Why the Efficient Frontier for Real Estate Is "Fuzzy"

By
Richard B. Gold
Director of Research
The REIS Reports, Inc.
11 East 36th Street
New York, NY 10016

Why the Efficient Frontier for Real Estate Is "Fuzzy"[1]

Executive Summary. **By moving asset allocation models from the classroom to the boardroom, institutional real estate investors are confronted with a number of practical problems. Many of these difficulties are directly related to both data quality and availability and have received much attention in the current literature. While important, these concerns are perhaps overshadowed by the impact of uncertainty on the allocation process. In this study, uncertainty is introduced by bootstrapping expected returns and standard deviations in a traditional Markowitz framework. The results show that the efficient frontier is not singular, but "fuzzy." That is, this study illustrates that numerous statistically dissimilar weighted portfolios can be equally attractive for any given combination of expected risk and return and that allocation ranges rather than specific targets are, more realistic.**

I. Introduction

Asset allocation models have become quite popular within the institutional real estate research community in recent years. While the role of diversification in reducing portfolio risk has long been recognized, real estate presents analysts with a host of problems not usually encountered with stocks or bonds. Not only has limited data prevented more widespread use of Modern Portfolio Theory (MPT) in conjunction with real estate portfolios, substantive issues regarding appraisal-based values, such as "lumpiness," illiquidity, and transaction costs, have also made use of this technique difficult.

This study will not revisit all the old arguments regarding the applicability of MPT to real estate. Rather the goal of this study is twofold. First, a brief discussion of MPT will be outlined, followed by a summary of the unique problems that real estate presents to practitioners of MPT. Finally, an alternative approach to applying MPT to asset allocation will be discussed.

[1] I wish to thank Dan Dibartolomeo of Northfield Information Services for his guidance and help in preparing this study as well as Brain Webb of Aetna Realty Investors and James R. Webb of Cleveland State University for their comments. Additional thanks go to Harriett Odlum and Sue Miller of Aetna Realty Investors for their editorial insights. An earlier version of this paper was published internally by Aetna Realty Investors in July 1993.

Specifically, by introducing uncertainty into the allocation process, it will be shown that numerous statistically dissimilar weighted portfolios can be equally attractive for any given combination of expected risk and return. Because of uncertainty, it is entirely possible that any two points on the efficient frontier are statistically indistinguishable from each other. Introducing uncertainty into the Markowitz MPT framework will help overcome the practical problem of balancing optimization's goals with real estate's illiquidity, transaction costs, and data problems by creating bundles of portfolios to which investors would be indifferent. Departing from MPT's traditional approach of creating singular combinations of assets for any given level of risk-adjusted returns, this alternative helps quantify the impact of uncertainty by generating acceptable allocation ranges, rather than specific, and difficult to obtain, point estimates.

II. How Does Asset Allocation Work?

In the world of asset allocation, risk is divided into two components: systematic and unsystematic. Systematic risk is the force to which all investments within a market are subject to and must be borne by some investor. Unsystematic risk is confined to the peculiarities of an individual investment and may be diversified away. Market vacancy and rental rates, production, and inflation are all part of systematic risk; while location, asbestos, building age, tenant mix, and the timing of lease rollovers are all examples of unsystematic risk. MPT takes advantage of the diversification effect provided by multiple and less-than-perfectly correlated assets, allowing minimization of unsystematic risk and choice of exposure to systematic risk to desired levels for any given level of return. Efficient market theories suggest that investors are not rewarded for bearing unsystematic risk, and therefore should seek to minimize its presence.

In their pursuit of reducing unsystematic risk, investors diversify their holdings across assets to take advantage of the offsetting effects that less than perfectly correlated asset returns have on volatility.

Of course, diversification is only relevant when risk/uncertainty is present. Sure bets require no diversification since the results are known, *a priori*. If an investor knows with certainty that a particular stock will outperform all other assets during his/her investment horizon, then there is no need to diversify. In the real world, however, there is almost always uncertainty and, thus, the on-going need for diversification.

Investors facing asset allocation decisions must always remember that their goal is to hold a portfolio that will be beneficial to future market conditions. Therefore, all explicit inputs into the asset allocation process need to be based on *expectations* of future returns, varability, and correlations. Historic information is only a benchmark against which forecasts of future performance can be modeled.

Recognizing that not every investor has the same yield hurdle or risk tolerance, MPT provides investors with an infinite number of efficient portfolios. Each of the portfolios along the efficient frontier provides the best expected returns for a given level of risk. It is the representation of the expected risk and returns of these portfolios in a two-dimensional space that forms what is known as the efficient frontier.

III. Data Limitations

Strict interpretation of MPT defines a portfolio efficient if, and only if, it lies directly on the frontier. Given this strict division between efficient and inefficient portfolios, any change in expected returns, risk, and/or correlation between assets could potentially "trigger" a reallocation decision. However, in the real world, real time changes in portfolio allocations are generally regarded as practical. First, data availability varies significantly across asset classes. While stocks are constantly being "marked to market," expectations for future equity real estate returns are usually derived from a much more limited data set. For example, the Russell/NCREIF indices that measure historical returns and risk are only available quarterly and are based on appraised values at best and "guesstimates" at worst.

Second, transaction times, while extremely short in the stock and bond markets, can be excruciatingly slow in other asset classes. For equity real estate, transactions are measured in months, not minutes. In a mixed-asset portfolio, transaction times can have a significant impact on expected returns. This is less important for portfolios assembled with long-term goals, but critical for portfolios with shorter-term strategies. As a result, institutional investors need to be convinced that expected shifts in risks and returns are permanent, not transitory. Clearly, it would be both frustrating, as well as costly, to make major changes to one's portfolio as a result of temporary changes in market conditions. Illiquidity also makes direct comparisons between real estate and other asset classes more difficult. How much of a return premium is required to compensate real estate investors for illiquidity and how much does it expand or contract during the business cycle?

Third, equity real estate transaction costs are significant. Therefore, asset allocation decisions need to carefully consider the costs, as well as the benefits, of any transaction.

Fourth, equity real estate transactions are not typically available for ownership units smaller than a single building or parcel of land. In comparison, individual shares of stocks can be bought or sold to achieve a precisely desired mix of stocks. While REITS and other forms of securitization try to address this problem, they are not perfect substitutes for equity ownership.

Finally, data quality issues remain a significant problem for equity real estate. For example, the appreciation component of the Russell/NCREIF Index is primarily derived from appraisals rather than transactions. Moreover, most properties are only appraised annually, further reducing real estate's apparent volatility as represented by the index. At the same time, this smoothing effect boosts real estate's apparent risk-adjusted returns and dampens its correlation with other assets. These biases, all else being equal, have a tendency to artificially raise equity real estate's share within a mixed-asset portfolio.

Irrespective of these measurement problems, investors recognize that real estate is too large an asset class to be ignored. Rather than dealing directly with the difficulty of using MPT-based asset allocations models because of the associated data problems, many institutional investors peg real estate's allocation within a mixed-asset portfolio at some constant percentage or fixed dollar amount. Within real estate, property-type and geographic allocation decisions are often accomplished by even more generic approaches. While expeditious, fixed allocation or intuitive strategies are risky. When the underlying dynamics inherent in mixed-asset portfolios are ignored, sub-optimal results are likely resulting in the investor bearing more risk than necessary to acheive a desired rate of reutrn. Therefore, the challenge is to find a methodology that incorporates existing return data but does not ignore the data's known limitations.

IV. Fuzzy Frontiers

MPT stipulates that only one portfolio can occupy any given point along the efficient frontier. However, what if the frontier was actually composed of multiple parallel and statistically indistinguishable curves? Moreover, what if these parallel curves contained portfolios with similar asset weightings as those found on the original unperturbed frontier? If this were the case, then multiple solutions would be possible for any given level of risk tolerance. With expectations so heavily dependent on an accurate accounting of past performance, it is difficult to believe that a single portfolio can compensate for the uncertainties of either the real estate return data or capture the "fuzzy" nature of the forecasts of returns and risk. "Fuzziness" is merely admitting to ourselves that our forecasting abilities are imperfect and, by definition, conditional forecasts of future returns and risk are more likely to be wrong than right over any significant time horizon. Since small errors in expected risk, return, or correlation have the potential to significantly change the composition of portfolios along the frontier, uncertainty casts a wide shadow over the entire allocation process.

Exhibit 1
Expected Returns and Standard Deviations For Appropriate Horizon

	RETURN	STANDARD DEVIATION
Office	9.5%	6.6%
Retail	8.2%	3.5%
Apartments	8.3%	3.3%
Industrial	8.5%	4.0%

To illustrate the extent to which uncertainty is a factor for asset allocators, the quarterly Russell/NCREIF total return data from 1978 to 1992 for office, retail, apartment, and industrial properties were recreated using a statistical technique called "bootstrapping". [2]

[2] While the Russell/NCREIF data is quarterly, it was converted to a monthly frequency for this study by linerally interpolating the quarterly observations. Because the resulting "fuzziness" factors were created from interpolated quarterly data, this had a smoothing effect on the volatility of the returns. Since the Russell/NCREIF data is smooth to begin with, it was felt that interpolation did little to significantly alter

Bootstrapping was considered the most appropriate methodology for recreating the true distribution of future returns, given the limited history of the NCREIF data, and therefore the limited volatitlity of the series. The problem is less acute for stock and bond returns, not only because they are high frequency, transaction-based series, but because data are available over a significantly longer period of time and across a wider range of economic conditions.

Exhibit 2
Expected Property-Type Correlations

	Office	Retail	Apartment	Industrial
Office	1.00			
Retail	.78	1.00		
Apartment	.64	.72	1.00	
Industrial	.87	.82	.67	1.00

Generally, the bootstrap methodology has a tendency to more evenly distribute the intra-asset weightings within a portfolio. Assets with the largest weights will have their allocations reduced, while assets with the smallest weights will receive a higher allocation. This is because in a world in which an investor knows nothing about the behavior of assets, he or she will tend to reduce risk by equally weighting their investments across assets. The bootstrap approach simply models an investor's response to uncertainty regarding asset behavior.

---

the results. Monthly data is also the most commonly used frequency for inter-asset (i.e. stocks, bonds, cash, and real estate) allocation models.

Exhibit 3 Base Case Results						
			Percent Portfolio Allocation			
RAP	Return	Std Dev	Office	Retail	Aptmt	Ind
1	8.2%	3.1%	0%	26%	64%	10%
2	8.3%	3.2%	0%	15%	64%	21%
3	8.3%	3.2%	0%	4%	64%	32%
4	8.4%	3.4%	7%	0%	63%	30%
5	8.5%	3.6%	19%	0%	61%	20%
6	8.7%	3.9%	31%	0%	59%	10%
7	8.8%	4.3%	43%	0%	57%	0%
8	8.9%	4.5%	51%	0%	49%	0%
9	9.0%	4.8%	58%	0%	42%	0%
10	9.1%	5.1%	66%	0%	34%	0%
11	9.2%	5.5%	74%	0%	26%	0%

In order to use the information generated by bootstrapping, a base case efficient frontier was created using a 60 month holding period with a purely illustrative data set of "expected" property-type performance and correlations (Exhibits 1 and 2). These returns, standard deviations, and correlations were created for the sole purpose of this simulation. In fact, to emphasis the illustrative nature of this base case scenario, the expected returns, standard deviations, and correlations were adjusted to insure a large number of mixed-asset portfolios at numerous points along the efficient frontier. The resulting allocations can be seen in Exhibit 3 where the first column represents a progressively higher risk acceptance point (RAP) along the estimated efficient frontier. The higher the RAP, the more agressive the investment strategy and with it a higher risk profile.

Next, using the bootstrap technique, 1,000 alternative representations of returns and standard deviations for each of the four property types using the previously described Russell/NCREIF total return data were created. These "pseudo" return series were created by randomly selecting, with replacement, 60 dates from the historic sample period. Since the selection process used replacement, it was entirely possible that the same date was chosen multiple times. Because the same selection pattern was used across property types for any given "pseudo" series, the underlying correlations between assets were undisturbed.

These "fuzzy" estimates can be viewed as dispersion measures of the original time series and provide an alternative representation of risk and return for each of the four property types. By observing the variability in the means and standard deviations, relative to their actual history, an investor is able to get some sense of where the observed mean fell, relative to the "boostrapped" mean.

Exhibit 4 shows the resulting "fuzziness" factors as measured by the standard deviation in the means and standard deviations of the historic return series. For example, mean returns for office properties had a "fuzziness" factor (i.e., standard deviation) of 23 basis points per month in its mean, and a 26 basis point difference in its standard deviation. The relative size of a "fuzzy" factor is a direct function of the underlying standard deviation of the original series.

While these "fuzziness" factors are interesting by themselves, their ultimate purpose was to measure their impact on property type allocation. With the help of an asset allocation software package, 250 alternative efficient frontiers were created changing each property type's base case expected return and standard deviation by an adjustment factor.[3]

Exhibit 4 Fuzziness Factors Standard Deviations From Base Case (Basis Points Per Month)		
	Mean	Standard Deviation
Office	23	26
Retail	12	13
Apartments	17	17
Industrial	13	14

The result is a series of 2,750 new portfolios. Each of these new portfolios has its own property type mix and risk/return profile. However, because they are parallel to the original unpeturbed frontier, cross sectional analysis is possible. For this study, the characteristics of both the base case and "fuzzy" portfolios across eleven progressively higher risk, and therefore higher return points along the frontier were analyzed. The further out along the frontier an investor moves, the more risk he/she is willing to bear to receive higher returns.[4]

[3] These factors were calculated by taking a property type's respective mean and standard deviation "fuzziness" factor from Exhibit 4 and multiplying it by a randomly generated, but normally distributed, set of scalars with a mean of zero and standard deviation of one. Because these scalars were normally distributed, most of the scalars, and by extension, adjustment factors centered around zero, with only a handful of the scalars and their corresponding adjustment factors approaching the bounded limits of the distribution. For each of the 250 alternative frontiers, two scalars were generated. The first was applied to the fuzzy factors for all the property type means, and the second was applied to all the fuzzy standard deviation factors. Northfield Information Services' PACO software was employed throughout this study.
[4] The PACO software produces eleven Risk Acceptance Points (RAPs) per frontier. Therefore, the study produced 250 per Risk Acceptance Point or 2,750 (11 X 250) in total.

The next step was to calculate the standard deviation of each property type's allocation for the 250 portfolios at each of the 11 Risk Acceptance Points (RAPS) as well as the standard deviations for the expected return and risk for each of these 11 points. This provided the necessary information to calculate return and property type-specific asset allocation ranges over which an investor would be indifferent. In this study, ranges are defined as the base case mean, plus or minus one standard deviation, as calculated from the previous step.

EXHIBIT 5

"Fuzzy" Mean Returns and Standard Deviations Relative to Base Case

	Returns		Standard Deviation	
RAP	**Base**	**Fuzzy**	**Base**	**Fuzzy**
1	8.2%	8.4%	3.1%	3.2%
2	8.3%	8.4%	3.2%	3.2%
3	8.3%	8.5%	3.2%	3.4%
4	8.4%	8.6%	3.4%	3.5%
5	8.5%	8.8%	3.6%	3.8%
6	8.7%	8.9%	3.9%	4.0%
7	8.8%	9.0%	4.3%	4.3%
8	8.9%	9.0%	4.5%	4.6%
9	9.0%	9.2%	4.8%	4.8%
10	9.1%	9.3%	5.1%	5.1%
11	9.2%	9.4%	5.5%	5.4%

Exhibits 5 and 6 show the baseline expected returns, standard deviations, and property type allocations as well as a one "fuzzy" standard deviation confidence interval around the base line results. For example, at RAP 6 the base case frontier shows an expected return of 8.7% with a standard deviation of 3.9%, and the expected mean "fuzzy" return and standard deviation are 8.9% and 4.0% respectively. The small absolute differences between the expected base and mean "fuzzy" returns and standard deviations are to be expected since the "fuzzy" returns and standard deviations in Exhibit 5 represent the means of distributions whose values fall both above and below the base case. What these small differences do show is the nonlinearity of the assest allocation process. Otherwise the "fuzzy" means would have approximated the base case returns and standard deviations.

However, significant differences between the base and "fuzzy" cases do arise if the standard deviations of the "fuzzy" expected returns and risk are employed. For example, the one "fuzzy" standard deviation confidence interval around the base line case shows an expected return range of 8.1% at the low end and 9.3% at the upper end.[5] At the same time the expected risk would have a confidence band of 3.5% to 4.3%. Exhibit 7 shows the one standard deviation risk/return confidence interval for all 11 risk acceptance points.

To receive the base case expected return at RAP 6, an investor's portfolio should be composed of 31% office, 0% retail, 59% apartments, and 10% warehouse using the expected returns, standard deviations, and correlations in Exhibits 1 and 2. However, by banding each one of these allocations by plus or minus one "fuzzy" standard deviation, an investor would target, as well as be indifferent to, an allocation range of between 11 and 51% for offices, 0 to 11% for retail, 43 to 75% for apartments, 0 to 35% for industrial properties.

[5] This is determined by calculating the standard deviation of expected return and risk at each Risk Acceptance Point using the results of the 250 bootstrapped frontiers. Therefore, while the mean bootstrapped returns and standard deviations shown in Exhibit 5 closely resemble the Base Case results, it is the standard deviation of these results which are used.

Since the ranges are not additive, investors can not arbitrarily pick points within the ranges which sum to 100% and necessarily be within the confidence interval. However, they could begin by targeting an allocation range for a particular property type and eliminating all portfolios from their set of base and "fuzzy" portfolios that do not have that initial allocation range for the property type in question. The process could then be repeated for a second property type and again for the third and fourth, if desired, using progressively smaller subsets of portfolios.

In our example for RAP 6, apartments are the dominant property type. Assuming an investor would prefer a portfolio composed of 50 to 60% apartments, he/she finds that there are 80 out of the 251 portfolios that meet this criterion. By choosing a target range of 20 to 30% offices, the second dominant property type, they find that the number of possible portfolios has been reduced to 13. Finally, similar screens could be performed for the remaining property types, if desired, to reduce the number of portfolios even further.

What are these bands and ranges telling us? First, it is apparent that uncertainty about the quality of expectations significantly reduces our ability to perform asset allocation with any precision. Within each banded risk acceptance point, there are an infinite number of portfolios to which an investor would be statistically indifferent. Second, these results provide investors with some additional insights into the problems created by appraisal-based series. Finlly, these results reinforce the importance and value of expanding our real estate knowledge base.

Exhibit 6
Base Case Risk/Return Confidence Interval
+/- 1 "Fuzzy" Standard Deviation

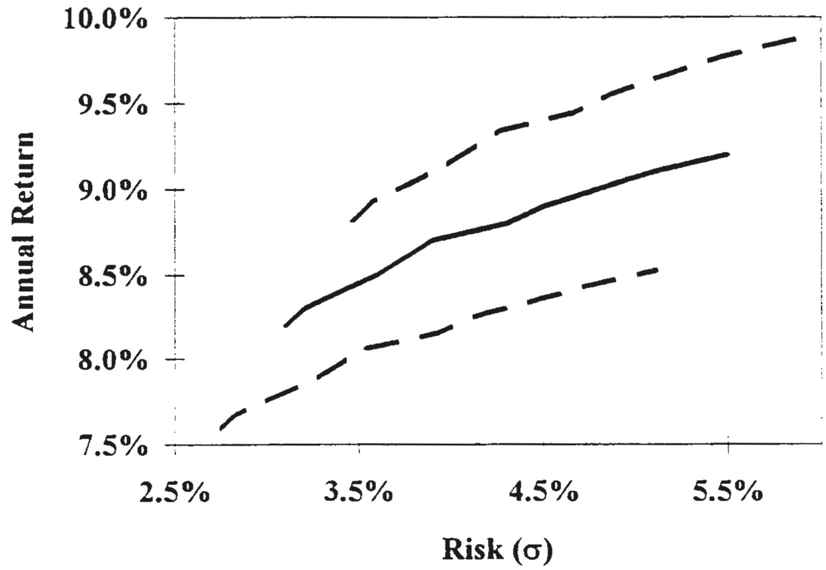

Clearly, for an investor who is holding nonsecuritized and "lumpy" assets such as real estate, investment ranges are easier to obtain than strict point estimates. Not only is it more practical, but it is consistent with the nature of the asset class, since the timing of sales and acquisitions can overlap and buildings are not available in divisible units. While some of these problems will ease with the increasing availability and acceptance of securitized real estate, equity ownership is certainly not on the verge of extinction. Therefore, "fuzzy" frontiers provide portfolio managers with needed flexibility within a rigorous investment framework.

Using risk and return ranges also partially addresses the problem of portfolio rebalancing in the light of real estate's significant transaction costs. For example, what if a portfolio is outside its respective RAP range or if new information becomes available that may trigger the need to rebalance? Is it return maximizing to rebalance? The answer is, it depends.