Asset-level risk and return in real estate investments

Jacob S. Sagi*

Abstract

In stark contrast with liquid asset returns, I find that commercial real estate idiosyncratic return means and variances do not scale with the holding period, even after accounting for all cash flow relevant events. This puzzling phenomenon survives controlling for vintage effects, systematic risk heterogeneity, and a host of other explanations. To explain the findings, I derive an equilibrium search-based asset-pricing model which, when calibrated, provides an excellent fit to transactions data. A structural model of transaction risk seems crucial to understanding real estate price dynamics. These insights extend to other highly illiquid asset classes, such as private equity and residential real estate.

*I am grateful to the National Council of Real Estate Investment Fiduciaries (NCREIF) for providing me with the data, and especially to Jeff Fisher for helping me understand it. I also benefited from feedback from Bob Connolly, Lynn Fisher, Andra Ghent, David Geltner, Adam Gurren, Dave Hartzell, Allen Head, Robert Kimmel, Crocker Liu, Greg MacKinnon, Ludo Phalippou, Tim Riddiough, Kanis Saengchote, Morten Sørensen, Chester Spatt, Richard Stanton, Nancy Wallace, seminar participants at Baruch College, Duke University, NC State University, the University of Calgary, 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, the 2017 Utah Winter Finance Conference, the March 2017 IPC Conference, and the 2017 Real Estate Price Dynamic Workshop at MIT. Finally, I wish to thank the editor, Stijn van Nieuwerburgh, and two anonymous referees for their many thoughtful comments. Please send correspondence by email to sagi@unc.edu.

1 Introduction

Real estate is an important investment class. As of 2019, MSCI estimated that roughly $9 trillion (USD) in global real estate assets were held for investment purposes under professional institutional management, while Savills estimated in 2016 that the potential stock of investable institutional-quality global commercial real estate (CRE) was $19 trillion.[1] Relatedly, CRE mortgages totaled over $4 trillion in the U.S. alone as of 2019 and comprised 35% of all loans for (mostly regional) banks with assets under $10B.[2] Despite the significance of CRE as an investment class, price dynamics for individual assets are not as well understood as those belonging to more liquid categories such as equities, fixed income, commodities, currencies, and derivatives. This gap is important to address because asset-specific price-dynamics can significantly impact the pricing of non-recourse mortgages, the concentrated portfolios held by many real estate investors, and the option-like features in investment management contracts prevalent among private equity funds.[3]

The few papers that attempt to quantify attributes of property level risk and return assume that an individual property's log-value evolves as a random walk with drift (RWD).[4]. This is consistent with the prevailing paradigm in asset pricing (Campbell, Lo, and MacKinlay, 1997) and is the modeling assumptions most widely adopted in the applied theory literature, going back to Williams (1993). This paper makes two contributions to current understanding of CRE asset price dynamics, and the insights should generalize to other highly illiquid asset investments. First, I demonstrate with an exhaustive series of empirical tests that the RWD assumption is far from an appropriate description for CRE asset prices. Second, I derive an equilibrium search-based model that is able to fit observed dynamics of prices and turnover statistics, as well as transaction dispersion relative to perceived market prices.

The literature is full of hints that real estate prices and those of other illiquid assets deviate from a RWD, but such findings have been typically dismissed or viewed as data artifacts (e.g., measurement

[1]See MSCI's "Real Estate Market Size 2018", and Savills' "Around the World in Dollars and Cents (2016)".

[2]See Table L.217 of the Financial Accounts Guide published by the Board of Governors of the Federal Reserve System and the Yardi Matrix Bulletin (July 2018), titled "Regional/Local Banks Eat More of the Commercial Mortgage Pie: When is Enough?".

[3]Real estate corresponds to roughly 15% of private equity assets (Preqin, 2019). Asset-level risk may also be material for roughly a third of CRE investors, likely undiversified, whom Geltner et. al (2013) identify as "...private investors, relatively small, typically locally oriented often family-based enterprises."

[4]See, for instance, Downing, Stanton, and Wallace (2008); Plazzi, Torous, and Valkanov (2008); Peng (2016).

error and/or missing variables). The calling card of a RWD is the scaling of the mean and variance of log-price changes with the time between changes (i.e., the holding horizon). Case and Shiller (1987), Abraham and Schauman (1991), Goetzmann (1993), Goetzmann and Spiegel (1995), and Calhoun (1996), find that return scaling is violated for house price increases based on repeat sales, and attribute the anomalies to missing variables (e.g., investment in the form of renovation activity or price mismeasurement). Using detailed property-level data from the National Council of Real Estate Investment Fiduciaries (NCREIF), which includes capital expenditures, I robustly demonstrate that CRE prices are inconsistent with RWD dynamics even after accounting for all cash flow events. Specifically, both risk-adjusted mean (i.e., "alpha") and variance of property log-returns fail to scale with the return horizon, and both exhibit large positive *atemporal* components: Return means and variances remain significantly positive even when the holding period is extrapolated to zero. The atemporal variance of roughly 3% is two to three times the idiosyncratic annual diffusion variance. The atemporal alpha is absurdly high, extrapolating to double-digit percentage points as the holding period vanishes.

With perfectly liquid assets, the RWD anomalies would present an arbitrage opportunity consisting of sequentially purchasing and then instantly selling the asset an arbitrary number of times over a finite time interval.[5] The illiquidity of real estate assets rules out the executability or profitability of such an exercise. Indeed, by providing strong evidence that their source is transactional, I am able to link the root cause of the return anomalies to illiquidity rather than missing variables or a host of other explanations. To my knowledge, aside from this paper and contemporaneous complementary work in residential real estate by Giacoletti (2017), no other paper has carefully demonstrated that the deviation from RWD is an inherent feature of real estate price data nor made definitive progress in pinning down the reasons. The fact that empirical deviations from RWD have been observed (though not as thoroughly investigated) with other illiquid assets serves to both robustify my findings with CRE assets and, more importantly, strongly suggest that transactional frictions are essential to the modeling of price dynamics across *all* highly illiquid asset types.[6]

[5]Because each instant two-way transaction is an independent draw with strictly positive average return, the accumulated return of this strategy would result in arbitrarily large profits at negligible risk.

[6]For a review of the financial econometrics used to analyze non-real estate private equity markets see, for instance, Ang and Sorensen (2012). For evidence of atemporal mean and variance in these markets see Axelson, Sorensen, and Stromberg (2015) and Lopez-de Silanes, Phalippou, and Gottschalg (2015).

The second main contribution of this paper is to provide the theoretical link between illiquidity and the observed deviations from RWD in real estate price data. I do this by deriving a search-based model featuring two key mechanisms. The first is dispersion in the relative valuations of randomly matched counterparties, which leads to idiosyncratic randomness in the transaction prices of otherwise identical properties. This accounts for atemporal variance and shows up in an asset's observed price series as variance that depends on the number of transactions observed and not only on the holding horizon. Because random matching and bargaining are common devices, many search models exhibit atemporal variance. The second key mechanism is the presence of "intermediate" valuation investors, coupled with the highly plausible assumption that private valuations change slowly.[7] If a recent property buyer is slow to change his or her valuation, the holding period will be short only if a much better offer happens to quickly come along. This is most likely to happen to intermediate valuation investors because "better offers" are not available to high valuation owners, while low valuation investors are rarely recent buyers. The observed positive 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.

In the limit of a liquid market the model exhibits RWD returns, highlighting the connection between illiquidity and the atemporal return anomalies. To my knowledge, among existing search models in real estate (see Han and Strange, 2015, for a review), mine is the first to incorporate the second key mechanism and is the only one able to explain the atemporal alpha (through selection bias in repeat sales holding periods).[8] From that perspective, the model can be viewed as a novel contribution to the prodigious real estate search literature and opens the door to new theoretical and econometric analyses of transaction-level data.

The model allows for cyclical dynamics in the property market and investors' cost of capital. Beyond the qualitative economics, the model's steady-state equilibrium can be calibrated to fit a large set of transaction-based moments from NCREIF property data, conditional on boom or bust market states. Among these are holding period return moments for properties bought/sold at different cycle points, quarterly turnover, the fraction of properties sold after five years, the average transacted income-to-

[7]The two key model mechanisms are distinct. It is theoretically possible for private valuations to become more dispersed even as intermediate valuation investors become rarer (or vice versa).

[8]Similar features are employed in contemporaneous papers by Lovo and Spaenjers (2018) and Hugonnier, Lester, and Weill (2018) studying, respectively, art auction markets and over-the-counter endogenous intermediation chains.

price ratio (or "cap rate"), and the distribution of transactions relative to perceived market prices (i.e., appraisals). The model fits the data very well and fully accounts for atemporal alpha and variance. For the best fit parameter set, as the market moves from an expansion to a contraction, investors' discount rates increase and strongly revert towards consensus. The former leads to aggregate asset devaluation and the latter to a dramatic decline in liquidity.

In contrast with a liquid market, sellers in the model must trade off price and certainty of execution. I quantify this notion of *transaction risk* in various ways: The probability of transacting at or above the average transaction price, expected or median time on market given a reservation price, and expected discount conditional on a binding liquidation horizon. The calibrated model implies a reduction in high-valuation prospective buyers during periods of market contractions, and this makes the tension between price and speed of execution more severe as reflected across all measures. The magnitude of model-imputed transaction risk and its pronounced cyclical variation has significant implications for the pricing of property derivative assets such as mortgage loans and mortgage backed securities. These could be important to incorporate into the current literature, where RWD assumptions are prevalent.

The paper is organized into an empirical (§2), model (§3), calibration (§4), applications (§5), discussion (§6), and concluding (§7) sections. Appropriate literature is discussed in each section, and online appendices provide supplementary details and analyses.

2 Evidence on CRE Holding Period Risk and Returns

This section and the associated Appendix A provide evidence for rejecting a random walk with drift (RWD) in commercial real estate (CRE) returns and trace the source to transactional frictions.

2.1 Property-level Data

Property-level CRE data comes from the National Council of Real Estate Investment Fiduciaries (NCREIF) and consists of quarterly financial and accounting information reported by member funds between 1978Q1 and 2017Q2. The dataset contains property acquisition dates and transaction prices. If a sale (or partial sale) took place, the sale date and net/gross prices (excluding/including selling expenses) are typically reported.[9] For each property and quarter (usually) starting from the acquisition quarter (or 1978Q1 if acquisition is earlier) until 2017Q2 (or until the property is sold or otherwise exits the dataset), market value appraisals and net operating income (NOI) are likewise reported. Also available are details about quarterly capital expenditures (CapEx), property location, age, property type, leverage, ownership structure, owning fund, and type of fund. A flag reports whether a 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 most would be viewed as "core assets" by professionals.

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_{iT}. Total holding-period returns depend on the reinvestment strategy of interim cash flow. 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 CRE generates considerable income, longer holding period return characteristics would be dominated by those of the reinvestment alternative rather than the property. For this reason, I mainly focus on price appreciation returns. As will be shown, including income reinvested at LIBOR, does not meaningfully alter any results or conclusions.

I calculate the excess log of price-appreciation return over property i's holding period as,

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

where $r_{I,s}$ is the quarter s return on investing the proceeds from partial sales. Equation (1) corresponds to an excess return because the denominator is capitalized to date T using a risk-free return. While the

[9] The NCREIF data contributor manual states that: "Selling expenses are directly attributable to the sale which are the seller's responsibility including, but not limited to, items such as commissions, disposition fees, legal fees, title insurance, escrow fees, etc."

numerator depends on a discretionary investment strategy for partial sales, in practice this corresponds to fewer than 5% of properties and the reinvestment strategy chosen has negligible impact on the analysis. The reinvestment return I use for $r_{I,s}$ is the corresponding quarter's NPI total index return for that property's major property type. *Henceforth, to avoid tedious expressions, unless otherwise indicated I employ the term "return" to refer to the quantity calculated in Eq.* (1).

The data contains properties with missing or inconsistent entries that I attempt to filter using a set of criteria detailed in Appendix A.1.[10] To mitigate any influence of tax motives on holding periods, I restrict attention to properties that qualified for the NPI when first acquired. 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 filtering results in 4,535 single-property repeat sale transactions, 17 portfolio-property dispositions, 40 property foreclosures, and 3,628 properties that had not exited as of 2017Q2. Of these, I exclude the 1% of return observations with the greatest discrepancy between appraisal returns and those calculated from Eq. (1).

Table 1 reports quarterly panel statistics for the filtered dataset described above. Differences between sold and unsold properties are mostly attributable to time period (the median year for sold/unsold properties is 2005/2013). NCREIF NPI properties tend to be medium to large in size, have high occupancy, and relatively low leverage. CapEx spending is lumpy but averages roughly 2% per year for sold properties. About 10% of these experience CapEx spending of 10% or more of initial appraisal value in the first two years and the standard deviation of total CapEx in the first two years is about 9%. Thus, consistent with Goetzmann and Spiegel (1995), CapEx can drive sizable changes to a property's gross market value appreciation in the first two years after acquisition.

Table 2 reports on additional property characteristics. NCREIF property ownership is organized through fund structures, some of which are sponsored by the largest private equity investment firms in the U.S. Roughly a quarter of properties experience some financing through a joint venture (JV) partnership. About 13% are owned through a private closed-end fund (CEF) structure whereby the managers raise capital commitments from limited partners for a period of time, deploy this capital over a subsequent limited period, and then liquidate assets to meet a contracted fund termination date.

[10] Appendix A.4 reports on the robustness of the main empirical analysis to alternative filters.

Table 1: Summary statistics from the NCREIF panel, 1978Q1-2017Q2. The data is cleaned according to the procedure outlined in Appendix A.1. Acquisition cap rates are the property's first four quarters of NOI divided by the reported purchase price. NCREIF reports quarterly returns (Qtrly Ret) based on quarterly appraised market value (MV), net operating income (NOI), and capital expenditures (CapEx). Loan to value (LTV) ratios are calculated using MV and the balance remaining on loans secured to the property. The CapEx ratio is that quarter's CapEx divided by MV. Selling transaction costs (Sale TCost) are calculated using reported differences between gross and net sales prices, divided by gross price. The variable "HP PApp Ret" is as calculated in Eq. (1) and the corresponding return calculated using appraisal values is "HP NPI PApp Ret". Turnover statistics are calculated using data from 1983q2, a year after the NCREIF began recording property data. Quarterly turnover is calculated as the number of sold properties each quarter divided by the number of properties in the dataset that quarter. An estimate of the distribution of holding periods is calculated from the ratio of properties sold h years after purchase relative to all properties purchased at or before 2017q2 $-h$ (the mean and variance are estimated from this distribution to avoid bias). To avoid the impact of extreme outliers, means and standard deviations are reported after dropping values below the 0.5 and above the 99.5 percentiles.

Variable	N	mean	sd	p1	p5	p10	p25	p50	p75	p90	p95	p99
	Unsold properties as of 2017Q2											
Acquisition cap rate	2,935	0.059	0.029	-0.016	0.019	0.033	0.046	0.057	0.070	0.091	0.107	0.225
Qtrly Ret	80,269	0.020	0.045	-0.154	-0.047	-0.010	0.011	0.017	0.031	0.062	0.089	0.173
LTV	75,195	0.26	0.29	0	0	0	0	0.099	0.495	0.640	0.755	1.011
MV ($ Millions)	80,269	56	70	1	3	6	15	35	70	137.000	219	548
SqFt (1000's)	80,191	268	253	-	7	41	95	183	335	554.663	846	1,435
Age	73,716	22	17	2	4	6	10	17	29	43.000	56.0	107
Occupancy	78,542	0.89	0.12	0.398	0.683	0.797	0.913	0.970	1.000	1.000	1.000	1.000
CapEx Ratio	80,269	0.0036	0.0084	-0.0030	0	0	0	0.001	0.003	0.010	0.018	0.056
	Sold properties											
Acquisition cap rate	4,672	0.073	0.033	-0.042	0.016	0.036	0.054	0.074	0.092	0.106	0.120	0.189
Qtrly Ret	120,936	0.017	0.057	-0.255	-0.090	-0.035	0.009	0.019	0.028	0.066	0.109	0.260
LTV	74,166	0.33	0.32	0	0	0	0	0.384	0.596	0.753	0.866	1.127
MV ($ Millions)	120,936	26	37	-	1	2	6	14	31	59.114	87	179
SqFt (1000's)	120,936	230	217	-	0	37	84	157	278	482.829	643	1,253
Age	94,304	20	13	2	4	5	10	17	26	34.000	42.0	80
Occupancy	98,258	0.87	0.13	0.330	0.620	0.732	0.870	0.950	1.000	1.000	1.000	1.000
CapEx Ratio	117,246	0.0045	0.0098	-0.0033	0	0	0	0.001	0.004	0.013	0.023	0.068
Sale TCost	5,002	0.025	0.025	0.000	0.000	0.003	0.010	0.018	0.032	0.052	0.070	0.167
HP PApp Ret	4,472	-0.22	0.47	-1.825	-1.215	-0.862	-0.462	-0.128	0.104	0.297	0.417	0.760
HP NPI PApp Ret	4,472	-0.22	0.54	-1.980	-1.332	-0.947	-0.505	-0.136	0.130	0.371	0.580	0.994
Qtrly Turn	179,708	0.030	0.019	0.002	0.008	0.011	0.017	0.027	0.038	0.054	0.068	0.100
Holding Period	8,788	8.2	5.2	0.6	1.4	2.1	3.8	7.3	11.3	15.5	17.9	22.7

Additional fund structures include separate account funds dedicated to a single client, and private open-end funds that permit capital contribution at any time as well as redemption of capital by existing investors. The Pareto Principle (the "80-20" Rule) applies in that roughly 80% of the properties are managed by 20% of the (LgMgr) funds. Heterogeneity in ownership structure proves important for understanding the property-level return characteristics of the data and for interpreting the model.

Table 2: Additional summary statistics reporting on ownership, type and location characteristics of properties in the NCREIF dataset (cleaned according to the procedure outlined in Appendix A.1). JV corresponds to the number of properties that have been part of a joint venture at any point during their tenure with the owning NCREIF member. CEF denotes properties held by closed-end private equity funds. LgMgr corresponds to properties held by a large fund (the set of largest funds owning 80% of properties). Property types (Apartments, Industrial, Office, or Retail) are denoted as A, I, O and R. The six cities with the greatest share of properties are denoted by their international airport codes.

	Unsold properties as of 2017Q2								
Not JV	JV			Not CEF	CEF			SmMgr	LgMgr
2,536	1,149			3,305	380			662	3,023
A	I	O	R	ATL	ORD	DFW	LAX	SFO	IAD
836	1,325	878	646	207	237	224	454	266	250
	Sold properties								
Not JV	JV			Not CEF	CEF			SmMgr	LgMgr
3,994	1,109			4,325	778			1,044	4,059
A	I	O	R	ATL	ORD	DFW	LAX	SFO	IAD
1,172	1,895	1,289	747	372	335	357	516	299	359

A plurality of the properties in the dataset correspond to industrial CRE (I). These include warehouse, manufacturing, research, and showroom facilities. Office (O) and multifamily (A) properties make up the majority of the remainder. Retail properties (R) form the smallest category. Table 2 documents all MSAs with more than 500 properties in the dataset, referring to each by its major airport code. These account for 44% of properties. Prominently missing from this list is New York City, highlighting the fact that NCREIF membership is only representative of CRE ownership and management through investment fiduciaries (as opposed to direct ownership). That said, the properties held by NCREIF members in 2017Q2 amounted to over $500 billion, and estimated to be around a sixth of all U.S. CRE professionally managed by investment institutions and funds.

Table 3: Simple holding period return analysis. The panels report means and variances for holding period returns as calculated in Eq. (1) (including and excluding income) for the NCREIF dataset (cleaned according to the procedure outlined in Appendix A.1). Extreme return observations (below the first and above the ninety-ninth percentiles) are dropped from each nearest-integer holding period bin, (half integer holding periods are randomly rounded either up or down). The last column is a GLS estimate of the intercept from regressing each row on the holding period. Under standard asset pricing assumptions the intercept should be zero.

	Holding Period (years)						
	1	2	3	4	5	6	GLS Intercept
	Annualized mean return						
Excluding income	0.072	0.078	0.040	-0.040	-0.203	-0.189	0.175**
Including income	0.124	0.170	0.191	0.182	0.086	0.165	0.155***
	Annualized return variance						
Excluding income	0.034	0.067	0.089	0.124	0.131	0.121	0.016*
Including income	0.031	0.061	0.082	0.096	0.101	0.088	0.024**
	Number of properties						
	219	429	579	440	391	361	

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

2.2 Reduced-form evidence

Table 3 reports mean and variance for holding period returns calculated both with and without capitalized income in the numerator of Eq. (1).[11] Here, and in subsequent regressions, I drop extreme return observations (below the first and above the ninety-ninth percentiles) from each nearest-integer holding period bin, (half integer holding periods are randomly rounded either up or down). The column "GLS Intercept" reports an estimate of the intercept from the regression $S_\tau = a + b\tau + \varepsilon_\tau$ where τ is the holding period, S_τ is the holding period statistic (mean or variance of returns), and ε_τ is a residual with heteroskedastic variance determined by the estimation error in S_τ. Under standard asset pricing assumptions, asset returns are the accumulation of independent shocks through time, and thus return mean and variance should both scale with the holding period (i.e., the intercept should be zero).

The top two rows document the first key anomaly: Average returns do not scale with horizon and when extrapolated to an arbitrarily short holding horizon exhibit a large and highly significant intercept. This is true regardless of the presence of income in returns, consistent with the view that the anomaly arises from price, rather than income, dynamics. In a liquid market where one can trade quickly, earning

[11] The additional income term is $\sum_{s=t}^{T-1} \text{inc}_{is} e^{\sum_{s'=s+1}^{T} r_{\text{LIBOR},s'}}$ where inc_{is} is the reported NOI for property i in period s, and the reinvestment rate, $r_{\text{LIBOR},s'}$, corresponds to 3-month LIBOR.

a finite return over an arbitrarily short period would constitute an arbitrage opportunity. The next two rows document the second key anomaly: Return variance also does not scale with horizon. Here too, the extrapolated intercept is significant. Naively, both anomalies appear to violate standard market efficiency and the RWD assumptions. Note that the variance anomaly is not related to estimation smoothing or missing investment because the property returns are actual and incorporate CapEx.

In a noteworthy though unpublished study, Ciochetti and Fisher (2002) examine holding-period internal rates of return (IRRs) in NCREIF data during 1978-2001. Their Table 10 reports average unadjusted IRRs by holding period. Over holding periods up to sixteen years, where the sample size is at least 30, short-term IRRs appear to dominate long-term IRRs (there is no test for significance and standard errors are not reported). While consistent with the average returns in Table 3, IRRs are biased relative to expected returns and the pattern they find could be mechanically induced.[12] Although they do not comment on their anomalous results, test whether they might be spurious, nor investigate their economic origin, their analysis provides early hints at an interesting and important anomaly.

2.3 A More Sophisticated Analysis

The analysis above does not control for the fact that properties with different holding periods are likely to have been held during different years and subjected to different market conditions (see Table A-II in Appendix A.1). To guard against spurious inference, I develop econometric tests that control for a host of influences. Consider the following assumptions concerning the returns in Equation (1).

Null hypothesis. *The logarithm of capital gains in property i's price over a short time interval,* Δ*, and net of a risk-free alternative takes the form*

$$r_{i,\Delta} = \alpha_{i,t}\Delta + \beta_{i,t} r_{m,\Delta} + \sigma_{i,t}\sqrt{\Delta}\varepsilon_{i,\Delta},$$

where (i) ε_{Δ} *is a standardized shock to property growth and is independent across distinct time intervals and properties; (ii)* $r_{m,\Delta}$ *represents the influence of systematic shifts in property market values over*

[12]In the simplest case in which cash flow per dollar of initial investment, $\tilde{x}$, is non-zero only at sale, average IRR is just $E[\tilde{x}^{1/n}]$ where n is the holding period. By Jensen's inequality, $E[\tilde{x}^{1/n}]$ (for $n > 1$) is biased *down* relative to $E[\tilde{x}]^{1/n}$, and the bias increases with both the holding period n and the variance of $\tilde{x}$.

Δ; and (iii) $\alpha_{i,t}, \beta_{i,t}$ and $\sigma_{i,t}$ are constant over the interval Δ, uniformly bounded across all distinct time intervals and properties, and independent of the history of $r_{m,\Delta}$ and ε_Δ. Moreover, transaction decisions are unrelated to the history of $r_{i,\Delta}$.

This, of course, is a joint null, of which parts (*i*)-(*iii*) essentially form the RWD hypothesis. By assumption, for any holding period, τ, the logarithm of excess capital gains, $r_{i,\tau}$, is a sum of incremental returns as defined above. Because $\alpha_{i,t}, \beta_{i,t}$ and $\sigma_{i,t}$ are uniformly bounded and independent of the ε_Δ's and $r_{m,\Delta}$'s, this sum can be written as

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

where $r_{m,\tau}$ is the change in the systematic risk factor over τ, and each of the constants has the same attributes as stated in the Null Hypothesis. $\varepsilon_{i,\tau}$ is an idiosyncratic mean-zero shock with variance one and uncorrelated with $r_{m,\tau}$. It is instructive to make a connection with the standard asset pricing literature that commonly assumes a simple log-normal structure.[13] Suppose that one identifies $r_{m,\tau}$ with the log of excess capital gains on some benchmark portfolio. Then in a standard setting, as derived in Appendix A.2, α_i can be decomposed into the following constituents.

$$\alpha_i = -\frac{\sigma_i^2}{2} + \beta_i \big(1 - \beta_i\big) \frac{\sigma_m^2}{2} + \big(\beta_i \mu_m - \mu_i\big) + \ell_i,$$

where μ_m and μ_i are the the income rates generated by, respectively, the benchmark portfolio and the property, σ_m is the price volatility of the benchmark portfolio, and ℓ_i is a liquidity premium denoting a return component that can only be generated in an imperfect market where deviations from the law of one price are hard to exploit. If (2) referred to a standard CAPM equation for simple total returns, α_i would be zero. The $\frac{\sigma_i^2}{2}$ is a Jensen's inequality adjustment to the idiosyncratic ($\varepsilon_{i,\tau}$) shock and is required because (2) employs log-returns of price appreciation rather than simple total returns. The component proportional to β_i is likewise a Jensen's adjustment term to the benchmark log-return contribution, and compensates for shocks to the benchmark portfolio. The difference between income

[13]See Campbell, Lo, and MacKinlay (1997) for a more complete review. Empirically estimating Equation 2, which may be viewed as a cross-sectional rather than a time-series decomposition of return attributes, has parallels in the labor economics literature (see Benzoni and Chyruk, 2015).

rates would be absent if (2) were to be stated using the log of total returns and appears because $r_{i,\tau}$ and $r_{m,\tau}$ are capital gain returns defined net of income.

Consider averaging across α_i's. If $r_{m,\tau}$ is a property market benchmark (a weighted average return across properties), the average beta should be close to one, and the average of $\beta_i\mu_m - \mu_i$ should be zero. If the β_i's are not too skewed, the average of $\beta_i(1-\beta_i)$ would be negative. Further assuming the average of ℓ_i is zero, implies a negative average α_i (approximately equal to the average of the $-\frac{\sigma_i^2}{2}$'s).[14]

For a property purchased at date t and sold at $t+\tau$, it is possible to observe $r_{i,\tau}, \tau$ and $r_{m,\tau}$. Under the null, $\varepsilon_{i,\tau}$ is unrelated across properties and independent of τ. Thus each property's holding period return from a repeat transaction is an independent random sampling from a standardized random variable (i.e., $\varepsilon_{i,\tau}$) and the distribution of property-specific coefficients (α_i, β_i and σ_i). 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 $\sigma\tilde{\varepsilon}$.[15] Importantly, under the Null, after controlling for the benchmark contribution, both mean and variance of returns should vanish with the holding period. Appendix A.2 demonstrates that is true even for returns expressed gross of the systematic benchmark return and/or asset income.

To capture deviations from scaling with the holding period in both mean and variance (as observed in the simple analysis of Section 2.2), I add a random atemporal term, $\tilde{\alpha}_0$ to $\tilde{r}$. To minimize collinearity in the regression specification to follow, I normalize the resulting equation by $\sqrt{\tau}$ and then add an intercept term, α_1, 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 emphasizing that $\tilde{\alpha}_0$ and α_1 are zero under the null. Further separating the expected values of the random coefficients above by setting $\tilde{\alpha}_0 \equiv \alpha_0 + \tilde{\epsilon}_0$, $\tilde{\alpha} \equiv \alpha + \tilde{\epsilon}$, and $\tilde{\beta} \equiv \beta + \tilde{\epsilon}_\beta$, yields

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

[14] A zero average ℓ_i is consistent with assuming that the aggregated property market liquidity premium is incorporated into the property market benchmark.

[15] Here, the residual $\tilde{\varepsilon}$ is redefined to absorb variations in σ_i.

Under the null, $\tilde{\varepsilon}$ and the $\tilde{\epsilon}$'s are uncorrelated with $r_{m,\tau}$ and τ, rendering (3) a regression equation with heteroskedastic residuals. Appendix A.3 describes a four-stage OLS procedure for estimating the coefficients and variance components in (3). Care is taken to control for the significant contributions of year fixed effects to α and σ because these may lead to spurious rejections of the null (e.g., if large τ properties are predominantly held during years of low volatility as compared with small τ properties).

I also exclude properties 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 because 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 4: The table reports estimates of $\beta, \alpha, \alpha_0, \alpha_1, \sigma^2$ and σ_0^2 from the regression $\frac{\tilde{r}}{\sqrt{\tau}} = \alpha_0 \frac{1}{\sqrt{\tau}} + \alpha_1 + \alpha\sqrt{\tau} + \beta\frac{r_{m,\tau}}{\sqrt{\tau}} + \left(\tilde{\epsilon}_0 \frac{1}{\sqrt{\tau}} + \sigma\tilde{\varepsilon}\right)$, where $\tilde{r}$ is a property holding period return, τ is the holding period, $r_{m,\tau}$ is the NPI index return for the corresponding property type over the holding period, $\tilde{\epsilon}_0$ is a residual with variance σ_0^2, and $\tilde{\varepsilon}$ is a residual with variance one. The procedure for the estimation , detailed in Appendix A.3, controls for year fixed effects in α and σ^2. The second and fifth columns report the baseline estimation using price appreciation returns for the property as calculated in Eq. (1), and the NPI property-type specific price appreciation return index for $r_{m,\tau}$ (capital expenditures are incorporated into the calculation of both return series). Under the conventional asset pricing null, in which average returns scale with the holding period, α_0, α_1 and σ_0^2 should be zero. The third and last columns report estimates when income is incorporated into the return series. The penultimate row reports the probability that the stated parameters are unchanged relative to their baseline estimates.

	Baseline	incl. income		Baseline	incl. income
β	0.997***	0.9453***	σ^2	0.0085***	0.0032**
	(0.0610)	(0.0603)		(0.0014)	(0.0011)
α	0.005	0.0045	σ_0^2	0.0326***	0.0352***
	(0.008191)	(0.0092)		(0.00445)	(0.0038)
α_0	0.201***	0.156***			
	(0.0383)	(0.0427)			
α_1	-0.126***	-0.0967*			
	(0.008191)	(0.043)			
Prb no change in β, α & α_1		0.1399	Prb no change in σ^2		<1E-6
Observations	4233	4224		4233	4224

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

For each property I use the corresponding type-specific NPI index 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. The second and fifth columns of Table 4 report coefficient estimates in the baseline case, which applies to price appreciation as calculated in Eq. (1). The average baseline β and α are very close to one and zero, respectively. The atemporal α_0, however, is economically very large, positive, and highly significant. A time-dependent decline in average idiosyncratic returns is captured by a negative, large, and significant α_1 coefficient. Under the null, α_0 and α_1 should be zero and this is solidly rejected by the data. The various contributions to average returns for a one-year holding period sum to roughly 8%, which is comparable to the simple one-year mean returns in Table 3. The GLS intercept estimate in the top row of Table 3 is likewise comparable with the estimate of α_0. The variance component estimates in the penultimate column of Table 4 indicate an annualized idiosyncratic diffusion variance of 0.0085, corresponding to an annualized idiosyncratic volatility of roughly 9% and representing about 85% of the total diffusion variance in price appreciation.[16] The atemporal variance component, $\sigma_0^2 = \text{VAR}[\tilde{\varepsilon}_0]$, is more than three times as large as the diffusion component and statistically comparable with the simple GLS estimate in Table 3 (third row). This too is a rejection of the null, which asserts that variance scales with holding period, and confirms the conclusions of the simple analysis in Table 3. Random variations in property α's and β's (corresponding to $\tilde{\epsilon}$ and $\tilde{\epsilon}_\beta$, and reported in Table A-III of Appendix A.3) are not significantly different from zero and are omitted from Table 4.

The third and last columns in Table 4 report the analysis when income is included in the property returns calculation and their associated market indices. The key results are qualitatively unchanged. The penultimate row reports a probability of 14% that the estimates of α, β and α_1 differ by chance from those in the baseline case. Although the atemporal variance does not significantly change when including income, the diffusion variance estimate, σ^2, significantly declines relative to the baseline case. This is likely because income (which is reinvested at prevailing LIBOR rates) is less volatile than prices, and the impact of the reinvestment strategy is more pronounced for long-duration returns. Regardless, the exclusion of income does not appear to be a factor in the rejection of the null and, as in the baseline case, the remaining analysis will exclude it.

[16] The total diffusion variance is calculated by estimating the specification without a benchmark return. Consistent with the view that it is transactional and idiosyncratic, the estimate of σ_0^2 is insensitive to inclusion of the benchmark return.

Summarizing, deviations from RWD in the form of atemporal variance and alpha are large and hold across various specifications and controls. Estimates of the atemporal variance are consistently around 3% and several times larger than estimates of the idiosyncratic annual diffusion variance. Likewise, the atemporal alpha is consistently estimated at double digit figures (in percentage points).

2.4 Robustness and evidence of a transactional source

The analysis reported in this section is qualitatively robust to various alternative specifications, such as: Excluding properties purchased before 1997, including non-NPI properties, using only non-NPI properties, including properties with holding period less than one year, and removing the constraint that the purchase date must not lag the initial reporting date for the property by more than a quarter. Details can be found in Appendix A.4.

When repeating the exercise of Subsection 2.3 with matched REIT returns substituted for property returns, the results are consistent with the Null, suggesting that the phenomenon is specific to how real estate is traded. Alternatively, when one randomly matches the return of a single property held over consecutive periods Δ_1 and Δ_2 with the compounded returns from two properties, one of which was held over Δ_1 and the other over Δ_2, the return variance of the single property is lower than the variance of the compounded returns from the two matched properties by essentially the estimate of σ_0^2 from Table 4. 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). These exercises suggest that the anomalous return behavior stems from transactional frictions.

External validity for the findings of this section is documented for other highly illiquid asset classes, such as residential real estate and individual private equity deals. In nearly all studies, atemporal variance and alpha are attributed to unobserved investment, dismissed as a curiosity, or noted but not investigated further. For instance, 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) attribute the findings to missing data on renovations (i.e., capital expenditures), and model it by adding a jump component with positive mean to explain holding period returns. In the context of private equity buyout deals, 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 (2015) 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.

By incorporating CapEx into the calculation of holding period returns, I am able to reject interim investment as the source of the anomalous return behavior.[17] In contemporaneous and highly complementary work, Giacoletti (2017) includes proxies for home improvement expenditures in his return calculations and continues to find that holding period returns deviate from the RWD prediction.[18]

Theoretically, it is possible that property prices follow a RWD process but that the repeat sales data suffers from systematic selection bias that arises because purchase and sale decisions depend on the return distribution (i.e., τ and $\tilde{\varepsilon}$ are correlated in Eq. (3)). One simple possibility is that institutions prefer to hold riskier assets for 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. Alternatively, the anomalous atemporal alpha may result from a disposition effect in which investors are eager to realize gains but reluctant to do so for losses. Appendix A.5 examines and provides evidence against these alternative explanations.

3 A model of holding-period returns for illiquid assets

The robustness exercises reported in Section 2.4 suggest a transactional source for the holding period return anomalies. Commercial real estate assets are highly illiquid, taking months to transact, and there is little possibility of systematically exploiting the documented anomalous short-term returns. A

[17] Peng (2016) also includes CapEx in his calculations of holding period commercial property returns and notes that the variance of idiosyncratic return does not strongly depend on the holding period. He does not delve into the possible reasons for this nor consider the implications for his assumption of a RWD process.

[18] Perhaps because he uses a local index to control for systematic price variation, Giacoletti (2017) finds that nearly all of the idiosyncratic risk is atemporal. My sample size of national CRE properties is a small fraction of his residential data, which limits use of a more granular index.

natural conjecture is that illiquidity borne of market frictions may help explain the peculiar properties of observed holding period returns. Search and matching models have been extensively used in the literature to shed light on the dynamics of and relationships between aggregate quantities in illiquid markets. It is therefore sensible to look to such a model for an explanation of asset-level return behavior.

3.1 Model details

Time is discrete and there are N infinitely-lived income-producing properties. An investor's *valuation type* is defined to be an element of an index set $\mathcal{A}$, and each property is held and managed by some investor of type $a \in \mathcal{A}$. At date $t+1$, a property owned by an investor of type a will pay income $\tilde{d}_{t+1} = d_t e^{(\mu_t - \frac{\sigma^2}{2}) + \sigma \tilde{\epsilon}_{t+1}}$, 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, and the volatility σ is constant. The Jensen's Inequality term, $-\frac{\sigma^2}{2}$ ensures that $E_t[\frac{\tilde{d}_{t+1}}{d_t}] = e^{\mu_t}$. The drift, μ_t, a Markov process to be specified soon, is the same across all properties and is independent of $\tilde{\epsilon}_{t'}$ for any t'. Thus, μ_t may be interpreted as a macro state, and assumed to be specified by an index set, $\mathcal{S}$. If at date t the growth state is $s \in \mathcal{S}$ then I abuse notation somewhat in referring to the corresponding growth rate as μ_s. I assume that $\tilde{\epsilon}_{t+1}$ 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. To avoid burdensome notation, I suppress reference to any specific property.

For trade to take place, there must be gains from trade to all parties and, therefore, heterogeneity in private valuations. This is achieved here by assuming that an investor of type $a \in \mathcal{A}$ discounts next period's expected income and private property value by a factor $e^{-r_{a,s}}$, which varies across investors and may depend on the macro state. Although this approach is chosen for its tractability, cross-sectional heterogeneity in investors' discount rates may be viewed as a reduced-form proxy for the effects of institutional liquidity and capital constraints, individual managerial beliefs and/or preferences, skill, unspecified portfolio effects or hedging needs, fiduciary contractual constraints, and agency concerns.

The sequence of events each period is depicted in Figure 1 and proceeds as follows: Income is first distributed to each property's owner. The macro state then transitions, followed by possible changes to individual investor types. Next, each property's distress status is determined and this is followed by a match with a prospective buyer. If a transaction is suitable for both parties, property ownership is transferred. At the beginning of the next period, the sequence repeats.

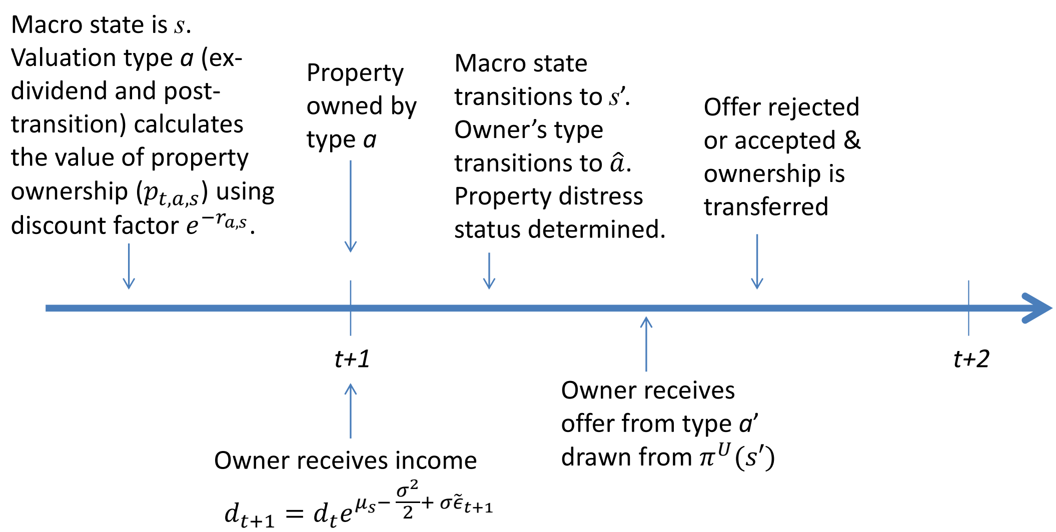

Fig. 1: Time line representing the sequence of events between dates $t+1$ and $t+2$. This is the relevant time line for an ex-dividend valuation of the property between dates t and $t+1$ by an investor or owner of type a in state s (post-transition).

States transition according to the Markov transition matrix, $\Pi_{\mathcal{S}}$, assumed to be regular to guarantee mean reversion. Note that if private values are heterogeneous but do not change, then eventually all properties would be owned by the investors with the lowest discount rate. This is avoided by assuming that individual types also transition according to a regular Markov matrix $\Pi_{\mathcal{A}}(s)$, which can depend on the prevailing macro state, say s, but is otherwise applied independently across investors. I employ the convention that the incumbent state corresponds to the row index of $\Pi_{\mathcal{S}}$, and similar for $\Pi_{\mathcal{A}}(s)$.

After transitioning from state s to s', a property may enter into "distress" with probability $\rho_{dist}(s)$, in which case it is sold at a distressed price of $Q_{dist}(s) \times d_t$ to the next bidder.[19] This captures the extreme left tail of transactions observed in the data. Each owner of a non-distressed property receives a purchase offer from some randomly chosen investor and must decide whether or not to sell at a cost, $c \times d_t$, $c \geq 0$, assumed for analytic convenience to be proportional to the property's income. Offers are sampled from the unconditional distribution of valuation types, denoted as π^U, conditioning on the

[19]It is assumed that distressed prices are below the lowest private valuation.

new state, s'.[20] This is consistent with absence of search costs or explicit constraints on the number of properties that an investor may hold.[21] In particular, the ratio of investors to properties is immaterial.

A sale at date $t+1$, when the macro growth state is s', takes place between an owner of type $\hat{a}$ and investor of type a' if and only if the owner's valuation of the property, $p_{t+1,\hat{a},s'}$, is smaller than the investor's valuation, $p_{t+1,a',s'}$, less cd_{t+1}. If the difference between bidder's and owner's private values exceeds the transaction costs, bargaining ensues and 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 $p_{t+1,\hat{a},s'} + \tilde{\lambda}\Big\{p_{t+1,a',s'} - p_{t+1,\hat{a},s'} - cd_{t+1}\Big\}^+$ (where $\{y\}^+ = \max\{0,y\}$). I assume that $\tilde{\lambda}$ is identically and independently distributed across time and buyers/sellers, and independent of $\tilde{\epsilon}_{t+1}$ and the Markov chain process underlying macro or valuation type transitions.

Combining the sequence of events outlined above, the ex-dividend and post-transition private value of a non-distressed property owner of type a in macro state s is calculated as

$$p_{t,a,s} = e^{-r_{a,s}} E\Bigg[\tilde{d}_{t+1} + (1-\tilde{\rho}_{dist}(s'))\Big(\tilde{p}_{t+1,\hat{a},s'} + \tilde{\lambda}\Big\{\tilde{p}_{t+1,a',s'} - \tilde{p}_{t+1,\hat{a},s'} - c\tilde{d}_{t+1}\Big\}^+\Big) + \tilde{\rho}_{dist}(s')\tilde{d}_{t+1}\tilde{Q}_{dist}(s')\Bigg], \quad (4)$$

where the tilde denotes random variables. The owner's and bidder's private values at date $t+1$ and macro state s', ex-dividend and post transition, are respectively denoted by $\tilde{p}_{t+1,\hat{a},s'}$ and $\tilde{p}_{t+1,a',s'}$. The last term in the expectation corresponds to the payoffs from a distressed sale. In words, the owner's valuation equals the expected continuation value of holding the income producing property plus the option value of selling (at a cost) to a prospective buyer, both discounted for time and the likelihood of distress, plus the present value of a potential distressed sale.

Definition. *An equilibrium is a positive and finite random variable $p_{t,a,s}$ that solves* (4) *for every*

[20] π^U is defined to be any one of the identical rows resulting from the limit transition matrix, $\lim_{n\to\infty}\Pi^n_{\mathcal{SA}}$, where $(\Pi_{\mathcal{SA}})_{s,a;s',a'} = (\Pi_{\mathcal{S}})_{ss'}(\Pi_{\mathcal{A}}(s'))_{aa'}$. Regularity and the so-called Fundamental Theorem of Markov Chains ensures existence and uniqueness of π^U in this limit. Conditioning on s' amounts to restricting attention to elements of π^U corresponding to s' (normalized so that they sum to one).

[21] Capital constraints are instead implicit in the variation across discount rates.

$a \in \mathcal{A}$ and $s \in \mathcal{S}$.

If $\mathcal{A}$ is a singleton set, then valuations are homogeneous across investors and $p_{t,s} = e^{-r_s}E[\tilde{d}_{t+1} + p_{t+1,s'}]$ defines the equilibrium in the absence of distress. Thus, if all investors are identical then prices are set as if the market is frictionless and each investor discounts cash flow at a (macro) state-dependent rate r_s. Liquidity has no role to play in such a market because there are no gains from trade.

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 at most a single potential buyer rather than a market of potential buyers. 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.[22] It is instructive to consider a situation where the owner faces multiple bidders, each arriving with independent probability and possessing different valuation and bargaining power. In this case, and temporarily suppressing the macro-state dependence or distressed sales, (4) becomes

$$\begin{aligned} p_{t,a} = e^{-r_a}E\Big[\tilde{d}_{t+1} + \tilde{p}_{t+1,\hat{a}} + \max\Big\{0, &\tilde{\lambda}\big(\tilde{p}_{t+1,a'} - \tilde{p}_{t+1,\hat{a}} - c\tilde{d}_{t+1}\big) \\ &\tilde{\lambda}'\big(\tilde{p}_{t+1,a''} - \tilde{p}_{t+1,\hat{a}} - c\tilde{d}_{t+1}\big) \\ &\tilde{\lambda}''\big(\tilde{p}_{t+1,a'''} - \tilde{p}_{t+1,\hat{a}} - c\tilde{d}_{t+1}\big), \ldots \Big\}\Big]. \end{aligned}$$

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 transactions will still only occur at the highest private valuation, but owners' private values will lie between this valuation and a lower bound determined by c.

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

[22] 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).

transition) of

$$p_{t,a,s} = d_t Q_{a,s}.$$

Define $\eta_{a,s} \equiv e^{r_{a,s}-\mu_s}$ to be the growth-adjusted private capitalization factor of investor a in state s. Then from (4) and the model assumptions, $Q_{a,s}$ must solve the following system of piece-wise linear equations for every $a \in \mathcal{A}$ and $s \in \mathcal{S}$:

$$\begin{aligned} \eta_{a,s} Q_{a,s} = 1 &+ \sum_{\substack{a' \in \mathcal{A} \\ s' \in \mathcal{S}}} (\Pi_{\mathcal{S}})_{ss'} (\Pi_{\mathcal{A}}(s'))_{aa'} (1 - \rho_{dist}(s')) \Big(Q_{a',s'} + \bar{\lambda} \sum_{b \in \mathcal{A}} \pi_b^U(s') \Big\{ Q_{b,s'} - Q_{a',s'} - c \Big\}^+ \Big) \\ &+ \sum_{s' \in \mathcal{S}} (\Pi_{\mathcal{S}})_{ss'} \rho_{dist}(s') Q_{dist}(s'), \end{aligned} \tag{5}$$

where $\bar{\lambda}$ is the mean of the random bargaining variable, $\tilde{\lambda}$. If $\rho_{dist}(s) = 0$ and $\eta_{a,s} \equiv e^{r-\mu}$ across macro and type states 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. Sufficient conditions for the existence of a unique equilibrium are provided by the following application of a famous result in Blackwell (1965):

Theorem 1. *Let $\eta_{a,s} > 1$ for every $(a, s) \in \mathcal{A} \times \mathcal{S}$. Then (5) has a unique solution.*

Proof. See Appendix B □

A transaction takes place at date t if and only if an arriving buyer's private value less the transaction cost exceeds the private value of the seller. This is true if and only if $Q_{a',s} - Q_{a,s} \geq c$, where $Q_{a',s}$ corresponds to the valuation of the prospective investor in state s, while $Q_{a,s}$ to that of the incumbent owner.[23] Thus, the realization of a transaction is a random variable whose distribution depends on the incumbent owner type. If a trade occurs between an owner of type a and an investor of type a' at date t and state s, then the observed transaction price is

$$p_t(a, a', s) = d_t \Big(Q_{a,s} + \tilde{\lambda} \big(Q_{a',s} - Q_{a,s} - c \big) \Big) \text{ s.t. } \; Q_{a',s} - Q_{a,s} \geq c.$$

This expression is a function of property market characteristics as well as the identities of the seller

[23] I assume that the buyer pays all transaction costs in a distressed sale. Here, and elsewhere in the ensuing equations, if the property is in distress, then the corresponding expression for the transaction price is obtained by setting the owner's valuation and bargaining power to $Q_{dist}(s)$ and zero, respectively.

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). If one interprets an appraisal as an average over potential transaction prices, then one should observe transaction dispersion around appraisals.

3.2 Holding Period Returns

Consider a property that at date t and state s 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 $t+\tau$, by which point the economy transitions to state s' while current owner has transitioned to type $\hat{a}$ and sells to a buyer of type $b \in \mathcal{A}$. Then the observed price appreciation return corresponding to the repeat transaction is:

$$\tilde{R}_{t,\tau}(o,a,s,\hat{a},b,s') = \frac{p_{t+\tau}(\hat{a},b,s')}{p_t(o,a,s)} = \frac{Q_{\hat{a},s'} + \tilde{\lambda}'\big(Q_{b,s'} - Q_{\hat{a},s'} - c\big)}{Q_{o,s} + c + \tilde{\lambda}\big(Q_{a,s} - Q_{o,s} - c\big)} \, \frac{\tilde{d}_{t+\tau}}{d_t}$$
$$= \frac{Q_{\hat{a},s'} + \tilde{\lambda}'\big(Q_{b,s'} - Q_{\hat{a},s'} - c\big)}{Q_{o,s} + c + \tilde{\lambda}\big(Q_{a,s} - Q_{o,s} - c\big)} \, e^{\sum_{j=0}^{\tau-1}(\mu_{s(t+j)} - \frac{\sigma^2}{2}) + \sigma\sqrt{\tau}\tilde{n}},$$

where $s(t+j)$ is the growth state of the economy at date $t+j$, $\tilde{n} = \frac{1}{\sqrt{\tau}} \sum_{j=1}^{\tau} \tilde{\epsilon}_{t+j}$ 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, $c = 0$ and there is only one valuation type, so $\tilde{R}_{t,\tau}(o,a,s,\hat{a},b,s') = \exp\Big(\sum_{j=0}^{\tau-1}(\mu_{s(t+j)} - \frac{\sigma^2}{2}) + \sigma\sqrt{\tau}\tilde{n}\Big)$. This is the standard (time-varying) geometric random walk with drift (RWD) 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, $\tilde{R}_{t,\tau}(o,a,s,\hat{a},b,s')$ 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 may not be able to infer the property market parameters from repeat sales without also modeling (implicitly or explicitly) transaction dynamics.

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

$$\ln \tilde{R}_{t,\tau}(o,a,s,\hat{a},b,s') = \underbrace{\ln\Big(Q_{\hat{a},s'} + \tilde{\lambda}'\big(Q_{b,s'} - Q_{\hat{a},s'} - c\big)\Big)}_{\text{Selling shock}} - \underbrace{\ln\Big(Q_{o,s} + c + \tilde{\lambda}\big(Q_{a,s} - Q_{o,s} - c\big)\Big)}_{\text{Purchasing shock}} \tag{6}$$
$$+ \underbrace{\sigma\sqrt{\tau}\tilde{n}}_{\text{Income shock}} + \sum_{j=0}^{\tau-1}\big(\underbrace{\mu_{s(t+j)}}_{\text{Macro risk}} - \frac{\sigma^2}{2}\big).$$

In equation (6), 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 a 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

$$\ln \tilde{R}^I_{t,\tau}(o,a,s,\hat{a},b,s') = \underbrace{\ln\left(\frac{Q_{\hat{a},s'} + \tilde{\lambda}'\big(Q_{b,s'} - Q_{\hat{a},s'} - c\big)}{Q_{o,s} + c + \tilde{\lambda}\big(Q_{a,s} - Q_{o,s} - c\big)}\right)}_{\text{Transaction risk}} + \sigma_I\sqrt{\tau}\tilde{n}_{I,i} - \frac{\sigma_I^2}{2}\tau, \tag{7}$$

where it is assumed that, over the holding period, a well-diversified portfolio or index of properties will exhibit an expected rate of log-price appreciation of $\sum_{j=0}^{\tau-1}(\mu_{s(t+j)} - \frac{\sigma_M^2}{2})$, so that risk-adjustment eliminates the macro growth component from (6).

Equation (7) describes the main object of interest in Section 2. In a perfectly efficient market, only the last two terms contribute to Eq. (2), and only the Jensen's term contributes to α_i. The idiosyncratic variance of the shock component, $\sigma_I\sqrt{\tau}\tilde{n}_{I,i}$, grows linearly with the holding period. Thus, in the limit of perfect liquidity, the model reduces to the standard RWD (or diffusion) price dynamics. The first term, corresponding to *transaction risk*, subsumes the return impact of selling and purchasing shocks which, in this model, arise from search frictions and transaction costs. Atemporal alpha and variance arise from the mean and variance of transaction risk.

To characterize the transactional return attributes corresponding to the purchase and sale shocks, one must characterize their joint distribution. It should be immediately clear that selection plays a role in determining this joint distribution. For instance, in the presence of transaction costs, the lowest valuation type would never purchase the asset and the highest valuation type would never sell.

A steady state is characterized by a distribution of ownership that is not expected to change. Let $\pi_{t,\tau}(o, a, s, \hat{a}, b, s')$ be the probability of realizing the following path in a generic property's history. The 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, when the macro state is s. The property is then held without being sold (despite the arrival of offers) until the new owner transitions to type $\hat{a}$ at date $t+\tau$ and state s'. Following this last type transition the owner receives a satisfactory bid from an investor of type b and the property is sold. Observing a holding period return is tantamount to observing one of these paths. Appendix B derives $\pi_{t,\tau}(o, a, s, \hat{a}, b, s')$ which, given the model parameters, can be used to calculate the joint distribution of the purchasing and selling shocks in (7). Other key steady state observables, useful in calibrating the model to data, are calculated in Appendix B. These include the average property turnover rates, proportional transaction costs, transaction cap rates, and the distribution of holding periods.

3.3 Qualitative Model Predictions

Transaction risk in Eq. (7) contributes to the holding period return volatility, even if the holding period is short. This is a simple implication of nearly all random matching and bargaining models, including those set in continuous time. Sufficient criteria are that expected gains from trade are finite even as the holding period vanishes, and that bargaining outcomes are random and uncorrelated.

There are various ways by which the model can produce negative average long holding period returns, as observed in Table 3 and implied by the estimates in Table 4. Suppose, for instance, that the arrival rate of high bids is lower than the rate at which valuations mean revert. Then, in the steady state, the average valuation of owners who have not managed to sell for a long time will tend to decline relative to the high average valuation of newly purchasing investors. This, together with the negative contribution from Jensen's term in Eq. (7), can lead to $E\left[\ln \tilde{R}^I_{t,\tau}(o, a, s, \hat{a}, b, s')\right] < 0$ for large τ.

More challenging to explain is the positive average return for short holding periods. To see the problem, consider the case where there is no persistence in private values and ignore variations in the macro state (which should not matter for average idiosyncratic returns). In this case, each row of $\Pi_{\mathcal{A}}(s)$ is identical, implying that selling and buying shocks are unrelated at all horizons. If $s = s'$, the selling shock component in (6) can be denoted as $\ln \tilde{A}$ and the purchasing shock contributes $-\ln\left(\tilde{A}' + c\right)$, where $\tilde{A}$ and $\tilde{A}'$ are identically and independently distributed.[24] Because $\tilde{A}' + c$ first-degree stochastically dominates $\tilde{A}$, it must be that $E[\ln \frac{\tilde{A}}{\tilde{A}'+c}] \leq 0$ and $E\big[\ln \tilde{R}^I_{t,\tau}(o, a, s, \hat{a}, b, s)\big] < 0$ for any horizon.

While persistence appears necessary for atemporal alpha, it is not sufficient. To see this, suppose that there are only two valuation 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 another high type. Any repeat transaction consists of buying high and selling low, leading to negative expected alphas in the observed transaction regardless of the horizon. Thus, positive atemporal alpha requires the presence of intermediate valuation types.

To see how the combination of persistence and intermediate valuation types leads to positive average returns in observed short-hold transactions, consider a situation with three highly persistent types, $Q_H > Q_M > Q_L$. When an asset is first purchased at date t by a type a investor, the owner's type is highly unlikely to change in the next period. A sale at $t + 1$ by a type $a = H$ will not be observed because transaction costs preclude a sale to another investor with the same valuation. Thus a sale after only one period of ownership is most likely to take place between a type M (who recently purchased from a type L) and a bidder with type H. In other words, when types are persistent, a short holding period is most likely to feature an initial purchase by an intermediate valuation type, followed by the arrival of and sale to a new buyer with an even higher valuation. This repeat transaction links large alphas with short holding periods, but the causality is reversed: The apparent "premium" reflects the chance arrival of a high offer shortly after an initial purchase and represents the upside of a range of outcomes. The downside is not observed because the property will be otherwise held longer. This example demonstrates selection bias in transactions data and highlights a role for intermediate value investors as property "flippers". The presence of such intermediaries, and the profits that they make in equilibrium, reflects information about structural illiquidity in the market, which is most pronounced

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

in short-term transactions.

3.4 Empirical validation of the Model Mechanisms

The previous subsection outlines how the existence of valuation dispersion and the presence of transitional investors can explain the return anomalies. Here, to provide empirical validation, I investigate proxies for these mechanisms.

For their NPI-qualifying properties, NCREIF members are required to report quarterly appraisal values where at least once every three years the appraisal is done by an independent (external) appraiser. In practice, roughly one third of the market values reported in the cleaned data (used in this study) correspond to external appraisals. If valuations differ across market participants, then an external appraiser, a willing buyer, and a willing seller will generally have different views of the property's value. This would be expressed as variation in the ratio of a sale price and an appraisal that closely precedes the sale.[25] Following this intuition, I first calculate the Sale-Appraisal Ratio as the property sale price divided by the (externally) appraised market value one quarter prior to the sale (available for 2,478 properties). The distribution of this ratio is fat-tailed, right-skewed, and its standard deviation is 12.0% (far exceeding dispersion in transaction costs). The median and mean of the ratio are 1.000 and 1.011, respectively, suggesting that appraisals are relatively unbiased measures. In each year, I then construct a "disagreement measure", Dis_{app}, equal to the variance of the Sale-Appraisal Ratio in that year. I do this annually because calculating a quarterly disagreement measure would exhibit too much noise in the first half of the sample when there are fewer transactions per quarter.

I hand-collect data from the Urban Land Institute's *Emerging Trends in Real Estate*, an annual survey of CRE stakeholders available from 2002. The survey consistently reports the distribution of buy/hold/sell views from its respondents for the major CRE asset categories and some subcategories in the overall U.S. property market.[26] The correlation between buy and sell recommendations is large in magnitude and negative (-0.75). I use Dis_{et}, the difference between the proportion of sell and buy recommendations divided by their sum, to proxy for the spread in private values. The idea is that, in

[25] Buyers, sellers, and financing stakeholders (e.g., banks, JV partners) will undertake their own appraisals. It is unlikely that appraisals solicited by different stakeholders would be performed by the same appraiser.
[26] An example is available at `https://goo.gl/D9U8G4`.

equilibrium, a high level of valuation disagreement across property owners can only be sustained if there are very few buyers able to take advantage of the discrepancy. Separately, I consider Trans_{et}, defined as one minus the proportion of hold recommendations, to proxy for the depth of the market and extent of transitional investor activity. An unusually high number of hold recommendations corresponds to a thin market and one where transitional investors are less likely to prosper. The two measures are matched to the return sample (by year and property subtype). It is noteworthy that the correlation between Dis_{et} and Trans_{et} is less than 10%.

A property owned through a joint venture (JV) partnership reflects a degree of consensus valuation among the partners.[27] Thus, on average, the effective JV private value of a property is less likely to be extreme. Consistent with that idea, JVs may be viewed as transitional investors in the model (at least while the JV is in place). Only 83% of JV properties enter the (cleaned) NCREIF dataset as JVs, and about 91% of properties that were owned as a JV at some point are sold as a JV. While 2,258 properties are JV-owned at some point, some move in and out of JV ownership (there are 2,750 spells of contiguous periods during which the ownership is classified as JV). JV spells terminated by a property sale are significantly shorter than the unconditional distribution of holding periods documented Table 1. Consistent with that, Table A-VI in Appendix A reports that a JV property is significantly more likely to be sold. When a property is purchased as a JV, the average purchase price as a proportion of next period's (independent) appraisal is significantly lower by 2.8% than non-JV purchases. By contrast, when a JV sells, there is no significant discount relative to non-JV sales. This suggests that JVs behave as transitional investors who purchase if the price is right and sell in the same market, and with the same market outcomes, as other sellers. The only difference with other sellers is that JVs are more open to disposition soon after a purchase. This seems to fit well the description of the intermediate value investor that is key to the model mechanics. Correspondingly, I use the proportion of properties owned by a JV, defined to be Trans_{JV}, as another measure for the presence of transitional investors.

The distributional properties of the proxies are reported in Table B-I of Appendix B. To formally test for a relationship between the proxies and anomalous return moments, one can interact the proxies with α_0 and σ_0^2 in a repeat of the analysis in Table 4. I de-mean and standardize the proxies using their

[27] JV partnerships tend to have buy-sell agreements allowing each partner to force a disposition if the partners disagree on strategy or value.

sample means and standard deviations, to more easily interpret the regression coefficients, and reports the results in Table 5.

Table 5: The table reports estimates of regression parameters in Table 4 including interaction terms for α_0 and σ_0^2 using proxies for investor disagreement ("Dis") and the presence of transitional investors ("Trans") — see Section 3.4 for an explanation of the proxies. Interaction variables are standardized (denoted by "std(·)") to facilitate interpretation. The penultimate row reports the probability that the stated parameters are unchanged relative to their baseline estimates.

	(1)	(2)	(3)		(1)	(2)	(3)
β	1.0141***	1.0750***	0.9731***	σ^2	0.0082**	0.0094*	0.0087***
	(0.0605)	(0.1345)	(0.0595)		(0.0029)	(0.0045)	(0.0014)
α	-0.0055	-0.0466**	0.0064				
	(0.0111)	(0.0161)	(0.0083)				
α_0	0.1506***	0.1559**	0.1995***	σ_0^2	0.033***	0.0397***	0.0304***
	(0.0401)	(0.0584)	(0.0386)		(0.0043)	(0.006)	(0.0041)
std(Dis_{app})	0.018**				0.0054		
	(0.0057)				(0.0032)		
std(Dis_{et})		-0.0124				0.0081**	
		(0.0065)				(0.0029)	
std(Trans_{et})		0.028***				-0.0021	
		(0.0068)				(0.0031)	
std(Trans_{JV})			0.0312***				0.0039
			(0.0048)				(0.002)
α_1	-0.0736	-0.0629	-0.1291***				
	(0.0377)	(0.0577)	(0.0361)				
Prb no change in β, α & α_1	0.5595	0.0003	0.9145	Prb no change in σ^2	0.917	0.8447	0.8799
Observations	3862	2159	4232		3862	2159	4232

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Both proxies for the presence of intermediate value investors are associated with positive atemporal alpha, and both disagreement proxies are associated with positive atemporal variance. These interactions are highly significant but for std(Dis_{app}), which has a t-statistic of 1.69.[28] The remaining coefficients $(\beta, \alpha, \alpha_1, \sigma^2)$ mostly remain the same across the different specifications. The notable exception is the regression with the survey-based proxies for which α is large and negative and α_1 has the smallest

[28] I include the survey-based variables together in one regression because they share a common data source and time-period. If one combines std(Dis_{app}) and std(Trans_{JV}) in the same regression, the results are *strengthened* (i.e., std(Dis_{app}) becomes significantly and positively related to atemporal variance).

magnitude relative to other specification. This could be because the time series panel for the survey data is shorter by twenty years than the other panels, making it hard to disentangle the effects of α from those of α_1 (the correlation of the two estimates in the second specification is -79%).

Appendix B similarly documents the impact of fund size and closed-end funds (CEFs) on the anomalous return moments. Consistent with the hypothesis that larger investors seek more liquidity, the atemporal variance associated with smaller funds is significantly higher by 250 basis points. Likewise, CEF managers may be viewed as skilled asset allocators with a limited holding horizon built into their contract. Confirming that, I find that CEF properties are associated with a high atemporal alpha separate from that observed for JVs. All of these results qualitatively hold when employing a simpler analysis, along the lines of Table 3.

3.5 Comparison with Search Literature

Han and Strange (2015) comprehensively review the vast literature on search-based transactions in real estate. Save for Fisher, Gatzlaff, Geltner, and Haurin (2003) which uses a heuristic rather than an equilibrium search framework, research in this area has focused exclusively on residential housing. Because such studies aim at understanding aggregated real estate market variables (e.g., supply, demand, prices, turnover, etc.), less attention has been paid to matching the transaction process to data or gauging the sensitivity of aggregate quantities (like repeat sales indices) to the transaction process. This paper helps to bridge that gap.

Recent models generally employ either random or directed matching of counterparties, followed by a "two-sided" transaction process.[29] As with the model developed here, most random matching search models feature heterogeneous private values and gains from trade that are randomly split between buyer and seller (Krainer and LeRoy, 2002, is an exception). Atemporal transaction variance naturally emerges from this. By contrast, directed search models do not typically allow for such transaction price uncertainty (and atemporal variance) as they assume that sellers commit to listing prices (See Diaz and

[29] In random matching the counterparties meet through a process that does not condition on attributes of buyers or sellers. In directed matching, at least one of the counterparties can take an action (e.g., commit to a listing price) that impacts the conditional probability of being matched. Two-sided transactions refers to a process in which a buyer and seller, after meeting, each have to solve a choice problem in order to determine whether to transact. In a one-sided transaction process, the decision to transact is exogenously specified for one of the counterparties.

Jerez, 2013; Albrecht, Gautier, and Vroman, 2016; Hedlund, 2016a,b; Garriga and Hedlund, 2017).

Existing models, however, cannot explain the atemporal alpha. Search models in the literature generally employ one of two mechanisms to induce gains from trade in a given transaction.[30] In the first, conditioning on a match, a buyer and seller split a random match surplus determined only after the match is consummated (E.g., Williams, 1995; Novy-Marx, 2009; Genesove and Han, 2012). In the second mechanism, property owners become potential sellers only after receiving a memory-less "disutility shock" to ownership (See Krainer, 2001; Krainer and LeRoy, 2002; Caplin and Leahy, 2011; Head and Lloyd-Ellis, 2012; Diaz and Jerez, 2013; Ngai and Tenreyro, 2014; Head, Lloyd-Ellis, and Sun, 2014; McQuade, Guren, et al., 2015; Hedlund, 2016b,a; Garriga and Hedlund, 2017). Holding the state of the economy constant, in the first mechanism there is no correlation between the owner's valuation at purchase and at sale, at any horizon. While in the second mechanism, a sale can only take place if the owner's valuation falls to a level that is independent of their valuation as a buyer. These assumptions imply that average short term returns, and therefore atemporal alpha, cannot be positive. The model I derive is therefore a unique contribution to the real estate literature.

Search is also prevalent in models of over the counter (OTC) markets (Duffie, Gârleanu, and Pedersen, 2005, 2007; Lagos and Rocheteau, 2009). These and related models employ a two-sided random matching paradigm with iid type switching. Thus atemporal variance will typically be a feature of holding period returns.[31] As explained in Hugonnier, Lester, and Weill (2018), theirs is the only decentralized OTC model that assumes more than two types of investors.

To the best of my knowledge only two models in the broader search literature, both contemporaneous with this paper, feature the mechanisms required to deliver atemporal alpha and variance (random matching and three or more persistent valuation types). Investors in the OTC model of Hugonnier, Lester, and Weill (2018) are constrained to holding either one or no asset, which in equilibrium keeps most high valuation investors out of the market. While this allows them to explain intermediation chains in decentralized over the counter markets it may not result in an adequate description of institutional

[30] Two exceptions are Wheaton (1990) and Albrecht, Anderson, Smith, and Vroman (2007). The former only examines aggregate variables rather than the specifics of individual transactions, and the latter does not feature repeat transactions (sellers enter the market exogenously).

[31] In the presence of competitive dealers, through which all trade must be transacted, the transaction price is set irrespective of private values thereby doing away with atemporal holding period variance.

CRE markets where trading costs are high and intermediation chains rarely occur.[32] In the art auction market model of Lovo and Spaenjers (2018), atemporal variance arises because the number of auction bidders is finite. Atemporal alpha arises because there are more than two persistent private valuation types. The specific transaction mechanisms (in the models and in practice) are different between CRE and art asset markets. This is also true of the structure of the underlying shocks to the economy and to types. Conceptually, however, the model in Lovo and Spaenjers (2018) is very similar to mine.

4 Fitting the model to the data

To explore whether the model delivers a plausible quantitative description of observed return and transaction moments, I calibrate a boom and bust version of the model to NCREIF data. In calibrating, I exclude closed-end fund properties from the data because they face contractual liquidation incentives that depend on their holding periods.[33] Appendix B.4 establishes that atemporal alpha and variance are robust to the exclusion of closed-end fund assets and discusses how, relative to other institutional owners, closed-end funds target different assets and asset management strategies.

4.1 Model parameters, data moments and calibration procedure

Model calibration choices are summarized below (Appendix C provides extensive details).

Model parameters To match the NCREIF panel, each period represents a quarter. I assume two macro states, corresponding to an expanding/contracting real estate market. To calibrate these, I estimate a two-state Markov switching model to the national NPI price appreciation index returns reported by the NCREIF. The estimated smoothed probabilities identify two episodes of national property market contractions between 1978 and 2017: One from 1990Q3 until 1994Q1, and one from 2008Q3 until 2010Q1.[34] The probabilities of switching from expansion to contraction, and vice versa, are 0.015 and

[32] Hugonnier, Lester, and Weill (2018) predict that assets will typically drift up the "value" chain through intermediation. In the NCREIF data, there is no significant relationship between the rank of a property's purchase price relative to appraisal and the subsequent rank of selling price relative to appraisal.

[33] Figure A-I(a) in Appendix A.4 illustrates the exceptionally low returns for holding periods of 8-10 years, much of which stem from asset liquidations by closed-end funds, likely at or near the end of their term. Without explicitly including an investor with horizon constraints it would be difficult for the model to fit to these specific data points.

[34] I identify a market expansion whenever the estimated smoothed probability of an expansionary state is greater than one half (otherwise, the market is in contraction).

0.097 per quarter. This fixes $\Pi_{\mathcal{S}}$.

For ease of notation, I refer to the expansion state as $s = exp$ and the contraction state as $s = cnt$. Using the Sale-Appraisal ratio, introduced in Section 3.4, I define a distressed sale as a transaction below two standard deviations of a recent appraisal.[35] Applying this definition, I estimate distress probabilities of $\rho_{dist}(exp) = 0.0467\%$ and $\rho_{dist}(cnt) = 0.0924\%$. The corresponding discounted price-earning ratios in the different regimes (i.e., the $Q_{dist}(s)$'s) are left to freely vary subject to $Q_{dist}(cnt) \leq Q_{dist}(exp)$. To give the model the best chance at capturing right-tail transactions, I set the number of types to 21. The calculation of model statistics, is simplified by assuming that the bargaining power parameter, $\tilde{\lambda}$, is one or zero with one half chance. Listing and bidding data, which I do not have, could shed more light on the bargaining process.

The quarterly type transition matrix, $\Pi_{\mathcal{A}}(s)$, assumed symmetric about the middle investor type, is akin to a discretized bounded Ornstein-Uhlenbeck process supplemented with a jump reversion to the middle type. The diffusion "volatility" rate parameter ($x_s \geq 0$) controls the persistence of types (decreasing x_s increases persistence). Increasing the mean reversion strength (z_s) increases the rate at which types that are further away revert towards the middle. The "jump" reversion process shifts any given type to the middle with probability ξ_s. The six type-transition parameters, together with $\Pi_{\mathcal{S}}$, also determine the unconditional distribution of types, $\pi^U(s)$ in each state.[36] The growth-adjusted capitalization factor, $\eta_{a,s} = e^{r_{a,s}-\mu_s}$, is assumed to increase in a, and $\eta_{a,cnt} - 1$ is modeled as a (two-parameter) log-normal distribution with quantiles defined by $\pi^U_a(cnt)$. The $\eta_{a,exp}$'s are modeled similarly but constrained so that in moving from an expansion to a contraction, after accounting for the drop in income growth, type-specific discount rates (the $r_{a,s}$'s) are weakly increasing. Overall, ten parameters are employed in modeling the state-dependent distribution of $\eta_{a,s}$'s and their corresponding type transitions. Transaction costs (c), idiosyncratic volatility (σ_I), together with the $Q_{dist}(s)$'s, bring the number of free parameters to 14.

[35] As discussed later, steep discounts relative to appraised values in the data are likely driven by liquidity events rather than defaults. This makes it difficult to clearly identify "distress" and justifies the somewhat ad-hoc definition.

[36] Roughly, z_s and ξ_s respectively determine the standard deviation and kurtosis of the unconditional distribution of types. As will soon be explained, the jump also helps capture negative skewness in transaction prices.

Data moments The free model parameters are varied to fit 31 data "moments" that I now describe. I divide repeat sales capital gains data, calculated as in Eq. (1), into four categories based on whether a property is bought/sold in an expansion or a contraction, and then exclude the *cnt/cnt* category because it only contains six repeat sales. For the other three categories, I estimate idiosyncratic return means and variances, netting out the market and year fixed-effects. For calibration, I use the mean and variance estimates at horizons $\tau = \{1, 5, 8\}$, resulting in 18 model moment restrictions.

Next, conditional on an expansion or a contraction, I estimate means for property turnover, proportional transaction costs, transaction cap rates, and the proportion of purchased properties that are sold within five years. These contribute eight additional moment restrictions. In the case of transaction cap rates, I remove the observed secular trend over the sample period by subtracting the spread between the one-year TBill and inflation. The latter can be viewed as a monetary policy instrument and is therefore outside the model.

There are 2478 transactions preceded by independent appraisals within one quarter. These are used to construct the Sale-Appraisal Ratio, described in Section 3.4. If, in the model, one associates the appraised value with the average transaction property cap rate, then one can construct a model-implied distribution of the Sale-Appraisal Ratio.[37] The mean, variance, skewness, and kurtosis of this distribution provide four additional calibration moments.

Lastly, I restrict the interpolated probability of an above-market offer arrival (during an expansion) to be 22.5% per quarter, with a standard deviation tolerance of 2%. This is the only moment that is not directly derived from data, and is adopted to ensure that, during an expansionary quarter, the probability of a "fair market" offer within the next year is roughly between 50% and 75%, consistent with recent CoStar reports (see `https://prn.to/2X8AyuL`).

4.2 Calibration results and comparison with the data

To calibrate, I minimize the sum of standardized squared deviations of model moments from their empirical estimates. Details of the procedure and results of the calibration, as well as a sensitivity analysis, are found in Appendix C. Two of the 14 model parameters are at their binding constraints

[37] Details are in Appendix C.2.

($x_{cnt} = 0$ and $Q_{dist}(cnt) = Q_{dist}(exp)$), and a third, z_{cnt}, has no influence when $x_{cnt} = 0$. Thus, in practice, only 11 parameters are meaningfully employed.

The type transition parameters are generally highly sensitive to the turnover moments, and vice versa. Consistent with the qualitative analysis in Section 3.3, type-transition persistence significantly impacts short-term alpha. Calibrated idiosyncratic diffusion volatility is modest at roughly 12% per year, while transaction costs are roughly 2%-3% of property value. Both compare well with estimates from the data. Not surprisingly, σ_I^2 is most sensitive to the variance of longer holding period returns and c is most sensitive to transaction costs.

Table 6: Calibrated transaction statistics. The panel reports aggregate transaction statistics in property market expansion and contraction states. Acquisition income to price ratios (cap rates) are adjusted by the spread between one-year treasury rates and inflation to remove the secular trend in cap rates since the 1980's. The column "data" reports point estimates from the data while "sd" denotes the associated standard error. The column "model" reports the corresponding calibrated model value.

	Expansion State			Contraction State		
Statistic	data	sd	model	data	sd	model
Quarterly turnover (during...)	0.025	0.004	0.024	0.009	0.01	0.009
Average proportional transaction costs (props sold in...)	0.025	0.005	0.024	0.032	0.015	0.026
Average adjusted acq cap rate (props acquired in...)	0.060	0.008	0.060	0.072	0.017	0.068
Fraction sold within 5 years (prop acquired in...)	0.203	0.078	0.274	0.335	0.165	0.357

The fit to the data moments is generally very good. The worst fitting model moment, still within 1.56 standard deviations of its observed value, corresponds to the idiosyncratic holding period return variance for properties purchased during a contraction and sold one year later in an expansion. Of the remaining 30 moments, the model is able to fit 27 within one standard deviation. As seen in Table 6, the model captures both pricing attributes (the average trend-adjusted cap rate in each regime), and transaction dynamics (turnover statistics).

Figure 2 demonstrates that the model quantitatively explains the "anomalous" risk and return patterns in holding-period returns. For the three holding period categories used in the calibration, the figure plots estimates of holding period returns net of market and year fixed effects (circles) against model predictions (solid lines). Larger red circles denote the eighteen calibration target return moments. At every horizon and combination of states, the model prediction is within the 95% confidence interval of the empirical estimates (vertical dashed lines) for both targeted and non-targeted moments.

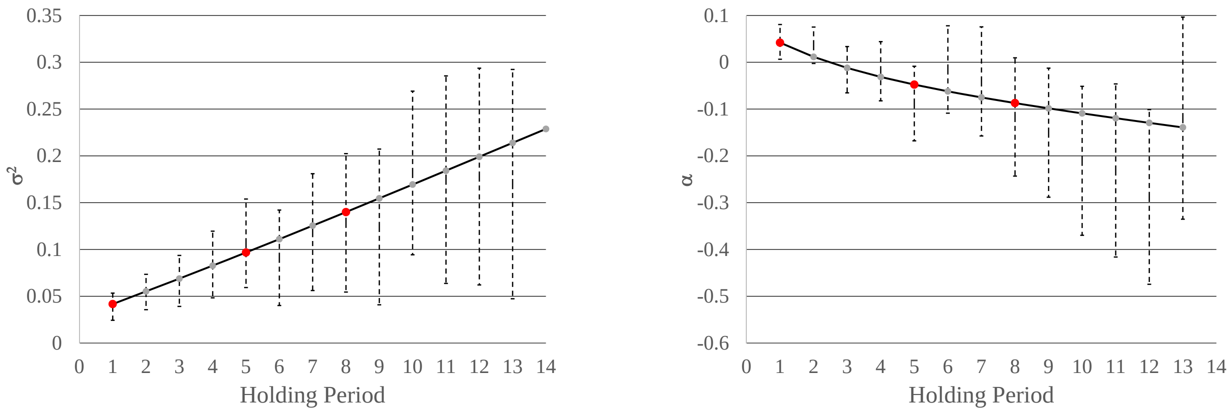

(a) Properties purchased and sold in expansion states

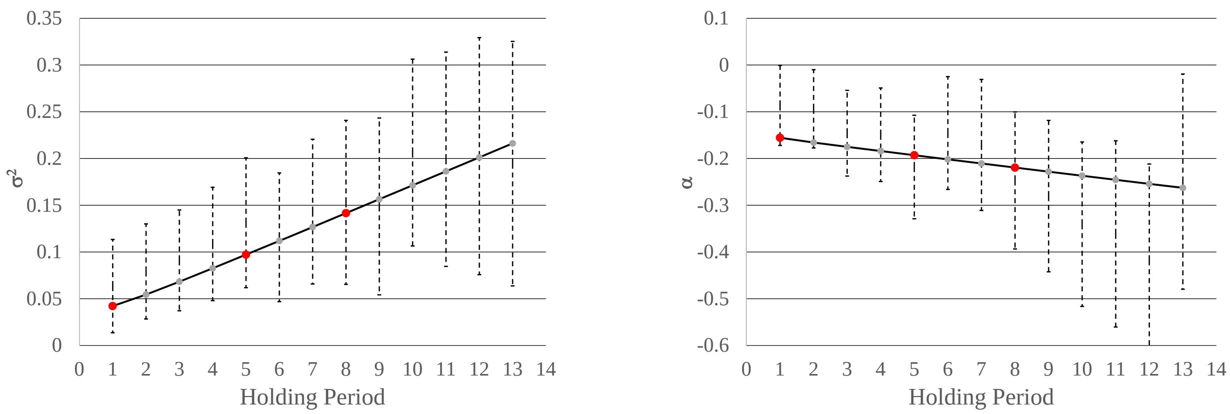

(b) Properties purchased in expansion states and sold in contraction states

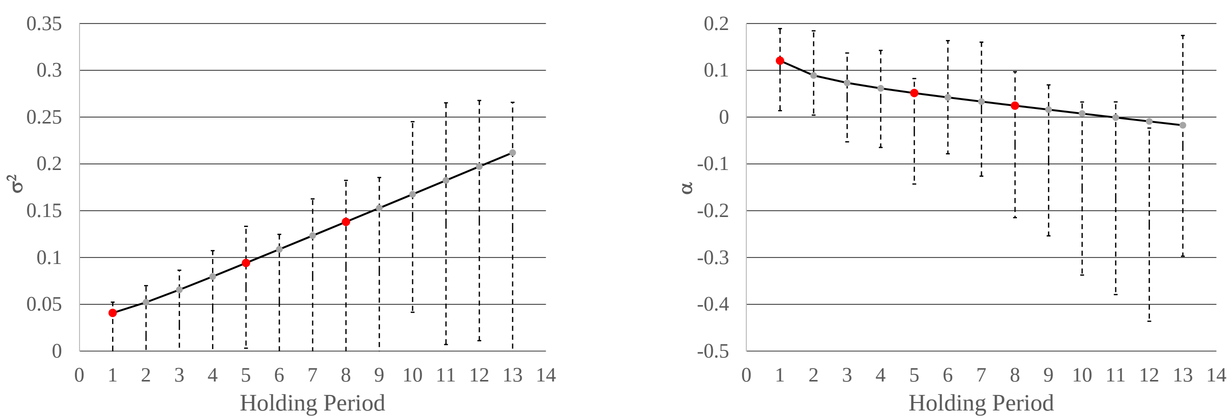

(c) Properties purchased in contraction states and sold in expansion states

Fig. 2: Holding period idiosyncratic variances (σ^2) and means (α) for different repeat sale regime categories. The dashed lines correspond to 95% confidence intervals from the NCREIF panel. The continuous thick line corresponds to the calibrated model predictions. Larger red points denote return moments used in the calibration (a total of 18). Because only six properties are purchased and sold in contraction states, the corresponding confidence intervals are too large to convey a meaningful sense of fit and plots for this regime category are not depicted.

Figure 3 demonstrates that the model can also produce realistic transaction heterogeneity in the form of model-versus-actual Sale-Appraisal Ratio distributions. Right tail transactions leading to positive skewness are well-captured, but details of the left tail are only coarsely reproduced. This might be addressed by allowing the jump diffusion to the middle type to instead be distributed across a wider spectrum of types below the middle. The extreme left tail is entirely absorbed by the distress shock and distributing its outcome more finely would likewise improve the fit to data.

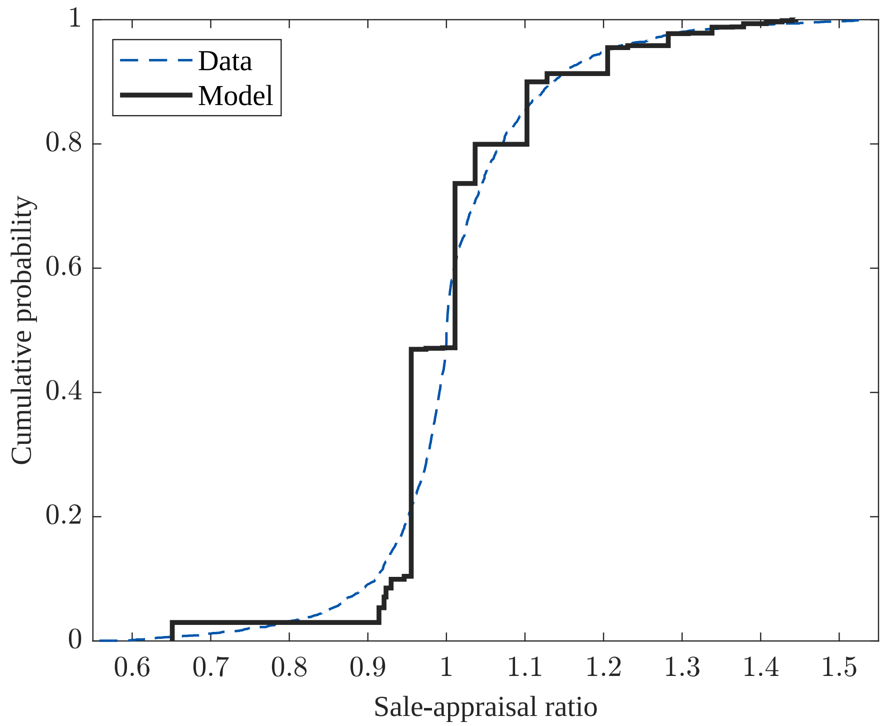

Fig. 3: Cumulative distribution of the Sale-Appraisal Ratio. The dashed line corresponds to transactions in the NCREIF data that are preceded, one quarter earlier, by independent appraisals. The solid line is the corresponding calculation from the calibrated model.

The calibrated type transition parameters determine the dynamics of an investor's *rank*, relative to other investors, in terms of her growth-adjusted discount rate. The $\eta_{a,s}$'s, determine the level of each rank. Jointly, these parameters define the distributions of bidders and owners in the steady state. To understand the key factors determining these parameters, first consider that if types do not transition, then all properties would eventually be owned by the highest valuation type. Correspondingly, without the jump transition component, persistence in the type diffusion concentrates ownership among high-valuation types. This, however, leads to *negatively* skewed transaction prices, inconsistent with data.

Adding the jump component ensures non-negligible middle-type ownership in the steady-state through a constant flow of owners who transition directly from high-valuation types. Fitting to low turnover and positive transaction price skewness, in this case, requires that above-middle type arrival is infrequent. Symmetry of the transition process subsequently forces below-middle type arrival to also be infrequent.[38] This suggests that, in fitting to the data, type arrival will concentrate around the middle rank.

Type transitions in the calibration are indeed found to be highly persistent, just as implied by the discussion in Section 3.3. Consistent with the argument above, the jump component is significant and the arrival distribution is concentrated around the middle type in both states. Although an expansion is characterized by a small jump component (about 1.25% per quarter), a randomly arriving investor is middle-type with 91% chance, and will be no more than three types away from the middle with 99.57% probability. During a contraction, the diffusion parameter is completely shut off ($x(cnt) = 0$) but the jump component is substantial (27.6% chance per quarter of shifting to the middle type). This further concentrates arrivals around the middle: A bid has a 98% chance of being middle-type and will be no more than one type away from the middle with 99.51% probability. The dramatic increase in concentration is needed to account for the large decline in transactions during a contraction.[39]

During contractions, the calibrated middle investor $\eta_{a,cnt}$ is 1.021, with the two nearest types at 1.017 and 1.027. During an expansion, the middle $\eta_{a,exp}$ is 1.018, with the six nearest types ranging from 1.006 to 1.025. Thus in moving from contraction to expansion, investors' growth-adjusted discount rates shift to the left and mass is redistributed to the tails.[40] To depict this graphically, I look for a continuous distribution that can be partitioned into "bins" so that both the bin masses and their means fit the point-mass distribution of $\eta_{a,s}$'s with respect to $\pi^U(s)$. Figure 4 plots the best-fitting log-Johnson (1949) distribution to the calibrated investor types in the expansion and contraction regimes.[41] Bins associated with distinct investor types are alternately shaded to assist with visualization, and the middle investor type bin is labeled as such. Save for the extreme left tail in the expansion regime, which is

[38] Recall that bid arrivals are determined by π^U, the unconditional distribution of types derived from the transition process. Symmetry in the transition process about the middle type results in symmetry in the bid arrival distribution.

[39] It is possible, without meaningfully changing the model results, to "rig" some of the below-middle types and their transition dynamics in order to distribute mass away from the middle type in the bid arrival distribution. I take the view that, absent bidding data to guide the exercise, the value of such a re-parameterization would primarily be aesthetic.

[40] Although the distribution of types is symmetric, the distribution of $\eta_{a,s}$'s assigned to the types is not. That's because $\eta_{a,s}$'s are assigned to types based on the quantiles of a log-normal (skewed) distribution rather than uniformly.

[41] The Johnson (1949) distribution is a flexible four-parameter generalization of the Normal distribution.

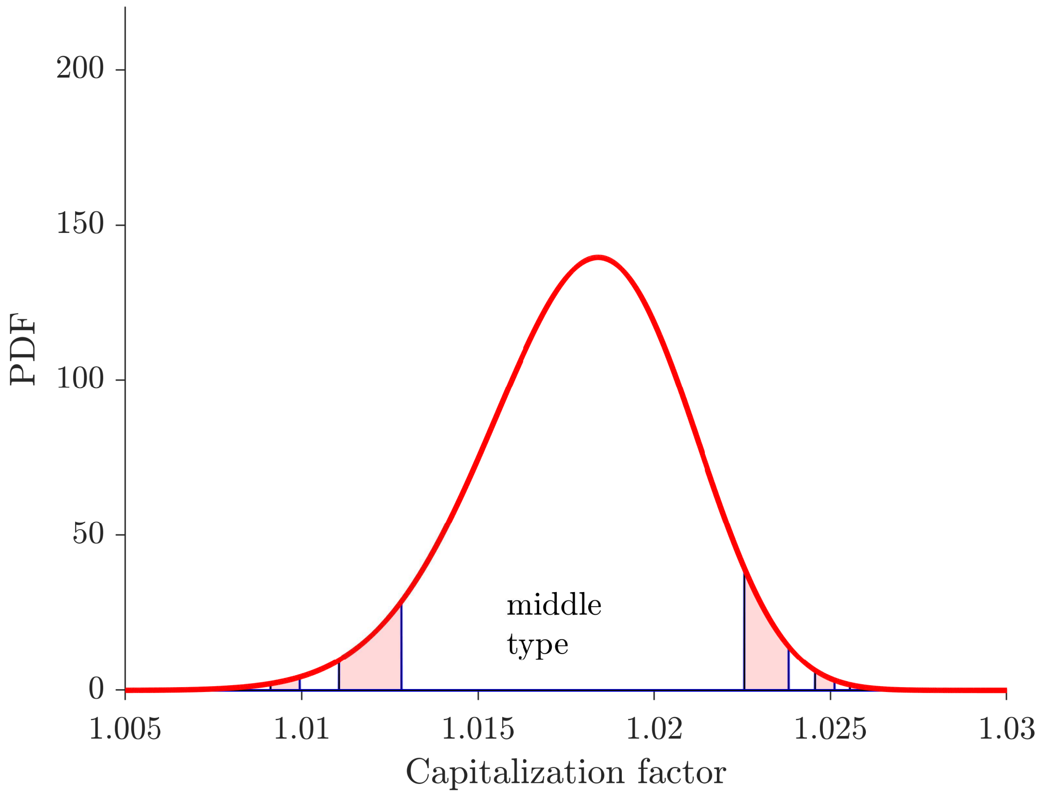

(a) Expansion state

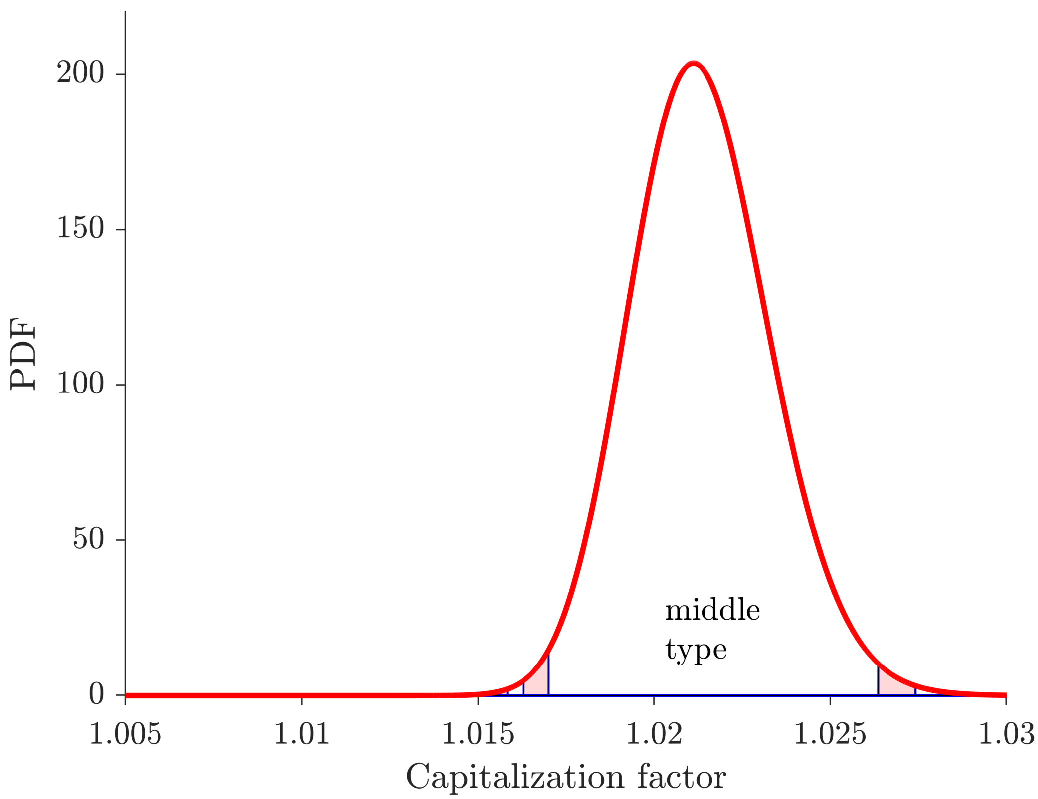

(b) Contraction state

Fig. 4: Continuous histogram representation of the distributions of growth-adjusted capitalization factors, $\eta_{a,s}$, in the calibrated model based on the estimated parameters. Types in the model are discrete and represented by the alternating shaded bins in the plotted distribution. The middle bin of 21 types is explicitly denoted. The unconditional mass, $\pi_a^U(s)$, corresponding to the random arrival of an investor of type a, is the area within each bin. The (discrete) value $\eta_{a,s}$ in each bin corresponds to the bin mean.

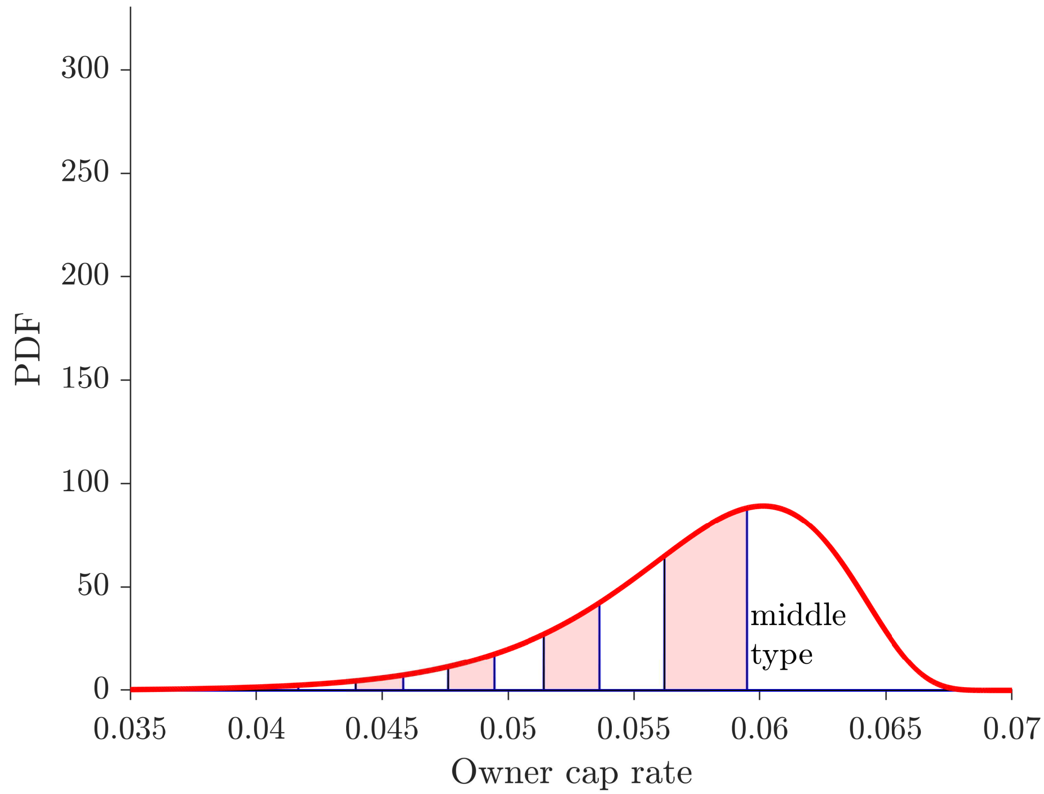

(a) Expansion state

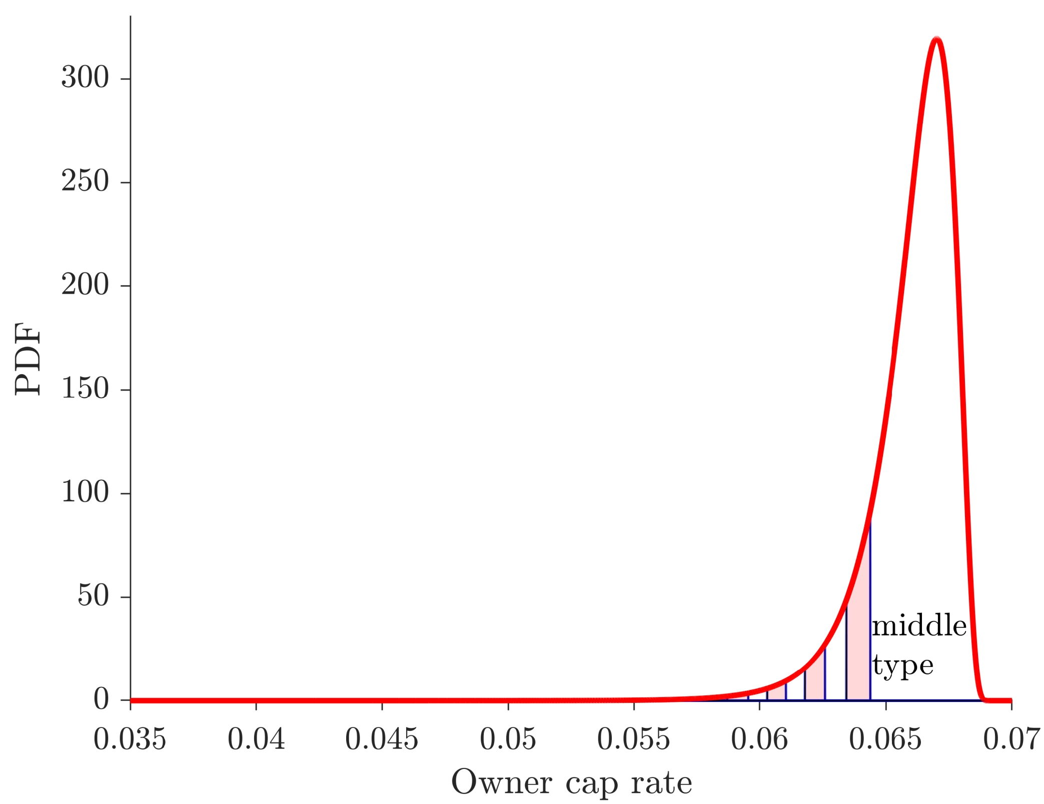

(b) Contraction state

Fig. 5: Continuous histogram representation of the distributions of steady state ownership in the calibrated model, as a function of the owners' private valuation (measured as a cap rate). Types in the model are discrete and represented by the alternating shaded bins in the plotted distribution. The middle bin of 21 types is explicitly denoted. The area within each bin is the share of the property market owned by the corresponds type. The (discrete) private valuation of associated with each type corresponds to the bin mean.

barely discernible, the bin masses and means provide a faithful rendering of how bids are allocated in the two regimes. As one might expect, the distribution of capitalization factors during contractions first-order stochastically dominates the distribution during expansions.

The distributions of the $\eta_{a,s}$'s in Figure 4 depict model arrival probabilities of the different buyer types and are therefore key drivers of transaction dynamics. Another key driver, depicted in Figure 5, is the distribution of owners' valuations in the steady state (expressed as cap rates). Here too I employ a binned continuous distribution where bin masses and their means fit the steady state point-mass distribution of ownership. As explained earlier, the jump transition component leads to substantial steady-state ownership by middle-valuation types. Because arrivals of high-valuation bids are rare, transaction prices are positively skewed. Relatedly, the arrival rate of middle-type investors is high in both regimes and this prevents types with above-middle cap rates from owning a meaningful share of the property market.

Owners' types and private valuations change dramatically in moving from expansion to contraction. The jump-transition results in a doubling of middle-type steady-state ownership, from 41% to 86%, and serves to "level the field" among investors. A typical owner's valuation likewise shifts from an expansion average cap rate of 5.77% to a contraction average of 6.61%. The only reason that the shift is not greater is the relatively short durations of contractions, averaging roughly 2.5 years. Were an owner forever "stuck" in a contraction as a middle type, (s)he would have a private valuation cap rate of roughly $4 \times (\eta_{mid,cnt} - 1) \approx 8.4\%$. With the substantial resetting of types during a contraction, one might expect a high turnover of properties. Instead, as documented in Table 6, the turnover *declines* by nearly a third. This is because the shifting of owner types to the middle during a contraction is accompanied by a factor of four reduction in the arrival rate of offers acceptable to middle-types.

5 Model Implications: Illiquidity and Transaction Risk

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 per-period 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, there is no single "fundamental" price. Instead, the price is probabilistically distributed in a manner that depends on investors' and owners' private valuations. This distribution drives transaction risk in the model and the calibration to the repeat sales data allows one to quantify it. In this subsection I characterize model transaction risk in several ways: The likelihood of selling within a given period, the time required to sell, and the price of immediacy. Finally, I turn to the related topic of distressed sales.

5.1 Transaction risk

For each macro state, Figure 6(a) plots a hypothetical probability distribution that a non-distressed owner will receive a satisfactory bid within a quarter, given their reservation cap rate gross of costs. Although I do not have bidding data to help validate the predicted distributions, the plots still serve to illustrate the power of the model to link to quantities that are, in principle, observable. To produce the plots using the calibrated model parameters, a discrete distribution is first calculated using the equilibrium solution to the $Q_{a,s}$'s and the probability $\pi_a^U(s)$ of bid arrivals. This is then extrapolated to a continuous log-Johnson distribution, as with Figures 4 and 5.[42]

For comparison, Figure 6(a) also plots the transaction probability for a perfectly liquid asset whose annualized income to price ratio coincides with the market average of adjusted *non-distressed* acquisition cap rates (5.95% in an expansion and 6.54% in a contraction). For a perfectly liquid asset, an above-market offer to sell would be instantly transacted with probability one, while an offer below-market would never be transacted. Transaction risk appears to be substantial and skewed towards lower cap rates (or higher prices). In an expansion, the probability of transacting at a "fair market" cap rate or better is in the neighborhood of 10% per quarter. The corresponding probability in a contraction is lower and corresponds to a more than a 10% reduction in price. The key message is that owners with a high cost of capital must be willing to accept below-market prices in order to transact with high probability. Another takeaway is that the dispersion in potential bids can be lower in a contraction, which in the model calibration happens because bidder discount rates are "squeezed" relative to the expansion state (see Figure 4).[43]

[42] In extrapolating to the continuum, I am assuming that a supporting transition process exists.

[43] The are too few transactions during contraction quarters to definitively confirm a decline in transaction dispersion. While the quarterly standard deviation of the Sale-Appraisal Ratio does decline during contractions, the difference is not statistically significant.

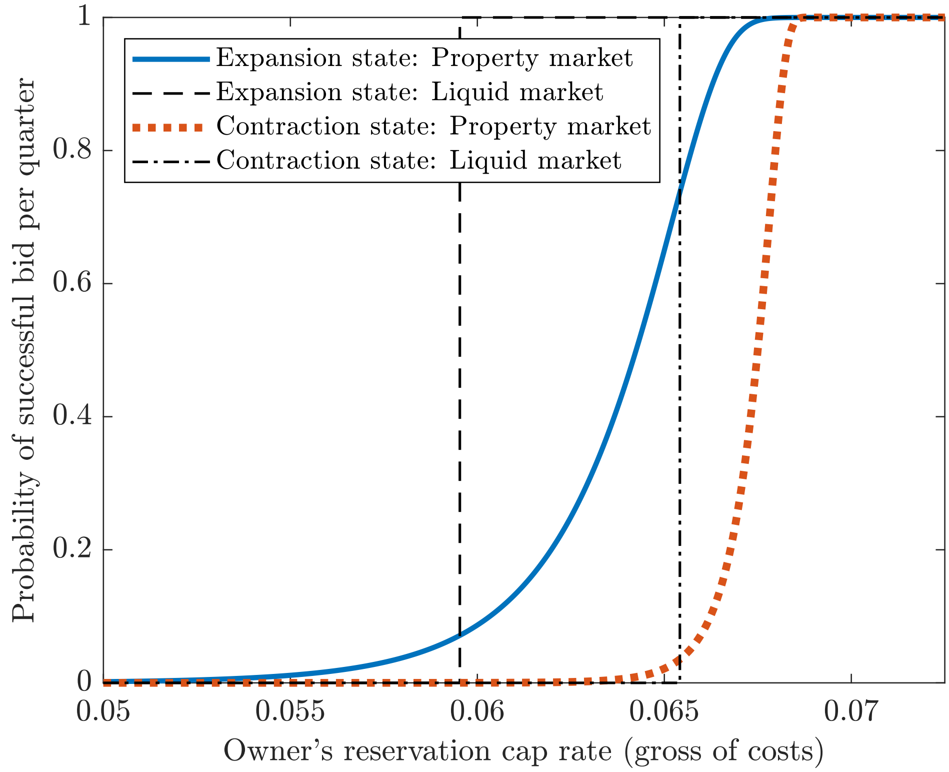

(a) Quarterly transaction risk in the expansion and contraction regimes

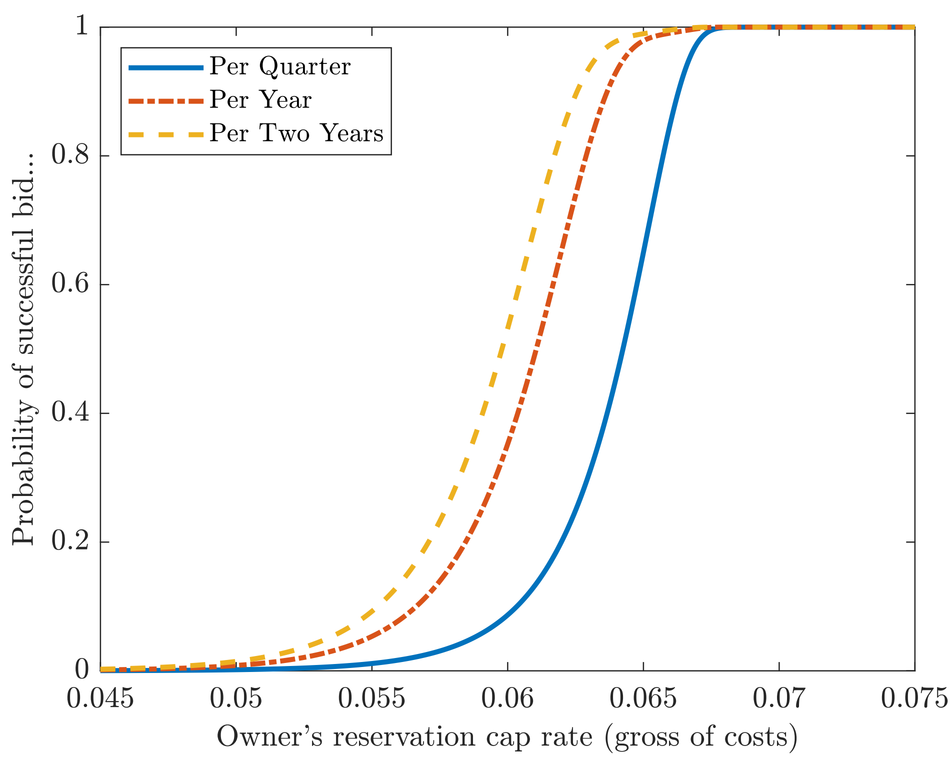

(b) Comparison of transaction risk in the expansion regime for various horizons

Fig. 6: Transaction Risk in Commercial Real Estate Properties. The top plot depicts the regime-dependent probability of transacting a non-distressed asset within one quarter as a function of the seller's reservation income-to-price ratio (i.e., "cap rate"). Each step function corresponds to a perfectly liquid asset with the same expected transaction income-to-price ratio as its non-distressed illiquid market counterpart. The bottom plot depicts the transaction probability over varying horizons, assuming the property market is currently in an expansion state.

Figure 6(b) illustrates how transaction risk changes with selling horizon, assuming that the current macro state is expansionary. As the selling horizon increases from one quarter to two years, the probability of transacting at a "fair" market cap rate of 5.95% increases. A longer selling horizon effectively deepens the market by allowing a greater number of offers to be entertained, thereby substituting for an instantaneous auction mechanism.

5.2 Time to sell.

Consider an owner with reservation cap rate c_r in state s_t. Figure 6(a) depicts $p(c_r, s_t)$, the probability that this owner will sell in the current quarter. From this, one can calculate the expected and median times to sell a property, conditional on a reservation cap rate. For an owner with a fixed reservation cap rate, the arrival of a satisfactory bid is a serially independent event. If the state of the economy does not change, then the expected and median times to sell (in years) are $\frac{1}{4p(c_r)}$ and $\frac{\ln .5}{4\ln(1-p(c_r))}$, respectively. Related, though more complicated, expressions apply when $p(c_r, s_t)$ is state dependent.

Figures 7(a) and (b) plot the mean and median times to sell given the seller's reservation cap rate in the two states. All of the curves are calculated from the continuous per-period non-distressed transaction likelihoods, $p(c_r, s_t)$, depicted in Figure 6(a), and account for potential macro state transitions. Also plotted, and corresponding to the right vertical axis, is the expected transaction cap rate for a seller, given their reservation cap rate. Note that the expected transaction cap rates are necessarily smaller than their corresponding reservation cap rates. Thus patient sellers can trade immediacy in favor of higher expected transaction prices.

The shaded region identifies reservation cap rates that result in non-distressed expected transactions (right vertical axis) between the model's unconditional transaction mean (6.03%) and the transaction mean conditional on the regime (5.95% in an expansion and 6.54% in a contraction). Achieving a target average transaction price in this range during an expansion amounts to setting the reservation price roughly 4% below the target. This potential discount is offset by the high premia that some buyers might be willing to pay. In other words, during an expansion a CRE portfolio manager targeting an average performance level can be fairly tolerant of dispositions by as much as 4% below market because (s)he knows that some properties will transact above market. Such tolerance roughly corresponds to

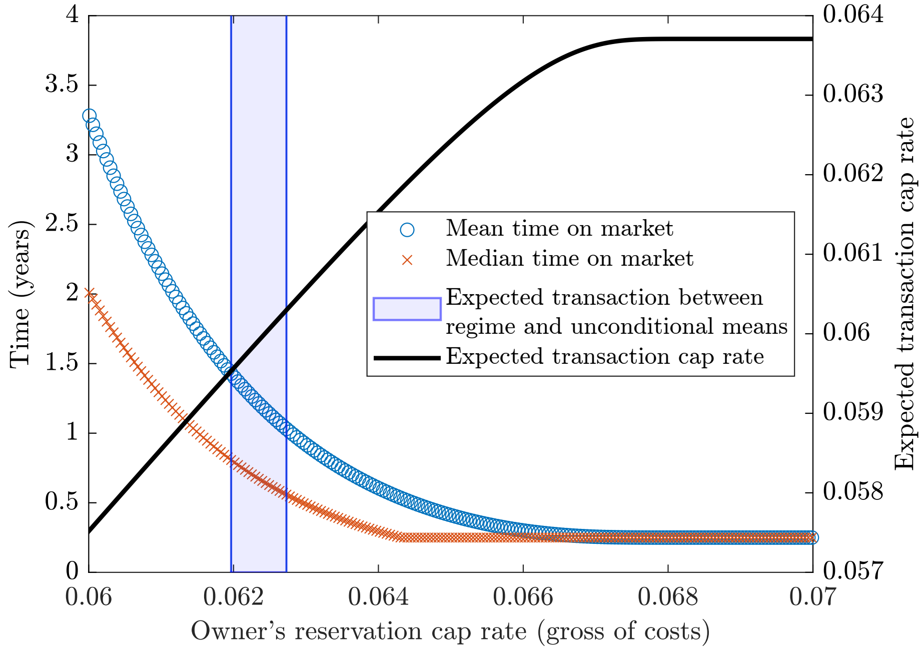

(a) Expansion state

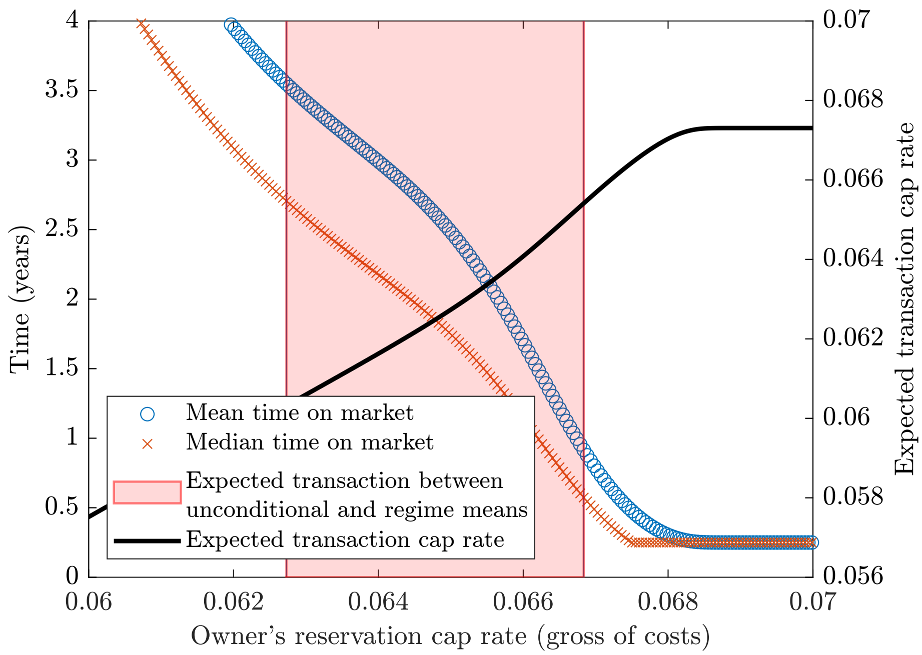

(b) Contraction state

Fig. 7: The blue circles and red ×'s respectively trace the mean and median times to sell (left vertical axis) given a non-distressed seller's reservation cap rate. The thick black line plots the corresponding expected transaction cap rate for non-distressed properties (right vertical axis). The shaded region denotes reservation cap rates that are expected to transact between the unconditional transaction cap rate mean and the mean conditional on the current state (expansion/contraction) of the non-distressed property market.

an expected (median) time on market of one year (six months). Greater patience is needed during a downturn to obtain the same performance as during an expansion: An expected transaction at the unconditional mean requires setting an aggressive reservation cap rate of 6.27%, associated with an expected time on market of over three years. On the other hand, sellers willing to part with their property a couple percentage points below prevailing contraction prices can expect to sell within a year.[44]

Although the NCREIF data does not report time to sell, such data is occasionally available elsewhere. For example, CoStar reports average "days on market" 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. Because about a third of listings are withdrawn, the CoStar estimates are significantly biased down. In addition, it is not clear how one might translate from the listing data to reservation cap rates. Still, given that the overall magnitudes reported by CoStar are roughly consistent with the model-implied time on market, one might anticipate that the model could be extended to fit detailed data on the listing and bidding process.

5.3 The 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 that the seller expects to experience. I calculate this quantity as follows. Let $Q^{(0)} \equiv 0$ and recursively define for $n \geq 1$,

$$Q_s^{(n)} = E\big[\{\tilde{Q}, Q^{(n-1)}\}^+\big] = \sum_{\substack{a \in \mathcal{A} \\ s' \in \mathcal{S}}} (\Pi_{\mathcal{S}})_{ss'} \pi_a^U(s') \{Q_a, Q_{s'}^{(n-1)}\}^+.$$

Then $Q_s^{(1)}$ is just the expected price to income ratio in macro state s if a seller must liquidate in one period. One can interpret $Q_s^{(2)}$ 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. A similar interpretation applies for $n > 2$.

[44]The hump shape in Figure 7(b) reflects that offers below a 6.5% cap rate are unlikely to arrive until after the contraction ends. Time on market steeply declines above 6.5% because such offers are likely to arrive during a contraction.

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. Also ignored in this calculation are distressed sales and transaction costs.

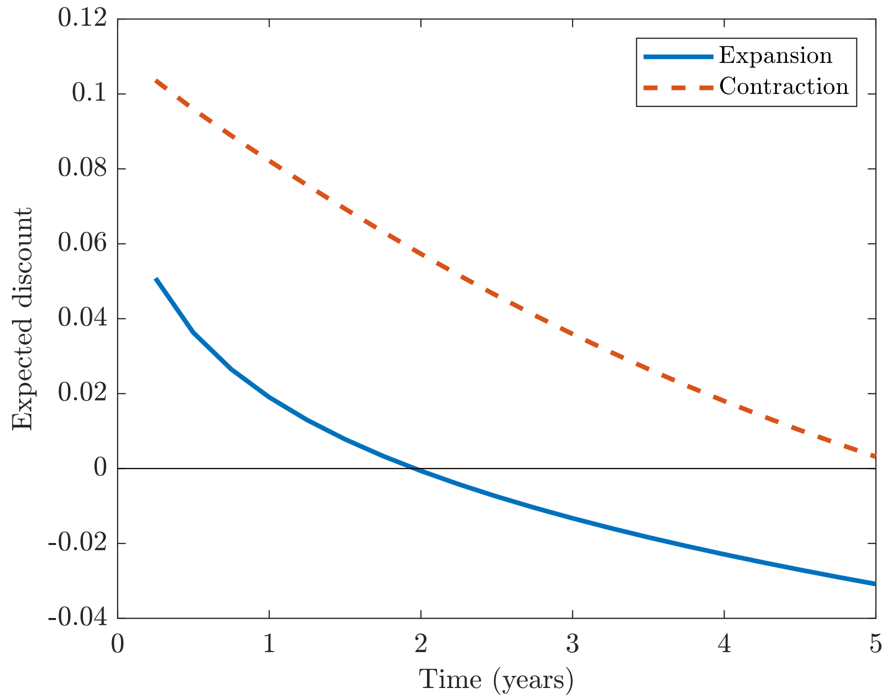

Fig. 8: An owner is said to have a need for immediacy of y years if (s)he must sell the property within y years. The graph depicts the discount relative to prevailing transaction prices that the seller expects to experience if using the disposition strategy described in Section 5.3. The possibility of a distressed sale is not incorporated into the discount in the plot.

The price of immediacy is defined as $1 - \frac{Q_s^{(n)}}{Q^*}$, where Q^* is the unconditional mean transaction price in the model. Figure 8 plots this for the continuous transaction probability distribution function, $p(c_r, s_t)$, in Figure 6. A seller can expect to experience a discount to average transactions if (s)he must liquidate within two (five) years during a market expansion (contraction). 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 one to two year option to wait out the market (Marra, Sagi, and Shaw, 2018). Such options ostensibly protect both the general partner and more patient limited partners from pressured sales by allowing for an "orderly liquidation". The plots in Figure 8 indicate that these options are valuable for the portfolio of remaining properties in the last year of a closed-end fund's legal term. During an expansion, exercising an extension option in the fund's penultimate year can eliminate the expected discount from pressured liquidation. Doing the same in a downturn nearly halves the expected discount.

5.4 Distressed sales

The analysis in the previous subsection does not incorporate distressed sales or what might be interpreted as an "extreme" need for immediacy. In the empirical Sale-Appraisal Ratio of Figure 3, roughly 3% of sales transact at more than a 20% discount to appraised value. In the calibrated model, this left tail is captured by allowing properties to enter into distress. It is instructive to explore the data for further insights into the sources of property distress.

Out of 75 distinct funds reporting to the NCREIF in the sample used for calibration, only 10 reported a total of 46 sales that fell below 80% of appraisal value for properties that were held for at least two years. Two thirds of these sales correspond to three distinct institutions, and ninety percent to institutions that sold at least two properties at a steep discount. Only 11 of the "distressed" sales corresponded to properties levered beyond 65%, and only one institution's distressed sales were all highly levered. The fact that leverage does not appear to drive extreme discounts is consistent with Genesove and Mayer (1997). The clustering of distressed sales across owners is accompanied by clustering in time, evidenced by the fact that 25 of the 46 sales involved more than one property in a single quarter by a single institution. The significant degree of clustering is consistent with heterogeneity in liquidity needs across institutions and time, which is the main mechanism in the model. All but one of the sales were by a private equity open-end fund, suggesting that redemption pressures forced liquidation. The latter conjecture is further supported by the fact that half of the sales took place either during the Great Financial Crisis or in the period 2015Q3-2016Q3 when private real estate open-end funds experienced heavy redemption pressures.[45]

All of the funds affected by distressed liquidation and discussed above were ongoing as of 2017Q2, and in possession of substantial holdings (median of $13B). Thus, among NCREIF open-end funds,

[45] See `https://goo.gl/21WpQd`. Some private equity real estate open end funds are required to begin forced liquidation once the redemption queue exceeds a threshold.

episodes of distressed asset liquidation are temporary and restricted to a few assets, possibly to meet rigid contractual requirements. This motivates modeling an episode of distressed liquidation as an asset-level shock rather than a type-transition. That is because the latter necessitates the liquidation of *all* assets that do not produce sufficient near-term cash flows, rather than enough to meet some threshold.

It is worth noting that CRE sales of "real estate owned" (REO, i.e., foreclosed and bank-owned) properties are associated with an *average* discount of 34% (see Table 3 in Chu, 2015). In terms of the magnitude of the discount, REO distressed sale data are roughly consistent with observed distressed sales in the NCREIF data and in the model. This might suggest a similar disposition mechanism (e.g., an auction with little scope for due diligence). A key difference, however, likely arises in the reason behind such sales. While property management may fall outside the core competency of a bank, making its management and maintenance of the asset costly, the same is not true for NCREIF member funds. Indeed, there is only weak evidence that NCREIF properties are sold at steep discounts because they can no longer be profitably operated by their owners.[46] As mentioned earlier, liquidation pressures on private equity funds (whether closed or open) may come from transient events that trigger fund contract covenants. By contrast, REO properties create a persistent significant drag on banks' deployable capital because of their associated reserve requirements (Ramcharan, 2019). Thus, in extending the model to incorporate lenders, it might be appropriate to model a bank as an investor with a very high cost of capital who only acquires properties through foreclosures.

6 Further Discussion

Here I discuss the variety of ways in which the model could be further applied and extended.

Source of valuation heterogeneity and policy implications. The model is currently silent on the sources of heterogeneity in growth-adjusted discount rates (i.e., the distribution of bids). These could arise from differences in cost of capital, beliefs, or skill. Indeed, different interpretations have markedly different policy implications.

[46]The discount is weakly (but significantly) related to smaller than average properties, a below average rise in market value since purchase, poor prior year's returns, and a weak market. The adjusted R^2 from a regression of positive discounts on these variables is about 4.5%.

In the calibrated model the average steady state transaction price is greater than the highest private *perpetual-hold* valuation. This is because a high valuation investor will eventually transition into a lower valuation investor and being able to sell the asset in that event improves the asset's current private value. Under a heterogeneous beliefs story, this is interpreted as a bubble.[47] Assuming asset bubbles are not desirable, the policy cure would be to introduce trading limitations. Under an institutional friction explanation, the ability to trade the asset can be interpreted as a welfare enhancing function of markets and policy might strive to enhance rather than restrict trade and liquidity.

In Appendix D.1 I argue that in the NCREIF context, the institutional friction story may be more compelling. That, however, may not be true in a broader context or with different data, and modeling the sources of heterogeneity may shed additional light on important policy questions. Such an exercise may well require endogenizing entry in the model (esp. of intermediaries) in the presence of time-varying capital constraints or other instruments of policy.

Improving realism and parameter identification. Detailed bidding data would help inform estimation and refinements of the model. For instance, one could incorporate the possibility of multiple bids as discussed in Section 3.1 or recast the model in continuous time. One could also hope to have direct measures of investor type transitions or the determinants of bargaining power (e.g., competition). The current calibration exhibits "nearby" parameter sets that also fit the data well but have significantly different implications for transaction risk and the distribution of bids. Bidding data could help better identify the deep model parameters and, importantly, improve the quantification of transaction risk discussed in Section 5.

Repeat sales index methodology. Francke (2018) finds that applying the repeat sales methodology (e.g., Calhoun, 1996) to distinct data subsamples, categorized by holding period, leads to substantial differences across the resulting indices: An index constructed using properties held over short periods exhibits greater average growth than an index constructed from properties held over a longer period. His findings raise the concern that changes in the index could result from a different mix of holding

[47] One could recast my model as a variant of Scheinkman and Xiong (2003) that includes "intermediate belief" types as well as random matching and bargaining.

period returns rather than changes to a fundamental factor. This issue is likely to also impact Ang, Chen, Goetzmann, and Phalippou (2018) who apply a sophisticated Bayesian version of the repeat sales index methodology to filter aggregate influences from sparse data.

The model of Section 3 exhibits a selection bias that leads to a similar pattern as that identified by Francke (2018). Thus the model also holds the promise of enabling econometricians to undo the selection bias and infer a more representative repeat sales market index. To do this in principle, one can incorporate the factor structure employed in Ang, Chen, Goetzmann, and Phalippou (2018) into μ_{t+j} in Eq. (6) and augment their Bayesian approach to structurally estimate an equilibrium model of transactions (such as the one offered here).

Additional Applications. It is clear that transaction risk carries important implications for mortgage and real estate derivative pricing. A model of transaction risk may similarly be useful in studying optimal contracting for skilled managers, including compensation and liquidation restrictions. Relatedly, it may be used to identify or develop performance metrics on which such contracts are based.

Although the NCREIF dataset lacks sufficient depth to finely resolve transaction risk by geography and asset type (see Appendix D.2), a more comprehensive dataset could be used to study cross-sectional attributes and dynamics (e.g., liquidity spillovers).

Finally, a model of market transaction risk could be employed to quantify banks' effective cost of capital for high reserve requirement assets (Ramcharan, 2019). This, accompanied with loan data and a model of mortgage pricing (that incorporates transaction risk for both borrower and lender) can potentially be used to estimate a lender's actual cost of capital.

7 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 standard asset pricing assumption of a random walk with drift in log-returns. This conclusion is robust to controlling for all cash flow events, as well as heterogeneity in random walk parameters across time and across properties.

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 transactional shocks that arise from the random matching and bargaining, and are nearly unrelated to the duration of the holding period (i.e., they are atemporal). If private valuations are persistent, observed short horizon holding periods in equilibrium will exhibit a positive "alpha" through selection bias: A short hold will only be observed when an owner is lucky enough to receive a bid significantly higher than the price (s)he recently paid for the property. This is consistent with the data and requires the significant presence of intermediate valuation investors. Both the model and the data imply that idiosyncratic property risk comes in two forms: a diffusion variance component similar to that exhibited by liquid assets measuring roughly 1% per year (annualized); and a purchasing/selling shock that, in a single optimal transaction, has a variance between 1% and 2%. Thus round-trip observed transaction returns exhibit a substantial "time-independent" variance component between 2% and 4% that arises from illiquidity rather than a missing variable (e.g., capital investment).

Some key takeaways are that (i) repeat transaction data, corresponding to the only available asset-specific return data for highly illiquid assets, inherently exhibit considerable selection bias, and (ii) transaction risk comprises a substantial portion of asset variance, and may be important to consider in any model of derivative (e.g., mortgage) pricing, portfolio choice, and delegated asset management. 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 whole loans, private equity deals, or other real assets.

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.

Albrecht, J., A. Anderson, E. Smith, and S. Vroman, 2007, "Opportunistic matching in the housing market," *International Economic Review*, 48(2), 641–664.

Albrecht, J., P. A. Gautier, and S. Vroman, 2016, "Directed search in the housing market," *Review of Economic Dynamics*, 19, 218–231.

Ang, A., B. Chen, W. N. Goetzmann, and L. Phalippou, 2018, "Estimating private equity returns from limited partner cash flows," *The Journal of Finance*, 73(4), 1751–1783.

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*.

Benzoni, L., and O. Chyruk, 2015, "The value and risk of human capital," *Annual Review of Financial Economics*, 7, 179–200.

Blackwell, D., 1965, "Discounted dynamic programming," *The Annals of Mathematical Statistics*, 36(1), 226–235.

Calhoun, C. A., 1996, "OFHEO house price indexes: HPI technical description," *Office of Federal Housing Enterprise Oversight*, pp. 1–15.

Campbell, J., A. Lo, and A. MacKinlay, 1997, *The econometrics of financial markets*. princeton University press.

Caplin, A., and J. Leahy, 2011, "Trading frictions and house price dynamics," *Journal of Money, Credit and Banking*, 43, 283–303.

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.

Chu, Y., 2015, "Asset fire sales by banks: evidence from commercial REO sales," *The Review of Corporate Finance Studies*, 5(1), 76–101.

Ciochetti, B., and J. Fisher, 2002, "The characteristics of commercial real estate holding period returns (IRRs)," *Real Estate Research Institute Working Paper*.

Diaz, A., and B. Jerez, 2013, "House Prices, Sales, and Time on the Market: A Search-Theoretic Framework," *International Economic Review*, 54(3), 837–872.

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), 269–303.

Francke, M., 2018, "The Role of Holding Periods in Repeat Sales Models," working paper, European Real Estate Society (ERES).

Garriga, C., and A. Hedlund, 2017, "Mortgage debt, consumption, and illiquid housing markets in the great recession," *FRB St. Louis Working Paper*, (2017-30).

Geltner, D., N. Miller, J. Clayton, and P. Eichholtz, 2013, *Commercial Real Estate Analysis and Investments*. OnCourse Learning, 3rd edn.

Genesove, D., and L. Han, 2012, "Search and matching in the housing market," *Journal of Urban Economics*, 72(1), 31–45.

Genesove, D., and C. J. Mayer, 1997, "Equity and time to sale in the real estate market," working paper 3.

Giacoletti, M., 2017, "Idiosyncratic risk in housing markets," .

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.

Han, L., and W. C. Strange, 2015, "The microstructure of housing markets: Search, bargaining, and brokerage," in *Handbook of regional and urban economics*. Elsevier, vol. 5, pp. 813–886.

Head, A., and H. Lloyd-Ellis, 2012, "Housing liquidity, mobility, and the labour market," *Review of Economic Studies*, 79(4), 1559–1589.

Head, A., H. Lloyd-Ellis, and H. Sun, 2014, "Search, liquidity, and the dynamics of house prices and construction," *American Economic Review*, 104(4), 1172–1210.

Hedlund, A., 2016a, "The cyclical dynamics of illiquid housing, debt, and foreclosures," *Quantitative Economics*, 7(1), 289–328.

———, 2016b, "Illiquidity and its discontents: Trading delays and foreclosures in the housing market," *Journal of Monetary Economics*, 83, 1–13.

Hugonnier, J., B. Lester, and P.-O. Weill, 2018, "Frictional Intermediation in Over-the-counter Markets," Working Paper 24956, National Bureau of Economic Research.

Johnson, N. L., 1949, "Systems of Frequency Curves Generated by Methods of Translation," *Biometrika*, 36(1/2), 149–176.

Krainer, J., 2001, "A theory of liquidity in residential real estate markets," *Journal of urban Economics*, 49(1), 32–53.

Krainer, J., and S. F. LeRoy, 2002, "Equilibrium valuation of illiquid assets," *Economic Theory*, 19(2), 223–242.

Lagos, R., and G. Rocheteau, 2009, "Liquidity in asset markets with search frictions," *Econometrica*, 77(2), 403–426.

Lopez-de Silanes, F., L. Phalippou, and O. Gottschalg, 2015, "Giants at the gate: Investment returns and diseconomies of scale in private equity," *Journal of Financial and Quantitative Analysis*, 50(3), 377–411.

Lovo, S., and C. Spaenjers, 2018, "A Model of Trading in the Art Market," *American Economic Review*, 108(3), 744–74.

Marra, H., J. Sagi, and I. Shaw, 2018, "A Glimpse into the Afterlife: An Investigation into Private Real Estate Fund Legal Duration, Extensions, and Liquidation," *Institutional Real Estate*, 30(9).

McQuade, T., A. Guren, et al., 2015, "How Do Foreclosures Exacerbate Housing Downturns?," in *2015 Meeting papers*, no. 40. Society for Economic Dynamics.

Ngai, L. R., and S. Tenreyro, 2014, "Hot and cold seasons in the housing market," *American Economic Review*, 104(12), 3991–4026.

Novy-Marx, R., 2009, "Hot and cold markets," *Real Estate Economics*, 37(1), 1–22.

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

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.

Ramcharan, R., 2019, "Banks' Balance Sheets and Liquidation Values: Evidence from Real Estate Collateral," *The Review of Financial Studies*.

Scheinkman, J. A., and W. Xiong, 2003, "Overconfidence and speculative bubbles," *Journal of political Economy*, 111(6), 1183–1220.

Wheaton, W. C., 1990, "Vacancy, search, and prices in a housing market matching model," *Journal of Political Economy*, 98(6), 1270–1292.

Williams, J. T., 1993, "Equilibrium and options on real assets," *The Review of Financial Studies*, 6(4), 825–850.

——— , 1995, "Pricing real assets with costly search," *The review of financial studies*, 8(1), 55–90.