



Netspar

Network for Studies on Pensions, Aging and Retirement

Giovanna Nicodano

Massimo Guidolin

Carolina Fugazza

Time and Risk Diversification in Real Estate Investments: Assessing the Ex Post Economic Value

Discussion Paper 2008 - 030

August, 2008

Time and Risk Diversification in Real Estate Investments: Assessing the Ex Post Economic Value*

Carolina FUGAZZA

Center for Research on Pensions and Collegio Carlo Alberto (CeRP-CCA)

Massimo GUIDOLIN**

Manchester Business School and St. Louis FED

Giovanna NICODANO

University of Turin, CeRP-CCA and Netspar

This draft: August 2008

Welfare gains to long-horizon investors may derive from time diversification that exploits non-zero intertemporal return correlations associated with predictable returns. Real estate may thus become more desirable if its returns are negatively serially correlated. While it could be important for long horizon investors, time diversification has been mostly investigated in asset menus without real estate and focusing on in-sample experiments. This paper evaluates ex post, out-of-sample gains from diversification when E-REITs belong to the investment opportunity set. We find that diversification into REITs increases both the Sharpe ratio and the certainty equivalent of wealth for all investment horizons and for both Classical and Bayesian (who account for parameter uncertainty) investors. The increases in Sharpe ratios are often statistically significant. However the out-of-sample average Sharpe ratio and realized expected utility of long-horizon portfolios are frequently lower than that of a one-period portfolio, which casts doubts on the value of time diversification.

JEL Classification Codes: G11, L85.

Keywords: real time asset allocation, real estate, ex post performance, predictability, parameter uncertainty.

* Doug Smith and Yu Man Tam provided excellent research assistance. We thank Frank De Jong and seminar participants at Erasmus University (Tinbergen Institute) and Collegio Carlo Alberto, Turin, for helpful comments. Financial support from Netspar to Massimo Guidolin and Giovanna Nicodano is gratefully acknowledged.

Introduction

Institutional investors diversify their portfolios by investing in public real estate. This practice is supported by empirical studies indicating that the risk return trade-off of optimal portfolios that include real estate improves vs. portfolios that include only standard financial asset classes, such as stocks and bonds. However, such evidence mostly refers to *in-sample* evaluation of portfolio performance, which assumes that portfolio managers know the return distribution far better than they do in the real world. The first goal of this paper is to assess whether public real estate improves portfolio performance *out-of-sample*, realistically assuming that the portfolio manager chooses asset allocation for the future only on the basis of past information on realized returns, which can be at best recursively updated.

Another feature of existing studies of public real estate is their reliance on a one-period mean variance approach, which only allows for a reduction in portfolio risks arising from an imperfect (i.e., low, possibly negative) contemporaneous correlation of asset returns. Thus, real estate becomes more desirable if its return tends to increase when the returns on other assets fall. This allows for static, *across-asset (contemporaneous) diversification*. However, a considerable reduction in risk to long-horizon investors may derive from *time diversification* that exploits non-zero intertemporal return correlations associated with complicated predictability patterns that rely on the linear association between news (shocks) to returns vs. predictor variables. Real estate may thus become more desirable if its returns are negatively autocorrelated in time, ensuring that times of higher returns will follow times of lower returns. While it could be especially important for pension funds and other long horizon investors, time diversification has been mostly investigated in asset menus without real estate (e.g., Campbell, Chan and Viceira, 2003). Since the effects of predictability on multi-period volatility obviously depend on the specific asset menu under examination (see Campbell and Viceira, 2005), our second – and perhaps most important – goal is to assess whether there is scope for time diversification of risk when public real estate is explicitly added to standard asset menus composed of stocks, bonds and bills.

Importantly, our paper investigates the presence of ex post gains from time diversification. Even with modest statistical evidence in favour of predictability, such gains can be substantial when changes in

portfolio performance are measured “in sample”, i.e. assuming that a given statistical framework of analysis correctly measures the features of the distribution of asset returns that are of importance. For instance, Hoevenaars et al. (2007) and Fugazza, Guidolin, and Nicodano (2007) argue that expanding the asset menu to include public real estate, while accounting for predictable asset returns, may significantly increase the *ex ante* welfare of a standard, expected-utility maximizing investor. However, these findings are retrospective in nature. We therefore compute *ex post* gains, mimicking a risk-averse investor who relies on past evidence concerning return predictability in order to decide on future multi-period asset allocation.

This experiment would be of negligible interest if it considered only one specific sample period, as the outcome could be due to bad or good luck. Hence, we average portfolio performance attained over 120 simulated portfolio allocations which are obtained with the following recursive method. We first use data from January 1972 up to December 1994 to estimate the parameters of our prediction model and to forecast multi-period means, variances, and covariances of returns on all asset classes, which allow us to determine optimal portfolio weights. This exercise is repeated the following month, using data up to January 1995 to compute afresh forecasts of return moments and select portfolio weights. Iterating this recursive scheme until November 2004 generates a sequence of realized portfolio returns from which ex-post performance measures of optimal portfolios are computed. Our evaluation of the role of real estate when returns are predictable thus averages times of good and bad performance for this as well as other asset classes. This experiment is repeated for six alternative investment horizons, ranging from one month to 5 years. Moreover, we allow not only for the standard, Classical approach, but also for a Bayesian approach (as in Barberis, 2000) in computing optimal portfolio composition. Indeed, welfare can substantially increase if the investor takes into account the uncertainty in forecasts by using Bayesian updating (Jorion, 1985; Kandel and Stambaugh, 1996), especially when return predictability appears weak according to statistical tests.¹

We use a simple vector autoregressive framework to capture predictable variations in the investment opportunity set (as in Campbell, Chan, and Viceira, 2003, Geltner and Mei, 1995, and Glascock, Lu, and

So, 2001) and solve a portfolio problem in which investor's preferences are specified as a power utility of terminal wealth. In most cases, the optimal average weight to be assigned to real estate vehicles (E-REITs) is large for a 1-month investor, given a high expected return-to-volatility ratio. As the horizon grows, the attractiveness of stocks and real estate improves relative to cash for Classical investors, given their favorable multi-period time diversification properties implied by predictability. However, in a Bayesian framework parameter uncertainty is so high that it reverses this pattern, echoing results in Barberis (2000). The optimal share invested in REITs by a Bayesian investor falls from 43% for a 1-month horizon to 33% for a five-year horizon, while the optimal average allocation to short-term T-Bills grows from 21% to 50%, because their return is precisely anticipated by several predictors. The optimal average allocation to bonds drops from over 25% for a short horizon to almost zero for a 5-year horizon. This result is similar to findings in previous experiments with different frequencies and asset menus that have been reported in the literature (see Campbell and Viceira, 2005), irrespective of the estimation method adopted. These changes in portfolio composition indicate that return predictability patterns suggest that a multi-period strategy should optimally exploit not only contemporaneous diversification but time diversification of risk as well.

Ex ante, an investor with a multi-period horizon would always be at least as well off as a $T = 1$ investor, because she can always choose to overlook predictability and the opportunities for time diversification. However, ex post, the difference in realized expected utility could be negative – for instance, because the prediction model is misspecified and mis-predicts the auto and serial cross-correlations that characterize asset returns. We therefore measure the differential ex-post performance of a one-month versus a 5-year investor, in order to assess whether there are gains from intertemporal hedging brought about by return predictability. Our results contribute more evidence to the skeptical view on the ex post value of predictability (e.g., Goyal and Welch, 2004) within linear frameworks, in that the out-of-sample average Sharpe ratio of a long-horizon portfolios is often lower than that of a one-period investor. For instance, it drops from 0.45 for a $T = 1$ horizon to 0.30 for a $T = 60$ horizon, when a Bayesian investor uses an asset menu that includes real estate. Similar results are obtained for horizons shorter than 60 months, as well as for the asset menu without real estate. One may argue that the Sharpe ratio, which is based on the mean and

the variance of returns only, can be a misleading indicator of performance when returns are not normally distributed (see Leland, 1999, Goetzman et al., 2002, 2004).² For this reason, we also study the ex post welfare gains from time diversification, that account from changes in the skewness and higher order moments of wealth. Results based on welfare gains confirm, however, the results based on Sharpe ratios. The average certainty equivalent of welfare for a Bayesian investor drops from 8.75% to 6.77% when the horizon grows to 5 years.

On the contrary, contemporaneous diversification into real estate vehicles increases both the Sharpe ratio and the certainty equivalent of wealth for all investment horizons and for both Bayesian and Classical investors. These results extend to an ex post setting the evidence in Fugazza, Guidolin, and Nicodano (2007) concerning European data and ex ante performances, as well as results obtained in mean variance models (see Seiler, Webb, and Myer, 1999, and Feldman, 2003).³ The annualized percentage increase in initial wealth that should be awarded to a Bayesian investor in order to compensate her for excluding REITs from her asset menu ranges from 1.39 to 2.59 percent of initial wealth, depending on her horizon. Ex post welfare gains are even larger, apart from the 1-month horizon case, for a Classical investor who overlooks estimation risk when choosing portfolio composition.

Several recent papers have explored whether predictability improves the ex-post performance for an investor with a one period horizon, who therefore only exploits market timing (when predictability is modeled) and contemporaneous diversification opportunities. Returns to market timing appear to be positive for a Bayesian investor in a mean variance framework (Avramov and Chordia, 2006; Abhyankar, Basu, and Stremme, 2005; Wachter and Warusawitharana, 2005) even though they can turn negative when an investor tries to guess ex-ante which are the best predictors for returns (Cooper, Gutierrez, and Marcum, 2005). Other papers also analyze out-of-sample performance of investment in public real estate with predictable returns. Some find that active strategies outperform passive ones (Liu and Mei, 1994), even after deducting transaction costs (Bharati and Gupta, 1992). This is no longer the case in more recent studies, such as Nelling and Gyourko (1996) and Ling, Naranjo, and Ryngaert (2000), who find it difficult

to exploit predictability ex post particularly in the 1990s. While these studies focus on short term portfolio strategies, our paper completes the picture by investigating the ex-post welfare gains of time diversification in a multi-period setting. We find that ex-post economic value of time diversification is negligible before accounting for transaction costs, even if linear prediction performs reasonably well. However, the contemporaneous diversification opportunities offered by real estate vehicles remain substantial and likely in excess of most sensible measures of transaction costs.

The plan of the paper is as follows. Section 1 briefly outlines the methodology of the paper. Section 2 describes the data and reports results on their statistical properties, revealing predictability in the dynamics of the investment opportunity set. Section 3 characterizes optimal portfolios which include real estate, and compares them to the case without real estate. It also assesses the gains from intertemporal risk diversification. In Section 5, we calculate the welfare costs of ignoring either predictability or real estate. Section 6 concludes. A final Appendix collects details on the statistical models and solution methods employed in the paper.

The Asset Allocation Model

In this paper we proceed to compute forecasts of the basic sample moments relevant to portfolio choice and, as a result, optimal portfolio shares using both Classical and Bayesian econometric approaches. In the classical case we estimate the parameters characterizing a set of simultaneous linear relationships that link returns to the predictors. We then use the current realization of the predictors to compute conditional (predictive) moments and distributions for future asset returns that take the estimated parameter values as given and use them in place of the true (and yet, unknown) parameter values. Since the expected utility maximization that describes the portfolio problem is solved using predictive densities that ignore the parameter estimators are themselves random variables, an important source of uncertainty (sometimes called *estimation risk*) is ignored. Such an approach is called classical, as typical in the literature (see e.g., Barberis, 2000). In the Bayesian case we specify a (weak, uninformative) set of beliefs concerning the parameters characterizing the linear relationships among asset returns and predictors that the investor might

have prior to viewing the data. A posterior distribution of such parameters is then obtained – by Bayes' rule – conditional on the observation of the predictors, which is used to generate a conditional, predictive distribution of returns and – as a result – a predictive distribution of future utility levels. By maximizing the expectation of such predictive utility density by selecting portfolio shares, the optimal asset allocation is computed and characterized. Sections 2.1 and 2.2 provide a few additional details on the asset allocation frameworks employed in this paper, while an Appendix gives a primer on the two asset allocation frameworks employed.

Classical Portfolio Choice

Consider an investor with constant relative risk aversion, $\gamma > 1$, who maximizes the expected utility derived from her final wealth, accumulated after T months, by choosing a vector of optimal portfolio weights at time t (ω_t),

$$\max_{\omega_t} E_t \left[\frac{W_{t+T}^{1-\gamma}}{1-\gamma} \right] \quad \gamma > 1,$$

and holding the asset composition of her portfolio constant until time $t+T$. Wealth can be invested in stocks, bonds, real estate, with continuously compounded *excess* returns between month $t-1$ and t denoted by r_t^s , r_t^b , and r_t^r , respectively. The asset menu is completed by the possibility of investing in cash (say, short-term government zero coupons). We realistically model the return on cash, r_t^f , as random over time; notice however that by construction (because it is characterized as an essentially default-free zero coupon bond) the short-term investment is free of risk over $[t-1, t]$, i.e., its yield is deterministic over a one-period holding interval.

When initial wealth W_t is normalized to one, the process for investor's terminal wealth is given by:

$$W_{t+T} = \omega_t^s \exp(R_{t,T}^s) + \omega_t^b \exp(R_{t,T}^b) + \omega_t^r \exp(R_{t,T}^r) + (1 - \omega_t^s - \omega_t^b - \omega_t^r) \exp(R_{t,T}^f),$$

where ω_t^j is the fraction of wealth invested in the j -th asset class and $R_{t,T}^j$ denotes the cumulative returns between t and T :

$$R_{t,T}^j \equiv \sum_{k=1}^T (r_{t+k}^j + r_{t+k}^f), \quad j = s, b, r; \quad R_{t,T}^f \equiv \sum_{k=1}^T (r_{t+k}^f)$$

Call n the number of asset classes included in the asset menu. Our baseline experiment concerns $n = 4$. If there are no-short sale constraints, the problem can be stated as:

$$\max_{\omega_t} E_t \left[\frac{\left\{ \omega_t^s \exp(R_{t,T}^s) + \omega_t^b \exp(R_{t,T}^b) + \omega_t^r \exp(R_{t,T}^r) + (1 - \omega_t^s - \omega_t^b - \omega_t^r) \exp(R_{t,T}^f) \right\}^{1-\gamma}}{1-\gamma} \right] \quad (1)$$

s.t. $1 > \omega_t^s \geq 0 \quad 1 > \omega_t^b \geq 0 \quad 1 > \omega_t^r \geq 0.$

Time-variation in (excess) returns is modeled using a simple Gaussian VAR(1) framework, as in most of the finance literature on time-varying investment opportunities (see the review in Campbell and Viceira, 2003):

$$\mathbf{z}_t = \boldsymbol{\mu} + \boldsymbol{\Phi} \mathbf{z}_{t-1} + \boldsymbol{\varepsilon}_t, \quad (2)$$

where $\boldsymbol{\varepsilon}_t$ is i.i.d. $N(\mathbf{0}, \boldsymbol{\Sigma})$, $\mathbf{z}_t \equiv [r_t^s \quad r_t^b \quad r_t^r \quad r_t^f \quad \mathbf{x}_t']'$, and \mathbf{x}_t represents a vector of economic variables able to forecast future asset returns. Model (2) implies that

$$E_{t-1}[\mathbf{z}_t] = \boldsymbol{\mu} + \boldsymbol{\Phi} \mathbf{z}_{t-1},$$

i.e. the conditional risk premia on the assets are time-varying and function of past excess asset returns, past short-term interest rates, as well as lagged values of the predictor variable \mathbf{x}_{t-1} . The Appendix provides further details on the characterization of the joint predictive density for asset returns in this case.

This problem can be solved by employing simulation methods similar to Kandel and Stambaugh (1996), Barberis (2000), and Guidolin and Timmermann (2005):

$$\max_{\omega_t} \frac{1}{N} \sum_{i=1}^N \left[\frac{\left\{ \omega_t^s \exp(R_{t,T}^{s,i}) + \omega_t^b \exp(R_{t,T}^{b,i}) + \omega_t^r \exp(R_{t,T}^{r,i}) + (1 - \omega_t^s - \omega_t^b - \omega_t^r) \exp(R_{t,T}^{f,i}) \right\}^{1-\gamma}}{1-\gamma} \right], \quad (3)$$

where $\{R_{t,T}^{s,i}, R_{t,T}^{b,i}, R_{t,T}^{r,i}, R_{t,T}^{f,i}\}_{i=1}^N$ are obtained simulating from the process in (2) N times. To obtain the results that follow, we have employed $N = 100,000$ Monte Carlo trials in order to minimize any residual random errors in optimal weights induced by simulations.

Bayesian Portfolio Choice

We incorporate estimation risk in the model by using a Bayesian approach as in Barberis (2000). This relies

on the principle that portfolio choice ought to be based on the multivariate predictive distribution of future asset returns that also “integrates over” (i.e., accounts for) the fact that estimated coefficients within the simple VAR framework in (2) do possess a distribution because they are just sample statistics.⁴ Such a predictive distribution is obtained by integrating the joint distribution of $\boldsymbol{\theta}$ and returns $p(\mathbf{z}_{t,T}, \boldsymbol{\theta} | \ddot{\mathbf{Z}}_t)$ with respect to the posterior distribution of $\boldsymbol{\theta}$, $p(\boldsymbol{\theta} | \ddot{\mathbf{Z}}_t)$:

$$p(\mathbf{z}_{t,T}) = \int p(\mathbf{z}_{t,T}, \boldsymbol{\theta} | \ddot{\mathbf{Z}}_t) d\boldsymbol{\theta} = \int p(\mathbf{z}_{t,T} | \ddot{\mathbf{Z}}_t, \boldsymbol{\theta}) p(\boldsymbol{\theta} | \ddot{\mathbf{Z}}_t) d\boldsymbol{\theta},$$

where $\ddot{\mathbf{Z}}_t$ collects the time series of observed values for asset returns and the predictor, $\ddot{\mathbf{Z}}_t \equiv \{\mathbf{z}_i\}_{i=1}^t$. The portfolio optimization problem becomes then:

$$\max_{\boldsymbol{\omega}_t} \int \frac{W_{t+T}^{1-\gamma}}{1-\gamma} p(\mathbf{z}_{t,T}) \cdot d\mathbf{z}_{t,T}.$$

In this case, Monte Carlo methods require drawing a large number of times from $p(\mathbf{z}_{t,T})$ and then “extracting” cumulative returns from the resulting vector. The Appendix provides further details on methods and on the Bayesian prior densities, which we simply assume to be of a standard uninformative diffuse type.⁵ In particular, since applying Monte Carlo methods implies a double simulation scheme (i.e., one pass to characterize the predictive density of returns, and a second pass to solve the portfolio choice problem), in the following N is set to a relatively large value of 300,000 independent trials that are intended to approximate the joint predictive density of excess returns and predictors.

Data and Descriptive Statistics

Our sample of monthly data runs from January 1972 to November 2004 for a total of 371 observations, as the public real estate data we use are available for this time span only. Importantly, the sample period includes several stock market cycles. The NAREIT website (www.NaReit.com) provides monthly returns on US equity REITS. Stock returns are derived from the value weighted CRSP index of listings on the NYSE, NASDAQ and the AMEX. The 10-Year constant maturity portfolio returns on US government Bonds as well as the 3-month T-bill come from the Federal Reserve Bank of St. Louis database (FREDII®).

We use continuously compounded total return market-capitalization indices, including both capital gains and income return components. Excess returns are calculated by deducting short-term cash returns from total returns. The short-term investment yield is expressed in real terms as the difference between the nominal yield and the seasonally-adjusted monthly rate of change in the consumer price index for urban consumers provided by FREDII®.

As for the choice of economic predictors, we follow Ling, Naranjo and Ryngaert (2000) by using the dividend yield computed on the CRSP index, the term and default spreads, and the realized inflation rate as predictors of asset returns.⁶ As commonly done, the dividend yield is computed as the ratio between the moving average of the 12 most recent monthly cash dividends paid out by companies in the CRSP universe, divided by the $t - 12$ value-weighted CRSP price index. The term spread is the difference between the yield on a portfolio of long term US government bonds (10 year benchmark maturity) and the yield on 3-month Treasury Bills; both series are downloaded from the FREDII® Database and both yield series are annualized. The default spread is measured as the yield difference of BAA corporate bonds and the 10-year constant-maturity Treasury bond yield series, both from FREDII® and annualized.⁷ Since much literature has documented a relationship between real estate returns and the rate of inflation (see e.g., Karolyi and Sanders, 1998), we also augment the space of predictor variables by the inflation rate, measured as the continuously compounded rate of change of the Consumer Price Index For All Urban Consumers (all Items, seasonally adjusted, again from FREDII®). Finally, (VAR(1) model) implies by construction that past asset returns forecast future returns as well as the future values of the four predictors.

In Table 1 we present summary statistics for the eight times series under investigation. In fact, to support interpretations that are offered in Section 4, the table entertains nine different series because it covers both nominal and real short-term interest rates. Panel A of Table 1 refers to our complete sample period (1972-2004), while Panel B concerns the sample used for estimating the initial parameters of the linear predictability model (1972-1994), with the purpose of initializing our recursive scheme of estimation, portfolio optimization, and ex-post realized performance evaluation.

Over the complete sample, securitized real estate dominates (in mean-variance terms) the stock market, in spite of the stock market boom that has characterized the mid and late 1990s. Public real estate performs better than equities in mean terms (6.0 and 4.0 percent in annualized terms and in excess of short-term yields, respectively), and is less volatile than stocks (their annualized standard deviations are 13.8 vs. 16.1 percent, for REITs and equities, respectively). Correspondingly, the (monthly) Sharpe ratio of real estate (0.13) is almost twice the ratio for equities (0.07). As one would expect, bonds have been less profitable (1.5 percent per year in excess of short-term yield) but also less volatile (7.8 percent in annualized terms) than stocks and real estate. The corresponding Sharpe ratio is, however, rather low, only 0.05. In real terms, short-term T-bills have given a non-negligible average yield of 1.9 percent per year with a very small annualized volatility of 1.2 percent only, as one would expect.

The lower part of Panel A displays simultaneous correlations. The performances of the four asset classes are only weakly correlated, with a peak correlation coefficient of 0.57 between excess stock and real estate returns. Excess bond returns are characterized by correlations vs. both equities and real estate which are lower than 0.2. Under these conditions, there is wide scope for contemporaneous portfolio diversification across risky assets. Even lower is the correlation of the real return on T-Bills with stocks and E-REITs, which never exceeds 0.12: therefore we expect a relatively large demand of T-Bills for hedging purposes.

Panel B shows summary statistics for a shorter, 1972-1994 sample period used to initialize our recursive portfolio experiments. We briefly discuss the features of our data series also as means to document the robustness (stationarity) over time of the main statistical features discussed above. The sub-sample 1972-1994 is qualitatively similar to the full sample, although investment opportunities during the 1970s and 1980s turned out to be significantly worse than in the 1990s and the following half decade. For instance, Sharpe ratios are lower (from 0.03 for bonds to 0.09 for public real estate) and all correlations increase, when compared to our full-sample period. This is because the first few years (1972-1976), which have a relatively larger weight in this shorter sample, are characterized by a well-known, supply-side induced recession that caused (ex-post) Sharpe ratios to turn negative for both stocks and E-REITs. However, it is

remarkable that the mean, volatility, and Sharpe ratios ranking across risky assets is entirely preserved vs. panel A: also during our initial sample, public real estate gave a considerably higher mean excess return than equities (in fact, almost double, 4.5 vs. 2.4 percent per year), a somewhat lower standard deviation (14.1 vs. 16.0 percent), and a more appealing reward-to-risk ratio. Correlations are generally similar to those reported in Panel A.

Predictability Patterns in a VAR Model

Table 2 reports estimates of the VAR coefficients in (2), for the case in which classical estimation methods are employed over our complete sample. Robust t-statistics are reported in parenthesis, under the corresponding point estimates. We boldface p-values equal to or below 0.05. The lower part of the table displays MLE estimates of the covariance matrix of the VAR residuals. The table shows that a number of variables are able to predict future real estate performance.⁸ Real estate returns are positively related to lagged stock and bond returns with positive and significant coefficients, as if wealth effects from financial securities were capable to systematically spill over to the real estate market. REITs returns are negatively related to both the lagged term spread and the short term real rate, with coefficients that are large and economically meaningful (-5.81 and -4.67 imply that a one standard deviation shock to either the termspread or to the real short rate will induce changes of 1.01 and 1.71 percent in monthly REIT returns). We also find evidence of forecasting power of both the dividend yield and inflation for real estate returns and also in this case the effects are not only statistically significant, but economically important as well (i.e., a one standard deviation shock to the dividend yield forecasts a change of REITs returns of 1.17 percent and in the same direction, while an identical shock to the inflation rate predicts a REIT return change of 2.16 percent, but in the opposite direction). Interestingly, higher current inflation forecasts lower future public real estate returns, which is a finding similar to the one typical in the equity literature.

The “amount” of predictability characterizing stocks is similar to the one found for REITs. The only two differences is that in the case of equities, lagged asset returns have no forecasting power, while the coefficients characterizing all the relevant predictors (once more, the dividend yield, the term spread, real

short term rates, and inflation) turn out to be remarkably larger than in the REIT case. For instance, while a one-standard deviation shock to the dividend yield was predicting a 1.17 percent increase in REIT returns, the matching prediction for equities is 1.73 percent in the case of equities. The signs of all these effects are completely consistent with the literature as far as the dividend yield (see Barberis, 2000), the term spread (see Avramov, 2002), and the effects of real short term rates (see Keim and Stambaugh, 1986). The largest economic difference concerns, consequently, the effect of shocks to the term spread, as the estimated coefficient for equities is almost double the estimated effect for public real estate. These results seem to validate the view that public real estate may be just a special type (sectorially characterized) of equity security. Bond returns are instead scarcely predictable, in the sense that only the dividend yield, the default spread, and inflation seem to have a marginally significant predictive power. However, only the economic importance of the default spread is non-negligible, in the sense that a one-standard deviation shock to the default spread triggers a 0.34 percent reaction in bond returns.

Interestingly, real one-month T-bill returns are precisely predicted by all predictors as well as by lagged bond and REIT returns. Additionally, as one should expect in the light of the finance literature debating whether short-term rates contain a unit root, real one-month T-bill returns contain also a substantial degree of persistence (the coefficient is in fact in excess of unity, although this has no implications for stationarity as the entire vector autoregressive system does turn out to be stationary).⁹ We also note that the dividend yield, the term spread, and inflation predict subsequent real short term rates with highly accurate, yet rather small coefficients (which do imply weak economic significance), but also with signs which are systematically opposite vs. those found in the case of public real estate and stocks. These are the same patterns documented by Campbell, Chan and Viceira (2003). In particular, while past inflation is negatively correlated with future stock, bond and real estate returns (with estimated VAR(1) coefficients of -9.5, -2.4, and -6.3, the only predictive influence that links an increase in a rate of return with higher inflation concerns 1-month T-bills. Therefore, short-term, essentially default risk-free investments are the only ones that provide a hedge against inflation shocks.

The lower panel of Table 2 shows the volatilities of the VAR residuals on the main diagonal, pair-wise covariances below the main diagonal, and pair-wise correlations above the main diagonal. As argued in Barberis (2000), such pair-wise shock correlations are crucial because the behavior of the variance of portfolio returns depends on them as the investment horizon grows. For instance, even focusing on the simple case of one risky asset with nominal return R_{t+1} (and excess return $r_{t+1} \equiv R_{t+1} - r^f$) predicted by a variable x_t , it is clear that a VAR(1) framework implies that while $Var_t[R_{t,1}] = Var_t[R_{t+1}] = \sigma^2$, with a two-period investment horizon:

$$\begin{aligned} Var_t[R_{t,2}] &= Var_t \left[\sum_{k=1}^2 (r_{t+k} + r^f) \right] \\ &= Var_t \left[\sum_{k=0}^1 (r^f + \mu + \phi_{11} r_{t+k} + \phi_{12} x_{t+k} + \sigma^2 \varepsilon_{t+k} + \rho \sigma_x \varepsilon_{t+k}^x) \right] \\ &= 2\sigma^2 + (\phi_{12})^2 \sigma_x^2 + 2\phi_{12} \rho \sigma \sigma_x, \end{aligned}$$

where ϕ_{12} is the VAR coefficient that measures the effect of x_t on r_{t+1} , σ_x^2 is the variance of the shocks to the predictor, and ρ is the correlation between shocks to the predictor and shocks to excess returns. It is now easy to show that the conditional variance of the asset return grows at a slower rate than the horizon if and only if $(\phi_{12})^2 \sigma_x^2 + 2\phi_{12} \rho \sigma \sigma_x < 0$, which may occur if and only if ρ and ϕ_{12} have opposite signs. This result makes the sign of the correlation between VAR shocks crucial, given ϕ_{12} . When $Var_t[R_{t,T}]/Var_t[R_{t,1}]$ declines as the horizon T grows and $\phi_{12} > 0$, the economic interpretation is that when the predictor falls unexpectedly (i.e. it is hit by some adverse shock), $\rho < 0$ implies that the news will be likely accompanied by a positive, *contemporaneous* shock to excess asset returns. On the other hand, since $\phi_{12} > 0$, a currently declining dividend yield forecasts *future* lower risk premia on stocks and real estate. Hence such a parameter configuration leads to a built-in element of negative serial correlation, as it is easy to show that processes characterized by negative serial correlations are less volatile in the long-run than in the short-run, due to mean-reversion effects.¹⁰

The lower panel of Table 2 does highlight a few large and negative pair-wise correlations between excess

asset returns and predictors. This happens between excess stock returns and the dividend yield (-0.89), between REIT returns and the dividend yield (-0.56), and between the real short term rate and the term spread (-0.44). Since the predictive relationship between these three couples is always positive, such negative correlations imply an element of mean reversion that can make stocks, REITs, and short-term T-bills increasingly attractive as the investment horizon grows. Such findings are ubiquitous in the literature analyzing US equity data (see e.g., Barberis, 2000), implying that stocks and real estate are a good hedge against adverse future dividend yield news. Fugazza, Guidolin, and Nicodano (2007) find an identical result on European real estate data. However, a few correlation coefficients are always positive and highly significant, especially between stock and REITs residuals. This means that unexplained (residual) returns in both stocks and public real estate will tend to appear together and this makes both real estate and equities riskier than they would otherwise be. Clearly, at an empirical level it remains unclear which of the two effects – i.e., the mean-reversion implied by the negative correlation between dividend yield and return news, or the mean-aversion implied by the positive correlation between real estate and equity return news – will prevail.

Additionally, REITs do not appear to be good hedges against news affecting bond returns, inflation and the real short-term rate. Consider for instance an inflationary surprise. This will be associated with lower contemporaneous returns on REITs, since the correlation coefficient is equal to -0.13. In turn, lower inflation predicts lower expected real estate risk premia, since the VAR coefficient is equal to 0.23. Thus, lower returns associated with inflationary surprises will tend to persist, making REITs relatively riskier in this respect. On the contrary, stocks are good hedges with respect to shocks to the term spread and to the real short rate. Thus, stock returns become relatively less risky than REITs return over a longer horizon.

One final remark about the lower panel of Table 2 concerns the behavior of the risk of the real short term rate as the investment horizon grows. On the one hand, the standard deviation of the short term real rate is unconditionally very low, as displayed in Table 1, a tiny 1.22 percent per year. On the other hand, this asset becomes even less risky over a longer horizon. Indeed, its VAR residuals display negative contemporaneous correlation with innovations to both the term spread (-0.44) and the inflation rate (-0.87),

which help in predicting future short term real return with a positive coefficient (1.17 and 0.95, respectively). Thus, an inflationary surprise -- while decreasing the contemporaneous return to short term assets -- is expected to be associated with higher future real short term rates.¹¹ Of course, these heuristic arguments concerning the behavior of the variance of the different asset classes as a function of the horizon hardly map in precise quantitative results on optimal portfolio choices as a function of the horizon. Below, we will see how these features have implications for multi-period portfolio choice.

All in all, the classical results in Table 2 reveal the presence of a considerable “amount” of predictability of non-negligible strength, even when a rather rudimentary VAR(1) model is used to capture predictability in asset risk premia. However, since the seminal paper by Jorion (1985), it is well known that ex-post performance may improve significantly when Bayesian estimation techniques – which are able to model and quantify the estimation risk implicit in a given econometric model – are deployed to support optimal portfolio choice. This is why we repeat the econometric analysis employing Bayesian estimation methods to derive the joint posterior density for the unknown parameters and hence the joint predictive density of asset returns that allows us to compute optimal portfolio choices. Table 3 reports the means of the marginal posteriors of each of the coefficients in predictive coefficient matrix \mathbf{C} (further defined in the Appendix) along with the standard deviation of the corresponding marginal posterior, which provides a measure of the uncertainty involved.

The posterior means in Table 3 only marginally depart from the MLE point estimates in Table 2, a fact which is consistent with previous findings in the financial econometrics literature on return predictability. However, it is possible to notice that the additional variance of the coefficients caused by the presence of estimation uncertainty reduces our empirical ability to accurately predict bond returns and, to a lesser extent, real estate returns, as in Avramov (2002). For completeness, we also report in the last panel of Table 3 the posterior means and standard deviations (in parenthesis) for the covariance matrix Σ . Most elements of Σ have very tight posteriors and all the implied correlations are identical (to the second decimal) to those found under MLE. Therefore similar comments about the economic meaning and implications of the

econometric estimates also apply to the Bayesian results in Table 3.

Optimal, Real-Time Portfolio Choice

The main exercise of this paper consists of a fully recursive scheme of model estimation and optimal portfolio optimization. In particular, we initialize our experiment using data from January 1972 up to December 1994 to estimate the parameters of our VAR(1) prediction model and to forecast multi-period means, variances, and covariances of returns on all asset classes, which allow to determine optimal portfolio weights in a classical framework. We also proceed to numerically characterize the Bayesian joint posterior densities (see the Appendix for details) for the coefficients in \mathbf{C} and $\mathbf{\Sigma}$ and the joint predictive density for future asset returns. We then compute optimal portfolio weights both in the classical framework that ignores estimation risk, as well as in the Bayesian one. This is done imputing to our “hypothetical” investor a range of alternative, potential investment horizons parameterized by T ; in fact we use six alternative horizons, ranging from 1 month to 5 years.¹² These recursive estimation and portfolio choice exercises are repeated on the following month, using data from January 1972 and up to January 1995 to compute afresh forecasts of return moments and Bayesian joint predictive densities, and to select (ex-ante) optimal portfolio weights. Iterating this recursive scheme until November 2004 generates a sequence of 120 sets of optimal portfolio shares – importantly, one for each possible investment horizon – as well as realized portfolio returns from such ex-ante optimal choices, from which ex-post performance measures for these alternative portfolio strategies and horizons may be computed.

Portfolio Diversification

Table 4 reports the optimal mean portfolio weights over the full sample period 1972-2004 for alternative investment horizons. Although tracking the dynamics of the weights over time may have some interest, in this paper we are mostly interested in the ex-post performance of optimal portfolios and therefore purely concentrate on the “average” picture concerning the behavior of the investor. We first focus on the effects on portfolio composition from enlarging the asset menu to REITs, for a one-month investor. When securitized real estate is exogenously ruled out from the asset menu, a short-term investor who ignores

estimation risk overweighs 10-year Treasury bonds (43%) and stocks (38%) and underweights cash (19%) relative to an equally weighted portfolio. The monthly Sharpe ratio of stocks and bonds is indeed considerable over both the initial sample (0.05 and 0.03) and the full sample (0.07 and 0.05), while the correlation between these two assets falls around rather moderate values of 0.31 to 0.18. The fact that an investor would underweight cash is further rationalized by the observation that its real return is positively correlated with excess stock and bond returns.

When REITs are re-introduced in the asset menu, public real estate ends up crowding out the other risky assets, due to its very high Sharpe ratio (0.13 in the full sample and 0.09 in the initial one) along with relatively low correlation with short-term bills.¹³ The share destined to stocks is particularly moderated by the introduction of real estate (from 38% to 13%), given the relatively high correlation between equity returns and the size of the portfolio share optimally allocated to real estate (50%). Bond holdings fall from 43% to 27%, while cash from 19% to 10%. Therefore the displacement effect caused by public real estate on the average holdings of other assets over time is rather substantial.

A Bayesian investor, by considering estimation risk, turns out to be more cautious, with lower holdings of all risky assets (especially, REITs, down to 43%), compensated by larger holdings of cash (21%). However the tendency to overweigh real estate relative optimally to the equally weighted portfolio remains of first-order magnitude.

Interestingly, as the investment horizon grows, in the asset menu that includes public real estate, the reaction of the REITs weight becomes entirely a function of whether the investor employs or does not employ a Bayesian estimation framework. A Bayesian investor would allocate a steeply declining weight to real estate (and, but only marginally, to bonds) as the horizon grows, with the shares of stocks and especially short-term deposits strongly increasing. For instance, a 5-year Bayesian investor would invest 33% in real estate (down from 43% for $T = 1$ month), 52% in cash (up from 21%), and 14% in stocks (up from 11%). This makes sense if assets imply an increasing/decreasing quantity of cumulative (“compounded”) estimation risk as the investment horizon changes. On the opposite, a classical investor

would increase the share invested in public real estate (from 50% to 59%) and stocks (from 13% to 29%), and to the contrary reduce the share of bonds to essentially nothing as the horizon increases. As a result, a classical long-run investor ends up choosing a portfolio that is considerably riskier than both a short- or long-run Bayesian investor who is concerned not only with the intrinsic risk of the assets, but also with estimation risk. These results generally agree with both typical results in the finance (see e.g., Barberis, 2000) and real estate (see e.g., Fugazza, Guidolin, and Nicodano, 2007) literatures. These wide differences between short- and long-run portfolios lead us to offer a few additional thoughts on opportunities of diversification over time offered by real estate investments.

Time Diversification

It is well known that a one-period asset allocation may differ substantially from a long-term one when returns are predictable, while the investor's planning horizon is irrelevant for portfolio choice when returns are independently and identically distributed (Samuelson, 1969, Merton, 1969). For instance, mean reversion in asset returns will lead to a positive intertemporal hedging demand in a multi-period dynamic setting, to be added to the speculative demand. The impact of time diversification on the allocation can therefore be gauged by comparing the allocation for $T = 1$ with the ones for longer horizons, in our case $T = 60$. Here we consider an investor who ignores parameter uncertainty.¹⁴

The percentages invested in bonds are 27% vs. 2% , 13% vs. 19% for stocks, 50% vs. 54% for real estate and 10% vs. 25% for cash. Overall, real estate and stocks – the riskier assets – account for 73% of a Classical long term portfolio versus an already surprising 63% for a short term portfolio. Thus predictability implies a shift out of bonds by 25% , and into stocks (+6%), real estate (+4%), and short term deposits (+15%).

Clearly, the assets whose long-run risk/return trade-off is mostly improved by the mean-reversion effects implied by (VAR(1) model) are in lower demand for a short horizon than for a longer horizon. Unreported simulations reveal that conditional variances of bond returns increase when accounting for predictability.

This occurs because shocks to bond returns are positively correlated with shocks to the default spread (0.35), which helps in predicting future bond returns with a large coefficient (7.05). Thus an unexpected reduction in the default premium is associated with both lower unexpected bond returns today and lower expected bond returns tomorrow, a fact that makes a two period bond portfolio riskier than in the absence of such intertemporal links. On the contrary, short term deposits become less risky over multi-period horizons as their return is negatively correlated with contemporaneous surprises to both the term spread (-0.44) and the inflation rate (-0.87), which help predicting future short term real return with a positive coefficient (1.17 and 0.95, respectively).

Due to its high persistence coupled with the strong negative correlation between shocks to returns and shocks to the dividend yield, Campbell, Chan, and Viceira (2003) find that the dividend yield generates a very large intertemporal hedging demand for stocks. Here, this effect is reinforced by the intertemporal link between stock returns on the one side and the term spread and the real short term rate on the other. The dividend yield also exerts a smoothing effect on multi-period REITs returns. However, this is counterbalanced by the impact of stock and bond returns, inflation and the real short-term rate, whose shocks are positively correlated with real estate shocks. This explains why the demand for public real estate does not increase as sharply as that for stocks as a function of the investment horizon.

Interestingly, a classical investor ignoring real estate would reduce the share of bonds by 33% anyway, while increasing the demand of stocks and cash by 9% and 24% respectively. It thus appears that the predictability patterns induced by real estate leave almost unaltered the percentage decrease in cash, while exacerbating the fall in bond holdings.

Diversification and Parameter Uncertainty

Estimation risk influences portfolio choice in two ways. Important modifications occur in the structure of the investment schedules as a function of the horizon. When parameter uncertainty is taken into account, the average holdings of real estate and equity respectively become decreasing and flatter in T , resulting in lower investment in riskier assets. A classical investor, however, chooses average weights for riskier assets

increasing in the investment horizon. For instance, the allocation to real estate decreases from 43% at a 1-month horizon to 33% at a 5-year horizon, while the allocation to stocks marginally increases from 11% to 12%. This effect on the average portfolio share of the riskier assets is well explained by the fact that the uncertainty deriving from estimation risk compounds over time, implying that the difficulty to predict is magnified over longer planning periods. It follows that the opposite effects of a reduction in long-run risk resulting from predictability – which would cause the investment schedules to be upward sloping – and of estimation risk roughly cancel out for a long-horizon investor, with the result of either flat or weakly monotonically decreasing schedules.

It is also clear that when an investor accounts for estimation risk, she develops a strong incentive to invest in short-term deposits, especially in the long-run, even if in our framework the real short term rate becomes risky for horizons exceeding one month. The average portfolio share in cash, which already doubles for a one-month horizon when estimation risk is accounted for, increases steadily in T reaching 49% for a 5-year horizon. A strategy that rolls over “cash” investments is not only the safest among the available assets in terms of overall variance, but also the one that remains predictable with high precision from its own lagged value, the lagged return on bonds and real state, the dividend yield, the term premium, the default spread as well as the inflation rate. Furthermore, it appears that shocks to the real short-term rate are negatively correlated with shocks to the term spread, with a large coefficient in absolute value (-0.44). At the same time, a higher value of the lagged term spread predicts a higher future short-term rate (1.17). This embeds some mean reversion in the return to short term deposits, making them comparatively safer as the investment horizon grows. Therefore short term deposits *de facto* preserve their role of safe assets in relative terms, even though their stochastic nature is fully recognized by our econometric set up.

Finally, remember that we had noticed in the previous section that the predictability of bond returns almost disappeared when we had allowed for parameter uncertainty. It is therefore not surprising that the optimal share in 10-year bonds drops from 25% to 2% as the horizon grows, and REITs -- with their superior Sharpe ratio and more precisely estimated predictive relationships -- are included in the asset menu.

Surprisingly, these patterns are not too dissimilar to those uncovered on European data, in simulated buy-

and-hold optimal portfolio allocations by Fugazza, Guidolin, and Nicodano (2007). The reduction over T associated with parameter uncertainty of the optimal shares invested in the most risky among the assets turns out to be optimal for a European investor, too. However, in Fugazza et al.'s results, the schedule for stocks is flatter than the one for securitized real estate. The schedules for cash and for bonds also have a similar pattern in optimal US and European portfolios, although the relevance of bonds in longer run portfolios remains higher for a European investor.

Ex-Post Performance Results

In the literature, econometric models of predictable assets returns produce good in-sample fits and optimal portfolios built on estimated mean returns and volatilities usually display good in-sample performance. Moreover, expanding the asset menu cannot reduce the *ex-ante* investor's welfare, as it is always possible to exclude – by assigning an optimizing zero weight – the additional asset(s) from the optimal portfolio. However, this is by no means a guarantee that expanding the asset menu will lead to improved *ex-post* portfolio performance. This problem may arise, for example, when the proposed model for asset returns is misspecified and/or there are large parameter estimation errors. To address this concern, we next explore how well real estate as an asset class performs when it is added to portfolios formed from 1995 onwards, on the basis of estimates obtained only on the basis of the available information on a recursive basis. We distinguish between ex-post results as measured by standard reward-to-risk ratio indicators (e.g., the classical Sharpe ratio) and performance in terms of realized, average utility (welfare) of the investor.

Realized Reward-to-Risk Ratios

Table 5 considers average performances over the period 1995-2004 when equity REITs are alternatively included or excluded from the asset menu. It reveals that the mean Sharpe ratio achieved by both a Classical and a Bayesian investor increases – irrespective of the investor's horizon – when real estate is added to the asset menu. Such a change is substantial, ranging from 0.12 to 0.42, depending on the horizon. This is caused by the fact that realized portfolio volatility falls (except for the horizon $T = 60$) as a result of declining investments in stocks. Such reduction in risk taking appears not to have damaged profitability in the out-of-sample period, as mean portfolio returns increase.

We also proceed to test hypotheses concerning the portfolio Sharpe ratios following Jobson and Korkie (1981). In our application, we are facing a pair of portfolio strategies, i.e., with and without real estate assets in the choice menu. We call $\mathbf{r}_{i,T}$ ($i=1, \dots, \tau$) the 2×1 vector that collects the excess returns on both strategies, when the selected investment horizon is T and τ is the total length of the time series of portfolio performances; $\mathbf{r}_{i,T}$ is assumed to be bivariate normal with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, with unbiased samples estimators given by the sample mean vector $\bar{\mathbf{r}}_T$ and sample covariance matrix \mathbf{S}_T . Jobson and Korkie show that in finite samples, computing the Sharpe ratio as a simple ratio between the sample mean excess return and sample variance ($SR_{i,T} = \bar{r}_{i,T} / s_{i,T}$, $i=1,2$) leads to a biased estimation of the Sharpe ratio; for instance, using a Taylor approximation to the order $1/\tau^2$ (i.e., of order $O(1/\tau^2)$), they show that

$$E[SR_{i,T}] \simeq \frac{\mu_i}{\sigma_i} \left(1 + \frac{3}{4\tau} + \frac{100}{128\tau^2} \right) > \frac{\mu_i}{\sigma_i},$$

i.e., that there is an over-estimation bias. As a result, $SR_{1,T} - SR_{2,T}$ is also affected by the bias (while $SR_{1,T} / SR_{2,T}$ is not).¹⁵ To test the null hypothesis that $H_o: SR_{1,T} - SR_{2,T} = 0$, Jobson and Korkie suggest instead to test the transformed null hypothesis $H'_o: \Delta T SR_T \equiv \bar{r}_{1,T} s_{2,T} - \bar{r}_{2,T} s_{1,T} = 0$ and show that:

$$z_{SR} = \frac{\Delta T SR_T}{\sqrt{\widehat{\text{Var}}[\Delta T SR_T]}} = \frac{\sqrt{\tau} (\bar{r}_{1,T} s_{2,T} - \bar{r}_{2,T} s_{1,T})}{\sqrt{2s_{1,T}^2 s_{2,T}^2 - 2s_{1,T} s_{2,T} s_{12,T} + \frac{1}{2} \bar{r}_{1,T}^2 s_{2,T}^2 + \frac{1}{2} \bar{r}_{2,T}^2 s_{1,T}^2 - \frac{\bar{r}_{1,T} \bar{r}_{2,T}}{2s_{1,T} s_{2,T}} (s_{12,T}^2 + s_{1,T}^2 s_{2,T}^2)}} \stackrel{asy}{\sim} N(0, 1).$$

where one can estimate both numerators and denominators by replacing $\mu_{i,T}$, $\sigma_{i,T}$ ($i=1,2$), and $\sigma_{12,T}$ with their sample estimates $\bar{r}_{i,T}$, $s_{i,T}$, and $s_{12,T}$.

We proceed first to compute 95% (asymptotic) confidence intervals for the bias-corrected Sharpe ratios,

$$SR_{i,T}^{bc} = \frac{\bar{r}_{i,T} / s_{i,T}}{\left(1 + \frac{3}{4\tau} + \frac{100}{128\tau^2} \right)}$$

from portfolio strategies with and without real estate, both in the classical and in the Bayesian framework.

To save space, we only report results for $T = 1, 3, 6, 24,$ and 60 months. In the classical case, they are as follows:

$T = 1$		$T = 3$		$T = 6$		$T = 12$		$T = 60$	
RE	No RE	RE	No RE	RE	No RE	RE	No RE	RE	No RE
0.443	0.186	0.384	0.128	0.359	0.124	0.306	-0.158	0.393	0.090
[0.25,0.63]	[0.00,0.37]	[0.20,0.57]	[-0.06,0.31]	[0.18,0.54]	[-0.07,0.31]	[0.12,0.49]	[-0.36,0.04]	[0.12,0.66]	[-0.13,0.31]

A comparison with Table 5 reveals that correcting for small sample bias hardly changes the point estimates of the ex-post Sharpe ratios. Even though the confidence intervals (CIs) are rather wide, one can immediately spot a case (for $T = 12$) in which the CIs for the two strategies are disjoint, i.e., the interval for the case with real estate lies completely to the right of the confidence interval for the strategy without real estate. In other four cases ($T = 1, 3, 6,$ and 60 months), the CI with real estate has a lower bound that exceeds the point estimate of the Sharpe ratio without real estate, which is also a powerful indication that – even taking sample uncertainty into account – including real estate in the asset menu strongly increases ex-post performance. In the case of the Bayesian portfolio strategies, we obtain:

$T = 1$		$T = 3$		$T = 6$		$T = 12$		$T = 60$	
RE	No RE								
0.448	0.031	0.320	0.041	0.262	0.005	0.222	0.002	0.306	0.078
[0.25,0.64]	[-0.15,0.21]	[0.13,0.51]	[-0.14,0.23]	[0.08,0.45]	[-0.18,0.19]	[0.03,0.41]	[-0.19,0.20]	[0.00,0.61]	[-0.16,0.31]

Results are similar in the Bayesian case: once more there is one case ($T = 1$) in which the CIs are disjoint, and three cases ($T = 3, 6, 12$) in which the CI with real estate has a lower bound that exceeds the point estimate of the Sharpe ratio without real estate.¹⁶

However, these informal comparisons of CIs cannot substitute for a formal test of the hypothesis that the Sharpe ratios are identical across strategies with and without real estate. We therefore test $H'_o: \Delta TSR_T \equiv \bar{r}_{1,T} s_{2,T} - \bar{r}_{2,T} s_{1,T} = 0$ and obtain the following Jobson and Korkie's test statistics (the corresponding p-values are parenthesis):¹⁷

	$T = 1$	$T = 3$	$T = 6$	$T = 12$	$T = 24$	$T = 60$
Classical	-1.887	-1.388	-1.301	-0.661	-2.793	-2.259
	(0.059)	(0.165)	(0.193)	(0.509)	(0.005)	(0.024)
Bayesian	-2.410	-3.465	-4.000	-2.271	-3.941	-4.220
	(0.016)	(0.001)	(0.000)	(0.023)	(0.000)	(0.000)

In this case there is a difference between classical and Bayesian results: while in the former case, the z_{SR} statistic has a U-shape pattern and the null of identical Sharpe ratios can only be rejected (at a standard 5% size) at $T = 24$ and 60, in the case of recursive Bayesian portfolios, we reject the null of identical Sharpe ratios for all investment horizons, although also in this case z_{SR} reaches its highest values for long investment horizons. All in all, this is rather compelling evidence that including real estate in the asset menu does lead to higher ex-post, realized values of the Sharpe ratios, even when sample uncertainty is taken into account.

As recently discussed in the literature (see e.g., Gallo, Lockwood, and Rodriguez, 2006) is that Sharpe ratios are highly sensitive to non-normally distributed returns. We find strong evidence of non-normal returns in our sample. For instance, with a one-month investment horizon, in the classical framework the skewness of realized portfolio returns is -0.39 (-0.62 without real estate) while in the Bayesian asset allocation scheme skewness is -1.51 (-0.90 without real estate). The corresponding estimates for the excess kurtosis of one-month realized portfolio returns are 1.77 (1.42 without real estate) and 5.37 (1.64 without real estate). One common remedy consists of supplementing the presentation of Sharpe ratios with related reward-to-risk measures that divide the numerator (mean excess return) of the Sharpe ratio by portfolio downside risk (semi-standard deviation), what is commonly called the Sortino ratio (also see Fishburn, 1977).¹⁸ We therefore present the Sortino measures in a format comparable to Table 5:

	$T = 1$		$T = 3$		$T = 6$		$T = 24$		$T = 60$	
	No RE	Δ	No RE	Δ	No RE	Δ	No RE	Δ	No RE	Δ
Classical	0.025	0.069	0.047	0.105	0.083	0.112	0.047	0.207	0.239	0.001
Bayesian	0.015	0.094	0.035	0.136	0.043	0.173	0.097	0.198	0.294	0.260

To save space, we have only reported the Sortino measures when real estate is *not* included in the asset

menu and the corresponding increase (Δ) when real estate is additionally included. Although the values of the reward-to-risk ratio are now different, the basic implication that portfolio strategies that include real estate yield an ex-post performance that is superior to strategies that exclude real estate continues to hold even when possible difference between overall and downside variance are taken into account.

Realized Utility Comparisons

We still have to assess the effects of real estate on the expected utility of an investor. Indeed, an increase in Sharpe ratio is not necessarily associated with higher welfare, if it is obtained at the cost of worse higher-order moment properties of portfolio returns. This is because investors are typically averse to negative skewness and excess kurtosis, and these preferences are fully captured by a power utility function. For instance (and differently from a simpler, mean-variance optimizer), a power utility investor may perceive a reduction in ex-post realized utility when a portfolio exposes her to higher negative skewness of realized wealth (i.e., higher probability of large deviations below the mean) even though the portfolio Sharpe ratio improves. In fact, comparing ex-post realized utility across different portfolio strategies is the only robust remedy to the presence of non-normalities in realized portfolio returns, as this performance measure perfectly aligns the ex-ante preferences of investors when they select their optimal portfolios to the way in which these portfolios are evaluated and compared ex-post.

We obtain estimates of the welfare cost of restricting the span of the asset menu available to our investors. Call $V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t)$ the realized utility of the unconstrained problem – i.e., when real estate belongs to the asset menu – and $V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t^R)$ the constrained realized utility, where $\hat{\boldsymbol{\omega}}_t^R$ is the vector of portfolio weights obtained when real estate investments are ruled out. Note that in ex-post, out-of-sample experiments, constrained realized utility may exceed the unconstrained one, while ex ante this is impossible. The compensatory premium, π_t^R , is the percentage of wealth that, when added to the investor's initial wealth, equates the realized utility from the constrained and unconstrained problems:

$$\pi_t^R = \left[\frac{V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t)}{V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t^R)} \right]^{\frac{1}{1-\gamma}} - 1. \quad (4)$$

The compensatory premium is therefore a measure of the ex-post welfare gain from enlarging the asset menu to real estate.

Table 5 reports our estimates of the average premium. Such estimates confirm earlier insights based on Sharpe ratios, in that they are always positive and far from negligible. An investor would be willing to pay an annual fee in excess of 1.39% of her initial wealth in order to improve her portfolio performance by accessing real estate indirect vehicles such as REITs. Welfare gains are larger, apart from the $T = 1$ month horizon, for a Classical investor who overlooks estimation risk when choosing portfolio composition. Such an investor would have been willing to pay a minimum annual fee of 1.63% of her initial wealth in order to diversify into real estate over our recursive period 1995-2004. This fee is larger for horizons of one year or longer, suggesting that real estate especially adds value to investment strategies for long-horizon Classical investors. Welfare gains are smaller, translating in fees between 1.39 and 2.59 percent of initial wealth, for a Bayesian investor who, by accounting for estimation risk, is substantially more cautious than a Classical one. Our results extend to an ex-post setting previous evidence on in-sample (ex-ante) performance gains from real estate, see e.g. Seiler, Webb and Myer (1999) and Feldman (2003).

Looking at the Certainty Equivalent return (CER) rows in Table 5 (for the classical case) shows that while asset menus that include real estate yield annualized CERs between 6.3 and 8.7 percent; when real estate is excluded from the asset menu, the CERs are between 4.1 and 6.8 percent.. The CER values are generally lower in the Bayesian case, ranging between 6.2 and 8.8 percent when real estate is included, and between 4.8 and 6.2 percent when real estate is excluded. One surprising feature of our results is that a Bayesian investor achieves lower ex-post expected utility than a Classical one. In Jorion (1985), a pioneering study of mean-variance portfolio choice, the opposite occurs. Furthermore, an opposite result has also been reported in one-period portfolio choice models that allow for predictable returns (Abhyankar, Basu, and Stremme, 2005; Avramov and Chordia, 2006; Wachter and Warusawitharana, 2005). This is all the more surprising because in the present paper we have allowed all predictors in our VAR to predict future returns, even the

ones that do not display statistical significance. This means that we have not performed any typical specification search of the predictors that best forecast excess returns on each type of asset class. By also considering imprecise predictors we are likely to magnify the difference between the classical and the Bayesian asset allocation, as the latter "discounts" the weight to be assigned to asset classes the more, the less precise is the prediction of future returns. As a consequence, we expected the Bayesian investors, that optimizes forecasting accuracy to do much better than a Classical investors, that relies too much on imprecisely estimated parameters. Despite this, we find that the ex post performance of Bayesian portfolios is worse than the performance of classical portfolios.¹⁹

Value of Time Diversification

Return predictability may affect asset allocation decisions in two ways. First, one-period-ahead expected returns and volatilities change with the arrival of new information, generating the possibility of market timing. Gains from market timing can be assessed by comparing performance and welfare of short term investors, when they alternatively use or ignore return predictability. A large literature, mentioned in the introduction, evaluates these gains. In our sample, the (unreported) out-of-sample value of market timing is negative, both for a classical investor and a Bayesian one- whether or not real estate is included.²⁰

The second consequence of predictability can be appreciated in multi-period portfolio problems only, as it entails a modification in the conditional variances and covariances of multi-period returns. It may thus happen that an asset class becomes riskier over a long horizon, because its returns are mean averting. Thus, longer-term investors may be able to enhance their risk-adjusted portfolio performance and welfare by taking into account these conditional multi-period volatilities in their portfolio optimizers.

We now turn to the assessment of ex post gains brought about by longer horizons, that allow the exploitation of intertemporal diversification. It is well known that the annualized volatility for a T -period investor is lower than T times the one period volatility when mean reversion in returns allows for risk diversification over time. Our experiment reveals that ex-post gains from risk diversification over time are not present, whether or not real estate is included. Mean annualized volatility of portfolio returns increases

from 8.5% for a $T = 1$ month horizon to 11.6% for a $T = 60$ horizon, when a Classical investor uses an asset menu that accounts for real estate and from 7.0% to 19.0% when real estate is excluded. It is also the case that annual mean portfolio returns also fall from 10.8% (8.2% without REITs) to 9.2% (7.1%). Similar results obtain for horizons shorter than 60 months. The average certainty equivalent of welfare drops from 8.7% (6.8% without REITs) to 7.6% (4.2% without REITs). This points out that exploiting linear predictability for the construction of expected utility maximizing portfolios has no value for $T > 1$. The very same pattern emerges for the less aggressive Bayesian investor, who accounts for parameter uncertainty.

A vast literature has already questioned the possibility to exploit linear predictability for improving on ex post investment performance for short-term portfolios (see Cooper, Gutierrez, and Marcum., 2005 and references therein), as well as when real estate is considered (see Ling, Naranjo, and Ryngaert, 2000). Our results cast doubt on the possibility of exploiting the risk return trade-off for long horizon investors in the way suggested by Campbell and Viceira (2005). Clearly, this does not imply that gains would not exist over other sample periods. It also leaves open the possibility that alternative prediction models may have yielded significant gains over our sample period. Finally, it remains possible that non-linear (e.g., of the Markov switching type models, as in Guidolin and Timmermann, 2005) econometric frameworks may lead to more precisely estimated and economically valuable predictability patterns, also in the presence of public real estate vehicles.

Discussion and Conclusions

Irrespective of the investment horizon and of the estimation method, we find positive ex-post welfare effects deriving from the inclusion of equity REITs in the asset menu. These gains are larger for a Classical investor who, by overlooking estimation risk when choosing her optimal portfolio composition, is less cautious than a Bayesian one. These results extend to an *ex-post* setting previous evidence on in-sample, *ex-ante* performance gains from real estate. Our second experiment, devoted to the assessment of ex-post gains from risk diversification over time, reveals that such gains are not present, as the average annualized

volatility of portfolio returns increases in the investor horizon, while it should fall when mean reversion allows for intertemporal risk diversification. It is also the case that annual mean portfolio returns fall as the investment horizon increases. This points out that exploiting *linear* predictability for the construction of expected utility maximizing portfolios has no value for longer horizon investors, whether or not real estate is included. Our results thus cast doubt on the possibility to exploit the risk return trade-off for long horizon investors in the way suggested by Campbell and Viceira (2005). Clearly, this does not imply that gains from time diversification would not exist over other sample periods. More importantly, it leaves open the possibility that alternative prediction models may yield significant gains, even over our sample period, as suggested by Lettau and Van Nieuwerburgh (2006).

Our results on time diversification – and hence on the impact of predictability on ex post, out of sample welfare – should be considered as a lower bound on its usefulness for another reason. In our VAR analysis, we allow all predictors to predict future returns, even the ones that do not display statistical significance. Recent work on one-period portfolio choice suggests that out-of-sample performance can be improved by carefully selecting the predictors (Pesaran and Timmerman, 1995; Bossaerts and Hillion, 1999) and by relying on CAPM-based prediction rather than on a linear vector-autoregression framework (Avramov Chordia, 2006; Handa and Tiwari, 2006). Similarly, in performing our Bayesian estimates we impose an uninformative prior, implying that the investor neither believes nor doubts *à priori* in the strength of predictability. Recent work on one-period portfolio choice suggests that better out-of-sample performance might actually be obtained by imposing ex-ante skepticism by the investor (Wachter and Warusawitharana, 2005). Future research might check whether this is also true in a multi-period portfolio model.

Finally, our paper has used NAREIT equity REIT index returns to proxy for the U.S. public real estate asset class. However, several investors with a potential interest in our results may also access real estate investments using commingled real estate funds (CREFs), which are essentially composed of unlevered properties held on behalf of investors in the portfolios of fiduciary firms (see e.g., Gallo, Lockwood, and Rodriguez, 2006). It would be possible and interesting to extend our analysis of optimal portfolio choices under predictable returns to include return indices proxying for private market CREF returns in addition to

the NAREIT series used in this paper.²¹

References

- Abhyankar, A., D. Basu and A. Stremme. 2005. The Optimal Use of Return Predictability. An Empirical Analysis, mimeo, Warwick Business School.
- Avramov, D. 2002. Stock Return Predictability and Model Uncertainty. *Journal of Financial Economics* 64: 423-458.
- Avramov, D. and T. Chordia. 2006. Predicting Stock Returns. *Journal of Financial Economics* 82: 387-415.
- Barberis, N. 2000. Investing for the Long Run When Returns Are Predictable. *Journal of Finance* 55: 225-64.
- Bharati, R. and M. Gupta. 1992. Asset Allocation and Predictability of Real Estate Returns. *Journal of Real Estate Research* 7: 469-484.
- Bossaerts, P. and P. Hillion. 1999. Implementing Statistical Criteria to Select Return Forecasting Models: What Do We Learn? *Review of Financial Studies* 12: 405-428.
- Brandt, M. 1999. Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach. *Journal of Finance* 54:1609-1645.
- Campbell, J., Y. Chan and L. Viceira. 2003. A Multivariate Model of Strategic Asset Allocation. *Journal of Financial Economics* 67:41-80.
- Campbell, J. and R. Shiller. 1988. Stock Prices, Earnings, and Expected Dividends. *Journal of Finance* 43: 661-676.
- Campbell, J. and L. Viceira. 1999. Consumption and Portfolio Decisions When Expected Returns Are Time Varying. *Quarterly Journal of Economics* 114: 433-495.
- Campbell, J. and L. Viceira. 2001. Who Should Buy Long-Term Bonds? *American Economic Review* 91: 99-127.
- Campbell, J. and L. Viceira 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.
- Campbell J. and L. Viceira. 2005. The Term Structure of the Risk-Return Tradeoff. *Financial Analysts Journal* 61: 34-44.
- Cho, H., Y. Kawaguchi and J. Shilling. 2003. Unsmoothing Commercial Property Returns: A Revision to Fisher-Geltner-Webb's Unsmoothing Methodology. *Journal of Real Estate Finance and Economics* 27: 393-405.
- Cooper, M., R. Gutierrez and B. Marcum. 2005. On the Predictability of Stock Returns in Real Time. *Journal of Business* 78: 469-499.
- Fama, E. and K. French. 1989. Business Conditions and Expected Returns on Stocks and Bonds. *Journal of Financial Economics* 25: 23-49.

- Feldman, B. 2003. Investment Policy for Securitized and Direct Real Estate. *Journal of Portfolio Management*, Special Real Estate Issue, 112-121.
- Fishburn, P. 1977. Mean-Risk Analysis with Risk Associated with Below Market Returns. *American Economic Review* 67: 116-126.
- Fisher, J., D. Geltner and R. Webb. 1994. Value Indices of Commercial Real Estate: A Comparison of Index Construction Methods. *Journal of Real Estate Finance and Economics* 9: 137-164.
- Fugazza, C., M. Guidolin and G. Nicodano. 2007. Investing for the Long-Run in European Real Estate. *Journal of Real Estate Finance and Economics* 34: 5-80.
- Gallo, J., L. Lockwood and M. Rodriguez. 2006. Differentiating CREF Performance. *Real Estate Economics* 34: 173-210.
- Geltner, D. 1993. Estimating Market Values from Appraised Values Without Assuming an Efficient Market. *Journal of Real Estate Research* 8: 325-346.
- Geltner, D. and J. Mei. 1995. The Present Value Model with Time-Varying Discount Rates: Implication for Commercial Property Valuation and Investment Decisions. *Journal of Real Estate Finance and Economics* 11: 119-135.
- Georgiev, G., Gupta B. and T. Kunkel. 2003. Benefits of Real Estate Investment. *Journal of Portfolio Management*, Special Real Estate Issue, 28-34.
- Glascok, J., C. Lu and R. So. 2001. Further Evidence on the Integration of REIT, Bond, and Stock Returns. *Journal of Real Estate Finance and Economics* 20: 177-194.
- Goetzmann, W., N. Ingersoll, J. Spiegel and I. Welch. 2002. Sharpening Sharpe Ratios. NBER Working Paper No. 9116.
- Goetzmann, W., N. Ingersoll, J. Spiegel and I. Welch. 2004. Portfolio Performance Manipulation and Manipulation – Proof Performance Measures. Yale ICF Working Paper No. 02-08.
- Goyal, A. and I. Welch. 2004. A Comprehensive Look at the Empirical Performance of the Equity Premium Prediction. Yale ICF Working Paper, No. 04-11.
- Guidolin, M. and A. Timmermann. 2005. Economic Implications of Bull and Bear Regimes in UK Stock and Bond Returns. *Economic Journal* 115: 111-143.
- Handa, P. and A. Tiwari. 2006. Does Stock Return Predictability Imply Improved Asset Allocation Performance? Evidence from the U.S. Stock Market (1954-2002). *Journal of Business* 79: 2423-2468.
- Hoevenaars R., R. Molenaar, P. Schotman and T. Steenkamp. 2007. Strategic Asset Allocation with Liabilities: Beyond Stock and Bonds. Working paper, Maastricht University.
- Jorion, P. 1985. International Portfolio Diversification with Estimation Risk. *Journal of Business* 58: 259-277.
- Kandel, S. and R. Stambaugh. 1996. On the Predictability of Stock Returns: An Asset-Allocation Perspective. *Journal of Finance* 51: 385-424.
- Karolyi, G. and A. Sanders. 1998. The Variation of Economic Risk Premiums in Real Estate Returns. *Journal of Real Estate Financial Economics* 17: 245-262.
- Leland, H. 1999. Beyond Mean – Variance Performance Measurement in a Non-symmetrical World.

Financial Analysts Journal 55: 27-36.

- Lettau, M. and S. Van Nieuwerburgh. 2006. Reconciling the Return Predictability Evidence. *Review of Financial Studies*, forthcoming.
- Liang, Y., N. Myer and J. Webb. 1996. The Bootstrap Efficient Frontier for Mixed-Asset Portfolios. *Real Estate Economics* 24: 247-256.
- Ling, D., A. Naranjo and M. Ryngaert. 2000. The Predictability of Equity REIT Returns: Time Variation and Economic Significance. *Journal of Real Estate Finance and Economics* 20: 117-136.
- Liu, C. and J. Mei. 1992. The Predictability of Returns on Equity REITs and Their Co-movement with Other Assets. *Journal of Real Estate Financial Economics* 5: 401- 418.
- Merton, R. 1969. Lifetime Portfolio Selection under Uncertainty: the Continuous Time Case. *Review of Economics and Statistics* 51: 247-257.
- Nelling, E. and J. Gyourko. 1998. The Predictability of Equity REIT Returns. *Journal of Real Estate Research* 16:1-18.
- Pesaran, M. and A. Timmermann. 1995. Predictability of Stock Returns: Robustness and Economic Significance. *Journal of Finance* 50: 1201-1228.
- Samuelson, P. 1969. Lifetime Portfolio Selection by Dynamic Stochastic Programming. *Review of Economics and Statistics* 51: 239-246.
- Seiler M., J. Webb and F. Myer. 1999. Diversification Issues in Real Estate Investment. *Journal of Real Estate Literature* 5: 163-179.
- Wachter, J. A. and M. Warusawitharana. 2005. Predictable returns and asset allocation: should a skeptical investor time the market? mimeo, University of Pennsylvania.

Appendix: Long Run Asset Allocation when Returns Are Predictable

In this section we review the structure and solution methods for a portfolio choice problem when returns are predictable and the uncertainty about the extent of predictability is taken into account. The methodology follows Kandel and Stambaugh (1996) and Barberis (2000) and so we only briefly discuss the main issues and technical details.

Long run portfolio strategies may be calculated under two alternative assumptions: buy-and-hold vs. optimal rebalancing. An investor who follows a buy-and-hold strategy chooses the optimal allocation at the beginning of the planning horizon (t) and does not modify it until the end-point ($t + T$) is reached. Clearly, when T is large, this represents a strong commitment not to revise the portfolio weights despite the receipt of news characterizing the investment opportunity set. Under a rebalancing strategy, the investor chooses the asset allocation at the beginning of the planning horizon taking into account that it shall be optimal to modify the portfolio weights at intermediate dates (rebalancing points), $t + \varphi$, $t + 2\varphi$, ..., $t + T - \varphi$. In the following we focus only on buy-and-hold portfolio strategies (see Fugazza, Guidolin, Nicodano, 2007, for

details on rebalancing strategies).

Classical Buy-and-Hold Investor

Call θ the vector collecting all the parameters entering (2), i.e. $\theta \equiv [\boldsymbol{\mu}' \text{vec}(\boldsymbol{\Phi})' \text{vech}(\boldsymbol{\Sigma})']'$. Under (2), the (conditional) distribution of cumulative future returns (i.e. the first four elements in $\mathbf{z}_{t,T} \equiv \sum_{k=1}^T z_{t+k}$) is multivariate normal with mean and covariance matrix given by the appropriately selected elements of:

$$\begin{aligned} E_{t-1}[\mathbf{z}_{t,T}] &= T\boldsymbol{\mu} + (T-1)\boldsymbol{\Phi}\boldsymbol{\mu} + (T-2)\boldsymbol{\Phi}^2\boldsymbol{\mu} + \dots + \boldsymbol{\Phi}^{T-1}\boldsymbol{\mu} + (\boldsymbol{\Phi} + \boldsymbol{\Phi}^2 + \dots + \boldsymbol{\Phi}^T)\mathbf{z}_{t-1} \\ \text{Var}_{t-1}[\mathbf{z}_{t,T}] &= \boldsymbol{\Sigma} + (\mathbf{I} + \boldsymbol{\Phi})\boldsymbol{\Sigma}(\mathbf{I} + \boldsymbol{\Phi})' + (\mathbf{I} + \boldsymbol{\Phi} + \boldsymbol{\Phi}^2)\boldsymbol{\Sigma}(\mathbf{I} + \boldsymbol{\Phi} + \boldsymbol{\Phi}^2)' + \\ &\quad \dots + (\mathbf{I} + \boldsymbol{\Phi} + \dots + \boldsymbol{\Phi}^{T-1})\boldsymbol{\Sigma}(\mathbf{I} + \boldsymbol{\Phi} + \dots + \boldsymbol{\Phi}^{T-1})', \end{aligned} \quad (5)$$

where \mathbf{I} is the identity matrix of dimension n and $\boldsymbol{\Phi}^k \equiv \prod_{i=1}^k \boldsymbol{\Phi}$. Since the parametric form of the predictive distribution of $\mathbf{z}_{t,T}$ is known, it is simple to approach the problem in (1), or equivalently

$$\max_{\boldsymbol{\omega}_t} \int \frac{W_{t+T}^{1-\gamma}}{1-\gamma} \phi(E_t[\mathbf{z}_{t,T}], \text{Var}_t[\mathbf{z}_{t,T}]) \cdot d\mathbf{z}_{t,T} \quad (6)$$

where $(\phi(E_t[\mathbf{z}_{t,T}], \text{Var}_t[\mathbf{z}_{t,T}]))$ is a multivariate normal with mean $E_t[\mathbf{z}_{t,T}]$ and covariance matrix $\text{Var}_t[\mathbf{z}_{t,T}]$, by simulation methods. Similarly to Kandel and Stambaugh (1996), Barberis (2000), and Guidolin and Timmermann (2005), this means evaluating the integral in (6) by drawing a large number of times (N) from $(\phi(E_t[\mathbf{z}_{t,T}], \text{Var}_t[\mathbf{z}_{t,T}]))$ and then maximizing the following functional:

$$\max_{\boldsymbol{\omega}_t} \frac{1}{N} \sum_{i=1}^N \left[\frac{\{\omega_t^s \exp(R_{t,T}^{s,i}) + \omega_t^b \exp(R_{t,T}^{b,i}) + \omega_t^r \exp(R_{t,T}^{r,i}) + (1 - \omega_t^s - \omega_t^b - \omega_t^r) \exp(R_{t,T}^{f,i})\}^{1-\gamma}}{1-\gamma} \right], \quad (7)$$

where $[R_{t,T}^{s,i} \ R_{t,T}^{b,i} \ R_{t,T}^{r,i} \ R_{t,T}^{f,i}]'$ represent the first four elements of $\mathbf{z}_{t,T}^i$ along a sample path $i = 1, \dots, N$. At this stage, the portfolio weight non-negativity constraints are imposed by maximizing (7) using a simple two-stage grid search algorithm that sets ω_t^j to 0, 0.01, 0.02, ..., 0.99, 0.9999 for $j = s, b, r$.

Bayesian Buy-and-Hold Investor

Given the problem

$$\max_{\boldsymbol{\omega}_t} \int \frac{W_{t+T}^{1-\gamma}}{1-\gamma} p(\mathbf{z}_{t,T} | \ddot{\mathbf{Z}}_t, \boldsymbol{\theta}) p(\boldsymbol{\theta} | \ddot{\mathbf{Z}}_t) \cdot d\mathbf{z}_{t,T},$$

the task is somewhat simplified by the fact that predictive draws can be obtained by drawing from the posterior distribution of the parameters and then, for each set of parameters drawn, by sampling one point from the distribution of returns conditional on past data and the parameters. At this point, (2) can be re-written as:

$$\begin{bmatrix} \mathbf{z}'_2 \\ \mathbf{z}'_3 \\ \vdots \\ \mathbf{z}'_t \end{bmatrix} = \begin{bmatrix} 1 & \mathbf{z}'_1 \\ 1 & \mathbf{z}'_2 \\ \vdots & \vdots \\ 1 & \mathbf{z}'_{t-1} \end{bmatrix} \begin{bmatrix} \boldsymbol{\mu}' \\ \boldsymbol{\Phi}' \end{bmatrix} + \begin{bmatrix} \boldsymbol{\varepsilon}'_2 \\ \boldsymbol{\varepsilon}'_3 \\ \vdots \\ \boldsymbol{\varepsilon}'_t \end{bmatrix},$$

or simply $\mathbf{Z} = \mathbf{X}\mathbf{C} + \mathbf{E}$, where \mathbf{Z} is a $(t-1, n+1)$ matrix with the observed vectors as rows, \mathbf{X} is a $(t-1, n+2)$ matrix of regressors, and \mathbf{E} a $(t-1, n+1)$ matrix of error terms, respectively. All the coefficients are instead collected in the $(n+2, n+1)$ matrix \mathbf{C} . If we consider the following standard uninformative diffuse prior:

$$p(\mathbf{C}, \boldsymbol{\Sigma}) \propto |\boldsymbol{\Sigma}|^{-\frac{n+2}{2}},$$

then the posterior distribution for the coefficients in θ , $p(\mathbf{C}, \boldsymbol{\Sigma}^{-1} | \ddot{\mathbf{Z}}_t)$ can be characterized as:

$$\boldsymbol{\Sigma}^{-1} | \ddot{\mathbf{Z}}_t \sim \text{Wishart}(t-n-2, \hat{\mathbf{S}}^{-1})$$

$$\text{vec}(\mathbf{C}) | \boldsymbol{\Sigma}^{-1}, \ddot{\mathbf{Z}}_t \sim N(\text{vec}(\hat{\mathbf{C}}), \boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X})^{-1})$$

where $\hat{\mathbf{S}} = (\mathbf{Z} - \mathbf{X}\hat{\mathbf{C}})'(\mathbf{Z} - \mathbf{X}\hat{\mathbf{C}})$ and $\hat{\mathbf{C}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Z}$, i.e. the classical OLS estimators for the coefficients and covariance matrix of the residuals.

We adopt the following simulation method. First, we draw N independent variates from $p(\mathbf{C}, \boldsymbol{\Sigma}^{-1} | \ddot{\mathbf{Z}}_t)$. This is done by first sampling from a marginal Wishart for $\boldsymbol{\Sigma}^{-1}$ and then (after calculating $\boldsymbol{\Sigma}$) from the conditional $N(\text{vec}(\hat{\mathbf{C}}), \boldsymbol{\Sigma} \otimes (\mathbf{X}'\mathbf{X})^{-1})$, where $\hat{\mathbf{C}}$ is easily calculated. Second, for each set $(\mathbf{C}, \boldsymbol{\Sigma})$ obtained, the algorithm samples cumulated returns from a multivariate normal with mean vector and covariance matrix given by (5). Given the double simulation scheme, in this case N is set to a relatively large value of 300,000 independent trials.

Table 1 – Descriptive Statistics

The upper part of both Panel A and B report descriptive statistics for monthly excess returns on stocks, bonds, E-REITs and returns on cash investments (real and nominal), as well as for predictor variables (dividend yield, inflation, default spread and term spread). The lower part of both panels report, for the various samples, contemporaneous correlations. The sample period considered is 1/1972 – 12/ 2004 in Panel A, and 1/1972 – 12/1994 in Panel B. Data on stocks and real estate are from CRSP (NYSE, NASDAQ, and AMEX value-weighted indices) and NaREIT respectively. We use FREDII[®] series for computing the returns on 10-year constant-maturity Treasury bonds and on 3-month Treasury bills. Inflation is derived from the Consumer Price Index For All Urban Consumers (all Items, seasonally adjusted).

Panel A: Sample 1972-2004

	Stocks	Bonds	Real Estate	Nominal Short Rate	Real Short Rate	Dividend Yield	Inflation	Def. Spread	Term Spread
Mean	4.002	1.469	5.946	6.549	1.867	36.948	4.667	2.039	1.247
Standard Deviation	16.069	7.842	13.781	1.004	1.217	4.152	1.189	0.167	0.602
Sharpe Ratio	0.072	0.054	0.125	-	-	-	-	-	-

Correlation Matrix

	Stocks	Bonds	Real Estate	Nominal Short Rate	Real Short Rate	Dividend Yield	Inflation	Def. Spread	Term Spread
Stocks	1								
Bonds	0.176	1							
Real Estate	0.570	0.167	1						
Nominal Short Rate	-0.069	0.127	-0.104	1					
Real Short Rate	0.121	0.237	0.074	0.440	1				
Dividend Yield	-0.071	-0.025	-0.061	0.680	0.074	1			
Inflation	-0.182	-0.136	-0.163	0.394	-0.652	0.498	1		
Def Spread	0.017	0.280	0.062	-0.032	0.232	-0.125	-0.264	1	
Term Spread	0.050	-0.347	0.086	-0.706	-0.362	-0.184	-0.227	-0.084	1

Panel B: Sample 1972-1994

	Stocks	Bonds	Real Estate	Nominal Short Rate	Real Short Rate	Dividend Yield	Inflation	Def Spread	Term Spread
Mean	2.755	0.715	4.489	7.595	1.953	44.604	5.620	1.960	1.238
Standard Deviation	15.974	8.113	14.090	1.002	1.306	2.794	1.216	0.159	0.673
Sharpe Ratio	0.050	0.025	0.092	-	-	-	-	-	-

Correlation Matrix

	Stocks	Bonds	Real Estate	Nominal Short Rate	Real Short Rate	Dividend Yield	Inflation	Def Spread	Term Spread
Stocks	1								
Bonds	0.305	1							
Real Estate	0.675	0.210	1						
Nominal Short Rate	-0.078	0.189	-0.089	1					
Real Short Rate	0.144	0.301	0.114	0.470	1				
Dividend Yield	-0.106	-0.005	-0.050	0.655	0.095	1			
Inflation	-0.218	-0.168	-0.196	0.319	-0.687	0.437	1		
Def Spread	0.139	0.351	0.097	0.226	0.413	0.175	-0.258	1	
Term Spread	0.051	-0.389	0.075	-0.777	-0.333	-0.356	-0.283	-0.168	1

Table 2 – Classical Parameter Estimates for a VAR(1) Model

The table reports the MLE estimation outputs for the Gaussian VAR(1) model:

$$z_t = \mu + \Phi z_{t-1} + \varepsilon_t$$

where z_t includes continuously compounded monthly excess asset returns and the dividend yield, and $\varepsilon_t \sim N(\mathbf{0}, \Sigma)$. t statistics are reported in parenthesis under the corresponding point estimates. Bold coefficients imply a p-value of 0.1 or lower.

	Stocks $_t$	Bonds $_t$	Real Estate $_t$	Dividend Yield $_t$	Term Spread $_t$	Def Spread $_t$	Real Short Rate $_t$	Inflation $_t$
μ'								
	0.009 (0.710)	-0.007 (-1.199)	0.001 (0.082)	0.000 (0.399)	0.001 (3.393)	0.000 (-0.157)	-0.004 (-4.889)	0.003 (3.608)
Φ'								
Stocks $_{t-1}$	-0.014 (-0.241)	-0.057 (-1.922)	0.101 (1.980)	0.001 (0.572)	0.001 (0.491)	0.000 (-1.139)	-0.005 (-1.312)	0.004 (1.308)
Bonds $_{t-1}$	-0.084 (-0.725)	0.017 (0.292)	0.208 (2.097)	0.004 (0.871)	0.002 (0.407)	-0.001 (-1.466)	0.030 (4.111)	-0.031 (-4.803)
Real Estate $_{t-1}$	0.113 (1.619)	-0.025 (-0.736)	-0.018 (-0.295)	-0.004 (-1.546)	0.009 (4.004)	-0.001 (-2.330)	-0.011 (-2.522)	0.003 (0.751)
Dividend Yield $_{t-1}$	1.446 (4.262)	0.349 (2.091)	0.974 (3.362)	0.954 (80.423)	0.013 (1.195)	-0.003 (-2.425)	-0.101 (-4.807)	0.089 (4.639)
Term Spread $_{t-1}$	-10.843 (-4.426)	-2.245 (-1.865)	-5.816 (-2.781)	0.432 (5.045)	0.314 (3.961)	0.025 (3.141)	1.168 (7.697)	-0.507 (-3.679)
Def Spread $_{t-1}$	4.736 (0.938)	7.053 (2.841)	6.730 (1.560)	-0.406 (-2.303)	-0.104 (-0.639)	0.971 (58.536)	0.905 (2.892)	-0.844 (-2.969)
Real Short rate $_{t-1}$	-6.563 (-3.575)	-1.316 (-1.459)	-4.666 (-2.977)	0.195 (3.043)	-0.111 (-1.874)	0.017 (2.895)	1.282 (11.271)	-0.185 (-1.795)
Inflation $_{t-1}$	-9.500 (-4.553)	-2.431 (-2.370)	-6.288 (-3.529)	0.303 (4.162)	-0.193 (-2.858)	0.023 (3.311)	0.952 (7.365)	0.233 (1.985)
Covariance matrix (for VAR(1) residuals)								
	Stocks $_t$	Bonds $_t$	Real Estate $_t$	Dividend Yield $_t$	Term Spread $_t$	Def Spread $_t$	Real Short Rate $_t$	Inflation $_t$
Stocks $_t$	0.002	0.137	0.547	-0.885	0.059	-0.199	0.163	-0.213
Bonds $_t$	0.000	0.000	0.134	-0.200	-0.400	0.359	0.224	-0.08
Real Estate $_t$	0.001	0.000	0.001	-0.558	0.082	-0.188	0.078	-0.129
Dividend Yield $_t$	-6.2E-05	-6.9E-06	-3.3E-05	2.4E-06	-0.067	0.170	-0.168	0.226
Term Spread $_t$	3.8E-06	-1.3E-05	4.6E-06	-1.5E-07	2.1E-06	-0.282	-0.443	-0.043
Def Spread $_t$	-1.3E-06	1.1E-06	-1.1E-06	3.9E-08	-6.0E-08	2.2E-08	0.062	0.015
Real Short rate $_t$	2.0E-05	1.4E-05	8.2E-06	-7.3E-07	-1.8E-06	2.5E-08	7.7E-06	-0.872
Inflation $_t$	-2.4E-05	-4.6E-06	-1.2E-05	8.9E-07	-1.6E-07	5.5E-09	-6.1E-06	6.3E-06

Table 3 – Bayesian Coefficient Estimates for a VAR(1) Model

The table reports the Bayesian posterior means for the coefficients of the Gaussian VAR(1) model:

$$z_t = \mu + \Phi z_{t-1} + \varepsilon_t$$

where z_t includes continuously compounded monthly excess asset returns, the real 1-month interest rate, the rate of inflation, the term spread, and the dividend yield; $\varepsilon_t \sim N(\mathbf{0}, \Sigma)$. The standard errors of the Bayesian posterior densities are reported in parenthesis under the corresponding posterior means. The posteriors are obtained from a standard uninformative prior, $p(\mathbf{C}, \Sigma) \propto |\Sigma|^{-(n+2)/2}$, where $\mathbf{C} = [\mu' \Phi']'$ is the matrix of the coefficients in the VAR model and n is the number of variables (4) in the multivariate system. The lower panel shows volatilities and covariances on the main diagonal and below it, and implied pairwise correlations in the upper triangular portion.

	Stocks_t	Bonds_t	Real Estate_t	Dividend Yield_t	Term Spread_t	Default Spread_t	Real Short Rate_t	Inflation_t
	μ'							
	0.009	-0.007	0.001	0.000	0.001	0.000	-0.004	0.003
	(0.014)	(0.007)	(0.012)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
	Φ'							
Stocks_{t-1}	-0.015	-0.057	0.101	0.001	0.001	0.000	-0.005	0.004
	(0.069)	(0.034)	(0.059)	(0.002)	(0.002)	(0.000)	(0.004)	(0.004)
Bonds_{t-1}	-0.083	0.016	0.209	0.004	0.002	-0.001	0.030	-0.031
	(0.133)	(0.066)	(0.114)	(0.005)	(0.004)	(0.000)	(0.008)	(0.007)
Real Estate_{t-1}	0.113	-0.025	-0.018	-0.004	0.009	-0.001	-0.011	0.003
	(0.080)	(0.039)	(0.068)	(0.003)	(0.003)	(0.000)	(0.005)	(0.004)
Dividend Yield_{t-1}	1.449	0.350	0.977	0.953	0.013	-0.003	-0.101	0.089
	(0.388)	(0.191)	(0.333)	(0.014)	(0.013)	(0.001)	(0.024)	(0.022)
Term Spread_{t-1}	-10.834	-2.259	-5.837	0.432	0.314	0.025	1.168	-0.507
	(2.790)	(1.375)	(2.384)	(0.097)	(0.090)	(0.009)	(0.173)	(0.158)
Def Spread_{t-1}	4.713	7.061	6.752	-0.406	-0.104	0.971	0.902	-0.841
	(5.786)	(2.833)	(4.935)	(0.202)	(0.186)	(0.019)	(0.358)	(0.326)
Real Short rate_{t-1}	-6.564	-1.322	-4.682	0.196	-0.111	0.017	1.282	-0.185
	(2.099)	(1.036)	(1.798)	(0.073)	(0.068)	(0.007)	(0.130)	(0.118)
Inflation_{t-1}	-9.501	-2.438	-6.307	0.304	-0.193	0.023	0.952	0.233
	(2.378)	(1.173)	(2.037)	(0.083)	(0.077)	(0.008)	(0.147)	(0.134)

Table 3 (continued) – Bayesian Coefficient Estimates for a VAR(1) Model

	Covariance matrix (for VAR(1) shocks)							
	Stocks _t	Bonds _t	Real Estate _t	Dividend Yield _t	Term Spread _t	Default Spread _t	Real Short Rate _t	Inflation _t
Stocks_t	2.6E-03 (1.8E-04)	0.138	0.547	-0.885	0.059	-0.199	0.164	-0.213
Bonds_t	1.8E-04 (6.6E-05)	6.3E-04 (4.4E-05)	0.134	-0.200	-0.400	0.349	0.224	-0.083
Real Estate_t	1.2E-03 (1.3E-04)	1.5E-04 (5.7E-05)	0.002 (1.3E-04)	-0.557	0.082	-0.188	0.078	-0.128
Dividend Yield_t	-8.1E-05 (5.6E-06)	-9.0E-06 (2.3E-06)	-4.3E-05 (4.4E-06)	3.2E-06 (2.0E-07)	-0.067	0.170	-0.168	0.226
Term Spread_t	5.0E-06 (4.3E-06)	-1.7E-05 (2.2E-06)	5.9E-06 (3.7E-06)	-2.0E-07 (1.5E-07)	2.7E-06 (1.7E-07)	-0.282	-0.443	-0.043
Def Spread_t	-1.7E-06 (4.4E-07)	1.5E-06 (2.3E-07)	-1.4E-06 (3.8E-07)	5.1E-08 (1.5E-08)	-7.8E-08 (1.5E-08)	2.8E-08 (1.8E-09)	0.062	0.015
Real Short rate_t	2.6E-05 (8.4E-06)	1.8E-05 (4.2E-06)	1.1E-05 (7.2E-06)	-9.5E-07 (2.9E-07)	-2.3E-06 (2.8E-07)	3.3E-08 (2.7E-08)	(1.0E-05 (6.6E-07)	-0.872
Inflation_t	-3.1E-05 (7.7E-06)	-6.0E-06 (3.7E-06)	-1.6E-05 (6.5E-06)	1.2E-06 (2.7E-07)	-2.0E-07 (2.4E-07)	7.2E-09 (2.5E-08)	-7.9E-06 (5.4E-07)	8.2E-06 (5.1E-07)

Table 4 – Optimal Average Portfolio Weights

This table reports mean optimal portfolio weights for stocks, bonds and cash, for the two alternative cases in which real estate is or is not in the asset menu. The investment horizon varies from 1 (first two rows) to 60 months (last two rows). For each horizon, means are computed over 120 monthly portfolio allocations on the recursive sample 1995-2004. In the Bayesian case, parameter uncertainty is accounted for. The coefficient of relative risk aversion is set equal to 5.

	Stock		Bond		E-REIT	Cash	
	<i>With E-REIT</i>	<i>Without E-REIT</i>	<i>With E-REIT</i>	<i>Without E-REIT</i>		<i>With E-REIT</i>	<i>Without E-REIT</i>
T=1							
Classical	0.13	0.38	0.27	0.43	0.50	0.10	0.19
Bayesian	0.11	0.33	0.25	0.27	0.43	0.21	0.40
T=3							
Classical	0.15	0.40	0.18	0.27	0.49	0.18	0.33
Bayesian	0.11	0.33	0.16	0.22	0.38	0.35	0.45
T=6							
Classical	0.17	0.43	0.15	0.25	0.50	0.18	0.32
Bayesian	0.12	0.33	0.10	0.17	0.38	0.40	0.50
T=12							
Classical	0.17	0.46	0.07	0.14	0.53	0.23	0.40
Bayesian	0.13	0.33	0.04	0.11	0.38	0.45	0.56
T=24							
Classical	0.19	0.47	0.02	0.10	0.54	0.25	0.43
Bayesian	0.12	0.31	0.02	0.07	0.37	0.49	0.62
T=60							
Classical	0.29	0.57	0.00	0.07	0.59	0.12	0.36
Bayesian	0.14	0.29	0.01	0.07	0.33	0.52	0.64

Table 5 – Ex-Post Performance (1995 - 2004 Sample)

The table shows the mean ex post performance of optimal portfolios recursively computed over the sample January 1995 - November 2004. Performance measures are computed for an investor with constant relative risk aversion equal to 5 and different investment horizons (from 1 to 60 months). Two alternative asset menus are considered, with and without E-REITs. Panel A reports the performance of classical optimal portfolios while panel B covers Bayesian portfolios.

Classical

	<i>T=1</i>			<i>T=3</i>			<i>T=6</i>			<i>T=12</i>			<i>T=24</i>			<i>T=60</i>		
	<i>With E-REIT</i>	<i>Without E-REIT</i>	Δ															
Sharpe ratio	0.442	0.186	0.256	0.383	0.128	0.256	0.358	0.123	0.235	0.304	-0.159	0.463	0.215	0.052	0.163	0.390	0.085	0.306
Certainty Equivalent	8.742	6.810	1.931	8.263	6.436	1.827	8.008	6.279	1.729	7.366	4.459	2.907	6.288	4.344	1.944	7.587	4.123	3.464
Annual Mean Returns	10.782	8.145	2.637	9.997	7.625	2.372	9.781	7.555	2.226	9.571	5.502	4.070	8.712	6.865	1.847	9.220	7.080	2.140
Annualized Volatility	8.452	7.040	-1.411	8.069	6.786	-1.283	8.394	7.050	1.343	9.948	6.592	-3.356	11.836	7.572	-4.265	11.619	18.972	7.353

Bayesian

	<i>T=1</i>			<i>T=3</i>			<i>T=6</i>			<i>T=12</i>			<i>T=24</i>			<i>T=60</i>		
	<i>With E-REIT</i>	<i>Without E-REIT</i>	Δ															
Sharpe ratio	0.447	0.031	0.416	0.319	0.041	0.278	0.261	0.005	0.256	0.220	0.002	0.218	0.137	-0.002	0.140	0.299	0.074	0.225
Certainty Equivalent	8.750	6.001	2.749	7.794	6.166	1.628	7.376	5.898	1.478	6.970	5.461	1.509	6.224	4.764	1.460	6.773	5.028	1.745
Annual Mean Returns	10.247	6.946	3.302	8.764	6.951	1.813	8.317	6.684	1.633	8.118	6.562	1.556	7.351	6.330	1.021	7.327	6.517	0.810
Annualized Volatility	7.266	5.933	-1.333	6.083	5.528	-0.555	6.131	5.559	0.572	7.123	6.650	-0.472	7.804	5.603	-2.201	6.465	8.350	1.885

Footnotes

1. Section 2 explains in some detail what the classical and Bayesian asset allocations models consists of. In essence, in a classical framework, the parametric model that is used to capture predictability patterns and to generate predictions for means, variances, and covariances of returns is (counter-factually) taken to correspond to the true, but unknown generating process. This completely ignores the presence of any parameter uncertainty. In a Bayesian framework, the parameters of the model are instead considered to be random variables and therefore the presence of estimation uncertainty is fully integrated out when deriving predictions of means, variances, and covariances.
2. In our paper, log returns on individual assets are assumed to be normally distributed. However the resulting optimal portfolio returns are not.
3. Giorgiev, Gupta and Kunkel (2003) find instead negligible increases in the Sharpe ratio over 1990-2002.
4. Indeed, an investor's welfare can substantially increase if she takes into account the uncertainty in forecasts by using Bayesian updating (see Jorion, 1985; Kandel and Stambaugh, 1996), especially when return predictability appears weak according to classical statistical tests.
5. Uninformative priors may be a suboptimal choice, even in in-sample exercises. Hoovernaars et al. (2007b) develop the concept of robust portfolio—the portfolio of an investor with a prior that has minimal welfare costs when evaluated under a wide range of alternative priors. We do not pursue this extension in the current paper.
6. The dividend yield is widely used in the literature as a predictor of future excess asset returns. See Campbell and Shiller (1988), Fama and French (1989), and Kandel and Stambaugh (1996). Karolyi and Sanders (1998) and Liu and Mei (1992) find that the dividend yield also helps predicting REIT returns.
7. Following Brandt (1999) and Campbell, Chan, and Viceira (2003), we allow the slope of the yield curve and the spread between high- and medium-rated debt, both anticipating business cycle dynamics, to predict future asset returns.
8. This result confirms previous evidence on the higher degree of predictability of REITs when compared to equities and bonds, see e.g., Liao and Mei (1998).
9. In fact, the dividend yield and the term spread also turn out to be rather persistent variables, although unreported tests confirm the stationarity of the VAR(1) system. These results are common to the existing literature, see e.g., Fugazza, Guidolin, and Nicodano (2007) on a different data set. Detailed results on the econometric estimates and related tests are available upon request.
10. Opposite interpretation applies when $\phi_{12} < 0$ and $\rho < 0$, in the sense that this configuration configures a built-in element of mean-aversion that makes an asset riskier, the longer the investment horizon. Crucially, these effects may be of first-order importance even when the standard errors associated with many of the coefficients in (VAR(1) model) are high and the estimated VAR(1) coefficients relatively small, as long as adequate covariance loadings come through estimates of the off-diagonal elements of the covariance matrix of the residuals.
11. In a similar way, bond returns tend to be positively serially correlated because of their positive shock correlation with the default premium. A fall in the default spread is associated, through a positive contemporaneous correlation (0.35), with both lower unexpected bond returns today and lower expected bond returns tomorrow, through a positive VAR coefficient (7.05). This makes holding a bond for two periods riskier than in the absence of such intertemporal links.
12. We perform these calculations for a range of alternative coefficients of relative risk aversion ($\gamma = 2$,

4, 5, and 10). Since the results do not appear to be overly sensitive – especially as far as welfare costs are concerned – to a specific choice of preferences, in what follows we only report results for the case $\gamma = 5$. Further, detailed results are available upon request.

13. Notice that our arguments are by construction simple partial equilibrium arguments concerning optimal investment policies.

14. We could also assess the effects of market timing by comparing the $T = 1$ allocation with predictable versus i.i.d. returns that ignores all forms of predictability. The allocation for $T = 1$ may differ from the i.i.d. benchmark for market timing effects, which derive from conditional moments of asset returns being different than unconditional ones due to predictable returns. However, for $T = 1$ there is no intertemporal hedging demand. Results are available upon request.

15. The asymptotic distribution of the statistic $SR_{i,T}$ is obtained by observing that it is a ratio of elements of $\bar{\mathbf{r}}_T$ and \mathbf{S}_T which have known small sample distributions when returns are assumed to be bivariate normal:

$$SR_{i,T} = \frac{\bar{r}_{i,T}}{s_{i,T}} \stackrel{asy}{\sim} N\left(\frac{\mu_i}{\sigma_i}, \frac{1}{\tau} \left[1 + \frac{\mu_i^2}{2\sigma_i^2}\right]\right),$$

where the variance is derived from $O(1/\tau)$ Taylor series approximations.

16. Notice that we are performing classical inference on Sharpe ratios obtained from portfolio strategies computed from a Bayesian statistical framework.

17. We set strategy 2 to include real estate and strategy 1 not to. Therefore we expect that if $SR_{2,T} > SR_{1,T}$, then $\Delta TSR_T \equiv \bar{r}_{1,T} s_{2,T} - \bar{r}_{2,T} s_{1,T} < 0$. Jobson and Korkie (1981) stress that at least a few dozen observations are needed for their tests to yield satisfactory power. This should be no problem in our application because we always have at least $\tau = 60$ observations available.

18. When returns are normally distributed, total variance and semi-variance (which conditions on returns being below their mean) are identical. Therefore deviations from normality are needed for total and downside variance to differ.

19. Tentatively, we ascribe the difference between Jorion's (1985) and our results to the fact that he considers estimation risk concerning mean returns only. On the contrary, our method accounts for estimation errors affecting all parameter estimates.

20. Details on the IID case are available from the authors. Interestingly, in this case only adding real estate can reduce ex-post welfare – if the investor does not take parameter uncertainty into account.

21. However such a choice presents a few issues. First, CREF indices are available for shorter periods than NAREITs and only at quarterly frequencies. Short time series are of course problematic for our time series approach to model predictability. Second, as it is well known in the literature (see Geltner, 1993, and Cho, Kawaguchi, and Shilling, 2003), CREF index returns (like NCREIF) should be unsmoothed to adjust for appraisal smoothing biases. Third, a portion of the literature has argued that comparisons between NAREITs (which are typically similar to equity stock, and therefore levered on both sides of the balance sheets of the underlying investment trusts) and CREF index returns would require that the former be un-levered (see e.g., Fisher, Geltner, and Webb, 1994). However, in our application that would create further problems of homogeneity with stock indices, which are themselves (often heavily) levered financial instruments.