1$, the distribution is respectively, strictly convex or strictly concave. Beyond η_L, η_H and v , the remaining three free model parameters are x , c , and σ_I .²² To calibrate to the

²¹The model in Lovo and Spaenjers (2016) similarly delivers a negative relationship between holding periods and returns.

²²As assumed in the derivation of (11), the adjustment for systematic risk in Section 2 obviates the need to calibrate μ .

data, I minimize a function of the following type:

$$F_O(x, c, \sigma_I, \eta_H, \eta_L, v) = \sum_{j=1}^M \frac{(m_j - \hat{m}_j(x, c, \sigma_I, \eta_H, \eta_L, v))^2}{s_j^2},$$

where m_j is one of M quantities estimated from the data, s_j is its standard deviation, and $\hat{m}_j(x, c, \sigma_I, \eta_H, \eta_L, v)$ is the model-calculated value of the quantity. The objective function $F_O(x, c, \sigma_I, \eta_H, \eta_L, v)$ is akin to a sum of squared standardized, but correlated, normal variables — or, in other words, a generalized χ^2 statistic. The analogy, however, is not exact because some of the quantities to which I calibrate are themselves estimates from the econometric model of Section 2.2. In other words, the exercise undertaken here should be understood as a calibration rather than an estimation. In particular, I do not derive standard errors for the calibrated model parameters.

Setting $M = 9$, I attempt fit the model to data statistics derived in Sections 3.1.1 and 3.2: Quarterly turnover, Average transaction costs, Average acquisition cap rate, Fraction sold within five years, the expect value of one, six, and eight year risk-adjusted holding period returns, and the variance of one and eight year risk-adjusted holding period returns. I choose three expected return horizons because of the non-linearity in the expected return detected in Section 2.2. I choose two return variance horizons because the empirical holding period return variance is roughly linear in horizon. The discussion in Appendix A.1 provides details for calculating these statistics from the data. It turns out that, at least in matching the selected quantities, decreasing v corresponds to a narrower interval defined by η_L and η_H . The objective function, $F_O(x, c, \sigma_I, \eta_H, \eta_L, v)$, is not very sensitive to $v \in [0.05, 1]$. I therefore fix $v = 0.50$ and minimize over the remaining unfixed parameters to arrive at $F_O \approx 3.02$.²³ The calibrated model parameters are reported in Table 3. The data and model statistics are compared in Table 4.

[Table 3 about here.]

[Table 4 about here.]

²³If F_O was indeed a generalized χ^2 distribution, then the corresponding p -value would be at least as high as a standard χ^2 with nine degrees of freedom. With $F_O = 3.02$, this translates to $p > 0.96$, suggesting that the fit to the data is excellent.

The model fit to the data is very good. Figure 3 plots the model predictions along with the predictions of the two-pass regression model from Section 2. At every horizon, the calibrated model prediction is within the 95% confidence interval about the econometric estimates. Although the calibrated model does not incorporate parameters that depend on calendar time, it fits the econometric model predictions which do include calendar year effects. This raises the possibility that the calendar year effects are significant because without them the econometric model is not sufficiently non-linear.

[Fig. 3 about here.]

Figure 4 is a histogram of the steady-state property ownership for the calibrated model. Types are denoted by their private valuation annualized cap rate, $\frac{4}{Q_a}$. These range from 6.1% for the highest valuation type to 9.9% for the lowest valuation type. Although the highest valuation type represents only 2.5% of the investors in the economy, this class of investors represents nearly 50% of the steady state ownership base. The histogram makes clear that virtually all properties are owned by the top four valuation types comprising 17.5% ($= \sum_{a=8}^{11} \pi_a^U = 5\% + 5\% + 5\% + 2.5\%$) of the investor base. Turnover is low because types are persistent — each investor has only about a 6.3% per quarter of changing type — and because property owners correspond to the highest valuation quintile.

[Fig. 4 about here.]

The weighted average private value of properties in the economy is given by $Q' \cdot \pi^O = 60.14$, corresponding to an annualized cap rate of 6.65%. This represents a premium of $\frac{60.14}{(4/7.04\%)} - 1 \approx 5.9\%$ over the average transaction. The reason for the premium is simply that the highest valuation types, comprising a large proportion of owners, never sell. They do, however, buy at prices that are on average below their private values. An implication is that an index based on sales will understate the true value of the property market to its steady state owners (i.e., the sum of private values across all owners).

4.2 Transaction Risk and Model Implications

In the limit of a liquid property market (e.g., $c = 0$ and $r_a = r$ for all investors), property transaction prices are consistent with a dividend-discount model and the income-to-price ratio, or cap rate, is objective and constant at $\frac{1}{Q} = \eta - 1 = e^{r-\mu} - 1$. In the presence of illiquidity, the cap rate will be a distribution corresponding to the private values of investors in the economy. This distribution drives the transaction risk in the model and the calibration to the holding period return data allows one to quantify it. In this subsection I attempt to do that in three different ways.

Transaction probability. Figure 5 plots the probability that an owner will receive a satisfactory bid, given a reservation cap rate gross of costs and within one period (a quarter). The probability is calculated using the model parameters from Table 3, the equilibrium solution to the Q_a 's, and the probability π_a^U of the arrival of a bid from type a . For instance, the probability that a bid will arrive at a cap rate lower than 9.033% is 45% because there are five types whose private values correspond to a cap rate lower than 9.033%, and $\sum_{\frac{4}{Q_a} \leq 9.033\%} \pi_a^U = 0.45$. The greater the reservation cap rate, the greater the probability that a bidder will arrive who will transact at a satisfactory price. Because the number of types in the model is finite, the model probability of transaction is a step function. The points at which the step function “jumps” are plotted as red squares in Figure 5. The continuous curve in blue is a least squares fit to the red squares taking the form

$$\text{Prob}(\text{cap rate bid} \leq x) = \exp\left(\sum_{n=0}^3 a_n x^n\right),$$

where x is the seller's reservation cap rate gross of costs (i.e., the seller has to agree to pay the transaction costs out of pocket).

For comparison, also plotted in Figure 5 is the corresponding transaction probability for a perfectly liquid asset whose annualized market income to price ratio is 0.0704, coinciding with the unconditional expected transaction cap rate in the property market calculated in (20) using the calibrated model parameters. For a perfectly liquid asset, an offer to sell at an income-to-price ratio higher than market would be instantly transacted with probability one. An offer to sell at a lower income-to-price ratio would never be transacted.

[Fig. 5 about here.]

The Figure clarifies that transaction risk is substantial. The probability of transacting at a “fair market price” of 0.0704 (or better), which might correspond to an appraised value in the property market, is only 12% per quarter. Owners with a high cost of capital would therefore be forced to list at higher cap rates (lower prices) in order to transact with a high probability. To achieve a probability of sale per quarter greater than $\frac{2}{3}$ a seller must be prepared to accept a cap rate bid as high as 0.0946. Such a “fire sale” amounts to a discount of 25% relative to average prevailing transaction prices. The ability to quantify such discounts can be useful in pricing mortgages or CMBS securities, in which the foreclosing entity has a high cost of capital and/or lacks the expertise to manage the asset while searching for a liquidation opportunity. In other words, the equilibrium curve in Figure 5 should be considered when calculating recovery rates and is linked to the cost of capital of the foreclosing institution. Correspondingly, it may be possible to use the model calibration to back out the implicit cost of capital of lending institutions.

A testable implication of the shape of the transaction risk curve in Figure 5 is that property transaction risk has a pronounced negative skew with respect to cap rates, which translates to a positive skew with respect to prices. A deeper exploration of this point may provide insights into property portfolio and risk management, the pricing of contingent claims such as mortgages, and managerial incentives in a private equity setting where waterfall structures are prevalent.

Time to sell. A related notion to the probability of transaction, is the expected or median time to sell. If p is the probability of a serially independent event per period, then the expected number of periods to first occurrence of the event is $\frac{1}{p}$. The model is based on quarterly periods so the annualized time to a sale is $\frac{1}{4p}$, where p is the probability of a sale per quarter. Likewise, the median time to sell is $\frac{\ln .5}{4 \ln(1-p)}$. Figure 6 plots the mean and median time to sell given the seller’s reservation cap rate (solid blue and dashed red curves). Also plotted, and corresponding to the right vertical axis, is the expected transaction cap rate for a seller, given their reservation cap rate. All of the curves are calculated off the continuous per-period transaction likelihood in Figure 5. For reservation cap rates below 0.0704, the prevailing transaction rate in the market, it can take years

to transact. The expected cap rate corresponding to such a reservation value is necessarily lower than prevailing cap rates. Thus patient sellers trade immediacy for lower cap rates (i.e., higher expected transaction prices).

[Fig. 6 about here.]

The large circle denotes the reservation cap rate at which the seller can expect to transact at prevailing market prices. When fit to the holding period return data, the model predicts that a seller expecting to transact at prevailing market prices must also be prepared to sell at a cap rate as high as 0.0835 — significantly above market rates. The time to sell curves predict that such a seller can expect to sell within 0.9 years. Although the NCREIF data does not report time to sell, the predicted figure is consistent with CoStar data reporting average “days on market” figures varying from 250 in January 2007 to as high as 450 in July 2012, and then coming back down to 284 as of November 2016.²⁴

Need for immediacy. Yet another measure of property market illiquidity that can be inferred from the model is the “price of immediacy”. I define a seller to have a need for immediacy of $q/4$ years whenever the seller must liquidate a position within q quarters. The corresponding price of immediacy is the discount relative to prevailing market prices which the seller expects to experience. I calculate this quantity as follows. Let $Q^{(0)} \equiv 0$ and recursively define

$$Q^{(n)} = E[\{\tilde{Q}, Q^{(n-1)}\}^+] = \sum_{a \in \mathcal{A}} \pi_a^U \{Q_a, Q^{(n-1)}\}^+.$$

Then $Q^{(1)}$ is just the expected price to income ratio if a seller must liquidate in one period. One can interpret $Q^{(2)}$ as the expected transaction price for a two-period strategy in which one sells in the first period only if a bid higher than $Q^{(1)}$ arrives; otherwise, the first-period bid is ignored and one accepts whatever bid comes along in the second period. For $q > 2$, $Q^{(n)}$ is similarly defined for an n -period strategy. There is no time (or risk) discounting in this calculation because discount factors are heterogeneous in the model and thus there is no objective strategy for accepting or rejecting a bid.

²⁴See <https://goo.gl/mc1OpM>.

[Fig. 7 about here.]

The price of immediacy is defined as $1 - \frac{Q^{(n)}}{Q^*}$, where Q^* is price corresponding to the average prevailing transaction cap rates ($\frac{4}{0.0704} = 56.8$ in the calibrated model). Figure 7 plots this for the model parameters in Table 3. A seller can expect to experience a discount to prevailing prices if (s)he must liquidate within three years or less. Given more flexibility to “wait out the market”, a seller can expect to perform at least as well as average prevailing transactions.

The closed-end private equity fund structure, frequently used by commercial real estate investors, commits the general partner to liquidate the portfolio by a certain date. It is quite common, however, for the general partner to exercise a two to four year option to wait out the market. Thus, here too the model appears able to provide sensible quantitative predictions.

5 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. Using purchase and sale data from the National Council of Real Estate Investment Fiduciaries (NCREIF) to compute and analyze holding period returns for commercial properties, I find that the data is not consistent with the joint hypothesis that individual real estate property prices follow a geometric random walk *and* that transaction decisions are unrelated to the underlying price process. This conclusion is robust to controlling for heterogeneity in random walk parameters across time and across properties. In addition, the data also appears inconsistent with the random walk assumption when accounting for possible selection biases in transaction data.

The data is consistent, however, with a calibrated search-based illiquid asset pricing model. In the model, owners periodically receive bids for their property from investors but gains from trade only exist if the valuation of bidders (net of transaction costs) exceeds that of owners. Holding period returns therefore exhibit selling and purchasing shocks that arise from the random matching and bargaining. If private valuations are persistent, observed short horizon holding periods will

exhibit a positive “alpha” in equilibrium because a short hold will only be observed when an owner receives a bid significantly higher than the price recently paid for the property. This is consistent with the data. Both the model and the data imply that idiosyncratic property risk comes in two forms: a volatility component similar to that exhibited by liquid assets measuring roughly 10% per year (annualized); and a purchasing/selling shock that, in an optimal transaction, has a variance of roughly 0.01. Thus round-trip observed transaction returns exhibit a “time-independent” variance component of roughly 0.02.

The model can be applied to the pricing of 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.

A Appendix: Cleaning the Data

I provide here details of how the raw NCREIF data is cleaned for use in Section 2. Table A.1 documents the sequential steps taken in culling properties from the dataset. The table reports the number of properties in the dataset left after the corresponding step is implemented. The steps are explained below.

1. The total number of properties in the dataset.
2. No missing appraisal returns in any but the last period in the panel. Sales are often reported in the “last period” that a property is reported in the dataset without reporting any cash flows or other economic performance measures. In such cases, the sale is assumed to take place at the end of the previous quarter.
3. Each property should be characterized by a unique “SaleCode” identifying its status as of the last reporting period. E.g., “F” if it is ongoing in the dataset, “S” if it is separately sold, “P” if it is sold as part of a portfolio, “R” if it is surrendered to the bank in foreclosure, “T” if it is transferred to another manager, etc. All but “F” correspond to some type of exit from the dataset and the SaleCode is updated and backfilled each quarter. Properties whose SaleCode varies while it is in the dataset are dropped. In addition, most of the analysis only includes “S”, “F”, “P” and “R” properties, the vast majority of which are “S” or “F”.
4. If an appraisal-based return measure (“total”, “price appreciation”, or “income”) in any quarter is below -0.999, then the property is dropped.
5. The acquisition date, redundantly reported each quarter, should be available and consistently reported (i.e., the same in each quarter and no later than the first period that the property appears in the dataset).
6. The property type should correspond to one of the major commercial real estate asset categories (Apartments, Industrial, Retail, or Office) and should not change across reporting periods.

7. The net sale price for sold properties should be reported as a number greater than \$1.
8. According 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 net sales are reported in the quarter of final disposition, they should coincide with the net proceeds from disposition. Reported partial sales should be positive.
9. It is standard accounting practice to report capital expenditures that are allocated but not actually spent. Funds that are not spent should be eventually recorded as negative capital expenditures. If a property is sold, then the sum of capital expenditures across all reporting periods should be positive.
10. The initial acquisition price should be positive.
11. The total price appreciation return, as defined in (1), should be well defined (i.e., the argument of the logarithm should be a positive number).
12. Out of concern that important initial capital expenditure data is missing when there is a gap between the acquisition quarter and the first reporting quarter, I drop properties that exhibit such gaps. This results in a loss of nearly half of the properties remaining after the previous filters are implemented. About a third of the properties report the acquisition date a year or more before the first reporting period.

[Table A.1 about here.]

A.1 Estimating turnover, cap rates, transaction costs, and ‘survival’ rates

Figure A.1 plots calendar year statistics for some of the different data attributes used to calibrate the model in Section 3: The proportion of properties sold within five years, acquisition cap rates, proportional transaction costs, and quarterly turnover. From the plots, it is clear that these characteristics of the commercial property market exhibit a great deal of time variation, suggesting the

presence of significant correlations across properties and time. Calculating unconditional moments for each characteristic would therefore drastically overstate the accuracy of the point estimate. To get around that, I calculate the time-series averages of the points in each of the plots, weighting each point depicted by the number of properties it represents. I then calculate the property number-weighted time-series estimation error. The results are used in Table 4.

[Fig. A.1 about here.]

The remaining statistics used for calibrating the model correspond to the expect value of one, six, and eight year risk-adjusted holding period returns, and the variance of one and eight year risk-adjusted holding period returns. For each of these, point estimates and corresponding standard errors are calculated using the predictions of the four-stage econometric model in Section 2.2.

B Appendix: Robustness Exercises

Here, I provide additional evidence that the results from Section 2 are not spurious. Firstly, I repeat the econometric estimation with randomly sampled real estate investment trust (REIT) returns substituted for the property holding period returns. In this case, I find that the null is not rejected, suggesting that any discrepancy is attributable to the properties of repeat transactions of individual real estate assets rather than the real estate market per se. Then, I demonstrate that the amount of transaction risk (i.e., σ_0^2) identified in 2 can be reproduced using an alternative specification, thus reinforcing the original findings. Finally, I survey some additional experiments involving the tightening or loosening of the filtering criteria discussed in Appendix A.

B.1 REIT-simulated holding period returns

Each property's return in the dataset is replaced with a matched unlevered return constructed from randomly selected publicly traded real estate investment trusts (REITs). For instance, if a property's first and last quarters in the dataset are 1997Q3 and 2001Q1 then each quarter, starting

with 1997Q3 and ending with 2001Q1, a quarterly return is randomly drawn from REITs trading at the beginning of that quarter. The sequence of randomly drawn REIT unlevered returns is then compounded to yield a matched holding period return for that property. Likewise, the market benchmark is based on a return index of equally weighted unlevered REIT returns in each quarter. To unlever, I used Moody’s BBB bond total return index. To reproduce the analysis in Table 2 and Figure 1 I simulate a panel and run the corresponding regressions. This is repeated 1000 times and the resulting statistics are pooled. The results are presented in Table B.1 and Figure B.1.

Presumably, any systematic changes in property market characteristics that might lead to spurious conclusions in the analysis of Section 2 would also affect unlevered REIT returns. Comparing the corresponding tables and figures, this seems not to be the case. The simulated REIT panel exhibits virtually no alpha in expected returns let alone one that is non-linear in the holding period, and the variance decomposition suggests that there is no significantly positive contribution coming from any source other than a standard diffusion. The exercise confirms that, as with the physical property market, the publicly traded property market exhibits time variation in the diffusion parameters.

The punchline is that the null hypothesis of Section 2.2 is confirmed, meaning that publicly traded real estate is priced consistent with a random walk with time varying coefficients. This suggests that the “anomalous” behavior of holding period returns in the physical market arises from endogenous considerations and/or illiquidity.

[Table B.1 about here.]

[Fig. B.1 about here.]

B.2 A test of matching return horizon

To further rule out vintage effects 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 purchase/sale decisions are unrelated to the price process, the random walk model predicts no difference between the actual and the synthetic eight year investment — on average both are exposed to the same time-varying economic conditions. For various values of k and τ , I use a linear program to “vintage-match” the largest possible number of properties having holding period $k \times \tau$ with k properties of (non-overlapping) holding periods τ . In the illustrative example, $k = 2$ and $\tau = 4$. Two equally-sized samples are created in this way: One containing the holding period returns of selected properties with maturity $k \times \tau$, and one containing the corresponding vintage-matched synthetic holding period returns. When the matching entails a strict subset of eligible properties, the subset is chosen randomly. The regression in (6) is then estimated for each one of the two matched samples, and the difference between the residual variances as well as between the intercepts is calculated. To increase the accuracy of the estimates, and allow for all property returns to enter the estimation (through randomization), the procedure is repeated 1000 times and the resulting statistics are averaged.

[Table B.2 about here.]

Table B.2 documents the results of the vintage-matched sample analyses with N corresponding to the size of the sample where each property has maturity of $k \times \tau$ years. The difference in regression residual variances and the difference in average returns (or “ α ’s”) are normalized by $k - 1$ to account for the fact that each time a property is rolled over the transaction risk component (if any) of return variance and mean should be accrued. Under the null, both differences should be zero. The t -statistics, which reflect pooled values of the coefficients and their standard errors, are likely to be conservative.

Despite the relatively low average number of matches, the test has sufficient power to measure the variance of the transaction risk component (“(Var Diff)/($k - 1$)”). All but one of the comparisons correspond to a positive difference between the property rolling strategy and the single property strategy. A χ^2 test that all the variance differences are zero is rejected with $p < 0.01$.

Moreover, the GLS estimate of the average normalized variance difference yields a number (0.0342) that is statistically consistent with the estimate of σ_0^2 from Table 2. Unfortunately the test does not have sufficient power to identify differences in the slope coefficients (the alphas). One cannot reject the null that they are all identical and equal to zero. Still, the upshot of this exercise is that the null hypothesis of zero variance difference is soundly rejected, providing an alternative channel of support to the findings of Section 2.

B.3 Additional tests

The application of the econometric specification in (4) seems quite robust to different ways of treating the NCREIF dataset. I examine the following alternatives (details of the estimates are available upon request).

1. Exclude all properties that appear in the dataset before 1997. This is done to address concerns about the quality of the data in the earlier period.
2. Include non-NPI properties.
3. Use *only* non-NPI properties.
4. Include properties that are sold less than a year after purchase.
5. Permit the acquisition quarter to be up to two quarters before the first reporting quarter (thus allowing a gap of one quarter in the property history).

In every case, the null hypothesis of 2.2 is rejected with a very high degree of confidence. In particular, each of α_0 , α_1 , and σ_0^2 is highly significant and the coefficients share the same signs across all specifications, meaning that the qualitative results are robust. Restricting analysis to non-NPI properties corresponds to the most dramatic magnification of the effect. This is intuitive given that NPI properties are generally deemed less risky than their counterparts in the data. Including properties that are sold almost immediately after purchase results in the most pronounced decline in the effect (though it remains significant). This too is consistent with the concern outlined earlier

that such properties are inherently “noisy”, perhaps because many are purchased as part of a larger deal and are quickly shed by their new institutional owners. In such instances, the purchase price and other accounting details may be subjectively assessed, possibly introducing either noise or bias.

C Appendix: Testing Alternative Hypotheses

Many studies in the literature model the price of a single real estate property as a geometric random walk. Section 2, and the additional robustness exercises in Appendix B, cast some doubt on such a modeling approach. The question remains whether it is possible for the underlying price process to be a geometric random walk but for the decision to purchase or sell to be dependent on the price history.²⁵ Specifically, the econometric tests reject a joint null but not the possibility that the property price is random walk *and* τ in equation (2) depends on $\varepsilon_{i,\tau}$ or $r_{m,\tau}$. In particular, I focus on alternatives that can explain the robust finding that as the holding period decreases the holding period variance does not proportionately decrease.

In this section, I demonstrate that property-level measures associated with higher risk have little ability to predict shorter holding horizons, thus dispensing with one potential form of sample selection bias. Then, while I find some evidence that selling is linked to poor performance, there is also evidence that this endogeneity is not enough to explain the large and significant time-independent variance component.

C.1 Endogeneity I: Investor risk preferences

Suppose that property prices follow a geometric random walk and that investors prefer to hold riskier properties for shorter periods of time. Then, when analyzing holding period returns, this might give rise to the non-zero intercept in Figure 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

²⁵See Fisher, Gatzlaff, Geltner, and Haurin (2003), Gatzlaff and Haurin (1997), Goetzmann and Peng (2006), and Hwang and Quigley (2004).

the dependent variable is set to one if a property is within four quarters of its final sale. I choose a lag of four quarters because the decision to sell a property typically takes place months before an actual sale is consummated. Both sold and unsold properties in the data are used. 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. This information would in principle 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 joint venture dummy variable. A joint venture ownership structure might be associated with a need for risk sharing and thus linked to the property's riskiness.

Age when acquired — Older properties may have more idiosyncratic risk because capital expenses are less predictable.

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 might be characterized by different risk attributes.

Region — East/Midwest/South/West. Supply and demand varies by region, potentially resulting in different growth risk.

Appraisal-based lagged return — The accumulated (appraisal based) total return on the property lagged four quarters. This is included to control for performance motives for selling.

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. Remaining lease duration, an important risk measure that is not available in the data, will be captured by the annual independent appraisal valuations and thus impact appraisal time series variance.

Fund Type — Many properties are held through funds. Closed-end funds (CEFs) are finite-lived entities and have target liquidations dates, while open-end funds (OEFs) are essentially infinite-life investment vehicles. The ODCE (Open End Diversified Core Equity) funds comprise a subset of the OEF properties in the NCREIF dataset.

[Table C.1 about here.]

Table C.1 reports the results of a logistic regression, with and without year, fund-type, and manager fixed effects. For continuous variables, the “Marginal impact” column reports the change in the probability of sale 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.

When significant, 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 commercial property type, and yet in the data their average holding period is one year shorter than other asset types. 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 marginal impact of the the variables with the “correct” sign (i.e., size, JV, age, occupancy, and relative loan spread) declines from 23.9% to 5.3% when year, fund-type, and manager fixed effects are added. Much of the explanatory power is taken up by the type of fund in which the property is held. Not surprisingly closed-end funds, typically committed to a liquidation date, are substantially more likely to sell than an open-end fund (ODCE or otherwise). When year, fund, and manager effects are taken into account the total impact of the variables with the “correct” sign is comparable to that of the ones without the correct sign (Idiosyncratic variance and Apartments) at 4.4%. Moreover, in the latter case the impact of both sets of characteristics on disposition is swamped by that of property-specific variables not obviously related to property-specific risk: The lagged total appraisal return and the percentage of systematic variance comprising the property’s total appraisal return. The sum of absolute marginal impact of these two variables is about 14%.

Once one controls for year, fund, and manager effects, 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 of returns relative to the corresponding NPI benchmark, the more likely it is to be sold.²⁶ That said, the effects studied in Section 2 are *net* of a systematic risk exposure, 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 in Table C.1 .

Summarizing, while some risk-related property-level variables do predict disposition, this is not consistent across measures and the economic magnitudes do not appear decisive. In particular, there is no compelling evidence demonstrating a link between the magnitude of idiosyncratic risk

²⁶A 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 regression. To further allay such concerns, the adjusted R^2 is only calculated when the number of observations available for the time-series regression equal to or exceeds 10. I substitute the median adjusted R^2 of the former for the remaining properties in the database. Whether it is followed or not, this procedure does not qualitatively alter the results of the logistic regression.

and shorter holding periods. Under the assumption of a random walk price process, it does not appear that the econometric model estimation in Section 2 is an artifact of a preference among managers for holding properties with less idiosyncratic risk over longer periods.

C.2 Endogeneity II: Strategic asset disposition and holding period returns

At any given point in time, 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 holding period returns exhibit a selection bias that masks the underlying geometric random walk price process. A natural economic link between τ and idiosyncratic return 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.²⁷ 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 C.1 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.²⁸ If properties are mostly sold because they underperform up to some threshold, then the risk-adjusted price paths associated with sold properties will only portray a strict subset of possible price paths and thus exhibit a lower variance. One possibility 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

²⁷Properties 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 financially constrained entities. The NCREIF member firms 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.

²⁸An alternative to the updating story is disposition due to tax loss selling (Constantinides, 1983, 1984; Dammon and Spatt, 1996; Williams, 1985). This, however, is unlikely to be the culprit here because NPI qualifying properties are held by or on behalf of tax-exempt institutional investors.

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. Assuming that dr_t is consistent with a random walk, the goal is to see whether selection bias can lead to a holding period return variance that is not proportional to the holding period.

Suppose that the investment's abnormal returns or $\alpha(t)$ is a constant that it 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

$$dr_t^O = a_t dt + \sigma_W dW_t \tag{22}$$

$$da_t = \frac{\sigma_W}{\kappa + t} dW_t, \quad \kappa = \frac{\sigma_W^2}{\eta^2}, \tag{23}$$

where $a_0 = \alpha_0$ and $r_0^O = W_0 = 0$. 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 and $r_0^U = 0$.²⁹ Thus, the evolution of the total returns realized when the property is sold is

$$dr_t = dr_t^O + dr_t^U = a_t dt + \sigma_W dW_t + \sigma_Z dZ_t, \tag{24}$$

²⁹One can also consider r_t^U to be observed by the manager but containing no information about $\alpha(t)$.

where $r_0 = 0$ and $Z_0 = W_0 = 0$, and where a_t evolves according to (23). 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.

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

$$r_\tau = \sigma_Z Z_\tau - \kappa \hat{\alpha} + \alpha_L \tau. \quad (25)$$

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

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 (25). The expected value and variance of r_τ follow from the fact that Z_t is independent of τ . \square

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. It may be possible, however, that variability in $\kappa \hat{\alpha}$ across investors leads to a perceived time-independent variability in the *cross-section* of holding period returns.³⁰ This, however, implies a counter-factual

³⁰Assume κ across investors is on average equal to ten, corresponding to an investor with ten years of experience for forming a prior, and that $\hat{\alpha}$ is on average 5%. Assume further that across investors $\kappa \hat{\alpha}$ is uniformly distributed from 25% to 75%. This would be sufficient to deliver a time-independent variability of 0.021, in line with the observed quantity.

prediction. Specifically, the additional variation should also appear in the observable cumulative returns (i.e., r_t^O) 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 C.1 depicts the results of repeating the variance analysis of Section 2.2 using appraisal-based returns accumulated until one year before the sale. It is evident that a positive intercept component, $\sigma_0^2 > 0$, is missing in comparison with the analogous Figure 1.

[Fig. C.1 about here.]

Another problem with the predictions of this model is that expected holding period returns are negative for short holding periods ($\tau \rightarrow 0$) in equation (25). This too is not in line with the findings of Section 2.2.

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

$$\text{VAR}[r_\tau] = \sigma_Z^2 \tau + p(\tau)(1 - p(\tau))(\tau + \kappa)^2(\alpha_H - \alpha_L)^2, \quad (26)$$

$$E[r_\tau] = -\kappa\alpha_0 + (\kappa + \tau)\left(p(\tau)\alpha_L + (1 - p(\tau))\alpha_H\right), \quad (27)$$

where $p(\tau)$ is the probability that a sale occurs because of underperformance rather than because of overperformance.³¹ The expression for the variance is non-zero at $\tau = 0$, thus mimicking the effect of σ_0^2 in (5). Given sufficiently high precision, κ , it may even be made to fit both Figure 1a and 1b, as long as $\alpha_0 - \alpha_L > \alpha_H - \alpha_0$. An immediate implication is the prediction of bi-modality: The holding period return distributions of underperforming and overperforming properties should separate with time.

[Table C.2 about here.]

³¹It is tedious to derive an expression for $p(\tau)$ and there is no need to do so in the ensuing analysis.

Table C.2 reports the results of the Hartigan and Hartigan (1985) “dip” test for multi-modality at each integer holding horizon and/or vintage-year combination for which at least 30 observations are available. All 37 tests, save for one, are unable to reject the null of a unimodal distribution at any conventional level. Visual inspection and kernel density estimates of the number of modes confirm the unimodal distribution of the holding period returns. The alternative model therefore appears to furnish an unlikely explanation for the findings in Section 2.2.

In summary, while models of strategic disposition under a random walk hypothesis can account for the time-independent effects, they also make counter-factual predictions about the appraisal-based and actual return distributions, thereby rejecting this joint hypothesis.


References

- Abraham, J. M., and W. S. Schauman, 1991, “New evidence on home prices from Freddie Mac repeat sales,” *Real Estate Economics*, 19(3), 333–352.
- Ang, A., and M. Sorensen, 2012, “Risks, returns, and optimal holdings of private equity: A survey of existing approaches,” *The Quarterly Journal of Finance*, 2(03).
- Axelson, U., M. Sorensen, and P. Stromberg, 2015, “Alpha and Beta of Buyout Deals: A Jump CAPM for Long-Term Illiquid Investments,” *Working Paper*.
- 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.
- Calhoun, C. A., 1996, “OFHEO house price indexes: HPI technical description,” *Office of Federal Housing Enterprise Oversight*, pp. 1–15.
- 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.
- Constantinides, G. M., 1983, “Capital market equilibrium with personal tax,” *Econometrica: Journal of the Econometric Society*, pp. 611–636.
- , 1984, “Optimal stock trading with personal taxes: Implications for prices and the abnormal January returns,” *Journal of Financial Economics*, 13(1), 65–89.
- Dammon, R. M., and C. S. Spatt, 1996, “The optimal trading and pricing of securities with asymmetric capital gains taxes and transaction costs,” *Review of Financial Studies*, 9(3), 921–952.
- Downing, C., R. Stanton, and N. Wallace, 2008, “Volatility, mortgage default, and CMBS subordination,” *University of California, Berkeley*.

- Duffie, D., N. Gârleanu, 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.
- Ghent, A., and R. Valkanov, 2014, "Comparing Securitized and Balance Sheet Loans: Size Matters," *SSRN Working Paper*, <http://ssrn.com/abstract=2152703>.
- Giacoletti, M., 2016, "Idiosyncratic Risk in Housing Markets," *Working Paper, Stanford Graduate School of Business*.
- 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.
- Goetzmann, W. N., and M. Spiegel, 1995, "Non-temporal components of residential real estate appreciation," *The Review of Economics and Statistics*, pp. 199–206.
- 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.

- Hugonnier, J., B. Lester, and P.-O. Weill, 2014, “Heterogeneity in decentralized asset markets,” working paper, National Bureau of Economic Research.
- 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.
- Lopez-de-Silanes, F., L. Phalippou, and O. Gottschalg, 2016, “Giants at the gate: Investment Returns and diseconomies of scale in private equity,” *Forthcoming in the Journal of Financial and Quantitative Analysis*.
- Lovo, S., and C. Spaenjers, 2016, “A Model of Trading and Price Indexes in the Art Market,” .
- Merton, R. C., 1974, “On the Pricing of Corporate Debt: The Risk Structure of Interest Rates,” *Journal of Finance*, 29(2), 449–70.
- Peng, L., 2015, “The Risk and Return of Commercial Real Estate: A Property Level Analysis,” *Forthcoming in Real Estate Economics*.
- Plazzi, A., W. Torous, and R. Valkanov, 2008, “The Cross-Sectional Dispersion of Commercial Real Estate Returns and Rent Growth: Time Variation and Economic Fluctuations,” *Real Estate Economics*, 36(3), 403–439.
- 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.
- Williams, J. T., 1985, “Trading and valuing depreciable assets,” *Journal of Financial Economics*, 14(2), 283–308.

Table 1: The histogram reports entries of properties into the “clean” dataset (see Appendix A) and includes both sold and unsold properties. The table reports the average vintage for various holding periods of sold properties. 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.



Holding Period	Average Vintage	Number of Props
1	2003	403
2	2001	634
3	2001	823
4	2000	685
5	1999	779
6	1999	715
7	1998	722
8	1996	685
9	1993	473
10	1992	382
11	1991	361
12	1991	283
13	1991	252
14	1991	182
15	1991	137
16	1991	124

Table 2: Panels A and B report the first and second stage regressions for estimating the model with random effects in Equation (4). The estimation procedure is described in the text. Under the null $\alpha_0 = \alpha_1 = \sigma_0^2 = 0$, while $v_{\alpha,y}$ and $v_{\sigma,y}$ are the idiosyncratic excess expected return and variance, or α and σ^2 , accumulated while holding the property for the entirety of year $y \in \{1978, \dots, 2013\}$. The terms $\frac{1}{36} \sum_y v_{\alpha,y}$ and $\frac{1}{36} \sum_y v_{\sigma,y}$ correspond to the average annualized α and σ^2 under the null. $P(v_{\alpha,y})$ (resp. $P(v_{\sigma,y})$) is the probability that all $v_{\alpha,y}$'s (resp. $v_{\sigma,y}$'s) are equal.

Panel A

	(1)	(2)
	$\frac{\hat{r}}{\sqrt{\tau}}$	Adj $\frac{\hat{r}}{\sqrt{\tau}}$
β	0.957*** (0.0632)	0.994*** (0.0593)
α_0	0.210*** (0.0384)	0.225*** (0.0401)
α_1	-0.171*** (0.0401)	-0.182*** (0.0393)
$\frac{1}{36} \sum_y v_{\alpha,y}$	0.0166 (0.01)	0.0189* (0.0094)
$P(v_{\alpha,y})$ const	<1E-6	<1E-6
Observations	3322	3322

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_α^2	-0.000432 (0.000363)	-0.000398 (0.000366)	-0.000178 (0.000277)	-0.000198 (0.000263)	
$\sigma_{\alpha\beta}$	-0.00632 (0.00619)	-0.00574 (0.00625)			
σ_β^2	-0.0391 (0.0498)	-0.0435 (0.0502)	-0.00758 (0.0316)		
σ_0^2	0.0237*** (0.00494)	0.0258*** (0.00499)	0.0276*** (0.00458)	0.0274*** (0.00453)	0.0298*** (0.00324)
$\frac{1}{36} \sum_y v_{\sigma,y}$	0.0131*** (0.0031)	0.013*** (0.0032)	0.012*** (0.003)	0.012*** (0.003)	0.0101*** (0.0015)
$P(v_{\sigma,y})$ const	0.000015	<1E-6	<1E-6	<1E-6	<1E-6
Observations	3322	3322	3322	3322	3322

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Model parameter choices

Parameter	Fixed value
$\tilde{\lambda}$	± 1 with 50/50 probability
v	0.50
Calibrated value	
σ	0.0112
x	0.00817
c	1.297
η_H	1.0737
η_L	1.004

Table 4: Comparison of statistics estimated from the data (point estimates and their standard errors) versus their counterparts calculated from the calibrated model. To control for the correlations across properties and time, the quarterly turnover, average transaction costs, average acquisition cap rates, and fraction sold within five years are first calculated by calendar year. The point estimates and standard deviations are then calculated using the corresponding time series statistics, weighted according to the number of properties in each calendar year (see discussion and plots in Appendix A.1).

Quantity	Point estimate	SD	Calibrated Model
Idiosyncratic diffusion volatility, σ_I^2	0.0101	0.0015	0.0112
Quarterly Turnover	0.031	0.015	0.031
Average Transaction Costs	0.0201	0.0062	0.0226
Average Acquisition Cap Rate	0.071	0.014	0.070
Fraction Sold Within 5 Years	0.33	0.16	0.43
Holding Period Return Statistics			
1yr Adj. Exp. Return	0.0401	0.0210	0.0397
1yr Adj. Variance	0.0310	0.0067	0.0308
6yr Adj. Exp. Return	-0.1057	0.0308	-0.0714
8yr Adj. Exp. Return	-0.1182	0.0313	-0.0882
8yr Adj. Variance	0.1050	0.0168	0.1134

Table A.1: This table reports the number of properties remaining in the NCREIF dataset, 1978Q1-2013Q4 after applying a sequence of filters. Details of each step are reported in Appendix A.

Step	All Properties	Started as NPI	“S”/“F”/“P” /‘R” NPI Properties
1	30,895	20,547	
2	30,362	20,473	
3	30,362	20,473	16,045
4	24,403	19,278	15,929
5	22,971	18,375	15,048
6	20,825	17,808	14,536
	All Properties	“S” NPI Properties	“S”/“F”/“P” /‘R” NPI Properties
Total Return Filters Only applicable to return analysis	10,145	9,038	14,536
7	10,144	9,037	
8	10,121	9,020	
9	9,832	8,791	
10	8,818	8,115	
11	8,817	8,115	
	“S” NPI Properties w/wo Return Filter	“F” NPI Properties	“P” /‘R” NPI Properties
	9,038 / 8,115	5,318	180
12	3,960	3,148	46

Table B.1: Panels A and B report the first and second stage regressions for estimating the model with random effects in Equation (4) using randomly sampled REIT holding period returns substituted for the property holding period returns in the panel. The estimation procedure is described in the text. Under the null $\alpha_0 = \alpha_1 = \sigma_0^2 = 0$, while $v_{\alpha,y}$ and $v_{\sigma,y}$ are the idiosyncratic excess expected return and variance, or α and σ^2 , accumulated while holding the property for the entirety of year $y \in \{1978, \dots, 2013\}$. The terms $\frac{1}{36} \sum_y v_{\alpha,y}$ and $\frac{1}{36} \sum_y v_{\sigma,y}$ correspond to the average annualized α and σ^2 under the null. $P(v_{\alpha,y})$ (resp. $P(v_{\sigma,y})$) is the probability that all $v_{\alpha,y}$'s (resp. $v_{\sigma,y}$'s) are equal.

	(1)		(2)	
	$\frac{\hat{r}}{\sqrt{\tau}}$		Adj $\frac{\hat{r}}{\sqrt{\tau}}$	
β	0.8837***	(0.1058)	0.902***	(0.0975)
α_0	0.008	(0.0769)	0.0033	(0.0172)
α_1	0.005	(0.0805)	0.04	(0.1208)
$\frac{1}{36} \sum_y v_{\alpha,y}$	0.0004	(0.0256)	-0.001	(0.0182)
$P(v_{\alpha,y})$ const	<1E-6		<1E-6	
Observations	3206		3191	

	(1)	(2)	(3)	(4)	(5)	
	(6)	Adj \hat{z}^2	Adj \hat{z}^2	Adj \hat{z}^2	Adj \hat{z}^2	Adj \hat{z}^2
	Adj \hat{z}^2					
σ_α^2	-0.0002 (0.0009)	-0.0003 (0.0009)	-0.0006 (0.0009)	-0.0003 (0.0009)		
$\sigma_{\alpha\beta}$	-0.0163 (0.0101)	-0.0169 (0.0102)				
σ_β^2	-0.0574 (0.0594)	-0.0574 (0.0601)	-0.0884 (0.057)			
σ_0^2	-0.0447** (0.0149)	-0.05** (0.0155)	-0.0452** (0.0152)	-0.0408** (0.015)	-0.0366*** (0.0106)	
$\frac{1}{36} \sum_y v_{\sigma,y}$	0.0793*** (0.0184)	0.0813*** (0.0186)	0.0814*** (0.0186)	0.0762*** (0.0183)	0.073*** (0.0164)	0.065*** (0.0162)
$P(v_{\sigma,y})$ const	<1E-6	<1E-6	<1E-6	<1E-6	<1E-6	<1E-6
Observations	3206	3191	3191	3191	3191	3191

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.2: 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 (6) regression residual variance and intercept 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 transaction risk 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 in the table.

k	τ	N	Avg Vintage	(Var Diff)/($k - 1$)	t -stat	(α Diff)/($k - 1$)	t -stat
2	4	79	2002	0.0588	2.18	0.015	0.26
2	3	97	2001	0.0355	1.74	-0.015	-0.28
2	2	128	2003	0.0387	2.50	-0.029	-0.76
2	1	75	2004	0.0382	2.60	0.026	0.57
3	3	57	2002	0.0104	0.58	0.020	0.26
2	5	45	1997	0.0396	0.88	-0.144	-1.30
2	6	27	1995	-0.0706	-0.71	-0.106	-0.46
2	7	23	1994	0.0115	0.13	-0.351	-1.33
3	5	23	1996	0.0426	1.04	-0.124	-0.58
2	8	21	1989	0.0834	0.83	-0.434	-0.53

p -value for $\chi^2(\text{Var Diff} = 0)$	0.0072
p -value for $\chi^2(\alpha = 0)$	0.68


Statistic	GLS Estimate	se
Var Diff	0.0342	0.0081
α	-0.009	0.020
p -value for $\chi^2(\text{Var Diff} = \text{GLS Estimate})$		0.90

Table C.1: 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. Quarterly data for sold and unsold properties and quarters are used, including properties that were purchased more than one quarter prior to the first reporting period. The explanatory variables are property-specific attributes discussed in the text. Vintage year and manager fixed effects are included but their coefficients not reported. The marginal impact reports the change in probability of a sale when the corresponding variable changes from its 10th to its 90th percentile (for dummy variables the change is from zero to one).


variable	No year or mgr FEs			Includes year and mgr FEs		
	coeff	<i>t</i> -stat	Marginal impact	coeff	<i>t</i> -stat	Marginal impact
SqFt	-0.0855	-13.49	-0.034	-0.057	-6.01	-0.012
JV	0.3540	26.09	0.068	0.029	1.20	0.003
Age	0.3797	45.10	0.131	0.225	18.58	0.042
Percent Leased	0.0242	0.54	0.001	-0.191	-3.03	-0.005
Relative Loan spread	4.7618	4.84	0.005	1.004	0.70	0.000
Apartments	0.1380	7.56	0.026	0.271	9.45	0.029
Industrial	0.1575	9.66	0.029	0.037	1.41	0.004
Office	-0.0691	-3.91	-0.013	-0.120	-4.47	-0.012
East	-0.0195	-1.25	-0.004	-0.026	-1.19	-0.003
Midwest	-0.1495	-8.42	-0.027	0.053	2.13	0.005
South	-0.0178	-1.20	-0.003	0.043	2.04	0.004
Lagged return	-0.6791	-27.59	-0.063	-0.820	-23.14	-0.040
Idiosyncratic variance	-3.1538	-11.58	-0.014	-6.374	-13.63	-0.015
R_a^2	-0.7867	-26.80	-0.076	-1.914	-39.57	-0.100
β	0.0232	15.77	0.012	0.022	9.15	0.006
CEF				0.666	24.14	0.080
ODCE				-0.150	-5.74	-0.015
Unclassified fund				1.579	13.99	0.266
Non-ODCE OEF				-0.176	-4.96	-0.017
Constant	-1.1110	-11.91		-0.896	-3.24	
Observations			174,115			146,302

Table C.2: Results of a Hartigan and Hartigan (1985) “dip” test for deviations from unimodality of the holding period returns. At the top of the table analyzed returns are grouped by their nearest integer holding periods (half-integer holding periods are randomly rounded). In the bottom part of the table, the analysis is extended to subgroups of the same vintage (acquisition year). In the latter case, the test is done only with groups containing 30 or more observations . 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		213	0.0173	0.961	9		172	0.0286	0.539
2		379	0.0162	0.803	10		101	0.0243	0.964
3		487	0.0115	0.97	11		99	0.0281	0.887
4		338	0.0122	0.992	12		75	0.031	0.911
5		353	0.0248	0.241	13		70	0.0258	0.985
6		305	0.0182	0.774	14		48	0.0351	0.934
7		351	0.015	0.914	15		48	0.0446	0.738
8		309	0.019	0.73	16		28	0.0604	0.641
1	2003	35	0.0387	0.955	4	2002	40	0.0597	0.425
1	2005	49	0.0314	0.973	4	2005	30	0.0683	0.422
2	2003	50	0.0486	0.586	4	2006	34	0.0699	0.348
2	2004	45	0.0423	0.827	4	2007	33	0.0421	0.929
2	2005	60	0.0367	0.826	5	2007	52	0.0448	0.681
2	2006	31	0.0799	0.251	6	2006	31	0.0479	0.867
3	2002	72	0.0973	0.004	6	2007	42	0.0334	0.967
3	2003	101	0.0554	0.109	7	2006	90	0.0264	0.954
3	2004	35	0.0642	0.427	7	2007	34	0.049	0.823
3	2005	49	0.0313	0.958	8	2005	57	0.0351	0.893
3	2007	37	0.0534	0.679					



(a)



(b)

Fig. 1: The top figure depicts point estimates and their 95% confidence intervals for property expected idiosyncratic excess returns. The black circles and thick gray confidence intervals correspond to the in-sample predictions obtained from the econometric specification in (4) (including vintage effects). The green triangles and thin dashed confidence intervals correspond to predictions obtained from the simple regression model in (6). The bottom figure repeats the exercise for the idiosyncratic excess return variance.

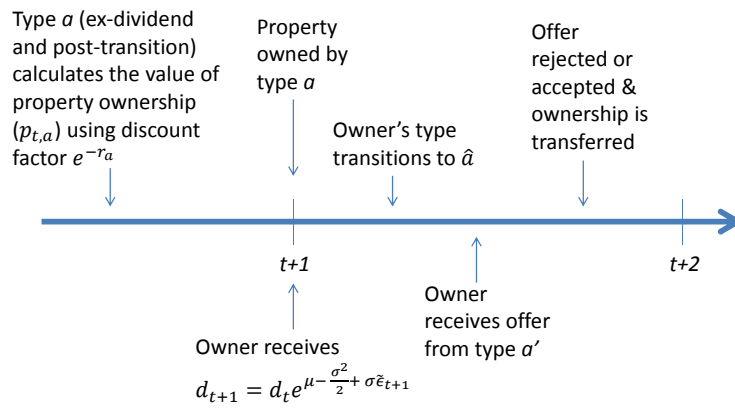




Fig. 2: Time line representing the sequence of events between dates $t + 1$ and $t + 2$. This is the relevant time line for an ex-dividend valuation of the property between dates t and $t + 1$ by an investor of type a (post-transition).



(a)



(b)

Fig. 3: The figures depict the predictions of the calibrated theoretical model (blue diamonds) against those of the estimated econometric model in (4) (red squares for the point estimates and thin dashed line for the 95% confidence interval). The top figure compares expected holding period returns while the bottom figure compares variance of holding period returns.




Fig. 4: The histogram plots the distribution of steady-state ownership for the model with the calibrated parameters in Table 3. The horizontal axis lists the owner type by the owner's annualized (private value) cap rate, $\frac{4}{Q_a}$. The average transaction cap rate in the model is 0.0704.




Fig. 5: Transaction Risk in Commercial Real Estate Properties. Each plot depicts the probability of transacting an asset within one quarter as a function of the seller’s reservation income-to-price ratio (i.e., “cap rate”). The step function corresponds to the probability distribution transacting a perfectly liquid asset. The calibrated model yields significant transaction risk for commercial properties, as shown by the distribution function labeled “Property Market”.




Fig. 6: The solid blue and dashed red lines trace the mean and median, respectively, of the time to sell given a seller's reservation cap rate (gross of costs). The black dash-dotted line graphs the expected transaction cap rate (corresponding to the right vertical axis). The large solid circle signifies the reservation cap rate at which a seller can expect to sell at prevailing market rates.




Fig. 7: An owner is said to have a need for immediacy of y years if (s)he must sell the property within y years. The graph depicts the discount relative to prevailing transaction prices that the seller expects to experience if using the disposition strategy described in the text.

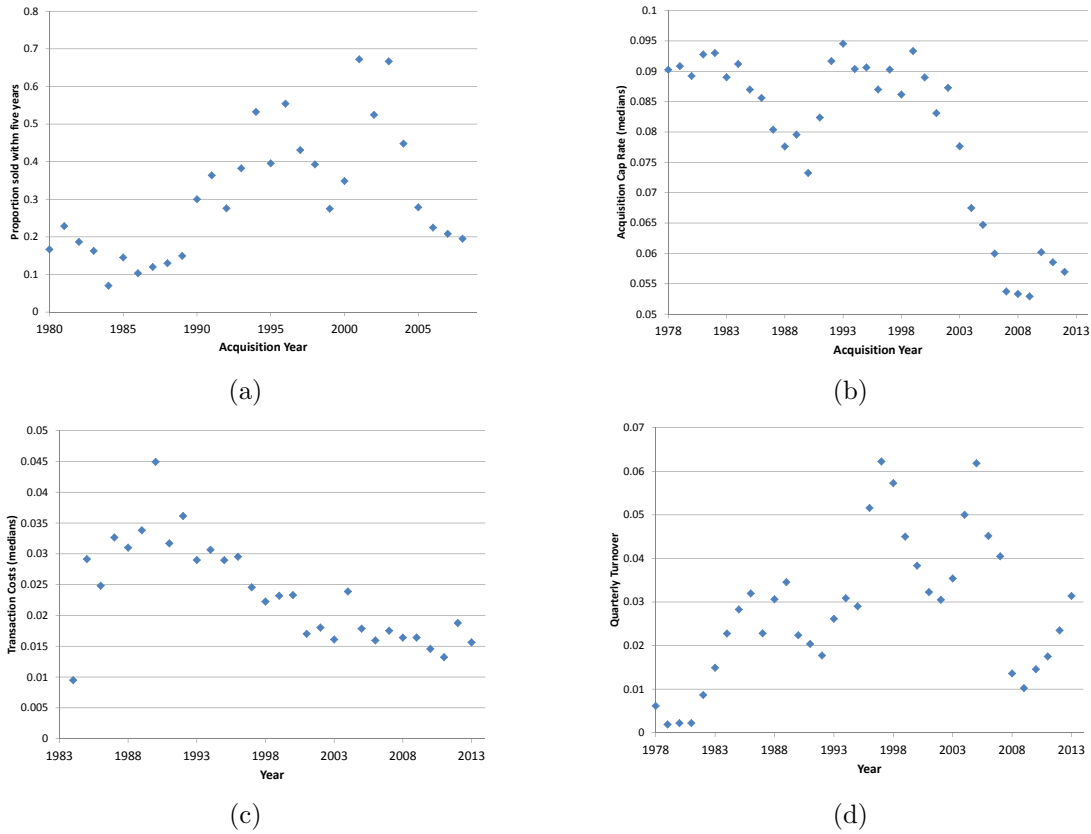




Fig. A.1: The four figures plot property statistics from the “cleaned” dataset described in Appendix A. Figure (a) depicts the proportion of properties sold within five years of purchase by the year they were acquired. Only properties that began reporting into the dataset no more than one quarter after acquisition are included. Figure (b) depicts the acquisition cap rate of these properties by the year of acquisition. The cap rate is calculated as sum of first four quarters of reported NOI divided by the purchase price. Medians are used to avoid bias from outliers. Figure (c) reports transaction costs for sold properties by year of sale. This is calculated as gross transaction price less net transaction price, divided by the gross transaction price. Here too medians are used. Figure (d) plots the quarterly turnover per year, calculated as the number of property-quarters in which a sale is reported in a given year divided by the number of property-quarters in that year. Property number-weighted time-series averages of each quantity graphed are used in the calibration of the model from Section 3.



(a)



(b)

Fig. B.1: The top figure depicts point estimates and their 95% confidence intervals for expected idiosyncratic excess returns when property holding period returns are replaced with randomly sampled REIT returns. The black circles and thick gray confidence intervals correspond to the in-sample predictions obtained from the econometric specification in (4) (including vintage effects). The green triangles and thin dashed confidence intervals correspond to predictions obtained from the simple regression model in (6). The bottom figure repeats the exercise for the idiosyncratic excess return variance.




Fig. C.1: The figure depicts point estimates and their 95% confidence intervals for the variance of commercial property idiosyncratic excess returns. The holding period returns are appraisal-based and are calculated up to one year prior to the sale of the property. The black circles and thick gray confidence intervals correspond to the in-sample predictions obtained from the econometric specification in (4) (including vintage effects). The green triangles and thin dashed confidence intervals correspond to predictions obtained from the simple regression model in (6).