Performance of Real Estate Portfolios

Jeffrey D. Fisher and William N. Goetzmann

A bootstrapping approach.

Portfolio managers rely heavily on historical information to make asset allocation decisions. The empirical application of neoclassical financial models—from portfolio optimization to value at risk—typically requires as inputs the time series of returns to the asset class. This approach presupposes the ability to invest in a portfolio of assets that mimics the performance of the index.

For some very liquid assets such as large-cap U.S. equities and government bonds, this is possible; one can purchase a diversified pool of stocks or bonds, rebalance or reinvest the proceeds as needed, and immediately liquidate the position at the end of a fixed time period. For other asset classes, the performance of a standard capital-weighted index may be a poor approximation of the investment experience, and indeed may limit the applicability of standard neoclassical models.

As institutional managers increase their allocations to alternative asset classes, the use of indexes as proxies has become increasingly unrealistic. For example, the risk and return characteristics of venture capital and private equity investments are not reliably measurable from historical data, yet modern managers regard these asset classes as important sources of return and diversification.

Real estate investors have long been aware of the challenges of translating the returns of property investment into reliable time series data. Investors, institutions, and consultants have addressed these challenges by developing statistical risk and return inputs to allocation models through the construction of indexes that reflect broad trends in a

Jeffrey D. Fisher
is the Charles H. and Barbara F. Dunn professor of real estate at Indiana University in Bloomington, and consulting director of research at NCREIF.

William N. Goetzmann
is the Edwin J. Beinecke professor of finance and management studies, and director of the International Center for Finance, at the Yale School of Management in New Haven.

diversified portfolio of investable properties.

Here we develop an alternative approach to understanding the long-term performance of real estate investment. Rather than assuming that a real estate investor's experience matches an industry index, we simulate the experience of a commercial real estate investor from the last quarter of 1977 through second quarter 2004 using actual property histories to imitate the total return an active investor in this time period would have experienced. To do this, we use a large database of assets actually purchased and sold by institutional investors. The National Council of Real Estate Investment Fiduciaries (NCREIF) has pooled property-level investment data of its members for more than 25 years. This rich database is now the definitive foundation for commercial real estate index data.

The premise of our analysis is that real estate investors over this period invested in properties, not in an index. We thus measure performance by the internal rate of return (IRR) to simulated portfolios of commercial properties, U.S. stocks, and U.S. bonds. The stock and bond portions of the portfolio are rebalanced to accommodate the positive and negative cash flows required by real estate investing. This simulated IRR approach allows us to examine the cross-sectional distribution of real estate returns over one time period under a variety of assumptions. We vary the number of properties in the portfolio, the approximate portfolio allocation to real estate, and the regional and

properties. In other words, diversification requires substantial capital. This has important implications for an active manager's ability to meet an IRR-based benchmark. Real estate investors cannot own a portion of every property in the benchmark portfolio, and the greater the cross-sectional volatility of IRRs for a portfolio, the greater the risk they will not achieve the same IRR as the benchmark.

We are also able to quantify the return effects of regional and property type focus. Diversification across property types reduces cross-sectional variation more than geographic diversification.

PERFORMANCE MEASUREMENT

We first describe the two measures of performance.

Time-Weighted Rate of Return

The use of IRR as a measure of returns is somewhat unconventional in portfolio analysis. Much of the current statistical analysis of the performance of asset classes is based upon the time-weighted rate of return, which relies on the periodic measurement of value accruing to the investor as a result of holding an asset or a portfolio. The time-weighted rate of return is calculated assuming that both income and capital appreciation of an asset may be realized (or at least accurately measured) each period. These erty investment are lumpy. This is a well-known and well-studied issue. The use of IRR mitigates the fundamental problem of periodic return calculation by using actual cash flows including the cash flow from acquisition and disposition rather than appraised values.[1]

A second drawback of the time-weighted rate of return that is somewhat subtle is that it fails to account for investment timing. One striking example is the measurement of equity mutual fund investor returns over the last 30 years. Although time-weighted rates of return display a healthy annual performance, returns on a dollar-weighted basis tell a different story.

The huge inflows during the late 1990s represent investment by individuals who lost money in the crash of 2000. In an economic sense, dollar-weighted rates of return are a better measure of the economic scale of the crash in equities following the 1990s.

Internal Rate of Return

The internal rate of return is the compound periodic discount rate that sets the present value of an investment's cash flows to zero. The yield to maturity on a non-callable bond, for example, is essentially its IRR—it is the interest rate that sets the present value of future coupons and principal repayment equal to the current price of the bond.

Hirshleifer [1970] points out that IRR has two important drawbacks. First, when future cash flows vary in sign, there may be multiple solutions to the IRR. Second, the IRR implicitly assumes that intermediate cash flows may be reinvested at the discount rate.[2]

Interestingly, our simulation analysis allows us to examine the practical problem of multiple solutions to the IRR. Through repeated application of the IRR, we can observe whether the measure converges in practice to a meaningful number and how often it gives nonsensical results. Because properties are purchased over time for our portfolios, there can be a period of negative cash flows following a period of positive cash flows as additional properties are purchased, and then positive cash flows again as properties are sold.

Despite these known limitations, the IRR has two attractive features for our analysis. First, it does not rely on intermediate valuations of property, and, second, it takes into account the timing of the investment, and hence captures the actual experience of investors over the period.

SIMULATION PROCEDURE—BOOTSTRAPPING

Our approach to the simulation of portfolio returns is the bootstrap. Bootstrapping has become one of the most widely used numerical techniques in statistical analysis over the past three decades. It has been specifically applied in financial research in the last 15 years to address issues such as the small-sample distribution of predictive regressions, or the stability of portfolio optimization solutions.[3]

Its advantage lies in its broad flexibility and the fact that it captures small-sample effects accurately. Although many distributions converge asymptotically, investors rarely have the luxury of an infinite number of trials. It is a matter of practical importance to understand the effects of investing in a fixed number of assets.

In the current analysis, the bootstrap is an ideal tool for analyzing the effects on investing in a set of actual properties rather than in a hypothetical portfolio of thousands. It allows us to set the sample number and replicate the outcome conditional upon this specific sample size. Investors do not hold the entire NCREIF portfolio—they instead invest in some tens or perhaps hundreds of properties. The bootstrap allows us to examine performance under these conditions.

The essential feature of the bootstrap is that it samples data repeatedly with replacement from an underlying distribution. When the samples are independent, the bootstrapped distribution of statistics has been shown to converge to the true distribution.

A simple example is the repeated draw of a sample of ten values from a large underlying population. As the number of draws increases, the standard deviation of the bootstrapped distribution of the mean of the ten observations approaches the standard error of the mean conditional upon a sample size of ten.

The beauty of the bootstrap is that it can be applied to understand the small-sample distribution of fairly complex statistics such as the internal rate of return.

Overview of the Procedure

The goal of our bootstrapping analysis is to simulate the experience of an institutional investor with a significant allocation to real estate from fourth-quarter 1977 through second-quarter 2004, using the actual cash flows from property investment during this period. To do this, we use data on the purchase, sale, and intermediate cash flows of specific properties in the NCREIF database.

The concept is relatively straightforward. We randomly sample subsets from the roughly 4,000 properties in the NCREIF database that were bought in or after 1977 Q4 and sold in or before 2004 Q2. Each random sample represents a hypothetical history of property investment experienced by an institutional investor. These property histories are included in a larger hypothetical portfolio of U.S. stocks and bonds according to some target level of allocation to the three asset classes. Property purchases are financed by shifting assets from the other asset classes, and property liquidations are used to either invest in another property (depending on the allocation rule) or to purchase stocks and bonds in policy proportions.

Given this basic approach, we then vary the number of properties per simulated history as well as the type and geographic focus of the property portfolio. For each variation, we perform 1,000 simulations, retaining some basic information about the portfolio history and calculating the IRR of the portfolio as a whole and the real estate portion of the portfolio specifically. For the stock and bond portions of the portfolios, we use the S&P 500 index to represent equities and the Ibbotson Associates U.S. government bond index to represent the bond investment.

Recognize that by using passive stock and bond indexes we are capturing variations in IRR due only to the changing composition and timing of the real estate portion of the portfolio. If, for example, we presume that stocks, bonds, and real estate could be held in fixed proportion (say, thirds) through the time period, and the real estate investment could be represented by the NCREIF index, each simulation would yield precisely the same IRR, rendering the exercise meaningless. Thus, the analysis is specifically designed to describe the range of investor experience due to the constraints of investing in individual properties as opposed to an index as well as the effects of financing this activity through rebalancing of the rest of the portfolio.

A useful way to understand the analysis is to treat portfolio performance as attributable to three components: long-term policy weights, time variation in policy weights, and security selection. We are in effect varying the selection and timing for one of the three asset classes while holding the selection constant for the other two. We can also vary the long-term policy weights, but the effect of this is somewhat arbitrary; the larger the component of real estate in the portfolio, the greater the variation in IRR explained by real estate factors.

One other constraint is important to recognize. We are sampling only properties with realized total returns within the period. While 27 years would seem to be long enough to evaluate performance, real estate is by nature a long-term asset. Pro forma analysis of returns to property investment is often conceived in decades. We are not including properties that are still in the NCREIF. Thus, if there are some systematic performance characteristics of sold properties, we could be biasing our results.

Finally, the experience of real estate investing at the beginning of the period, i.e., the 1970s and early 1980s, will be effectively underrepresented, as there were relatively few properties in the database at its inception. By the same token, the experience of the 2000s will be underrepresented in the database because relatively few properties that have been bought in the last few years have been sold. Hence, the simulation analysis reflects the middle part of the time period to a greater degree.

Given that the NCREIF index exhibited gains in the early 1980s and in the 2000s and had more modest returns in the early 1990s, the IRRs in this study may be largely determined by one part of the long-term real estate cycle. On the other hand, one distinct advantage of our approach is that it is based entirely upon realized property returns and portfolios of actual histories of properties in a manner never really possible before.

To examine the extent of any bias due to including only sold properties in our sample, we also calculate the IRR on a portfolio that consists of all NCREIF properties, and assume that unsold properties were sold at their appraised value. To evaluate the real rate of return of realized property investment, we develop a method for calculating the IRR in real terms. This analysis provides an estimate of real estate's realized premium over inflation.

Thus, while our goal is to understand the particular effects of investing in an asset class subject to integrality constraints (the purchase and sale of whole properties), limitations on the number of properties in the portfolio, and the effects of purchase and sale decisions occurring at different times, we also provide some interesting evidence on real estate as an inflation hedge, as well as a quantitative benchmark for developing an expectation of its real return going forward.

Our approach allows us moreover to examine the cross-sectional standard deviation of IRRs, an important measure of the ability of an active manager to meet a benchmark. The greater the cross-sectional standard deviation, the greater the risk of not meeting the benchmark because of particular properties held by the manager.

EXHIBIT 1
Distribution of Purchase Dates of NCREIF Properties in Sample

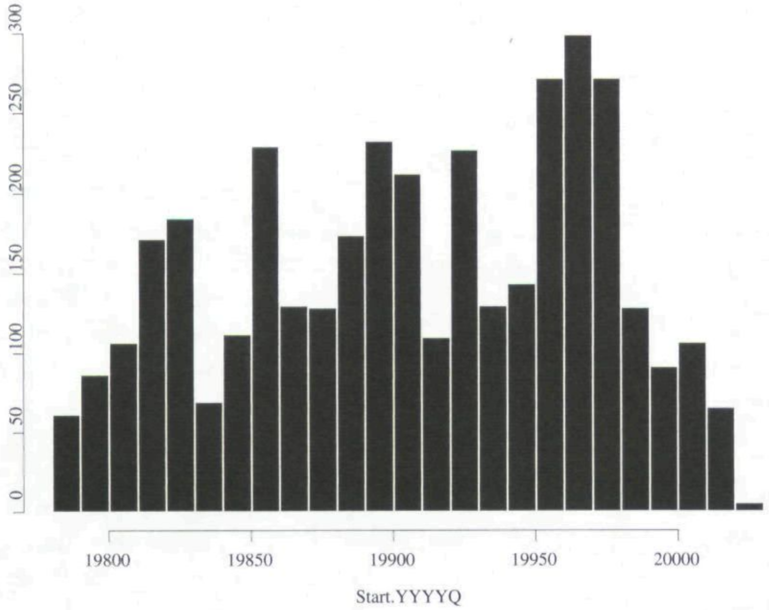

EXHIBIT 2
Distribution of Sales Dates of NCREIF Properties in Sample

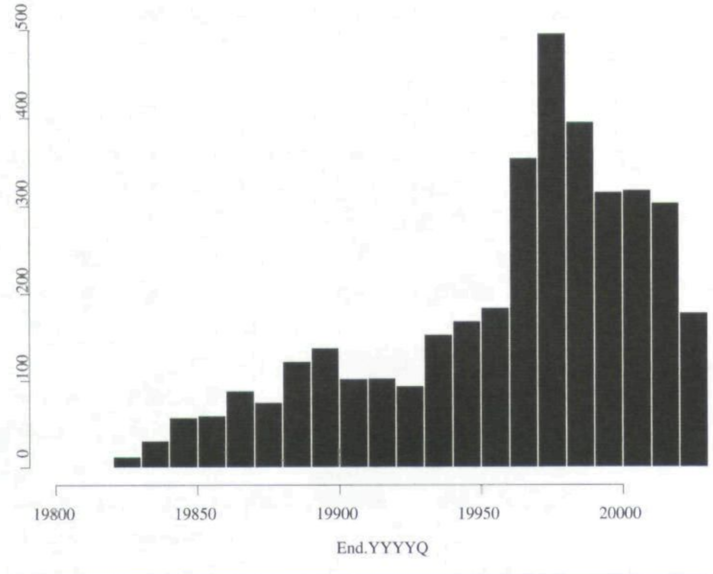

Data Description

The NCREIF database consists of properties that are held by investment managers on behalf of tax-exempt pension funds. Members of NCREIF contribute individual property data quarterly that include the acquisition price, net operating income, capital expenditures, and market value of each property. When a property is sold, the sale price of the property is reported. Although this information is usually used to calculate quarterly time-weighted returns, all the necessary data are available to calculate the IRR on the property from acquisition to disposition.

The data start in the fourth quarter of 1977, with quarterly data available since that time. Data are made available to members in a manner that ensures confidentiality with respect to individual property performance.

The database allows us to analyze cash flows from individual properties, including purchase and sale prices as well as intermediate cash investment and income. Appraisal information, property type, and regional location are also provided for this study. All return data are calculated on an unlevered basis.

Size

The NCREIF index currently includes information on approximately 4,000 properties, although the total number of properties reported at one time or another to the NCREIF database since inception is approximately double this number. We are sampling from the 4,773 properties that have *left* this database due to sale of the property.[4]

The total appraised value of the current NCREIF index properties as of 2004 Q2 was approximately $140 billion. The sample properties sold had an initial acquisition cost totaling about $84 billion and an aggregate sale price of $91 billion.

Time Period

Exhibits 1 and 2 show the distribution of the starting and ending dates of the properties in our sold sample. Notice that the sale dates suggest that many of the properties we analyze are those for which returns were realized in the late 1990s and 2000s. Exhibit 3 shows the distribution of holding periods for properties in the sample. The half-life of a typical property in the sample is roughly five years; this implies a holding period of

EXHIBIT 3
Histogram of Holding Periods for Properties in Sample

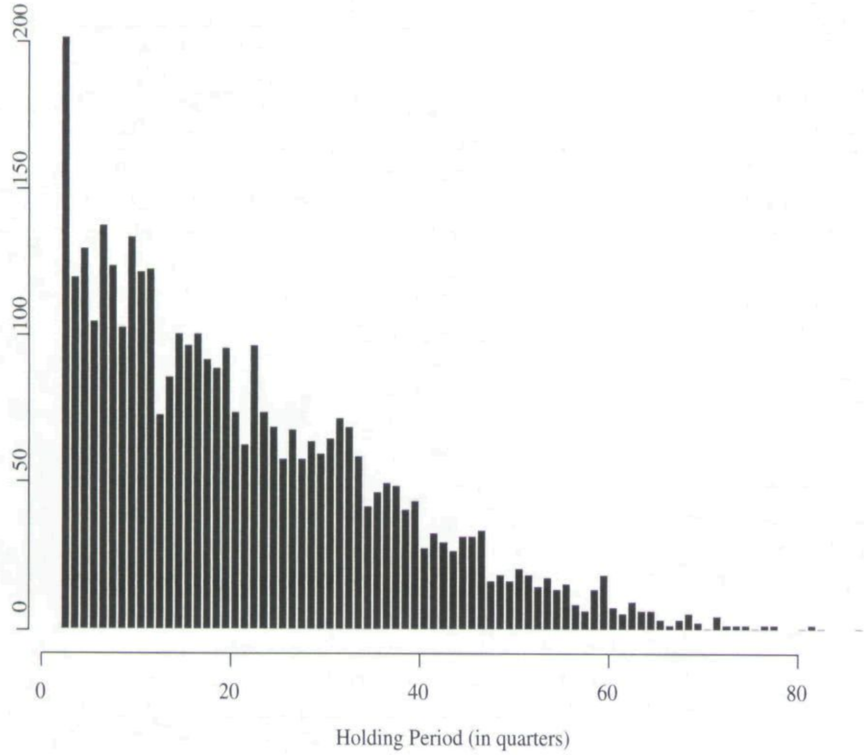

returns (or different characteristics from the retained properties), sampling from this group is not likely to be biased in a direction toward high or low returns over the time period covered in this study, which includes a full real estate cycle.[5]

It may be reasonable to expect that, on average, institutional real estate managers have been astute enough to sell before poorer-than-average performance. A test of this proposition is a potentially useful thing to perform in order to calibrate the potential bias. (In the current stage of our analysis, we have not addressed this issue.)

One important feature of the NCREIF sample is that it does not include development properties. All properties are complete and have reached at least 60% occupancy when added to the index. This implies that the NCREIF data capture returns over a particular stage in a property's history. This makes the analysis of returns cleaner for the simulation, but, to the extent that development returns might be higher due to commensurately higher risk, this should be taken into consideration in interpretation of the results. One might expect development properties, for example, to have higher unlevered rates of return.

Five different property types are represented in the sample: office, industrial, retail, apartment, and hotel. Institutional investors of course own other types of property. We are not including farmland or timberland, for example.

Whether the proportion of each type reflects the underlying rate of institutional investment as a whole is not known, nor is whether the distribution of properties by region within the sample is representative of investor activity. The fact that NCREIF is an industry organization intended to collect and provide investment-relevant information to its membership, however, is a strong indication that the sample is a representative one.

Exhibit 4 shows the current NCREIF portfolio by region and property type, reporting dollar values of properties and numbers of properties by category. For any single region or property type, it is clear that the current portfolio includes enough properties to allow a meaningful bootstrap.

about ten years. This is consistent with a study by Fisher and Young [2000a,b] that finds a median holding period for all NCREIF properties of about 11 years.

Because we do not know the holding period for properties that have not sold, we cannot use the holding period for only sold properties to calculate an average holding period. The time until half the properties are sold is an accurate measure of the median holding period because the time it takes for the other half to sell does not affect the median. Thus, although we are conditioning upon sale by the first quarter of 2004, we are capturing a holding period that is very similar to the median holding period for all NCREIF properties.

Sample Issues

Although a statistician might prefer otherwise, the basis for including a property in the NCREIF portfolio is not necessarily a random selection from an institutional portfolio. Yet there is no reason to believe that the sold properties are unrepresentative of institutional real estate investments during the past 27 years. While ex post the sold properties in the portfolio might have had different

Exhibit 5 shows the composition of the sold property sample.

Implementation

Whatever the various sampling and measurement issues, there is no doubt that NCREIF has the largest, the richest, and the most complete database of commercial property investment ever created. There is simply no comparable source of information. It has the particular advantage of including properties actually purchased and sold by institutional investors with the intention of generating returns for their beneficiaries.

Implementation of the bootstrapping procedure requires some detailed choices about the relative scale of the portfolios. We make a variety of assumptions in calculating the IRRs of diversified portfolios of real estate, stock, and bond assets. First, the initial investment amount depends on the number of real estate assets included in the diversified portfolio. We choose dollar amounts so that for roughly 95% of the diversified portfolios the maximum exposure to real estate would range from 10% to 40% of the total portfolio value during the 1977 Q4-2004 period as follows:

Number of Real Estate Properties	Initial Investment in $M
10	50
30	150
60	300
100	500

Initially all the available cash is split between stocks and bonds with *p.stocks* and *p.bonds* proportions (where *p.stocks* + *p.bonds* = 1). After each quarter, the stock and bond positions (but not real estate) are rebalanced to preserve the set proportions, regardless of real estate position.

Second, real estate assets are included in the diversified portfolio as follows. For each number of properties 10 through 100, NCREIF data provide us with cash flows of 1,000 real-estate only portfolios that were generated by sampling with replacement from more than 4,000 properties from 1977 Q4 through 2004 Q2. Then, for a given real estate portfolio, when investments into real estate are made or cash flows are generated from the real estate investments, the money is taken from (or put back into) stocks and bonds to preserve their initial proportions.

At the beginning of the sampling period (1977 Q4) and at its end (2004), all the diversified portfolios have $0 in real estate, by construction. Only properties that were bought and sold during the 1977 Q4-2004 period are included in real-estate only NCREIF portfolios.

These rules for portfolio construction result in time series of real estate portfolio weights that are somewhat lumpy and reach their maximums toward the middle of the sample period, since our procedure effectively requires all real estate positions to be closed out by second-quarter 2004.

EXHIBIT 4
NCREIF Database Current Holdings by Region and Property Type

Region	Dollar Value in Millions	Number of Properties
Midwest	19,636	668
East	44,567	920
South	28,991	1,101
West	48,640	1,336
Total	141,834	4,025

Property Type	Dollar Value in Millions	Number of Properties
Retail	30,831	563
Apartment	27,527	806
Industrial	26,676	1,465
Office	54,465	1,135
Hotel	2,335	56
Total	141,834	4,025

Source: NCREIF.

EXHIBIT 5
NCREIF Sold Properties

Region	Dollar Value in Millions	Number of Properties
Midwest	14,235	890
East	27,371	968
South	20,721	1,346
West	28,510	1,569
Total	90,836	4,773

Property Type	Dollar Value in Millions	Number of Properties
Retail	24,314	902
Apartment	16,136	793
Industrial	15,371	1,715
Office	32,653	1,306
Hotel	2,360	57
Total	90,836	4,773

EXHIBIT 6

Summary Statistics for Bootstrapped Distributions of IRR Values —1997 Q4–2004 Q2

Number of Properties	Mean	StDev	Min	Max	Median	2.5% below	2.5 % above
10	0.0756	0.0376	-0.0777	0.2298	0.0756	-0.0024	0.1520
30	0.0750	0.0233	-0.0204	0.1873	0.0751	0.0280	0.1212
60	0.0748	0.0158	-0.0039	0.1430	0.0745	0.0451	0.1068
100	0.0749	0.0127	0.0376	0.1354	0.0746	0.0506	0.0996

RESULTS

We present results in several ways.

Real Estate Only

Exhibit 6 summarizes the results of the basic simulation based on 1,000 bootstrapped samples each of portfolios of 10, 30, 60, and 100 properties. Notice that the mean and median IRR of roughly 7.5% is consistent, regardless of portfolio size. This is not surprising, given the law of large numbers.

The time-weighted rate of return of the NCREIF NPI, constructed from the entire database using periodic appraisal data, is 9.35% per year. Ciochetti and Fisher [2002] find an equal-weighted IRR for NCREIF properties from 1980 to 2001 of 8.73%. Note that this does not consider the differences in the amount invested for different properties and the exact timing of the cash flows. The IRR we calculate is based on pooling the cash flows from all the properties and considering exactly in what quarter the cash flow occurs for the entire portfolio. This is a better measure of the IRR actually earned on a portfolio represented by the NCREIF sold properties.

One possible explanation for the difference between the mean IRR of 7.5% for the sold properties and the time-weighted return from the entire database (sold and unsold based on appraised values) of 9.35% is that the sold property sample underrepresents the performance of properties in recent years that have not been sold. This in turn could result in a sample selection bias.

Test for Sample Selection Bias

To test for the possibility of sample selection bias, we create a portfolio that includes all the properties that have ever been in the NCREIF database. We assume that the properties not actually sold were sold at their final appraised value. That is, all properties still held at the end of the first quarter of 2004 were sold at their appraised value. Similarly, properties that left the database for other reasons over time are assumed to have been sold at their last reported appraised value. The acquisition price and quarterly cash flows for these properties are then included in the calculation of the IRR for these properties.

The IRR for this portfolio of all 11,729 properties that were ever in the NCREIF database is 7.82%, just slightly higher than the average 7.50% IRR for the sold property sample. This suggests that the sold property sample is quite representative of all the properties in the NCREIF database, and there is no significant selectivity bias in using just the sold properties. According to actual dollar inflows and outflows for tax-exempt institutional investors over the 27-year history of the NCREIF index, the average IRR ranges from about 7.5% to 7.8%. This is significantly lower than the since-inception time-weighted return of about 9.4% that does not consider the capital flows into and out of real estate over this time period.

Inflation-Adjusted Return

The second half of the sample period was also somewhat unusual, in that it was characterized by low inflation compared to the first half. Inflation protection is one of the main reasons institutions invest in real estate. Thus, it is helpful to develop a measure of the premium that realized property returns provided over inflation.

To do this, we cannot simply subtract inflation over the time period. Instead, we take the quarterly cash flows from the aggregated portfolio of sold properties and discount each by the realized inflation back to the first time period. We then calculate the IRR on the basis of these inflation-adjusted flows.

The real IRR was 3.99%—close to a 4 percentage point premium over inflation. For the portfolio that includes all the properties (sold and unsold) that we have discussed, the real rate of return was 4.57%.

This slightly higher return makes sense because it puts more weight on properties that were not sold before the strong real estate market in recent years of low inflation. This historical experience is potentially useful for estimating the future expected return for real estate in a building block model of risk premiums for exposure to relevant systematic risk factors. Inflation plus 4 percentage points appears to be a reasonable rule of thumb.

Cross-Sectional Standard Deviation and Number of Properties

Of even more interest is the spread of the distribution. With only ten properties held in the portfolio, the standard deviation is roughly half of the mean. Thus one of 20 of these portfolios had returns of lower than zero. This risk of loss drops significantly as the number of properties increases to 100. At this scale of investment, 95% of the realized IRRs lie within the 5% to 10% range.[6]

Exhibit 6 provides a benchmark for how individual institutional portfolio experiences might be expected to have varied over the period because of random variations in the properties in the portfolio. A portfolio of substantial size, diversified across region and property type, with an historical realized IRR of more than 10% can be viewed as an exceptional performer. By the same token, a return of this scale for a small real estate portfolio might easily have occurred by chance.

Because managers cannot hold a proportion of every property in the benchmark, there is a risk that the portfolio they hold will perform differently from the benchmark simply because of the unique characteristics of the properties held. The greater the cross-sectional volatility, the greater the risk of differing from the benchmark. These results suggest that rather large portfolios are necessary to significantly reduce this benchmark risk.

Exhibit 7 reports the IRR values for portfolios of properties categorized by region and by property type. For example, the first cell in the first column of the table is created by restricting the randomization only to apartments in the East. The last row of the table is based on random draws across properties in the region (*diversified*). This result will of course depend on the relative base frequencies of property types by region.

The last column is based upon random draws from single property types across the country. This too will depend on the relative base frequency of regional location by property type across the country.

These focused portfolio IRRs range from 3.8% (Southern office) to 10.8% (Western apartment), which represent statistically different means for portfolios of 100 properties. An analysis of variance of Exhibit 7 indicates significant differences across property types and insignificant differences across regions.

EXHIBIT 7
Mean IRR Values for Subsamples by Region and Property Type

	East	Midwest	South	West	US
Apartment	9.5%	9.6%	9.5%	10.8%	9.9%
Industrial	8.2%	7.7%	6.8%	9.6%	8.4%
Office	8.4%	5.1%	3.8%	4.6%	5.9%
Retail	7.8%	8.5%	7.3%	7.9%	7.8%
Diversified	8.3%	7.2%	6.6%	7.5%	7.5%

EXHIBIT 8
Standard Deviation of IRR Values by Region and Property Type

Ten Properties in the Portfolio

	East	Midwest	South	West	US
Apartment	1.95%	1.93%	3.66%	9.35%	5.12%
Industrial	2.81%	2.26%	3.66%	3.17%	3.22%
Office	4.56%	2.87%	3.64%	4.15%	4.16%
Retail	4.15%	2.93%	3.81%	3.98%	3.92%
Diversified	3.96%	3.12%	3.92%	3.98%	3.76%

One Hundred Properties in the Portfolio

	East	Midwest	South	West	US
Apartment	0.61%	0.58%	0.81%	0.77%	0.73%
Industrial	0.93%	0.70%	0.96%	1.04%	0.96%
Office	1.26%	0.97%	0.96%	1.08%	1.23%
Retail	0.97%	0.94%	1.14%	1.52%	1.24%
Diversified	1.19%	1.17%	1.12%	1.46%	1.27%

Exhibit 8 reports the standard deviations of the IRRs for the subportfolios. Some types display much greater cross-sectional variation than others. This provides some evidence on a long-standing issue in real estate portfolio analysis: whether regional diversification is as important as diversification across property type.

To date, this question has relied upon time series data and either correlation or clustering analysis. By bootstrapping IRR values, we are able to directly observe the relative importance of regional diversification versus property type diversification.

Our approach suggests that, at least ex post, property type explains more of the variation in performance than regional classification. There are smaller interregional than interproperty type differences across IRR by property type. This is consistent with results in Fisher and

Exhibit 9

IRR and Standard Deviation of IRR versus Holding Period

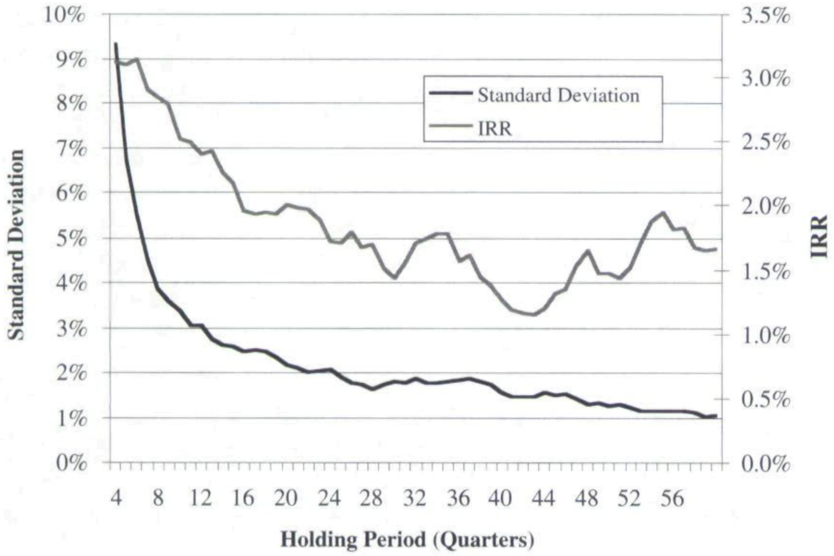

Liang [2000], who find that property type diversification is more important than regional diversification (according to time-weighted returns from the NCREIF database).

One caveat to these results is that they are sensitive to the number and cross-sectional value distributions of properties within property type. For example, suppose there are only a few properties in one asset class in one region, and one or two performed very well. This would lead to a skewed IRR distribution for that property type for that region, because some of the bootstrapped samples would include those properties and some would not.

There is some evidence of this pattern. The Southern apartment IRR distributions are highly positively skewed, even in samples of 100 (with a maximum of 1,000 draws of 16.15%). For samples of 30, the Western apartment IRR distributions are unusually positively skewed (with a maximum of 1,000 draws of 87.65%).

Differences in skewness for apartment subsamples do not appear to be driving the diversification results. In Exhibit 8, the cross-sectional variation for apartments is lower than the mean, not greater than the mean. If skewness were responsible for these results we would expect the opposite.

Cross-Sectional Standard Deviation and Holding Period

We also examine the cross-sectional standard deviation based on the length of time a property was held. That is, how does the volatility compare for a strategy of shorter versus longer holding periods? The holding periods include all properties held for a particular holding period, no matter what year the property was purchased, so they are an average of the results for various holding periods over the entire sample period. This allows us to evaluate how the benchmark risk might differ for a strategy involving more active management of properties by holding for shorter periods versus more of a buy-and-hold strategy.

Exhibit 9 shows the cross-sectional standard deviation for each possible holding period up to 60 quarters (15 years).[7] The equal-weighted IRR is also shown for each holding period.

It is interesting that both the IRR and standard deviation drop fairly steadily over time until about a 40-quarter (10-year) holding period. For holding periods longer than 10 years, the IRR appears to start to rise, although the standard deviation is still falling. This may be an anomaly because of the smaller number of properties with holding periods of longer than 10 years.

What is clear is that the cross-sectional standard deviation drops as the holding period lengthens. Or, to put this differently, the standard deviation increases as the holding period is shortened but the IRR also increases—so investors who had a strategy of shorter holding periods took on more benchmark risk but also on average achieved a higher IRR.

Real Estate in the Diversified Portfolio

While the results of the analysis thus far are of interest to real estate portfolio managers, the goal of investing in any asset class is to add value to the portfolio as a whole. Thus, it is particularly useful to examine the effects of the simulated real estate investments on the broader portfolio. We thus look at the IRR distributions of the whole investment portfolio.[8]

These distributions are conditioned on the actual historical performance of an index-based stock-bond-real estate portfolio return from 1977 Q4 through 2004 Q2. Our procedure introduces variation only in the relative proportions of these assets and in the composition of the

EXHIBIT 10
Proportion of Stocks, Bonds, and Real Estate in Given Portfolio

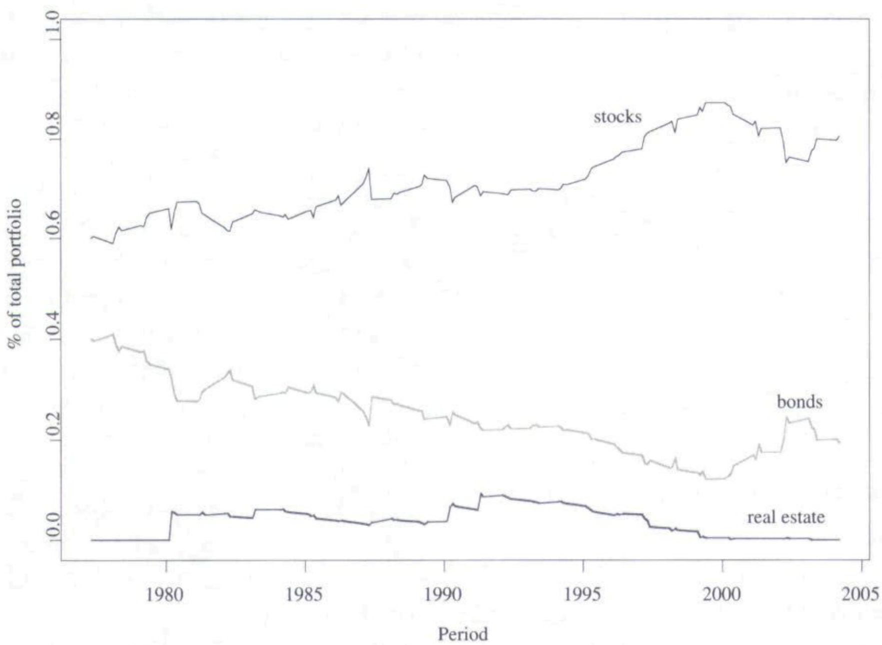

property portion of the portfolio. It does not allow us to assess the expected diversification effects provided by real estate due to correlations. It does allow us to examine the effects of rebalancing timing and integrality.

The best way to appreciate the effect of assumptions on the smooth rebalancing of a portfolio of stocks, bonds, and real estate versus the practicalities of modeling a portfolio when properties enter and exit the portfolio at different times and in different amounts is to examine the time series of approximate portfolio weights for a random portfolio in our analysis. Exhibit 10 is one of 1,000 sample portfolio paths in our analysis. It is based upon an initial value of $100 million divided 60/40 between stocks and bonds.

Notice that, as the randomly sampled properties enter the portfolio, stocks and bonds are liquidated in the same 60/40 proportion to finance the purchase, and as properties are sold, the proportion of real estate is reduced. Real estate is not marked to market, but valued at cost. The lumpiness of this process is apparent from the stepped pattern of the real estate component.

Exhibit 11 reports summary statistics about the IRRs of the portfolios of stocks, bonds, and real estate. The first panel provides results for 100 property portfolios; the second panel provides results for 10 property portfolios.

The median return for the simulated portfolios is about 11% for simulations in both panels. The cross-sectional variation based on 100 properties (32 basis points) is about half of that based on 10 properties (74 bp). These effects appear modest in the context of the overall investment portfolio, but this is because the proportion of real estate in the portfolio itself is modest in the simulations.

The median maximum percentage of real estate in a portfolio with 100 properties is 23%, while the median maximum for simulations with 10 properties is 28%. This means that, at some time in the period real estate achieved a maximum in the typical portfolio of just over 20%.

Of course each portfolio begins with no real estate and ends with no real estate. This proportion is driven by the lumpiness of properties—sometimes the simulation samples a large property that requires a major drawdown from the stock/bond holdings. When the drawdown exceeds the other holdings, that simulation is omitted from the analysis.

Because of the manner of construction, 10-property portfolios are smaller and lumpier than 100-property portfolios. This is consistent with the challenges facing a small investment fund that maintains a portfolio of investment properties, and thus is a desirable feature of the simulation. For larger portfolios, a single purchase has a minimal effect—integrality is less of an effect. As a consequence, the difference between the distributions of 10 versus 100 property IRRs is largely due to integrality and not to investment timing.

Exhibit 12 shows four histograms based on the simulations with 100 properties. It breaks the 1,000 IRRs generated by the diversified portfolios into four groups according to the maximum percentage of real estate represented.

The first group, the highest proportion, represents portfolios with maximums of over 29%. It clearly has the widest spread of the four groups, since it has the most variation introduced by the simulation. As the proportion in real estate drops, the distributions converge toward 11.5%, the IRR provided by 60/40 stocks and bonds over the period.

The differences between the real estate portfolio

EXHIBIT 11

Summary Statistics of Bootstrapped Distributions of IRR Values for Diversified Portfolios

...%	...182%	1.270%
1.076%	22.970%	7.463%
1.754%	74.373%	13.544%
9.018%	8.988%	3.756%

...folio IRR	Maximum % in Real Estate	Real Estate IRR
	10 Properties	
0.747%	20.535%	3.706%
1.320%	28.095%	7.559%
2.757%	99.189%	22.981%
6.079%	5.117%	-7.372%

...l the 60/40 portfolio returns are driven by two ...t, there were very different realized returns across ...s in this period. The 1980s and 1990s represent ...greatest decades in the history of the U.S. equity ...d bonds enjoyed remarkable unanticipated cap- ...ue to falling rates in this period as well.

...nd, real estate in the simulation is unlevered, ...ties, to the extent a firm had debt on its bal- ..., were levered. Thus the cross-sectional spread ...tocks and properties can be scaled somewhat ...through leverage.

Cross-Sectional Volatility of a Portfolio of Real Estate and Stocks

folio as the holding period shortens. In a final experiment, we examine how the cross-sectional volatility differs as a function of holding period when the portfolio consists of both real estate and stocks in equal proportions. This lets us examine the diversification benefits of real estate across different time horizons. Lower cross-sectional volatility means less unsystematic risk as well as a greater likelihood of tracking the benchmark, i.e., less tracking error.

Perhaps not surprisingly, the holding periods have to be significantly longer for stocks in order for the cross-sectional volatility to decline to the same level as for real estate. More important, the cross-sectional volatility declines more rapidly for the mixed-asset portfolio consisting of both stocks and real estate than for either a pure stock or a pure real estate portfolio. This suggests that multiasset portfolios that include real estate are more likely to track the benchmark.

This also exemplifies the diversification benefits of adding real estate to a stock portfolio by using IRRs rather than time-weighted returns—a topic that warrants further research.

...BIT 12

...ons of IRRs According to Proportion in Real Estate

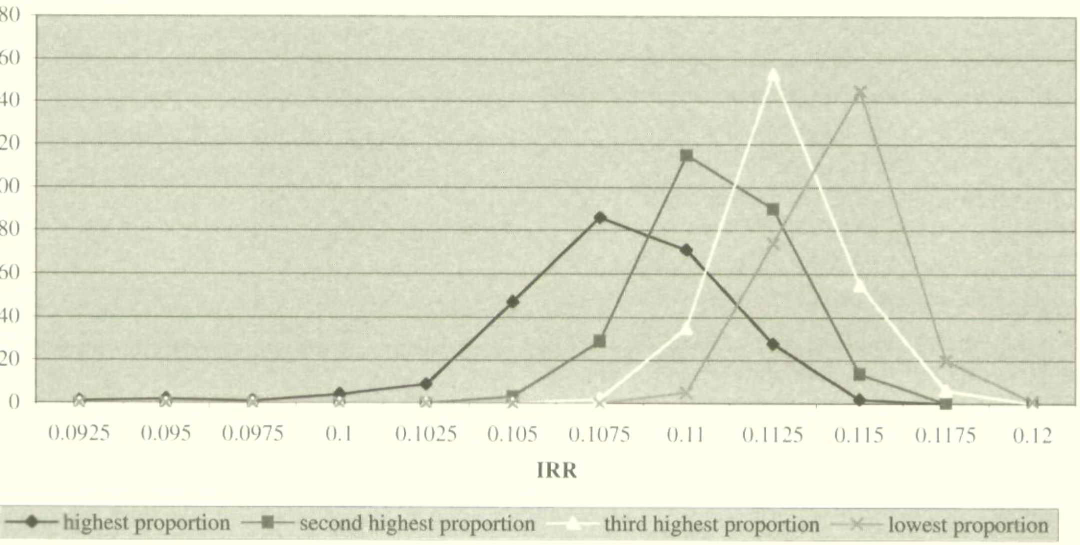

Median 1
Max 1
Min

Port...

Std
Median 1
Max 1
Min

returns and... factors. Firs... asset classes... two of the... markets, an... ital gains d...

Seco... while equ... ance sheet... between s... arbitrarily...

EXHI...

Distributi...

CONCLUSION

Investment analysis relies heavily upon indexes for portfolio choice as well as for performance evaluation. Even indexes for frequently traded asset classes can be misleading if they are not in fact investable.

Commercial real estate is one of the most important asset classes in institutional investment portfolios. Institutional investors typically hold it through commingled investment funds or real estate investment trusts, or in separate accounts. Even large-scale investment portfolios, however, confront integrality constraints and the need to invest in properties, not indexes.

We simulate the experience of an institutional investor active in the period between last-quarter 1997 and second quarter 2004. Instead of using appraisal data, we use transactions and intermediate cash flows to compute the returns to property portfolios that are subsets of the actual investable universe, and that taken together effectively cover the range of realistically achievable performance. In simple terms, these simulations allow us to look at the realized returns that real estate generated during this period, without the potentially confounding effects of appraised values and time-weighted return calculations.

The unlevered realized IRR of commercial property investment was about 7.5% in this period—lower than the time-weighted rate of return of 9.4% due to issues of investment timing. This is not surprising, given the strong positive returns of property investment at the beginning and the end of the sample period and given changes in realized rates of inflation. To address the obvious macroeconomic regime change, our methodology accounts for shifts in inflation. We estimate the inflation-adjusted rate of return for real estate investment over the period at 3.99%.

Also interesting is the range of variation in IRR attributable to investing in properties as opposed to an index. For small portfolios holding 10 properties, there are non-trivial probabilities of a negative IRR. For portfolios with 100 properties, the variation declines significantly but still ranges from 5% to 10%, with an IRR over 10% as an exceptional benchmark.

This is important because it suggests that individual portfolio experiences with real estate investing will vary because of individual property differences and the timing of investments. These differences are greatest when the real estate portfolio holdings are small, but they are significant even for large portfolios. The practical implications of this finding are that large pools of properties are required to achieve returns similar to the returns of the population of commercial properties—i.e., scale is required for effective diversification and to reduce the risk of failing to meet a benchmark due to the unique characteristics of the properties held by the manager.

We also examine the effect of shorter holding periods and thus more active management versus longer holding periods that would represent more of a buy-and-hold strategy. Shorter holding periods show more cross-sectional volatility (benchmark risk) but also generally higher rates of return. Furthermore, the cross-sectional volatility of a multiasset portfolio consisting of stocks and real estate is lower for a given holding period than for portfolios of a single asset type.

In analysis of subportfolios, we find that mean returns differ by property type more than might be expected, and that cross-sectional variation in IRR performance is explained well by property type. This is an interesting result that merits more analysis. There is also some evidence that diversification by property sector is more important than diversification by property location.

Finally, we place our simulated real estate portfolios into the context of a stock-bond allocation and simulate the practical process of rebalancing. This involves the financing of real estate purchase at the expense of the other asset classes, and the reinvestment of sales proceeds into tradable securities. This generates "modified" IRRs for investment portfolios that range from 10% to 12%, depending on the proportion of real estate and number of properties. This approach demonstrates that portfolio IRRs can vary significantly, depending on the timing and the diversification of the real estate investment.

We should emphasize that time-weighted returns are still appropriate to measure the volatility of real estate versus stocks and bonds. This is because the simulations really capture only returns, not risk. Yet our results show that the IRR is a powerful tool for measuring the expected return and examining its variations. Particularly useful is an estimate of the premium that real estate provides over inflation.

The full portfolio simulation procedure also highlights the role that real versus liquid portfolio securities play in a dynamic investment portfolio. Swensen's *Pioneering Portfolio Management* [2000], for example, makes a strong case for less-liquid asset classes as a source of superior investment returns. Swensen argues that institutions such as pension funds and endowments that can afford to wait for the opportunity to buy or sell less-liquid assets depending upon evolving market conditions over years can take advantage of potentially attractive valuations. The ability to do this depends on financing lumpy real asset transactions with more-liquid capital resources.

Such a structural relationship between marketable securities and real assets in the portfolio is not captured by the usual approaches to portfolio analysis. Our simulations explicitly take these synergies into account.

ENDNOTES

The authors thank the Real Estate Research Institute for funding this project through directed grants by the TIAA-CREF Institute and the Pension Real Estate Association, and thank NCREIF for access to the data. They also thank Eduardas Valaitis for programming and research support.

[1]See Geltner [1977], Fisher and Geltner [2000], and Clayton, Geltner, and Hamilton [2001] for a discussion of the unsmoothing problem.

[2]An explicit reinvestment assumption is not necessary to calculate an IRR, but if cash flows must be reinvested at a lower rate than the IRR the realized return from the portfolio will be lower than the IRR.

[3]For application in real estate, see Goetzmann [1992].

[4]There is some ambiguity about the number reported, because if a property is sold by one institution to another, it may leave the database for the first and then be added by the second owner. Additional properties leave the database for other reasons than sale, such as transfer to another manager, but this has little practical effect on our analysis.

[5]Fisher et al. [2003] do find some evidence of selling losers during the down market of the early 1990s and selling winners during the up market of the late 1990s. Fisher et al. [2004] use a probit analysis to examine what properties are more likely to sell as a function of economic and property characteristics.

[6]This convergence of the distribution with an increase in portfolio size tells us something about IRR as a summary measure of investment as well. If pathological rates of return for cash flow sequences that reverse sign were a major problem with IRR, we would not observe this reduction in range. The sequences with many properties have just as many sign reversals as those with fewer properties.

[7]Beyond 15 years there were fewer than 10 properties sold in any given quarter, so we include only holding periods up to 15 years.

[8]As an aside, this addresses one of the known limitations of the IRR, which is the assumption that intermediate cash flows are reinvested at the IRR, despite the uniqueness of the project. By mechanically drawing from a 60/40 stock-bond portfolio and reinvesting in it, we are effectively calculating a particular kind of the modified IRR—substituting a risky portfolio for the riskless asset.

REFERENCES

Ciochetti, B., and J. Fisher. "The Characteristics of Commercial Real Estate Holding Period Returns (IRRs)." Working paper, Real Estate Research Institute, 2002.

Clayton, J., D. Geltner, and S. Hamilton. "Smoothing in Commercial Property Valuations: Evidence from Individual Appraisals." *Real Estate Economics,* 29(3) (2001), pp. 337-360.

Fisher, J., D. Gatzlaff, D. Geltner, and D. Haurin. "An Analysis of the Determinants of Transaction Frequency of Institutional Commercial Real Estate Investment Property." *Real Estate Economics,* 32(2) (2004), pp. 239-264.

——. "Controlling for the Impact of Variable Liquidity in Commercial Real Estate Price Indices." *Real Estate Economics,* 31(2) (2003), pp. 269-303.

Fisher, J., and D. Geltner. "De-Lagging the NCREIF Index: Transaction Prices and Reverse Engineering." *Real Estate Finance,* 17(1) (2000), pp. 7-22.

Fisher, J., and Y. Liang. "Is Sector Diversification More Important than Regional Diversification?" *Real Estate Finance,* Fall 2000.

Fisher, J., and M. Young. "Holding Periods for Institutional Real Estate in the NCREIF Database." *Real Estate Finance,* Fall 2000a.

——. "Institutional Property Tenure: Evidence from the NCREIF Database." *Journal of Real Estate Portfolio Management,* 6(4) (2000b), pp. 1-8.

Geltner, D. "The Use of Appraisals in Portfolio Valuation and Index Construction." *Journal of Property Valuation and Investments,* 15(5) (1997), pp. 423-447.

Goetzmann, William N. "The Accuracy of Real Estate Indices: Repeat Sales Estimators." *Journal of Real Estate Finance and Economics*, 5(1) (March 1992), pp. 5-53.

Hirshleifer, Jack. *Investment, Interest and Capital.* Englewood Cliffs, NJ: Prentice-Hall, 1970.

Swensen, D. *Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment.* New York: Free Press, 2000.