Do Discount Rates Predict Returns?
Evidence from Private Commercial Real Estate

Liang Peng*

Smeal College of Business
The Pennsylvania State University
University Park, PA 16802
Phone: (814) 863 1046
Fax: (814) 865 6284
Email: liang.peng@psu.edu

December 2016

Abstract

This paper analyzes whether investors' discount rates predict ex post returns and Jensen's alpha in the private commercial real estate market. Tracking 4,430 properties in the U.S. that were worth 127 billion dollars at acquisition over the 1997 to 2014 period, I find that individual properties' acquisition cap rates, which measure discount rates, have significant predicting power for ex post returns and Jensen's alpha. The power is robust across property types and metro areas, is stronger in the short term, and persists when I control for sample selection, latent factors, heterogeneous factor loadings, and the pricing of some non-systematic risk.

Key words: Return predictability, market efficiency, and commercial real estate

JEL classification: G12, R33

* I thank John Cochrane, William Goetzmann, Benjamin Golez, Alberto Plazzi, Robert Shiller, and Eva Steiner for very constructive suggestions and comments. I thank RERI for a research grant and thank NCREIF for providing commercial real estate data. I thank participants of the 2016 RERI conference, a seminar at Lehigh University, the 2016 Real Estate Finance and Investment Symposium at Cambridge University, and the 2016 ReCapNet conference for valuable comments. All errors in this paper are my sole responsibility.

1. Introduction

In his AFA presidential address, Cochrane (2011) states that discount rates should and do predict investment returns: "Previously, we thought returns were unpredictable, with variation in price-dividend ratios due to variation in expected cash flows. Now it seems all price-dividend variation corresponds to discount-rate variation." Further, "high prices relative to current dividends entirely forecast low returns. This is the true meaning of return forecastability."

The main evidence supporting this statement is the predicting power of price-dividend ratios for stock returns (e.g. Campbell and Shiller (1988), Campbell and Ammer (1993), Cochrane (1992), Cochrane (1994), Ang and Bekaert (2007), among others). However, it is challenging to have clean tests on whether discount rates predict returns using stock market data (see, e.g. Stambaugh (1999), Goyal and Welch (2003), Goyal and Welch (2008), and others). This is largely because price-dividend ratios are a noisy measure for investors' discount rates (Campbell and Shiller (1988) and Fama and French (1988)). Price dividend ratios are affected by both discount rate news and cash flow news, and there is a debate on which news has greater impact on price-dividend ratios (See, e.g. Larrain and Yogo (2008), Chen (2009), Jules H. V. Binsbergen and Koijen (2010), Chen, Da and Zhao (2014)). The problem is particularly pronounced when discount rates and expected dividend growth are correlated (Menzly, Santos and Veronesi (2004) and Lettau and Nieuwerburgh (2008)). Therefore, it is of great importance to use more precise measures of discount rates to test return predictability. Golez (2014) substantiates this by showing that the price-dividend ratio corrected with a forward-looking measure of dividend growth extracted from S&P 500 futures and options has stronger predicting power for stock returns. Furthermore, the literature has been focusing on the predictability of returns, but the predictability of risk-adjusted returns seems also important. In an efficient capital market with rational investors, discount rates should predict returns, but returns should be just sufficient to compensate investors for the risk they take. Discount rates, therefore, should not predict risk-adjusted returns.

This paper analyzes whether the acquisition capitalization ratio (cap rate), which is the ratio of net operating income (NOI) of a property to its acquisition value, predicts individual properties' ex post annualized returns and Jensen's alpha in the private commercial real estate market. There are two advantages in studying private commercial real estate. First, cap rates may more accurately measure real estate investors' ex ante discount rates than price-dividend ratios do for stock investors (Plazzi, Torous and Valkanov (2010)). While firms' dividends are leveraged cash flows (Belo, Collin-Dufresne and Goldstein (2015)), properties' NOI is an unleveraged cash flow and tends to be more stable than dividends. Moreover, rents are not discretionary and is paid by tenants, while dividends are paid at the discretion of the firm's management and they are either actively smoothed, catering to particular clienteles, or the result of management's reaction to perceived mispricing (Shefrin and Statman (1984), Stein (1996), and Baker and Wurgler (2004)). Further, commercial real estate is often sold with existing leases with known rent arrangements, which leads to income growth rates that are much easier to predict than dividend growth rates. Therefore, compared with price-dividend ratios, cap rates are likely a more precise measure of discount rates, and the commercial real estate market may allow cleaner tests on whether discount rates predict returns.

The second advantage is that the private commercial real estate market is a large component of the economy (Plazzi, Torous and Valkanov (2012) and Peng (2016)). It differs from the stock market in many aspects, including transparency, liquidity, and transaction mechanisms. While there is well known research on the efficiency of other real estate markets, such as the housing market (Case and Shiller (1989)), there is little evidence on efficiency of the commercial real estate market. Therefore, analyzing return predictability in this market would help improve our understanding of this important component of the economy, and also provide new insights on whether return predictability is a universal phenomenon across asset classes.

Due to difficulty in obtaining high quality data, the literature on return predictability of private commercial real estate is almost nonexistent. A notable exception is Plazzi, Torous and Valkanov (2010). They analyze metro-level average cap rates of traded properties and find that, among other original results, cap rates predict returns for some property types. However, their data have limitations. First, the average properties are not investable, so their return predictability does not appear to be directly relevant for investors. Second, different properties are traded in different quarters, so the composition of traded properties changes over time. This makes it tricky to interpret their results (Case and Shiller (1989)),[1] particularly as traded properties are likely selected samples (Goetzmann and Peng (2006)). Further, they only study predictability of returns, not that of risk-adjusted returns, which is an important component of this paper.

This paper leverages a novel proprietary dataset of 33,338 private commercial properties in the U.S., which was worth about 950 billion dollars at acquisition.[2] This dataset is the *universe*, not a sample, of properties invested by members of the National Council of Real Estate Investment Fiduciaries (NCREIF) from the third quarter of 1977 to the fourth quarter of 2014. It is the largest dataset covering the longest sample period that any academic research has ever used to study private commercial real estate. A data set of so many properties makes for an authoritative sample for testing return predictability of commercial real estate.

The dataset provides quarterly reports of detailed financial and operational information at the property level, which allows me to calculate both acquisition cap rates and ex poste investment returns for individual properties. Cleaning the data leads to a sample of 4,430 properties that were worth about 127 billion dollars at acquisition, for which I can calculate both cap rates and ex post returns. The sample consists of 2,706 properties that were sold, which have varying holding periods and actual returns, and 1,727 properties that were still held at the end of the sample period, for which I estimate the five-year returns since acquisition using appraised values.

[1] The large literature on real estate price indices is built on the notion that average/median variables of traded properties do not track the same properties across time, so they may lead to biased results.

[2] The total value is estimated from properties with observed acquisition values. I multiply the average acquisition value of such properties with the total number of properties to estimate the total value at acquisition.

I find strong evidence for return predictability in the private commercial real estate market. Specifically, individual properties' acquisition cap rates have strong predicting power for properties' ex post returns, which are measured with annualized modified internal rate of return (MIRR) in both total returns and capital appreciation rates. The predicting power is significant both statistically and economically: an increase of 100 basis points in the cap rate would increase the ex post annualized total return MIRR by about 106 basis points. Using placebo tests, I show that such predicting power is not likely driven by unknown mechanical relationships between cap rates and ex post returns, because random cap rates have no such predicting power. Further, the predicting power is not likely subject to sample selection bias, which might happen if investors choose to sell properties when their realized returns equal acquisition cap rates, because the power remains strong for properties there were never sold. I also find that the predicting power is robust across the four main property types: apartment, industrial, office, and retail, and across metro areas with different market thinness.

I also find that cap rates have stronger predicting power in the short term than in the long term. While this seemingly contrasts with the conventional wisdom that price-dividend ratios predict stock returns more accurately in the long term than in the short term, it is consistent with the notion that cap rates more accurately measure discount rates in the short term as NOI growth is easier to predict in the short term (e.g. due to existing leases), and price-dividend ratios more accurately measure discount rates in the long term as dividend growth may be less volatile and easier to forecast in the long term (Belo, Collin-Dufresne and Goldstein (2015)).

I further investigate whether cap rates predict Jensen's alpha. The main challenge is that I only observe the annualized return during the entire holding period for each property, not its returns in each period. I use a holding-period log-linear factor model to overcome this problem. Such a model is first adopted by Cochrane (2005) to estimate the beta of venture capital investments, and also used by Korteweg and Sorensen (2010), Driessen, Lin and Phalippou (2012), Franzoni, Nowak and Phalippou (2012), and Peng (2016) for the estimation of factor loadings for private equity and commercial real estate. This model essentially regresses properties' holding-period aggregate risk premium against aggregate factors during the same periods. Results from estimating a standard four-factor (Fama and French (1993) factors and the Pastor and Stambaugh (2003) liquidity factor) holding-period model indicate that cap rates have strong predicting power for properties' alpha but not their factor loadings. Specifically, an increase of 100 basis points in cap rates would increase the quarterly alpha by about 21 basis points.

It is important to note that standard factor models may produce biased results for a variety of reasons. First, should there be unknown factors correlated with cap rates, omitting those factors may produce a spurious relationship between cap rates and alpha. To mitigate this problem, I conduct the tests in a latent factor holding-period model, which uses period dummies to capture the average impact of all known or unknown factors. Results from estimating such a model still suggest that cap rates predict alpha. Further, the results are robust across property types, metro areas with different market thinness, and properties with short or long holding periods. I also find that the predicting power of cap rates for alpha is stronger for properties with shorter holding periods.

Second, should properties' loadings of latent factors be correlated with their cap rates, the latent factor model discussed above may still provide biased results. To mitigate this problem, I define and estimate a real estate factor, which captures the common component of properties' returns that are not explained by the four factors, and validate it using out of sample tests conducted with Monte Carlo simulations. I then estimate a modified latent-factor model that allows properties' loadings on the real estate factor to be correlated with their cap rates. Such a model provides very robust evidence that cap rates predict alpha.

Third, if properties' non-systematic risk were priced in cap rates, factor models that omit such risk may provide biased results. I consider three types of non-systematic risk and analyze whether they affect the predicting power of cap rates for alpha. The first is the idiosyncratic component of each property's holding-period return. The second is temporal volatility of the real estate factor returns during each property's holding period.

The third is time varying market thinness. I then augment the latent factor model with measures of such non-systematic risk. Results suggest that (1) all three types of non-systematic risk are priced in properties' risk premium, and (2) cap rates still provide strong predicting power for alpha, for the whole sample and across property types.

The predictability of alpha does not necessarily suggest arbitrage opportunities or contradicts the efficiency of the private commercial real estate market. Since non-systematic risk is priced and my three measures of such risk may not be perfect, I am unable to rule out the possibility that the predictability of alpha might be spurious. Specifically, if investors price non-systematic risk in cap rates and they receive higher returns for taking higher non-systematic risk, a model that does not perfectly control for non-systematic risk may capture the higher returns with positive alpha. This will lead to a spurious relationship between cap rates and alpha. More theoretical guidance on how to measure non-systematic risk properly for commercial real estate and better datasets that allow such measures seem crucial for further analysis on this issue in future research.

This paper makes a few novel contributions to the literature. It is the first to show that individual properties' acquisition cap rates predict both ex post returns and risk adjusted returns in the private commercial real estate market. It also finds that the predictability is stronger in the short term than in the long term. These results are important because commercial real estate likely allows cleaner tests of return predictability, as cap rates may measure discount rates more accurately. Furthermore, commercial real estate is an important part of the economy itself. The results seem credible as the dataset of a large number of properties used in this paper makes for an authoritative sample for testing return predictability of commercial real estate.

The rest of this paper is organized as follows. Next section describes the data. The third section investigates whether cap rates predict ex post investment returns. The fourth section tests whether cap rates predict Jensen's alpha using holding-period factor models. The last section concludes.

2. Data

This paper uses the proprietary dataset of the National Council of Real Estate Investment Fiduciaries (NCREIF). NCREIF is a not-for-profit real estate industry association, which collects, processes, and disseminates information on the operation and transactions of commercial real estate. Its members are typically large investment companies, pension funds, and life insurance companies.[3] This paper uses the 2014:Q4 release of the database, which consists of 33,338 properties owned or managed by NCREIF members in a fiduciary setting during the period from the third quarter of 1977 to the fourth quarter of 2014. The database contains information on property attributes, such as property type, street address, square footage, etc., as well as quarterly property level operational and transactional information, including net operating income (NOI), capital expenditures, acquisition cost (if applicable), net proceeds from selling the property (if applicable), appraised values, etc. All cash flow variables are on an unlevered basis. Peng (2016) uses an earlier release of the property level NCREIF database (2012:Q3) to analyze the loadings of commercial real estate on standard stock market factors. Plazzi, Torous and Valkanov (2012) use a subsample (1998 to 2012) of the data and find that property attributes, including cap rates, seem to contain information that helps predict future returns, as the attributes help improve the performance of real estate portfolios.

I calculate the acquisition cap rate for each property whenever the data allow. The cap rate of property i acquired at the end of quarter t, denoted by $C_{i,t}$, is defined as

$$C_{i,t} = \frac{\sum_{s=t+1}^{t+4} NOI_{i,s}}{P_{i,t}}, \tag{1}$$

where $P_{i,t}$ is the acquisition price and $NOI_{i,s}$ is the quarterly net operating income. I am able to calculate acquisition cap rates for 15,617 properties but not for others due to missing information on either the acquisition price or net operating income.

[3] Examples of NCREIF members are Blackrock, Citi group, TIAA, New York Life, Invesco, Heitman/JMB, and Cornerstone real estate advisers.

For each property that had been sold by 2014:Q4, I calculate the annualized total return Modified Internal Rate of Return (MIRR) during its entire holding period whenever the database allows. I calculate MIRRs instead of IRRs because IRRs are often not well defined for commercial real estate investments as the present value equations often have multiple solutions, mainly due to the long holding periods and irregular cash flows of real estate investments. I also calculate the annualized capital appreciation MIRR for each property whenever possible. Note that I also repeat all analyses throughout this paper with IRRs,[4] and the results are robust.

To calculate the MIRRs, I first construct the quarterly cash flow series for each property. In the acquisition quarter of a property, the cash flow is simply the acquisition cost.[5] In each of the subsequent quarters before disposition, the cash flow is the NOI minus capital expenditures for the calculation of total return MIRR, and minus capital expenditures only for the calculation of capital appreciation MIRR. If there is a partial sale in that quarter, I also add net proceeds from the partial sale to the cash flow. In the disposition quarter, the cash flow is net sale proceeds plus NOI and then minus capital expenditures for the calculation of total-return MIRR, and net sale proceeds minus capital expenditures for the calculation of capital appreciation MIRR.[6]

After constructing the quarterly cash flow series, I calculate a simple total return index for each type of properties and use the index's quarterly returns as both the financing rate and the reinvestment rate to calculate the MIRRs for the same type of properties. When constructing such indices, I first use market values (or appraised values if market values are not available) at the beginning and the end of each quarter and the net cash flow (NOI plus partial sale minus capital expenditures) for each quarter to calculate the quarterly total return for each property. The index's return in that quarter simply equals the equal-

[4] When there are multiple solutions for total return IRRs, I select the smallest one from all solutions that are higher than the capital appreciation IRRs.

[5] I assume that all acquisitions and dispositions take place at the end of quarters. For a small number of properties, the database shows positive net operating income in the recorded acquisition quarters, possibly because their acquisitions took place in the middle of those quarters. For these properties, I assume the acquisitions took place at the end of the previous quarters.

[6] For a small number of properties, the net operating income in the disposition quarter is 0. I then assume that the dispositions took place at the end of the previous quarters.

weighted average of properties' returns.[7] Finally, I use the quarterly cash flow series and the series of the financing and reinvestment rates to calculate the annualized holding period total return and capital appreciation MIRRs for each property.

If disposition decisions were related to investment performance, sold properties would be a selected sample and analyses based on them may produce biased results. To mitigate this problem and to increase the sample size, for properties that were not sold, I calculate annualized five-year holding period total return and capital appreciation MIRRs using appraised values five years after acquisition (minus a selling cost calculated from the average ratio of net sale proceeds to gross sale proceeds for sold properties) as the net sale proceeds. I call these estimated MIRRs. This paper analyzes the pooled sample as well as the actual and estimated MIRRs separately.

Table 1 counts properties according to their final disposition status, which are true sales (arm's length transactions), other sales (e.g. transfer of ownership to another member, split into multiple properties, consolidation into existing properties, returned to lender, property destroyed, etc.), and being held by investors at the end of the sample period (2014:Q4), and whether I am able to calculate acquisition cap rates as well as total return MIRRs for them. For 13,398 properties that had been sold in arm's length transactions, I am able to calculate both cap rates and total return MIRRs for 6,834 properties. For other properties, I am able to calculate cap rates and estimate MIRRs for 3,800 properties. The total number of properties with cap rates and total return MIRRs, either actual or estimated, is 10,634.

I further apply some filtering rules to clean the data. First, it appears that NCREIF members started to report capital expenditures in 1997:Q2. Therefore, I limit my sample to properties acquired on or after 1997:Q1, which are 24,055 properties. I then focus on the four main commercial property types: apartment, industrial, office, and retail, which consist of 21,598 properties, 7,643 of which have both cap rates and total return MIRRs.

[7] I also use size-weighted and appraised-value-weighted index, and results throughout this paper are robust.

I further clean the data by excluding extreme outliers of cap rates and MIRRs, which are likely due to data errors, and requiring properties to have highly correlated total return and capital appreciation MIRRs, which is to mitigate the impact of problematic NOI values on my analyses.[8] Specifically, each of the 7,643 properties will stay in the final sample if (1) its cap rate is between 1% and 15%; (2) its annualized total return MIRR and capital appreciation MIRR are between -10% and 40% and are highly correlated. I deem a property to have highly correlated MIRRs if its residual from a linear regression of capital appreciation MIRRs against total return MIRRs is within three standard deviations from the mean of all regression residuals. Note that results in this paper are robust to minor variation of the above filtering rules.

Table 2 reports basic statistics of the final sample of 4,433 properties. 2,706 of them have actual total return MIRRs and 1,727 have estimated MIRRs. The final sample consists of 1,134 apartment, 1,573 industrial, 1,056 office, and 670 retail properties, which are respectively located in 106, 95, 88, and 134 metro areas. The table also reports the quartiles, mean, and standard deviation of cap rates and total return MIRRs. Figures 1 and 2 plot the histograms of the cap rates and total return MIRRs respectively. Both distributions seem reasonable and do not appear to have extreme outliers.

3. Predicting ex post returns

3.1. Cap rates and ex post investment returns

I first analyze whether investors' discount rates, which are measured with acquisition cap rates, predict properties' ex post investment returns, which is measured with the annualized total return MIRRs, using the following cross-sectional regression.

$$R_i = \alpha + \beta C_i + \sum_{k=1}^{K} \rho_k D_{i,k} + \varepsilon_i \tag{2}$$

In equation (2), R_i is the annualized total return MIRR of property i during its holding period, C_i is the acquisition cap rate, $D_{i,k}$ for $k = 1, \cdots, K$ are K fixed effect dummies,

[8] Extremely large NOI will cause a property's total return IRR to differ significantly from its capital appreciation IRR.

and ε_i is the error term. The null hypothesis is $\beta = 0$. Rejecting it would indicate that discount rates help predict ex post returns.

There are three issues in estimating the model in (2). First, this paper focuses on the return predictability of individual properties. Note that possible heterogeneity in risk-return characteristics across property types, location, and different phases of economic cycles may lead to return predictability. For example, hypothetically, investors may correctly perceive that retail properties have higher risk than apartments; as a result, they use higher discount rates to value retail properties, which lead to higher acquisition cap rates, and receive higher ex post returns for retail than for apartment. A regression using a pooled sample of both retail and apartment properties would find a positive β, even if individual property returns within each type are not predictable. Similarly, heterogeneity in risk-return characteristics across metro areas or different phases of economic cycles (e.g. before and after the financial crisis in 2007) may also lead to similar return predictability. While such predictability is of interest itself, it is not the focus of this paper. Therefore, I use fixed effects to control for heterogeneity in risk and returns related to property types, metro areas, and economic conditions in acquisition periods.

Second, it is important to note that a possible sample selection problem may bias results. Specifically, if there is no return predictability at all but investors tend to sell properties that have realized ex post returns similar with their acquisition cap rates, estimating (2) using sold properties only would falsely suggest return predictability. I mitigate this issue by analyzing properties that were not sold. Note that 1,727 properties in the sample were not sold and thus are not selected according to investors' disposition behavior. Therefore, return predictability of such properties, if substantiated, would help establish that the results are not driven by selected samples.

Third, it is important to distinguish true predicting power of cap rates due to information from artificial relationships due to unknown mechanisms or outliers. I achieve this by conducting placebo tests. Specifically, for each specification of (2) that I estimate, I conduct 1,000 rounds of placebo tests. Each round is a regression that is otherwise identical to the true regression of (2). The only difference is that, for each property, instead of using its own cap rate, I randomly draw a cap rate from those of the entire sample (without replacement) and regress each property's MIRR against this random cap rate. Should the predicting power of cap rate be due to information about individual properties instead of unknown mechanisms or outliers, such placebo tests should find no predicting power of random cap rates. The 1,000 rounds of placebo tests also allow me to construct empirical distributions of the coefficients of random cap rates as well as regression summary statistics such as adjusted R2 and means of squared regression errors. Such statistics serve as benchmarks from uninformative regressions, to which I can compare the same statistics from true regressions that use properties' own cap rates.

Panel A of Table 3 reports results of four specifications of (2). The first does not include any fixed effects and uses all the 4,430 properties in the sample. The second, third, and fourth include fixed effects of property types, metro areas, and acquisition periods. The difference among these specifications is that the second uses all 4,430 properties; the third uses the 2,706 sold properties with actual MIRRs; and the fourth uses the 1,727 properties with estimated MIRRs that were never sold.

All four specifications provide very strong results that cap rates have significant predicting power for ex post investment returns. For example, the coefficient of the cap rate is 1.063 in the second specification, which means that if the cap rate increases by 100 basis points, the ex post annualized total return MIRR increases by about 106 basis points. Note that the strong result of return predictability is very robust when fixed effects are included and when I analyze both sold properties and those that were never sold. The fact that cap rates have significant predicting power for properties that were not sold suggests that sample selection does not seem to drive our results.

Panel A also reports the mean of squared regression errors (MSE) for each of the four specifications. To evaluate the magnitude of the MSEs, I report the percentage of placebo tests that have greater MSEs for each specification, which is always 100%. Also, while not reported in the table, the MSE of each specification is always 3 standard deviations away from and lower than the mean of placebo test MSEs.[9] This suggests that each property's cap rate has strong predicting power for its own ex post investment returns.

Panel B reports the mean and the standard deviation of the intercept term, the coefficient of a random cap rate, the adjusted R2, and the MSEs of the 1,000 rounds of placebo tests for each specification in Panel A. It is apparent that random cap rates have no predicting power for properties' ex post investment returns, which is strong evidence that the predicting power of cap rates is due to information on individual properties instead of unknown mechanisms or outliers.

3.2. Robustness checks

Note that the cap rate and the total return MIRR are correlated for a mechanical reason: net operating income in the year after acquisition is used to calculate both the acquisition cap rate and the total return MIRR. Therefore, it is worthwhile to investigate whether the return predictability found in Table 3 is solely driven by this mechanical relationship. To do so, I replicate all regressions and placebo tests in Table 3 but use the capital appreciation MIRR instead of the total return MIRR as the dependent variable, and then report results in Table 4. Since capital appreciation MIRR is not affected by net operating income, results in Table 4 is not affected by the mechanical relationship.

Table 4 presents results that are consistent with those in Table 3. First, Panel A of Table 4 shows that cap rates provide statistically significant predicting power for holding-period capital appreciation MIRRs in all four specifications. The predicting power is also economically significant. For example, in specification II, an increase of 100 basis points in the cap rate would increase the ex post annualized capital appreciation MIRR by about 65 basis points. Note that this effect is weaker than the impact of cap rates on the total return MIRRs. This is partly because capital appreciation is only one component of the total return. Second, Panel B of Table 4 shows that random cap rates have no predicting

[9] Both the mean and the standard deviation of the placebo test MSEs are calculated from the 1,000 rounds of placebo tests.

power for capital appreciation MIRRs. Comparing the mean squared errors of the true regressions in Panel A and those of placebo tests in Panel B continues to show that MSEs of true regressions are significantly lower that those of placebo tests. Overall, Table 4 suggests that the predicting power of cap rates for ex post returns is not solely driven by the mechanical correlation between cap rates and total returns.

I then analyze whether the prediction power of cap rates is robust across property types. Practitioners often conduct separate market analyses for different property types, assuming that they have different risk and return characteristics. This is corroborated by Peng (2016), which shows that different property types have different loadings on conventional asset pricing factors. Since the private commercial real estate market seems to be segmented by types, it is useful to investigate whether the predicting power of cap rates presents for all the four main property types.

I estimate the model in (2) for apartment, industrial, office, and retail properties separately, and report the results in Table 5. Since I run regressions for each property type separately, I no longer include the type fixe effects. Table 5 provides strong evidence that cap rates predict annualized total return MIRRs for all four types, and the predicting power is statistically significant. It is interesting to note that the magnitude of the predicting power seems to vary across types. An increase of 100 basis points in the cap rate would increase the total return MIRR by 153 basis points for apartment, 79 basis points for industrial, 105 basis points for office, and 149 basis points for retail properties. In unreported placebo tests, random cap rates have no predicting power for any types. I also conduct the same regressions in Table 5 for capital appreciation returns, and the results are robust.

I then investigate whether the predicting power varies across metro areas with different market thinness/liquidity. It is possible that information might be easier to gather and properties might be easier to value in thicker or more liquid markets, due to more transactions and more comparable properties (Hochberg and Mühlhofer (2015)); as a result, those markets might be more efficient and thus have higher return predictability. I investigate this using five different regressions and report results in Table 6.

I measure market thinness for the metro area for each property in the sample using the following continuous and discrete variables. The first is “total volume”, which is the number of properties in the same metro area of the property that have ever been held by NCREIF members during the entire sample period from 1977 to 2014. The second is “type volume”, which is the number of properties of the same type in the same metro area ever held by NCREIF members. For instance, for an office building in Washington, D.C., the “total volume” for this property is the number of all unique properties in Washington, D.C. that ever appeared in the NCREIF database, and the “type volume” is the number of all unique office properties in Washington, D.C. that ever appeared in the database. The third is a dummy variable that equals 1 for the four gateway markets: New York, San Francisco, Washington D.C., and Houston.[10] There are 554 properties located in these gateway markets. There are certainly alternative ways to define gateway markets, but our results do not depend on this measure of market thinness only. The fourth is a dummy called “total top 10”, which equals 1 if the metro area is among the top 10 areas with the highest “total volume”. 1,893 properties are located in the “total top 10” areas. The fifth is a dummy called “type top 10”, which equals 1 if the metro area is among the top 10 areas with the highest “type volume”. 2,438 properties are located in the “type top 10” metro areas.

The first regression in Table 6 augments the model in (2) with the interaction term between the cap rate and the “total volume”. The second regression augments the model with the interaction of the cap rate with the “type volume”. The third contains the interaction with the gateway dummies. The fourth and the fifth regressions interact the cap rate with the “total top 10” and “type top 10” dummies respectively. Note that, throughout this paper whenever I interact cap rates with other variables, I use the demeaned cap rates (cap rates minus the average of all cap rates in the sample) in the

[10] This definition of gateway markets is provided by NAIOP, the Commercial Real Estate Development Association.

interaction terms. This makes it easy to interpret the coefficients of the cap rate and the variables it interacts with.

Results in Table 6 are very robust. First, the predicting power of cap rates for ex post total returns remains strong both statistically and economically in the presence of all interaction terms. This suggests that return predictability is persistent across all markets regardless their market thinness. Second, the predicting power of cap rates do not seem to be related to market thinness, except that returns are slightly more predictable in the four gateway areas. In unreported regressions, I replicate all regressions using capital appreciation MIRRs and results are similar.

It is conventional wisdom that stock returns are more predictable in the long term (see, e.g. Cochrane (2011)),[11] as price-dividend ratios may measure investors' discount rates more accurately in the long time (Belo, Collin-Dufresne and Goldstein (2015). For commercial real estate, rents are easier to forecast in the short term due to existing leases so cap rates may more precisely measure investors' discount rates in the short term; therefore, the predicting power of cap rates for returns may be stronger in the short term.

To investigate this, I run two types of regressions using the 2,706 properties that were sold. The first augments the model in (2) by interacting a property's cap rate with the duration of its holding period. This interaction term allows me to capture possible linear or monotonic relationship between the predicting power of cap rates and the duration of the holding period. Should the predicting power increase (decrease) with the duration, the interaction term should have a positive (negative) coefficient.

The second type augments the model in (2) by interacting properties cap rates with dummies of short and long holding-period duration. Such interaction terms allow the predicting power to differ across properties with short, middle, and long duration, and thus allow me to capture non-linear relationship between the predicting power of cap

[11] There is also evidence for possibly stronger return predictability in the short term, see, e.g. Ang and Bekaert (2007).

rates and holding period duration. I define short duration as duration shorter than 16 quarters (four years), long duration as duration longer than 28 quarters (7 years). This is to simply split all sold properties into three roughly equal groups: 920 with short duration, 957 in the middle, and 829 with long duration.

Table 7 reports results of these two types of regressions, with Panel A for total return MIRRs and Panel B for capital appreciation MIRRs. In Panel A, the first regression contains the interaction term between the cap rate and duration. Two results are apparent for this regression. First, the predicting power of cap rates remains significant both statistically and economically in the presence of the interaction term. Second, the interaction term has a significant negative coefficient, which indicates that the predicting power declines with duration. Specifically, the coefficient is -0.036, which suggests that the predicting power of cap rates decreases by 3 basis points for each extra quarter. The second regression augments the model in (2) with the interaction between the cap rate and the dummy for short duration. The result indicates that, first, the predicting power of cap rates remains significant in the presence of the interaction term; and second, the predicting power is significantly stronger for properties with short duration. The third regression includes the interaction of the cap rate with the dummy for long duration, and the result shows that the predicting power of cap rates remains strong and the predicting power is weaker for long duration. The fourth regression includes both the interaction terms of the cap rate with the dummies for short and long duration. The results are consistent with that of earlier regressions. Overall, Table 7 provides consistent results that the predicting power of cap rates is stronger in the short term than in the long term.

As discussed earlier, net operating income in the year after the acquisition is used to calculate both the cap rate and the total return MIRR. When the duration is shorter, the net operating income is a larger portion of the cash flows during the holding period, and thus its impact on the total return MIRR is stronger. Therefore, it is worthwhile to investigate whether this mechanical relationship causes the stronger predicting power of cap rates in the short term. To so do, I re-run the regressions in Panel A but use capital appreciation MIRRs as dependent variables, which are not affected by net operating income. I report results in Panel B of Table 7. It is clear that results from all four regressions in Panel B are consistent with those in Panel A. This provides strong evidence that the stronger predicting power of cap rates in the short term is not due to the use of NOI in the calculation of both cap rates and total return MIRRs. In another unreported robustness check, I follow Golez (2014) and adjust cap rates with the average NOI growth rate in the 5-year period following acquisition so that cap rates are more comparable with discount rates. The predicting power of adjusted cap rates remains statistically significant.

4. Predicting Jensen's alpha

4.1. A holding-period return model

Return predictability may or may not be consistent with market efficiency, depending on whether high returns are compensating investors for taking high risk. This section analyzes whether individual properties' acquisition cap rates predict ex post Jensen's alpha. The market of private commercial real estate would seem efficient if cap rates predict ex post returns but not Jensen's alpha.

It is infeasible to directly estimate Jensen's alpha for individual properties as I only observe the holding period MIRR for each property. I overcome this challenge by using a holding-period return model that was first adopted by Cochrane (2005) in estimating the beta of venture capital investments, then by Korteweg and Sorensen (2010), Driessen, Lin and Phalippou (2012), and Franzoni, Nowak and Phalippou (2012) to estimate factor loadings for private equity, and also by Peng (2016) to estimate factor loadings of private commercial real estate.

The model is built on a single-period log-linear factor model. Consider a property i that was acquired in period buy_i and sold in period $sell_i$. I assume the unobserved single-period investment return for this property in period t, $R_{i,t}$ (a gross return), is generated from the following log-linear factor model,

$$\log\left(R_{i,t}\right)-\log\left(T_t\right)=\alpha_i+\sum_{k=1}^{K}\beta_k F_{k,t}+\upsilon_{i,t}\ , \tag{3}$$

where T_t is the risk-free interest rate (a gross return), α_i is Jensen's alpha, $F_{k,t}$ are k factors, and $\upsilon_{i,t}$ is an error term.

Apparently the dependent variable in (3), the single period return, is unobserved for non-traded assets such as commercial properties. To obtain observed dependent variable, I aggregate both sides of equation (3) across periods within the property's holding period, which leads to the following equation.

$$\begin{aligned}&\sum_{t=buy_i+1}^{sell_i}\log\left(R_{i,t}\right)-\sum_{t=buy_i+1}^{sell_i}\log\left(T_t\right)\\&=\alpha_i\left(sell_i-buy_i\right)+\sum_{k=1}^{K}\left(\beta_k\sum_{t=buy_i+1}^{sell_i}F_{k,t}\right)+\sum_{t=buy_i+1}^{sell_i}\upsilon_{i,t}\end{aligned} \tag{4}$$

I simplify the notation by defining the duration of the holding period, U_i, as

$$U_i=sell_i-buy_i. \tag{5}$$

I denote by R_i the total return (a gross return) of the property during its entire holding period, which can be calculated using the total return MIRR as follows.

$$\log\left(R_i\right)\triangleq\sum_{t=buy_i+1}^{sell_i}\log\left(R_{i,t}\right)=U_i\times\log\left(1+MIRR_i\right). \tag{6}$$

I further simplify the notation for the error term as follows.

$$\sum_{t=buy_i+1}^{sell_i}\upsilon_{i,s}=\varepsilon_i. \tag{7}$$

The model is now

$$\begin{aligned}&\log\left(R_i\right)-\sum_{s=buy_i+1}^{sell_i}\log\left(T_t\right)\\&=\alpha_iU_i+\sum_{k=1}^{K}\left(\beta_k\sum_{s=buy_i+1}^{sell_i}F_{k,t}\right)+\varepsilon_i.\end{aligned} \tag{8}$$

Apparently, the model in (8) does not allow the estimation of property-specific alpha or factor loadings. However, it allows me to test whether each property's alpha is correlated with its acquisition cap rate. Specifically, I parameterize the property's alpha as

$$\alpha_i=\alpha+\rho C_i, \tag{9}$$

and then test whether $\rho=0$. If cap rates predict alpha, ρ should be positive.

For the test to provide unbiased results, it is important to allow properties' factor loadings to be correlated with their cap rates. If each property's true loadings is correlated with its cap rate but the model is incorrectly specified and forces loadings to be identical across properties, the heterogeneous effects of factors on properties' holding-period risk premium would be picked up by the estimated alpha. Consequently, alpha may be estimated with bias. To allow loadings to be correlated with cap rates, I parameterize property i's loading for a factor F_k, denoted by $\beta_{k,i}$, as

$$\beta_{k,i} = \beta_k + \lambda_k C_i, \tag{10}$$

and allow λ_k to differ from 0.

I plug both (9) and (10) into (8) and have the following estimable holding-period model.

$$\begin{aligned} &\log\left(R_i\right) - \sum_{T=buy_i+1}^{sell_i} \log\left(T_t\right) \\ &= \alpha U_i + \rho C_i U_i + \sum_{k=1}^{K}\left(\beta_k \sum_{s=buy_i+1}^{sell_i} F_{k,t}\right) + \sum_{k=1}^{K}\left(\lambda_k C_i \sum_{s=buy_i+1}^{sell_i} F_{k,t}\right) + \varepsilon_i \end{aligned} \tag{11}$$

If ρ, the coefficient of the interaction term between the (demeaned) cap rate and the holding period duration, significantly differs from 0, I would conclude that cap rates help predict Jensen's alpha. Similarly, if λ_k, the coefficient of the interaction term between the (demeaned) cap rate and the kth factor, significantly differs from 0, I conclude that properties' cap rates predict their loadings on the kth factor.

There is an econometric detail in estimating (11). Should the error term in the single-period model (3) be i.i.d., the variance of the error term in (11) should increase with the duration of the holding period (Goetzmann (1992) and Korteweg and Sorensen (2010)). A standard three-stage approach (e.g. Case and Shiller (1989)) to address this is to estimate the model using OLS as the first stage, regress squared OLS residuals against the duration in the second stage, and then use the fitted values of squared residuals as weights to estimate the model again using weighted OLS in the third stage. However, throughout my analyses, I find that squared OLS residuals are not or slightly negatively related to duration, and the negative relationship, if exists, is not economically significant (coefficient is less than 0.001). Therefore, in all reported results, I estimate the model with OLS and calculate and report White's heteroscedasticity-consistent standard deviations.

4.2. Cap rates and alpha in a four-factor model

I first test whether cap rates predict alpha in a four-factor model, and report results in Table 8. The four factors are the Fama and French (1993) factors and the Pastor and Stambaugh (2003) liquidity factor. The first regression in Table 8 serves as a benchmark, and does not include any interaction terms between cap rates and duration or factors. The second regression includes the interaction terms and uses the whole sample of 4,430 properties. The third and the fourth are identical to the second but use the 2,706 properties that were sold and the 1,727 properties that were not sold respectively.

Table 8 shows that, first, the benchmark regression suggests that private real estate has an insignificant alpha, a small positive equity market beta (0.204), a small positive loading on SMB (0.186), an insignificant loading on HML, and a small positive loading on Liquidity (0.336). These results are generally consistent with the findings in Peng (2016).[12] It is also worth noting that the factor model fits the data reasonably well, as the adjust R2 is 0.401. Second, cap rates have strong predicting power for alpha. The coefficient of interaction term between the cap rate and duration (in quarters) is 0.173, 0.228, and 0.133 in the second, third, and fourth regressions respectively and statistically significant. The coefficients are also economically significant. For example, 0.173 means that an increase of 100 basis points in the cap rate would increase alpha by about 17.3 basis points per quarter, or 69.2 basis points per annum. Third, there is no robust evidence that properties' factor loadings are correlated with cap rates. While properties' loadings seem to be correlated with cap rates for the liquidity factor in the second regression, for SMB in the third regression, and for the stock market risk premium and HML in the fourth regression, these results are not robust across regressions. Finally, by

[12] Note that this paper focuses on testing whether cap rates predict alpha, not evaluating the level of alpha itself. Consequently, I do not discuss the average alphas throughout the remaining part of this paper.

comparing the adjusted R2 of the second regression to that of the first regression, it is clear that allowing properties' alphas to be correlated with their cap rates helps the model fit the data better, as the adjusted R2 increases from 0.421 to 0.476.

As a robustness check, Table 9 repeats the second regression in Table 8 for each of the four property types separately. The results are generally robust: cap rates have very strong predicting power for alpha for all types except office. Further, the results seem to suggest that the predicting power varies across types: the coefficient varies from being insignificant for office to 0.413 for apartment properties. This seems consistent with the conventional wisdom that different types of commercial real estate may have different risk characteristics, which, however, is not the focus of this paper. In another unreported robustness check, I adjust cap rates with the average NOI growth rate in the 5-year period following acquisition. The predicting power of adjusted cap rates for alpha remains statistically significant.

4.3. Cap rates and alpha in a latent factor model

It is debatable whether the four-factor model is correctly specified. In fact, it is *always* debatable whether *any* factor models are correctly specified. It is certainly possible that some factors are latent, and such factors may be correlated with cap rates. Consequently, models that omit these latent factors may estimate the coefficient of the interaction term between the cap rate and duration with biases. The biased estimates can be incorrectly interpreted as evidence for the predictability of alpha.

To mitigate possible biases due to latent factors, I run a latent factor version of the holding period model in (11). In this version, I use a period dummy to capture the impact of all factors, including latent ones, in each period. Specifically, the model is

$$\begin{aligned} &\log\left(R_i\right) - \sum_{T=buy_i+1}^{sell_i} \log\left(T_t\right) \\ &= \alpha U_i + \rho C_i U + \sum_{t=buy_i+1}^{sell_i} M_t + \varepsilon_i \end{aligned} \tag{12}$$

where the period dummy variable M_t is defined as

$$M_t = \sum_{k=all} \left(\beta_k F_{k,t} \right). \tag{13}$$

I estimate the latent factor model in (12) for all properties as well as each property type separately, and present results in Table 10. First and foremost, it is apparent that cap rates have strong predicting power for alpha. Second, the latent factor model seems to fit the data better than the four-factor model. The adjusted R2 increases from 0.476 to 0.512 for the whole sample, from 0.449 to 0.507 for apartment, from 0.548 to 0.611 for industrial, from 0.384 to 0.421 for office, and from 0.611 to 0.702 for retail properties. This seems to suggest that there might be factors missing in the four-factor model but captured by the latent factor model.

Note that the specification in (12) is correct only if properties' loadings for each factor are identical. To mitigate this concern, I also estimate a variation of (12) that further includes interactions between cap rates and the four factors. These interaction terms capture heterogeneous effects of factors across properties that are correlated with cap rates, and the period dummies would capture common impact of the four factors and latent factors. Results from estimating this variation, while not reported, are robust.

I also investigate whether the predictability of alpha is persistent across thick and thin markets, as well as across properties with long and short holding periods. I first estimate (12) for properties located in the top 10 metro areas with the highest "type volume", which I call the thick markets, and those located in other metros, which I call the thin markets. I then estimate (12) for properties with duration longer than or equal to the median of duration, which is 20 quarters, and for those with duration shorter than 20 quarters. Table 11 reports results from these four regressions, which are robust and consistent with those in Table 10: cap rates have significant predicting power for alpha in both thick and thin markets and for properties with long and short duration. It is worth noting that cap rates have stronger predicting power for alpha in the short term than in the long term, which is consistent with the predictability of total return MIRRs. In unreported regressions, I use other definitions of market thinness, such as "total volume", and split the sample into different groups with different duration, and the results are very robust. I also estimate the variation that includes interaction terms between cap rates and the four factors, and results are robust.

4.4. Cap rate and alpha with a real estate factor

The latent factor model in (12) is correctly specified if all properties have identical loadings for each of the latent factors. However, if properties' loadings of a latent factor are correlated with cap rates, the latent factor model may still provide biased results. To overcome this problem, I consider a model that includes a factor that captures the common component of returns not explained by these four factors, which I call the "real estate factor". I allow properties' loadings on the real estate factor to be correlated with their cap rates, and investigate whether cap rates predict alpha in this model. This model still includes period dummies to control for the common component of the impact of all factors on returns.

I define the real estate factor in each period as the common component of all properties' risk premium in the period that is not explained by the four factors. As a result, the real estate factor is orthogonal to the four factors and essentially the residual from estimating the four-factor model. However, estimating the four-factor model provides residuals for each property's entire holding period, not for each period. I then need to estimate the real estate factor in each period from properties' holding-period residuals. Note that this is the same classical econometric challenge economists need to overcome to construct real estate price indices.

I first obtain holding period residuals for each property i, S_i, from estimating a simple four-factor model as specified in (4) that assumes identical loadings across properties for each factor.[13] Following the classical repeat sales regression, I assume that the holding period residual S_i is the sum of residuals in each period,

$$S_i = \sum_{t=buy_i+1}^{sell} S_{i,t} \tag{14}$$

[13] Results are robust when I also include interaction terms between cap rates and the four factors.

and the residual in each period contains a common component, which is the real estate factor, and an error,

$$S_{i,t} = E_t + \varepsilon_{i,t} \,. \tag{15}$$

Combining (14) and (15) leads to the familiar repeat sales regression below.

$$S_i = \sum_{t=buy_i+1}^{sell_i} E_t + \sum_{t=buy_i+1}^{sell_i} \varepsilon_{i,t} \tag{16}$$

Two things are worth noting. First, (16) is essentially a simple dummy regression and the real estate factors are coefficients of dummies for each of the holding period for each property. Second, the Case and Shiller (1989) three-stage approach should be used if the variance of the error term in (16) indeed increases with the duration of each property's holding period. However, I find no relationship between the variance and the holding period duration; therefore, I estimate (16) with OLS. Table 12 presents the real estate factors from 1997 Q2 to 2014 Q4, and the factors are also plotted in Figure 3.

It is important to validate the real estate factor before using it, particularly on whether they indeed capture information or they are simply noise. To do so, I conduct out-of-sample tests as follows. I first randomly split all properties into two samples, say A and B. I then use return residuals of properties in sample A to estimate the real estate factor according to (16), and then test whether the real estate factor estimated from sample A, $\hat{E}_t^A$, helps explain return residuals of properties in sample B, S_i^B, in the following regression.

$$S_i^B = \rho \sum_{t=buy_i+1}^{sell_i} \hat{E}_t^A + \upsilon_i^B \tag{17}$$

A significant and positive ρ would suggest that the real estate factor contains information.

I conduct 1,000 rounds of the out-of-sample test, randomly splitting the sample each time. I then plot the histogram of ρ from the 1,000 rounds in Figure 4. It is apparent that ρ is positive and significantly different from 0, which is also confirmed by formal t-tests. In fact, ρ is positive in all the 1,000 rounds. This is strong evidence that the real estate factor contains information, not merely noise, on common components of properties risk premium that are not explained by the four factors.

I then estimate the following latent factor model, which allows properties' real estate factor loadings to be correlated with their cap rates.

$$\begin{aligned} &\log\left(R_i\right)-\sum_{T=buy_i+1}^{sell_i}\log\left(T_t\right) \\ &=\alpha U_i+\rho C_i U+\pi C_i\sum_{t=buy_i+1}^{sell_i}E_t+\sum_{t=buy_i+1}^{sell_i}M_t+\varepsilon_i \end{aligned} \tag{18}$$

The model in (18) allows me to investigate whether cap rates predict alpha in the presence of the real estate factor for which properties' loadings may be correlated with cap rates. I estimate (18) for the whole sample, as well as for each of the four property types, and report results in Table 13. It is apparent that cap rates still have strong predicting power for alpha for all properties and all types, with the coefficient varying between 0.159 for industrial to 0.302 for apartment properties. Further, there is no evidence that properties' loadings on the real estate factor are correlated with their cap rates. In unreported regressions, I also include interaction terms between cap rates with the four financial factors in (18), and the results are robust.

4.5. Cap rate and alpha with non-systematic risk

Finally, I investigate whether the predictability of alpha is a spurious relation due to the pricing of non-systematic risk. If investors use higher discount rates to value properties that they perceive have higher non-systematic risk, there may be a positive relationship between cap rates and ex post returns that are not related to factors. Such a relationship might be picked up by alpha in factor models that do not account for non-systematic risk, and thus lead to a spurious relation between cap rates and alpha.

I analyze three types of non-systematic risk. The first is the idiosyncratic component of properties' returns, which I call idiosyncratic risk. As Plazzi, Torous and Valkanov (2008) and others recognize, individual properties often represent a non-trivial share of investors' portfolios; therefore, investors may price the idiosyncratic component of properties' returns.

To construct a measure for the idiosyncratic risk of each property, I run the following simple holding-period latent factor model.

$$\log\left(R_i\right) - \sum_{T=buy_i+1}^{sell_i} \log\left(T_t\right) = \alpha U_i + \sum_{t=buy_i+1}^{sell_i} M_t + \varepsilon_i \quad (19)$$

The regression residual of (19) captures the component of the holding period risk premium that is not explained by the average effects of all factors, and thus helps measure the idiosyncratic risk. I calculate the squared regression residual for each property, $\hat{\varepsilon}_i^2$, and then normalize it by dividing it with the duration of the holding period. I use the normalized squared residual, denoted by I_i, to measure idiosyncratic risk.

Investors may also price temporal volatility of property returns. Because properties' returns are not observed in each period, I measure the temporal volatility of each property's returns during its holding period with the standard deviation of the real estate factor I estimated earlier during the holding period. While this measure is not perfect, it captures the temporal variation of the common component of all properties' returns.

Lack of recent transactions of similar properties, i.e. market thinness, may make it difficult for investors to value properties. As a result, investors might use higher discount rates to price properties and expect higher ex post returns. I measure the market thinness of a property in its acquisition quarter with the negative number of acquisitions of the same type properties in the previous two quarters in the NCREIF database. I use negative numbers because the market is thinner when there are a smaller number of recent transactions. In unreported robustness checks, I use two alternative measures: (1) the negative number of acquisitions of the same type properties in the *same* quarter and (2) the negative number of acquisitions of the same type properties in the previous two quarters in the *same* metro area. The results are robust to the alternative measures and thus not reported.

I estimate the latent factor model in (12) with the idiosyncratic risk, temporal volatility of real estate factor returns, and time-varying market thinness included, for the whole sample as well as the four property types. Results are reported in Table 14. First, it is apparent that cap rates still have strong predicting power for alpha. The coefficient is 0.211 for the whole sample, and varies between 0.162 for office and 0.311 for apartment, all being statistically significant. Second, idiosyncratic risk is significantly correlated with the holding period risk premium for the whole sample and for all property types except office. Third, temporal volatility is also significantly correlated with the holding period risk premium for the whole sample and for all types except retail. Fourth, time-varying market thinness is significantly correlated with the holding period risk premium for the whole sample, and for industrial and office properties. Overall, Table 14 shows that the predicting power of cap rates for alpha is robust when the pricing effect of noon-systematic risk is accounted. In addition, it provides evidence that real estate investors earn higher returns for bearing more non-systematic risk. In unreported regressions, I further include the interaction terms between the cap rate and the four financial market factors as well as the real estate factor, and the results are robust.

5. Conclusions

Do discount rates predict ex post investment returns? Current evidence in the literature mostly comes from the stock market. This paper is the first to study return predictability of individual assets in another major capital market – the private commercial real estate. Studying commercial real estate has two main advantages. First, as rents are more stable/predictable than dividends, properties' cap rates may measure real estate investor's discount rates more accurately than price-dividend ratios do for stock investors. This may lead to cleaner results of testing the predicting power of discount rates. Second, the commercial real estate market differs from the stock market in many aspects, including liquidity, transparency, and trading mechanisms. As a result, testing return predictability in this market would be a valuable robustness check to the large body of research that uses stock market data. In addition to the above advantages, studying private commercial real estate itself is important due to the large size of this market in the economy and the lack of knowledge on its basic risk and return characteristics.

Leveraging detailed property level information from a proprietary dataset of about 950 billion dollar worth of properties in the U.S. from 1977 to 2014, I find that acquisition cap rates have significant predicting power for ex post returns of individual properties. This result is robust across property types, metro areas with different market thinness, and properties with different duration of holding periods, and not likely due to sample selection bias or unknown mechanical relationships. Further, I find that cap rates have stronger predicting power in the short term, which is an interesting contrast with results from the stock market, in which returns seem to be better predicted in the long term. This distinction is consistent with the notion that cap rates measure discount rates more accurately in the short term but price-dividend ratios of stocks may measure discount rates better in the long term.

Using a holding-period four-factor model that allows heterogeneous property loadings on factors, I find strong evidence that cap rates predict Jensen's alpha. This result is robust across property types, metro areas, and properties with different holding period duration. I then estimate a latent-factor holding-period model and show that the predicting power of cap rates is not likely due to latent factors. To mitigate a possible bias due to heterogeneous loadings across properties for latent factors, I estimate a real estate factor that captures the common component of returns not explained by the four financial market factors, and allow properties to have heterogeneous loadings on this factor. Estimating such a model provides robust evidence that cap rates still predict alpha. Finally, I construct measures for non-systematic risk and include them in latent factor models. Results suggest that non-systematic risk is priced, and the predicting power of cap rates for alpha remains strong in the presence of the non-systematic risk. However, since my measures may not completely capture all non-systematic risk, I am unable to rule out the possibility the predictability of alpha is spurious. Specifically, if investors price non-systematic risk in cap rates and they receive higher returns for taking higher non-systematic risk, the higher returns would be captured by alpha in a model that does not perfectly control for all non-systematic risk. More theoretical guidance and better datasets may help shed more light on this issue in future research.

References

Ang, Andrew, and Geert Bekaert, 2007, Stock return predictability: Is it there?, *Review of Financial Studies* 20, 651-707.

Baker, Malcolm, and Jeffrey Wurgler, 2004, A catering theory of dividends, *Journal of Finance* 59, 1125–1165.

Belo, Frederico, Pierre Collin-Dufresne, and Robert S. Goldstein, 2015, Dividend dynamics and the term structure of dividend strips, *Journal of Finance* 70, 1115–1160.

Campbell, John Y., and John Ammer, 1993, What moves the stock and bond markets? A variance decomposition for long-term asset returns, *Journal of Finance* 48, 3-37.

Campbell, John Y., and Robert Shiller, 1988, The dividend price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Case, Karl E., and Robert Shiller, 1989, The efficiency of the market for single family homes, *American Economic Review* 79, 125-137.

Chen, Long, 2009, On the reversal of return and dividend growth predictability: A tale of two periods, *Journal of Financial Economics* 92, 128-151.

Chen, Long, Zhi Da, and Xinlei Zhao, 2014, What drives stock price movements?, *Review of Financial Studies* 26, 841-876.

Cochrane, John, 2005, The risk and return of venture capital, *journal of Financial Economics* 75, 3-52.

Cochrane, John, 2011, Discount rates, *Journal of Finance* 66, 1047-1108.

Cochrane, John H., 1992, Explaining the variance of price-dividend ratios, *Review of Financial Studies* 5, 243-280.

Cochrane, John H., 1994, Permanent and transitory components ofgnpand stock prices, *Quarterly Journal of Economics* 109, 241–266.

Driessen, Joost, Tse-Chun Lin, and Ludovic Phalippou, 2012, A new method to estimate risk and return of non-traded assets from cash fows: The case of private equity funds, *Journal of Financial and Quantitative Analysis* 47, 511-535.

Fama, Eugene F, and Kenneth R French, 1988, Dividend yields and expected stock returns. Journal of financial economics, *Journal of Financial Economics* 22, 3-27.

Fama, Eugene, and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, journal of financial economics, *Journal of Financial Economics* 33, 3-36.

Franzoni, Francesco, Eric Nowak, and Ludovic Phalippou, 2012, Private equity performance and liquidity risk, *Journal of Finance* 67, 2341–2373.

Goetzmann, William N., 1992, The accuracy of real estate indices: Repeat sale estimators, *Journal of Real Estate Finance and Economics* 5, 5-53.

Goetzmann, William N., and Liang Peng, 2006, Estimating house price indices in the presence of seller reservation prices, *Review of Economics and Statistics* 88, 100-112.

Golez, Benjamin, 2014, Expected returns and dividend growth rates implied by derivative markets, *Review of Financial Studies* 27, 790-822.

Goyal, Amit, and Ivo Welch, 2003, Predicting the equity premium with dividend ratios, *Management Science* 49, 639-654.

Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.

Hochberg, Yael, and Tobias Mühlhofer, 2015, Capital-market competitiveness and managerial investment decisions: Evidence from commercial real estate, *Working paper*.

Jules H. V. Binsbergen, and Ralph S. J. Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance* 65, 1439-1471.

Korteweg, Arthur, and Morten Sorensen, 2010, Risk and return characteristics of venture capital-backed entrepreneurial companies, *Review of Financial Studies* 23, 3738–3772.

Larrain, Borja, and Motohiro Yogo, 2008, Does firm value move too much to be justified by subsequent changes in cash flow? 87:200–26, *Journal of Financial Economics* 87, 200-226.

Lettau, Martin, and Stijn Van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607-1652.

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding predictability, *Journal of Political Economy* 112, 1-47.

Pastor, Lubos, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642-685.

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

Plazzi, Alberto, Walter Torous, and Rossen Valkanov, 2008, The cross-sectional dispersion of commercial real estate returns and rent growth: Time variation and economic flucuations, *Real Estate Economics* 36, 403-439.

Plazzi, Alberto, Walter Torous, and Rossen Valkanov, 2010, Expected returns and the expected growth in rents in commercial real estate, *Review of Financial Studies* 23, 3469-3519.

Plazzi, Alberto, Walter Torous, and Rossen Valkanov, 2012, Exploiting property characteristics in commercial real estate portfolio allocation, *Journal of Portfolio Management* 35, 39-50.

Shefrin, Hersh M., and Meir Statman, 1984, Explaining investor preference for cash dividends, *Journal of Financial Economics* 13, 253–282.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Stein, Jeremy, 1996, Rational capital budgeting in an irrational world, *Journal of Business* 69, 429–455.

Table 1. Summary of NCREIF properties

This table counts properties in the NCREIF database in different categories.

	True sales	Other sales	Held	All
Cap rate and true MIRR	6,834	0	0	6,834
Cap rate and estimated MIRR	10	973	2,817	3,800
Cap rate but no MIRR	52	1,851	3,080	4,983
No cap rate	6,502	5,645	5,574	17,721
All	13,398	8,469	11,471	33,338

Table 2. Data summary

This table reports the number of properties (all, with actual annualized holding-period total return MIRRs, and with estimated annualized holding-period total return MIRRs), the number of metro areas where the properties are located, and summary statistics of acquisition cap rates and holding-period total return MIRRs for all properties and each property type.

	All	Apartment	Industrial	Office	Retail
All properties	4,433	1,134	1,573	1,056	670
Properties true MIRRs	2,706	759	869	702	376
Properties est. MIRRs	1,727	375	704	354	294
Metro areas	181	106	95	88	134
Cap rate minimum	0.010	0.011	0.010	0.011	0.012
Cap rate 25%	0.052	0.045	0.060	0.053	0.054
Cap rate median	0.069	0.057	0.077	0.071	0.068
Cap rate 75%	0.086	0.072	0.092	0.089	0.086
Cap rate maximum	0.150	0.139	0.150	0.147	0.150
Cap rate mean	0.069	0.059	0.075	0.071	0.071
Cap rate std.	0.025	0.021	0.025	0.025	0.022
MIRR minimum	-0.094	-0.078	-0.082	-0.094	-0.073
MIRR 25%	0.021	0.005	0.029	0.019	0.026
MIRR median	0.069	0.066	0.073	0.065	0.074
MIRR 75%	0.118	0.109	0.116	0.116	0.143
MIRR maximum	0.395	0.395	0.395	0.390	0.388
MIRR mean	0.079	0.071	0.082	0.077	0.089
MIRR std.	0.084	0.087	0.079	0.086	0.085

Table 3. Predicting ex post total returns

Panel A of this table reports results of four specifications of the regression of the holding-period annualized total return MIRR against the acquisition cap rate and other variables across properties. White's heteroscedasticity-consistent standard deviations are reported in parentheses. The mean squared error (MSE) is the mean of squared regression residuals. The "Placebo with larger MSEs" reports the percentage of placebo test rounds that have MSEs greater than the MSE of the regression. Panel B summarizes results of 1,000 rounds of placebo tests corresponding to each specification in Panel A. Each round uses cap rates randomly drawn from all properties (without replacement) as the acquisition cap rates, but is otherwise identical to the corresponding regression in Panel A. Panel B reports the mean and standard deviation of the intercept term, the coefficient of the cap rate, the adjusted R2, and the MSE of regression residuals across the 1,000 rounds of placebo tests. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

Panel A. Predicting total return MIRRs				
	I (All)	II (All)	III (True)	IV (Estimated)
Intercept	-0.017***	0.054***	0.041***	0.042***
	(0.003)	(0.014)	(0.018)	(0.017)
Cap rate	1.381***	1.063***	1.062***	1.043***
	(0.050)	(0.059)	(0.082)	(0.080)
Type fixed effect	No	Yes	Yes	Yes
Metro fixed effect	No	Yes	Yes	Yes
Acquisition period fixed effect	No	Yes	Yes	Yes
Sample size	4,433	4,433	2,706	1,727
Adjusted R2	0.164	0.306	0.291	0.374
MSE	0.006	0.005	0.006	0.002
Placebo with larger MSEs	100%	100%	100%	100%

Panel B. Placebo tests (Monte Carlo simulations)				
	I (All)	II (All)	III (True)	IV (Estimated)
Simulation rounds	1,000	1,000	1,000	1,000
Intercept	0.079***	0.136***	0.125***	0.117***
	(0.004)	(0.003)	(0.005)	(0.004)
Cap rate	-0.002	0.000	0.003	-0.002
	(0.052)	(0.045)	(0.070)	(0.055)
Adjusted R2	0.000	0.239***	0.239***	0.253***
	(0.000)	(0.000)	(0.000)	(0.001)
MSE	0.007***	0.005***	0.006***	0.003***
	(0.000)	(0.000)	(0.000)	(0.000)

Table 4. Predicting ex post capital appreciation returns

Panel A of this table reports results of four specifications of the regression of the holding-period annualized capital appreciation MIRR against the acquisition cap rate and other variables across properties. White's heteroscedasticity-consistent standard deviations are reported in parentheses. The mean squared error (MSE) is the mean of squared regression residuals. The "Placebo with larger MSEs" reports the percentage of placebo test rounds that have MSEs greater than the MSE of the regression. Panel B summarizes results of 1,000 rounds of placebo tests corresponding to each specification in Panel A. Each round uses cap rates randomly drawn from all properties (without replacement) as the acquisition cap rates, but is otherwise identical to the corresponding regression in Panel A. Panel B reports the mean and standard deviation of the intercept term, the coefficient of the cap rate, the adjusted R2, and the MSE of regression residuals across the 1,000 rounds of placebo tests. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

Panel A. Predicting capital appreciation MIRRs				
	I (All)	II (All)	III (True)	IV (Estimated)
Intercept	-0.050***	0.014	0.002	-0.019
	(0.004)	(0.015)	(0.018)	(0.018)
Cap rate	0.929***	0.648***	0.631***	0.648***
	(0.050)	(0.058)	(0.081)	(0.058)
Type fixed effect	No	Yes	Yes	Yes
Metro fixed effect	No	Yes	Yes	Yes
Acquisition period fixed effect	No	Yes	Yes	Yes
Sample size	4,433	4,433	2,706	1,727
Adjusted R2	0.078	0.252	0.238	0.252
MSE	0.006	0.005	0.006	0.002
Placebo with larger MSEs	100%	100%	100%	100%
Panel B. Placebo tests (Monte Carlo simulations)				
	I (All)	II (All)	III (True)	IV (Estimated)
Simulation rounds	1,000	1,000	1,000	
Intercept	0.015***	0.063***	0.052***	0.029***
	(0.003)	(0.003)	(0.005)	(0.004)
Cap rate	-0.001	0.001	0.003	-0.002
	(0.049)	(0.047)	(0.065)	(0.054)
Adjusted R2	-0.000	0.225***	0.219***	0.264***
	(0.000)	(0.000)	(0.000)	(0.001)
MSE	0.007***	0.005***	0.006***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)

Table 5. Property type and the prediction of ex post total returns

This table reports results of the regressions of the holding-period annualized total return MIRR against the acquisition cap rate and other variables across properties for each of the four property types. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

	Apartment	Industrial	Office	Retail
Intercept	0.012	0.100***	0.050*	-0.004
	(0.040)	(0.025)	(0.028)	(0.022)
Cap rate	1.530***	0.789***	1.053***	1.478***
	(0.153)	(0.086)	(0.118)	(0.142)
Metro fixed effect	Yes	Yes	Yes	Yes
Acquisition period fixed effect	Yes	Yes	Yes	Yes
Sample size	1,134	1,573	1,056	670
Adjusted R2	0.330	0.332	0.242	0.590

Table 6. Market thinness and the prediction of ex post total returns

This table reports results of five specifications of the regression of the holding-period annualized total return MIRR against the acquisition cap rate, its interaction with variables that measure the thinness of the commercial real estate market of the metro area where the property is located, and other variables across properties. "Total volume" for a property is a continuous variable that equals the number of unique properties of any types located in the same metro area with the property that have ever been held by NCREIF members during the entire sample period from 1977 to 2014. "Type volume" is a continuous variable that equals the number of properties of the same type located in the same metro area with the property that have ever been held by NCREIF members during the entire sample period from 1977 to 2014. "Gateway" is a dummy variable that equals 1 for New York, San Francisco, Washington, D.C., and Houston. There are 554 properties with "Gateway" being 1. "Total top 10" is a dummy variable that equals 1 for the top 10 metro areas with the highest "Total volume". There are 1,893 properties with "Total top 10" being 1. "Type top 10" is a dummy variable that equals 1 if the metro area is among the top 10 with the highest "Type volume" for the type of the property. There are 2,438 properties with "Type top 10" being 1. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

	I	II	III	IV	V
Intercept	0.062***	0.062***	0.061***	0.062***	0.062***
	(0.011)	(0.011)	(0.011)	(0.014)	(0.014)
Cap rate	1.059***	1.060***	1.011***	1.033***	1.045***
	(0.055)	(0.054)	(0.052)	(0.060)	(0.059)
Cap rate * Total volume	-0.000				
	(0.000)				
Cap rate * Type volume		-0.000			
		(0.000)			
Cap rate * Gateway			0.193***		
			(0.046)		
Cap rate * Total top 10				-0.002	
				(0.031)	
Cap rate * Type top 10					-0.037
					(0.031)
Type fixed effect	Yes	Yes	Yes	Yes	Yes
Acquisition fixed effect	Yes	Yes	Yes	Yes	Yes
Sample size	4,430	4,430	4,430	4,430	4,430
Adjusted R2	0.287	0.287	0.290	0.287	0.287

Table 7. Duration and the prediction of ex post returns

This table reports results of four specifications of the regression of the holding-period annualized total return MIRR (Panel A) and capital appreciation MIRR (Panel B) against the acquisition cap rate, its interaction with the duration of the holding period (in quarters), its interaction with the dummy for short duration (shorter than or equal to 16 quarters or four years), its interaction with the dummy for long duration (longer than 28 quarters or seven years), and other variables across properties. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

Panel A. Predicting total return MIRRs				
	I	II	III	IV
Intercept	-0.070***	0.086***	0.051***	0.083***
	(0.017)	(0.017)	(0.017)	(0.014)
Cap rate	1.830***	0.680***	1.234***	0.837***
	(0.087)	(0.078)	(0.080)	(0.081)
Cap rate * Duration	-0.036***			
	(0.002)			
Cap rate * Short duration		0.952***		0.778***
		(0.052)		(0.057)
Cap rate * Long duration			-0.733***	-0.372***
			(0.041)	(0.042)
Type fixed effect	Yes	Yes	Yes	Yes
Metro fixed effect	Yes	Yes	Yes	Yes
Acquisition period fixed effect	Yes	Yes	Yes	Yes
Sample size	2,706	2,706	2,706	2,706
Adjusted R2	0.393	0.395	0.353	0.407

Panel B. Predicting capital appreciation MIRRs				
	I	II	III	IV
Intercept	-0.031	0.046***	0.012	0.043**
	(0.045)	(0.016)	(0.022)	(0.017)
Cap rate	1.389***	0.254***	0.802***	0.412***
	(0.083)	(0.079)	(0.078)	(0.082)
Cap rate * Duration	-0.036***			
	(0.002)			
Cap rate * Short duration		0.940***		0.764***
		(0.053)		(0.059)
Cap rate * Long duration			-0.728***	-0.374***
			(0.043)	(0.044)
Type fixed effect	Yes	Yes	Yes	Yes
Metro fixed effect	Yes	Yes	Yes	Yes
Acquisition period fixed effect	Yes	Yes	Yes	Yes
Sample size	2,706	2,706	2,706	2,706
Adjusted R2	0.343	0.345	0.303	0.358

Table 8. Predicting alpha in a four-factor model

This table reports results of testing whether the cap rate predicts a property's alpha in a four-factor holding-period model for all properties as well as those were sold. "Duration" is the number of quarters of the holding period. "Rm-Rf", "SMB", and "HML" are the three Fama-French factors, and Liquidity is the Pastor and Stambaugh liquidity factor. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

	I	II	III	IV
	(All)	(All)	(True)	(Estimated)
Duration	-0.000	0.003**	0.001	0.006***
	(0.001)	(0.001)	(0.001)	(0.001)
Duration * Cap rate		0.173***	0.228***	0.133***
		(0.029)	(0.039)	(0.047)
Rm-Rf * Cap rate		1.068	-1.197	3.230***
		(0.915)	(1.373)	(1.220)
SMB * Cap rate		-1.086	-5.191**	4.484
		(1.848)	(2.517)	(2.922)
HML * Cap rate		-2.023	-0.388	-3.905**
		(1.488)	(2.121)	(1.900)
Liquidity * Cap rate		2.567**	2.167	2.456
		(1.001)	(1.402)	(1.511)
Rm-Rf	0.204***	0.163***	0.229***	0.112***
	(0.021)	(0.021)	(0.032)	(0.029)
SMB	0.186***	0.130***	0.193***	0.019
	(0.041)	(0.041)	(0.059)	(0.055)
HML	0.034	0.038	0.004	0.090**
	(0.033)	(0.034)	(0.048)	(0.045)
Liquidity	0.336***	0.189***	0.246***	0.090***
	(0.023)	(0.024)	(0.034)	(0.035)
Sample size	4,430	4,430	2,706	1,727
Adjusted R2	0.401	0.476	0.488	0.446

Table 9. Predicting alpha in a four-factor model by property types

This table reports results of testing whether the cap rate predicts a property's alpha in a four-factor holding-period model for each of the four property types. "Duration" is the number of quarters of the holding period. "Rm-Rf", "SMB", and "HML" are the three Fama-French factors, and Liquidity is the Pastor and Stambaugh liquidity factor. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

	Apartment	Industrial	Office	Retail
Duration	0.007***	0.002**	0.004***	0.004***
	(0.001)	(0.001)	(0.001)	(0.002)
Duration * Cap rate	0.413***	0.115***	0.060	0.254***
	(0.069)	(0.043)	(0.063)	(0.079)
Rm-Rf * Cap rate	-1.682	1.121	3.443*	1.032
	(2.309)	(1.366)	(2.035)	(2.598)
SMB * Cap rate	-3.953	-2.047	2.190	0.755
	(4.654)	(2.841)	(4.181)	(4.493)
HML * Cap rate	-0.154	-1.794	-5.233*	0.112
	(2.974)	(2.563)	(2.850)	(3.765)
Liquidity * Cap rate	-1.477	4.110***	5.569***	-0.771
	(2.464)	(1.505)	(1.996)	(2.938)
Rm-Rf	0.150***	0.142***	0.118**	0.181***
	(0.050)	(0.033)	(0.049)	(0.051)
SMB	0.186**	0.111	-0.040	0.298***
	(0.083)	(0.070)	(0.098)	(0.088)
HML	-0.035	0.004	0.080	0.309***
	(0.058)	(0.063)	(0.070)	(0.083)
Liquidity	0.099*	0.235***	0.137***	0.135**
	(0.057)	(0.036)	(0.053)	(0.063)
Sample size	1,134	1,573	1,056	670
Adjusted R2	0.449	0.548	0.384	0.611

Table 10. Predicting alpha in a latent factor model

This table reports results of testing whether the cap rate predicts a property's alpha in a latent-factor holding-period model for all properties and for each property type respectively. "Duration" is the number of quarters of the holding period. Latent factor dummies are dummies for each period from 1997:Q2 to 2014:Q4. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

	All	Apartment	Industrial	Office	Retail
Duration	0.002	0.069	0.077	-0.161	0.098
	(0.051)	(0.093)	(0.082)	(0.108)	(0.131)
Duration * Cap rate	0.196***	0.304***	0.158***	0.190***	0.245***
	(0.010)	(0.027)	(0.014)	(0.020)	(0.030)
Latent factor dummy	Yes	Yes	Yes	Yes	Yes
Sample size	4,430	1,134	1,573	1,056	670
Adjusted R2	0.512	0.507	0.611	0.421	0.702

Tale 11. Market thinness, duration, and predicting alpha in a latent factor model

This table reports results of testing whether the cap rate predicts a property's alpha in a latent-factor holding-period model for properties in thick markets (the top 10 metro areas with the highest "Type volume") and thin markets (not in the top 10) and those with long holding periods (duration longer than the median duration, which is 20 quarter) and short holding periods (duration shorter than 20 quarters). "Duration" is the number of quarters of the holding period. Latent factor dummies are dummies for each period from 1997:Q2 to 2014:Q4. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

	Thick market	Thin market	Long term	Short term
Duration	0.193**	-0.168***	-0.013	0.090
	(0.075)	(0.065)	(0.057)	(0.097)
Duration * Cap rate	0.200***	0.196***	0.198***	0.240***
	(0.016)	(0.013)	(0.011)	(0.028)
Latent factor dummy	Yes	Yes	Yes	Yes
Sample size	1,995	2,438	3,277	1,156
Adjusted R2	0.537	0.510	0.501	0.685

Table 12. Real estate factor

This table reports the quarterly series of the estimated real estate factor.

	Q1	Q2	Q3	Q4
1997		0.0469	0.0105	-0.1132
1998	0.0657	0.0088	0.0361	-0.0759
1999	0.1729	0.0257	0.0688	-0.1478
2000	-0.0914	0.2052	-0.0281	-0.0107
2001	-0.0654	0.0242	0.1688	-0.3751
2002	0.1998	-0.0445	0.1085	-0.1085
2003	0.1016	-0.0232	-0.0531	-0.0869
2004	0.0799	0.0533	0.0104	-0.1149
2005	0.0706	0.0531	-0.0375	-0.0308
2006	0.1159	-0.0882	0.1142	-0.1188
2007	-0.0454	0.0225	0.0304	-0.0859
2008	0.0664	-0.0088	0.0114	0.0352
2009	0.1034	-0.0769	-0.0252	-0.1760
2010	0.0389	0.1605	-0.1513	0.0090
2011	-0.0395	0.0129	0.1663	-0.1918
2012	0.0735	0.0678	-0.0067	-0.0073
2013	-0.0905	0.1633	-0.0419	-0.0368
2014	-0.0539	0.1219	0.0054	0.0184

Table 13. Real estate factor and predicting alpha in a latent factor model

This table reports results of testing whether the cap rate predicts a property's alpha in a latent-factor holding-period model for all properties and for each property type respectively. "Duration" is the number of quarters of the holding period. "Real estate" is the real estate factor. Latent factor dummies are dummies for each period from 1997:Q2 to 2014:Q4. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

	All	Apartment	Industrial	Office	Retail
Duration	0.002	0.065	0.075	-0.162	0.098
	(0.051)	(0.093)	(0.082)	(0.108)	(0.132)
Duration * Cap rate	0.196***	0.302***	0.159***	0.190***	0.246***
	(0.010)	(0.027)	(0.014)	(0.020)	(0.030)
Real estate * Cap rate	-2.106	-5.884	-3.286	2.009	0.540
	(2.240)	(5.517)	(3.556)	(4.789)	(6.123)
Latent factor dummy	Yes	Yes	Yes	Yes	Yes
Sample size	4,430	1,134	1,573	1,056	670
Adjusted R2	0.512	0.507	0.611	0.420	0.701

Table 14. Non-systematic risk and predicting alpha in a latent factor model

This table reports results of testing whether the cap rate predicts a property's alpha in a latent-factor holding-period model for all properties and for each property type respectively. "Duration" is the number of quarters of the holding period. "Idiosyncratic risk" is the squared regression residual from a latent factor model with homogeneous loadings estimated using all properties (normalized by divided with Duration). "Volatility" is the standard deviation of quarterly real estate factor during the holding period. "Thinness" is the negative number of acquisitions (of the same property type) in the previous two quarters. Its coefficients and standard deviations are multiplied by 10,000 when being reported. Latent factor dummies are dummies for each period from 1997:Q2 to 2014:Q4. White's heteroscedasticity-consistent standard deviations are reported in parentheses. ***, **, and * indicate significant levels of 1%, 5%, and 10% respectively.

	All	Apartment	Industrial	Office	Retail
Duration	0.028	0.029	0.125	-0.150	0.045
	(0.051)	(0.096)	(0.081)	(0.100)	(0.134)
Duration * Cap rate	0.211***	0.311***	0.162***	0.218***	0.219***
	(0.010)	(0.028)	(0.014)	(0.020)	(0.030)
Idiosyncratic risk	5.787***	3.616*	8.819***	3.091	12.094***
	(1.213)	(1.887)	(1.840)	(2.343)	(2.477)
Volatility	1.436***	1.469***	1.605**	2.544***	-0.378
	(0.169)	(0.378)	(0.364)	(0.452)	(0.501)
Thinness	1.578***	1.754	3.165***	3.664*	-6.536
	(0.332)	(2.249)	(0.815)	(2.103)	(3.599)
Latent factor dummy	Yes	Yes	Yes	Yes	Yes
Sample size	4,430	1,134	1,573	1,056	670
Adjusted R2	0.541	0.527	0.637	0.465	0.734

Figure 1. Histogram of property cap rates

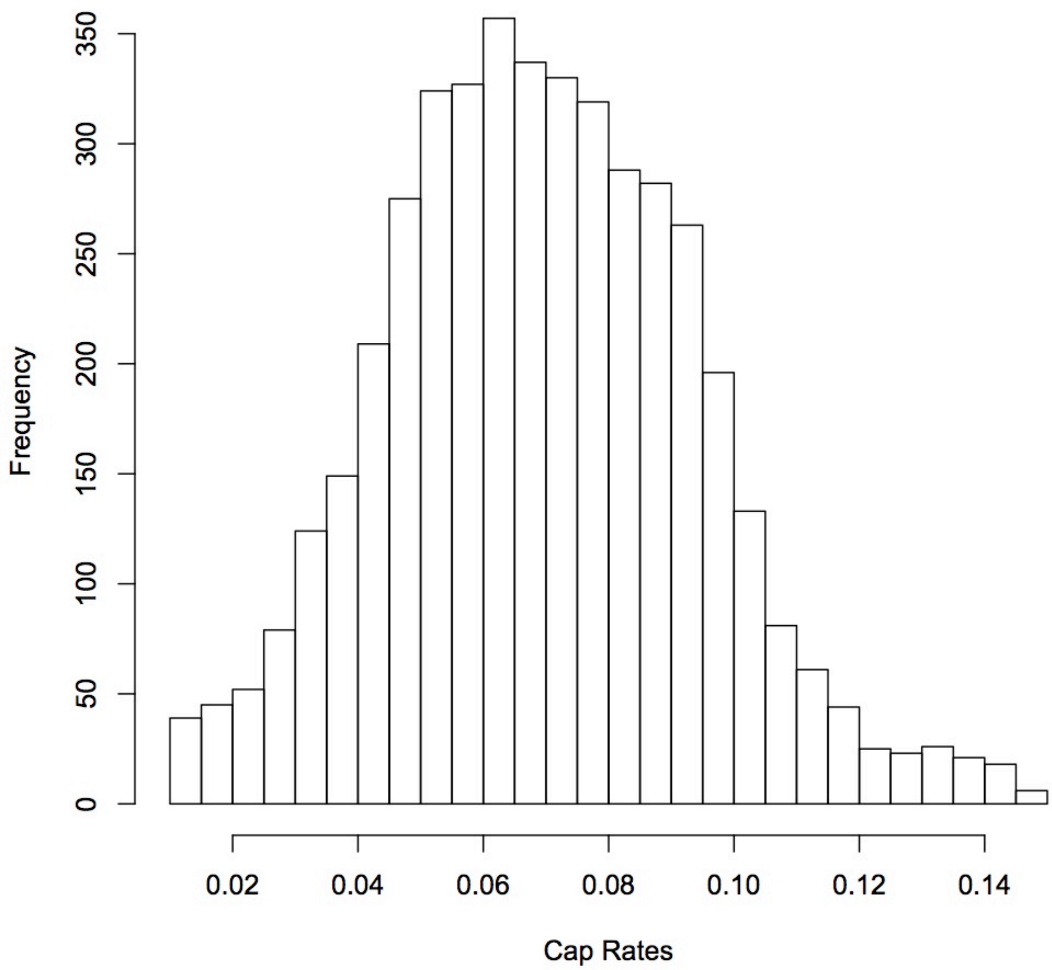

Figure 2. Histogram of annual total return MIRRs

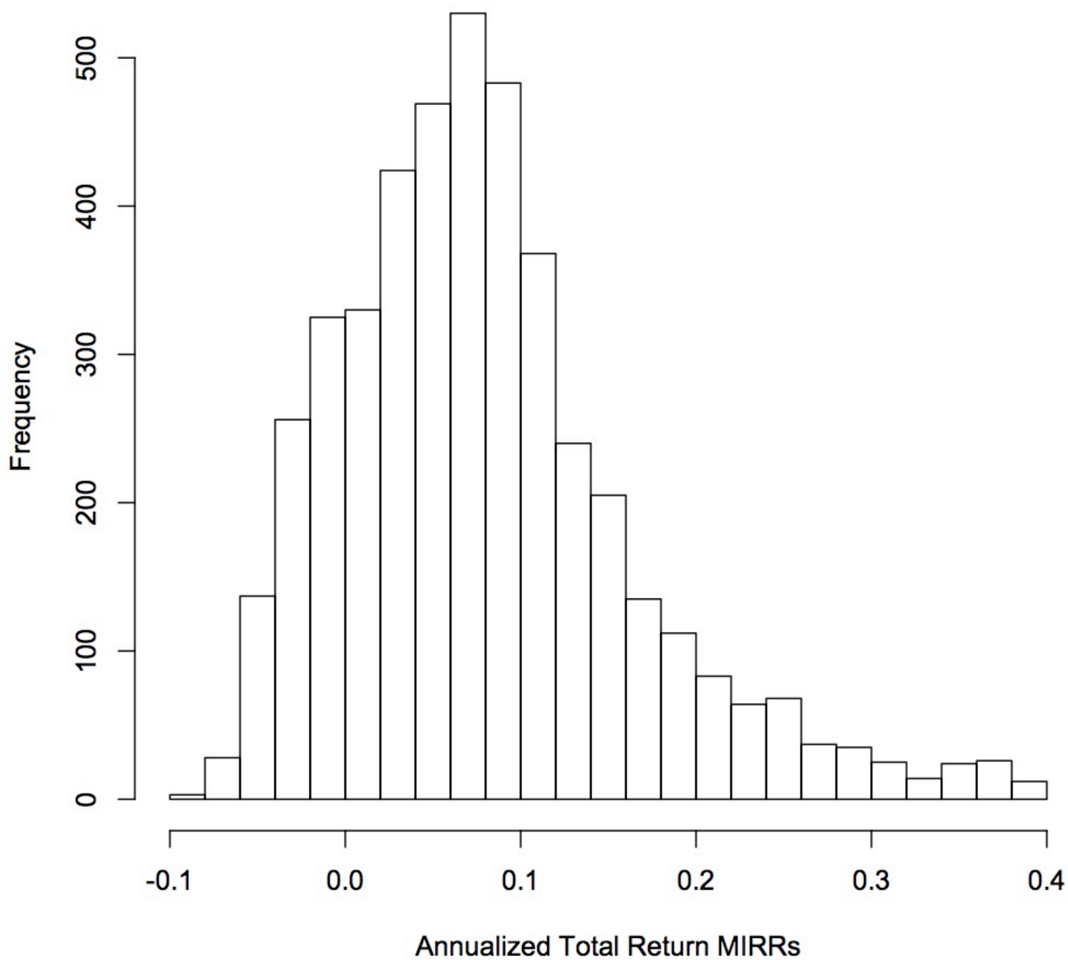

Figure 3. Real estate factor

Figure 4. Histogram of the coefficient of the real estate factor out of sample tests