Investing for the Long-Run in European Real Estate. Does Predictability Matter?*

Carolina FUGAZZA[†]
Center for Research on Pensions and Welfare Policies

Massimo GUIDOLIN[‡]
Federal Reserve Bank of St. Louis

Giovanna NICODANO[§]
Center for Research on Pensions and Welfare Policies
and University of Turin

November 2004

Abstract

We calculate optimal portfolio choices for a long-horizon, risk-averse European investor who diversifies among stocks, bonds, real estate, and a liquid asset when excess asset returns are predictable. Simulations are performed for scenarios involving different risk aversion levels, horizons, and statistical models employed to capture predictability in risk premia. Importantly, under one of the scenarios, the investor takes into account the parameter uncertainty implied by the use of estimated coefficients to characterize predictability. We find that real estate ought to play a significant role in optimal portfolio choices, with weights between 10 and 50%. Under plausible assumptions, welfare costs of either ignoring predictability or restricting portfolio choices to financial assets only are found to be in the order of at least 100 basis points per year. These results are robust to changes in the benchmarks and in the statistical framework.

1. Introduction

Predictability of asset returns is kown to have powerful effects on the structure and dynamics of optimal portfolio weights for a long-horizon investor. This conclusion holds across alternative models for predictability, different data sets, and heterogeneous asset allocation frameworks, see e.g. Brennan, Schwartz, and Lagnado (1997), Campbell and Viceira (1999, 2001), and Detemple, Garcia, and Rindisbacher (2003). However, most of this evidence has been obtained in asset menus limited to traditional financial indexes only, i.e. stocks, bonds, and short-term liquid assets.

Our paper provides further evidence on the effects of predictability on long-run portfolio choice when the asset menu includes one of the most exploited types of asset in long-horizon portfolios: real estate. Furthermore, our asset allocation results are based on predictability patterns characterizing a European data set that has been left unexplored thus far. Clearly, both extensions are crucial to make the results

---

found in the literature relevant to the operational goals of long-horizon investors (like pension funds and insurance companies) that commonly employ asset menus not limited to financial securities only, and that fail to limit their portfolio choices to North American assets only. Obviously, among them European institutional investors occupy a leading position.

From an empirical viewpoint, contributions available to pension funds are invested not only in equity and bonds, but in real estate property as well (see Hudson-Wilson, Fabozzi, and Gordon (2003)). For instance, as of the mid 1990s, in the UK 75 and 7.8 per cent were held in stocks and real estate respectively, while the corresponding weights were 6.6 and 4.2 in Germany, and 26.9 and 2.2 in France. In the last two countries, long-term bonds represented 42.3 and 59.0 of long term portfolios (see Miles, 1994, p.23), while bonds were given a negligible weight in the UK. So it appears that considerable heterogeneity exists in the relative weights assigned to stocks, bonds and real estate in three major European countries. Although our paper is normative, tracing out the implications of predictability for *optimal* portfolio composition, we report results that make it clear under what conditions one may obtain rational choices consistent with either the German-French pattern (dominated by bonds) or with the British one (dominated by stocks). Additionally − since the evidence is for real estate weights between 4 and 8 percent − we will understand whether existing data support the notion that real estate ought to be included in long-horizon portfolios, although with rather moderate weights.

We use a simple vector autoregressive framework to capture predictable time variations in the investment opportunity set, similarly to Campbell, Chan, and Viceira (2003), Lynch (2001), and others. We solve a standard portfolio problem with isoelastic ut of terminal wealth) and find that − contrary to the bulk of the existing evidence − results on the importance of predictability for optimal portfolio weights are mixed. On one hand, predictability of mean excess returns (i.e. predictable time variation in risk premia) has first-order effects for the optimal allocation between equities and bonds, similarly to Campbell and Viceira (2002) and Guidolin and Timmermann (2004), that are robust to taking parameter uncertainty into account, see Barberis (2000). On the other hand, predictability in risk premia produces only second-order effects for the optimal diversification between financial securities vs. real estate assets, in the sense that results do not depart significantly from the weights obtained under the (false) assumption of independently and identically distributed returns over time. In all cases, the optimal weight to be assigned to real estate is between 10 and 30 percent of the initial wealth.

However, the overall picture is consistent with the general finding that long-run investors with an interest in European assets ought to carefully consider the effects of time-varying risk premia in their portfolio choices. In fact, the estimated welfare costs from ignoring predictability are large, in the order of 100 basis points per year for a long-run (10-year) investor with a plausible coefficient of relative risk aversion of 5. Additionally, we compute the costs of restricting the available asset menu to financial securities only, thus ignoring real estate. We find that for long-horizon investors the resulting damage would be substantial, once more in the approximate order of 100 basis points per year for a long-run, intermediate risk-averse investor. Such figure may however climb up to more than 200 basis points under some configurations of the predictability model and assuming a coefficient of relative risk aversion of 10. Therefore we report evidence that the utility gains from investing in real estate are of first-order magnitude, although the associated portfolio weights may be moderate and only weakly affected by time-variations in conditional risk premia.

Our paper contributes to three distinct literatures. Several studies have compared the risk and return

characteristics of stocks, bonds, and cash to real estate and analyzed optimal portfolio choice in a mean variance framework, see e.g. Li and Wang (1995). However, there is still considerable uncertainty on the optimal weight one should assign real estate. Among the others, Hudson-Wilson, Fabozzi, and Gordon (2003), Karlberg, Liu, and Greig (1996), Liang, Myer, and Webb (1996), and Ziobrowski, Caines, and Ziobrowski (1999) calculate optimal mean-variance US optimal portfolios when real estate belongs to the asset menu and is measured by direct (appraisal-based) indices. They find that in the US, real estate ought to have a rather negligible weight, although its importance increases when bootstrap methods are employed to account for the uncertainty surrounding the actual distribution of returns. On the opposite, Chandrashakaran (1999) and Liang and Webb (1996) find much larger weights using longer time series and/or different data (e.g. hedged REITs). However, none of these papers examine the predictability patterns of the returns on these asset classes, which affect the risk premia and variance of cumulative returns and hence their desirability in a multi-period setting. On the other hand, it is well known that while the investor's planning horizon is irrelevant for asset allocation when returns are independently and identically distributed (Samuelson, 1969; Merton, 1969, 1971), when returns are predictable the mean-variance asset allocation may differ substantially from the long-term one (see Samuelson, 1994, Bodie, 1995, Siegel, 1998 and Campell and Viceira, 2002, among others). Therefore, by taking predictability into account, our paper departs from the earlier literature on portfolio management when real estate is available. Moreover, at least to our knowledge, our paper is the first attempt at taking parameter uncertainty into account in a framework that includes real estate (and in which parameter estimation may substantially affect economic behavior).

Another literature has shown that stock returns predictability may affect long-term portfolio choice in two ways (e.g. Campbell, Chan, and Viceira, 2003). First, an investor would have powerful incentives to regularly rebalance his portfolio as he receives new information on the conditional risk premium of the available assets. Interestingly, Balduzzi and Lynch (1999) have shown that the expected utility gains derived from exploiting conditional information are robust to the inclusion of transaction costs incurred at the rebalancing points. Secondly (and assuming preferences different from log-utility), even a buy and hold investor would modify his asset holdings in order to hedge future adverse predictable changes in investment opportunities. When the asset menu is restricted to financial assets and a vector autoregressive (VAR) system captures return predictability in the US, Campbell and Viceira (1999, 2002) and Barberis (2000) have shown that mean-reversion in stock returns implies that average stock holdings generally increase in the investors' horizon for low risk aversion. Average nominal bond holdings typically increase in the degree of risk aversion, as their lower Sharpe ratio is accompanied by lower volatility. We extend these results to an asset menu that includes real estate and find that also real estate investment schedules are generally upward sloping, although to a lesser extent than stock schedules. We also document the effects of dynamic portfolio choices (rebalancing) and conclude that even in the presence of real estate, portfolio choices ought to strongly react to information that predicts future asset returns.

Finally, there is a recent econometric literature that has modeled returns on European assets and pinned down predictability patterns potentially useful for portfolio choice. For instance, extensive work has concerned the properties of UK stock and bond returns (see among the others Clare, Thomas and Wickens, 1994; Black and Fraser, 1995; Pesaran and Timmermann, 2002), although the recent debate has mostly focused on whether non-linear models (Flavin and Wickens, 2001; Guidolin and Timmermann, 2003)

3

may be needed to calculate optimal portfolio weights, see Guidolin and Timmermann (2004). However, to the best of our knowledge our paper is the first one to adopt an explicit aggregate, European focus and to jointly model the VAR structure present in financial and real estate returns. We find that in a simple linear, VAR framework, there is evidence of forecastability of returns compatible with the traditional notion that risky assets (i.e. real estate and stocks) are less risky in the long-run than they are over short periods of time.

The two papers most closely related to ours are Barberis (2000) and Bharati and Gupta (1992; 2004). Barberis carefully investigates the portfolio choice effects of (VAR) predictability when the latter is characterized through parametric models that are subject to estimation uncertainty (see also Brandt, 1999, on this point). The uncertainty about parameters, i.e. estimation risk, can be taken into account when solving long-run portfolio problems by adopting a Bayesian approach and integrating over the posterior density of the parameters to obtain the (multivariate) predictive density of future asset returns. Barberis (2000) goes on to show that optimal portfolio weights obtained ignoring parameter uncertainty (i.e. conditional on fixed parameter estimates) may be drastically different from those that account for parameter uncertainty. We adopt the same approach here. However, our asset allocation problem is more realistic and rich than Barberis', including bonds and real estate besides stocks and cash. Moreover, we use different (European) data and perform a number welfare calculations that illustrate the importance of both accounting for predictability and of expanding the asset menu to include real estate.

Bharati and Gupta (1992) model predictability in US asset returns − including real estate, measured as REITs returns − by using predictive regressions that employ typical variables such as the 1-month T-bill rate, the term spread, the default spread, monthly dummies etc. (see Pesaran and Timmermann (1995)). Long-horizon portfolio models are used to calculate optimal portfolio choices. They find that predictability and real estate as an asset class are both important, in the sense that active strategies involving real estate holdings outperform passive ones, even in presence of transaction costs. However their paper disregards parameter uncertainty, and uses a predictability framework that maximizes predictive R-squares through increases in the number of state variables that make it difficult to apply dynamic portfolio optimization methods. No welfare costs from omitting real estate are calculated, making their criteria of evaluation possibly inconsistent with the investors' objectives.

The plan of the paper is as follows. Section 2 describes the methodology of the paper, in terms of both statistical modeling and solution methods applied to the asset allocation problem. Section 3 describes the data and reports results on their statistical properties, revealing the existence of exploitable predictable patterns in the dynamics of the investment opportunity set. Section 4 is the core section of the paper. We characterize optimal portfolios including real estate, and compare them to the case without predictability and parameter uncertainty. In Section 6, we calculate welfare costs to quantify the importance of predictability and of expanding the asset menu to include real estate. Section 7 contains a few robustness checks involving both the statistical model and the choice of the benchmarks for welfare cost calculations. Section 8 concludes.

2. **Long Run Asset Allocation when Returns Are Predictable**

In this section we review the structure and solution methods for a portfolio choice problem over the long run when returns are predictable and when 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 based on two rather different approaches: buy-and-hold vs. optimal rebalancing strategies. 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 notwithstanding the possible receipt of important news characterizing the available investment opportunities. 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), say $t + \varphi$, $t + 2\varphi$, ..., $t + T$. In the following we separately describe the relevant methods distinguishing between buy-and-hold and rebalancing.

Another important dimension of the portfolio problem is whether the statistical uncertainty surrounding the parameters is taken into account or not. In the former case, the approach is usually a Bayesian one, in which conditional expectations are calculated employing the predictive density of future (cumulated) asset returns. In the latter case, the approach is based on the tools commonly employed by classical statisticians. Therefore a second, crucial distinction in the explanations to follow revolves around these aspects.

2.1. Buy-and-Hold Investor

Consider the time t problem of an investor who maximizes expected utility from terminal wealth over a planning horizon of T months by choosing optimal portfolio weights $(\boldsymbol{\omega}_t)$, when preferences are described by a standard, power isoelastic utility function:

$$\max_{\boldsymbol{\omega}_t} E_t \left[\frac{W_{t+T}^{1-\gamma}}{1-\gamma} \right] \quad \gamma \neq 1.^{1}$$

Wealth can be invested in three risky asset classes: stocks, bonds and real estate. The menu is completed by a cash, riskless (in the short-run) investment (1-month deposits). The continuously compounded monthly real return on the free risk asset, r^f, is assumed to be constant as in Barberis (2000), Campbell and Viceira (1999), and Lynch (2001). The continuously compounded excess returns (i.e. $R_t \equiv \ln P_t - \ln P_{t-1}$ where P denotes asset prices) between month $t-1$ and t on stocks, bonds and real estate are denoted by R_t^s, R_t^b, and R_t^r, respectively. The fraction of wealth invested in stocks, in bonds, and in real estate are ω_t^s, ω_t^b, and ω_t^r, respectively, so that $\boldsymbol{\omega}_t \equiv [\omega_t^s \ \omega_t^b \ \omega_t^r]'$. The investor's terminal wealth, when initial wealth W_t is equal to one, is given by:

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

where $R_{t,T}^s$, $R_{t,T}^b$, and $R_{t,T}^r$ denote the cumulative returns between t and T on stocks, bonds and real estate, respectively:

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

In general, we call n the number of risky asset classes modeled. Our baseline experiment concerns $n = 3$. Furthermore, we follow the bulk of the literature imposing no-short sale constraints, i.e. $1 \geq \omega_t^s, \omega_t^b, \omega_t^r \geq 0$.

Therefore the buy-and-hold problem may be simply re-written as

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

$$\text{s.t. } 1 \geq \omega_t^s \geq 0 \ \omega_t^b \geq 0 \ \omega_t^r \geq 0. \tag{1}$$

Time-variation in the excess returns are modelled using a Gaussian VAR(1) framework, for instance:

$$\mathbf{z}_t = \boldsymbol{\mu} + \Phi_R \mathbf{z}_{t-1} + \boldsymbol{\epsilon}_t \text{ or} \tag{2}$$

$$\begin{bmatrix} R_t^s \\ R_t^b \\ R_t^r \\ x_t \end{bmatrix} = \begin{bmatrix} \mu^s \\ \mu^b \\ \mu^r \\ \mu^x \end{bmatrix} + \begin{bmatrix} \phi^{ss} & 0 & 0 & \phi^{sx} \\ 0 & \phi^{bb} & 0 & \phi^{bx} \\ 0 & 0 & \phi^{rr} & \phi^{rx} \\ 0 & 0 & 0 & \phi^{xx} \end{bmatrix} \begin{bmatrix} R_{t-1}^s \\ R_{t-1}^b \\ R_{t-1}^r \\ x_{t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_t^s \\ \epsilon_t^b \\ \epsilon_t^r \\ \epsilon_t^x \end{bmatrix}$$

where $\boldsymbol{\epsilon}_t$ is i.i.d. $N(\mathbf{0}, \Sigma)$ and x_t represents any economic variable (e.g. industrial production growth, the term structure of interest rates, default risk premia indicators, etc.) able to forecast future asset returns. Model (2) implies that

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

i.e. the conditional risk premia on the assets are time-varying and function of past excess asset returns as well as lagged values of the predictor variable x_{t-1}. In particular, as in Barberis (2000), the zero restrictions appearing in the vector autoregressive matrix Φ_R will normally imply that most of the predictability in asset risk premia is explained by lagged values of x_t, besides possible persistent components captured by (positive) ϕ^{ss}, ϕ^{bb}, and ϕ^{rr}.

Although for most North-American data sets, the restrictions to Φ_R are generally supported by formal statistical tests, and papers like Campbell, Chen, and Viceira (2003) and Guidolin and Timmermann (2004) have shown that relaxing the restrictions does not radically affect optimal portfolio weights, in this paper we also deal with the unrestricted version of (2) in which $\mathbf{z}_t = \boldsymbol{\mu} + \Phi \mathbf{z}_{t-1} + \boldsymbol{\epsilon}_t$, with Φ full matrix.[2] For future reference, we call $\boldsymbol{\theta}$ the vector collecting all the parameters entering (2), i.e. $\boldsymbol{\theta} \equiv [\boldsymbol{\mu}' \ vec(\Phi)' \ vech(\Sigma)']'$.

Under models like (2), the (conditional) distribution of cumulative future returns (i.e. the first n elements in $\mathbf{z}_{t,T} \equiv \sum_{k=1}^{T} \mathbf{z}_{t+k}$) is multivariate normal with mean and covariance matrix given by:

$$\begin{aligned} E_{t-1}[\mathbf{z}_{t,T}] &= T\boldsymbol{\mu} + (T-1)\Phi\boldsymbol{\mu} + (T-2)\Phi^2\boldsymbol{\mu} + ... + \Phi^{T-1}\boldsymbol{\mu} + (\Phi + \Phi^2 + ... + \Phi^T)\mathbf{z}_{t-1} \\ Var_{t-1}[\mathbf{z}_{t,T}] &= \Sigma + (\mathbf{I} + \Phi)\Sigma(\mathbf{I} + \Phi)' + (\mathbf{I} + \Phi + \Phi^2)\Sigma(\mathbf{I} + \Phi + \Phi^2)' + \\ &\quad ... + (\mathbf{I} + \Phi + ... + \Phi^{T-1})\Sigma(\mathbf{I} + \Phi + ... + \Phi^{T-1})', \end{aligned} \tag{3}$$

where $\mathbf{I}$ is the identity matrix of dimension n and Φ^k is a short-hand for $\prod_{i=1}^{k} \Phi$. 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

$$\max_{\boldsymbol{\omega}_t} \int \frac{W_{t+T}^{1-\gamma}}{1 - \gamma} p\left(E_{t-1}[\mathbf{z}_{t,T}], Var_{t-1}[\mathbf{z}_{t,T}]\right) \cdot d\mathbf{z}_{t,T} \tag{4}$$

---

[2]We also experiment relaxing the first-order VAR constraint but find that for all exercises performed in this paper, a first-order VAR provides the best trade-off between fit and parsimony, i.e. it minimizes a few standard information criteria (AIC and BIC).

(where $p\left(E_{t-1}[\mathbf{z}_{t,T}], Var_{t-1}[\mathbf{z}_{t,T}]\right)$ is a multivariate normal with mean $E_{t-1}[\mathbf{z}_{t,T}]$ and covariance matrix $Var_{t-1}[\mathbf{z}_{t,T}]$), by simulation methods. Similarly to Kandel and Stambaugh (1996), Barberis (2000), and Guidolin and Timmermann (2005), this means evaluating the integral in (4) by drawing a large number of times (N) from $p\left(E_{t-1}[\mathbf{z}_{t,T}], Var_{t-1}[\mathbf{z}_{t,T}]\right)$ and then maximizing the following functional:

$$\max_{\boldsymbol{\omega}_t} \frac{1}{N} \sum_{i=1}^{N} \left[\frac{\left[\omega_t^s \exp(R_{t,T}^{s,i}) + \omega_t^b \exp(R_{t,T}^{b,i}) + \omega_t^r \exp(R_{t,T}^{r,i}) + (1 - \omega_t^s - \omega_t^b - \omega_t^r) \exp(r^f T)\right]^{1-\gamma}}{1-\gamma} \right], \quad (5)$$

where $[R_{t,T}^{s,i} \ R_{t,T}^{b,i} \ R_{t,T}^{r,i}]'$ represent the first three elements of $\mathbf{z}_{t,T}^i$ along a sample path $i = 1, ..., N$. At this stage, the portfolio weight non-negativity constraints are imposed by maximizing (5) using a simple two-stage grid search algorithm that sets ω_t^j to 0, 0.01, 0.02, ..., 0.99, 0.9999 for $j = s, b, r$.[3] In the results that follow, we employ $N = 100,000$ Monte Carlo trials in order to minimize (essentially eradicate) any residual random errors in optimal weights induced by our simulation approach.

2.2. Parameter Uncertainty

Since the true values of the 'parameters' in $\boldsymbol{\theta}$ are unknown, the uncertainty on the actual strength of predictability induced by estimation risk may substantially affect optimal portfolio rules, especially over the long run. As in Barberis (2000), parameter uncertainty is incorporated in the model by using a Bayesian framework that relies on the principle that portfolio choices ought to be based on the multivariate predictive distribution of future asset returns. Such a predictive distribution for future returns (and predictors), conditioned on observed data only, is obtained by integrating the joint distribution of $\boldsymbol{\theta}$ (in this case, interpreted as a random vector) and future 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, $\ddot{\mathbf{Z}}_t \equiv \{\mathbf{z}_i\}_{i=1}^t$. When parameter uncertainty is taken into account, the maximization problem involves the integral:[4]

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

In this case the Monte Carlo methods require drawing a large number of times from $p(\mathbf{z}_{t,T})$ and then 'extracting' cumulative returns only from the resulting vector. However, the task is further 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

---

[3]As explained by Kandel and Stambaugh (1996), $\omega_t^j = 1$ cannot be considered as when 100% of the wealth is invested in some asset and $R_{t,T}^j \rightarrow -\infty$ (bankruptcy) is feasible (as it is the case under a multivariate normal distribution), $\omega_t^j = 1$ would lead to zero or negative wealth which for $\gamma > 1$ would deliver a (meaningless) maximizer for (5).

[4]As it is well known from the Bayesian econometrics literature (see e.g. Hamilton (1993) for an introduction), integrating the joint posterior for $\mathbf{z}_{t,T}$ and $\boldsymbol{\theta}$ with respect to the posterior for $\boldsymbol{\theta}$ delivers a density with fatter tails which simply reflect the additional (estimation) uncertainty implied by $\boldsymbol{\theta}$ being random.

on past data and the parameters. At this point, (2) can be re-written as:

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

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 $\mathbf{C}$. We apply Bayesian methods, we follow Barberis (2000) who extends Zellner's (1971) Bayesian analysis of a multivariate regression model with exogenous regressors to the dynamic regression framework with endogenous regressors. If we consider the following standard uninformative diffuse prior:

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

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

$$
\begin{aligned}
\Sigma^{-1}|\ddot{\mathbf{Z}}_t &\sim \text{Wishart}(t - n - 2, \hat{\mathbf{S}}^{-1}) \\
vec(\mathbf{C})|\Sigma^{-1}, \ddot{\mathbf{Z}}_t &\sim N\left(vec(\hat{\mathbf{C}}), \Sigma \quad (\mathbf{X}'\mathbf{X})^{-1}\right)
\end{aligned}
$$

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 standard OLS estimators for the coefficients and covariance matrix of the residuals. Also for the Bayesian case, we adopt a simulation method that: First, draws N independent variates from $p(\mathbf{C}, \Sigma^{-1}|\ddot{\mathbf{Z}}_t)$. This is done by first sampling from a marginal Wishart for Σ^{-1} and then (after calculating Σ) from the conditional $N\left(vec(\hat{\mathbf{C}}), \Sigma \quad (\mathbf{X}'\mathbf{X})^{-1}\right)$, where $\hat{\mathbf{C}}$ is easily calculated. Second, for each set $(\mathbf{C}, \Sigma)$ thus obtained, samples cumulated excess returns from a multivariate normal with mean vector and covariance matrix given by (3). Given the double simulation scheme, in this case N is set to a relatively large value of 300,000 independent trials.

2.3. *Dynamic Rebalancing Strategies*

The solution method is in this case based on standard dynamic programming principles and on a discretization of the state space. In particular, divide the interval $[t, T]$ into $B \geq 1$ intervals $[t, t + \varphi]$, $[t + \varphi, t + 2\varphi]$, ..., $[t + (B-1)\varphi, t + B\varphi]$, where $B = T/\varphi$ and assume the rebalancing occurs at regular intervals B times over $[t, T]$. The problem is then similar to (1), with the only difference that the objective ought to be maximized by choosing the entire sequence $\{\boldsymbol{\omega}_{t+\varphi}\}_{\varphi=0}^{B-1}$:

$$
\max_{\{\boldsymbol{\omega}_{t+k\varphi}\}_{k=0}^{B-1}} E_t\left[\frac{W_{t+T}^{1-\gamma}}{1-\gamma}\right]
$$

$$
\text{s.t. } W_{t+(k+1)\varphi} = W_{t+k\varphi}[\omega_{t+k\varphi}^s \exp(R_{t+k\varphi, t+(k+1)\varphi}^s) + \omega_{t+k\varphi}^b \exp(R_{t+k\varphi, t+(k+1)\varphi}^b) +
$$

$$
+ \omega_{t+k\varphi}^r \exp(R_{t+k\varphi, t+(k+1)\varphi}^r) + (1 - \omega_{t+k\varphi}^s - \omega_{t+k\varphi}^b - \omega_{t+k\varphi}^r) \exp(r^f \varphi)]
$$

$$
1 \geq \omega_{t+k\varphi}^s \geq 0 \quad 1 \geq \omega_{t+k\varphi}^b \geq 0 \quad 1 \geq \omega_{t+k\varphi}^r \geq 0, \quad k = 0, 1, ..., B-1,
$$

where cumulated returns $R^j_{t+k\varphi,t+(k+1)\varphi} \equiv \sum_{i=1}^{\varphi}(R^j_{t+i} + r^f)$ $(j = s, b, r)$ are defined similarly to Section 2.1. Standard arguments (see e.g. Ingersoll (1987)) show that under a power utility function the value function of the problem is homogeneous in wealth, i.e.

$$V(W_{t+k\varphi}, \mathbf{z}_{t+k\varphi}) \equiv \max_{\{\boldsymbol{\omega}_{t+k\varphi}\}_{k=0}^{B-1}} E_{t+k\varphi} \left[\frac{W_{t+T}^{1-\gamma}}{1-\gamma}\right] = \frac{W_{t+k\varphi}^{1-\gamma}}{1-\gamma} Q(\mathbf{z}_{t+k\varphi}).$$

This fact suggests solving the problem by backward induction, starting at time $t + (B-1)\varphi$ and working to time t. The solution is approximate in the sense that it is based on the discretization of the space for the state vector $\mathbf{z}$ on a discrete grid of J points, say $\mathbf{z}^j$, $j = 1, ..., J$. In fact, at time $t + (B-1)\varphi$ the problem simplifies as $Q(\mathbf{z}^j_{t+(B-1)\varphi}) = 1 \; \forall j$, i.e. at the end of the investment horizon the investor ought to solve:

$$\max_{\boldsymbol{\omega}_{t+(B-1)\varphi}} E_{t+(B-1)\varphi} \left[\frac{W_{t+T}^{1-\gamma}}{1-\gamma}\right], \tag{6}$$

which is a simple buy-and-hold problem with horizon of φ periods. If the process of excess asset returns is correctly described by (2) and parameter uncertainty is ignored, then (6) has a simple solution that can be found using the results in Section 2.1, as the multivariate density for $\mathbf{z}$ remains normal $p\left(E_{t-1}[\mathbf{z}_{t+(B-1)\varphi,T}], Var_{t-1}[\mathbf{z}_{t+(B-1)\varphi,T}]\right)$ with moments given by (3) when $\mathbf{z}_{t+(B-1)\varphi} = \mathbf{z}^j_{t+(B-1)\varphi}$, $j = 1, ..., J$.[5] For instance, provided N is large enough, an approximate solution will be found by maximizing $N^{-1}\sum_{i=1}^{N}\left(W^i_{t+T}\right)^{1-\gamma}/(1-\gamma)$, where W^i_{t+T} is found on the simulated path i. Define then $Q(\mathbf{z}^j_{t+(B-1)\varphi}) \equiv$ argmax$_{\boldsymbol{\omega}_{t+(B-1)\varphi}} E_{t+(B-1)\varphi} \left[\frac{W_{t+T}^{1-\gamma}}{1-\gamma}\right]$ when $p\left(E_{t-1}[\mathbf{z}_{t,T}], Var_{t-1}[\mathbf{z}_{t,T}]\right)$ is conditional on $\mathbf{z}_{t+(B-1)\varphi} = \mathbf{z}^j_{t+(B-1)\varphi}$. Then for $j = 1, ..., J$, $\hat{\boldsymbol{\omega}}^j_{t+(B-2)\varphi}$ will be found by solving (once, more by simulation, using a multivariate normal conditional on $\mathbf{z}^j_{t+(B-2)\varphi}$)

$$\max_{\boldsymbol{\omega}^j_{t+(B-2)\varphi}} E_{t+(B-2)\varphi} \left[\frac{W_{t+(B-1)\varphi}^{1-\gamma}}{1-\gamma} Q(\mathbf{z}^j_{t+(B-1)\varphi})\right],$$

thus yielding J new values, $Q(\mathbf{z}^j_{t+(B-2)\varphi})$ $j = 1, ..., J$. The process is to be continued until $t+(B-B)\varphi = t$, i.e. until a vector $\hat{\boldsymbol{\omega}}^j_t$ $j = 1, ..., J$ emerges from expected utility maximization. By construction, each $\hat{\boldsymbol{\omega}}^j_t$ is matched to a $\mathbf{z}^j_t$. Although in general the observed $\mathbf{z}_t$ differs from $\mathbf{z}^j_t$ on the grid, simple interpolation algorithm will then be used to determine $\hat{\boldsymbol{\omega}}_t$ using the two closest values of $\hat{\boldsymbol{\omega}}^j_t$. For the calculations that follow, we use $J = 15$ discretization points and a number of Monte Carlo trials $N = 100,000$.

We also incorporate parameter uncertainty in a framework with dynamic rebalancing. In principle, this is a formidable problem as an investor ought to rationally recognize that in the future $-$ based on future realizations of $\mathbf{z}$ $-$ she will revise her posterior density for the coefficients in $\boldsymbol{\theta}$, and that such revisions will influence her portfolio choices and hence the stochastic process for wealth. Essentially, parameter uncertainty in a dynamic setting turns $\boldsymbol{\theta}$ into an additional vector of state variables. However, similarly

---

[5]This means that conditional moments have to be calculated by initializing $\mathbf{z}_{t+(B-1)\varphi}$ to $\mathbf{z}^j_{t+(B-1)\varphi}$ on the grid, $j = 1, ..., J$:

$$\begin{aligned} E_{t-1}[\mathbf{z}_{t+(B-1)\varphi,T}] &= \varphi\boldsymbol{\mu} + (\varphi - 1)\Phi\boldsymbol{\mu} + ... + \Phi^{\varphi-1}\boldsymbol{\mu} + (\Phi + \Phi^2 + ... + \Phi^{\varphi})\mathbf{z}^j_{t+(B-1)\varphi} \\ Var_{t-1}[\mathbf{z}_{t+(B-1)\varphi,T}] &= \Sigma + (\mathbf{I} + \Phi)\Sigma(\mathbf{I} + \Phi)' + ... + (\mathbf{I} + \Phi + ... + \Phi^{\varphi-1})\Sigma(\mathbf{I} + \Phi + ... + \Phi^{\varphi-1})', \end{aligned}$$

to Gennotte (1986) and Barberis (2000), we simplify the problem in the following way: we suppose that although the investor acknowledges that she is uncertain about model parameters, she ignores future revisions of her posterior density for the parameters themselves. Once this simplifying assumption is in place, the same dynamic programming, discretization-based approach followed in the classical case can be implemented in the Bayesian framework, with the only difference that instead of approximating expectations by simulation under a multivariate normal with moments described by (3), Monte Carlo methods are applied drawing from an appropriate joint predictive density for cumulative returns over the intervals $[t, t+\varphi]$, $[t+\varphi, t+2\varphi]$, ..., $[t+(B-1)\varphi, t+B\varphi]$.

3. Estimation Results

3.1. *The Data*

Since one of the contributions of this paper is to expand the asset menu to real estate, we start by providing a sense for what the related data issues may be. Real estate performance can be measured using two types of indices. *Direct* indices are derived from either transaction prices or the appraised value of properties, while *indirect* indices are inferred from the behavior of the stock price of property companies that are listed on public exchanges. Indirect real estate index returns normally show higher volatility than direct returns, and − being subject to similar common market factors − tend to display higher correlations with standard stock index returns. In this sense, indirect indices are biased towards a finding of simultaneous correlation of real estate returns with financial returns. On the other hand, the reliability of transaction-based, direct indices is often made problematic by both the fact that properties may be wildly heterogeneous and by the poor transparency of transaction conditions. Additionally, direct, appraisal-based data are known to be affected by many biases (see Giliberto 1988; 1991). For instance, the standard deviation of appraisal indices has been shown to represent (by construction) a downward biased estimate of the true value.[6]

Confronted with these pros and cons of direct vs. indirect real estate indices, our paper employs an indirect index of the 40 most liquid property companies in Europe, the monthly European Liquid 40 Index prepared by the European Public Real Estate Association (EPRALiquid40). EPRA indices track the performance of stocks of real estate companies engaged in the ownership, disposure and development of income-producing real estate. We select this index among the many indirect alternatives available (even including the EPRA universe only) with the intent of maximizing the homogeneity of the asset classes studied in the paper in terms of transaction costs − in the sense that only for the most liquid real estate property companies an assumption of similar frictions relative to stocks or bonds is a sensible one.

The remaining assets entering the investment opportunity set are European short-term deposits, long-term bonds, and stocks. We use monthly data, for the period March 1993 - March 2004, for a total of 132 observations. The sample period seems to be well balanced, including at least one complete bull (the 1990s) and one complete bear (2000-2002) market cycles. In particular, stock returns are calculated from the Datastream European price index; the Citigroup European World Government Bond Index is used to capture the behavior of European bond returns for maturities exceeding five years; money-market, short-term instruments are proxied by the JP Morgan Euro cash (three-month) index obtained from daily

---

[6]A comparison of direct appraisal-based vs. indirect indices is provided by Giliberto (1993) for the U.S. and by Maurer et al. (2003) for Germany.

compounding of the Euribor (the rate on euro-currency deposits) with a three-month maturity. [7] All indices are continuously compounded (i.e. $R_t^j \equiv \ln P_t^j - \ln P_t^j$, $j = s, b, r$), total return market-capitalization indices, including both capital gains and income return components, expressed in euros. Excess returns are calculated by deducting short-term cash returns from R_t^j. We express returns in real terms, which seems appropriate given the long-run horizons we are interested in. However, since our econometric models will concern excess returns, this is equivalent to deal with real short-term interest rates, calculated as a difference between nominal yields and the monthly rate of change in a European harmonized index of consumer prices (EICP). Finally, the predictor x_t is identified with the dividend yield on the Datastream stock price index.[8]

In table 1 we present summary statistics for the variables (returns and dividend yields) discussed above. Over our sample period, the European real estate market fails to be 'dominated' (in mean-variance terms) by the stock market, in spite of the euphoria characterizing the so-called New Economy period of 1995-2000: real estate investments performed similarly to equities in mean terms (around 6.5 percent per year in excess of the riskless real rate), but were strictly less volatile than stocks (their annualized standard deviation is 13% vs. 17% for equities).[9] As one would expect, bonds have been less profitable (4%) but also less volatile (5.4%) than stocks and real estate. However an annualized real return of approximately 7% remains remarkable for bonds and is explained by the declining short-term interest rates during the 1990s. The third and fourth central moments for excess returns stress that the series are moderately left skewed and leptokurtic. The Jarque-Bera test for normality induces to strongly reject the null hypothesis of normality for excess stock returns only, mainly on account of their negative skewness. Rejections of normality are weak for excess bond and real short-term returns. Ljung-Box portmanteau tests are then applied to test the null hypothesis of no serial correlations in levels and squares of returns. As it is well known, serial correlation tests applied to squares are useful to detect volatility clustering, i.e. ARCH effects. European asset returns appear to be approximately serially uncorrelated (there is weak evidence for the short-term real rate only) and − with the exception of excess stock returns − also statistically independent, in the

---

[7]The EPRALiquid40 index is based on prices of European quoted property company shares in the following countries: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom.

The WGBI Europe is a market-capitalization weighted portfolio that tracks the performance of bonds issued in the countries above as well as Poland.

The Datastream index covers the stock markets in the countries above as well as those in Prague, Budapest, Bucarest, Moscow, Stockholm and Instanbul.

[8]This is suggested by empirical findings on the predictability of stock returns reported for instance in Bekaert and Hodrick (1992), Campbell and Shiller (1988), Fama and French (1988, 1989), Ferson and Harvey (1991), Goetzmann and Jorion (1993) and Kandel and Stambaugh (1996). Due to its high persistence coupled with the strong negative correlation between shocks to returns and shocks to the dividend yield, Lynch (2001) and Campbell, Chan, and Viceira (2003) find that the dividend yield generates the largest hedging demand among a wider set of predictor variables.

[9]Most of the evidence reports lower mean returns for real estate appraisal-based returns coupled with lower volatility relative to stocks both in the US (Ibbotson and Siegel (1984); Maurer, Reiner and Sebastian (2003)) and in the UK and Germany (Maurer, Reiner and Sebastian (2003)). A similar pattern emerges in

Maurer, Reiner and Rogalla (2003), using German indirect real estate indices.

sense that volatility clustering is weak.[10] Finally, the dividend yield displays all the usual features that have appeared in the North-American literature, namely it is a playkurtic and highly serially correlated series.[11]

These findings of rare departures from (univariate) normality and of weak serial correlation and volatility clustering effects make it plausible to think that a simple, homoskedastic Gaussian VAR(1) model as in (2) may be not be badly misspecified. Table 2 provides a first impression of the main features of a multivariate Gaussian model, at least in terms of simultaneous correlations. The table shows that the performance across the three markets is only weakly correlated, with a peak correlation coefficient of 0.56 between excess stock and real estate returns.[12] Under these conditions, there is wide scope for portfolio diversification across financial and real assets. In particular, excess bond returns are characterized by insignificant correlations relative to both stock and real estate, and therefore we should expect a large demand for bonds (see e.g. Campbell and Viceira (2002)) for purely hedging reasons.

3.2. *Predictability in Excess Asset Returns*

Table 1 reveals only weak evidence of predictability in asset returns, essentially connected to the fact that past excess stock returns forecast future, subsequent excess stock returns. However, a complete picture of the degree to which risk premia may be forecast can be derived only from the estimation of (2). Results are reported in table 3 for the case in which classical estimation methods are employed; robust t-stats are reported in parenthesis, under the corresponding point estimates. There is strong statistical evidence that lagged excess returns on real estate forecast subsequent excess bond returns, and that lagged excess stock returns forecast excess real estate returns. In particular, an increase in time t real estate returns forecasts a decline in bond returns at time $t+1$, while a time t increase in stock returns predicts a time $t+1$ increase in real estate returns.[13] This is consistent with stories according to which real estate markets adjusts to the equity market swings (e.g., booming prices cause wealth effects that spread over the real estate market) and in which real property forecasts future economic conditions (hence adjustments in the slope of the term structure of interest rates), see e.g. C.H Liu and J. Mei (1992) and H. Liao and J. Mei (1998).

We also find weaker evidence of the forecasting power of the dividend yield for European excess stock returns: although the coefficient is interesting for its magnitude (it implies that a one standard deviation increase in the dividend yield forecasts an increase in stock returns by 74 basis points), its statistical significance is marginal (the corresponding p-value is 15% only). Therefore, while on US (and UK, see Guidolin and Timmermann (2005)) data there is robust evidence that the dividend yield forecasts asset

---

[10]The serial correlation for short real rates is stronger and fits the common perception of near-integration for short term rates. The absence of volatility clustering (ARCH) at monthly frequencies is not surprising. Among others, Campbell and Viceira (1999, 2002), Barberis (2000), and Lynch (2001) all ignore ARCH effects and focus on the predictability of the risk premium only.

[11]The evidence of volatility clustering in table 1 for the dividend yield is illusory: once an AR(1) model is fitted to the dividend yield, the residuals appear to be approximately homoskedastic.

[12]However this measure is likely to incorporate an upward bias as almost all the equities covered by the EPRA40 enter in the Datastream European Stock Index. However, the latter covers 2,431 different stocks, so the effect on the correlation coefficient is likely of second-order magnitude.

[13]The effects are economically important: a one standard deviation increase in monthly excess real estate returns forecasts a 41 basis points decline in excess bond returns; a one standard deviation increase in monthly excess stock returns predicts an 85 basis point increase in excess real estate returns.

12

returns, such evidence is more doubtful on European data, at least for our sample period. One last remark concerns the MLE estimates of the (simultaneous) covariance matrix of excess asset returns residuals:

$$\hat{\Sigma} = \begin{bmatrix} 0.168 & 0.068 & 0.540 & -0.951 \\ 0.001 & 0.052 & 0.228 & -0.055 \\ 0.011 & 0.001 & 0.117 & -0.574 \\ -0.001 & -0.000 & -0.000 & 0.005 \end{bmatrix},$$

where the elements on and below the main diagonal are annualized volatilities and pairwise covariances, respectively, while the elements above the main diagonal are pairwise correlations. In particular, notice the relatively high correlation (0.54) between excess stock and real estate returns residuals, indication that shocks unexplained by the VAR(1) model tend to appear simultaneously for the stock and real estate markets. Moreover, the simultaneous sample correlations between news affecting stock and real estate markets and news involving the dividend yields are negative and significant (-0.95 and -0.57, respectively): when shocks hit the dividend yield (i.e. the way in which company profits are transformed into asset prices), our estimates imply a contemporaneous negative effect on excess stock and real estate returns. Such findings are ubiquitous in the literature analyzing US financial returns data (see e.g. Barberis (2000)), but they are novel with reference to European and − more important − real estate markets. As a result, the finding that the dividend yield is only a weak predictor of future excess asset returns is only apparent, in the sense that we do find that news to the dividend yield have important and precisely estimated correlation structure with financial return news. As we will see in Section 4, these features have major portfolio choice implications.

In table 4, we repeat the econometric analysis employing Bayesian estimation techniques that − as stressed in Section 2.2 − allow us to derive a joint posterior density for the 'coefficients' collected in $\boldsymbol{\theta}$. The tails of this density also measure the amount of estimation risk present in the data. In fact, table 3 reports the means of the marginal posteriors of each of the coefficients in $\mathbf{C}$ along with the standard deviation of the corresponding marginal posterior, which gives an idea of its spread and therefore of a measure of the uncertainty involved. As typically found in the financial literature, the posterior means in table 4 only marginally depart from the MLE point estimates in table 3.[14] Standard errors are relatively high, confirming the presence of important amounts of estimation risk in this application. However, it remains clear that the effects of lagged excess stock returns on real estate returns and of lagged real estate on bond returns are characterized by tight posteriors which suggest a non-zero effect. Also in this case, the effect of the dividend yield on subsequent returns seems to be rather strong in terms of location of the posterior density, although the tails are thick enough to cast some doubts on the precision with which effects can be disentangled. For completeness, we also report the posterior means and standard deviations

---

[14]For instance, the most important difference (in absolute terms) concerns the coefficients measuring the effects of the lagged dividend yield on excess stock returns, which goes from 1.262 under MLE to a Bayesian posterior mean of 1.256.

(in parenthesis) for Σ:

$$
\Sigma^{Post} = \begin{bmatrix}
\underset{(0.063)}{0.175} & 0.068 & 0.540 & -0.951 \\
\underset{(0.001)}{0.001} & \underset{(0.019)}{0.054} & 0.228 & -0.055 \\
\underset{(0.002)}{0.012} & \underset{(0.001)}{0.002} & \underset{(0.044)}{0.122} & -0.574 \\
\underset{(0.000)}{-0.001} & \underset{(0.000)}{-0.000} & \underset{(0.000)}{-0.000} & \underset{(0.002)}{0.005}
\end{bmatrix}.
$$

Volatilities and covariances are again reported in annualized terms, while the coefficients above the main diagonal are correlations. Clearly, most elements of Σ have very tight posteriors, and importantly all the implied correlations are identical (to the third decimal) to those found under MLE. Therefore also in this case, the seemingly weak predictability from dividend yields must be supplemented by the evidence of strong and precisely identified simultaneous correlations between the dividend yield and excess stock and real estate returns.

4. **Optimal Asset Allocation with Real Estate**

4.1. *Baseline, Buy-and-Hold Classical Weights*

We start with the simplest of the portfolio allocation exercises: we consider an investor who commits her initial, unit wealth for T years and who ignores parameter uncertainty, i.e. replaces the unknown VAR coefficients with their MLE parameter estimates, ignoring the corresponding standard errors. Initially we do not impose any restrictions on the predictability model, i.e. report results for a full VAR(1) model (see Campbell, Chan, and Viceira (2003)). Throughout, we set $\mathbf{z}_{t-1}$ to the full-sample mean values for excess returns and the dividend yield.[15]

Figure 1 reports optimal portfolio weights for horizon between 1 month and 10 years, which is assumed to represent the typical long-horizon objective. The exercise is repeated for three alternative values of the coefficient of relative risk aversion, $\gamma = 2$, 5, and 10. Clearly, the latter represents rather a high risk aversion coefficient and we expect that its adoption would completely bias our results against the riskier assets and in favor of relatively safe bonds and cash investments.

Modeling linear predictability patterns seem to make a noticeable difference for portfolio choice, in the sense that short- and long-horizon portfolio weights radically differ for all levels of γ. For instance, for $\gamma = 2$ (not an atypical choice in the asset allocation literature, see Brennan et al. (1997) and Barberis (2001)), the 1-month allocation is 57% in real estate, 41% in stocks, and 2% in bonds; however, under a 10-year horizon, the optimal weights become 50%, 42%, and 8%, respectively. This means that there is an important (8%) shift out of the riskier assets and in favor of the safe bonds. Remarkably, cash is never demanded, i.e. the presence of relatively safer bonds and predictability combine into an asset menu that can satisfy the risk-return trade-offs of even highly risk-averse investors ($\gamma = 10$) without involving zero-variance assets.

Generally speaking, and consistently with results in the literature (see Kandel and Stambaugh (1996) and Barberis (2000)), the weight invested in riskier assets appears to be a monotone increasing function of

---

[15]Our results therefore qualify as 'simulations' representative of the average properties of our estimated model rather than actual end-of-sample (March 2004) portfolio recommendations.

the investment horizon. This is particularly evident for $\gamma = 5$ and 10, although the schedule characterizing real estate is rather flat. The explanation is that predictability in the risk premium makes risky assets less risky than what is conveyed by their standard deviations; this effect becomes stronger, the longer the horizon an investor has over which to exploit the forecastability patterns. The best way to see this is to focus on the multiperiod conditional variance of (say) the risky stocks implied by our VAR(1) framework:

$$
\begin{aligned}
Var_t[R_{t,2}^s] \quad &= Var_t\left[\sum_{k=1}^{2}(R_{t+k}^s + r^f)\right] \\
&= Var_t\left[\sum_{k=0}^{1}\left(r^f + \mu^s + \phi^{ss}R_{t+k}^s + \phi^{sb}R_{t+k}^b + \phi^{sr}R_{t+k}^r + \phi^{sx}x_{t+k} + \sigma_s^2\varepsilon_{t+k}^s + \sigma_{sb}\varepsilon_{t+k}^b + \sigma_{sr}\varepsilon_{t+k}^r + \sigma_{sx}\varepsilon_{t+k}^x\right)\right] \\
&= 2\sigma_s^2 + (\phi^{sb})^2\sigma_b^2 + (\phi^{sr})^2\sigma_r^2 + (\phi^{sx})^2\sigma_x^2 + 2\phi^{sb}\sigma_{sb} + 2\phi^{sr}\sigma_{sr} + 2\phi^{sx}\sigma_{sx},
\end{aligned}
$$

which can be compared to $Var_t[R_{t,1}^s] = Var_t[R_{t+1}^s] = \sigma_s^2$. Given our MLE estimates, it happens that

$$
\frac{Var_t[R_{t,2}^s]}{2Var_t[R_{t,1}^s]} < 1,
$$

i.e. that conditional variance grows at a slower rate than the horizon because the following terms are negative:

$$
(\phi^{sb})^2\sigma_b^2 + (\phi^{sr})^2\sigma_r^2 + (\phi^{sx})^2\sigma_x^2 + 2\phi^{sb}\sigma_{sb} + 2\phi^{sr}\sigma_{sr} + 2\phi^{sx}\sigma_{sx} < 0.
$$

This results in an increasing allocation to stocks as T gets larger.[16] The economic interpretation is that when dividends fall unexpectedly (i.e. they are hit by some adverse shock), $\sigma_{sx}, \sigma_{rx} < 0$ imply that the news will be likely accompanied by a positive, *contemporaneous* shock to excess stock and real estate returns. On the other hand, since $\phi^{sx}, \phi^{rx} >> 0$, a currently diminished dividend yield forecasts *future* lower risk premia on stocks and real estate. Hence the parameter configuration implied by the data leads to a built-in element of negative serial correlation. It is well known that processes characterized by negative serial correlations are less volatile in the long- than in the short-run, due to mean-reversion effects. These effects may be of first-order importance even when the standard errors associated to many of the coefficients in (2) are high and the estimated VAR(1) coefficients relatively small, as long as adequate covariance loadings come through $\hat{\Sigma}$ (in this case, $\hat{\sigma}_{sx} << 0$).

   Finally, figure 1 shows two another key results. First, the optimal allocation to bonds is generally monotone decreasing with T. However, this is fully explained by the statistical properties of the vector $\mathbf{z}_t$

---

[16]In fact, replacing unknown coefficients with their MLE estimates gives:

$$
(\hat{\phi}^{sb})^2\hat{\sigma}_b^2 + (\hat{\phi}^{sr})^2\hat{\sigma}_r^2 + (\hat{\phi}^{sx})^2\hat{\sigma}_x^2 + 2\hat{\phi}^{sb}\hat{\sigma}_{sb} + 2\hat{\phi}^{sr}\hat{\sigma}_{sr} + 2\hat{\phi}^{sx}\hat{\sigma}_{sx}
$$
$$
= (0.07)^2 \cdot 0.0002 + (0.07)^2 \cdot 0.0011 + (1.26)^2 \cdot 1.7e^{-06} +
$$
$$
2 \cdot 0.07 \cdot 5e^{-05} + 2 \cdot 0.07 \cdot 0.0009 - 2 \cdot 1.26 \cdot 6e^{-05} = -9.1e^{-06} < 0.
$$

This number is only apparently negligible: for instance it is more than 5 times larger than the unexplained variance of the dividend yield.

   Similar calculations may be performed for real estate. Since $(\hat{\phi}^{rs})^2$, $(\hat{\phi}^{rb})^2$, $\hat{\sigma}_x^2 \simeq 0$, while $\hat{\sigma}_{rx}$ implies a correlation of $-0.57 < 0$, $\hat{\phi}^{rx} = 0.65 > 0$, and $\hat{\phi}^{rs} = 0.14 \simeq 0$ and $\hat{\phi}^{rs} = 0.17 \simeq 0$, then also for real estate, risk grows slower than the investment horizon. However in this case the effect is smaller, which supports our findings of relatively flat investment schedules for real estate assets.

we have modeled in Section 3 (see table 3). In particular, notice that bonds display a negligible covariance with the dividend yield ($\hat{\sigma}_{bx}$ implies a correlation of -0.06 only) and actually positive covariances with shocks to stocks and real estate. This means that news affecting the variables (for us, the dividend yield) characterizing investment opportunities will essentially leave current, realized bond returns unchanged and then forecast future changes in risk premia of the same sign as the news. Therefore bonds will either be characterized by a variance that grows at the same rate as T, or even by increasing risk over longer and longer time horizons. This makes them rather safe assets for $T = 1$-2 years, although they become increasingly risky and less attractive the longer the horizon.

Figure 1 shows some non-monotonic shapes, especially for $\gamma = 2$ and with reference to the demand for bonds, which seems to increase between $T = 1$ month and 1 year and smoothly declines for the reasons explained above. This feature is fully explained by the simulation nature of the results in figure 1, i.e. the fact that in (2) we initialize $z_{t-1} = t^{-1} \sum_{i=1}^{t} z_i$. Notice that while the full-sample means for real returns are

$$[0.0079 \ 0.0060 \ 0.0080 \ 0.0263]',$$

the estimated VAR(1) model implies unconditional means of:

$$E[\mathbf{z}_t] = (\mathbf{I}_4 - \hat{\Phi})^{-1} \hat{\boldsymbol{\mu}} = [0.0004 \ 0.0047 \ 0.0019 \ 0.0233]'.$$

Although these values make sense, as they imply positive mean real returns (and risk premia), notice that all mean returns are systematically below the full-sample means. Therefore our initialization corresponds in fact to an economy in a period of above-mean returns and dividend yields. Since most of the coefficients in $\hat{\Phi}$ are positive (we do notice only one exception), initially high returns imply some persistence in the assumed bull market conditions. In particular, it is easy to check through simulation experiments, that using $t^{-1} \sum_{i=1}^{t} x_i = 0.0263 > E[x_t] = 0.0233$ implies forecasts of high, above-mean excess returns on stocks and bonds at horizons up to 6-12 months. This explains a transient effect by which the demand for stocks and real estate is particularly high at a 1-month horizon, and the demand for bond relatively modest. Since an investor ought to aggressively exploit these predicted effects in the mean, the lower her risk aversion, such non-monotonicities are more obvious for $\gamma = 2$ and 5.[17]

Another way to quantify the importance of predictability in determining rational portfolio choices when real estate belongs to the asset menu can be derived by comparing the results in figure 1 with those one can calculate assuming a simple model without any predictability patterns, i.e.

$$\mathbf{z}_t = \boldsymbol{\mu} + \boldsymbol{\epsilon}_t \quad \boldsymbol{\epsilon}_t \ \text{i.i.d.} \ N(\mathbf{0}, \Sigma), \tag{7}$$

with constant covariances as well as risk premia. We find that long-run asset allocations in the presence of predictability are rather different than those obtained under the IID benchmark.[18] For instance, when $\gamma = 5$, the percentages to be invested in bonds are 42% vs. 54% under no predictability, 32% vs. 18% for stocks, and 26 vs. 28% for real estate.[19] Hence, also in this metric predictability implies a shift

---

[17]This effect can be quantified in a diminished bond weight of roughly 20% when $\gamma = 2$ and of 5% when $\gamma = 5$.

[18]As discussed in the Introduction, under (7) the optimal portfolio weights become independent of the horizon. In the following we compare asset allocations under IID with those obtained under (2) for the $T = 10$ years case.

[19]The corresponding numbers are (for bonds, stocks, and real estate, respectively) 8-2%, 42-39%, 50-59% when $\gamma = 2$; 54-72%, 28-11%, 18-17% when $\gamma = 10$.

out of bonds by 10% or more and into the riskier assets, mostly stocks. Obviously, the interpretation is that the assets whose long-run risk is mitigated by the mean-reversion effects implied by (2) are in lower demand under IID, the opposite for assets that are riskier the longer the horizon. Overall, it seems that ignoring predictability altogether would in fact lead to inappropriate asset allocations. Section 5.1 further investigates the welfare losses resulting from disregarding predictability.

In conclusion, a classical analysis obtained by 'plugging-in' our parameter estimates from (2) implies that real estate ought to have an important role in simple, buy-and-hold portfolio choices, especially when investors have a long-run perspective. Depending on the assumed coefficient of relative risk aversion, we have found optimal real estate weights between 18 and 50% of the available wealth. These results are substantially higher than typical findings in the North American literature based on simple mean-variance static portfolio theory: for instance, using annual US data, Karlberg, Liu, and Greig (1996) and Ziobrowski, Caines, and Ziobrowski (1999) find that the optimal fraction of wealth to be allocated in real estate is around 9%, in the range of 3-15%; in fact, for moderate enough risk aversion levels, most of the investor's wealth ought to be assigned to stocks, while bonds carry a negligible weight.[20]

4.2. *Parameter Uncertainty Under Buy-and-Hold*

We next proceed to calculate optimal portfolio weights for the case in which the investor adopts a Bayesian approach. For simplicity, we deal with the buy-and-hold case first.[21]

Figure 2 reports portfolio weights as a function of T. The effects of estimation risk manifest themselves with varying intensity at two distinct levels. On one hand, accounting for parameter uncertainty does not radically change the structure of portfolios, especially for low risk aversion and short horizons. Optimal weights still depend (as one ought to expect) from the coefficient of relative risk aversion. For instance, for $\gamma = 2$, the 1-month allocation is *identical* to the one derived ignoring parameter uncertainty as an investor who is not strongly risk averse will not be affected by additional estimation risk, over short-horizons. The 10-year allocation is instead 31% in real estate, 34% in stocks, and 35% in bonds, i.e. very close to an equally-weighted portfolio, [22] different from the 50%, 42%, and 8% weights found in Section in 4.1 for this case. The only major difference in portfolio composition obtains under $T = 10$ years and high risk-aversion, and consists in the appearance of positive weights invested in the riskless asset , as much as 40% . This makes sense as cash is not only the safest among the assets but also the only one that completely escapes the additional uncertainty created by estimation risk.

On the other hand, important modifications occur in the structure of the investment schedules as a function of the horizon: while a classical investor will be characterized by weights to stocks and real estate increasing with the investment horizon, when parameter uncertainty is taken into account these schedules become either flat or weakly monotone decreasing. For instance, when $\gamma = 5$ the allocation to real estate declines from 26% at 1-month to 16% at 10 years; the allocation to stocks is essentially flat at 20-21%. Notice that the uncertainty deriving from estimation risk compounds over time, i.e. if it is difficult to assess whether the predictability patterns help in asset allocation over a 1-month horizon, this

---

[20]However Liand and Webb (1996) do estimate higher weights.

[21]Also in this case, we set $\mathbf{z}_{t-1}$ to the full-sample mean values for excess returns and the dividend yield.

[22]Gold (1993) reports a similar results for US optimal portfolios in a mean-variance framework in which investors protect themselves against estimation risk by bootstrapping real estate returns.

effect is magnified for longer planning periods. This means that the contrasting effects of the 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.[23] Additionally, it seems that the effects of parameter uncertainty are equally strong for bonds as for the riskier assets, stocks and real estate. Notice in fact that the posterior standard errors characterizing all the bond coefficients in table 4 are rather high. As a result, bond investment schedules turn now strongly downward sloping, in fact giving way to a positive demand for short-term cash deposits for T in excess of 4-5 years.

Once more, we provide a preliminary (see also Section 5) measure of the importance of predictability under parameter uncertainty by calculating optimal portfolio weights under the no-predictability benchmark (7) and by comparing results to (2). In practice, such an exercise is also helpful to document and quantify the pure effects of parameter uncertainty on optimal portfolio weights.[24] Figure 3 displays results through the usual set of plots. Although they generally confirm the importance of real estate for portfolio choice purposes, these plots also put further emphasis on our discussion of the separate effects of predictability vs. estimation risk. In a system with no predictability but in which there is parameter uncertainty, the investment schedules for both stocks and real estate are monotonically decreasing, i.e. as initially guessed estimation risk is compounded and magnified by longer and longer investment horizons. Interestingly, the demand for cash is completely absent, also for investors with high risk aversion. Moreover, the bond investment schedules are now upward sloping, which confirms that there exists a differential of estimation risk that favors bonds over riskier instruments. Finally, by comparing figures 2 and 3, one has an implicit proof of the fact that the estimation risk may be particularly important when the number of estimated coefficients grows, i.e. when an investor tries to capture the uncertain predictability patterns using some parametric framework. In particular, the schedules for bonds appear to change slope and a positive demand for cash investments appears when moving from figure 3 to 2. This is consistent with larger degrees of parameter uncertainty plaguing (2) vs. (7) and with the existence of first-order effects from the high posterior standard errors characterizing ϕ^{bs}, ϕ^{bb}, ϕ^{br}, and ϕ^{bx}.

In conclusion, adding back parameter uncertainty to the asset allocation problem changes a few of the results found in Section 4.1, but leaves the overall picture intact: real estate is an important class that − when predictability is measured and put to use through a Bayesian approach − ought to receive an optimal long-run weight between 10 and 30%, depending on the assumed coefficient of relative risk-aversion.[25]

---

[23]In fact the real estate schedules turn downward sloping as a reflection of the weaker predictability effects found for this asset class in Section 4.1.

[24]The coefficients in μ have posterior means of 0.005, 0.003, and 0.005 (the dividend yield is irrelevant in this case), and standard errors of 0.005, 0.001, 0.003. Details on the posterior distribution of the coefficients in Σ are available upon request.

[25]Notice that when estimation risk is accounted for, our implications become rather implausible for $\gamma = 10$, when roughly 75% of the available wealth ought to be invested in either cash or bonds when the horizon is 10 years. Clearly, such implications are at odds with the actual structure of the majority of observed managed portfolios, and simply reject $\gamma = 10$ as a plausible specification for preferences. If $\gamma = 2$ and 5 are the most plausible relative risk aversion levels, then the optimal real estate weight climbs to 20-30% at least. Similar conclusions apply to equity investments.

4.3. *Dynamic Rebalancing*

We assess the effects of dynamic rebalancing on the optimal real estate weights by computing portfolio allocations for the most realistic case in which $\varphi = 12$ months, i.e. adjustment to the asset allocation are admitted every year and the investors takes that into account even when formulating long-horizon strategies with $T >> \varphi$. In particular, we start with the case in which the statistical model is simplified to Barberis' (2000) framework,

$$
\begin{bmatrix} R_t^s \\ R_t^b \\ R_t^r \\ x_t \end{bmatrix} = \begin{bmatrix} \mu^s \\ \mu^b \\ \mu^r \\ \mu^x \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & \phi^{sx} \\ 0 & 0 & 0 & \phi^{bx} \\ 0 & 0 & 0 & \phi^{rx} \\ 0 & 0 & 0 & \phi^{xx} \end{bmatrix} \begin{bmatrix} R_{t-1}^s \\ R_{t-1}^b \\ R_{t-1}^r \\ x_{t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_t^s \\ \epsilon_t^b \\ \epsilon_t^r \\ \epsilon_t^x \end{bmatrix}, \tag{8}
$$

i.e. when predictability is restricted to come from the dividend yield only and − consistently with available results on North American markets (see e.g. Campbell, Lo, and MacKinlay (1997)) − past excess returns fail to predict future returns.[26] Given that this model simplifies the state space to one variable only (the dividend yield), for the calculations that follow we use $J = 20$ discretization points for increased accuracy. Once more, the initial dividend yield is set to its full-sample mean.

Column (a) of Figure 4 shows results for the classical case, when parameter uncertainty is ignored. A comparison with Figure 1 reveals only two important changes. First, rebalancing makes the resulting portfolio investment schedules (as a function of the horizon) rather flat. Careful analysis of the plots reveals in fact that differences between optimal weights at $T = 5$ and at $T = 10$ are generally negligible. This makes sense as two investors with either a five or a ten-year horizon that rationally discount the possibility to make (unrestricted) changes to their optimal portfolios in only 12 months are unlikely to drastically differ in their current portfolio choices, as all they know about asset returns between month $t + \varphi + 1$ and $T = 60$ or 120 months later is that their joint dynamics is well-described by the unconditional distribution. The fact that rebalancing tends to flatten optimal investment schedules for $T >> \varphi$ has been observed already by Brandt (1999) and Guidolin and Timmermann (2004, 2005) in related applications. Second, differences between the rebalancing and buy-and-hold cases are modest but exist: in general, rebalancing opportunities when the system is initialized at the full-sample mean dividend yield tend to penalize the investments in the riskier assets, stocks and real estate.[27] However, the demand for cash investments remains nil for all levels of risk aversion. It seems that a dividend yield equal to its full-sample mean may suggest some caution to a rational investor, who then waits for an improvement in the investment opportunities by going longer in the relatively safe (especially over a relatively short period) bonds. In general, even when rebalancing is admitted, real estate remains an important asset class, receiving a weight always close to 20% or higher for all horizons and risk-aversion coefficients. On the other hand, when rebalancing is possible, an investor aggressively tries to exploit the available information: for instance, when $\gamma = 5$ and the horizon is long (ten years), the real estate weight is only 10% when $dy_t = 2.2\%$, it increases to 21% for $dy_t = 2.6\%$, and

---

[26]In this case, the dividend yield predicts with some accuracy (i.e. the related coefficients are significant at 10%) both stocks and bonds, and the higher the dividend yield, the higher the risk premia on assets. The dividend yield remains considerably persistent and there is substantial negative correlation between shocks to excess stock and to real estate returns and dividend yield shocks. Complete results are available upon request.

[27]For instance, at $T = 3$ years, the optimal weights in real estate under rebalancing are 42% (vs. 43% under buy-and-hold) for $\gamma = 2$, 21% (vs. 23%) under $\gamma = 5$, and 14% (vs. 16%) under $\gamma = 10$.

becomes 28% for $dy_t = 3.2\%$. Importantly, all these values of the dividend yield were observed over our sample period.[28]

Column (b) of Figure 4 reports instead results for the Bayesian case. For this exercise, the number of simulations is increased from 100,000 to 200,000 independent trials to accommodate the complex drawing loops needed to approximate the Bayesian predictive density and reduce randomness deriving from the use of simulations. Once more, we start by entertaining the simplified model in (8). Introducing rebalancing opportunities in a Bayesian framework has two major effects. First, it brings down to zero the demand for 3-month deposits: while under buy-and-hold and $\gamma \geq 5$ the optimal cash weight was positive for horizons exceeding 4-5 years, rebalancing does flatten the demand for the less risky asset to nothing. This is consistent with the idea that rebalancing makes a rational investor more aggressive (i.e. less risk-averse) as the possibility to switch out of risky portfolio positions well ahead of the planning horizon exists.[29] A similar, albeit less strong effect is obtained with reference to bonds, for which the long-run demand declines from 35-45% under buy-and-hold to less than 20% (zero assuming $\gamma = 2$). Second, rebalancing shifts up the demand for stocks and (to a lesser extent) real estate and makes the corresponding investment schedules upward sloping for horizons in excess of 1 year (i.e. those affects by our annual rebalancing assumption). Also in this case, most of the effects occur at horizons between 1 and 4 years, as for long-run investors schedules are typically flat, as observed before.

5. Welfare Cost Analysis

Even though Section 4 has provided abundant and compelling evidence that (i) real estate enters the optimal long-run portfolio composition with non-negligible weights when asset returns are predictable, and that (ii) predictability has important effects on portfolio weights, it remains very difficult to reach definitive conclusions on the economic importance of the implied effects for expected utility maximization. Therefore we follow Ang and Bekaert (2002), Ang and Chen (2002), Guidolin and Timmermann (2004, 2005), and Lynch (2001), and obtain estimates of the welfare implications of restricting the asset allocation problem in any of its dimension, like the breadth of the asset menu or the richness of the statistical model used to describe the multivariate process of asset returns.

Call $\hat{\boldsymbol{\omega}}_t^R$ the vector of portfolio weights obtained imposing restrictions on the portfolio choice problem. For instance, $\hat{\boldsymbol{\omega}}_t^R$ may be the vector of optimal asset demands when the investor is forbidden from investing in real estate. We aim at comparing the investor's expected utility under the unrestricted models entertained in Section 4 − and leading to some optimal set of controls $\hat{\boldsymbol{\omega}}_t$ − to that derived assuming the investor is constrained to choose at time t the restricted optimum, $\hat{\boldsymbol{\omega}}_t^R$. Define now $V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t)$ the optimal value function of the unconstrained problem, and $V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t^R)$ the 'constrained' optimal value function. Since a restricted model is by construction a special case of a more general, unrestricted model, the following holds:

$$V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t^R) \quad V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t).$$

---

[28]Section 6.2 provides further analysis on the effects of changing predictor variables on optimal portfolio weights.

[29]This remark applies only when the simulations are initialized in (relatively) good market conditions, such as those associated with a mean, 2.6% dividend yield. For instance, when $dy_t = 1.2\%$ (i.e. prices are relatively high), the demand for cash by a 10-year horizon investor with $\gamma = 5$ is 68%, i.e. rebalancing makes it rational to wait on the sidelines and invest more than two-thirds of the initial wealth in the safe asset.

We compute the compensatory premium, π_t^R, that an investor with relative risk aversion coefficient γ is willing to pay to obtain the same expected utility from the constrained and unconstrained asset allocation problems as:

$$\pi_t^R = \left. \frac{V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t)}{V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t^R)} \right]^{\frac{1}{1-\gamma}} - 1. \tag{9}$$

The interpretation is that an investor, if endowed with an initial wealth of $(1 + \pi_t^R)$, would tolerate to be constrained to solve some kind of restricted asset allocation problem leading to $V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t^R)$ $V(W_t, \mathbf{z}_t; \hat{\boldsymbol{\omega}}_t)$ only. Several types of restrictions are analyzed in what follows. For simplicity, we limit ourselves to consider simple buy-and-hold strategies.[30]

5.1. *Cost of Ignoring Predictability*

We first proceed to calculate the π_t^R implied by forcing an investor to ignore predictability altogether, i.e. to pretend that (7) is a correctly specified model while we have evidence (see Section 3) that in fact (2) is a better model. As observed in Section 5, this would lead to 'incorrect' portfolio choices. We now quantify the costs of these mistakes. These are presented in figure 5, once more as a function of the alternative levels of the horizon T and of γ. In particular, panel (a) refers to the classical case, when parameter uncertainty is ignored. Clearly, the implied welfare costs from model misspecifications strongly depend on the assumed coefficient of relative risk aversion, in the sense that they are the higher the lower γ is.[31] Generally speaking, especially if a Reader considers $\gamma = 2 - 5$ as the most plausible range for risk aversion, the implied welfare costs are far from negligibile, and in the case of long-horizon investors they range between 3 and 10% of the initial wealth. This means that a rational investor with $\gamma = 2$ would require a (riskless) annual increase in the returns generated by her portfolio in the order of approximately 95 basis points, for him to accept to work under and make portfolio decisions based on a misspecified IID model that disregards predictability altogether.

Panel (b) of figure 5 presents instead results for the Bayesian portfolio choice case, when estimation risk is incorporated in optimal decisions. In this case, the implied utility loss seems to be somewhat larger and especially rather independent of the coefficient of relative risk aversion. At long horizons, the welfare loss is always between 10 and 12% of the initial wealth, which corresponds to a required, (riskless) annual increase in the returns generated by the portfolio of approximately 100 basis points. The reasons of the larger utility losses under parameter uncertainty vs. the classical case in panel (a) are related to the fact that while in panel (a) $\hat{\boldsymbol{\omega}}_t^{IID}$ is anyway very close to the optimal, long-run $\hat{\boldsymbol{\omega}}_t$, in panel (b) this is not the case: as shown by figures 2 and 3, for large T portfolio choices are structurally different, especially because in figure 2 implies substantial cash investments that are not found when predictability is ignored. All in all, we interpret the evidence in figure 5 as consistent with the idea that ignoring predictability is associated with welfare losses of substantial magnitude.

---

[30]These clearly give lower bounds to the implied welfare costs, see e.g. Guidolin and Timmermann (2005). As a matter of fact, under dynamic rebalancing predictability gives an investor a chance to aggressively act upon the information on $\mathbf{z}_t$; therefore ignoring predictability when rebalancing is possible implies even higher (sometimes enormous) utility costs. A similar reasoning applies to restrictions on the asset menu, in the sense that depriving investors of useful assets hurts them the most the highest is the frequency with which they can switch in and out of the assets themselves.

[31]This finding is not new, see e.g. Campbell and Viceira (1999) and Guidolin and Timmermann (2005).

5.2. *Cost of Excluding Real Estate*

Although the evidence on the economic value of predictability is important, our key question remains open: How important is real estate in asset allocation problems, i.e. what is the welfare loss associated with restricting portfolio choice from involving real estate? Given the finding that ignoring predictability may be largely suboptimal, especially when estimation risk is taken into account, we analyze the utility losses when the investor exploits the evidence of predictability provided by (2).

As a first step, table 5 presents classical MLE estimates along with t-ratios for the case in which $n = 2$, i.e. the asset menu is limited to stock, bonds, and cash. Obviously, the expressions and results in Section 2.1 still apply (e.g. (3)). Table 5 stresses that on a restricted asset menu, the evidence of predictability becomes extremely weak, in the sense that only stock returns exhibit some degree of persistence and the dividend yield is a somewhat imprecise predictor of subsequent excess stock returns; the marginal distribution of excess bond returns seems to be well approximated by a Gaussian i.i.d.. In a way, real estate may help in asset allocation problems as it seems to be capable to bring out more precise predictability structure in risk premia. The left column of plots in figure 6 shows 'classical' asset allocation results under this restricted asset menu. Results are consistent with the general patterns isolated in Section 4.1: the demand for stocks increases with the investment horizon as their riskiness actually declines thanks to the predictability implied by (2); on the opposite, the bond investment schedule is downward sloping. There is no demand for cash assets independently of the degree of relative risk aversion.[32] Figure 7, panel (a), shows the implied welfare costs from excluding real estate from the portfolio problem. In this case, the utility loss is clearly increasing in the coefficient of relative risk aversion. Even in the case of intermediate risk aversion ($\gamma = 5$) the welfare loss at a short horizon of one year only is of 70 basis points; the loss for a long horizon investor exceeds a total of 9% of the initial wealth, i.e. almost 90 basis points a year. Such figures double or even triple if one considers a highly risk-averse investor; e.g., under a 10-year horizon and assuming $\gamma = 10$, the loss is in the order of 200 basis points per year. This means that especially under long planning horizons, including real estate in the asset menu should represent a primary concern for all portfolio managers.

Table 6, and figures 6 and 7 complete the picture by reporting results for the Bayesian portfolio choice problem, when parameter uncertainty is kept into consideration. Table 6 gives Bayesian posterior means and standard deviations for the restricted VAR model in which $n = 2$, i.e. excluding real estate excess returns. Posterior means are very close to MLE estimates, and standard errors confirm the results in table 5: the evidence of predictability is weak and only stock returns exhibit some degree of persistence. Figure 6 plots instead optimal asset allocations and obtains differences between classical and Bayesian portfolio weights consistent with our comments in Section 4: positive weights on liquid investments appear for $\gamma \geq 5$ under parameter uncertainty, as a protective measure against the additional uncertainty deriving from the fact the coefficients are random; while the equity investment schedule is generally upward sloping in a classical framework (an effect of predictability), the Bayesian allocation to stocks tends to decline with the investment horizon. Finally, figure 7 displays the percentage compensatory variation from excluding real estate from the asset menu. In this case, results are similar to the classical ones, i.e. the loss from ignoring real estate may easily exceed 100 basis points per year for either highly risk-averse and/or for long-horizon investors.

---

[32]For moderate levels of relative risk aversion, some non-monotonic features appear for horizons between 1 month and 2 years. Once more, these are explained by the way in which our simulations were initialized.

6. Robustness Checks

We conclude by performing a few additional exercises that aim at corroborating our results and show that they scarcely depend on specific assumptions concerning the statistical model, how the state variables are initialized, and the measurement of the welfare loss implied by solving 'standard' portfolio problems in which the predictable nature of risk premia is ignored and the attention is focused on a standard mean-variance benchmark (cf. Elton et al. (2004)).

6.1. *Restricted VAR Results*

One naturally wonders whether our results on the importance of predictability and real estate in rational portfolio choices for a European portfolio manager depend in any way on two choices we have made: (i) similarly to Campbell, Chan, and Viceira (2003) (and differently from other papers, e.g. Barberis (2000) and Lynch (2001)) we have considered an unrestricted VAR(1) model in which past excess returns are allowed to forecast future excess returns as well as dividend yields; (ii) within such a model, we have disregarded the statistical significance of the estimates obtained, i.e. predictive densities were defined on the basis of a number of parameters characterized by relatively large standard errors.[33] We therefore run experiments in which both these aspects are dealt with.

First of all, we apply the same statistical restrictions employed in Barberis (2000) and many other papers in the literature, and prevent past excess asset returns from forecasting both future returns and dividend yields. Equivalently, we restrict the model so that only the dividend yield forecasts subsequent excess returns in a linear fashion. Clearly, this is just model (8) employed in Section 4.3 for our dynamic portfolio simulations. The buy-and-hold calculations are repeated for both the classical and Bayesian cases. In both frameworks, results are qualitatively similar to those reported in figures 1 and 2. However, the restrictions in (8) seem to slightly inflate the optimal weights assigned to stocks and real estate, i.e. to improve their mean-risk trade-off. For instance, assuming $\gamma = 5$ and $T = 10$ years, the stock and real estate weights go from 32 and 26 percent, respectively, under model (2) to 36 and 33 percent under (8). Clearly, an increase of 7 percentage points in the real estate holdings over the long-run is rather important, at least in economic terms. However, our early conclusion that at least one-fifth or more of an optimal, European portfolio should be held in real estate assets continues to hold, and is actually strengthened by these experiments.[34]

Secondly, we enforce a second type of statistical restrictions: we start from the general model in (2) but proceed to set to zero all the VAR coefficients that fail to be statistically significant according to a standard, classical t-ratio test with a significance level of 10%. The resulting VAR model is the one in Table 3 with all the non-boldfaced coefficients in Φ set to zero. Although formal tests would imply a different technology (and language) based on odds ratios, for comparison purposes identical restrictions are also enforced in the Bayesian set up. Also in this case, we focus on buy-and-hold results only. Classical results are virtually indistinguishable from those in figure 1. If any, restricting our attention to predictability parameters that are 'statistically significant' implies an increased weight to be assigned to real estate (the returns on which

---

[33]Notice that this second choice is natural given that we want to focus on both classical and Bayesian frameworks, where the latter does not (easily) entertains the notion of 'statistical significance'.

[34]Complete results are available from the Authors upon request.

are predictable from lagged excess bond returns).[35] Interestingly, Bayesian results are different from those reported in figure 2 in one specific respect: while in Figure 2 schedules were essentially flat − a finding that we have interpreted as a result of the compensation between predictability and parameter uncertainty − when non-significant parameters are eliminated, investment schedules as a function of the horizon become generally *downward sloping*. This means that the (negative, in terms of slope) effects of the reduced risk premia predictability induced by the elimination of many VAR coefficients are stronger than the (positive) effects of the reduction in parameter uncertainty caused by the elimination of just those coefficients that in table 3 happened to be estimated less precisely. In any event, also in this case it remains true that roughly one-fifth of the portfolio ought to be held in the form of real estate assets. In fact, calculations of welfare costs of expelling real returns from the asset menu corroborate this claim: real estate should be a first-order concern in applied portfolio management.

6.2. *Sensibility of Optimal Weights to the Dividend Yield*

With limited exceptions, all of our simulation experiments so far have been based on initializing the dividend yield to its full-sample mean of 2.6%. Under this assumption on a plausible value for dy_t, we have found that under predictable risk premia, real estate represents an important asset class that ought to complement stocks and bonds in optimal long-run portfolios. However, it remains natural to ask whether this conclusion is robust to different assumptions concerning the dividend yield, especially given the observation that table 1 implies that all yields in the interval [0.9 4.3] percent have to be considered 'plausible' as they fall in a 90% confidence interval.[36] Figure 8 plots the resulting optimal asset allocation choices as the dividend yield changes over a wide range of possible initial (current) values in simulation experiments. Several exercises are performed corresponding to alternative assumptions concerning the coefficient of relative risk aversion ($\gamma = 2$ or 5), the horizon ($T = 1$ and 10 years), and the rebalancing frequency ($\varphi = 1$ year and $\varphi = T$, i.e. buy-and-hold). However, to simplify calculations, we only report classical results that ignore parameter uncertainty. The qualitative insights are similar for the Bayesian case.

Results for stocks and bonds are those one would expect from first principles: as a higher dividend yield forecasts good future (equity) investment opportunities, the optimal stock weight is monotone increasing in dy_t for all γs, Ts, and φs; on the opposite, the rational bond weight is generally monotonically decreasing in the dividend yield. Such a statement admits one qualification: for $\gamma = 5$, the bond schedule is actually increasing for dy_t between 1 and 1.5%. This is explained by the fact that for intermediate risk aversion and extremely poor investment opportunities (i.e. assuming a very low initial dividend yield, an indication of initially high asset prices) an investor would actually reduce the demand for bonds and invest in cash in percentages between 10 and 50 percent. However, for dividend yields in excess of 1.5%, the demand for cash disappears entirely, consistently with the results reported in Section 4.1. More importantly, the optimal demand for real estate is positive and not negligible for a wide range of dividend yield values, especially those around the full-sample mean (i.e. between 2 and 3%). Visibly, the optimal real estate weight is especially high (20-50 percent) for long-run investors with $T = 10$ years. Interestingly, over the interval 2-3 percent, the demand for real estate seems to only weakly depend on the exact value taken by the initial dividend yield, while a few inverted U-shapes appear. Since most of the observed sample values

---

[35]Assuming $\gamma = 5$ and $T = 10$, the restricted real estate weight is 28% vs. an unrestricted 26%.
[36]In fact, over our sample period dividend yields between 1.7 and 3.7 percent were observed.

for the dividend yield do fall in the range 2-3 percent, Figure 8 implies in principle a very stable demand for real estate, potentially in the range 10-30%.[37] Figure 8 makes us speculate that real estate may somehow represent the pillar, the 'backbone' of strategic asset allocation for European portfolio managers.

6.3. *Welfare Costs of Excluding Real Estate Under Constant Investment Opportunities*

Section 5 has assessed the long-run welfare costs of ignoring predictability in approximately 100 basis points per year, i.e. a rational investor should request a riskless increase of the annual returns from his wealth in the order of 1 percent a year. A similar result was obtained for the compensatory variation that should be required to ignore real estate as an asset class in addition to stocks, bonds, and cash financial investments. However, this last estimate has been obtained assuming that predictability should and would not be ignored by a rational portfolio manager. Even though we believe that the statistical evidence in favor of the existence of predictability patterns in mean excess returns is strong and that a welfare cost of roughly 1% a year ought to be a major incentive for investors not to ignore predictability, it remains interesting to repeat the calculations of Section 6.2 when predictability is ignored and excess returns are generated by the simple model (7).

In the classical case, we obtain a picture that is very similar to Figure 7(a): the cost of excluding real estate grows with the coefficient of risk aversion; for long-run investor with $\gamma = 5$, the cost is in the order of 80 basis points a year, and this estimate grows to exceed 200 basis points when $\gamma = 10$ is considered. This implies that the cost of ignoring real estate scarcely depends on the whether predictability is modeled or not, although it is clear that the welfare gains from doing so remain substantial.[38] Similar calculations are performed in the Bayesian case. We find that welfare costs of ignoring real estate are actually higher when predictability is ignored altogether.[39] In fact, Figure 3 has shown that the optimal real estate weight is higher by 10-15 percent vs. the case in which predictability is taken into account (figure 2). This is easily explained by the fact that in tables 2 and 3, real estate excess returns were characterized by rather weak predictability patterns so that it is not difficult to think that the demand for real estate may be hurt in net terms by the existence of estimation uncertainty. In this sense, restricting our exercise to the case of i.i.d. excess asset returns may bring − when parameter uncertainty is taken into account − to a higher estimate of the utility loss deriving from ignoring real estate. In conclusion, the welfare losses reported and discussed in Section 6.2 represent at best *a lower bound* for the utility costs of omitting real estate when choosing optimal European portfolios.

7. Conclusions

In this paper we have documented the existence of linear predictability patterns − described by a simple VAR(1) framework − in an asset menu that involves both financial and real estate excess returns. In

---

[37]In fact, more than 60% of the available observations fall in the 2-3 percent range.

[38]This remark is consistent with the finding in Figure 8 that over a relatively wide range of values for the dividend yield, optimal portfolio weights fail to depend on the state variable that determines the strongest predictability patterns (see Lynch (2001) on this point). Full results on the welfare costs estimates as a function of the investment horizon are available upon request.

[39]For instance, assuming $\gamma = 5$, the cost for a 10-year horizon investor exceeds 120 basis points a year. When VAR predictability is modeled, such an estimate is 100 basis points 'only'.

particular, lagged excess stock and real estate returns predict subsequent real estate and bond excess returns respectively. Moreover, real estate and stock return shocks are negatively correlated with dividend news shocks. Both facts make stock and real estate returns less risky as the horizon grows. As a result, when we calculate optimal portfolio weights based on the MLE estimates of the VAR coefficients, we find portfolio weights for stocks and real estate that are monotone increasing in the planning horizon. Stocks, bonds, and real estate do not appear excessively risky to a long-run investor, so that the demand for cash is rather limited or even absent. These findings are robust to the adoption of a Bayesian approach that incorporates estimation risk into the formal portfolio problem, although the trade-off between predictability and parameter uncertainty makes for flatter investment schedules.

In this framework, we find that real estate has a considerable importance, both in absolute terms − as measured by the portfolio weights resulting from the problem − and in welfare terms: the compensatory variation required by an investor to do without real estate is easily in excess of 100 basis points per year. Although the welfare costs deriving from ignoring predictability would be of similar importance, it is interesting to notice that the conclusions above concerning the utility loss from expelling real estate from the asset menu do not depend on the finding of predictability. As a matter of fact, our robustness checks in Section 6 suggest that our estimates for the optimal real estate weights and welfare losses from restricting the asset menu are probably only a lower bound for higher estimates obtainable under alternative assumptions. For instance, it would be interesting to perform some of the calculations in this paper using indirect indices, that traditionally imply a much smaller correlation with stock and bond returns and hence offer greater diversification opportunities.

In policy terms, our paper implies that preferences and horizons exist that justify both the British and the continental European patterns of long-horizon portfolio diversification for pension institutions. In particular, the German-French pattern is consistent with the rational choices of portfolio managers that display high risk aversion, worry about estimation risk, and have investment horizons between 5 and 10 years. On the contrary, investing four-fifths of one's wealth in risky assets - as in the Anglo-Saxon pattern-may be optimal for investors with low risk aversion disregarding parameter uncertainty. However, in this case real estate ought to receive a weight at least as important as stocks, which is not the case.

Our results on optimal asset allocation stress the potential rationality of current practices with pension fund managers of limiting their exposure to real estate risk to roughly 10%-15% of their overall managed assets.

Of course, there are many issues that this paper merely touches upon. First and foremost, we have ignored transaction costs throughout. On one hand, this is consistent with our choice to use an indirect real estate return index based on the behavior of the market price of the equity issued by companies involved in real estate operations: it makes sense to entertain the assumption that the frictions associated in trading in these companies may not be structurally different from the market average. Furthermore, even assuming that real estate implies higher transaction costs than financial securities, it remains difficult to think that such a differential may easily exceed a full one percent of total initial wealth of the investor (also considering that only 10-30 of this wealth ought to be invested in real estate). On the other hand, recent papers (e.g. Balduzzi and Lynch (1999)) have shown how dynamic portfolio choices may be computed in the presence of transaction costs. Such an effort seems to be particularly appropriate for assets in the real estate class. We leave these further exciting explorations for future research.

References

[1] Ang, A., and G., Bekaert, 2002, "International Asset Allocation with Regime Shifts", *Review of Financial Studies*, 15, 1137-87.

[2] Ang, A., and J., Chen, 2002, "Asymmetric Correlations of Equity Portfolios", *Journal of Financial Economics*, 63, 443-94.

[3] Balduzzi, P., and A., Lynch, 1999, "Transaction Costs and Predictability: Some Utility Cost Calculations", *Journal of Financial Economic*s, 52, 47-78.

[4] Barberis, N., 2000, "Investing for the Long Run When Returns Are Predictable", *Journal of Finance*, 55, 225-64.

[5] Bawa, V. S., S. Brown, and R. Klein, 1979, *Estimation Risk and Optimal Portfolio Choice*, North Holland, Amsterdam.

[6] Bekaert, G., and R., Hodrick, 1992, "Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets", *Journal of Finance*, 47, 467-509.

[7] Bharati, R., and M., Gupta, 1992, "Asset Allocation and Predictability of Real Estate Returns', *Journal of Real Estate Research*, 7, 469-484.

[8] Black, A., and P., Fraser, 1995, "U.K. Stock Returns: Predictability and Business Conditions", *Manchester School*, 63, 85-102.

[9] Bodie, Z., 1995, "On the Risk of Stocks in the Long Run", *Financial Analysts Journal*, May-June,18-22.

[10] Brandt, M., 1999, "Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach", *Journal of Finance*, 54, 1609-1645.

[11] Brennan, M., E., Schwartz, and R., Lagnado, 1997, "Strategic Asset Allocation", *Journal of Economic Dynamics and Control*, 21, 1377-1403.

[12] Campbell, J., Y., Chan, and L., Viceira, 2003, "A Multivariate Model of Strategic Asset Allocation", *Journal of Financial Economics*, 67, 41-80.

[13] Campbell, J., J., Cocco, F., Gomes, P., Maenhout, and L., Viceira, 2001, "Stock Market Mean Reversion and the Optimal Equity Allocation of a Long-Lived Investor", *European Finance Review*, 5, 269-292.

[14] Campbell, J.Y., A. Lo, and A. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press.

[15] Campbell, J., and R., Shiller, 1988, "Stock Prices, Earnings, and Expected Dividends", *Journal of Finance*, 43, 661-676.

[16] Campbell, J., and L., Viceira, 1999, "Consumption and Portfolio Decisions When Expected Returns Are Time Varying", *Quarterly Journal of Economics*, 114, 433-495.

[17] Campbell, J., and L., Viceira, 2001, "Who Should Buy Long-Term Bonds?", *American Economic Review*, 91, 99-127.

[18] Campbell, J., and L., Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors.* Oxford University Press,

[19] Chandrashakaran, V., 1999, "*The Time Series Properties and Diversification Benefits of REIT Returns*", Journal of Real Estate Research, 17, 91-112.

[20] Clare, T., A., Thomas, and M., Wickens, 1994, "Is the Gilt-Equity Yield Ratio Useful for Predicting UK Stock Returns?", *Economic Journal*, 104, 303-315.

[21] Detemple, J., R., Garcia, and M., Rindisbacher, 2003, "A Monte Carlo Method for Optimal Portfolios", *Journal of Finance*, 58, 401-446.

[22] Fama, E., and K., French, 1988, "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics*, 22, 3-25.

[23] Fama, E., and K., French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds", *Journal of Financial Economics*, 25, 23-49.

[24] Ferson W., and C. R. Harvey, 1991, "The Variation of Economic Risk Premiums", *Journal of Political Economy,* 99, 385-415.

[25] Flavin, T., and M., Wickens, 2001, "Tactical asset allocation: a multivariate GARCH approach", mimeo, University of York.

[26] Geltner, D., 1989, "Bias in Appraisal-Based Returns American", *Real Estate and Urban Economics Association Journal*, 17, 338-352.

[27] Gennotte, G., 1986, "Optimal Portfolio Choice under Incomplete Information", *Journal of Finance*, 41, 733-746.

[28] Giliberto, M., 1993, "Measuring Real Estate Returns: The Hedged REIT Index", *Journal of Portfolio Management,* 19, 94-99.

[29] Goetzmann, W., and P., Jorion, 1993, "Testing the Predictive Power of Dividend Yields", *Journal of Finance*, 48, 663-79.

[30] Gold, R., 1993, "Why the efficient frontier for real estate is 'Fuzzy,", *Journal of Real Estate Portfolio Management*, 1, 59-66.

[31] Guidolin, M., and A., Timmermann, 2003, "Recursive Modeling of Nonlinear Dynamics in UK Stock Returns", *Manchester School*, 71, 381-395.

[32] Guidolin, M. and A., Timmermann, 2004, "An Econometric Model of Nonlinear Dynamics in the Joint Distribution of Stock and Bond Returns". *Journal of Applied Econometrics,* forthcoming.

[33] Hamilton, J.D. (1994). *Time Series Analysis*, Princeton University Press.

[34] Hodrick, R., 1992, "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement", *Review of Financial Studies*, 5, 357-386.

[35] Hudson-Wilson, S., F., Fabozzi, and J., Gordon, 2003, "Why Real Estate? An Expanding Role for Institutional Investors", mimeo.

[36] Ibbotson, R., and L., Siegel, 1984, "Real Estate Returns: A Comparison with Other Investments", *American Real Estate and Urban Economics Association Journal*, 12, 219-242.

[37] Ingersoll, J., 1987, *Theory of Financial Decision Making*", Rowman and Littlefield.

[38] Kandel, S. and R. F. Stambaugh. 1987, "Long Horizon Returns and Short Horizon Models", CRSP Working Paper No. 222. University of Chicago, Chicago IL.

[39] Kandel, S., and R., Stambaugh, 1996, "On the Predictability of Stock Returns: An Asset-Allocation Perspective", *Journal of Finance*, 51, 385-424.

[40] Karlberg, J., C., Liu, and W. Greig, 1996, "The Role of Real Estate in the Portfolio Allocation Process", *Real Estate Economics*, 24, 359-377.

[41] Klein, R.W. and V.S. Bawa,1976, "The Effect of Estimation Risk onOptimal Portfolio Choice", *Journal of Financial Economics* 3, 215-231.

[42] Li, Y., and K., Wang, 1995, "The Predictability of Real Estate Returns and Market Segmentation", *Journal of Real Estate Research*, 10, 471-482.

[43] Liang, Y., N., Myer, and J., Webb, 1996, "The Boostrap Efficient Frontier for Mixed-Asset Portfolios", *Real Estate Economics*, 24, 247-256.

[44] Liang, Y., and J., Webb, 1996, "The Hedged REIT Index and Mixed-Asset Portfolios", *Journal of Real Estate Portfolio Management*, 2, 55-61.

[45] Liao, H. and J. Mei,1998, "*Risk characteristics of real estate related securities - An extension of Liu and Mei (1992)*", Journal of real estate research, volume 16, 279-290.

[46] Liu, C. H. and J. Mei, 1992, "*The predictability of returns on equity REITs and their co-movement with other assets*", Journal of real estate finance and economics, 5, 401-408.

[47] Lynch, A., 2001, "Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands Induced by Return Predictability", *Journal of Financial Economics*, 62, 67-130.

[48] Maurer R., F. Reiner and R. Rogalla, 2003, "Risk and Return of Open-End Real-Estate Funds: The German Case," Working Paper Series: Finance and Accounting 114, Department of Finance, Goethe University Frankfurt am Main.

[49] Maurer R., F. Reiner, and S. Sebastian, 2004, "Characteristics of German Real Estate Return Distributions: Empirical Evidence from Germany and Comparison to the U.S. and U.K," Working Paper Series: Finance and Accounting 108, Department of Finance, Goethe University Frankfurt am Main.

[50] Merton, R., 1969, "Lifetime Portfolio Selection under Uncertainty: the Continuous Time Case", *Review of Economics and Statistics*, 51, 247-257.

[51] Merton, R.C.,1971,Optimal Consumptionand Portfolio Rules in a Continuous-Time Model, *JournalofEconomicTheory*, 3, 373-413.

[52] Miles D., 1994, *Housing, Financial Markets and the Wider Economy*, John Wiley.

[53] Pesaran, M., and A., Timmermann, 1995, "Predictability of Stock Returns: Robustness and Economic Significance", *Journal of Finance*, 50, 1201-1228.

[54] Pesaran, M. H. and A. Timmermann, 2002, "Market timing and return prediction under model instability," *Journal of Empirical Finance*, 9, 495-510.

[55] Samuelson, P., 1969, "Lifetime portfolio selection by dynamic stochastic programming", *Reveiw of Economics and Statistics*, 51, 239-246.

[56] Siegel, J., 1998, *Stocks for the Long Run: The Definitive Guide to Financial Market Returns and Long-Term Investment Strategies*, McGraw-Hill.

[57] Zellner, A., 1971, *An Introduction to Bayesian Inference in Econometrics,* John Wiley and Sons.

[58] Ziobrowski, A., R., Caines, and B., Ziobrowski, 1999, "Mixed-Asset Portfolio Composition with Long-Term Holding Periods and Uncertainty", *Journal of Real Estate Portfolio Management*, 2, 139-144.

Table 1

Summary Statistics for Asset Returns and the Stock Dividend Yield

The table reports summary statistics for monthly total excess return series (including dividends, coupon distributions, adjusted for splits, etc.) for the 3-month JP Morgan Euro Cash Index, stocks, bonds, and real estate investments. The sample period is March 1993 – March 2004. All returns are expressed in euros. Return data for stocks, bonds, and real estate are in excess of the real risk-free rate. The sources of the data are: DataStream (European Equity Index), Citigroup (bond index, a weighted basket with maturities of five years and longer), and EPRA (Real Estate Liquid 40 Index). The real risk-free rate is calculated by subtracting the European harmonized consumer price inflation index from nominal returns. Means, medians, and standard deviations are annualized by multiplying monthly moments by 12 and $\sqrt{12}$, respectively. LB(j) denotes the j-th order Ljung-Box statistic. The last row reports statistics for the dividend yield calculated on the DataStream Equity Return Index.

Portfolio/Asset Class	Mean	Median	St. Dev.	Skewness	Kurtosis	Jarque-Bera	LB(4)	LB(4)-squares
3-month Euro cash (real)	0.030	0.025	0.013	0.515	3.748	8.969*	12.201*	8.959
Excess stock returns	0.065	0.115	0.172	-0.751	3.874	16.723**	4.018	22.933**
Excess bond returns	0.042	0.065	0.054	-0.550	3.422	7.701*	3.700	4.244
Excess real estate returns	0.066	0.080	0.125	-0.368	3.439	4.077	10.889*	1.272
Dividend Yield	0.026	0.027	0.017	-0.158	2.016	5.921	474.35**	465.38**

* denotes 5% significance, ** significance at 1%.

Table 2

Correlation Matrix of Excess Asset Returns

The table reports linear correlation coefficients for monthly excess total return series. The sample period is March 1993 – March 2004. All returns are expressed in euros.

	Excess stock returns	Excess bond returns	Excess real estate returns	Dividend Yield
Excess stock returns	1	0.067	0.558	-0.133
Excess bond returns		1	0.199	0.103
Excess real estate returns			1	-0.063
Dividend Yield				1

Table 3

Classical Parameter Estimates for a VAR(1) Model of Excess Returns and the Dividend Yield

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

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

where $\mathbf{z}_t$ 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$_t$
μ'				
	-0.029	-0.004	-0.014	**0.001**
	(-1.235)	(-0.562)	(-0.854)	**(1.963)**
Φ'				
Stocks$_{t-1}$	0.109	0.028	**0.143**	-0.001
	(1.028)	(0.864)	**(1.947)**	(-0.249)
Bonds$_{t-1}$	0.072	0.066	0.165	-0.002
	(0.253)	(0.749)	(0.832)	(-0.257)
Real Estate$_{t-1}$	0.067	**-0.114**	0.110	-0.003
	(0.448)	**(-2.478)**	(1.055)	(-0.796)
Dividend Yield$_{t-1}$	1.262	0.290	0.651	**0.952**
	(1.428)	(1.064)	(1.058)	**(40.238)**

Table 4

Bayesian Coefficient Estimates for a VAR(1) Model of Excess Returns and the Dividend Yield

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 $\mathbf{z}_t$ includes continuously compounded monthly excess asset returns and the dividend yield, and $\varepsilon_t \sim N(\mathbf{0}, \Sigma)$. 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.

	Stocks$_t$	Bonds$_t$	Real Estate$_t$	Dividend Yield$_t$
μ'				
	-0.029	-0.004	-0.014	0.001
	(0.024)	(0.007)	(0.017)	(0.001)
Φ'				
Stocks$_{t-1}$	0.109	0.028	0.143	-0.001
	(0.108)	(0.033)	(0.074)	(0.003)
Bonds$_{t-1}$	0.075	0.066	0.166	-0.002
	(0.293)	(0.090)	(0.204)	(0.008)

Real Estate$_{t-1}$	0.068	-0.114	0.111	-0.003
	(0.151)	(0.047)	(0.105)	(0.004)
Dividend Yield$_{t-1}$	1.256	0.289	0.653	0.952
	(0.896)	(0.278)	(0.625)	(0.024)

Table 5

Classical Parameter Estimates for a VAR(1) Model of Excess Returns and the Dividend Yield – Restricted Asset Menu

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

$$y_t = \mu + \Phi y_{t-1} + \eta_t$$

where $\mathbf{z}_t$ includes continuously compounded monthly excess asset returns and the dividend yield (but not excess real estate returns), and $\eta_t \sim N(\mathbf{0}, \Lambda)$. 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$_t$	Dividend Yield$_t$
	μ'		
	-0.029	-0.004	**0.001**
	(-1.248)	(-0.612)	**(1.980)**
	Φ'		
Stocks$_{t-1}$	**0.135**	-0.017	-0.002
	(1.558)	(-0.603)	(-0.840)
Bonds$_{t-1}$	0.096	0.025	-0.003
	(0.349)	(0.289)	(-0.420)
Dividend Yield$_{t-1}$	1.257	0.298	**0.952**
	(1.449)	(1.086)	**(40.932)**

Table 6

Bayesian Coefficient Estimates for a VAR(1) Model of Excess Returns and the Dividend Yield – Restricted Asset Menu

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 $\mathbf{z}_t$ includes continuously compounded monthly excess asset returns and the dividend yield (but not excess real estate returns), and $\varepsilon_t \sim N(\mathbf{0}, \Sigma)$. 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, $p(C, \Sigma) \propto |\Sigma|^{-(n+2)/2}$, where $\mathbf{C} = [\boldsymbol{\alpha}' \ \mathbf{B}']'$.

	Stocks$_t$	Bonds$_t$	Dividend Yield$_t$
	μ'		
	-0.029	-0.004	0.001
	(0.024)	(0.008)	(0.001)
	Φ'		

|			
Stocks$_{t-1}$	0.135	-0.017	-0.002
	(0.089)	(0.028)	(0.002)
Bonds$_{t-1}$	0.095	0.025	-0.003
	(0.285)	(0.090)	(0.008)
Dividend Yield$_{t-1}$	1.259	0.298	0.952
	(0.895)	(0.283)	(0.024)

Figure 1

Buy-and-Hold Optimal Allocation – Ignoring Parameter Uncertainty

The graphs plot the optimal portfolio weights as a function of the investment horizons when returns follow a Gaussian VAR(1) model and parameter uncertainty is ignored (i.e. classical MLE estimates are employed). Three alternative coefficients of relative risk aversion are employed.

Risk aversion (γ) of 10

Horizon (in years)

– – – Stocks ——— Bonds – – – Real Estate —·—·— Cash
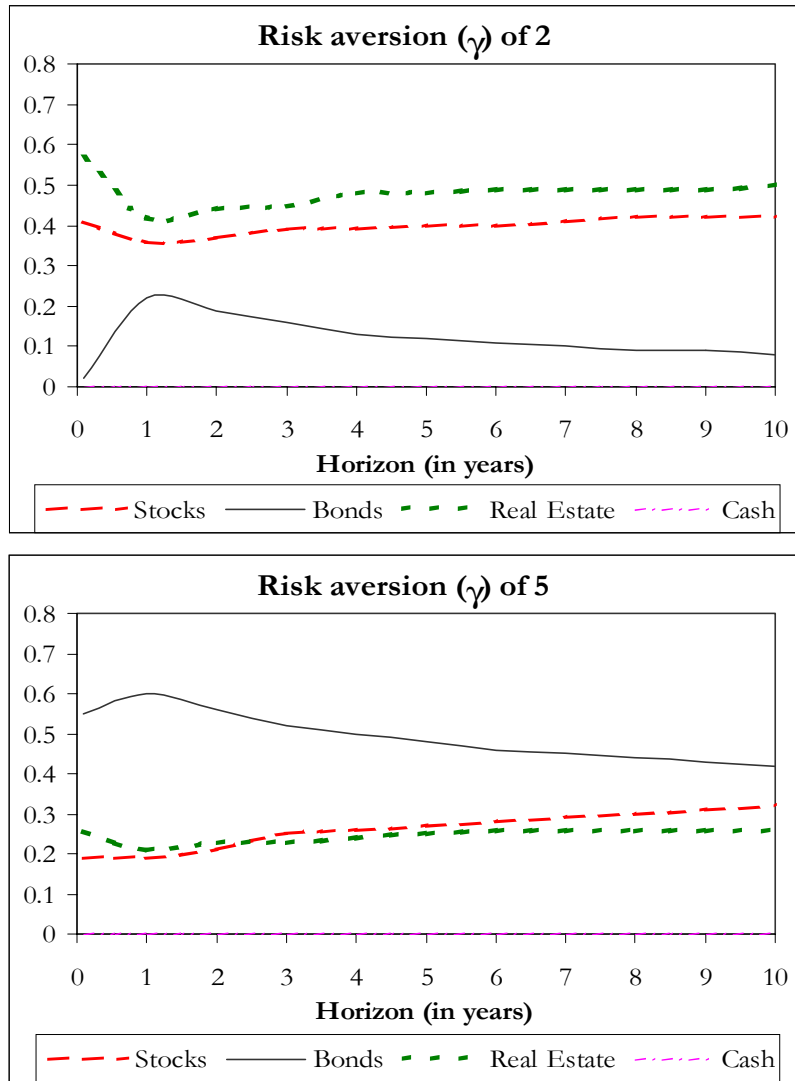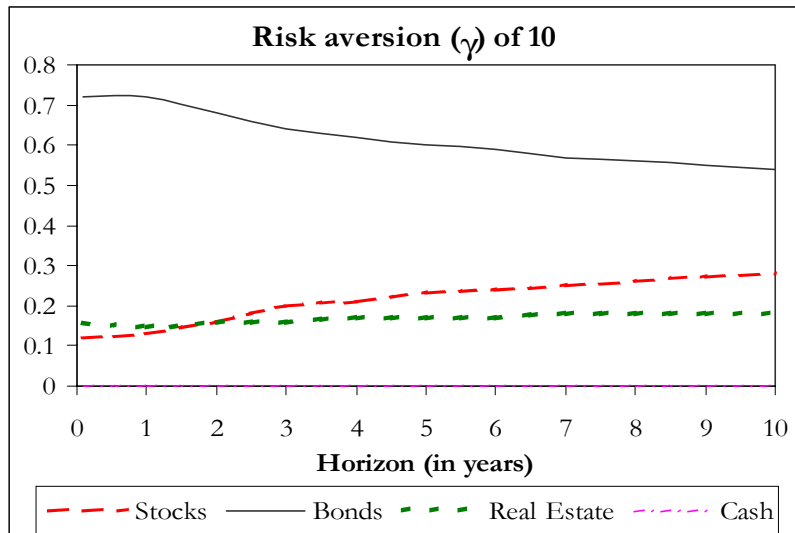
Figure 2
Buy-and-Hold Optimal Allocation – Effects of Parameter Uncertainty

The graphs plot the optimal portfolio weights as a function of the investment horizons when returns follow a Gaussian VAR(1) model and parameter uncertainty is accounted for (i.e. Bayesian predictive densities are employed). The posteriors are obtained from a standard uninformative prior, $p(C, \Sigma) \propto |\Sigma|^{-(n+2)/2}$, where $\mathbf{C} = [\boldsymbol{\alpha'}\ \mathbf{B'}]'$ is the matrix of the coefficients in the VAR model and n is the number of variables (4) in the multivariate system.

Risk aversion (γ) of 10
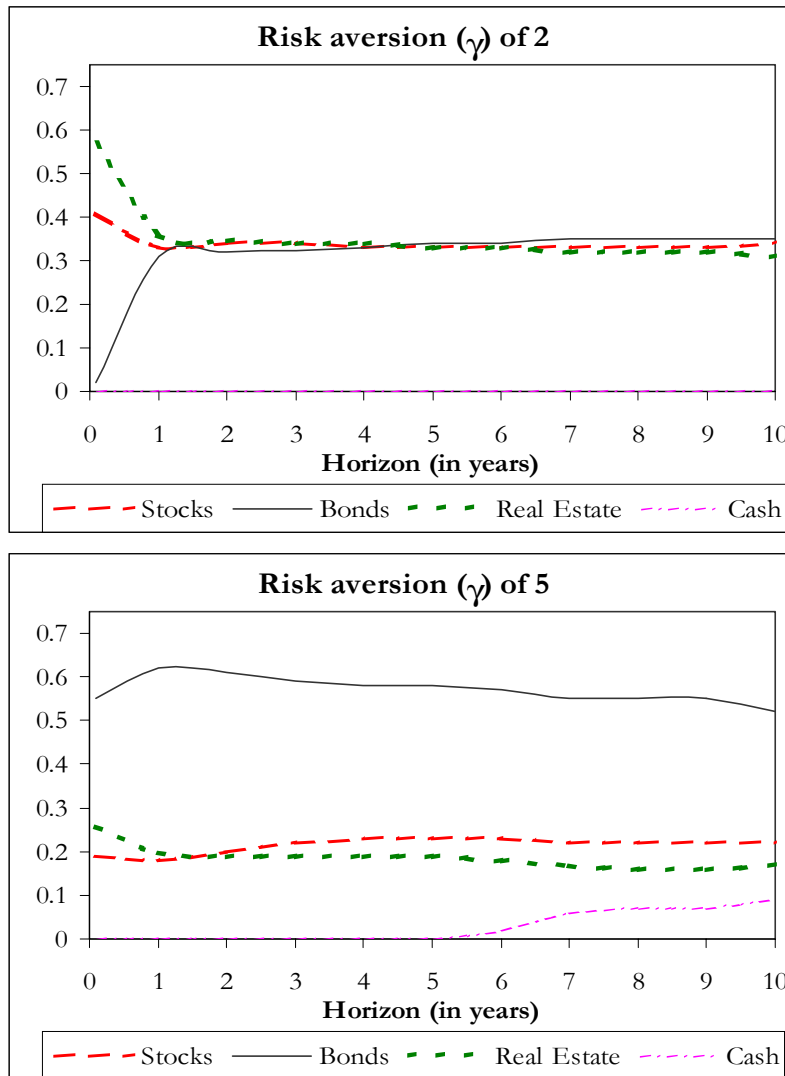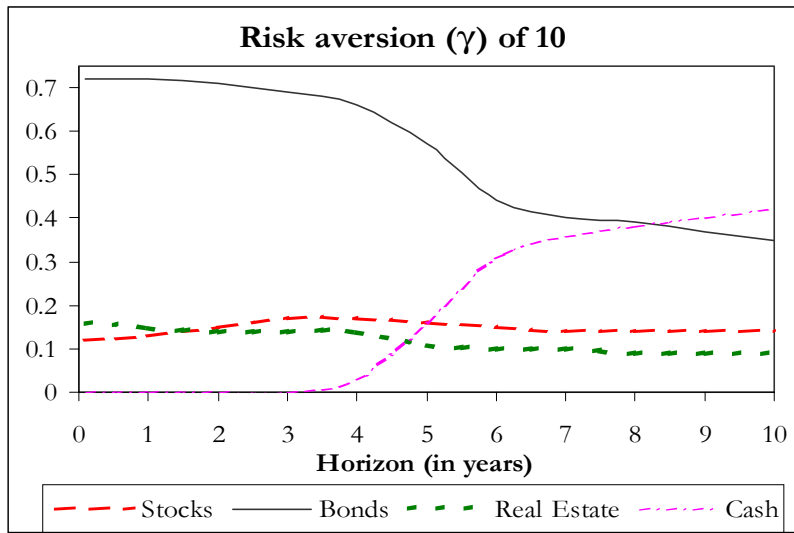
Figure 3
Buy-and-Hold Optimal Allocation – No Predictability and Parameter Uncertainty

The graphs plot the optimal portfolio weights as a function of the investment horizons when returns follow a Gaussian IID model and parameter uncertainty is accounted for (i.e. Bayesian predictive densities are employed). The posteriors are obtained from a standard uninformative prior, $p(\mu, \Sigma) \propto |\Sigma|^{-(n+2)/2}$, where n is the number of variables in the system.

Risk aversion (γ) of 10
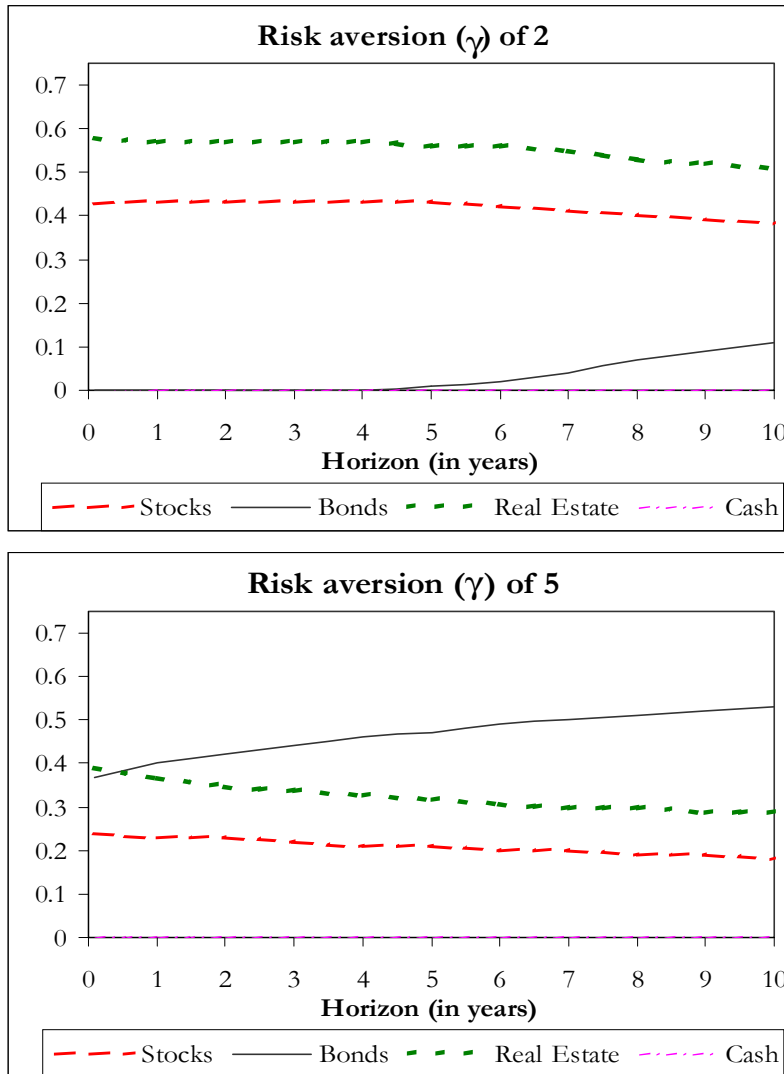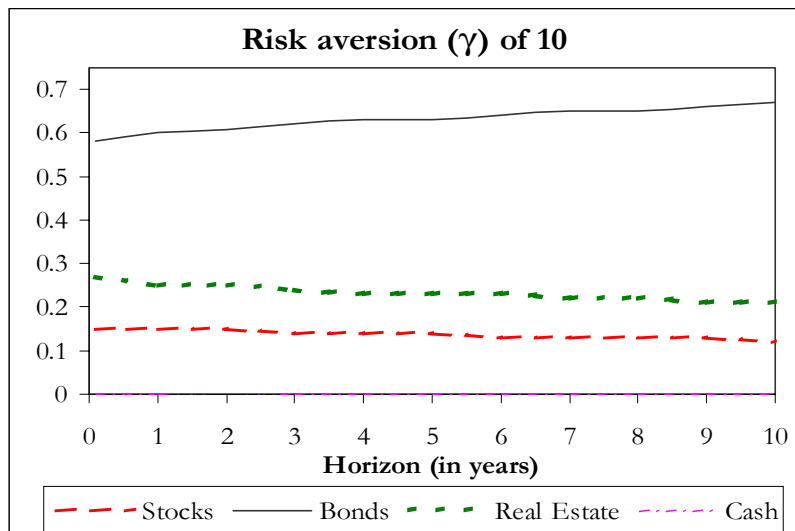
Figure 4
Optimal Dynamic Portfolio Allocation under Predictable Returns

The graphs plot the optimal portfolio weights as a function of the investment horizons when returns follow a Gaussian VAR(1) model and the investor rebalances her portfolio once a year The VAR is restricted so that only the dividend yield predicts future risk premia. Column (a) refers to the case in which parameter uncertainty is ignored (i.e. classical MLE estimates are employed), column (b) to the Bayesian case in which estimation risk is taken into account. When the horizon T exceeds or is equal to the rebalancing frequency $\varphi = 12$ months, optimal weights coincided by construction with those obtained under buy-and-hold strategies.

(a) Classical Asset Allocation   **(b) Bayesian Asset Allocation**

Risk aversion (γ) of 10

Horizon (in years)

Stocks — Bonds — Real Estate — Cash

Risk aversion (γ) of 10

Horizon (in years)

Stocks — Bonds — Real Estate — Cash
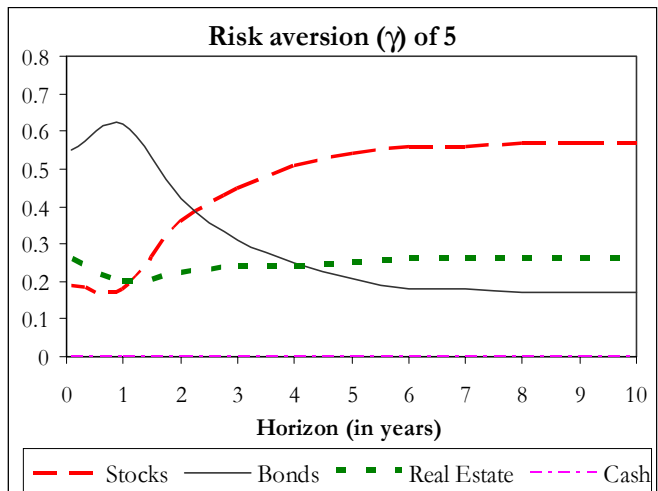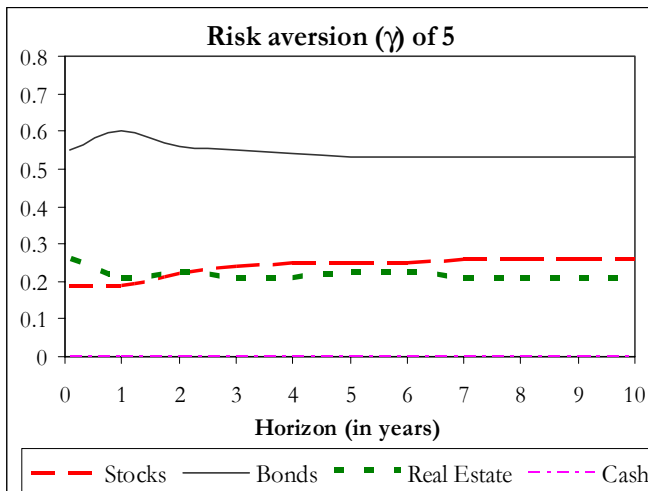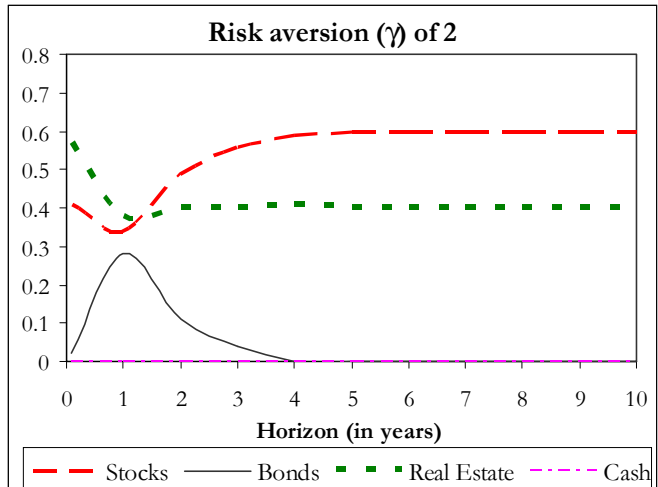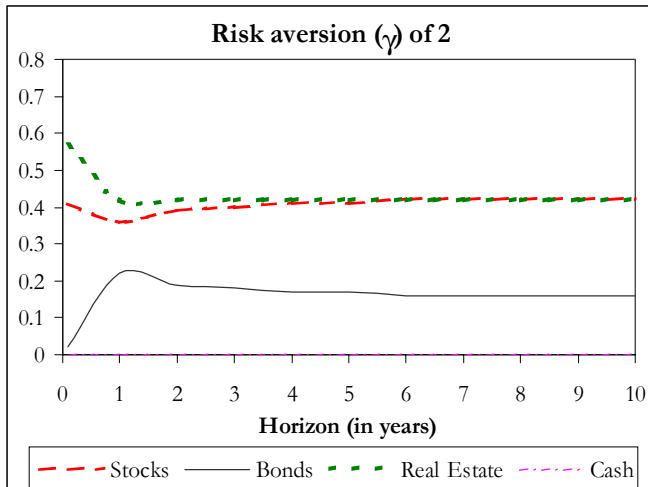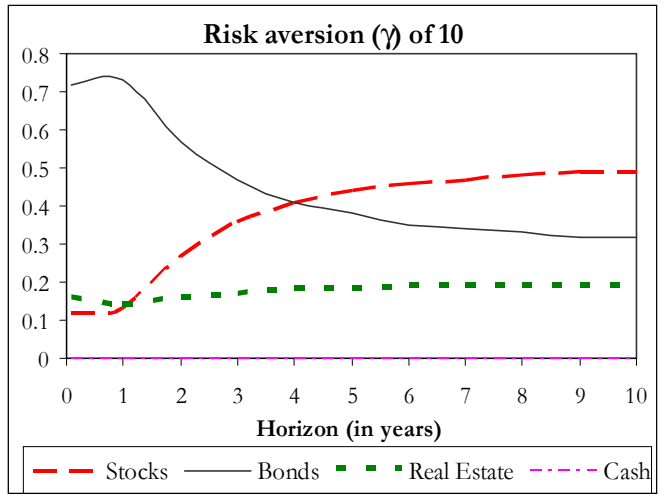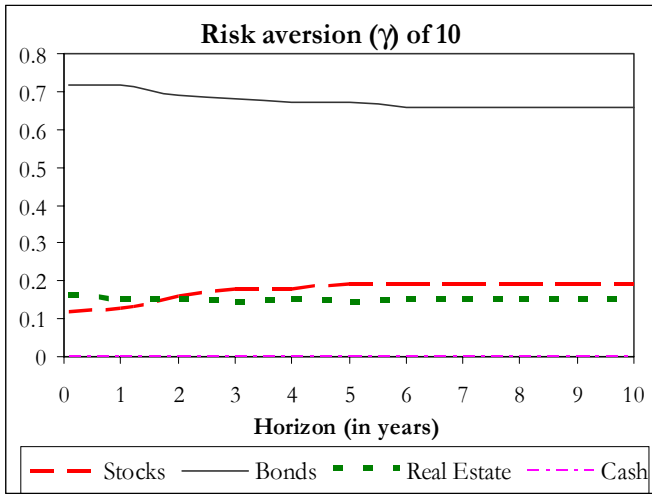
Figure 5
Welfare Costs – Ignoring Predictability

The graphs plot the percentage compensatory variation associated with ignoring the existence of predictability patterns in the data, i.e. with using a Gaussian IID model instead of a VAR(1) model. Panel (a) concerns the classical case in which MLE parameter estimates have replaced the unknown coefficients; panel (b) the Bayesian case in which parameter uncertainty is accounted for (i.e. Bayesian predictive densities are employed).

(a) Classical Estimates
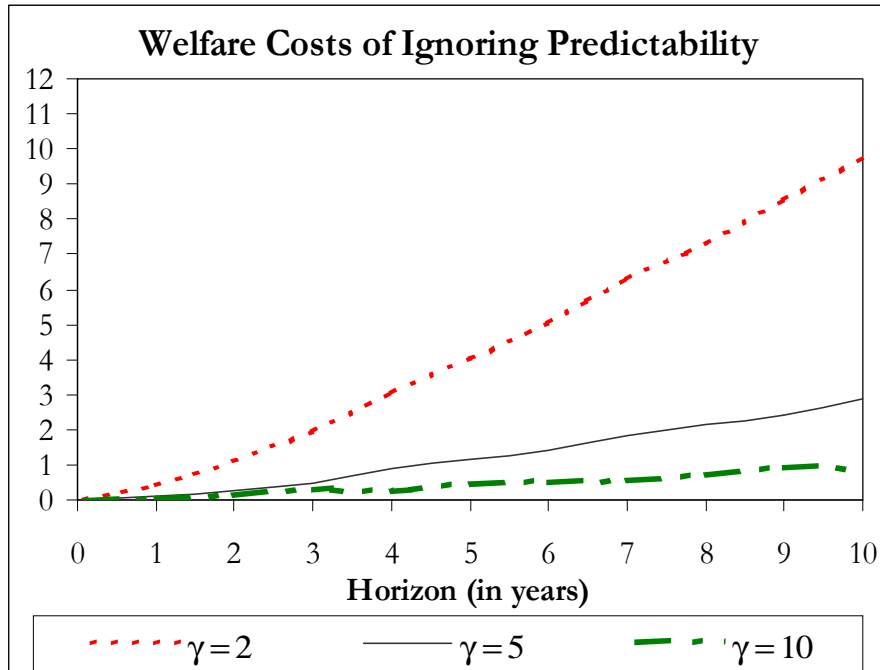
(b) Bayesian Estimates
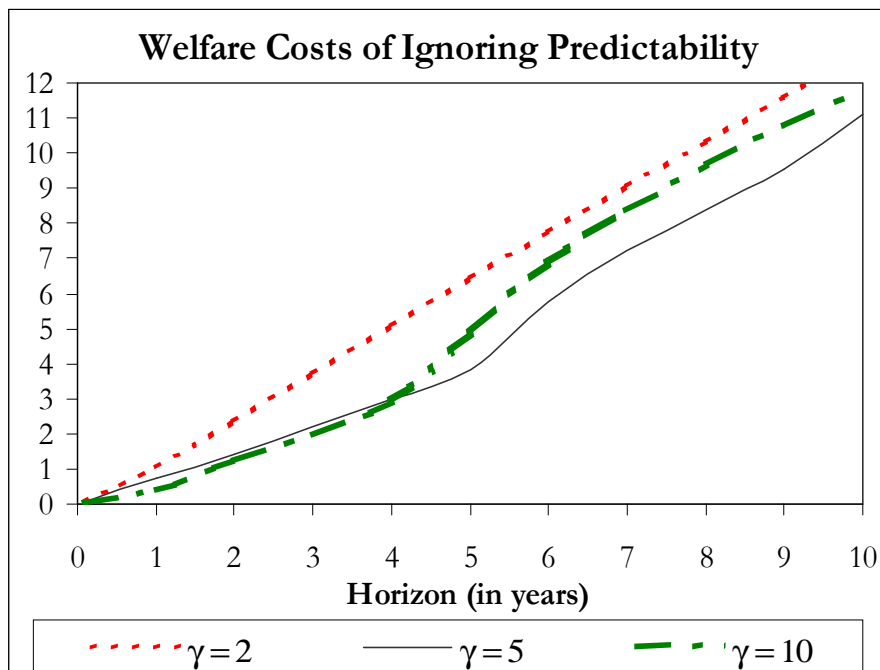
Figure 6
Buy-and-Hold Optimal Allocation – Restricted Asset Menu

The graphs plot the optimal portfolio weights as a function of the investment horizons when returns follow a Gaussian VAR(1) model. Three alternative coefficients of relative risk aversion are employed. The asset menu is restricted to the riskless asset, stocks, and bonds only. Column (a) refers to the case in which parameter uncertainty is ignored (i.e. classical MLE estimates are employed), column (b) to the Bayesian case in which estimation risk is taken into account.

(a) Classical Asset Allocation

(b) Bayesian Asset Allocation

Risk aversion (γ) of 10

Risk aversion (γ) of 10

Horizon (in years)

Horizon (in years)

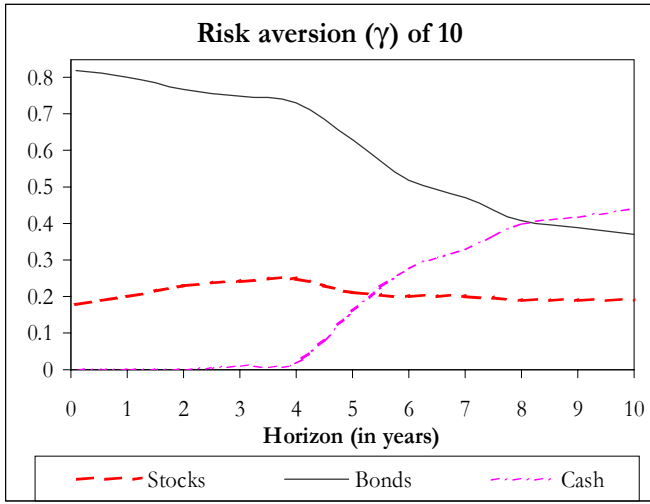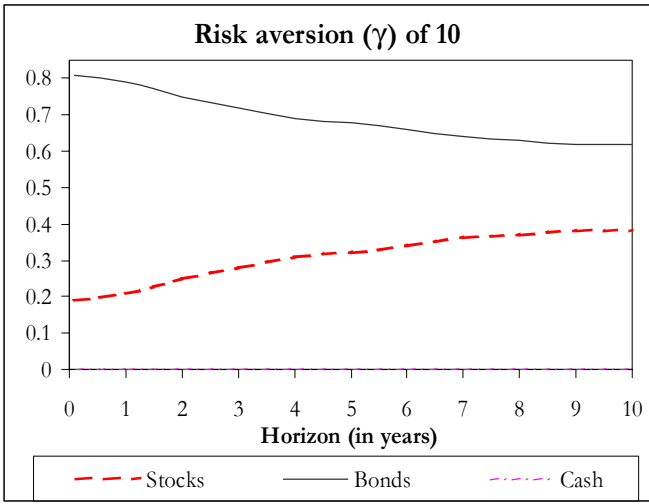– – – Stocks ——— Bonds –·–·– Cash

– – – Stocks ——— Bonds –·–·– Cash
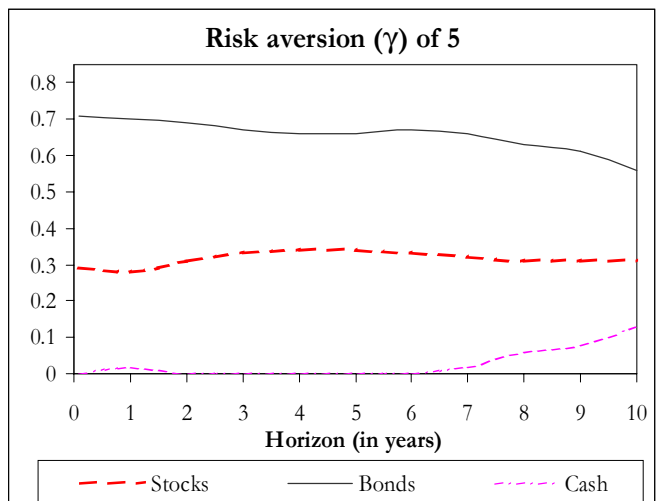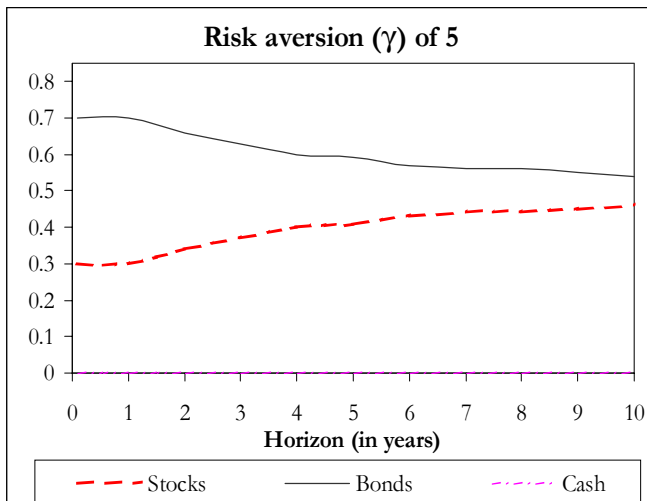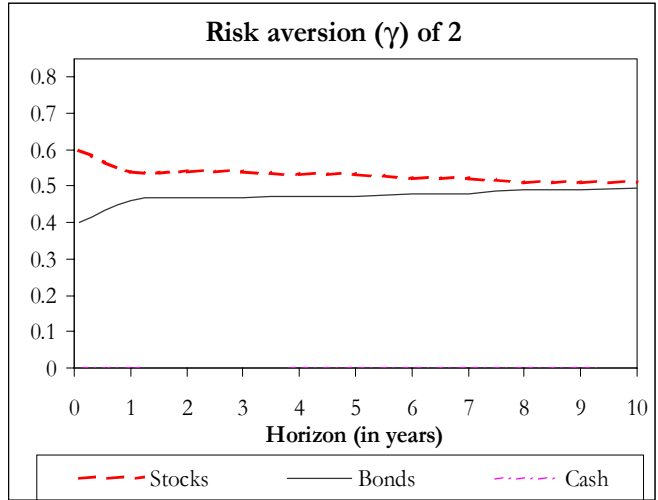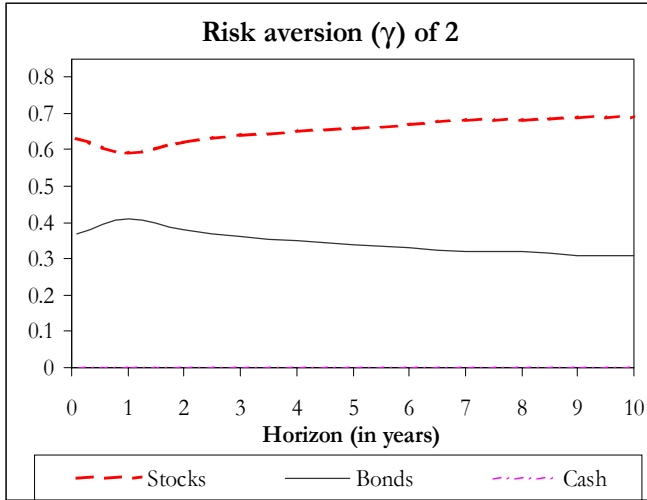
Figure 7
Welfare Costs – Ignoring Real Estate

The graphs plot the percentage compensatory variation associated with ignoring real estate as an asset class, i.e. with limiting an investor's portfolio choice to stock, bonds, and cash. Panel (a) concerns the classical case in which MLE parameter estimates have replaced the unknown coefficients; panel (b) the Bayesian case in which parameter uncertainty is accounted for (i.e. Bayesian predictive densities are employed).

(a) Classical Estimates
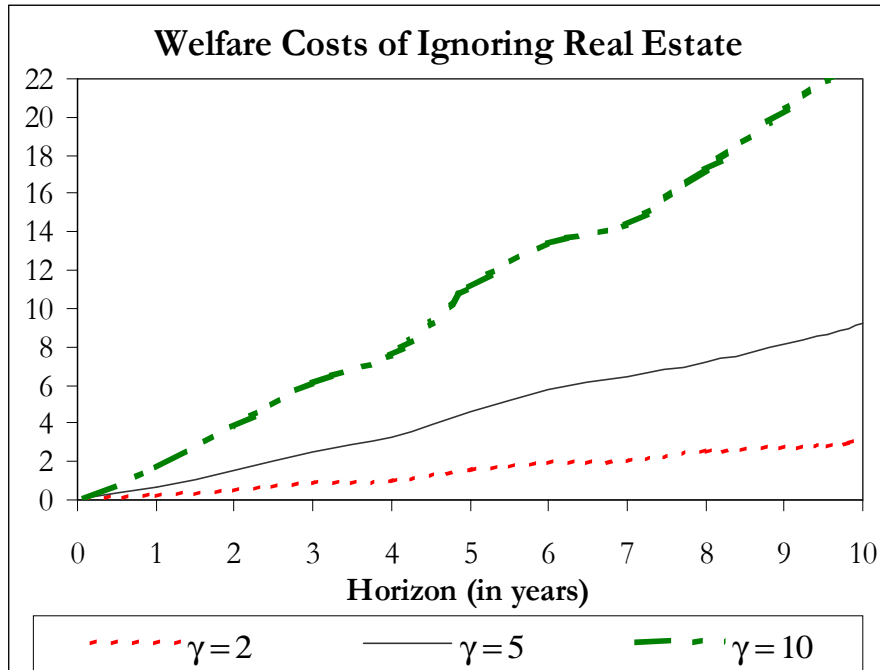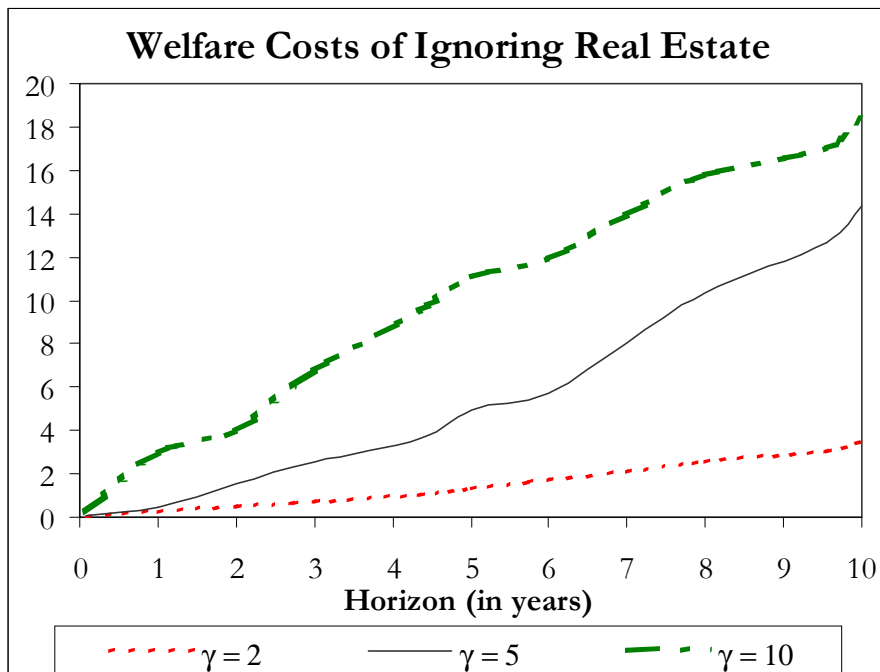

(b) Bayesian Estimates

Figure 8

Sensibility of Optimal Portfolio Weights to the Dividend Yield

The graphs plot the optimal portfolio weights to each asset class as a function of the initial dividend yield under a variety of assumptions on the coefficient of relative risk aversion (2 or 5), the investment horizon (1 or 10 years), and the rebalancing frequency (annual or identical to the investment horizon, the buy-and-hold case). Excess asset returns are assumed to follow a Gaussian VAR(1) model and estimation is performed either with classical (MLE) (i.e. disregarding parameter uncertainty) methods. In each plot, the dotted vertical bar indicates the full-sample mean dividend yield and hence locates the portfolio results previously reported in figures 1, 2, and 4.