

On Indexing Commercial Real Estate Properties and Portfolios

Walter I. Boudry
N. Edward Coulson
Jarl G. Kallberg
Crocker H. Liu

Commercial real estate indices play an important role in performance evaluation and overall investment strategy. However, the issue of how representative they are of the returns on portfolios of commercial properties is an open issue. Our study addresses this topic by analyzing a sample of 12,427 repeat sales transactions between Q4 2000 and Q2 2011. We find that the aggregate real estate indices (Moody's REAL CPPI) do a good job of tracking real returns when portfolios of more than 20 properties are considered. At this level, tracking is somewhat less effective than our benchmark of the S&P500 and its component stocks. Compared to the average root mean squared deviation (RMSD) from one asset, randomly selected portfolios with 20 assets reduce the RMSD by 75 % for the S&P500 compared to 66 % for the aggregate index. These results suggest that the aggregate indices can be effective in hedging and evaluating the performance of direct real estate investment. We further find that tracking at the property type level provides little benefit over using an aggregate index. However, indexing using a property type and location matched index provides lower tracking error for any level of diversification.

Introduction

The development of viable indices for commercial real estate is essential because investment benchmarks play a pivotal role in performance evaluation, hedging decisions, and overall investment strategy. Although practitioners and academics have invested considerable effort in this area, the issue of how well these indices track the returns on “real” portfolios of commercial properties remains an open issue.

We approach this question in a very traditional manner utilizing the concept of tracking error: the difference between the returns on a given portfolio and its benchmark. Tracking error measures how closely the return on a given portfolio follows the return on the index that it is benchmarked to.

From its origins in Markowitz (1952) and Sharpe (1964), one of the central theories in finance concerns systematic and idiosyncratic risk. An area of particular relevance to this study is how quickly idiosyncratic risk falls as the number of assets in the investment portfolio increases because tracking error measures the square root of idiosyncratic risk. In one of the most cited articles in finance, Evans and Archer (1968) use simulated returns from stocks in the S&P500 to estimate how a randomly selected, equally weighted portfolio's standard deviation decreases as the number of stocks in the portfolio increases. From their conclusion: The results also raise doubt concerning the economic justification of increasing portfolio sizes beyond 10 or so securities ... Elton and Gruber (1977) analytically show a related, but less binding result: 51 % of an equally weighted portfolio's diversifiable risk can be eliminated using random portfolios containing 10 stocks. On the other hand, Statman (1987) argues that a well diversified portfolio must contain at least 30 stocks. More recently, Campbell et al. (2001) state that about 50 randomly selected securities can achieve "relatively complete portfolio diversification." They also document an increase in the idiosyncratic risk of individual stocks; see also Brown and Kapadia (2007).¹

There are several sources of tracking error and basis risk, notably style differences, that lead to systematic variation between the returns on a given portfolio and its benchmark. Basis risk is present in our study since the portfolios we analyze are somewhat different from the assets that constitute the indices (as we will discuss further in Section 3). The source of tracking error that forms the primary focus of this study is the error stemming from relatively small numbers of assets in the investment portfolio. While this issue is significant for stock market investment, as the above references suggest, it is also relevant for direct real estate investment for a number of additional reasons: (i) the indivisibility of commercial real estate assets makes it impossible to take small positions; (ii) the transactions costs typically incurred in real estate investment are extremely high²; and (iii) the lack of effective vehicles for short selling and hedging.

This study uses a sample of 12,427 repeat sales from 2000 to 2011 to create a set of "real" property returns over different holding periods. We then compare the returns on the Real Capital Analytics repeat sales indices (Moody's REAL CPPI) to returns on randomly generated portfolios of properties. Our objective is twofold: first, to determine how well the index tracks the performance of

¹ Of course there is an extensive literature dealing with the issue of idiosyncratic risk and stock returns; see, for example, Fu (2009).

² For example, Collett et al. (2003) show how high transaction costs and illiquidity lead to very long holding periods for institutional real estate.

real portfolios with varying numbers of properties; second, to determine whether indices based on location or on property type have better tracking characteristics than a more aggregated index.

Our main findings are the following: We find that the aggregate real estate indices do a good job of tracking real returns when portfolios of more than 20 properties are considered. At this level, risk reduction is somewhat less effective than our benchmark of the S&P500 and its component stocks on a relative basis. Compared to the average root mean squared deviation (RMSD) from one asset, randomly selected portfolios with 20 assets reduce the RMSD by 75 % for the S&P500 compared to 66 % for the aggregate CRE index. However, in an absolute sense the average RMSD is lower for the real estate indices, partially due to the smaller values on the individual asset level.

We also find that tracking at the property type level provides little benefit over using an aggregate index. However, we provide some evidence that indexing using a property type and location matched index provides lower tracking error for any level of diversification. This result, while limited in scope, suggests that location is a more important factor in performance than property type. We find that returns on properties with shorter holding periods are much higher; specifically, properties with a holding period between 4 and 8 quarters have an average return of 17.6 % compared to 10.5 % for properties with longer holding periods. We believe our results should be important in determining investment strategies for direct investment in real estate and for assessing the value of these indices for hedging and performance measurement.

The remainder of this paper is organized as follows. Section 2 presents a brief overview of the literature most relevant to the construction of real estate indices. Section 3 presents the data and our econometric approach. Section 4 presents our empirical results. The conclusions are presented in Section 5.

Literature Review

Although the construction of transaction based commercial real estate indices is a relatively recent phenomenon, these indices had their genesis in residential real estate. The studies on the construction of residential real estate indices include the seminal papers by Rosen (1974) on hedonic methods based on homeowners' derived utility from specific house attributes and Bailey et al. (1963) on repeat sales indices (RSI). Case and Shiller (1989) provide further major modifications to the RSI.

While earlier research on housing indices focused on how to account for quality changes when constructing price indices, later research has dealt with various pricing biases and refinements arising from the representativeness of the sample (see for example, Clapp and Giaccotto 1992, Gatzlaff and

Haurin 1994, 1997), instability of house attributes (Dombrow et al. 1997), heterogeneous appreciation rates (Goodman and Thibodeau 2003) and shifting reservation prices (Goetzman and Peng 2006), among several others.

The development of transaction based commercial real estate indices was hampered by the unavailability of large commercial real estate databases until fairly recently (early 2000s.)³ With the advent of CoStar and Real Capital Analytics, two large database vendors, the construction of CRE indices commenced in earnest, although Hoag (1980) constructed the earliest chronicled transaction based hedonic index of commercial property using industrial properties. Early attempts to develop an institutional real estate index using sales data from the National Council of Real Estate Investment Fiduciaries (NCREIF) include Miles et al. (1990) and Webb et al. (1992). Using a different methodology, Fisher et al. (2007) developed a quarterly transaction-based index (TBI) of property-level investment performance for major property types included in NCREIF. Their methodology extends Geltner's earlier work on commercial real estate indices.⁴

A strand of literature that is closely related to this study addresses the issue of unsystematic (idiosyncratic) and systematic (market) risk. Starting with Evans and Archer (1968), the financial literature shows that naive diversification results in a decline in portfolio risk when the number of assets included in a portfolio increases. There are several challenges in attempting to address how many assets are necessary to reduce unsystematic risk in real estate given the lumpiness of the asset,⁵ no short sales, and other real estate market imperfections. Early studies in real estate did not consider these imperfections in looking at the decline in portfolio risk. An exception to this is the study by Kallberg et al. (1996), which takes into account the indivisibility of real assets and no short sales. They show that a 9 % allocation to real estate is optimal using the modified internal rate of return as the performance measure in conjunction with a branch and bound algorithm. Given the limited size of their sample, the authors were unable to address the question of the number of properties required to reduce unsystematic risk. In one of the few studies to examine this issue, using UK property funds Byrne and Lee (2003) suggest that hundreds if not thousands of properties are required to significantly reduce the

³ The National Association of Real Estate Investment Trusts (NAREIT) started to publish their REIT series in the early 1970's. While REITs represent securitized commercial real estate, they are a hybrid between stock and underlying real estate. While studies such as Liu and Mei (1992) and Gyourko and Keim (1992) have shown that REITs are a leading indicator of the underlying private real estate market, controversy remains as to the extent to which REITs reflect the underlying real estate. Some notable studies include, but are not limited to, Barkham and Geltner (1995), Geltner and Goetzmann (2000), and Boudry et al. (2012).

⁴ See Fisher et al. (2003, 2004), Gatzlaff and Geltner (1998), and Geltner and Goetzmann (2000).

⁵ In contrast to stocks where an investor can purchase a single share of stock, a real estate investor cannot typically purchase one square foot of a property.

unsystematic risk and to achieve a well-diversified portfolio. The authors find this puzzling given UK institutional investors have a median holding of 45 properties.

Recognizing the market imperfections in real estate, Geltner and Kluger (1998), Riddiough et al. (2005), and Horrigan et al. (2009) use a technique similar to attribution analysis to construct real estate indices by identifying a linear combination of attributes in one index that mimic the attributes in the other index. In other words, investment performance is decomposed into relevant firm or asset attributes and then those attributes are adjusted to create similar risk indices. The creation of the resulting indices and corresponding returns allow an investor to compare the performance of publicly versus privately held commercial real estate investments. In addition to this, the constructed portfolios are useful for hedging, especially to the extent that they can replicate the return distribution of the benchmark portfolio. Horrigan et al. (2009), for example, show that constructed REIT pure play indices display volatilities similar to transaction-based indices such as the Moody's/REAL CPPI.

One further consideration in using the constructed indices is the investment horizon. MacKinnon and Zaman (2009) show that returns to direct real estate are mean reverting and risk decreases with the investment horizon. Feng and Geltner (2011), using a variation of property level performance attribution based on the internal rate of return, find a wide dispersion in property price performance. They argue that improving property-level operational management is a key to better overall investment performance.

Data and Methodology

Data

Our sample of repeat sales from the CoStar database consists of 12,427 repeat sales occurring between Q4 2000 and Q2 2011. This time restriction is imposed in order to match the time period for which we are able to obtain repeat sales indices from Real Capital Analytics. The second filter we impose is that properties must have a purchase price greater than \$1 million.⁶ We include only apartment, office, industrial, and retail properties in our sample. While other property types exist, the resulting small sample sizes make estimation of property type specific indices and the formation of portfolios problematic.

⁶ Our results do not appear overly sensitive to this cut off. We have tried running the aggregate sample with a \$2 million filter and the results are similar. We use the smaller cut off to have more properties available for portfolio formation.

To be included in the sample, non-apartment properties must be greater than 2,500 square feet and apartment properties must have greater than 10 units, so that the most economically significant assets are included in our sample. We exclude all transactions identified by CoStar as being distressed or non-arm's length.

We create repeat sales pairs by matching sales at the same property address over time. We exclude repeat sales where property characteristics, such as property type and property size, have changed. As is conventional, we also exclude properties that have a resale window of less than 1 year. Properties that have annualized price appreciation of less than -40 % and greater than 60 % are also excluded.⁷

Table 1 reports descriptive statistics for the repeat sales used in our analysis. To make properties that have different holding periods comparable, we calculate the annualized holding period return for each property as $\left(\frac{P_s}{P_B}\right)^{\frac{1}{i}} - 1$, where i is the number of years between the end of the quarter of purchase and the end of the quarter of sale, P_s is the price at sale and P_B is the price at purchase. The left set of columns report descriptive statistics of the annualized holding period return for the property, while the right columns report descriptive statistics for the sample matched Moody's REAL CPPI. Aggregate is the aggregate sample including apartment, office, industrial, and retail properties. Apartment, Office, Industrial, and Retail are property type subsamples. SoCal Apartments is a sample of apartment properties in San Diego and Los Angeles. Panel A reports results for the aggregate sample, while Panels B and C report results for properties with holding periods greater than 8 quarters and less than or equal to 8 quarters respectively.

As expected, the standard deviation at the property level is much larger than the variation observed at the index level. In most cases the standard deviation is twice as large at the individual property level.⁸ The range is also much larger at the property level. Both maximums and minimums are more extreme at the property level. Mean returns are larger at the property level. There are several possible sources for this bias. First, we have excluded distressed properties from our sample. Second, we have included some properties with a lower value than those used in the index: we have a minimum value of \$1 million versus the index cutoff of \$2.5 million). Third, the minimum holding period for our

⁷ An examination of a sample of these "extreme" observations suggests they are much more likely to be data errors than actual transactions.

⁸ Using S&P500 component stocks, we find a similar relationship: the average standard deviation is twice as large at the individual stock level (22 %) as at the index level (11 %).

returns is lower than for the index: 12 versus 18 months. Finally, the disposition effect suggests that investors are more likely to sell winners than losers.

Table 1 Annualized holding period returns. Table reports descriptive statistics of the annualized holding period returns of 12427 repeat sales in the CoStar database from Q4 2000 to Q2 2011. Annualized holding period returns are calculated as $\left(\frac{P_S}{P_B}\right)^{\frac{1}{i}} - 1$, where i is the number of years between sales, P_B is the price at time of purchase and P_S is the price at the time of sale. Office, Industrial, Multifamily and Retail property types were included and properties had to be larger than 2500 sq.ft. or have more than 10 units for multifamily and have a price greater than \$1 m. All properties must have been held for longer than 1 year and properties with annual appreciation of less than -40 % or greater than 60 % were excluded. Transactions include only arm's length non-distressed sales. Aggregate is the aggregate sample including apartment, office, industrial, and retail properties. Apartment, Office, Industrial and Retail are property type subsamples respectively. SoCal Apartments is a sample of apartment properties in San Diego and Los Angeles. Annualized returns for the Index are the sample matched index return for the appropriate Moody's REAL CPPI. Panel A is the aggregate sample. Panel B reports results for properties with holding periods of greater than 8 quarters, while Panel C reports results for properties with holding periods between 4 and 8 quarters

Sample	Obs	Property				Index			
		Mean	StdDev	Min	Max	Mean	StdDev	Min	Max
Panel A: All									
Aggregate	12427	0.129	0.127	-0.393	0.599	0.097	0.068	-0.173	0.223
Apartment	4601	0.134	0.120	-0.393	0.528	0.104	0.069	-0.146	0.255
Office	2494	0.129	0.127	-0.380	0.599	0.098	0.056	-0.137	0.209
Industrial	2205	0.125	0.128	-0.347	0.573	0.083	0.068	-0.134	0.227
Retail	3127	0.127	0.135	-0.387	0.571	0.094	0.065	-0.146	0.200
SoCal Apartments	1596	0.158	0.112	-0.189	0.522	0.135	0.068	-0.065	0.230
Panel B: Hold>8 Quarters									
Aggregate	8141	0.105	0.112	-0.380	0.532	0.089	0.072	-0.173	0.184
Apartment	2840	0.105	0.104	-0.341	0.505	0.094	0.066	-0.146	0.208
Office	1715	0.104	0.113	-0.380	0.504	0.091	0.058	-0.137	0.169
Industrial	1525	0.103	0.112	-0.315	0.504	0.083	0.070	-0.134	0.180
Retail	2061	0.106	0.120	-0.351	0.532	0.084	0.067	-0.146	0.170
SoCal Apt	891	0.128	0.098	-0.171	0.475	0.126	0.072	-0.065	0.213
Panel C: Hold≤8 Quarters									
Aggregate	4286	0.176	0.140	-0.393	0.599	0.113	0.057	-0.097	0.223
Apartment	1761	0.180	0.128	-0.393	0.528	0.119	0.072	-0.128	0.255
Office	779	0.183	0.140	-0.370	0.599	0.104	0.047	-0.097	0.209
Industrial	680	0.175	0.147	-0.347	0.573	0.108	0.063	-0.071	0.227
Retail	1066	0.166	0.153	-0.387	0.571	0.101	0.062	-0.063	0.200
SoCal Apt	705	0.196	0.117	-0.189	0.522	0.148	0.062	-0.062	0.230

Notice that in our analysis this bias will be evident as the tracking error in very well diversified portfolios. Forming portfolios will remove the tracking error due to idiosyncratic risk, but will be unable to remove tracking error due to these systematic differences. From a practical perspective, an investor is likely to face these same systematic differences between their portfolio and the sample used to create any benchmarking index. While we could reduce the systematic portion of tracking error by using only the sample of properties used to create the index, it would provide an unrealistic setting to explore the benchmarking ability that a real world investor would experience.

Although we report annualized figures to make returns across properties with different holding periods comparable, it may be the case that properties that were held for shorter holding periods are economically different from those held for longer periods. Because we only observe prices when properties transact, the length of holding period is an endogenous choice of the investor. This differs from the case of dealing with stocks, where we observe essentially continuous prices because the shares of the same company held by many investors trade frequently.

Since holding period and performance may be related,⁹ an understanding of the motives for differing holding periods is valuable. The most likely explanation for short holding periods (other than financial distress since we have removed these properties from our sample) is that the investment was opportunistic in nature. To provide some evidence of this difference, for a subsample of properties where we are able to obtain the growth rate of Net Operating Income (NOI), the annualized NOI growth for properties held for less than 2 years was 11.1 %, while it was 3.9 % for properties held for longer than 2 years. This difference is both highly economically and statistically significant. The fact that NOI growth for these properties is dramatically higher lends some credence to the idea that these were not typical investments on average, and may have been opportunistic in nature. Properties held for shorter time periods are on average slightly smaller (both in terms of square feet, and purchase and sale prices) and slightly older than those held for longer time periods. However, the economic magnitude of the differences is quite small.

The results in Panels B and C are also consistent with this notion. Returns are higher for the short holding period properties. Standard deviations are also higher; minimums and maximums tend to be more extreme. In fact, the properties that are held for greater than two years in Panel B look much more like their corresponding index than do the properties held for less than 8 quarters.

To provide further insight into the indices we use, Fig. 1 shows the repeat sales indices used in our analysis. The indices we use are the Moody's REAL CPPI. These indices are available from Q4 2000 to Q2 2011 and are available from Real Capital Analytics and MIT.¹⁰ They are created using a repeat sales methodology similar to that proposed by Case and Shiller (1989). In our analysis we use the Aggregate index, Apartment, Office, Industrial, and Retail property type indices, as well as the Southern California

⁹ Crane and Hartzell (2010) find evidence that a disposition effect exists for REITs with REIT managers more likely to sell properties that have performed well. Fisher et al. (2004) similarly find that properties that have outperformed a national commercial real estate index have a greater likelihood of being sold. In a related study, Bokhari and Geltner (2011) find evidence that commercial property sellers exhibit loss aversion behavior, e. g., ask higher prices than otherwise similar sellers who do not face a loss.

¹⁰ Data we obtained from <http://web.mit.edu/cre/research/credl/rca.html>. The indices are now located at http://www.rcanalytics.com/Public/rca_indices.aspx

Apartments index. As can be seen in the figure, the indices capture the last real estate cycle including the global financial crisis. We see an expansion in all property types with prices rising and reaching a peak in late 2008 followed by the subsequent collapse in prices during the global financial crisis.

Portfolio Analysis

Creating portfolio returns from property data poses some obvious limitations that one would not encounter using stock data. We do not continuously observe prices, so we are unable to form portfolios in a continuous fashion or to easily form portfolios that are rebalanced. However, we can make progress examining buy and hold portfolios. To do this we create buy and hold portfolio windows based on the purchase and sale quarter of the properties. That is, all properties purchased in Q1 2001 and sold in Q4 2006 would form one window. Properties purchased in Q1 2001 and sold in Q1 2007 would form another window. In our sample we have 342 such windows that are greater than one year in length and also contain more than 20 property repeat sales. The smallest window in our sample is 4 quarters in length and the longest is 28 quarters. The maximum number of properties we observe in any given window is 87.

By using quarters to define when properties are bought and sold, we are implicitly assuming that although the property may be purchased during the quarter, we don't start indexing the portfolio until the end of the quarter in which all the properties were purchased and similarly we index the portfolio until the end of the quarter in which the properties are sold. There are obvious limitations to forming portfolios in this way. Results may be dependent on the particular properties that are in a given window or the particular calendar period the portfolio was held for. By averaging across many windows and selecting windows that include many properties, we can mitigate these property specific and window specific concerns. It is thus unlikely that our results are driven by any single calendar period or subsample of properties examined.

Based on the buy and hold windows defined above, it is obvious that we can examine the two extremes of either individual properties or a portfolio containing all the properties in the window. Given that each window will differ in the number of properties it contains, making inferences based on just these two portfolios per window would be problematic. To provide additional insights into the relationship between the index and portfolios containing different numbers of properties, we will need to employ a random sampling scheme.

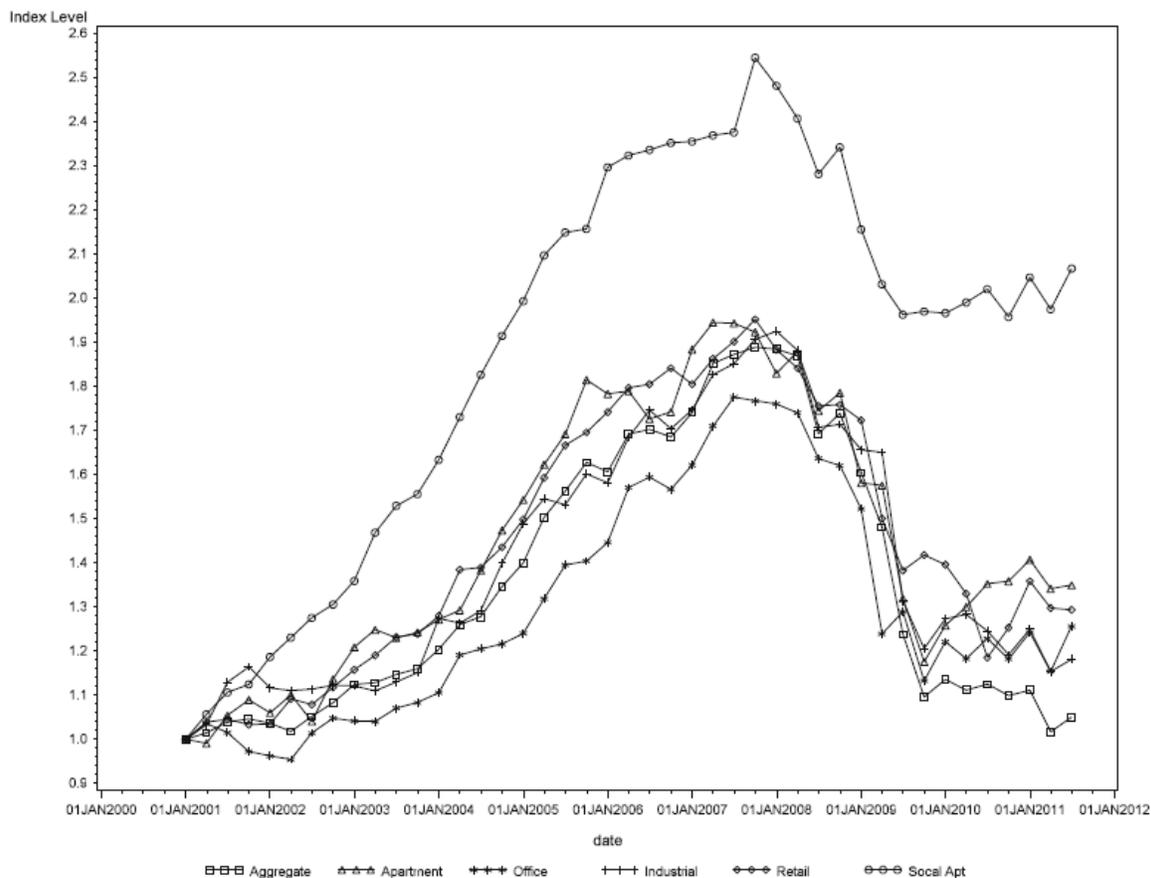


Fig. 1 Moody's REAL CPPI. Figure reports the index levels for the Moody's REAL Commercial Property Price Indices (CPPI). Indices are available from Q4 2000 to Q2 2011. Aggregate is the aggregate index, while Apartment, Office, Industrial, and Retail are property type specific indices. SoCal Apt is the index for apartments in Los Angeles and San Diego. Data are available from: http://www.rcanalytics.com/Public/rca_indices.aspx

We define k as being the number of properties in a given portfolio. When $k=1$, we are dealing with individual assets, when $k=10$ there are 10 properties in a given portfolio. Ideally we want to examine portfolios over a wide range of values of k in order to obtain greater insight into the benefits of diversification. We are restricted in doing this by the number of properties in a given portfolio window, n . That is, we cannot generate portfolios of 40 properties in a window that contains only 30 properties without resampling.¹¹ In this sense we are constrained by the repeat sales data that we have. Using quarters to define when properties are bought and sold, the largest number of transactions we observe in any given window is 87. Although a limitation, we can still examine fairly large portfolios using this sample. To relax this constraint we will allow for semi-annual purchase and sale windows in the later analysis.

¹¹ There are $n!/k!(n-k)!$ random portfolios containing k properties that can be drawn from a window containing n total properties without repetition.

Index by j the number of buy and hold windows in the sample. For each of these windows, define n as the number of property repeat sales we observe in window j . As above, k is defined as the number of properties in a given portfolio with the restriction that $k < n - 10$.¹² For each window j , and for each feasible value of k , we generate 5,000 random portfolios of k properties from the n available properties in window j .¹³

For each of the 5,000 portfolios in each of the (j, k) windows, we calculate the equally weighted portfolio return using the annualized holding period returns. Equally weighted portfolio returns are simply the average of the annualized holding period returns of the properties in the portfolio.

Having calculated the 5,000 equally weighted portfolio returns for each (j, k) window, we match these portfolio returns to the annualized holding period return on the Moody's REAL index over the same calendar window. The index return thus represents the return implied by buying the index at the end of the quarter in which the properties were purchased and selling the index at the end of the quarter in which the properties were sold. Thus it represents the natural benchmark that investors would use for either benchmarking or hedging their portfolios.

To examine how well the portfolio returns track the index, we calculate the root mean squared deviation RMSD between the 5000 portfolios and the index for each (j, k) window. The RMSD is calculated as:

$$RMSD = \sqrt{\frac{1}{5000} \sum_{i=1}^{5000} (R_i^p - R^I)^2},$$

where R_i^p is the annualized return on portfolio i and R^I is the annualized return on the matched index in the $(j-k)$ window.

For each value of k , we then average across all the available windows to calculate the average RMSD. Given that we are averaging across many different windows and thousands of random portfolios within each window, the average RMSD represents the average tracking error that an investor is likely to face on average when holding a random portfolio of k properties.

¹² This restriction is imposed so that we aren't repeatedly resampling the same portfolio of k assets. This is because as k approaches n , the number of unique portfolios declines.

¹³ The value of 5,000 was chosen because it is computationally feasible to generate this many portfolios repeatedly; see, e.g., Fama and French (2010). Our results do not change significantly in the aggregate sample if we increase this number to 100,000.

Results

Figure 2 plots the average RMSD across the 342 portfolio windows (All), for windows less than or equal to 8 quarters (Short) and greater than 8 quarters (Long) in length, respectively. For an average RMSD to be reported for a given number of properties, k , we require at least 15 available windows for that value of k . In this sense, as k increases, the number of windows available declines. We impose this restriction so that the results for any given k are unlikely to be driven by any one calendar window analyzed.

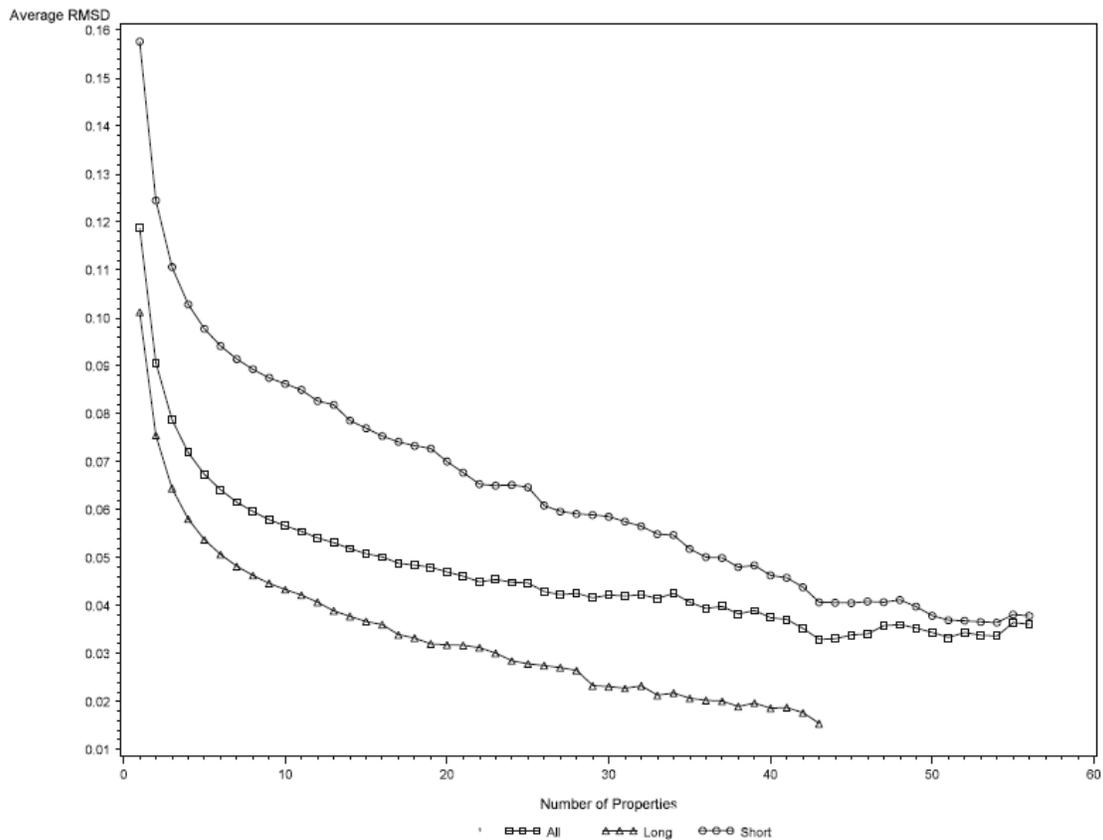


Fig. 2 Aggregate sample. Figure reports the average root mean squared deviation (Average RMSD) between the return on the aggregate Moody's CPPI and random portfolios containing different numbers of properties, k . The average is taken over 342 calendar windows between Q4 2000 and Q2 2011 that are greater than 1 year in length and that contain at least 20 property repeat sales. Each window contains all the properties whose buy and sell dates match the window's holding period. In each window 5000 portfolios containing k properties were formed from the n available repeat sales, and the RMSD calculated for each window for each level of k . The reported Average RMSD is the average taken across windows, where at least 15 windows are available for that value of k . Underlying data are 12,427 repeat sales occurring between 2000 and 2011 in the CoStar database. Office, Industrial, Multifamily, and Retail property types were included and properties had to be larger than 2500 sq.ft. or have more than 10 units for multifamily and have a price greater than \$1 m. All properties must have been held for longer than 1 year and properties with annual appreciation of less than -40 % or greater than 50 % were excluded. All transaction had to be arm's length and non-distressed

A few points are evident from Fig. 2. First, individual assets are not well matched by the index. Average RMSD is high, suggesting that there is substantial dispersion in price appreciation at the individual property level. Second, as the number of properties in the portfolio increases, the average RMSD declines quite rapidly. At 10 properties the average RMSD is more than halved. Suggesting the index matching even for small portfolios shows a marked improvement over individual assets. The average RMSD declines over all the values of k for which we are able to form portfolios at the quarterly level, suggesting that the idiosyncratic risk is reduced slowly.

Third, the average RMSD is very different across long and short windows. This is consistent with the results shown in Table 1. For any given level of diversification, the average RMSD for the short windows is much larger than the average RMSD of the long windows. This suggests that indexing unstabilized or opportunistic properties will be difficult with currently available indices. Notice that we are not suggesting that the length of the holding period drives the ability to index. Rather, the length of the ex post holding period is correlated to the opportunistic nature of the investment and this factor drives the inability to index effectively.

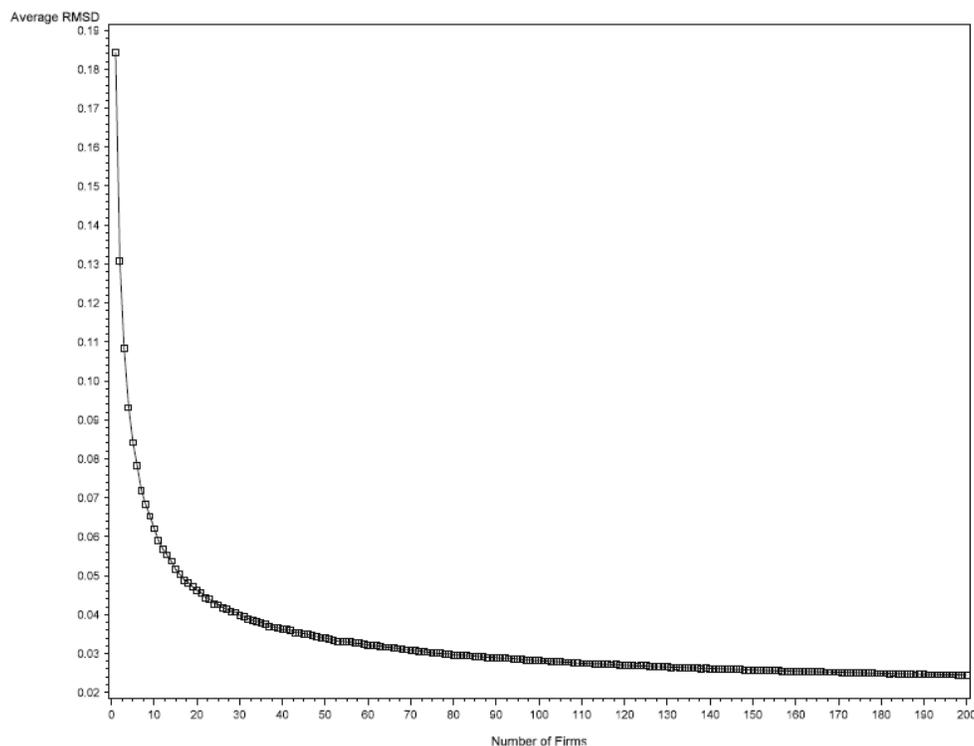


Fig. 3 S&P 500 constituents. Figure reports the average root mean squared deviation (Average RMSD) between the return on an equally weighted index containing S&P 500 constituent companies and portfolios containing different numbers of S&P 500 constituent stocks, k . The average is taken over the same 342 calendar windows from Fig. 1. In each window 5000 portfolios containing k stocks were formed and the RMSD calculated for each window for each level of k . The reported Average RMSD is the average taken across windows, where at least 15 windows are available for that value of k . Price and index constituent data were obtained from CRSP

Comparison with Equities

The results in Fig. 2 suggest that portfolio indexing is feasible for moderately sized portfolios. We do not however, have a clear benchmark to evaluate the index's tracking ability. To achieve this, we repeat the analysis above using the constituents of the S&P 500 and the return on an equally weighted index created from those constituent stocks.¹⁴ In order to make the analysis comparable to Fig. 2, we conduct this experiment using exactly the same holding period windows used in the previous analysis.

Figure 3 reports the results from our analysis of the S&P 500 constituents. A few results are evident. First, in a similar fashion to Fig. 2, there is high variation at the individual stock level. The average RMSD is higher than for the property sample. Second, the average RMSD declines rapidly once we start forming portfolios.

Nonetheless it takes large portfolios for the average RMSD to remain fairly constant as we increase the number of stocks in the portfolio.¹⁵ Finally, the level of average RMSD for diversified portfolios appears quite similar between the stock and property sample. This suggests that at least for the aggregate sample examined thus far, indexing a large portfolio of stabilized properties using an aggregate repeat sales index is likely to be as effective as indexing a large portfolio of stocks using an aggregate stock index.¹⁶ While we will elaborate further on the individual results below, Table 2 summarizes the reduction in average RMSD in our basic seven data sets. Using the S&P500 as our benchmark, we can see that a randomly selected portfolio containing 10 stocks reduces the average RMSD to a third of the initial (one stock) level; a portfolio containing 20 randomly selected stocks bring the RMSD to a quarter of its initial level. Column 3 shows the reduction for the aggregate sample. Here portfolios of 10 and 20 properties reduce the RMSD to 48 % and 40 % of the initial level, respectively. For the Apartment, Office, Industrial and Retail sample the diversification effects are quite similar. Portfolios of 10 and 20 properties on average reduce the RMSD to 41 % and 35 % of the initial level, respectively. These figures indicate that the tracking is marginally better when using property types. The last column shows the results for Southern California Apartments. Here the results are stronger and are

¹⁴ We use an equally weighted index because the repeat sales indices are equally weighted.

¹⁵ To provide a different context for this result, if we calculate the average R2 from the monthly time series regression of portfolio returns on index returns for each of the buy and hold windows, we observe that the average R2 starts at around 25 % for individual stocks and increases to 85 % at 40 stocks and reaches 92 % when we increase the number of stocks to 200.

¹⁶ This tracking error would be reduced if we made the selection of stocks sector neutral, just as we will see in the better tracking results when matching location in Table 2.

comparable to the results obtained for the S&P500: portfolios of 10 and 20 properties on average reduce the RMSD to 35 % and 27 % of the initial level, respectively.

Value Weighted Portfolios

Notice that our results only holds for equally weighted portfolios. If we consider value weighted portfolios the analysis becomes much more problematic. The reason for this is that if we have assets of very different sizes, then in creating value weighted portfolios we are effectively creating portfolios that are dominated by individual assets. Although the portfolio may contain many assets, the return of the portfolio can behave like the return of an individual asset. Furthermore our analysis in both Figs. 2 and 3 shows that individual assets are poorly matched by an index on average. Given that most practitioners would be interested in indexing a value weighted portfolio because that is the portfolios they hold, our results require some caution. In order for the index to match a value weighted portfolio, the value weighted portfolio needs to approximate an equally weighted portfolio. That is, you would need to create portfolios of many properties that have properties of roughly equal value. Indexing portfolios dominated by an individual asset would still be, on average, difficult.

The results in Fig. 2 suggest that we extend our analysis in two directions. First, the declining nature of the average RMSD as the number of properties increases leads us to consider if we could do better with larger portfolios. Second, although at the aggregate level property portfolio indexing appears to be as effective as stock indexing, can we do better by indexing with a property type matched index or even a property type and location matched index?

Table 2 Summary of figures. Table presents the average root mean squared deviation plotted in Figs. 3, 4, 5, 6, 7, 8 and 9 for portfolios containing various numbers of constituent assets. Figure 3 represents the S&P 500 constituents sample, Fig. 4 is the Aggregate sample using semi-annual windows, Fig. 5 is the Apartment sample, Fig. 6 is the Office sample, Fig. 7 is the Industrial sample, Fig. 8 us the Retail sample, and Fig. 9 is the Southern Californian Apartment sample

Properties	Figure 3 SP500	Figure 4 Aggregate	Figure 5 Apartment	Figure 6 Office	Figure 7 Industrial	Figure 8 Retail	Figure 9 SoCal Apt
1	0.184	0.103	0.096	0.110	0.105	0.112	0.079
5	0.084	0.052	0.048	0.054	0.059	0.058	0.038
10	0.062	0.041	0.038	0.041	0.049	0.046	0.028
15	0.052	0.036	0.033	0.036	0.048	0.038	0.024
20	0.046	0.034	0.031	0.033	0.048	0.035	0.021
25	0.043	0.032	0.029	0.031	0.049	0.034	
30	0.040	0.030	0.027	0.029		0.034	
35	0.038	0.029	0.026			0.029	
40	0.036	0.029	0.026			0.029	
50	0.034	0.028	0.025				
75	0.030	0.026					
100	0.028	0.024					

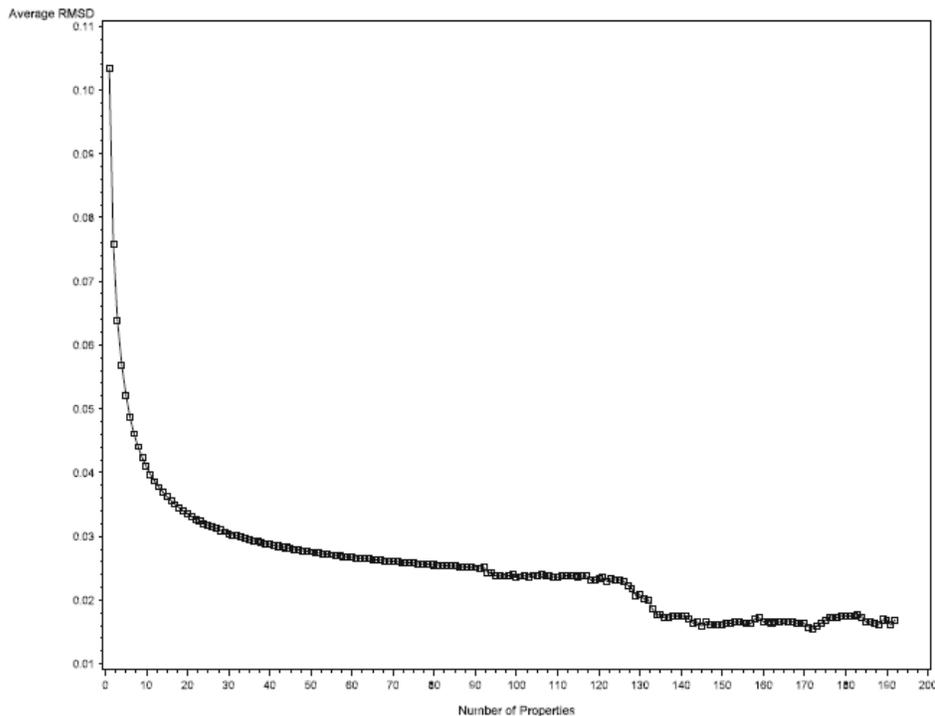


Fig. 4 Aggregate index semi-annual windows. Figure reports the average root mean squared deviation (Average RMSD) between the return on the aggregate Moody's CPPI and random portfolios containing different numbers of properties, k . The average is taken over 170 windows between Q4 2000 and Q2 2011 that are greater than 2 years in length and that contain at least 100 property repeat sales. Each window contains all the properties that were bought and sold in the same 6 month periods. In each window 5000 portfolios containing k properties were formed from the n available repeat sales, and the RMSD calculated for each window for each level of k . The reported Average RMSD is the average taken across windows, where at least 15 windows are available for that value of k . Underlying data are 12,427 repeat sales occurring between 2000 and 2011 in the CoStar database. Office, Industrial, Multifamily, and Retail property types were included and properties had to be larger than 2500 sq.ft. or have more than 10 units for multifamily and have a price greater than \$1 m. All properties must have been held for longer than 1 year and properties with annual appreciation of less than -40 % or greater than 60 % were excluded. All transaction had to be arm's length and non-distressed

Semi-Annual Windows

To examine these issues we modify the analysis slightly. In order to have more properties in each buy and hold window, we will need to expand the purchase and sale periods from one quarter to two quarters. That is, windows will be formed from all properties purchased in a given two quarter period (say Q1 and Q2 2001) and sold in a given two quarter period (say Q2 and Q3 2005). Once again this implies that we buy the index at the end of the purchase period and sell the index at the end of the sale period. We then repeat the analysis in exactly the same manner as used in Fig. 2.

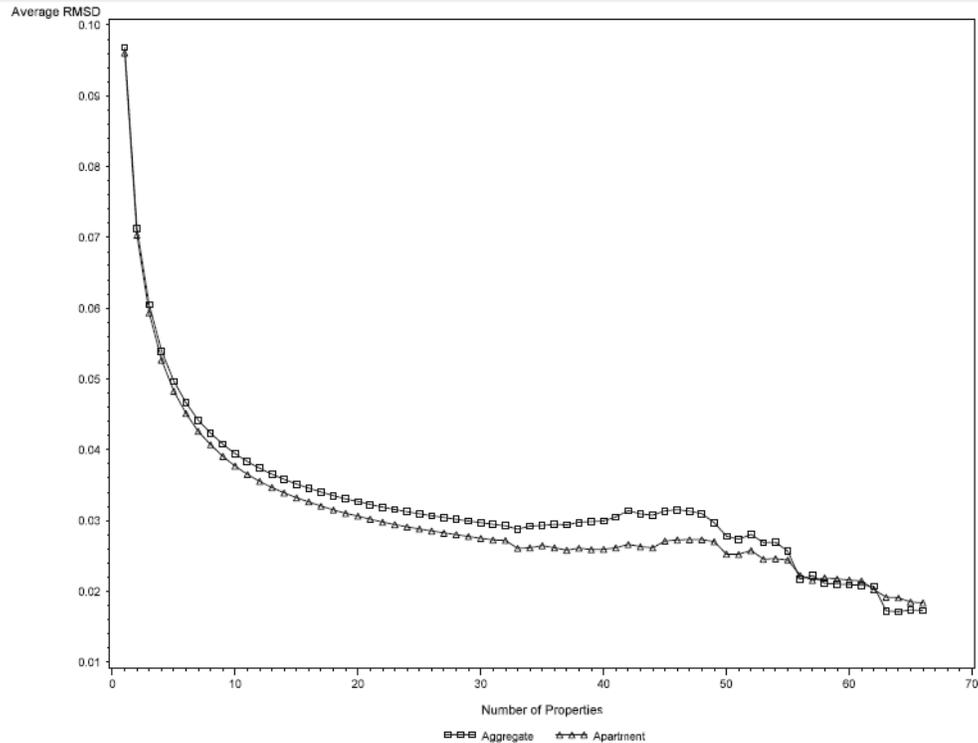


Fig. 5 Apartments. Figure reports the average root mean squared deviation (Average RMSD) between the return on the aggregate Moody’s CPPI or Moody’s apartment CPPI and random portfolios containing different numbers of apartment properties, k . The average is taken over 116 windows between Q4 2000 and Q2 2011 that are greater than 2 years in length and that contain at least 20 property repeat sales. Each window contains all the properties that were bought and sold in the same 6 month periods. In each window 5000 portfolios containing k properties were formed from the n available repeat sales, and the RMSD calculated for each window for each level of k . The reported Average RMSD is the average taken across windows, where at least 15 windows are available for that value of k . Underlying data are 4601 apartment repeat sales occurring between 2000 and 2011 in the CoStar database. Properties had to have more than 10 units and a price greater than \$1 m. All properties must have been held for longer than 1 year and properties with annual appreciation of less than -40% or greater than 60% were excluded. All transaction had to be arm’s length and non-distressed

In the analysis that follows we will focus only on windows that are greater than 8 quarters in length, because our previous analysis suggests that these are more likely to represent the stabilized properties and portfolios that investors may be trying to index. Figure 4 reports the average RMSD for the semi-annual windows. Although somewhat noisier than the results from Fig. 3 for stocks, Fig. 4 demonstrates that the average RMSD continues to decline even with very diversified portfolios. While even small portfolios are far better than single assets when it comes to matching the index, highly diversified portfolios of over 150 properties are significantly better than portfolios of 20 properties. The second interesting result from Fig. 4 is that the average RMSD for properties tends to lie below that for stocks. In this sense it appears we may do a better job indexing properties relative to stocks.

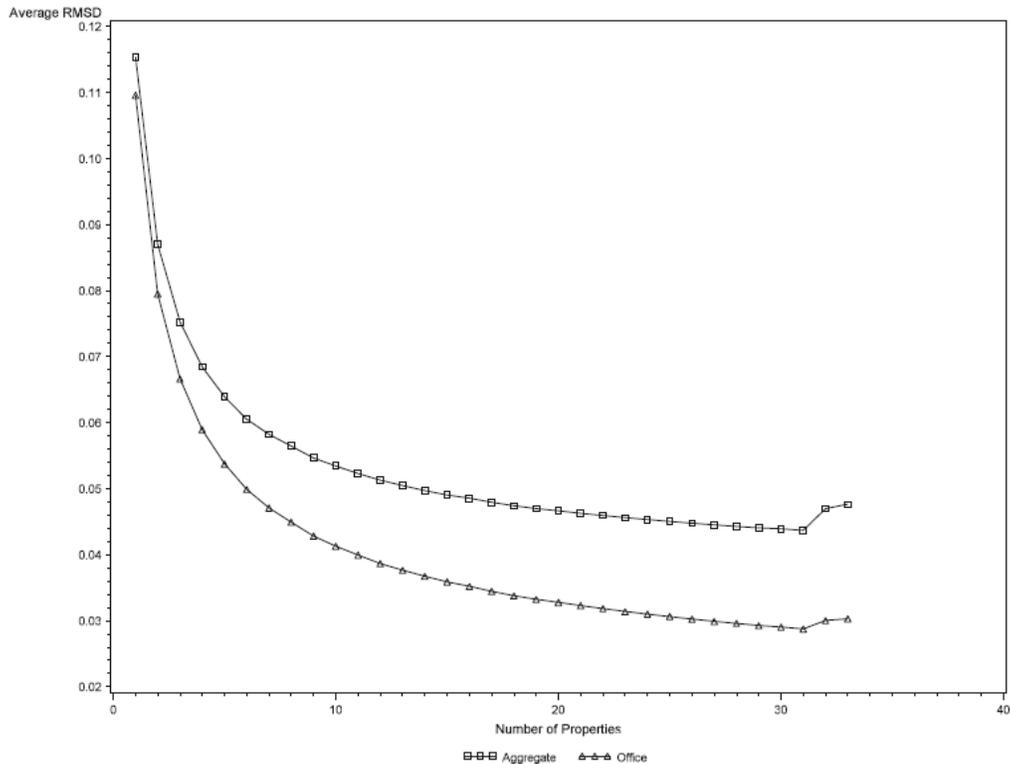


Fig. 6 Office. Figure reports the average root mean squared deviation (Average RMSD) between the return on the aggregate Moody’s CPPI or Moody’s office CPPI and random portfolios containing different numbers of office properties, k . The average is taken over 21 windows between Q4 2000 and Q2 2011 that are greater than 2 years in length and that contain at least 20 property repeat sales. Each window contains all the properties that were bought and sold in the same 6 month periods. In each window 5000 portfolios containing k properties were formed from the n available repeat sales, and the RMSD calculated for each window for each level of k . The reported Average RMSD is the average taken across windows, where at least 15 windows are available for that value of k . Underlying data are 2494 office repeat sales occurring between 2000 and 2011 in the CoStar database. Properties had to be larger than 2500 square feet and a price greater than \$1 m. All properties must have been held for longer than 1 year and properties with annual appreciation of less than -40 % or greater than 60 % were excluded. All transaction had to be arm’s length and non-distressed

Property Type Indices

So far we have focused on aggregate indices. Given the focused nature of real estate investment, it is unlikely that an investor is trying to index a completely diversified portfolio. That is, they are unlikely to hold a portfolio of multiple property types in multiple locations. The simplest non-random portfolios that an investor is likely to hold are specific property type portfolios. To examine the ability to index these portfolios, we repeat the analysis from Fig. 4 using property type subsamples matched to specific property type indices. As our benchmark indices we use the Apartment, Office, Retail, and Industrial property type indices from Moody’s REAL.

Figures 5, 6, 7 and 8 report the average RMSD for Apartment, Office, Industrial, and Retail subsamples, respectively. Each figure reports the average RMSD when the index is the property type

matched index (denoted by triangles) and also the average RMSD if the aggregate index (denoted by squares) is used. In each case we only consider holding windows greater than 8 quarters in length and require 15 windows to be available to form the average.

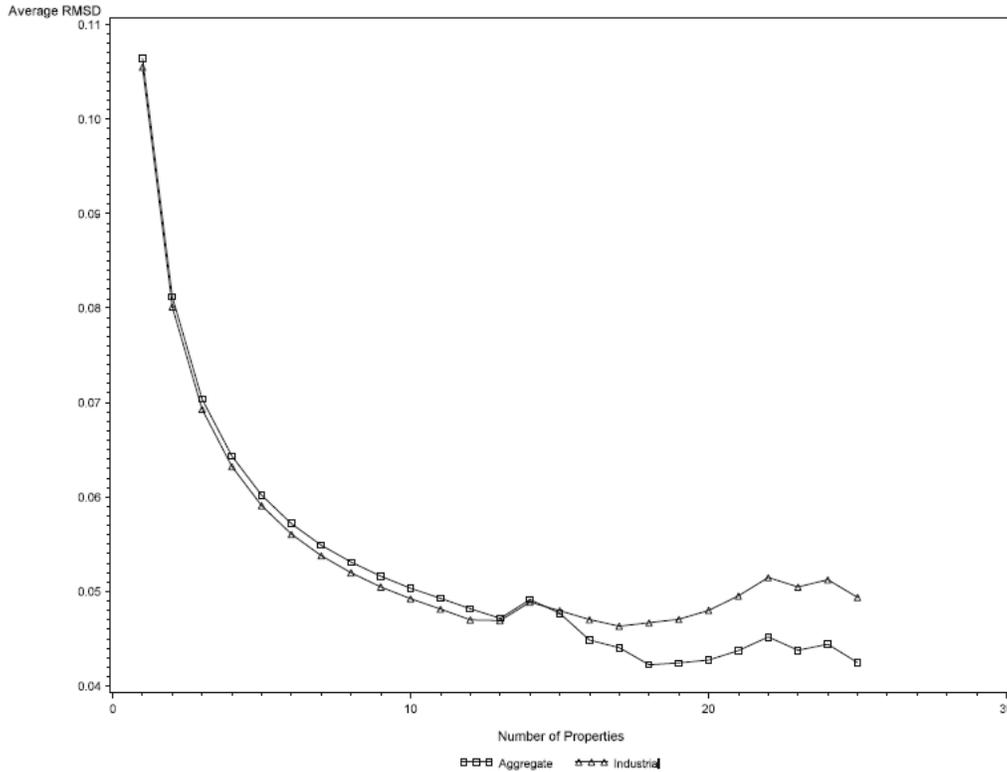


Fig. 7 Industrial. Figure reports the average root mean squared deviation (Average RMSD) between the return on the aggregate Moody’s CPPI or Moody’s industrial CPPI and random portfolios containing different numbers of industrial properties, k . The average is taken over 155 windows between Q4 2000 and Q2 2011 that are greater than 2 years in length and that contain at least 20 property repeat sales. Each window contains all the properties that were bought and sold in the same 6 month periods. In each window 5000 portfolios containing k properties were formed from the n available repeat sales, and the RMSD calculated for each window for each level of k . The reported Average RMSD is the average taken across windows, where at least 15 windows are available for that value of k . Underlying data are 2205 industrial repeat sales occurring between 2000 and 2011 in the CoStar database. Properties had to be larger than 2500 square feet and a price greater than \$1 m. All properties must have been held for longer than 1 year and properties with annual appreciation of less than -40 % or greater than 60 % were excluded. All transaction had to be arm’s length and non-distressed

The most striking result from Figs. 5, 6, 7 and 8, is that, apart from the office subsample in Fig. 6, using a property matched index does not appear to improve our ability to index specific property type portfolios over the aggregate index. This observation is consistent with the data presented in Table 2. In each case we observe that diversified portfolios are better matched by the index than individual assets or concentrated portfolios. However, it tends not to be the case that using a property type matched index significantly outperforms using the aggregate index. For example, with portfolios containing 25 properties, the aggregate index reduces the RMSD to 32 % of its initial value; for the 4 property type

indices, the reductions are 30 %, 28 %, 47 % and 30 %, respectively. A possible explanation of this is that because we are employing buy and hold returns over long periods of time, property types tend to look like the aggregate index. This is evident from the cyclical nature of the CPPI indices in Fig. 1.

Consequently, although property type performance may vary significantly over shorter periods of time, they tend to be more similar over a 3 or 4 years period.

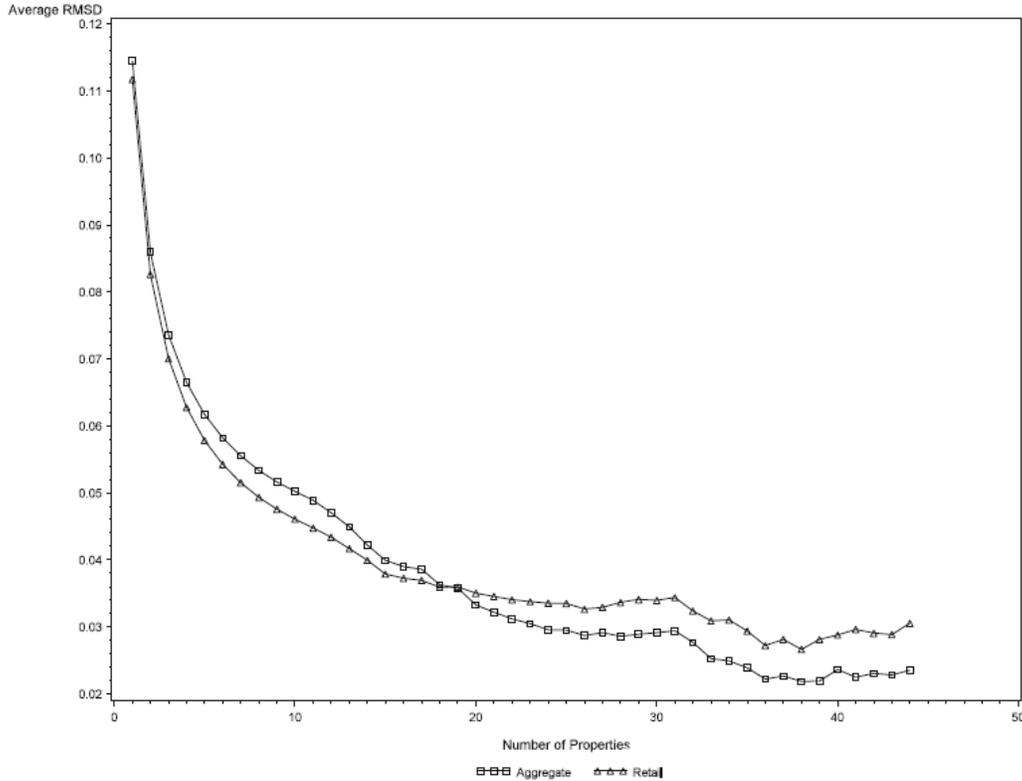


Fig. 8 Retail. Figure reports the average root mean squared deviation (Average RMSD) between the return on the aggregate Moody's CPPI or Moody's retail CPPI and random portfolios containing different numbers of retail properties, k . The average is taken over 212 windows between Q4 2000 and Q2 2011 that are greater than 2 years in length and that contain at least 20 property repeat sales. Each window contains all the properties that were bought and sold in the same 6 month periods. In each window 5000 portfolios containing k properties were formed from the n available repeat sales, and the RMSD calculated for each window for each level of k . The reported Average RMSD is the average taken across windows, where at least 15 windows are available for that value of k . Underlying data are 3127 retail repeat sales occurring between 2000 and 2011 in the CoStar database. Properties had to be larger than 2500 square feet and a price greater than \$1 m. All properties must have been held for longer than 1 year and properties with annual appreciation of less than -40 % or greater than 60 % were excluded. All transaction had to be arm's length and non-distressed

Southern California Apartments

If matching by property type does not tend to provide important benefits over using an aggregate index, the question still remains if using an even more focused index does. Intuitively one may expect this to be the case. While the property type level indices may appear similar over long periods of

time, if one examines a property type in a given MSA there is reason to suspect that differences in performance could persist over long periods of time. Examining this issue becomes difficult due to data limitations. In order to form diversified portfolios we need to observe a large number of repeat sales of the same property type in a given location through time. This necessitates examining very liquid real estate markets. A similar problem arises with constructing the underlying index—a lot of repeat sale transactions are required to be able to estimate a repeat sales model at the property type MSA level. We are able to examine one market using our data: the Southern California Apartment market. Moody's REAL reports an apartment index for the San Diego and Los Angeles MSAs. Based on this market and property type we are able to find sufficient transactions over enough holding period windows to employ our portfolio sampling strategy.

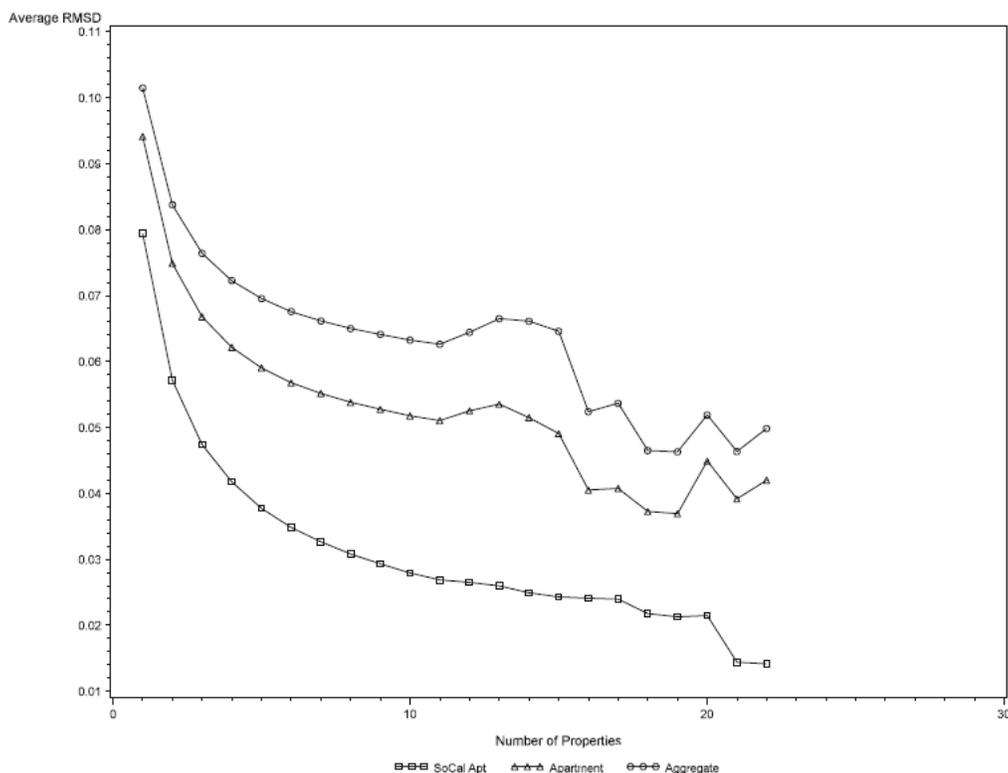


Fig. 9 Southern California apartments. Figure reports the average root mean squared deviation (Average RMSD) between the return on the aggregate Moody's CPPI, Moody's apartment CPPI, or Moody's SoCal Apt CPPI and random portfolios containing different numbers of Los Angeles and San Diego apartment properties, k . The average is taken over 32 windows between Q4 2000 and Q2 2011 that are greater than 2 years in length and that contain at least 20 property repeat sales. Each window contains all the properties that were bought and sold in the same 6 month periods. In each window 5000 portfolios containing k properties were formed from the n available repeat sales, and the RMSD calculated for each window for each level of k . The reported Average RMSD is the average taken across windows, where at least 15 windows are available for that value of k . Underlying data are 1596 Los Angeles and San Diego apartment repeat sales occurring between 2000 and 2011 in the CoStar database. Properties had to be larger than 2500 square feet and a price greater than \$1 m. All properties must have been held for longer than 1 year and properties with annual appreciation of less than -40% or greater than 60% were excluded. All transaction had to be arm's length and non-distressed

Figure 9 reports the average RMSD for the SoCal Apartment market. We once again only examine windows longer than 8 quarters. The figure reports average RMSD for three different benchmark indices: 1) Aggregate—the aggregate index (denoted by a circle); 2) Apartment—the apartment index (denoted by a triangle); and 3) SoCal Apt—the Southern California apartment index (denoted by a square). The most striking result from Fig. 9 is the difference in average RMSD between the SoCal Apartment index and the property type and aggregate indices. The match using a specific property type and location index appears significantly better than using either a property type or aggregate index. Even with a portfolio of 20 properties, on an absolute basis, the average RMSD is below what we observe for the S&P500. On a percentage basis, the reduction in average RMSD is 75.0 % for the S&P500 and 73.4 % for the SoCal index. Matching at this more granular level appears to decrease the average RMSD at a faster rate and also results in a lower total level of average RMSD than observed in the previous figures and Table 2. Although we are able to perform this experiment in only one market, the striking match between the index and the portfolio suggests that indexing at this level may be the most appropriate exercise.

Conclusions

This study examines the relative performance of portfolios of real returns using the Moody's REAL CPPI indices. Our real returns are generated from repeat sales in the CoStar data base between 2000 and 2011. We find that the aggregate real estate indices do a good job of tracking real returns when portfolios of more than 20 properties are considered. At this level, risk reduction is somewhat less effective than our benchmark of the S&P500 and its component stocks. Compared to the average RMSD from one asset, randomly selected portfolios with 20 assets reduce the average RMSD by 75 % for the S&P500 compared to 66 % for the aggregate index. This result suggests that the aggregate CPPI indices can be effective in hedging direct real estate investment and for performance measurement. We also find that tracking at the property type level provides little benefit over using an aggregate index. Using a smaller sample of Southern California apartment returns, we provide some evidence that indexing using a property type and location matched index provides lower tracking error for any level of diversification.

We also find that returns on properties with shorter holding periods are much higher than average; specifically, properties with a holding period less than 8 quarters have an average annual return of 17.6 % compared to 10.5 % for properties with longer holding periods.

References

- Bailey, M., Muth, R., & Nourse, H. (1963). A regression method for real estate price index construction. *Journal of the American Statistical Association*, 58, 933–942.
- Barkham, R., & Geltner, D. (1995). Price discovery in American and British property markets. *Real Estate Economics (formerly AREUEA Journal)*, 23(1), 21–44.
- Bokhari, S., & Geltner, D. (2011). Loss aversion and anchoring in commercial real estate pricing: Empirical evidence and price index implications. *Real Estate Economics*, 39(4), 635–670.
- Boudry, W. I., Coulson, N.E., Kallberg, J.G. and Liu, C.H. (2012). On the Hybrid Nature of REITs. *Journal of Real Estate Finance and Economics*, forthcoming.
- Brown, G., & Kapadia, N. (2007). Firm specific risk and equity market development. *Journal of Financial Economics*, 2(84), 358–388.
- Byrne, P., & Lee, S. L. (2003). An exploration of the relationship between size, diversification and risk in UK real estate portfolio: 1989–1999. *Journal of Property Research*, 20(2), 191–206.
- Campbell, J., Lettau, M., Malkiel, B., & Xu, Y. (2001). Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *Journal of Finance*, 56, 1–43.
- Case, K., & Shiller, R. (1989). The efficiency of the market for single-family homes. *The American Economic Review*, 79, 125–137.
- Clapp, J., & Giaccotto, C. (1992). Estimating price trends for residential property: A comparison of repeat sales and assessed value methods. *Journal of Real Estate Finance and Economics*, 5, 357–374.
- Collett, D., Lizieri, C., & Ward, C. (2003). Timing and the holding periods of institutional real estate. *Real Estate Economics*, 31(2), 205–222.
- Crane, A. and Hartzell, J. (2010). Is There a Disposition Effect in Corporate Investment Decisions? Evidence from Real Estate Investment Trusts. Working paper, University of Texas
- Dombrow, J., Knight, J., & Sirmans, C. F. (1997). Aggregation bias in repeat-sales indices. *Journal of Real Estate Finance and Economics*, 14, 75–88.
- Elton, E., & Gruber, M. (1977). Portfolio risk reduction: An analytical solution. *Journal of Business*, 50(4), 415–437.
- Evans, J. L., & Archer, S. H. (1968). Diversification and the reduction of dispersion: An empirical analysis. *Journal of Finance*, 23(5), 761–767.
- Fama, E., & French, K. (2010). Luck versus skill in the cross-section of mutual fund returns. *Journal of Finance*, 65, 1915–1947.

- Feng, T., & Geltner, D. (2011). Property-level performance attribution: Investment management diagnostics and the investment importance of property management. *Journal of Portfolio Management*, 37(5), 110–124.
- Fisher, J., Gatzlaff, D., Geltner, D., & Haurin, D. (2003). Controlling for the impact of variable liquidity in commercial real estate price indices. *Real Estate Economics*, 31(2), 269–303.
- Fisher, J., Gatzlaff, D., Geltner, D., & Haurin, D. (2004). An analysis of the determinants of transaction frequency of institutional commercial real estate investment property. *Real Estate Economics*, 32(2), 239–264.
- Fisher, J., Geltner, D., & Pollakowski, H. (2007). A quarterly transactions-based index of institutional real estate investment performance and movements in supply and demand. *Journal of Real Estate Finance and Economics*, 34, 5–33.
- Fu, F. (2009). Idiosyncratic risk and the cross-section of expected stock returns. *Journal of Financial Economics*, 91, 24–37.
- Gatzlaff, D., & Geltner, D. (1998). A transaction-based index of commercial property and its comparison to the NCREIF index. *Real Estate Finance*, 15(1), 7–22.
- Gatzlaff, D., and Haurin, D. (1994). Sample Selection and Biases in Local House Value Indices. Working paper, The Ohio State University.
- Gatzlaff, D., & Haurin, D. (1997). Sample selection bias and repeat-sales index estimates. *Journal of Real Estate Finance and Economics*, 14, 33–50.
- Geltner, D., & Goetzmann, W. (2000). Two decades of commercial property returns: A repeated-measures regression-based version of the NCREIF index. *Journal of Real Estate Finance and Economics*, 21(1), 5–21.
- Geltner, D., & Kluger, B. D. (1998). REIT-based pure-play portfolios: The case of property types. *Real Estate Economics*, 26(4), 581–612.
- Goetzman, W., & Peng, L. (2006). Estimating house price indices in the presence of seller reservation prices. *The Review of Economics and Statistics*, 88, 100–112.
- Gyourko, J., & Keim, D. B. (1992). What does the stock market tell Us about real estate returns. *Real Estate Economics*, 20(3), 457–485.
- Goodman, A. C., & Thibodeau, T. G. (2003). Housing market segmentation and hedonic prediction accuracy. *Journal of Housing Economics*, 12, 181–201.
- Hoag, J. W. (1980). Towards indices of real estate value and return. *Journal of Finance*, 35(2), 569–580.

- Horrigan, H., Case, B., Geltner, D., & Pollakowski, H. O. (2009). REIT-based property return indices: A New Way to track and trade commercial real estate. *Journal of Portfolio Management*, 35(5), 80–91.
- Kallberg, J. G., Liu, C. H., & Greig, D.W. (1996). The role of real estate in the portfolio allocation process. *Real Estate Economics*, 24(3), 359–377.
- Liu, C. H., & Mei, J. (1992). The Predictability of Returns on Equity REITs and Their Co-movement with Other Assets. *Journal of Real Estate Finance and Economics*, 5, 401–418.
- MacKinnon, G. H., & Zaman, A. A. (2009). Real estate for the long term: The effect of return predictability on long-horizon allocations. *Real Estate Economics*, 37(1), 117–153.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- Miles, M. E., Cole, R. E., & Guilkey, D. (1990). A different look at commercial real estate returns. *Real Estate Economics*, 18(4), 403–430.
- Riddiough, T. J., Moriarty, M., & Yeatman, P. J. (2005). Privately versus publicly held asset investment performance. *Real Estate Economics*, 33(1), 121–146.
- Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82, 34–55.
- Sharpe, W. (1964). Capital asset prices. *Journal of Finance*, 19(3), 425–442.
- Statman, M. (1987). How many stocks make a diversified portfolio? *Journal of Financial and Quantitative Analysis*, 22(3), 353–363.
- Webb, R. B., Miles, M. E., & Guilkey, D. (1992). Transactions-driven commercial real estate returns: The panacea to asset allocation models? *Real Estate Economics*, 20(2), 325–357.