YALE
INTERNATIONAL
CENTER
FOR FINANCE

Yale ICF Working Paper No. 05-07
April 2005

THE PERFORMANCE OF REAL ESTATE PORTFOLIOS: A SIMULATION APPROACH

Jeffrey Fisher

Indiana University

William N. Goetzmann

Yale School of Management, Yale University

The Performance of Real Estate Portfolios: A Simulation Approach[1]

Jeffrey Fisher, Indiana University
William N. Goetzmann, Yale School of Management

Current Draft: 4/13/2005

Abstract:

In this paper we simulate the performance of real estate portfolios using cash flows from commercial properties over the period 1977 Q4 through 2004 Q2. Our methodology differs from analyses that rely upon historical time-weighted rates of return on property. We relax implicit rebalancing and mark to market assumptions inherent in time-series analysis. We use the distribution of internal rates of return to analyze the performance distribution of commercial property investment. We examine the performance of real estate in the context of portfolios of stocks and bonds over the same period.

JEL: G11, R33

Keywords: Asset Allocation, Real Estate

---

[1] Acknowledgements: We thank the Real Estate Research Institute for funding the project through directed grants by TIAA-CREF Institute and the Pension Real Estate Association. We thank NCREIF for access to the data. We thank Eddie Valaitis for his research support.

I Introduction

Portfolio managers rely heavily on historical information to make asset allocation decisions. The empirical application of neo-classical 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 re-invest the proceeds as needed and immediately liquidate the position at the end of a fixed time period. For other asset classes, the performance of an index – at least a standard, capital-weighted index – may be a poor approximation to the investment experience, and indeed may limit the applicability of standard neo-classical models.

As institutional managers increase their allocation 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 and yet these asset classes are regarded as important sources of return and diversification by modern managers.

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 indices that reflect broad trends in a diversified portfolio of investable properties. Indeed the authors of this paper have both conducted research on the accurate measurement of time-weighted rate of return indices on property.

In this paper, however, 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 over the period 1977 Q4 through 2003 using actual property histories. 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 among its membership for more than 25 years. This rich database is now the definitive basis for commercial real estate index data. In this study, we rely upon the NCREIF data to imitate the total return an active investor this time period would have experienced.

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 comprised of commercial properties, U.S. stocks and U.S. bonds. The stock and bond portions of the portfolio are re-balanced 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 the time period under a variety of assumptions. In our study, we vary the number of properties in the portfolio, the approximate portfolio allocation to real estate and the regional and property-specific focus.

Ciochetti and Fisher (2002) calculate IRRs for properties in the NCREIF database but they calculate equal weighted IRRs (weighting the IRRs the same for all properties regardless of the amount invested in the property. This does not capture the IRR for a portfolio of properties for which the IRR depends on the amount of investment and cash flow from each property and the timing of that cash flow. Ciochetti and Fisher did not evaluate real estate in a mixed asset portfolio with stocks and bonds and did not employ bootstrapping.

The results of the study are at once familiar and surprising. For diversified property portfolios, the median IRR differs significantly from the compound time-weighted rate of return from the quarterly NCREIF index. We attribute this difference to the timing of most of the property investment in the time period – the IRR measure is investment-weighted, not time-weighted. The distribution of the returns is also instructive. We find that significant reductions in the cross-sectional dispersion of returns is still achievable in portfolios of 100 properties. In other words, diversification requires substantial capital. 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 geographical diversification.

Before a comprehensive presentation of our results, it is important to explain the premises of our approach as well as the details of the simulation procedure. The next section addresses the measurement of returns. Section three summarizes the simulation procedure, section four presents the results of our analysis and section five concludes.

II Performance Measurement

II.1 Time-weighted rate of return

The use of IRR as measure of returns is somewhat unconventional in the context of portfolio analysis. As discussed above, much of the current statistical analysis of the performance of asset classes is based upon the time-weighted rate of return. The time-weighted rate of return relies upon 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 under that assumption that both income and capital appreciation of the asset may be realized (or at least accurately measured) each period. These gains are then scaled by the amount of initial investment at the beginning of the period. Time-weighted rates of return have the advantage of allowing the estimation of standard deviation and correlation, and they are invariant to the magnitude of the investment at any period. They treat a gain of $1 on $100 the same as the gain of $1 million on $100 million. This is particularly useful when the investment manager cannot control the amount of money under management but can influence its percentage gains or losses. It is also useful when the investment technology can be easily scaled.

An important drawback to time-weighted rates of return is that capital appreciation may not be regularly observable – not all assets can be marked-to-market. For commercial real estate properties this problem is typically addressed by periodic appraisals. Appraisal values have long been criticized for having the effect of artificially smoothing returns, and being potentially poor measures of property liquidation values.[2] Real estate does not generate capital appreciation returns in a regularly observable and immediately realizable process – gains to property investment are "lumpy." This is a well-known and well-studied issue, and the use of IRR mitigates the fundamental problem of periodic return calculation.

A second, more subtle drawback of the time-weighted rate of return is that it fails to account for investment timing. The most striking example of this is the measurement of equity mutual fund investor returns over the period of the last 30 years. Although time-weighed rates of return display a healthy annual performance, returns on a dollar-weighted basis tell a different story. The huge inflows during the late 1990's represent investment by individuals who lost money in the crash of 2000. In an economic sense, dollar-weighted rates of returns are a better measure of the economic scale of the crash in equities following the 1990's.

II.2 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. Irving Fisher, the modern inventor of the net present value criterion for investment decision-making regarded IRR and NPV as equivalent measures with respect to choosing between alternative investments. Hirshleifer (1970) pointed out that IRR actually 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 re-invested at the discount rate. This second condition would seem at first to be inconsistent with an equilibrium intertemporal market rate of interest. The bond yield example above demonstrates why this is not so. Implicit in the market price of a long-term bond is the yield curve – a market expectation of future reinvestment rates on intermediate coupon payments. This argument does not work for unique projects, however. In practice real estate analysis has developed the use of a modified IRR which assumes that intermediate cash flows are re-invested at the short-term riskless rate. Interestingly, our simulation analysis allows us to example 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 being purchased over time for our portfolios,

[2] See Geltner (1977), Fisher and Geltner (2000), and Clayton, Geltner & Hamiltin (2001) for a discussion of the unsmoothing problem.

there can be a period of negative cash flows following a period of positive cash flows as additional properties are purchased, 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 upon 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.

III Simulation Procedure

Our approach to the simulation of portfolio returns is to use 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, for example, In the current analysis, the bootstrap is an ideal tool for analyzing the effects on investing in a set of 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 repeatedly samples data with replacement from an underlying distribution. Under conditions in which 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 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.

III.1 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 in the period 1977 Q4 through 2004 Q2, 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 sub-

[3] For application in real estate, see Goetzmann (1991)

sets from the roughly 4,000 properties in the NCREIF database which have been bought after 1977 Q4 and sold on 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 comprised 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 upon the allocation rule) or to purchase stocks and bonds in policy proportions. Given this basic approach we then vary the number of properties per "history" as well as the type and geographical 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.

It is important to point out that by using passive stock and bond indices we are only capturing variations in IRR due to the changing composition and timing of the real estate portion of the portfolio. If, for example, we presumed that stocks, bonds and real estate could be held in fixed proportion (say, 1/3, 1/3, 1/3) 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, and the effects of financing this activity through re-balancing 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 greater is the component of real estate in the portfolio, the greater is the variation in IRR explained by real estate factors.

One other constraint is important to recognize. We are only sampling from properties with realized total returns within the period. While 27 years would seem to be a long enough time period over which to evaluate performance, real estate is by nature a long-term asset. Pro-forma analysis of returns to property investment is often thought of 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 1970's and early 1980's will be effectively under-represented, since relatively few properties existed in the database at its inception. By the same token, the experience of the 2000's will be under-represented in the database because there are relatively few properties that have been bought in the last few years and then quickly 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 1980's and in the 2000's 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, the distinct advantage of our approach is that it is entirely based upon realized property returns and based upon portfolios comprised 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. In order 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 in this analysis is to understand the particular effects of investing in an asset class subject to integrality constraints (the purchase and sale of whole properties), the limitations on the number of properties in the portfolio, and the effects of purchase and sale decisions occurring at different points in time, we also have some interesting evidence on real estate as an inflation hedge, and a quantitative benchmark for developing an expectation of its real return going forward.

III.3 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 data quarterly on individual properties that includes the acquisition price, net operating income, capital expenditures and market value of each property. When a property is sold, the sale price for the property is reported. Although this data is usually used to calculate quarterly time-weighted returns, all the necessary data is available to calculate the IRR on the property from acquisition to disposition. The data starts in the fourth quarter of 1977 with quarterly data available since that time. Data is made available to members in a manner that insures confidentiality with respect to individual property performance. For the purpose of this study, the database has been standardized to effectively allow the analysis of 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 was also provided for this study. All property information is used for this study was calculated on an unlevered basis.

III.1.a size

The NCREIF database currently includes information on approximately 4,000 properties however the total number of properties reported at one time or another in the NCREIF database since inception is approximately double this number.[4] We are sampling from

---

[4] There is some ambiguity about this number because if a property is sold by an institution to another, it may "leave" the database and then be "added" by the second owner. There are also additional properties that leave the database for other reasons than

the 4,773 properties that have *left* this database due to sale of the property. The total appraised value of the current NCREIF portfolio as of 2004 Q2 was approximately $140 billion. The dollar value of properties we sample from that were sold had an initial acquisition cost totaling about $84 billion and an aggregate sale price of $91 billion.

III.1.b time period

Figures 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 1990's and 2000's. Figure 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 about 10 years. This is consistent with a study by Fisher and Young who found the median holding period for all NCREIF properties was about 11 years. Because we don't 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. But 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 doesn't affect the median. Thus, although we are conditioning upon sale by the first quarter on 2004, we are capturing a holding period that is very similar to the median holding period for all NCREIF properties.

III.1.c 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. On the other hand, there is no reason to believe that the sold properties are un-representative of institutional real estate investments during the past 27 years. Although ex-post the sold properties in the portfolio might have had different returns (or different characteristics from the retained properties) sampling from this group is not likely to be biased in a direction towards 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 lower than average performance. A test of this proposition is a potentially useful exercise to perform in order to calibrate the potential bias. In the current stage of our analysis we have not addressed this issue.

---

sale such as a transfer to another manger. This has little practical effect on our analysis, however.

[5] Fisher, Gatzlaff, Geltner and Haurin (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, Gatzlaff, Geltner and Haurin (2004) use a probit analysis to examine what properties are more likely to sell as a function of economic and property characteristics.

One important feature of the NCREIF sample is that it does not include development properties. All properties are complete when added to the portfolio. This implies that the NCREIF data capture returns over a particular stage in the trajectory 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 companies, 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 in our study. Whether the proportion of each type reflects the underlying rate of institutional investment as a whole is not known, nor do we know whether the distribution of properties by region within the sample is representative of investor activity. On the other hand, the fact that NCREIF is an industry organization intended to collect and provide investment relevant information to its membership is a strong indication that the sample is a representative one.

Table 1A breaks down the current NCREIF portfolio by region and property-type, reporting dollar-values of properties as well as numbers of properties by category. For any single region or property-type it is clear that the current portfolio contains a sufficient number of properties to allow a meaningful bootstrap. Table 1 B shows the breakdown of our sold property sample.

III.1.d discussion

With all of these caveats in mind with regard to sampling and measurement, there is no doubt that NCREIF has the largest, richest and most complete database of commercial property investment ever created. There is simply no comparable source of information like it. It has the particular advantage of including properties actually purchased and sold by institutional investors with the intention of generating returns for their beneficiaries.

III.3 Implementation

Implementation of the bootstrapping procedure requires some detailed choices about the relative scale of the portfolios. The following assumptions were made in calculating the IRRs of diversified portfolios containing real estate, stock and bond assets:

First, the initial investment amount depends on the number of real estate assets that were included in the diversified portfolio; the following dollar amounts were chosen 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:

# 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;

Real estate assets are included in the diversified portfolio in the following fashion: 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 so that aforementioned proportions *p.stocks* and *p.bonds* are preserved.

In 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 maxima towards the middle of the sample period, since our procedure effectively requires all real estate positions to be closed out by the 2004 Q2.

IV Results

IV.1 Real estate only

Table 2 summarizes the results of the basic simulation based on a 1,000 bootstrapped samples each of portfolios comprised 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 index, constructed from the entire database using periodic appraisal data is 9.35% per year. Ciochetti and Fisher (2002) found an equal weighted IRR for NCREIF properties from 1980 to 2001 was 8.73%. As discussed previously, this does not consider the differences in the amount invested for different properties and the exact timing of the cash flows. By contrast, the IRR we calculate is based on pooling the cash flows from all the properties and considering exactly 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.

As discussed earlier, one possible explanation for the difference between the mean IRR discussed above 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 under-represents the performance of properties in recent years that have not been sold. This, in turn, could result in a sample selection bias.

IV.1.a Test for Sample Selection Bias

To test for the possibility of sample selection bias, we created a portfolio that includes all the properties that have ever been in the NCREIF database. We assumed that those properties that were not actually sold were sold at their final appraised value. That is, all properties still being 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 such as transfer to a different manager were assumed to be sold at their last reported appraised value. The acquisition price and quarterly cash flows for these propertieswas then included in the calculation of the IRR for these properties. The IRR for this portfolio that consists of all 11,729 properties that were ever in the NCREIF database was 7.82%. This is just slightly higher than the 7.5% 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 not a significant selectivity bias in using just the sold properties. Thus, based on 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 less than the since inception time weighted return of about 9.4% over the same time period that does not consider the capital flows into and out of real estate over this time period.

IV.1.b Inflation-Adjusted Return

The second half of the sample period was also somewhat unusual in that it was characterized by low inflation relative to the first half.   Inflation protection is one of the main reasons that 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 off inflation over the time period. Instead, we took the quarterly cash flows from the aggregated portfolio of sold properties and discounted each by the realize inflation back to the first time period. We then calculated the IRR based on these inflation-adjusted flows.   The real IRR was 3.99% -- close to a 4% premium over inflation. For the portfolio that includes all properties (sold and unsold) discussed in the previous section, the real rate of return was 4.57%." The slightly higher return makes sense because this 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 premia for exposure to relevant systematic risk factors. Inflation plus 4% appears to be a reasonable rule of thumb.

IV.1.c Spread and Number of Properties

Of even more interest than the difference in IRR vs. time-weighted return *per se* 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 out twenty of these portfolios had returns less than zero. This risk of loss drops significantly as the number of properties increases to 100. With this scale of investment, 95% of the realized IRRs lie within the 5% to 10% range.[6] Table 2 provides a benchmark for how individual institutional portfolio experiences might be expected to have varied over the period 1977 Q4 through 2004 Q2 due to random variations in the properties included 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.

Table 3 reports the IRR values for portfolios comprised solely of properties categorized by region and by property type. For example, the first cell in the left column of the table was created by restricting the randomization only to apartments in the East. The final row of the table is based on random draws across properties in the region. This result will of course depend upon the relative base frequencies of property types by region. The final column is based upon random draws from single property types. This likewise will depend upon the relative base frequency of regional location by property type. These focused portfolios range from 3.8% IRR (Southern Office) to 10.8% IRR (Western Apartment) which are statistically different differences in means for portfolios of 100 properties. An analysis of variance of the table indicates significant differences across property types and insignificant differences across regions.

Table 4 contains the standard deviation of the IRRs for the sub-portfolios. Some types display much greater cross-sectional variation than others. This provides some evidence on a longstanding issue in real estate portfolio analysis, namely, whether regional diversification is as important as diversification across property type. To date, this question has relied upon the 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 vs. property-type diversification. Our approach suggests that, at least ex-post, property type explained more of the variation in performance than did regional classification. The inter-regional differences across IRR by property type were smaller than the inter-property type differences. This is consistent with the results found by Fisher and Liang (2000) who found that property type diversification was more important than regional diversification using 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

---

[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 reversed 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.

well. This would lead to a skewed IRR distribution for that property-type for that region because some of the bootstrap 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 out of 1,000 draws of 16.15%). For samples of 30, the western apartment IRR distributions are unusually positively skewed (with a maximum out of 1,000 draws of 87.65%). Differences in skewness for apartment sub-samples do not appear to be driving the diversification results. In Table 4, the cross-sectional variation for apartments is less than the mean, not greater than the mean. If skewness were responsible for these results we would expect the opposite.

IV.2 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. For this portion of the analysis, we will focus on the IRR distributions of the whole investment portfolio. [7]

As pointed out above, these distributions will be conditioned on the actual historical performance of an index-based stock-bond-real estate portfolio return over the time period from 1977 Q4 through 2/2004. Our procedure only introduces variation in the relative proportions of these assets and in the composition of the 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 re-balancing, 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 vs. the practicalities of modeling a portfolio with properties entering and exiting the portfolio at different times and in different magnitudes is to examine the time series of approximate portfolio weights for a random portfolio in our analysis. Figure 4 is one of a thousand sample portfolio paths from out 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 decreases. Real estate is not market to market, but valued at cost. The lumpiness of this process is apparent from the stepped pattern of the real estate component.

---

[7] As a side-note, this addresses one of the known limitations of IRR, which is the assumption that intermediate cash flows are re-invested 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 modified IRR – substituting a risky portfolio for the riskless asset.

Table 5 reports summary statistics about the IRRs of the portfolios comprised of stocks, bonds and real estate. The top panel contains results for 100 property portfolios, the bottom panel contains results for 10 property portfolios.

Notice that the median return for the simulated portfolios is about 11% for simulations in the top and bottom panels. The cross-sectional variation based on 100 properties (32 BP) is about half of that based on 10 properties (74BP). 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 22%, while the median maximum for simulations with 10 properties is 27%. What this means is that, at some point in the period 1977 Q4 through 2004 Q2 real estate achieved a maximum in the typical portfolio 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 draw-down from the stock/bond holdings. When the drawdown exceeded the other holdings that simulation was omitted from the analysis. Because of the manner of construction, ten-property portfolios are smaller and "lumpier" than 100-property portfolios. This is consistent with the challenges of a small investment fund maintaining a portfolio of investment properties and thus a desirable feature of the simulation. For larger portfolios, and single purchase has a minimal effect – integrality is less of an effect. As a consequence, the difference between the distributions of 10 vs. 100 property IRRs are largely due to integrality and not due to investment timing.

Figure 5 shows four histograms based upon the simulations with 100 properties. It breaks the 1,000 IRR's generated by the diversified portfolios into four groups according to the maximum percentage of real estate represented. The first group, "highest proportion" contains portfolios with maxima over 29%. It clearly has the largest 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 towards what would be a point mass around 11.5%, the IRR provided by 60/40 stocks and bonds over the time period. It is important to point out that the differences between the real estate portfolio returns and the 60/40 portfolio returns are driven by two things. First, there were, in fact, different realized returns across asset classes in this period. The 80's and 90's were two of the greatest decades in the history of the U.S. equity markets, and bonds enjoyed remarkable unanticipated capital gains due to falling rates in this time-period as well. Second, real estate in the simulation is unlevered, while equities, to the extent they had debt on their balance sheets, were levered. Thus the cross-sectional spread between stocks and properties can be scaled somewhat arbitrarily through leverage.

V 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 co-mingled investment funds, 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 indices. The approach in this paper is to simulate the experience of an institutional investor active over the period between 1997 Q4 and 2004 Q2. Instead of using appraisal data, we are able to use transactions and intermediate cash flows to compute the returns to property portfolios that are sub-sets of the actual investable universe, and which, 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 time period, without the potentially confounding effects of appraised values and time-weighted return calculations.

In our study, we find that 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 changes in the realized rates of inflation. To address the obvious macro-economic regime change, we developed a methodology that accounts for shifts in inflation. Our estimate of the inflation-adjusted rate of return for real estate investment over the period was 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 ten properties, there were non-trivial probabilities of a negative IRR. For portfolios with 100 properties the variation decreased significantly but still ranged from 5% to 10%, with an IRR over 10% being an exceptional benchmark. This is important because it suggests that individual portfolio experiences with real estate investing will vary due to the effects of individual property differences and the timing of investment. These differences are largest 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.

To further address the issue of diversification, we also studied the performance of sub-portfolios. In our analysis, we found that mean returns differed by property type more than might be expected, and that cross sectional variation in IRR performance was explained well by property-type. This is an interesting result and merits more analysis.

Finally, we placed these simulated real estate portfolios into the context of a 60/40 stock-bond allocation and simulated the practical process of re-balancing. This involved the financing of real estate purchase at the expense of the other asset classes, and the re-investment of sales proceeds into tradable securities. This generated “modified” IRR’s for investment portfolios of the time period which ranged from 10% to 12%, depending upon the proportion of real estate and number of properties. This approach demonstrated that portofolio IRR’s could vary significantly, depending upon the timing and the diversification of the real estate investment.

It is important to point out that the time-weighted returns are still appropriate to measure the volatility of real estate versus stocks and bonds. This is because the simulations only really capture returns, not risk. However, 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 provided over inflation.

The full portfolio simulation procedure also highlights the role that real vs. liquid portfolio securities play in a dynamic investment portfolio. David Swensen's *Pioneering Portfolio Management*, for example, makes a strong case for less-liquid asset classes as a source of superior investment returns. He argues that institutions such as pension funds and endowments who 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 capacity to do this relies upon the ability to finance lumpy real asset transactions with more liquid capital resources. This structural relationship between marketable securities and real assets in the portfolio is not captured by the usual approaches to portfolio analysis. However the simulations in this paper explicitly take these synergies into account.

References

Geltner, David and William N. Goetzmann, 2000, “Two Decades of Commercial Property Returns: A Repeated-Measures Regression-Based Version of the NCREIF Index” *The Journal of Real Estate Finance and Economics,* 21 (1): pp. 5-21.

Fisher, Jeffrey and Youguo Liang. 2000. “Is Sector Diversification More Important Than Regional Diversification?” Real Estate Finance, Fall issue.

Fisher, Jeffrey and Mike Young. 2000 “Holding Periods for Institutional Real Estate in the NCREIF Database.” Real Estate Finance, Fall issue.

Fisher, Jeffrey and Michael Young. 2000. "Institutional Property Tenure: Evidence from the NCREIF Database” Journal of Real Estate Portfolio Management (6, 4) 1-8.

Fisher, Jeffrey, Dean Gatzlaff, David Geltner and Donald Haurin, 2003, “Controlling for the Impact of Variable Liquidity in Commercial Real Estate Price Indices," Real Estate Economics 31 (2): 269-303.

Fisher, Jeffrey, Dean Gatzlaff, David Geltner and Donald Haurin, 2004, “An Analysis of the Determinants of Transaction Frequency of Institutional Commercial Real Estate Investment Property,” Real Estate Economics 32(2): 239-264

Ciochetti, Brian and Jeffrey Fisher, 2002, “The Characteristics of Commercial Real Estate Holding Period Returns (IRRs)” Real Estate Research Institute Working Paper.

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

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

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

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

Table 1A: NCREIF database. Current holdings by region and property type. Source: NCREIF website.

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

Table 1 B: 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

Table 2: Summary statistics for bootstrapped distributions of IRR values for number of property holdings ranging from 10 over the period 1977 Q4 through 2004 Q2 to 100 over the period. One thousand bootstraps were performed for each row of the table. Every row represents a simulation based on a different number of properties in the portfolio. In each simulation, the given number of properties were drawn from the sample of realized property histories and the cash flows aggregated quarter by quarter. From this, an internal rate of return [IRR] was calculated. Issues of multiple IRR measures were rare. Mean and median values provide measures of the central location of the calculated IRRs. The minimum and maximum IRRs from the thousand simulations are provides, as well as the 5% critical values in the distribution (2.5% above and 2.5% below).

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

Table 3: Mean IRR values for sub-samples by region and by property type. Each cell represents the results of 1,000 bootstrapped portfolios containing 100 properties over the period 1977 Q4-2004 Q2. The final row and final column are diversified with respect to the relative categories. An analysis of variance indicates significant differences across property types and insignificant differences across regions.

	East	Midwest	South	West	USA
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%

Table 4: Standard deviation of IRR values by region and property type. The top panel reports values for portfolios of ten properties. The bottom panel reports values for portfolios of 100 properties. Each cell represents the results of 1,000 bootstrapped portfolios over the period 1977 Q4-2004 Q2. The final row and final column are diversified with respect to the relative categories. An analysis of variance indicates insignificant differences across property types and regions for the ten property panel. The rows are significantly different for the 100 property panel.

Ten Properties in the Portfolio

	East	Midwest	South	West	USA
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	USA
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%

Table 5: Summary statistics of bootstrapped distributions of IRR values for diversified portfolios that include real estate. Top panel reports results for portfolios of 100 properties; bottom panel reports results for portfolios of 10 properties. Stock and bond holdings are set initially at 60/40 and rebalanced quarterly. These results are comparable to the last row in Table 2.

	Portfolio IRR	Maximum % in Real Estate	Real Estate IRR
Simulations with 100 properties			
std	0.324%	8.182%	1.270%
median	11.076%	22.970%	7.463%
max	11.754%	74.373%	13.544%
min	9.018%	8.988%	3.756%

	Portfolio IRR	Maximum % in Real Estate	Real Estate IRR
Simulations with 10 properties			
std	0.747%	20.535%	3.706%
median	11.320%	28.095%	7.559%
max	12.757%	99.189%	22.981%
min	6.079%	5.117%	-7.372%

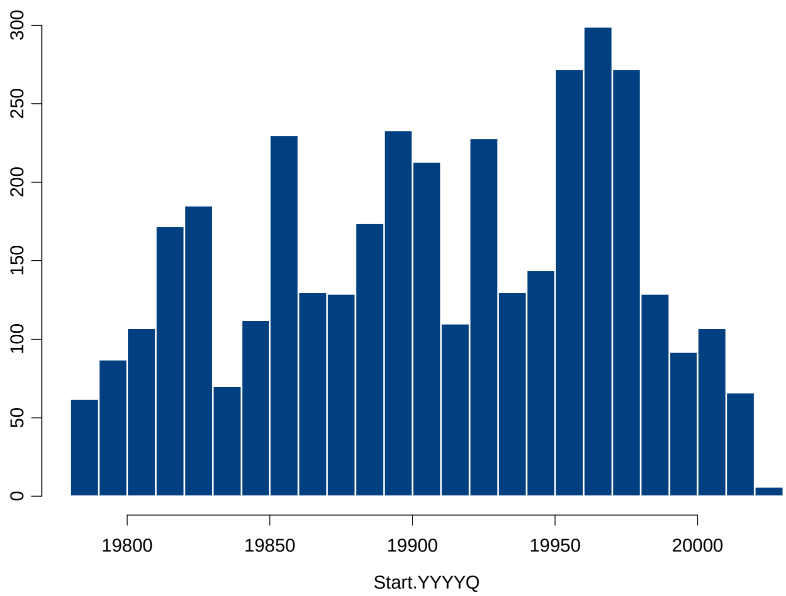

Figure 1 : The distribution of purchase dates of NCREIF properties in the sample used in the study. The format of the number is YYYYQ with 0 indicating start of first quarter.

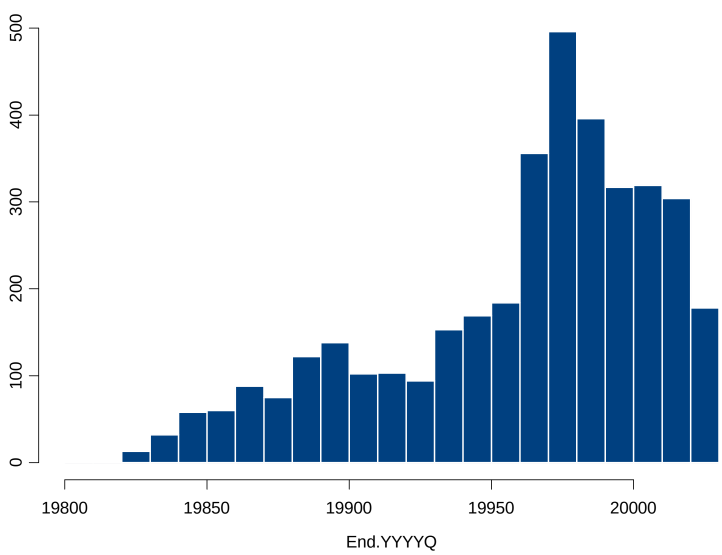

Figure 2 : The distribution of sale dates of NCREIF properties in the sample used in the study. The format of the number is YYYYQ with 0 indicating start of first quarter.

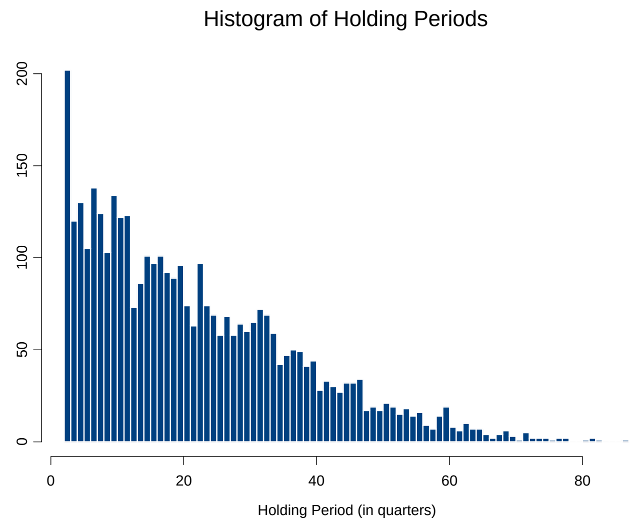

Figure 3: Histogram of holding periods for properties in the sample.

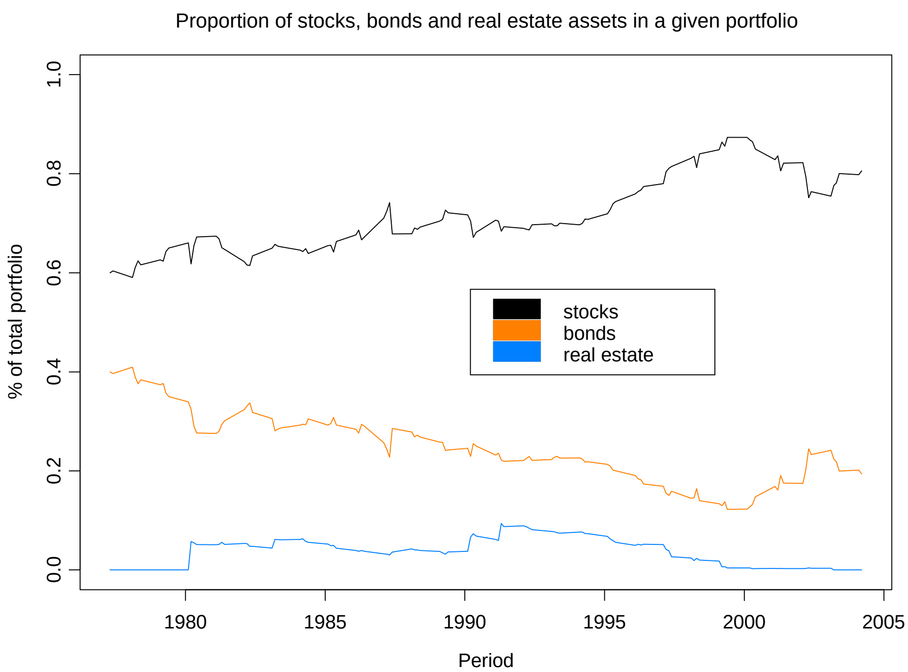

Figure 4: A sample portfolio based on a simulation with ten properties and $100 million initial portfolio value. Stocks and bonds are marked to market while property values are retained at cost.

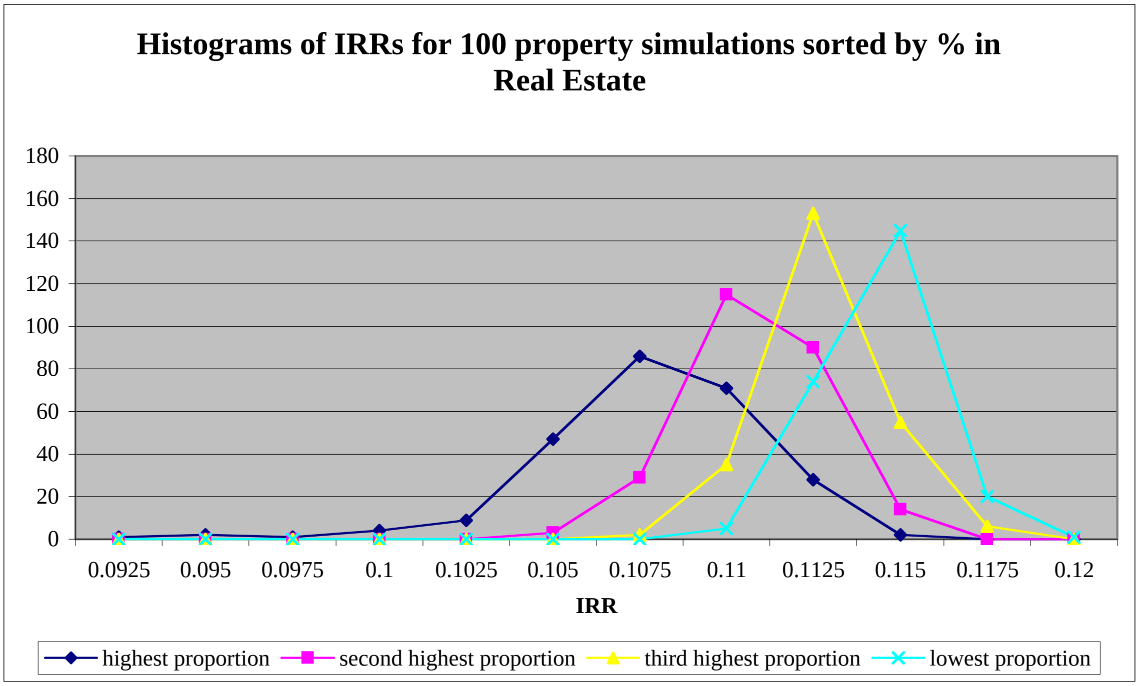

Figure 5: Distributions of IRR’s according to the proportion in real estate. Simulations based on 100 properties in the portfolio. Each histogram is based on 250 observations. The Y axis is cut off at 180 observations. The simulations resulting in the highest proportion of real estate in the portfolio are represented by the category “highest proportion.”