2009 V37 1: pp. 117–153
REAL ESTATE
ECONOMICS

Real Estate for the Long Term: The Effect of Return Predictability on Long-Horizon Allocations

Gregory H. MacKinnon* and Ashraf Al Zaman**

We examine how the predictability of real estate returns affects the risk of, and optimal allocations to, real estate for investors of differing investment horizons. Returns to direct real estate are mean reverting, and risk decreases with horizon. This is driven by a tendency for property transaction prices to overshoot inflation. Mean reversion in real estate returns is weaker than that of equities, resulting in real estate having similar risk to equities for long-term investors. However, optimal portfolios have large allocations to direct real estate at all horizons, and the allocation increases with horizon. Finally, we find that real estate investment trusts are a redundant asset class for investors with access to direct real estate as an asset class, but they do have a role in optimal allocations when direct property investment is not feasible.

What is the optimal allocation to real estate for an investor? Significant resources have been devoted to this question by institutional investors and academics alike. Seiler, Webb and Myer (1999) review the early academic literature on optimal real estate allocations and find answers that range from 0% to 67%. More recently, Feldman (2003) estimates optimal allocations to commercial real estate (direct plus real estate investment trusts (REITs)) ranging from 0% to 42% depending on the direct real estate index and target return chosen. Looking at the actual allocations of pension funds, Clayton (2007) reports that funds have a mean allocation to real estate of 7.3% as of 2006 and that both actual and target allocations have been increasing. Given the range of answers and the discrepancy between institutional investors' actual allocations and many of the prior studies, this remains a question of interest and importance.

Most prior studies of optimal real estate allocations rely on traditional mean-variance optimization using short-term returns (*e.g.*, monthly or quarterly). If one assumes that returns are i.i.d., then risk (variance) and return will both

*Department of Finance and Management Science, Sobey School of Business, Saint Mary's University, Halifax, Nova Scotia, Canada B3H 3C3 or greg.mackinnon@smu.ca.
**Department of Finance and Management Science, Sobey School of Business, Saint Mary's University, Halifax, Nova Scotia, Canada B3H 3C3 or ashraf.zaman@smu.ca.

scale with time so that the optimal portfolio is independent of the horizon and conclusions based on short-term returns can be applied to any investment horizon. However, it is widely acknowledged that the real estate market is not informationally efficient and returns are not i.i.d. Barkham and Geltner (1995) show that securitized real estate markets lead direct markets and conclude that direct markets are to some extent informationally inefficient. Looking at publicly traded real estate securities, Mei and Liu (1994) find evidence of predictability and show that market timing strategies can generate excess returns. This return predictability means that optimal real estate allocations may be different for long- and short-term investors, an issue ignored by traditional mean-variance optimization.

We adopt the methodology of Campbell and Viceira (2005a) to examine optimal real estate allocations. This approach explicitly accounts for predictability in asset class returns using a vector autoregressive (VAR) framework. Accounting for predictability has two main effects: (1) risk is defined based on only the unpredictable component of returns, reducing the risk level of predictable asset classes, and (2) asset class risk (and therefore the optimal allocation) differs depending on investment horizon. Campbell and Viceira (2005a) refer to the second effect as the "term structure of risk." This issue is of particular importance for institutional investors whose horizons are typically long term and do not match the short-term returns traditionally used in portfolio optimization. Based on the Campbell and Viceira (2005a) methodology we examine the predictability of commercial real estate returns and whether this translates into significant changes in real estate risk and optimal allocations as the investment horizon lengthens. Note that, despite the fact that we are examining the effects of predictability, our interest is not in market timing. Rather, we examine whether predictability of real estate returns affects the optimal asset allocation of a buy-and-hold investor with a long investment horizon.

Given the continuing debate on the relationship between direct real estate and REITs as well as the role of REITs as a proxy for direct real estate investment in a portfolio, we also examine the optimal allocations to REITs under two conditions: when both direct and securitized real estate can be included in portfolios and when direct real estate is not a feasible asset class for an investor.

Our results show that direct real estate returns exhibit a mean reverting pattern. Mean reversion seems to be driven by property transaction prices overshooting inflation. The effect of the mean reversion is that real estate risk decreases with investment horizon. However, mean reversion for real estate is weaker than for equities. This means that, for longer-term investors (10-year or greater investment horizon), real estate has risk equal to that of equities. This contradicts the typical view of real estate as a low-risk asset class, at least for long-term investors. Optimal allocations to real estate, however, are quite large at all investment horizons and largest for longer-term horizons. This is driven by the fact that the diversification benefits of real estate grow with investment horizon.

The correlation between direct real estate and REITs increases with horizon, although it never passes 0.54. The correlation between REITs and direct real estate is high enough to make REITs redundant as an asset class when direct real estate investment is an option for investors. However, for small or short-term investors for whom direct property may not be a feasible investment, an allocation to REITs does play a role in optimal portfolios.

The remainder of the article is organized as follows. The next section provides background on the issue and reviews the literature. The third section provides details on the methodology. The fourth section describes the data and provides some initial results on the predictability of real estate returns. The fifth section contains our main results and the sixth concludes.

Background and Literature Review

The Campbell and Viceira (2005a) methodology for estimating long-horizon optimal allocations is an intuitively appealing one as it is simply a variant of the familiar mean-variance analysis. We assume that investors are mean-variance optimizers, choosing a portfolio that minimizes variance subject to attaining a target expected return. Investors act myopically in that they do not consider multiple periods in their decision making. Based on a one-period analysis, investors choose an optimal buy-and-hold portfolio. This is the same as the familiar Markowitz (1952) portfolio optimization discussed in any introductory finance textbook. The key difference, however, is the definition of a "period" used in the analysis.

The appropriate length of the period coincides with the investment horizon of the investor, and therefore different investors will use different definitions of a period. A short-term investor, for instance, might define a period as six months or a year whereas a long-term investor might define a period as 25 years. Both are assumed myopic and optimize over a one-period interval. It is apparent, however, that "one period" can in some cases refer to decisions over quite long horizons. The resulting optimal allocations are for buy-and-hold investors, implying no rebalancing of the portfolio. While this is a simplification, we do not believe it to be entirely unrealistic. Our main emphasis is on the role of direct real estate in the portfolio. Given the difficulty and cost involved in property transactions, it is unlikely that investors, *ex ante*, would intend to rebalance real estate allocations frequently.

Whereas different investors have different horizons, all investors will have the same optimal portfolio if asset returns through time are i.i.d. Conversely, predictability of returns can drive a wedge between the optimal allocations of long- and short-horizon investors.

Consider risk on a per-period basis (return variance over the horizon divided by the number of periods in the horizon). If returns are positively autocorrelated then the asset will have greater per-period risk for long-term investors than for short-term investors.[1] The "term structure of risk" will slope up for this type of asset, and long-term investors will tend to hold lower allocations to it than short-term investors. This would be the result of momentum in returns over time (mean aversion). If returns are negatively autocorrelated, then per-period risk will be less for longer-term investors and they will tend to hold more of the asset in their portfolios relative to short-term investors. This is the result of mean reversion in returns; the statistical analogy to the common financial advice that the "ups and downs of the market will cancel out in the long run."

While the discussion above is simplified, it illustrates the basic principle. Rather than simple autocorrelations, Campbell and Viceira (2005a) define predictability based upon a VAR estimated on asset classes and state variables. Still, the essential result holds, namely, that mean reversion will act to decrease risk with investment horizon and mean aversion will act to increase it. This approach therefore allows us to determine if the optimal real estate allocation within a mixed-asset portfolio differs for long- and short-term investors.

We are not the first to consider the effect of long investment horizons on optimal real estate allocations. Lee and Stevenson (2005) examine the optimal allocation to REITs in a mixed-asset portfolio over 5-, 10-, 15- and 20-year time intervals. However, they employ annual returns for all estimations so that the implied investment horizon is always 1 year but estimates are calculated over data sets of differing lengths. Mueller and Mueller (2003) examine allocations to both direct and securitized real estate over 5-, 10-, 15-, 20- and 25-year intervals but again annual (or quarterly) returns are used in estimation.

Geltner and Rodriguez (1997) and Geltner, Rodriguez and O'Connor (1995) both discuss how predictability in real estate returns can lead to optimal

[1] Consider a two-period example with returns in each period, R_1 and R_2, in continuously compounded form. The total return over two periods is $R_1 + R_2$. Expected return therefore scales linearly with time. A short-term investor with a one-period investment horizon will measure risk as $\mathrm{Var}[R_1] = \sigma^2$. A long-term investor with a two-period investment horizon will measure risk as $2\sigma^2 + 2\mathrm{Cov}[R_1, R_2]$. Hence, risk will scale linearly with horizon only if returns are independent through time.

allocations differing with investment horizon. They derive optimal portfolios over a 5-year horizon, including the effect of predictability. A regression relating direct real estate returns to public market factors and lagged direct real estate returns is estimated and the parameters used to construct expected returns and variances for real estate over a 5-year horizon. They find that both direct and securitized real estate appears in the optimal portfolio for almost all target returns. More conservative portfolios have a larger weight on direct real estate and lower on public real estate, but this reverses as more aggressive portfolios are considered. The methodology of the Geltner, Rodriguez and O'Connor (1995) article is similar in spirit to that of Campbell and Viceira (2005a); however, they measure risk as the variance of 5-year returns rather than of the unpredictable component only, and they assume that all public markets are efficient and do not exhibit predictability. In this article a VAR system is employed which incorporates predictability in all markets as well as predictability due to external state variables. Also, Geltner and Rodriguez (1997) and Geltner, Rodriguez and O'Connor (1995) examine only one investment horizon, 5 years, whereas our main interest is in how allocations may change over different horizons.

Both Fugazza, Guidolin and Nicodano (2007) and Hoevenaars, Molenaar, Schotman and Steenkamp (2008) employ VAR systems to examine optimal allocations, including real estate. Fugazza, Guidolin and Nicodanna (2007) examine allocations among European stocks, bonds, cash equivalents and publicly traded European real estate securities. For a buy-and-hold investor ignoring parameter uncertainty, they conclude that real estate is more important to the portfolios of long-horizon European investors. The optimal allocation to real estate grows from 9% of the portfolio at a 1-year horizon to 44% at a 10-year horizon (for an intermediate level risk aversion parameter). Investigating the causes of this, they find that real estate exhibits mean aversion in that real estate risk grows with horizon faster than it would if returns were i.i.d., but that expected per-period excess returns on real estate increase with the horizon. On balance the return effect dominates and optimal allocations to real estate grow with investment horizon.

Hoevenaars *et al.* (2008) begin with the Campbell and Viceira (2005a) approach to return predictability and examine horizon effects on portfolios in both asset-only and asset-liability contexts. They examine portfolios of cash, bonds and stocks along with various alternative investment classes, including REITs. They find that the term structure of risk for REITs displays mean aversion for horizons up to 4 years, but, unlike the Fugazza, Guidolin and Nicodanna (2007) finding for European listed real estate, mean reversion dominates in the long run with per-period risk decreasing beyond a 4-year horizon. They also report that risk for public market real estate is greater than that for stocks, especially at long horizons. This is unlike the common finding in the literature that real estate is less risky than equities. Hoevenaars *et al.* (2008) report optimal allocations to real estate that are either zero or very small in all cases and horizons examined.

While Hoevenaars *et al.* (2008) and Fugazza, Guidolin and Nicodanna (2007) employ very similar methodologies to us, they both examine only publicly traded real estate securities. Direct and indirect real estate investments are known to have significantly different return properties, at least for short horizons (see, e.g., Clayton and MacKinnon 2003). Our emphasis in this article is on investment horizon effects on direct real estate investments, as well as on the roles played by the two types of real estate investment in a portfolio and whether those roles change for long- versus short-horizon investors.

Data and Methodology

Our methodology for looking at the term structure of risk and estimating optimal allocations over long horizons is that of Campbell and Viceira (2005a). We provide an overview here; technical details can be found in Campbell and Viceira (2005b).

We begin with a buy-and-hold investor making a portfolio decision over an investment horizon that is k periods in length. The investor chooses a mean-variance optimal portfolio where portfolio moments are estimated over the horizon. Let R_{0t} be the simple return to cash equivalents in period t and R_{it} be the simple return to asset class i in period t. In the model we work with returns in real terms and, to facilitate the summation of returns over time, in continuously compounded form. Let $r_{0t} = \ln(1 + R_{0t}) - \ln(1 + \pi_t)$ be the real return to cash where π_t is inflation. Although cash is used as a benchmark, its return does vary over time so there is no entirely risk-free asset. Returns for asset class i are defined as excess returns over cash, $x_t = \ln(1 + R_{it}) - \ln(1 + \pi_t) - r_{0t} = r_{it} - r_{0t}$. Note that x_{it} is the excess return to asset class i in both real and nominal terms.

All variables of interest are aggregated in an $m \times 1$ vector $\mathbf{z_t}$, where bold indicates a vector or matrix.

$$\mathbf{z_t} = \begin{bmatrix} r_{ot} \\ \mathbf{x_t} \\ \mathbf{s_t} \end{bmatrix}. \tag{1}$$

$\mathbf{x_t}$ is the $n \times 1$ vector of excess return on the n asset classes. Hence, including cash equivalents the investor chooses a portfolio from amongst $n + 1$ risky assets. $\mathbf{s_t}$ is a vector of $m - n - 1$ state variables used to predict asset class returns.

Predictability in returns is modeled as a VAR on $\mathbf{z_t}$ of the form:

$$\mathbf{z_{t+1}} = \mathbf{\Phi_0} + \mathbf{\Phi_1 z_t} + \mathbf{v_{t+1}} \tag{2}$$

where $\mathbf{\Phi_0}$ is an $m \times 1$ vector on constants and $\mathbf{\Phi_1}$ is an $m \times m$ matrix of slope coefficients. The error vector, $\mathbf{v_{t+1}}$, is assumed to be i.i.d. normally distributed with zero mean and a covariance matrix of $\mathbf{\Sigma_v}$.

The covariance matrix of the errors is important to the optimization process as it indicates the contemporaneous relationships among the unexpected shocks to the asset classes and state variables. Of special importance are the first $n + 1$ diagonal terms which give the variance of unexpected shocks to cash and the other asset classes. As risk is defined over the unpredictable component of returns; these diagonal terms represent the relevant one-period risks for each of the assets.

Estimation of the VAR parameters in (2) and the corresponding covariance matrix, $\mathbf{\Sigma_v}$, are assumed to capture the predictability of asset class returns. While estimation of (2) is done using data of a particular periodicity (*e.g.*, we use quarterly data in our estimation), assuming that these estimates capture predictability in asset returns and that the parameters are constant through time allows us to calculate return moments over any horizon. In other words, although we use quarterly data to estimate the model, the resulting parameters can be used to calculate conditional moments for a horizon of *k* quarters, where *k* can be arbitrarily large.

For an investor with a horizon of *k* periods the variable of interest as of time *t* is $\sum_{i=1}^{K} \mathbf{z_{t+i}}$ as the asset class excess returns over a *k*-period horizon are the summation of subperiod excess returns (because returns are continuously compounded). Campbell and Viceira (2005b) derive the conditional expectation of $\sum_{i=1}^{K} \mathbf{z_{t+i}}$ over a *k*-period horizon and show that it is dependant on initial conditions (*i.e.*, on the realization of $\mathbf{z_t}$). Hence, as would be expected in the presence of predictability, expected returns depend on the most recent observed returns. Fugazza, Guidolin and Nicodanna (2007) take this into account and find that the expected per-period return to European real estate securities increase with investment horizon for their sample. However, using this conditional expectation makes any optimal portfolios time specific as they depend on the most recently observed realizations of the asset classes and state variables.[2] We follow the approach of Campbell and Viceira (2005a) and use the unconditional mean across our sample as the estimate of per-period expected excess return for each asset class. For each asset class, the unconditional expectation of

[2] Fugazza, Guidolin and Nicodanna (2007) account for this in their work by initializing returns to their unconditional mean.

the k-period horizon return is simply this mean times k. Thus, we examine optimal allocations under "average" market conditions, an approach which may provide "sensible choices for policy portfolios of long-term institutional investors" (Campbell and Viceira 2005b, p. 27).

The covariance matrix of $\sum_{i=1}^{K} \mathbf{z_{t+i}}$ can be shown to be

$$\begin{aligned}\mathrm{Var}\left[\sum_{i=1}^{k} \mathbf{z_{t+i}}\right] = {} & \boldsymbol{\Sigma}_\mathbf{v} + (\mathbf{I} + \boldsymbol{\Phi}_\mathbf{1})\, \boldsymbol{\Sigma}_\mathbf{v}\, (\mathbf{I} + \boldsymbol{\Phi}_\mathbf{1})^T + (\mathbf{I} + \boldsymbol{\Phi}_\mathbf{1} + \boldsymbol{\Phi}_\mathbf{1}\boldsymbol{\Phi}_\mathbf{1}) \\ & \times \boldsymbol{\Sigma}_\mathbf{v}\, (\mathbf{I} + \boldsymbol{\Phi}_\mathbf{1} + \boldsymbol{\Phi}_\mathbf{1}\boldsymbol{\Phi}_\mathbf{1})^T \\ & + \cdots + \left(\mathbf{I} + \boldsymbol{\Phi}_\mathbf{1} + \boldsymbol{\Phi}_\mathbf{1}\boldsymbol{\Phi}_\mathbf{1} + \cdots + \boldsymbol{\Phi}_\mathbf{1}^{\mathbf{k-1}}\right) \\ & \times \boldsymbol{\Sigma}_\mathbf{v} \left(\mathbf{I} + \boldsymbol{\Phi}_\mathbf{1} + \boldsymbol{\Phi}_\mathbf{1}\boldsymbol{\Phi}_\mathbf{1} + \cdots + \boldsymbol{\Phi}_\mathbf{1}^{\mathbf{k-1}}\right)^T \end{aligned} \tag{3}$$

where $\mathbf{I}$ is an identity matrix. Note the covariance matrix does not depend on initial conditions. We define the upper-left term of the covariance matrix as σ_0^2, the variance of real cash returns. $\Sigma_\mathbf{xx}$ is defined as the $n \times n$ submatrix representing the covariance matrix of the other asset class excess returns. $\Sigma_\mathbf{xx}$ depends on the horizon (*i.e.*, k) that is chosen. An examination of the terms in diag$\{\Sigma_\mathbf{xx}\}$ for different values of k will reveal the term structure of risk: how the risk of each asset class varies with the horizon of the investor.

Based on the conditional covariance matrix and unconditional means described earlier, the mean-variance optimal portfolio can be solved for over any horizon, in a manner analogous to the traditional mean-variance optimization. Details on the optimization problem and solution are given in the Appendix; further technical details can be found in Campbell and Viceira (2005b).

Data and Initial Results

The study is based on quarterly data from Q2 1984 to Q1 2007. The length of the data series is constrained by the availability of data on direct real estate returns. Portfolios are chosen from amongst a cash equivalent benchmark and either three or four other asset classes, depending on whether REITs are included as an asset class. The logged real return to cash, r_{0t}, is calculated using the real effective yield per quarter on 90-day Treasury Bills, with the Consumer Price Index representing inflation. Equities are represented by the total return to the S&P 500 and bonds by the total return to 5-year Treasuries. Total returns to the variable liquidity version of the transaction based index (TBI) described in Fisher, Geltner and Pollakowski (2007) are obtained from the MIT Centre for Real Estate Web site and used to represent direct real estate. Returns to REITs, which are initially used as a state variable forecasting direct real estate returns and later as an asset class themselves, are from the Financial Times Stock Exchange (FTSE)/National Association of Real Estate Investment Trusts (NAREIT) Equity REIT Index.

Table 1 ■ Quarterly return statistics, Q2 1984–Q1 2007.

	Simple Returns		Log Real Excess Returns	
	Mean	Standard Deviation	Mean	Standard Deviation
T-Bills	0.0122	0.0054	n/a	n/a
Equities	0.0336	0.0783	0.0179	0.0777
Bonds	0.0189	0.0283	0.0062	0.0272
Real Estate	0.0235	0.0371	0.0105	0.0373
REITs	0.0352	0.0685	0.0203	0.0669

Adding real estate as an asset class to those (equities, bonds, cash) studied by Campbell and Viceira (2005a) implies that we allow lagged real estate returns to potentially affect returns to other asset classes and lagged returns on other classes to affect real estate returns. While we do not explore potential reasons for connections between real estate and other classes, possible explanations include wealth effects (*e.g.*, large positive (negative) returns in one asset class may result in investors rebalancing their portfolios and investing more (less) in another asset class) and exogenous economic factors that affect different asset classes at different lags (*e.g.*, an economic expansion may be reflected in stock or bond returns initially and then in real estate with a lag, or conversely an upswing in economic activity could result first in an increased demand for commercial real estate and a subsequent increase in equity markets as investors become aware of the increase in economic activity). While these explanations are speculative on our part, they do not affect our results as our purpose is to incorporate any potential linkages between real estate and other asset classes in forming optimal portfolios; we leave formal explanations of any linkages to other researchers.

Table 1 provides descriptive statistics for the asset class returns (in both simple and log real excess (*i.e.*, x_{it}) forms) over our sample period. The period was a good one for equities, which had an average quarterly simple return of 3.36% and an average log real excess return of 1.79%. As has historically been the norm, real estate had an average return (2.35% per quarter in simple form, 1.05% average excess return) and a standard deviation between those of equities and bonds. REITs performed well relative to other equities during our sample period, with the highest average return of all asset classes at 3.52% and volatility below that of stocks.

A set of state variables with predictive power for asset class returns is required. Following Campbell and Viceira's (2005a) study of cash/equity/bond portfolios, we use nominal T-Bill yield, equity dividend yield and term spread as state variables. We then augment this with state variables related to the predictability of real estate returns. The nominal T-Bill yield is defined as the log of one plus the yield on 90-day T-Bills. Because the VAR includes both the real T-Bill yield (representing return to cash) and nominal yield, the nominal will essentially capture inflation dynamics. We use the nominal yield rather than inflation directly to be consistent with Campbell and Viceira (2005a). Dividend yield is the log of one plus the trailing annual dividend yield on the S&P 500. Term spread is defined as $\ln(1 + R_{5yr}) - \ln(1 + R_{90})$, where R_{5yr} is the yield (expressed as effective per quarter) on a 5-year zero coupon T-Bond. Data for 5-year zero Treasury yields is initially from the data set of Campbell, Chan and Viceira (2003) (available on John Campbell's Web site as Campbell, Chan and Viceira (2007)). The Campbell *et al*. data is only available up to the end of 1999, so we augment the series with 5-year zero yields up to Q1 2007 from DataStream. Due to calculational differences, the two sources for 5-year zero yields do not match exactly in the quarters which the series overlap (Q1 1997 to Q4 1999). For consistency through time, the DataStream numbers are transformed based on regression parameters relating the two series in the overlap period.

The time period in this study is much shorter than that used by Campbell and Viceira (2005a), who are unconstrained by the availability of real estate data and use a 50-year time series of capital market returns to estimate their VAR. To ensure that our time period is "typical" and not substantively different than the longer series, we replicate the Campbell and Viceira results based on cash, equities, bonds and their three state variables. The six-equation VAR is estimated, and the covariance matrix in (3) is calculated. From the diagonal terms of (3) the annualized standard deviations at various investment horizons from one quarter to 25 years are calculated for each asset class.[3] These are shown in Figure 1.

The annualized standard deviations for cash, bonds and equities are somewhat smaller than those in Campbell and Viceira (2005a) at longer investment horizons. However, the important trends are the same. Equities and bonds both exhibit significant declines in riskiness at longer investment horizons, with the effect being very strong for equities. The standard deviation of excess equity returns declines from 9.1% at a 1-year horizon to 1.6% at a 25-year horizon. Thus, equities are less risky for long-term investors than for short-term, due to

[3] Annualization of the standard deviations is required to enable comparison across different horizons. Annualization of the raw numbers in the covariance matrix, Equation (3), is done by first dividing the variance by k, the number of quarters in the horizon, and then multiplying by four to create an annual variance. The square root of this is the annualized standard deviation.

Figure 1 ■ Annualized standard deviations of cash, equities and bonds.

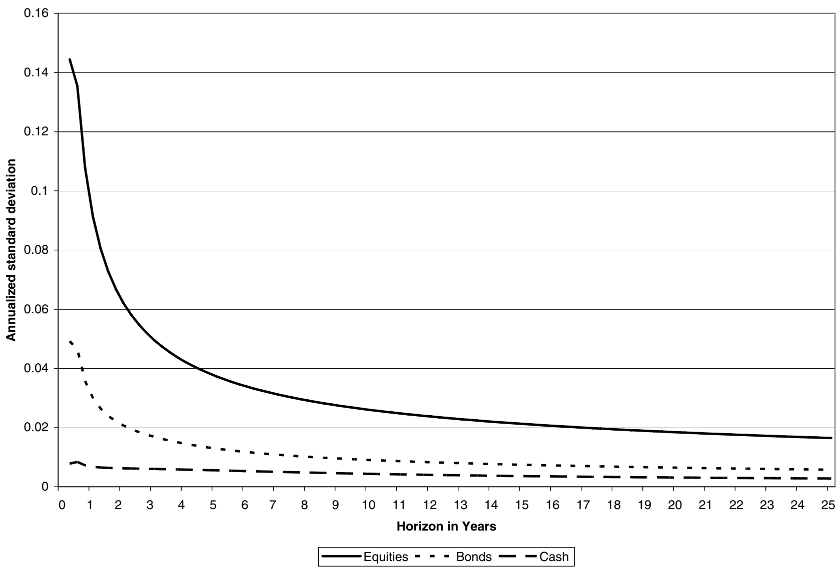

The figure shows annualized standard deviations at investment horizons from one quarter to 25 years. Standard deviations are calculated based on the estimated parameters of a six-variable vector autoregression on real returns to cash, excess returns to equities, excess returns to bonds, the nominal T-Bill yield, the term spread and the equity dividend yield. The vector autoregression is estimated based on quarterly data from Q2 1984 to Q1 2007.

equities' mean reversion qualities. Bonds also exhibit mean reverting tendencies, resulting in bonds over a long horizon having a risk level close to that of cash.

The term structure of risk for equities and bonds from our sample is similar to that from Campbell and Viceira's (2005a) longer time series. We are therefore confident that our sample time period is not an atypical one for capital markets.

Given our confidence that we are able to capture previously identified predictabilities in capital markets, we now turn to predictability of direct real estate returns. While we rely on the state variables identified by Campbell and Viceira (2005a) to predict capital market asset classes, there exist numerous possibilities for state variables to predict real estate returns. However, parsimony is an important consideration in estimating our VAR. Even with three asset classes plus a cash benchmark and three state variables (thus far) there are 56 VAR parameters and 49 covariance terms to estimate. Given our limited time series, care must be taken to preserve degrees of freedom in adding additional state variables to the system. Given this, and the central importance of real

Figure 2 ■ Partial autocorrelations of real estate returns.

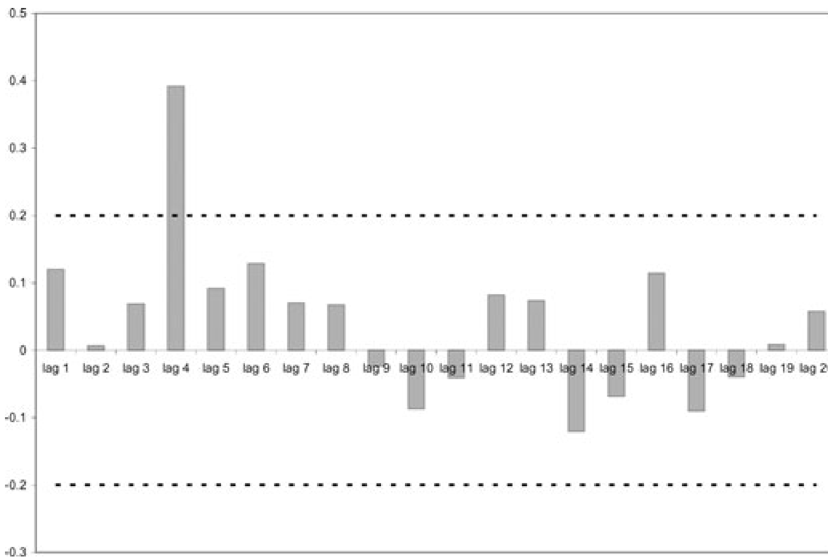

Vertical bars in the figure are the partial autocorrelations of real estate returns at lags from 1 to 20 quarters. Estimation is based on quarterly total returns to real estate from Q2 1984 to Q1 2007. Dashed lines indicate plus and minus two standard errors from zero.

estate to the study, we conduct a preliminary investigation of the predictability of real estate returns before deciding on the final form of the VAR.

Predictability of Real Estate Returns

Given that the informational efficiency of the real estate market is known to be substantially different from that of other asset classes, we begin by examining the time series properties of real estate returns. Figure 2 shows the partial autocorrelations of the TBI of total returns. Quarterly real estate returns show a strong (and significant) periodicity at an annual level. Given that the TBI is transaction rather than appraisal based, this is most likely an artifact of the actual real estate market rather than one of index construction. Because of the findings of Figure 2, in our VAR we use the fourth lag of real estate returns as a right-hand-side variable rather than the first lag as is done for the other asset classes. In our VAR estimation we lose four data points due to the use of fourth lag real estate returns, and therefore our estimation period is Q2 1985 to Q1 2007.[4]

[4] It is possible that the autocorrelation at a four-quarter lag is not due to momentum in real estate returns at an annual periodicity, but rather simply a time-of-the-year effect

Numerous state variables that might help predict real estate returns are possible. In the interest of parsimony we consider only three, each at lags from one to four quarters. The first potential state variable is employment growth. This is defined as the natural log of one plus the percentage change in full-time civilian employment (from the Bureau of Labor Statistics). This is chosen because of its sensitivity to broad macroeconomic factors, as well as the fact that employment directly relates to fundamental cashflows in most real estate property types. We also consider the lagged returns to REITs as a possible state variable predicting real estate, measured by the logged excess returns to the FTSE/NAREIT Equity REIT index. Barkham and Geltner (1995) conclude that public markets lead private real estate market because of greater informational efficiency. Finally, Liu and Mei (1992) show that cap rates have some predictive power for equity REITs, so we also explore the use of cap rates to predict private market returns.

To be consistent with the trailing dividend yield used for equities, we develop an analogous trailing cap rate. The difference between the percentage total return and the percentage capital appreciation return is defined as the real estate income return and represents the income earned through the quarter as a percentage of the beginning of period value. We multiply the beginning-of-period price index level by the income return for that period to get income in a pseudo-dollar form, and we then divide this dollar income by the end-of-period price index level to get an income return expressed as a percentage of the end-of-period price level. This then represents a form of cap rate on a historical basis. To annualize, for each quarter we sum that quarter's figure with the previous three quarters to get an annual trailing cap rate. The natural log of one plus this number is used in the analysis.

Table 2 shows correlations between real estate returns and lagged values of our three potential predictive state variables. The only statistically significant correlation is that between real estate returns and REIT returns lagged two quarters. Because of possible interactions among the explanatory variables, we also estimated a distributed lag model with real estate returns as the dependant variable and four lags of each of the explanatory variables as right-hand-side variables. The results are presented in Table 3. The only significant individual coefficient is that on REITs at a two-quarter lag. Further, F tests indicate that we cannot reject the null that all coefficients on employment growth and on trailing cap rates are equal to zero.

(*e.g.*, certain quarters tending to have high or low returns). To test for this we regressed (with no constant) real estate returns on dummy variables for Q2, Q3 and Q4 and took the residuals as "de-seasoned" real estate returns. The de-seasoned returns also exhibit a significant partial autocorrelation at a four-quarter lag (again, the first three lags are insignificant). Hence, the autocorrelation observed in the original series is not a time-of-the-year effect.

Table 2 ■ Correlations of real estate returns with lagged explanatory variables.

Lag on Variable	Correlation of Real Estate Returns with Lagged Values of: REIT Returns	Employment	Trailing Cap Rate
1	0.080	0.053	0.026
2	0.315***	0.086	0.044
3	0.075	0.028	0.091
4	0.143	−0.062	0.134

The table shows correlations between total returns to real estate and potential explanatory variables at lags from 1 to 4 quarters. The potential explanatory variables are REIT returns, employment and the trailing cap rate on real estate. Correlations are based on quarterly data from Q2 1984 to Q1 2007 (initial data points are lost when lagging the variables). * Significant at a 10% level, ** at a 5% level and *** at a 1% level.

Hence it seems that the best performing predictor for real estate returns is the lagged return to REITs, with cap rates and employment growth showing little predictive power. In our VAR we therefore include REITs as an additional variable in the system, with REIT returns lagged two quarters as the appropriate right-hand-side variable. The semi-annual lag between REIT returns and direct real estate returns is somewhat shorter than that discussed by Barkham and Geltner (1995), who report that the public market leads the private market by 1 to 2 years. The shorter lead time in our data may be due to the more recent time period of our study, or to the fact that we utilize a transaction-based index rather than an "unsmoothed" appraisal-based index as do Barkham and Geltner (1995).

Given our findings on the predictability of real estate returns, our final VAR model includes eight variables: real cash returns, and excess returns to equities, bonds and real estate as our asset classes; nominal T-Bill yield, equity dividend yield, term spread and excess returns to REITs as our state variables. These variables enter at lag 1 on the right-hand side of each equation, except real estate and REIT returns which enter at lag 4 and lag 2, respectively.

Results

The results of estimating our eight-equation VAR are presented in Table 4.[5] As measured by the R^2, real estate returns ($R^2 = 32\%$) exhibit a higher degree

[5] We tested for stationarity by examining the eigenvalues of the VAR parameter matrix. All eigenvalues are between -1 and $+1$, indicating our VAR system is stationary.

Table 3 ■ Distributed lag regression of real estate returns.

	Coefficient	t statistic	F statistic
Constant	−0.0353	−1.15	
Employment, lag 1	0.6174	0.66	Test of all employ. coefficients equal to zero: $F = 0.441$
Employment, lag 2	0.6478	0.69	
Employment, lag 3	0.1281	0.14	
Employment, lag 4	−0.6581	−0.72	
REIT Returns, lag 1	0.0232	0.39	Test of all REIT coefficients equal to zero: $F = 2.46^{*}$
REIT Returns, lag 2	0.1595	2.72^{***}	
REIT Returns, lag 3	0.0362	0.62	
REIT Returns, lag 4	0.0739	1.20	
Cap Rate, lag 1	2.5319	0.95	Test of all cap rate coefficients equal to zero: $F = 0.238$
Cap Rate, lag 2	−7.2962	−1.60	
Cap Rate, lag 3	3.7578	0.84	
Cap Rate, lag 4	1.6790	0.62	

The table shows the results of a distributed lag regression with total returns to real estate as the dependent variable. Independent variables are employment, total returns to REITs and the trailing cap rate, all at lags from 1 to 4 quarters. Estimation is based on quarterly data from Q2 1985 to Q1 2007. * Significant at a 10% level, ** at a 5% level and *** at a 1% level.

of predictability in their returns than do equities or bonds. As expected given the preceding analysis, real estate returns four quarters prior and REIT returns two quarters prior have significant predictive capability for real estate. Table 5 presents the cross-equation correlations of the VAR residuals. These indicate the contemporaneous effects of unexpectedly high or low values of each variable on the others. Taken together, Tables 4 and 5 will help explain the results presented later on the term structure of risk and optimal portfolio allocations.

Based on the VAR results the variance of each asset class's excess returns can be calculated as in Equation (3). Figure 3 presents the annualized standard deviations for cash, equities, bonds and real estate that result. First note that equities and bonds both exhibit mean reversion as the annualized risk decreases at longer investment horizons. This is especially evident for equities that see their annual standard deviation decrease from 14.6% at a quarterly horizon, to 9.5% at an annual horizon and 1.7% at a 25-year horizon. Equities, then, are far less risky for long-term investors, consistent with the findings of Campbell and Viceira (2005a). As in the Campbell and Viceira work, the driving factor in equity's mean reversion is its relationship to the dividend yield. Table 5 shows

Table 4 ■ Vector autoregression results.

Equation	Constant	Cash $(t-1)$	Equities $(t-1)$	Bonds $(t-1)$	Real Estate $(t-4)$	REITs $(t-2)$	T-Bill Yield $(t-1)$	Term Spread $(t-1)$	Div. Yield $(t-1)$	R^2
(1) Cash	−0.0031	0.0511	0.0040	0.0411	−0.0044	0.0068	0.8014	0.1665	−0.1369	0.44
	(−1.70)*	(0.45)	(0.67)	(2.17)**	(−0.33)	(0.99)	(4.45)***	(0.59)	(−1.63)	
(2) Equities	0.0216	1.9016	−0.0757	0.2123	0.1098	−0.0432	−7.1397	−12.2413	4.5861	0.14
	(0.65)	(0.89)	(−0.68)	(0.60)	(0.45)	(−0.34)	(−2.13)**	(−2.34)**	(2.94)***	
(3) Bonds	−0.0242	−0.0432	−0.0312	0.0975	0.0555	0.0481	2.5040	5.2119	−0.7083	0.15
	(−2.17)**	(−0.06)	(−0.84)	(0.83)	(0.68)	(1.14)	(2.24)**	(2.99)***	(−1.36)	
(4) Real Estate	0.0296	0.0933	−0.0515	0.1235	0.3573	0.1389	−1.4172	−1.3481	−0.1592	0.32
	(2.09)**	(0.10)	(−1.10)	(0.83)	(3.46)***	(2.59)**	(−1.00)	(−0.61)	(−0.24)	
(5) REITs	0.0386	0.1351	0.0060	0.0581	0.1359	0.0127	−4.3322	−1.0562	1.3576	0.06
	(1.28)	(0.07)	(0.06)	(0.18)	(0.62)	(0.11)	(−1.44)	(−0.22)	(0.97)	
(6) T-Bill Yield	0.0005	−0.0292	0.0006	−0.0127	0.0009	−0.0012	0.9264	−0.0559	0.0258	0.94
	(0.95)	(−0.82)	(0.31)	(−2.16)**	(0.21)	(−0.56)	(16.55)***	(0.52)	(0.99)	
(7) Term Spread	0.0010	0.0292	0.0018	0.0085	−0.0043	−0.0016	−0.0882	0.7953	0.0179	0.74
	(1.92)*	(0.86)	(1.02)	(1.51)	(−1.10)	(−0.77)	(−1.65)	(9.54)***	(0.72)	
(8) Div. Yield	0.0001	−0.0330	0.0017	−0.0101	0.0010	0.0014	0.0920	0.1773	0.9208	0.96
	(0.15)	(−0.62)	(0.61)	(−1.15)	(0.16)	(0.45)	(1.10)	(1.36)	(23.61)***	

The table contains the results of the estimation of an eight equation vector autoregression (VAR). The dependent variables in the VAR are: returns to cash as proxied by the logged real returns to T-Bills; the returns to equities, bonds, real estate and REITs, all in logged excess return form; the logged nominal T-Bill yield; term spread defined as the difference, in logged form, between the yield of 5 year zero coupon Treasuries and 90 day T-Bills; and the trailing annual dividend yield on the S&P 500 (again in logged form). Independent variables in the VAR are lags of the dependant variables, with all variables except real estate and REITs lagged one period. Real estate returns are lagged four periods, and REIT returns lagged two periods in the VAR. The VAR is estimated based on quarterly data from Q2 1985 to Q1 2007. t-statistics are in parentheses. * Significant at a 10% level, ** at a 5% level and *** at a 1% level.

Table 5 ■ Cross-correlations of residuals from VAR.

	Cash	Equities	Bonds	Real Estate	REITs	T-Bill Yield	Term Spread	Div. Yield
Cash	1							
Equities	0.281	1						
Bonds	−0.128	−0.133	1					
Real Estate	−0.028	0.217	−0.002	1				
REITs	0.092	0.510	0.177	0.152	1			
T-Bill Yield	0.264	0.204	−0.670	0.201	−0.099	1		
Term Spread	−0.179	−0.064	−0.517	−0.263	−0.141	−0.264	1	
Div. Yield	−0.206	−0.886	−0.028	−0.168	−0.565	−0.086	0.141	1

The table contains the correlations of residuals across the eight equations of a vector autoregression (VAR). The dependent variables in the VAR are: returns to cash as proxied by the logged real returns to T-Bills; the returns to equities, bonds, real estate and REITs, all in logged excess return form; the logged nominal T-Bill yield; term spread defined as the difference, in logged form, between the yield of 5 year zero coupon Treasuries and 90 day T-Bills; and the trailing annual dividend yield on the S&P 500 (again in logged form). Independent variables in the VAR are lags of the dependant variables, with all variables except real estate and REITs lagged one period. Real estate returns are lagged four periods, and REIT returns lagged two periods in the VAR. The VAR is estimated based on quarterly data from Q2 1985 to Q1 2007.

that equity returns and the dividend yield are strongly negatively correlated. When stocks increase, the dividend yield decreases. However, the VAR results in Table 4 show that the lagged dividend yield is positively related to equity returns; stock returns tend to be higher when last quarter's dividend yield was low. Together the results indicate a mean reverting process for equities that lowers risk for longer-term investors: high returns are associated with a lower contemporaneous dividend yield, but this lower yield predicts lower returns. This is consistent with Campbell and Viceira (2005a). Bonds also exhibit mean reversion, but to a lesser extent than equities. At long horizons, bonds have risk quite close to that of cash (cash and bonds have standard deviations of 0.3% per year and 0.6% per year, respectively, at a 25-year horizon), although cash remains the least risky asset class.

Of special interest for our purposes, of course, is the term structure of risk for real estate. The most obvious result from Figure 3 is that real estate also exhibits mean reversion. Real estate is less risky for longer-horizon investors. At short horizons, real estate risk is between that of stocks and bonds, an unsurprising result as this is often taken as a given in descriptions of real estate investment. While the risk of real estate decreases at longer horizons, the degree of mean reversion is much less than that of equities. In fact, at longer horizons the risks of real estate and equities are virtually identical. At a 10-year (40 quarter)

Figure 3 ■ Annualized standard deviations of cash, equities, bonds and real estate.

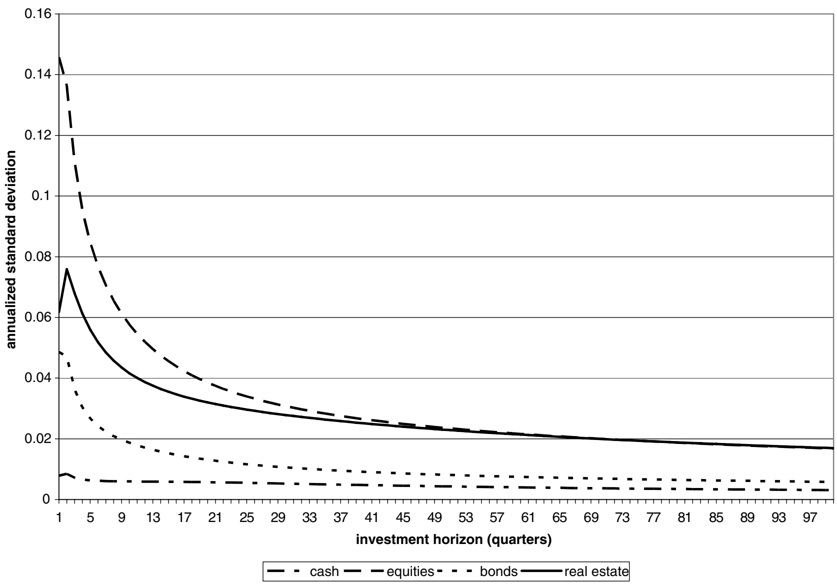

The figure shows annualized standard deviations at investment horizons from one quarter to 25 years. Standard deviations are calculated based on the estimated parameters of an eight variable vector autoregression on real returns to cash, excess returns to equities, excess returns to bonds, excess returns to real estate, excess returns to REITs, the nominal T-Bill yield, the term spread and the equity dividend yield. Independent variables in the VAR are lags of the dependant variables, with all variables except real estate and REITs lagged one period. Real estate returns are lagged four periods, and REIT returns lagged two periods in the VAR. The VAR is estimated based on quarterly data from Q2 1985 to Q1 2007.

investment horizon, the standard deviation of real estate is 2.5% per year and of equities is 2.6% per year. At longer horizons the risks become even closer.

Two major conclusions come out of Figure 3. First, real estate exhibits mean reversion and is less risky for longer-term investors. Second, and perhaps most surprising, for long-term investors (say horizons of 10 years or more) real estate is just as risky as equity investment. This belies the typical argument for real estate as a low-risk investment relative to stocks, at least for long-term investors.

To explain the mean reversion properties of real estate, one can look to the VAR results in Tables 4 and 5. Real estate returns have a certain amount of natural momentum as evidenced by the positive autocorrelation at a four-quarter lag. This tends to be reinforced by real estate's relationship to REIT returns. From Table 5, real estate and REIT returns are contemporaneously positively correlated; unexpectedly high REIT returns tend to coincide with high returns to direct real estate. But Table 4 shows that high REIT returns also predict higher real estate returns two quarters later. Hence high (low) levels of REIT returns are associated with higher (lower) real estate returns today and also in the future, reinforcing the momentum (or mean aversion) tendencies of real estate.

However, momentum in returns should result in increasing risk over longer time horizons, the opposite of what is observed for real estate. The observed long-term mean reversion is explained by the relationship between real estate returns and the nominal T-Bill yield, which induces mean reversion in real estate. Nominal T-Bill yields and real estate are contemporaneously positively correlated (Table 5). Given that, as noted previously, nominal T-Bill yield in the VAR proxies for inflation dynamics, this is consistent with the notion of real estate as an inflation hedge. However, Table 4 indicates that nominal T-Bill yields are negatively associated with real estate returns one period later. Hence, whereas high inflation is associated with high real estate returns this quarter, it predicts lower real estate returns next quarter. The result is mean reversion in real estate returns.[6]

The momentum of real estate returns (due to autocorrelation and the REIT relationship) explains the initial spike up in real estate risk observed in Figure 3, wherein real estate risk peaks at a two-quarter investment horizon, but this effect is counteracted by the T-Bill relationship. The T-Bill yield is a more persistent phenomenon (note the high autocorrelation and predictability of T-Bill yields from regression 6 of Table 4), and this effect outweighs the momentum aspects in the long run and results in the observed downward sloping term structure of risk for real estate.

Of key importance, however, is that the existence of some underlying momentum means that real estate's mean reversion is weaker than that of equities, the result being that equity risk decreases much faster than real estate's and the two asset classes end up with equal levels of risk over long horizons.

While the contemporaneous positive correlation between real estate returns and nominal T-Bill yields (*e.g.*, inflation) seems intuitive, the negative relationship between the nominal yield and future real estate returns seems at odds with

[6] The conditional standard deviations of each asset class as presented in Figure 3 are calculated using the coefficient estimates from Table 4 as given. Table 4 shows that the coefficient on nominal T-Bill yield in the real estate excess return equation is negative, but insignificant at conventional levels. However, Kandel and Stambaugh (1996) show that effects judged insignificant by traditional statistical tests may in fact have economically significant effects on asset allocation decisions.

the traditional view of real estate as an inflation hedge. To explore this more fully, we separate real estate return into two components: income return and capital appreciation return. Regressions of the same form as the VAR equations are estimated with those two types of returns (in logged excess return form) as the dependant variables. The results are presented in Table 6. Panel A of Table 6 presents the estimated regression coefficients, as well as the coefficient from the regression on T-Bill yields (which are repeated from Table 4 for comparison purposes). Panel B contains the correlations of residuals from the three regressions.

As might be expected, the momentum aspects of real estate returns (via the fourth lag autocorrelation and REIT relationship) are strongly evident in appreciation returns, but not in the income return regression. Momentum is a property value phenomenon. Panels A and B together indicate that the mean reversion observed for total returns is driven by the capital appreciation component, and it appears that transaction prices for properties "overshoot" inflation. Panel B shows that high nominal T-Bill yields (*e.g.*, high inflation) are associated with higher appreciation returns that period (unlike the correlation between T-Bill and income return residuals, which is negative). But Panel A shows that high nominal T-Bill yields are associated with a lower appreciation return the following period. This tendency for property prices to rise with higher inflation, but then fall back, explains the mean reverting tendency of real estate returns.

Combining our estimate of the covariance from Equation (3) and the unconditional mean per period of each asset class, we can solve for the optimal portfolio allocation (Equation (A3) in the Appendix shows the solution). The optimal allocation is contingent on the target return chosen. We choose four representative target returns (in simple return per quarter): 1.2%, 1.9%, 2.8% and 3.4%. These are chosen to represent a range of risk tolerances and correspond to the mean returns to T-Bills, bonds, a 60/40 stock/bond portfolio and equities, respectively. The results for these target returns, at horizons of 1, 5, 10 and 25 years, are shown in Table 7.

The optimal portfolios for the three highest target returns include large short positions in cash equivalents at all investment horizons. Highly leveraged positions are optimal for these target returns. The effects of mean reversion in asset class returns is evident when looking at the standard deviations of the optimal allocations. For all target returns the standard deviations decrease with investment horizon. Furthermore, the size of the decrease is very large and economically very significant. For the most aggressive target return (3.4%) the annual standard deviation is 57% for an investor with a 1-year horizon but decreases to only 14% for a long-term investor with a 25-year horizon.

Table 6 ■ Regression results for real estate income and appreciation returns separately.

Panel A: Parameter Estimates

	Constant	Cash ($t-1$)	Equities ($t-1$)	Bonds ($t-1$)	Real Estate ($t-4$)	REITs ($t-2$)	T-Bill Yield ($t-1$)	Term Spread ($t-1$)	Div. Yield ($t-1$)	R^2
(1) Real Estate Income	0.0131	0.1158	0.0003	0.0326	−0.0051	−0.0015	−0.6375	0.9736	−0.2817	0.79
	(10.69)***	(1.48)	(0.06)	(2.52)**	(−0.57)	(−0.31)	(−5.18)***	(5.07)***	(−4.91)***	
(2) Real Estate	0.0161	0.0059	−0.0530	0.1053	0.3660	0.1435	−1.7050	−2.2767	0.0934	0.32
Appreciation	(1.13)	(0.01)	(−1.12)	(0.70)	(3.52)***	(2.66)***	(−1.19)	(−1.02)	(0.14)	
(3) T-Bill Yield	0.0005	−0.0292	0.0006	−0.0127	0.0009	−0.0012	0.9264	−0.0559	0.0258	0.94
	(0.95)	(−0.82)	(0.31)	(−2.16)**	(0.21)	(−0.56)	(16.55)***	(0.52)	(0.99)	

Panel B: Correlations between Residuals

	Real Estate Income	Real Estate Appreciation	T-Bill Yield
Real Estate Income	1		
Real Estate Appreciation	−0.060	1	
T-Bill Yield	−0.620	0.217	1

The table shows the results of two regressions in which the dependent variables are the logged excess income return to direct real estate and the logged excess appreciation return to direct real estate. The independent variables are: returns to cash as proxied by the logged real returns to T-Bills; the total returns to equities, bonds, real estate and REITs, all in logged excess return form; the logged nominal T-Bill yield; term spread defined as the difference, in logged form, between the yield of 5-year zero coupon Treasuries and 90-day T-Bills; and the trailing annual dividend yield on the S&P 500 (again in logged form). All independent variables are lagged one period except for the total return to REITs and real estate which are lagged two periods and four periods, respectively. Results for the regression labeled T-Bill yield are the same as for that equation in the VAR presented in Table 4 and are presented again here for comparison purposes. Estimation is based on quarterly data from Q2 1985 to Q1 2007. t-statistics are in parentheses. * Significant at a 10% level, ** at a 5% level and *** at a 1% level.

Table 7 ■ Optimal allocations.

		Optimal Allocations					Relative Allocations to Non Cash Classes		
	Investment Horizon	Cash	Equities	Bonds	Real Estate	Std. Dev.	Equities	Bonds	Real Estate
Target = 3.4%	1 year	−2.07	0.54	1.92	0.61	0.57	0.18	0.63	0.20
	5 years	−2.40	0.75	2.03	0.62	0.28	0.22	0.60	0.18
	10 years	−2.46	0.75	1.91	0.80	0.21	0.22	0.55	0.23
	25 years	−2.34	0.62	1.68	1.04	0.14	0.19	0.50	0.31
Target = 2.8%	1 year	−1.46	0.42	1.54	0.49	0.51	0.17	0.63	0.20
	5 years	−1.74	0.59	1.64	0.50	0.25	0.22	0.60	0.18
	10 years	−1.79	0.59	1.55	0.65	0.19	0.21	0.56	0.23
	25 years	−1.70	0.49	1.38	0.84	0.13	0.18	0.51	0.31
Target = 1.9%	1 year	−0.54	0.25	0.98	0.31	0.40	0.16	0.64	0.20
	5 years	−0.74	0.35	1.07	0.32	0.20	0.20	0.61	0.19
	10 years	−0.79	0.35	1.02	0.41	0.15	0.20	0.57	0.23
	25 years	−0.75	0.29	0.92	0.53	0.10	0.17	0.53	0.30
Target = 1.2%	1 year	0.17	0.12	0.54	0.17	0.29	0.14	0.65	0.20
	5 years	0.03	0.17	0.62	0.19	0.14	0.17	0.64	0.19
	10 years	−0.01	0.17	0.61	0.23	0.11	0.17	0.60	0.23
	25 years	0.00	0.14	0.57	0.29	0.08	0.14	0.56	0.29

The table shows the optimal allocations to cash, bonds, equities and real estate for various target returns and at various investment horizons. For each optimal portfolio, the relative allocations to the noncash asset classes are also shown. For instance, the relative allocation to equities is calculated as optimal allocation to equities divided by the sum of the equity, bond and real estate optimal allocations. Allocations may not sum exactly to one due to rounding.

Real estate has a large positive allocation for all target returns and all investment horizons. Hence, including the predictability of asset returns does not negate the benefits of including real estate in a mixed-asset portfolio. In fact, the optimal allocations to real estate are quite large with the lowest being 17% for a very conservative target return of 1.2% at a short 1 year investment horizon, and ranging up to over 100% in the case of an aggressive portfolio at a long investment horizon.

The benefits to having real estate within a portfolio can be measured by comparing, for a given target return, the annualized standard deviation reported in Table 7 to that which would result if real estate were not included in the portfolio. Using the same VAR parameters, we calculated the optimal portfolios and resulting standard deviations when only cash, equities and bonds were included as asset classes.[7] As an example, for a 25-year horizon and a 2.8% target return, Table 7 reports an annualized standard deviation of 0.13; the optimal portfolio using only cash, stocks and bonds has a standard deviation of 0.18. Hence, including real estate as an asset class results in a 28% decrease in risk for the same expected return. While we do not report the results in a table (the results are available on request), the diversification benefits of real estate measured in this way increase with investment horizon and range from a 5% reduction in risk for a 5-year horizon and 1.9% target return up to a 30% reduction in risk for a 25-year horizon and the same target return. The benefits of real estate as an asset class are apparent, especially for longer-term investors.

The last three columns of Table 7 show the optimal portfolios in terms of relative allocation to the noncash asset classes. For example, dividing the optimal allocation to equities by the total allocation to equities, bonds and real estate gives the optimal relative allocation to equities. An examination of the optimal noncash allocations reveals only slight differences across the target returns. At each investment horizon, the optimal noncash allocations for different targets are very close or equal. This implies that the only significant differences across the target returns is the degree of leverage used in the portfolio. This is consistent with the traditional textbook approach to modern portfolio theory in which investors of all risk tolerances choose the same portfolio of risky assets and use leverage to achieve the desired risk level.

Given that investors' risk tolerance only affects the optimal degree of leverage to use, we can examine the optimal noncash allocations for one target return without loss of generality. Figure 4 shows the optimal allocations to equities, bonds and real estate over all horizons from one quarter up to 25 years for a target return of 3.4%. Optimal allocations to equities are generally steady over varying horizons, typically between 18% and 22%. Optimal allocations to bonds decline for longer investment horizons. The optimal bond allocation peaks at 64% of the noncash allocation for a 1.5-year investment horizon. It then declines for longer horizons, until only 50% of the optimal portfolio is invested in bonds at a 25-year horizon.

The optimal allocation to real estate declines initially at very short horizons. The maximum real estate allocation, 38%, is at the shortest horizon of one quarter. This then declines until hitting 17.5% at a horizon of 3 years. However, given

[7] Even though we do not include real estate as an asset class, we use the same VAR parameters so that we do not confuse the diversification benefits of real estate with differences in estimated VAR parameters. Intuitively, we are comparing an investor who includes real estate in their portfolio (*i.e.*, Table 7) to one who knows of the relationships between real estate and other asset classes but does not, or cannot, allocate any funds to it.

Figure 4 ■ Optimal allocations to noncash asset classes for a target return of 3.4%.

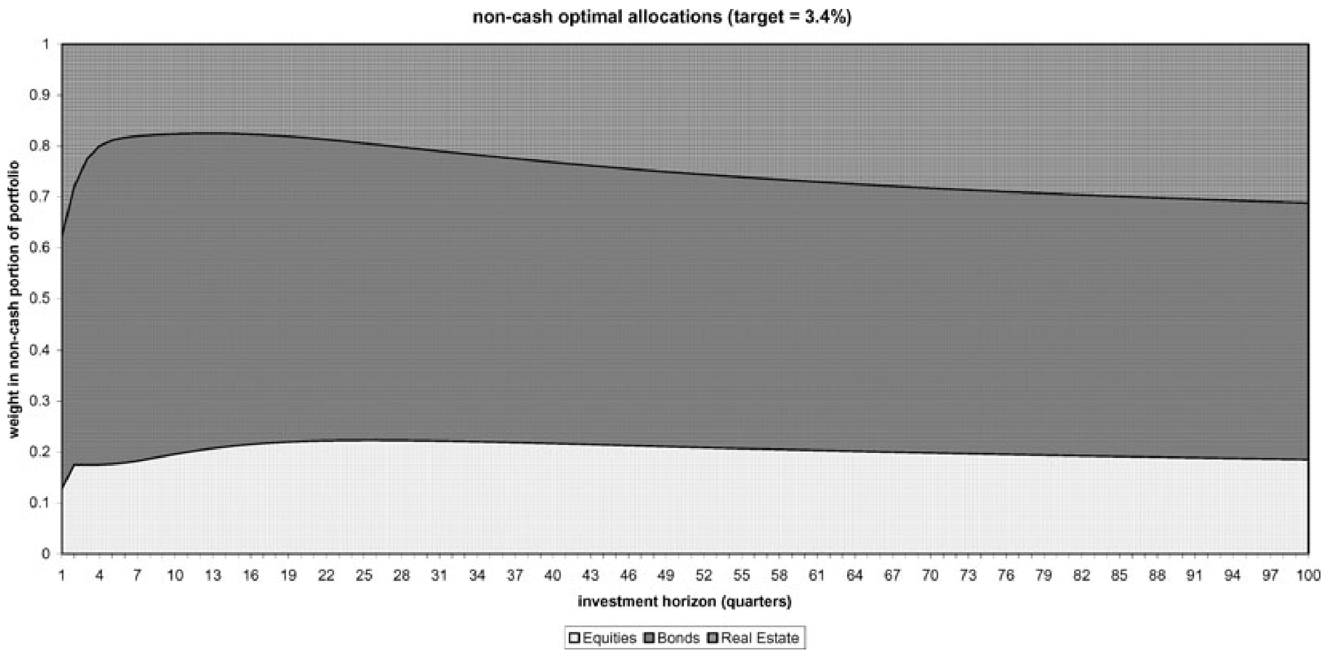

The figure shows the optimal allocations to noncash asset classes over varying horizons for a target return of 3.4% per quarter. Inputs for the optimization are based on the estimated parameters of an eight variable vector autoregression on real returns to cash, excess returns to equities, excess returns to bonds, excess returns to real estate, excess returns to REITs, the nominal T-Bill yield, the term spread and the equity dividend yield. Independent variables in the VAR are lags of the dependant variables, with all variables except real estate and REITs lagged one period. Real estate returns are lagged four periods, and REIT returns lagged two periods in the VAR. The VAR is estimated based on quarterly data from Q2 1985 to Q1 2007.

that the illiquidity of real estate over short time intervals has not been included in this analysis, these results for very short horizons may be unrealistic. We return to this issue later in the article when we allow investment in REITs as an asset class.

The most obvious trend from Figure 4 is the general trend toward higher real estate allocations at longer investment horizons. From horizons of 3 years and on, the optimal allocation to real estate increases monotonically with horizon, hitting 31% of the noncash allocation at a 25-year horizon. The increase in optimal real estate allocation comes mostly at the expense of the bond allocation, which gradually decreases with horizon. The conclusion is that real estate should be a more important asset class for longer-term investors. Given that most institutional investors, such as pension funds, have very long horizons, this reinforces the common argument (see, *e.g.*, Chun, Sa-Aadu and Shilling 2004) that pension funds underinvest in real estate.

We must emphasize from Table 7 and Figure 4 the importance of proper consideration of investment horizon in portfolio choice. It is not uncommon to use short-term returns to estimate optimal portfolios even when the true investment horizon is long, under the assumption that optimal portfolios do not change with horizon (which is equivalent to the assumption that returns are i.i.d.). Our results illustrate the dangers inherent in such an approach. For instance, an investor with a long-term (25-year) horizon and a 3.4% target return has an optimal noncash allocation of 19%, 50% and 31% to equities, bonds and real estate, respectively. If that investor were to, mistakenly, use short-term returns (1-year) to estimate an optimal allocation he/she would allocate 18%, 63% and 20% to equities, bonds and real estate, respectively. This short-term optimal portfolio is actually far from the optimal for the long-term investor; specially, he/she would significantly overinvest in bonds and underinvest in real estate compared to the true optimal allocation.

Given that the difference between the risk of real estate and that of equities narrows dramatically with investment horizon, it seems strange that the optimal allocation to real estate increases with horizon.[8] An explanation of this can be found in Figure 5, which shows the conditional correlations between real estate and equities and between real estate and bonds at various horizons (from the off-diagonal elements of the covariance matrix in Equation (3)). Although correlations increase at short intervals, the trend for longer horizons is quite distinctly downwards. Hence, real estate becomes better as a diversification tool at longer investment horizons. The diversification benefits of real estate

[8] Recall that we use the unconditional mean returns in the optimizations, so that mean return per period remains constant at all horizons.

Figure 5 ■ Correlations between real estate and equities/bonds.

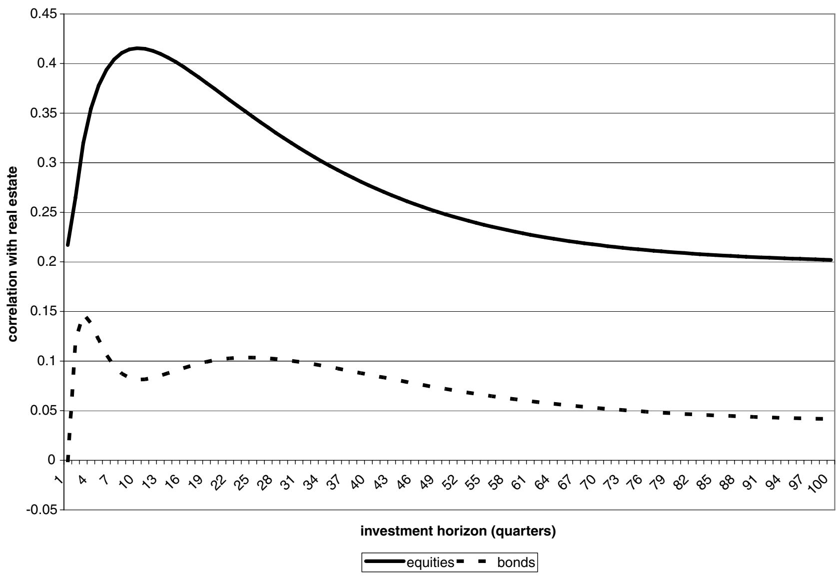

The figure shows the correlation between real estate and equity returns, and between real estate and bond returns (all returns are in logged excess return form). Correlations are shown at various investment horizons from one to 100 quarters. Correlations are estimated based on the estimated parameters of an eight-variable vector autoregression on real returns to cash, excess returns to equities, excess returns to bonds, excess returns to real estate, excess returns to REITs, the nominal T-Bill yield, the term spread and the equity dividend yield. Independent variables in the VAR are lags of the dependant variables, with all variables except real estate and REITs lagged one period. Real estate returns are lagged four periods, and REIT returns lagged two periods in the VAR. The VAR is estimated based on quarterly data from Q2 1985 to Q1 2007.

are therefore sensitive to the investment horizon, and this explains not only the decreasing allocation to real estate at very short horizons (less than 3 years), but also the general trend to higher real estate allocations at longer investment horizons.

Our analysis of optimal real estate allocations has not considered the effect of the illiquidity of direct real estate and transaction costs. Given that transaction costs for direct real estate are large, it is unlikely that an investor would invest in such an asset class on a short-term basis, and our results on optimal real estate allocations for short horizons should be considered with that caveat in mind. Collett, Lizieri and Ward (2003) find that the median realized holding period for institutional real estate investors varies with property type and year of purchase, but is generally between 7 and 14 years, consistent with transaction costs making real estate a suitable investment only for investors with medium- to long-term horizons.

Formal consideration of the effect of transaction costs is beyond the scope of this article; however, we do consider the role played by an asset class sometimes proposed as a more liquid alternative to direct real estate investments, namely REITs.

REITs as an Asset Class

The preceding analysis utilized REITs as a state variable predicting direct real estate returns. However, investors have access to REITs as an asset class themselves. An allocation to REITs could be either in addition to, or in lieu of, an allocation to direct property. Arguments for an allocation to REITs in lieu of private market real estate investment hinge on three issues: liquidity, divisibility and management costs and whether public market real estate returns serve as an adequate proxy for private market returns.

Investments in REITs are associated with higher liquidity and lower transaction costs than investments in direct real estate. As noted, the illiquidity of direct real estate likely makes it unsuitable for short-term investors. Given the liquidity of the public markets, it is possible that an allocation to REITs could fill the role of real estate over shorter horizons. For instance, the spanning tests of Chiang and Lee (2007) show that REITs are a redundant asset class when direct real estate is part of the portfolio, yet they conclude that "smaller investors with shorter time horizons" may still find REITs attractive.

A second issue limiting certain investors' ability to invest in direct real estate is the size of investment required. Direct property investment is typically indivisible and requires a large capital outlay. Hence, only portfolios of a significant size will be able to reasonably include an allocation to direct real estate (especially if the real estate portfolio is to be diversified across properties). Conversely, an investment in REITs can be made in any size and REITs may therefore be able to serve as a proxy for real estate investment for smaller investors. A related advantage of REITs for smaller investors is that the professional management of REITs allows investors to allocate to real estate without the need to develop an in-house real estate management team, a prospect not likely to be cost-effective for smaller portfolios.

While REITs provide an alternative to direct real estate for small investors, their effectiveness as a substitute depends on the degree to which REIT returns proxy for direct property returns. The relationship between REITs and real estate over the long term has been debated in the literature. Pagliari, Scherer and Monopoli (2005) report that means and variances of public and private market real estate are indistinguishable after controlling for property type and leverage. Conversely, Riddiough, Moriarty and Yeatman (2005) report a significant difference in public and private market performance.

We now examine optimal asset allocations when REITs are included as an asset class. Given the discussion earlier, we consider our analysis of REITs from the points of view of three classes of investors: (1) large, long-horizon investors examining whether *both* private and public market real estate should be included in a portfolio, (2) small, long-horizon investors examining whether an allocation to REITs should be included in a portfolio in lieu of direct property investment and (3) short-horizon investors considering an allocation to REITs. The definition of what constitutes a short- and long-horizon investor is necessarily somewhat arbitrary. Given the results of Collett, Lizieri and Ward (2003) that show average property holding periods are between 7 and 14 years, we take the midpoint and assume for discussion purposes that a horizon less than 10 years indicates a short-term investor for whom direct real estate investment is infeasible.

We repeat the previous analysis on optimal portfolio allocations, but this time we allow investment in REITs. The VAR system used in this instance is identical to that presented in Table 4. We add no new variables to the system; our only change is to treat REITs as an asset class rather than as a state variable.[9]

Figure 6 shows the correlation between excess returns to direct real estate and REITs at various horizons based on our VAR estimation. The correlation between real estate and REIT returns increases rapidly with horizon in the beginning but levels off, never going higher than 0.54. The correlation reaches the tangent level quite quickly; at a horizon of 10 years (40 quarters) the correlation is 0.53 and even at 5 years the correlation is already 0.5.

The rise in correlation supports the argument that REITs are more like real estate at longer horizons, over which stock-market-induced volatility will have less

[9] Given our desire to preserve degrees of freedom, we wish to be careful in adding additional variables. While premium to net asset value might be considered as an additional predictor, Gentry, Jones and Mayer (2004) report that it is not effective in predicting aggregate REIT returns through time (although it does have power cross-sectionally). We explored the time-series properties of REIT returns (not reported) but found no significant autocorrelations. We also found no predicative ability in the REIT dividend yield. Note that in Table 4 REITs have the lowest predictability, in terms of R^2 of all the variables; the addition of further state variables related to REITs seems to add little to this quite unpredictable asset class. The lack of predictability in REIT returns is consistent with Ling, Naranjo and Ryngaert (2000) who report very little out of sample predictability for equity REITs, especially in the 1990s. We therefore choose not to add additional state variables and rely on the same VAR system as was used previously.

Figure 6 ■ Correlation of direct real estate and REITs.

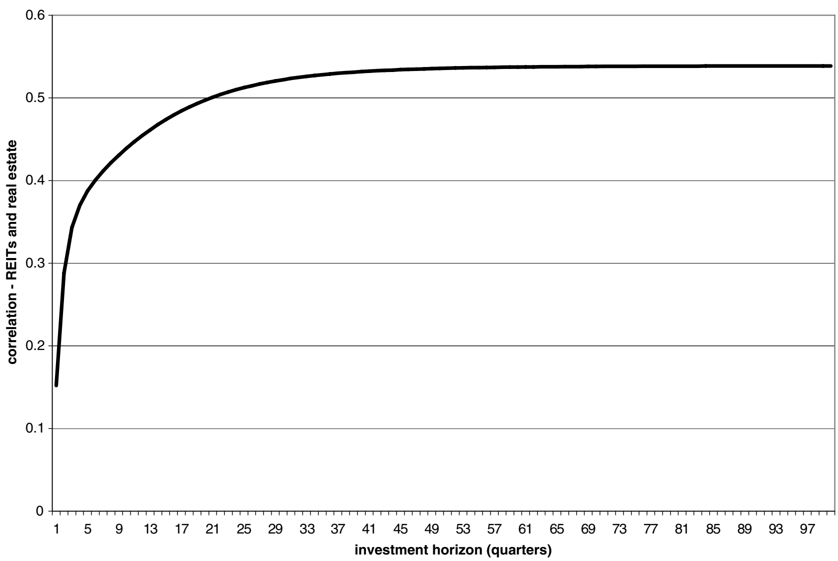

The figure shows the correlation between excess returns to direct real estate and REITs at investment horizons from one quarter to 25 years. Correlations are calculated based on the estimated parameters of an eight variable vector autoregression on real returns to cash, excess returns to equities, excess returns to bonds, excess returns to real estate, excess returns to REITs, the nominal T-Bill yield, the term spread and the equity dividend yield. Independent variables in the VAR are lags of the dependant variables, with all variables except real estate and REITs lagged one period. Real estate returns are lagged four periods, and REIT returns lagged two periods in the VAR. The VAR is estimated based on quarterly data from Q2 1985 to Q1 2007.

effect. While our results support this line of reasoning, we can only support it to a point given that the correlation tops out at 0.54. The relatively low correlation at even very long horizons is likely due in part to the fact that we do not adjust REIT returns, as do Pagliari, Scherer and Monopoli (2005) and Riddiough, Moriarty and Yeatman (2005), for leverage and property type composition. In our case we are interested in the investment characteristics of REITs as an asset class, that is, the role of a diversified REIT position within a larger portfolio. An investment in a diversified portfolio of REITs would be characterized by certain property type tendencies and leverage levels; adjusting for these characteristics would change the investment nature of REITs.

Figure 7 shows the term structure of risk for REITs, with direct real estate also displayed for comparison purposes. Contrary to the findings of Fugazza, Guidolin and Nicodanna (2007) for European real estate securities, we find that REITs, like the other asset classes, exhibit mean reversion. The result is

Figure 7 ■ Annualized standard deviations of real estate and REITs.

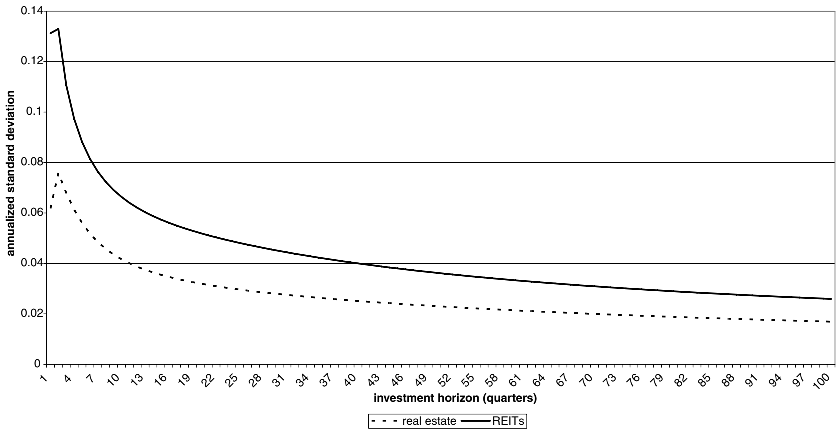

The figure shows annualized standard deviations for direct real estate and for REITs at investment horizons from one quarter to 25 years. Standard deviations are calculated based on the estimated parameters of an eight variable vector autoregression on real returns to cash, excess returns to equities, excess returns to bonds, excess returns to real estate, excess returns to REITs, the nominal T-Bill yield, the term spread and the equity dividend yield. Independent variables in the VAR are lags of the dependant variables, with all variables except real estate and REITs lagged one period. Real estate returns are lagged four periods, and REIT returns lagged two periods in the VAR. The VAR is estimated based on quarterly data from Q2 1985 to Q1 2007.

that risk declines with investment horizon. However, REITs remain riskier than direct real estate at all horizons.

Table 8 shows the result of the portfolio optimization at different horizons, when the investor can allocate to both public and private market real estate (as well as the other asset classes). As discussed earlier, the situation addressed in Table 8 corresponds to a large, long-horizon investor, while results for short horizons are presented for completeness the emphasis is on horizons of 10 years or greater. To conserve space, only the relative allocations to the noncash asset classes are shown (*i.e.*, cash allocations are not shown, although the optimization allowed for cash positions). As before, the relative allocations to the risky asset classes are very similar across target returns; the only significant difference is the amount invested in cash equivalents (not shown).

The addition of REITs as an asset class does little to improve overall portfolio performance. This can be seen by comparing the return standard deviations of the optimal portfolios in Table 8, in which REITs are included, to those in Table 7, in which REITs are not included as an asset class. The standard deviations are identical (at least to the two digit level as shown). The optimal allocations to REITs also reveal that long-term investors gain little from exposure to REITs. For horizons of 10 years or more, only the most conservative

Table 8 ■ Optimal allocations including both real estate and REITs.

	Investment Horizon	Relative Allocations to Noncash Classes: Equities	Bonds	Real Estate	REITs	Std. Dev.
Target = 3.4%	1 year	0.16	0.60	0.19	0.05	0.57
	5 years	0.22	0.60	0.19	−0.01	0.28
	10 years	0.22	0.57	0.24	−0.02	0.21
	25 years	0.19	0.52	0.31	−0.02	0.14
Target = 2.8%	1 year	0.16	0.60	0.19	0.05	0.51
	5 years	0.22	0.60	0.19	−0.01	0.25
	10 years	0.21	0.57	0.24	−0.02	0.19
	25 years	0.18	0.52	0.31	−0.02	0.13
Target = 1.9%	1 year	0.15	0.61	0.19	0.05	0.40
	5 years	0.20	0.61	0.18	0.01	0.20
	10 years	0.20	0.57	0.23	0.00	0.15
	25 years	0.17	0.53	0.30	0.00	0.10
Target = 1.2%	1 year	0.13	0.63	0.20	0.05	0.29
	5 years	0.16	0.62	0.18	0.03	0.14
	10 years	0.16	0.58	0.22	0.03	0.11
	25 years	0.14	0.54	0.29	0.04	0.08

The table shows optimal allocations to bonds, equities, real estate and REITs. Although optimization is based on the use of cash as an asset class, for each optimal portfolio only the relative allocations to the noncash asset classes are shown. For instance, the relative allocation to equities is calculated as optimal allocation to equities divided by the sum of the optimal allocations to equities, bonds and real estate. Allocations are shown for various target returns and at various investment horizons. Std. Dev. is the annualized standard deviation of the allocation (including the effect of optimal level of cash equivalents).

portfolios have any positive allocation to REITs, and even this is only 3%–4% of the total noncash allocation. These results are contrary to both Geltner, Rodriguez and O'Connor (1995) and Feldman (2003), who find that private and public market real estate are both typically represented in optimal portfolios. When considered within the return predictability paradigm used here, however, it seems clear that REITs add little to portfolio performance beyond what can be accomplished with direct real estate.

It seems that the correlation between REITs and direct real estate is high enough for REITs to serve as a proxy in many ways for direct real estate investment. However, for investors with access to both asset classes (such as large pension funds) this means that REITs become redundant. It seems clear given these results why large institutional investors have continued to construct most of their real estate allocation via direct investment and not via public market real estate securities.

While we do not explicitly include transaction costs in our analysis, they are unlikely to affect our result that REITs are redundant for large, long-horizon investors. Amihud and Mendelson (1986) model the interaction of investment horizon and transaction costs and show that long-horizon investors will tend to gravitate toward low-liquidity investments precisely because liquidity is less important for long-term investors but is valued by short-term investors (who drive expected returns up for illiquid investments). Hence, liquidity differences between REITs and direct real estate may actually reinforce our results on the redundancy of REITs for large, long-horizon investors.

We now turn to the case of investors who do not have access to direct real estate as an asset class: small or short-horizon investors. Table 9 presents the results of the portfolio optimization when only cash, stocks, bonds and REITs are available as asset classes (again, only the relative noncash allocations are presented). For short-term investors (1- or 5-year horizon) the optimal allocation to REITs ranges from 4% to 8% of the noncash portfolio. While small in absolute terms, relative to the allocation to equities for these short-term portfolios the REIT allocation is significant. Short-term investors, for whom direct real estate is not viable, would benefit from some allocation to REITs.

The optimal allocation to REITs increases for long investment horizons.[10] For a 25-year investment horizon the REIT allocation ranges from 16% to 20%, with the high end being for the more conservative target returns. For smaller investors (for whom direct real estate is not feasible) with long horizons REITs play a significant role in their optimal asset allocation. This is especially true for more conservative investors for whom the optimal REIT allocation is not much smaller than the allocation to equities overall.

Hence, while it is apparent from the results when both direct real estate and REITs are included that public market real estate does not replace private market investment in a portfolio, when direct investment is infeasible then the public markets can play a significant role in an optimal real estate allocation.[11] This is especially true for long-horizon investors.

[10] Note that the optimal allocation to REITs is not monotonically increasing with investment horizon, except for a target return of 1.2%. Generally, the optimal allocation initially drops from the shortest to the medium term horizons and then rises for the longest term.

[11] Note, however, that the allocation to REITs when direct real estate is not included as an asset class is far smaller than the allocation to real estate when it is included as an asset class.

Table 9 ■ Optimal allocations including REITs, no direct real estate.

	Investment Horizon	Relative Allocations to Noncash Classes			Std. Dev.
		Equities	Bonds	REITs	
Target = 3.4%	1 year	0.21	0.70	0.08	0.60
	5 years	0.29	0.67	0.04	0.30
	10 years	0.30	0.63	0.07	0.24
	25 years	0.29	0.54	0.16	0.18
Target = 2.8%	1 year	0.21	0.70	0.08	0.53
	5 years	0.29	0.67	0.05	0.27
	10 years	0.30	0.63	0.08	0.22
	25 years	0.29	0.55	0.16	0.17
Target = 1.9%	1 year	0.20	0.71	0.08	0.42
	5 years	0.27	0.67	0.06	0.21
	10 years	0.28	0.63	0.09	0.17
	25 years	0.27	0.55	0.17	0.13
Target = 1.2%	1 year	0.18	0.73	0.08	0.30
	5 years	0.23	0.68	0.08	0.15
	10 years	0.24	0.64	0.12	0.13
	25 years	0.24	0.56	0.20	0.10

The table shows optimal allocations to bonds, equities and REITs. Although optimization is based on the use of cash as an asset class, for each optimal portfolio only the relative allocations to the noncash asset classes are shown. For instance, the relative allocation to equities is calculated as optimal allocation to equities divided by the sum of the optimal allocations to equities, bonds and REITs. Allocations are shown for various target returns and at various investment horizons. Std. Dev. is the annualized standard deviation of the allocation (including the effect of optimal level of cash equivalents).

Conclusions

Real estate has traditionally been viewed as low-risk asset class with good diversification properties. Our results indicate that these characteristics depend crucially on the investor's time horizon. Given that most institutional investors have very long time horizons, typical portfolio optimizations based only on short-term returns, without accounting for return predictability, will produce results quite different from true optima for these investors.

Our results show that predictability of the returns to direct real estate leads to mean reversion in those returns over time. Mean reversion seems to be caused by a tendency for commercial property transaction prices to overshoot inflation; higher inflation results in higher prices in a given period, but prices tend to revert the following period. The result of mean reversion is that the risk of real estate investment is much less for long-term investors than for those with shorter investment horizons.

However, mean reversion in real estate is weaker than it is for equities. The weaker mean reversion of real estate returns means that for investors with long (10 years or more) investment horizons real estate has the same level of risk as do equities; a surprising result given the usual description of real estate as a low-risk asset class.

Despite this, optimal allocations to real estate increase with investment horizon (except for very short horizons). This is a result of the fact that real estate-equity and real estate-bond correlations decrease at longer horizons and hence real estate is a better diversifier for long-term portfolios than for short-term. The optimal allocations to real estate reported here are quite high at all investment horizons, but especially so for long-term portfolios. They are certainly larger than the current allocations reported by pension funds, which typically have very long-term horizons.

REITs also exhibit mean reversion and risk that decreases with investment horizon, as well as a correlation with direct real estate that increases with horizon. However, at all horizons REITs remain riskier than direct real estate. When both direct real estate and REITs are available as asset classes REITs play little or no role in optimal portfolios. REITs appear to be a redundant asset class for large, long-horizon investors who can invest in both private and public market real estate.

However, REITs do have a role to play in portfolios in certain circumstances. For small or short-horizon investors, for whom direct real estate is not a feasible investment, REITs do appear in the optimal portfolio. In fact, the size of the optimal REIT allocation is quite large (close to that for equities in some cases) for long-term portfolios when direct real estate is not available as an asset class. This last point is of importance to retail and smaller institutional investors who have a long horizon but may lack the size to invest in the property markets directly.

Our results are based on buy-and-hold portfolios. Given the difficulties in rebalancing direct real estate holdings, this has some intuitive appeal. REITs, however, are liquid enough to be rebalanced on an ongoing basis. An interesting avenue for future research is to examine the interaction between optimal holdings of direct and indirect real estate when portfolios are dynamic and the indirect holdings can be rebalanced. While it is possible that REITs have a greater role to play in long-term portfolios under these conditions, we leave that question to future work.

References

Amihud, Y. and H. Mendelson. 1986. Asset Pricing and the Bid-Ask Spread. *Journal of Financial Economics* 17(2): 223–249.

Barkham, R. and D. Geltner. 1995. Price Discovery in American and British Property Markets. *Real Estate Economics* 23(1): 21–44.

Campbell, J.Y., Y.L. Chan and L.M. Viceira. 2003. A Multivariate Model of Strategic Asset Allocation. *Journal of Financial Economics* 61(1): 42–80.

——. 2007. Replication data for: A Multivariate Model of Strategic Asset Allocation. Murray Research Archive [Distributor]. Available at http://www.economics.harvard.edu/faculty/campbel.

Campbell, J.Y. and L.M. Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press: New York.

——. 2005a. The Term Structure of the Risk-Return Trade-Off. *Financial Analysts Journal* 61(1): 34–44.

——. 2005b. Long-Horizon Mean-Variance Analysis: User Guide. Appendix to Campbell and Viceira (2005a). The Term Structure of the Risk-Return Trade-Off. *Financial Analysts Journal* 61(1): 34–44.

Chiang, K.C.H. and M.-L. Lee. 2007. Spanning Tests on Public and Private Real Estate. *Journal of Real Estate Portfolio Management* 13(1): 7–15.

Chun, G., J. Sa-Aadu and J. Shilling. 2004. The Role of Real Estate in an Institutional Investor's Portfolio Revisited. *The Journal of Real Estate Finance and Economics* 29(3): 295–320.

Clayton, J. 2007. PREA Plan Sponsor Research Report: 2005 and a First Look at 2006 Survey Results. Pension Real Estate Association: Hartford, CT.

Clayton, J. and G. MacKinnon. 2003. The Relative Importance of Stock, Bond and Real Estate Factors in Explaining REIT Returns. *The Journal of Real Estate Finance and Economics* 27(1): 39–60.

Collett, D., C. Lizieri and C. Ward. 2003. Timing and Holding Periods of Institutional Real Estate. *Real Estate Economics* 31(2): 205–222.

Feldman, B. 2003. *Investment Policy for Securitized and Direct Real Estate*. Ibbotson Associates: Chicago.

Fisher, J., D. Geltner and H. Pollakowski. 2007. A Quarterly Transaction-based Index of Institutional Real Estate Investment Performance and Movements in Supply and Demand. *The Journal of Real Estate Finance and Economics* 34(5): 5–33.

Fugazza, C., M. Guidolin and G. Nicodanna. 2007. Investing for the Long Run in European Real Estate. *The Journal of Real Estate Finance and Economics* 34(1): 35–80.

Geltner, D. and J.V. Rodriguez. 1997. Public and Private Real Estate: Performance Implications for Asset Allocation. R.T. Garrigan and J.F.C. Parsons, editors. *Real Estate Investment Trusts: Structure, Analysis and Strategy*. McGraw-Hill: New York.

Geltner, D., J.V. Rodriguez and D. O'Connor. 1995. The Similar Genetics of Public and Private Real Estate and the Optimal Long-Horizon Portfolio Mix. *Real Estate Finance* 12(3): 13–25.

Gentry, W., C. Jones and C. Mayer. 2004. Do Stock Prices Really Reflect Fundamental Values? The Case of REITs. Working Paper. Columbia University, New York.

Hoevenaars, R.P.M.M., R.D.J. Molenaar, P.C. Schotman and T.B.M. Steenkamp. 2008. Strategic Asset Allocation with Liabilities: Beyond Stocks and Bonds. *Journal of Economic Dynamics and Control* 32(9): 2939–2970.

Kandel, S. and R. Stambaugh. 1996. On the Predictability of Stock Returns: An Asset-Allocation Perspective. *Journal of Finance* 51(2): 385–424.

Lee, S. and S. Stevenson. 2005. The Case for REITs in the Mixed-Asset Portfolio in the Short and Long Run. *Journal of Real Estate Portfolio Management* 11(1): 55–80.

Ling, D., A. Naranjo and M. Ryngaert. 2000. The Predictability of Equity REIT Returns: Time Variation and Economic Significance. *The Journal of Real Estate Finance and Economics* 20(2): 117–136.

Liu, C. and J. Mei. 1992. The Predictability of Returns on Equity REITs and Their Co-movement with Other Assets. *The Journal of Real Estate Finance and Economics* 5:401–418.

Markowitz, H. 1952. Portfolio Selection. *The Journal of Finance* 7:77–91.

Mei, J. and C.H. Liu. 1994. The Predictability of Real Estate Returns and Market Timing. *The Journal of Real Estate Finance and Economics* 8:115–135.

Mueller, A.G. and G.R. Mueller. 2003. Public and Private Real Estate in a Mixed-Asset Portfolio. *Journal of Real Estate Portfolio Management* 9(3): 193–203.

Pagliari, J., K. Scherer and R. Monopoli. 2005. Public versus Private Real Estate Equities: A More Refined, Long Term Comparison. *Real Estate Economics* 33(1): 147–187.

Riddiough, T., M. Moriarty and P.J. Yeatman. 2005. Privately versus Publicly Held Asset Investment Performance. *Real Estate Economics* 33(1): 121–146.

Seiler, M.J., J.R. Webb and F.C.N. Myer. 1999. Diversification Issues in Real Estate Investment. *Journal of Real Estate Literature* 7(2): 163–179.

Appendix

Here we present the formal portfolio optimization problem used to solve for optimal buy-and-hold asset allocations over various horizons. The use of continuously compounded returns significantly simplifies the analysis because long-term returns are simply the summation of shorter term returns. However, it also introduces a complication as the continuously compounded return to a portfolio is not a simple weighted average of the continuously compounded returns to each asset. Campbell and Viceira (2002) suggest the following approximation to the return to the k-period return to a portfolio:

$$r_{p,t+k} = r_{0,t+k} + \boldsymbol{\alpha}^{\mathbf{T}}\mathbf{x}_{\mathbf{t}+k} + \frac{1}{2}\boldsymbol{\alpha}^{\mathbf{T}}\left(\mathrm{diag}\{\boldsymbol{\Sigma}_{\mathbf{xx}}\} - \boldsymbol{\Sigma}_{\mathbf{xx}}\boldsymbol{\alpha}^{\mathbf{T}}\right) \tag{A1}$$

where $\boldsymbol{\alpha}$ is a vector of portfolio weights on the noncash asset classes. $r_{p,t+k}$ is the return to a portfolio from period $t+1$ to $t+k$, and other variables are also defined over a k-period horizon. Hence, the optimal weights in $\boldsymbol{\alpha}$ will be a function of k.

Based on this, the portfolio optimization problem faced by an investor is to minimize per-period variance of portfolio returns (over a k-period horizon) subject to attaining a target return per period. More formally

$$\min_{\alpha}\left\{\frac{1}{2}\frac{\mathrm{Var}\left[r_{p,t+k}\right]}{k}\right\} \tag{A2}$$

subject to

$$\frac{E\left[r_{p,t+k}\right]+\frac{1}{2}\mathrm{Var}\left[r_{p,t+k}\right]}{k}=\mu$$

where the denominator in both objective and constraint converts risk and return into per-period equivalents. μ is the target return per period for the portfolio in simple (*i.e.*, not logged) form. The addition of the variance term in the numerator of the constraint converts the continuously compounded expected portfolio return into an expectation in simple return form (see Campbell and Viceira 2002).

The solution to (A2) is

$$\boldsymbol{\alpha}_{\mathbf{k}}=\lambda\boldsymbol{\Sigma}_{\mathbf{xx}}^{\mathbf{-1}}\left(E\left[\mathbf{x}_{\mathbf{t+k}}\right]+\frac{1}{2}\mathrm{diag}\left\{\boldsymbol{\Sigma}_{\mathbf{xx}}\right\}+\boldsymbol{\sigma}_{\mathbf{ox}}\right)-\boldsymbol{\Sigma}_{\mathbf{xx}}^{\mathbf{-1}}\boldsymbol{\sigma}_{\mathbf{ox}} \tag{A3}$$

where σ_{0x} is the vector of covariances between the returns to cash and the excess returns to the other asset classes and λ is the Lagrange multiplier on the constraint.

Note in (A3) that we define the optimal portfolio weights with a subscript k to signify that it varies with the chosen horizon. In our implementation we use the unconditional mean per period of each asset class to calculate $E[x_{t+k}]$. The variance and covariance terms in (A3) are taken from Equation (3) and vary (on a per-period basis) with investment horizon. The sum of the terms of $\boldsymbol{\alpha}_{\mathbf{k}}$ represents the total invested in noncash asset classes; the amount optimally invested in cash is therefore one minus this sum.