

School of Economics and Finance

Forecasting Government Bond Yields with Large Bayesian VARs

Andrea Carriero, George Kapetanios and Massimiliano Marcellino

Working Paper No. 662

April 2010

ISSN 1473-0278



Queen Mary
University of London

Forecasting Government Bond Yields with Large Bayesian VARs*

A. Carriero

Queen Mary, University of London

G. Kapetanios

Queen Mary, University of London

M. Marcellino

European University Institute and Bocconi University

March 30, 2010

Abstract

We propose a new approach to forecasting the term structure of interest rates, which allows to efficiently extract the information contained in a large panel of yields. In particular, we use a large Bayesian Vector Autoregression (BVAR) with an optimal amount of shrinkage towards univariate AR models. Focusing on the U.S., we provide an extensive study on the forecasting performance of our proposed model relative to most of the existing alternative specifications. While most of the existing evidence focuses on statistical measures of forecast accuracy, we also evaluate the performance of the alternative forecasts when used within trading schemes or as a basis for portfolio allocation. We extensively check the robustness of our results via subsample analysis and via a data based Monte Carlo simulation. We find that: i) our proposed BVAR approach produces forecasts systematically more accurate than the random walk forecasts, though the gains are small; ii) some models beat the BVAR for a few selected maturities and forecast horizons, but they perform much worse than the BVAR in the remaining cases; iii) predictive gains with respect to the random walk have decreased over time; iv) different loss functions (i.e., "statistical" vs "economic") lead to different ranking of specific models; v) modelling time variation in term premia is important and useful for forecasting.

Keywords: Bayesian methods, Forecasting, Term Structure.

JEL: C11, C53, E43, E47.

*We are grateful to Jens Christensen, Oscar Jorda, Jose Lopez, Glenn Rudebusch, Eric Swanson and seminar participants at the San Francisco FED for useful comments on a previous draft.

1 Introduction

Producing accurate forecasts of the term structure of interest rates is crucial for bond portfolio management, derivatives pricing, and risk management. Unfortunately, all the forecasting models proposed so far in the macroeconomic and financial literature have a hard time in producing forecasts more accurate than a simple no-change forecast (i.e. a random walk forecast). The existing methods can be roughly categorized in three groups. The first two groups of models have the clear advantage of being grounded on finance theory, while the third group is the one that so far has produced the best results in out of sample forecast accuracy.

The first group contains models based on forward rate regressions. Such models try to forecast the future yields by extracting the information contained in the present forward rates. Prominent examples of this approach are e.g. Fama and Bliss (1987), and Cochrane and Piazzesi (2005). Even though these papers document the existence of a predictive content in the forward rates, the out of sample forecasts produced by these models are typically outperformed by a simple no-change forecast (see e.g. Diebold and Li 2006).

The second group contains models based on the No-Arbitrage paradigm. Typically the practical implementation of these models involves imposing an affine specification on a set of latent factors. Affine term structure models perform extremely well in fitting the yield curve in sample (see, e.g. De Jong 2000 and Dai and Singleton 2000) but the performance in out of sample forecasting is quite poor. Duffee (2002) has shown that beating a random walk with a traditional no arbitrage affine term structure model is difficult. Ang and Piazzesi (2003) show that imposing no-arbitrage restrictions and an essentially affine specification of market prices of risk improves out-of-sample forecasts from a VAR(12), but the gain with respect to a random walk forecast is small. Slightly more favourable evidence in this respect has been found by Almeida and Vicente (2009) but this evidence is still mixed and based on a rather short dataset, because the focus of this paper is not specifically forecasting. Favero et al. (2007) and Moench (2009) document a rather good performance of the ATSM models, but in both cases such models are complemented with a large macroeconomic dataset.

A third group of papers uses the Nelson and Siegel (1987) exponential components framework (Diebold and Li, 2006), possibly also imposing on it the no-arbitrage restrictions (Christensen et al. 2007). The forecasting results obtained by these models are better, with the Diebold and Li (2006) model producing one-year-ahead forecasts that are noticeably more accurate than standard benchmarks. Still, the gains are small at shorter forecast horizons.

In this paper we propose a new strategy for forecasting the term structure of interest

rates, which exploits the information contained in a large panel of government bond yields. Focusing on the U.S., we show that our proposed strategy produces systematically better forecasts than all the methods outlined above. The starting point of our strategy is the consideration that the yield curve can be thought of as a vector process composed of yields of different maturities. In that light, a straightforward approach to forecast would be to simply run a Vector Autoregression (*VAR*). However, such a strategy soon encounters the so-called "curse of dimensionality" problem, as the number of parameters to estimate rapidly reduces the degrees of freedom of the *VAR* system. As a result, the forecasts produced by a *VAR* are typically poor. To overcome this difficulty, one can either chose to sacrifice completely the cross-sectional information, and estimate e.g. a simple AR model on each yield, or try to summarize the information in an efficient but manageable manner. This latter possibility can be pursued by using a Bayesian Vector Autoregression (*BVAR*).

A *BVAR* is a *VAR* whose coefficients are random variables on which the researcher can impose some a-priori information. Without entering into philosophical disputes about the Bayesian and the classical approach in econometrics, we think it is worth stressing here that *BVARs* can be taught of simply as a selection device. Consider the first equation of a large *VAR*: there are many regressors, and the researcher needs to solve the trade-off between using as much information as possible, and the loss in degrees of freedom coming from having too many parameters to estimate. An intuitive way to proceed would be to start with an empty model, then adding a candidate regressor and performing a test of significance for that regressor. If the null is rejected, then the regressor is kept. The procedure can be repeated until the last candidate regressor is considered. What this procedure implicitly does is selecting the regressors on the basis of how much valuable information they contain. Information is valuable if it is able to significantly increase the likelihood of the model. The Bayesian algorithm works similarly. A priori the coefficient attached to a given candidate regressor is set to 0, and only if the information contained in the data is valuable enough to influence the likelihood the posterior mean will be far from 0. More precisely, rather than acting as a selection device which either includes or excludes a regressor, the *BVAR* includes all the regressors but it assigns a different weight to each. The weight is higher the higher is the informational content of a given regressor. The advantage of the *BVAR* over the simple step-wise procedure outlined above is that the former is a fully blown approach grounded on statistical theory, while the latter is not, and as a result it might lead to incorrect inference. For example, in the simple step-wise procedure the order with which the various candidate regressors are examined can significantly influence the final outcome, and the overall size of

the step-wise procedure is unknown.

BVARs have a long story in econometrics. Although the good forecasting performance of *BVARs* has been documented years ago by Litterman (1986) and Doan et al. (1984), only recently they have started to be used more systematically for policy analysis and forecasting macroeconomic variables (Kadiyala and Karlsson 1997, Banbura et al. 2007, Carriero et al. 2009, 2010). One of the major stumbling blocks that prevented the use of *BVARs* as a model for forecasting and policy analysis has typically been the large computational burden they pose. Indeed, the computation of nonlinear functions of the parameters such as impulse-response functions and multi-step forecasts need to be performed via time consuming simulations. As we will discuss below in detail, in this paper we solve this problem by directly estimating the relevant coefficients for each forecast horizon, which allows us to compute the forecasts at all forecast horizon without resorting to simulation. As a result, the production of a full set of (Bayesian) forecasts for horizon 1- to 12- month ahead takes seconds.

We compare our proposed approach against all the major forecasting models used so far in the literature, including forward rate regression (Fama and Bliss (1987), Cochrane and Piazzesi (2005)), Affine term structure Models (Ang and Piazzesi (2003)), and models based on the exponential components framework (Diebold-Li (2006), Christensen et al. (2007), Stock and Watson (2002a,b)). For reference, we include in the comparison also a set of linear models (Random Walk, Autoregressive Models, Vector Autoregressions). De Pooter et al. (2007) also propose to use Bayesian techniques to forecast the term structure, but their paper differs from this as they use Bayesian model averaging over some term structure models, while we use a large *BVAR* to extract efficiently the information contained in a large cross section of yields.

Besides introducing the new approach to forecasting the yields, we extend the available empirical evidence in three directions: First, for all the models considered, we provide results for an homogeneous dataset. We use the novel dataset by Gurkaynak et al. (2007), publicly available on the website <http://www.federalreserve.gov> and updated regularly. Secondly, while most of the existing evidence evaluates forecast accuracy only in terms of statistical measures such as Root Mean Squared Forecast Errors, we also evaluate forecasts on the basis of "economic" criteria. In particular, we provide Sharpe Ratios arising from simple trading rules based on the alternative forecasts, and we use the alternative forecasts to perform optimal portfolio allocation. Finally, we provide extensive robustness checks of the results. In particular, we provide subsample results and a simulation study in which we simulate and forecast a set of "artificial" term structures. In doing such a Monte Carlo simulation

the researcher has to choose the Data Generating Process (DGP). Obviously the choice of a particular DGP over another would influence the results, advantaging one rather than the other model. Therefore, rather than concentrate on simulated data and an inevitably arbitrary data generation process, we carry out our simulation by bootstrapping the actual term structure dataset. The use of a real dataset as a basis for such a robustness analysis is referred to as a ‘data based Monte Carlo method’ and discussed further in, e.g., Ho and Sørensen (1996).

We find that: i) our proposed BVAR approach produces forecasts systematically more accurate than the random walk forecasts, though the gains are small; ii) some models beat the BVAR for a few selected maturities and forecast horizons, but they perform much worse than the BVAR in the remaining cases; iii) predictive gains with respect to the random walk have decreased over time; iv) different loss functions (i.e., "statistical" vs "economic") lead to different ranking of specific models; v) modelling time variation in term premia is important and useful for forecasting.

The paper is structured as follows. Section 2 develops our BVAR approach. Section 3 introduces the competing forecasting models under comparison. Section 4 describes the data, the forecasting exercise and the alternative criteria we shall use in evaluating the alternative forecasts. Section 5 presents the main results, and Section 6 the robustness analysis. Finally, section 7 concludes.

2 Bayesian VARs (BVAR)

The yield curve can be thought of as a vector process composed of yields of different maturities. In that light, a straightforward approach to forecasting it, is to simply run a Vector Autoregression. However, such a strategy encounters an overparameterization problem, as the number of estimated parameters rapidly reduces the degrees of freedom of the VAR system. As a result, the forecasts produced by a VAR are typically poor. To overcome this difficulty, one can either chose to sacrifice completely the cross-sectional information, and estimate, e.g., a simple AR model on each yield, or try to summarize the information in an efficient but manageable manner. This latter possibility can be pursued by modelling the system of yields as a multivariate model in which yields are a-priori following a univariate autoregressive process.

2.1 Notation and Preliminaries

We denote the log-price of a τ -year discount bond at time t as $p_t^{(\tau)}$, so that the log yield is $y_t^{(\tau)} = -\frac{1}{\tau}p_t^{(\tau)}$. The log forward (per-period) rate at time t for loans between $t+h$ and $t+h+\tau$ is given by $f_t(h, \tau) = \frac{1}{\tau}(-p_t^{(\tau+h)} + p_t^{(\tau)}) = \frac{1}{\tau}((\tau+h)y_t^{(\tau+h)} - \tau y_t^{(\tau)})$. The log holding period return is given by $r_{t+1}^{(\tau)} = p_{t+1}^{(\tau-1)} - p_t^{(\tau)}$. We have a cross-section of N different maturities which we denote $\tau_1, \tau_2, \dots, \tau_N$. Whenever we consider vectors of yields, returns, forward rates, or prices we use bold letters, e.g. $\mathbf{y}_t = (y_t^{(\tau_1)} \ y_t^{(\tau_2)} \ \dots \ y_t^{(\tau_N)})'$ is the N -dimensional vector of yields for each of the selected maturities. The h -step ahead forecasts (made at time t) are denoted with a hat, e.g. the h -step ahead forecast of the vector of yields is $\hat{\mathbf{y}}_{t+h}$.

Consider the following Vector Autoregression:

$$\mathbf{y}_t = \Phi_{0,h} + \Phi_{1,h}\mathbf{y}_{t-h} + \varepsilon_t; \ \varepsilon_t \sim i.i.d.N(0, \Sigma). \quad (1)$$

Note that in the above model the vector of yields \mathbf{y}_t is regressed directly onto \mathbf{y}_{t-h} , which means that for each forecast horizon h a different model is employed. Such an approach, is known as "direct" forecasting, and it focuses on minimizing the relevant loss function for each forecast horizon, i.e. the h -step ahead forecast error. Alternatively, one could use the traditional "powering up" approach, which consists in regressing \mathbf{y}_t onto \mathbf{y}_{t-1} , and then to compute recursively the h -step ahead forecasts. For a discussion and a comparison of these alternative methods see, e.g., Marcellino et al. (2006) and Pesaran et al. (2009). In general, the powering up approach is more efficient, as it uses more datapoints, but is dangerous in the presence of misspecification, because the misspecification will inflate with the forecast horizon when the forecasts are computed recursively. On the other side, the direct approach is less efficient but is more robust to misspecification. More importantly for our purposes, the direct approach implies that the h -step ahead forecast is a linear function of the coefficients: $\hat{\mathbf{y}}_{t+h} = \hat{\Phi}_{0,h} + \hat{\Phi}_{1,h}\mathbf{y}_t$, while in the traditional powering up approach the multi step forecasts will be a highly nonlinear function of the estimated coefficients. As a result, there is an exact closed form solution for the distribution of the h -step ahead forecasts computed using (1), while computing the forecasts resulting from the powering up strategy would require the use of time-demanding simulation methods.

2.2 A Normal-Inverted Wishart AR(1) Prior

It is clear that yields, regardless of maturity, are very persistent processes. Indeed both a simple Auto Regressive (AR) model and the Random Walk (RW) produce very good forecasts of the yields. Therefore it is reasonable to think a priori that each of the yields in (1) obeys

a univariate AR with high persistence, or equivalently, that the expected value of the matrix $\Phi_{1,h}$ is $E[\Phi_{1,h}] = 0.99 \times I$. We also need to assess how strong is the belief we have in such a prior, i.e. we need to set a variance around the prior mean. More formally, we assume that $\Phi_{1,h}$ is conditionally (on Σ) normal, with first and second moments given by:

$$E[\Phi_{1,h}^{(ij)}] = \begin{cases} 0.99 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}, \quad Var[\Phi_{1,h}^{(ij)}] = \theta \sigma_i^2 / \sigma_j^2, \quad (2)$$

where $\Phi_{1,h}^{(ij)}$ denotes the element in position (i, j) in the matrix $\Phi_{1,h}$, and where the covariances among the coefficients in $\Phi_{1,h}$ are zero. The shrinkage parameter θ measures the tightness of the prior: when $\theta \rightarrow 0$ the prior is imposed exactly and the data do not influence the estimates, while as $\theta \rightarrow \infty$ the prior becomes loose and the prior information does not influence the estimates, which will approach the standard *OLS* estimates. We will discuss in detail the choice of this parameter below, since it is a key element of our proposal. The factor σ_i^2 / σ_j^2 is a scaling parameter which accounts for the different scale and variability of the data. To set the scale parameters σ_i^2 we follow common practice (see e.g. Litterman, 1986; Sims and Zha, 1998) and set it equal to the variance of the residuals from a univariate autoregressive model for the variables. The prior specification is completed by assuming a diffuse normal prior on $\Phi_{0,h}$ and an Inverted Wishart prior for the matrix of disturbances $\Sigma \sim IW(v_0, S_0)$, where v_0 and S_0 are the prior scale and shape parameters, and are set such that the prior expectation of Σ is equal to a fixed diagonal residual variance $E[\Sigma] = diag(\sigma_1^2, \dots, \sigma_N^2)$. This provides us with a Normal-Inverted Wishart prior, which is conjugate, i.e. it features a Normal-Inverted Wishart posterior distribution.

We can re-write more compactly the VAR as:

$$Y = X_h \Psi_h + E, \quad (3)$$

where $Y = [y_{h+1}, \dots, y_T]'$ is a $T \times N$ matrix containing all the data points in y_t , $X_h = [1 \ Y_{-h}]$ is a $T \times M$ matrix containing a vector of ones (1) in the first columns and the h -th lag of Y in the remaining columns, $\Psi_h = [\Phi_{0,h} \ \Phi_{1,h}]'$ is a $M \times N$ matrix, and $E = [\varepsilon_{h+1}, \dots, \varepsilon_T]'$ is a $T \times N$ matrix of disturbances. As only one lag is considered we have $M = N + 1$. The prior distribution can then be written as:

$$\Psi_h | \Sigma \sim N(\Psi_0, \Sigma \otimes \Omega_0), \quad \Sigma \sim IW(v_0, S_0). \quad (4)$$

Note that $\Psi_h | \Sigma$ is a matrix-variate normal distribution where the prior expectation $E[\Psi_h] = \Psi_0$ and prior variance $Var[\Psi_h] = \Sigma \otimes \Omega_0$ are set according to equation (2). The prior variance matrix has a Kronecker structure $Var[\Psi_h] = \Sigma \otimes \Omega_0$ where Σ is the variance matrix of the

disturbances and the elements of Ω_0 are given by $Var[\Phi_{1,h}^{(ij)}]$ in (2). As the used N-IW prior is conjugate, the conditional posterior distribution of this model is also Normal-Inverted Wishart (Zellner 1973):

$$\Psi_h|\Sigma, Y \sim N(\bar{\Psi}, \Sigma \otimes \bar{\Omega}), \quad \Sigma|Y \sim IW(\bar{v}, \bar{S}), \quad (5)$$

where the bar denotes that the parameters are those of the posterior distribution. Defining $\hat{\Psi}$ and \hat{E} as the OLS estimates, we have that $\bar{\Psi} = (\Omega_0^{-1} + X'X)^{-1}(\Omega_0^{-1}\Psi_0 + X'Y)$, $\bar{\Omega} = (\Omega_0^{-1} + X'X)^{-1}$, $\bar{v} = v_0 + T$, and $\bar{S} = \hat{\Psi}'X'X\hat{\Psi} + \Psi_0'\Omega_0^{-1}\Psi_0 + \Psi_0 + \hat{E}'\hat{E} - \hat{\Psi}'\bar{\Omega}^{-1}\hat{\Psi}$.

2.3 Estimation

In order to perform inference and forecasting one needs the full joint posterior distribution and the marginal distributions of the parameters $\bar{\Psi}$ and Σ . One could use the conditional posteriors in (5) as a basis of a Gibbs sampling algorithm that drawing in turn from the conditionals $\Psi_h|\Sigma, Y$ and $\Sigma|Y$ would eventually produce a sequence of draws from the joint posterior $\Psi_h\Sigma|Y$ and the marginal posteriors $\Psi_h|Y, \Sigma|Y$, as well as the posterior distribution of any function of these coefficients (e.g. multi-step forecasts or impulse responses).

Still, if one is interested only in the posterior distribution of Ψ_h (rather than in any nonlinear function of it) there is an alternative to simulation: by integrating out Σ from (5) Zellner (1973) has shown that the marginal posterior distribution of Ψ_h is a matricvariate t :

$$\Psi_h|Y \sim MT(\bar{\Psi}, \bar{\Omega}^{-1}, \bar{S}, \bar{v}). \quad (6)$$

The expected value for this distribution is given by:

$$\bar{\Psi} = (\Omega_0^{-1} + X'X)^{-1}(\Omega_0^{-1}\Psi_0 + X'Y), \quad (7)$$

which is obviously extremely fast to compute. Recalling that $\hat{\Psi}$ is the OLS estimator, and using the normal equations $(X'X)^{-1}\hat{\Psi} = X'Y$ we can rewrite this as:

$$\bar{\Psi} = (\Omega_0^{-1} + X'X)^{-1}(\Omega_0^{-1}\Psi_0 + X'X\hat{\Psi}), \quad (8)$$

which shows that the posterior mean of Ψ_h is a weighted average of the OLS estimator and of the prior mean Ψ_0 , with weights proportional to the inverse of their respective variances. In the presence of a tight prior (i.e., when $\theta \rightarrow 0$) the posterior estimate will collapse to $\bar{\Psi} = \Psi_0$, while with a diffuse prior (i.e., when $\theta \rightarrow \infty$) the posterior estimate will collapse to the unrestricted OLS estimate.

Given the posterior mean $\bar{\Psi} = [\bar{\Phi}_{0,h} \ \bar{\Phi}_{1,h}]'$, it is straightforward to produce forecasts up to h steps ahead simply by setting:

$$\hat{\mathbf{y}}_{t+h} = \bar{\Phi}_{0,h} + \bar{\Phi}_{1,h}\mathbf{y}_t, \quad (9)$$

As shown by Banbura et al. (2009) it is also possible to implement the prior described above using a set of dummy observations. Consider adding T_d dummy observations Y_d and X_d such that their moments coincide with the prior moments: $\Psi_0 = (X_d'X_d)^{-1}X_d'Y_d$, $\Omega_0 = (X_d'X_d)^{-1}$, $v_0 = T_d - M - N - 1$, $S_0 = (Y_d - X_d\Psi_0)'(Y_d - X_d\Psi_0)$. Augmenting the system in (3) with the dummy observations gives:

$$Y^+ = X_h^+\Psi_h + E^+, \quad (10)$$

where $Y^+ = (Y' \ Y_d)'$ and $E^+ = (E' \ E_d)'$ are $(T + T_d) \times N$ matrices and $X^+ = (X' \ X_d)'$ is a $(T + T_d) \times M$ matrix. Then it is possible to show that the OLS estimator of the augmented system (given by the usual formula $(X_h^{+'}X_h^+)^{-1}X_h^{+'}Y^+$) is numerically equivalent to the posterior mean $\bar{\Psi}$.

2.4 Prior tightness

To make the prior operational, one needs to choose the value of the hyperparameters θ . The marginal data density of the model can be obtained by integrating out all the coefficients, i.e., defining Θ as the set of all the coefficients of the model, the marginal data density is:

$$p(Y) = \int p(Y|\Theta)p(\Theta)d\Theta. \quad (11)$$

Under our Normal-Inverted Wishart prior the density $p(Y)$ can be computed in closed form (see Bauwens et al. 1999). At each point in time we choose θ by maximizing:

$$\theta_t^* = \arg \max_{\theta} \ln p(Y) \quad (12)$$

Figure 2 depicts the value of $\ln p(Y)$ for the case with $h = 1$ as a function of θ and time. The results are based on the dataset of Gurkaynak et al. (2007) which is displayed in Figure 1, and which we will describe in detail in Section 4. We consider the values of θ_t in the grid $\theta \in \{2e-16, 4e-16, 6e-16, 8e-16, 1e-15, 0.00001, 0.0001, 0.001, 0.01, 0.1, 1, 10\}$. Interestingly, it turns out that for all forecast horizons $p(Y)$ is a hump-shaped function¹ of θ , which shows that

¹Figure 2 depicts the natural logarithm $\ln p(y)$, which is rather flat for intermediate values of θ , and it might seem to signal that the function is not sensitive to this parameter. However, this is a feature of the log-transformation, and the surface $p(y)$ is much steeper.

imposing some shrinkage (as opposed to no shrinkage or infinite shrinkage) is optimal, with the optimal value being around 0.00001. Such value looks rather low, and it implies a rather high amount of shrinkage. However one should bear in mind that the value of the shrinkage parameter θ^* is not interpretable per se but only in relation to the dimension of the system. The same value for the shrinkage parameter would imply different amounts of shrinkage if it is applied to a small or to a large system. De Mol et al. (2009) have formalized this point and have shown that the optimal shrinkage parameter θ^* goes to 0 when the cross-sectional size of the system increases.

As we estimate a different model for each forecast horizon h , we will select an optimal θ_t^* for each h . The time series path of the optimal θ_t^* for $h = 1, \dots, 12$ is depicted in Figure 3. As is clear from the figure, the optimal θ_t^* always hovers around 0.00001 for all forecast horizons, and only for long horizons it becomes occasionally higher. We have also experimented based on a θ kept fixed throughout the sample to 0.00001 and found that such a strategy produces very similar results to those obtained using the optimally selected θ_t^* .

3 Forecasting Models

In this section we describe all the competing forecasting models under evaluation. The models include linear models (autoregressive models and vector Autoregressions), forward rate regressions (along the lines of Fama and Bliss (1987) and Cochrane and Piazzesi (2005)), factor models (along the lines of Diebold and Li (2006), Christensen et al. (2007), and Stock and Watson (2002a,b)), and affine term structure models (along the lines of Ang and Piazzesi (2003)). Our proposed model, i.e. the Bayesian VAR, has been described in Section 2.

3.1 Random Walk (RW)

We include a simple Random Walk forecast (RW). The RW forecast of the yield of maturity τ at time $t + h$ is given by:

$$\hat{y}_{t+h}^{(\tau)} = y_t^{(\tau)} \tag{13}$$

Duffee (2002) and Diebold and Li (2006) have shown that beating a Random Walk forecast of the yield curve is difficult, therefore we will use the RW forecasts as the benchmark with respect to which we will compare the forecasts of all the competing models.

3.2 Univariate Autoregressive model (AR)

We also consider univariate autoregressive models. Forecasts from such process can be produced in two alternative ways . The first way is known as powering-up approach. In this approach the econometrician estimates the model:

$$y_t^{(\tau)} = \alpha + \beta y_{t-1}^{(\tau)} + \varepsilon_t \quad (14)$$

for the generic maturity τ . The 1-step ahead forecast is produced as $\hat{y}_{t+1}^{(\tau)} = \hat{\alpha} + \hat{\beta} y_t^{(\tau)}$, while the forecasts for h -step ahead horizon are obtained as:

$$\hat{y}_{t+h}^{(\tau)} = \hat{\alpha} + \hat{\beta} \hat{y}_{t+h-1}^{(\tau)}.$$

Alternatively, one can use the direct approach. As we discussed before, this approach directly optimizes the relevant loss function, i.e. the h -step ahead mean squared forecast error and has proved to be more robust but less efficient see, e.g., Marcellino, Stock and Watson (2006) and Pesaran et al (2009). The estimated model is:

$$y_t^{(\tau)} = \alpha_h + \beta_h y_{t-h}^{(\tau)} + \varepsilon_t, \quad (15)$$

i.e. the variable at time t is projected directly onto its past value in period $t - h$, so the estimated coefficients summarize the h -step ahead effect. Note that a different α_h and β_h are obtained for each forecast horizon. The forecasts are then derived as:

$$\hat{y}_{t+h}^{(\tau)} = \hat{\alpha}_{\tau h} + \hat{\beta}_{\tau h} y_t^{(\tau)} \quad (16)$$

We provide results for both these approaches, labeling them respectively $AR(pu)$ and $AR(di)$. Note that for the 1-step ahead horizon the two methods produce the same results.

Another relevant dimension of choice when using autoregressive models is the choice of the lag length. The models described above are AR(1), but in principle one could choose any lag length p . The specification with one lag is consistent with the prior mean of our BVAR, which we have described in Section 2. We have also considered richer dynamic specifications and re-estimated both the AR models (with power-up and direct approach) by choosing the lag length by means of the Bayesian Information Criterion (BIC). The chosen lag length oscillates between one and four depending on the time period, but the final results are close to those obtained with the simpler AR(1) specifications. Therefore, we present the latter results only, while we have the former available upon request.

3.3 Vector Autoregression (VAR)

VAR forecasts are produced along the same lines of univariate AR forecasts. Specifically, VAR forecasts based on powering up are produced as follows. The regression model is:

$$\mathbf{y}_t = \mathbf{A} + \mathbf{B}\mathbf{y}_{t-1} + \varepsilon_t, \quad (17)$$

where $\mathbf{y}_t = (y_t^{(\tau_1)} \ y_t^{(\tau_2)} \ \dots \ y_t^{(\tau_N)})'$. The 1-step ahead forecast is produced as $\hat{\mathbf{y}}_{t+1} = \hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{y}_t$, while the h -step ahead forecasts are obtained as:

$$\hat{\mathbf{y}}_{t+h} = \hat{\mathbf{A}} + \hat{\mathbf{B}}\hat{\mathbf{y}}_{t+h-1} \quad (18)$$

For the direct approach, forecasts are obtained as follows. The regression model is:

$$\mathbf{y}_t = \mathbf{A}_h + \mathbf{B}_h\mathbf{y}_{t-h} + \varepsilon_t, \quad (19)$$

where $\mathbf{y}_t = (y_t^{(\tau_1)} \ y_t^{(\tau_2)} \ \dots \ y_t^{(\tau_N)})'$. Note that different \mathbf{A}_h and \mathbf{B}_h matrices are used for each forecast horizon. The forecasts are then derived as:

$$\hat{\mathbf{y}}_{t+h} = \hat{\mathbf{A}}_h + \hat{\mathbf{B}}_h\mathbf{y}_t \quad (20)$$

We provide results for both these approaches, labeling them respectively. $VAR(pu)$ and $VAR(di)$. Note that for the 1-step ahead horizon the two methods produce the same results. We have also experimented with longer lag specifications and BIC lag selection. This provides rather poor results which are available upon request.

3.4 Fama and Bliss (1987) (FB)

Fama and Bliss (1987) found that the spread between the τ -year forward rate and the one-year yield predicts the one-year excess return of the τ -year bond, with R^2 of about 18 percent. The natural extension of the Fama-Bliss approach in our context uses the following regression model for each maturity τ :

$$y_t^{(\tau)} - y_{t-h}^{(\tau)} = \alpha_h + \beta_h(f_{t-h}(h, \tau) - y_{t-h}^{(\tau)}) + \varepsilon_t, \quad (21)$$

where $f_t(h, \tau)$ as the forward rate at time t for loans between time $t+h$ and $t+h+\tau$. The h -step ahead forecasts $\hat{y}_{t+h}^{(\tau)}$ are then produced by projecting the change in yields from time t to time $t+h$ on the forward-spot spread at time t :

$$\hat{y}_{t+h}^{(\tau)} - y_t^{(\tau)} = \hat{\alpha}_{\tau h} + \hat{\beta}_{\tau h}(f_t(h, \tau) - y_t^{(\tau)}) \quad (22)$$

3.5 Cochrane and Piazzesi (2005) (CP)

Cochrane and Piazzesi (2005) show that the excess returns on one- to five-year maturity bonds are predicted (with R^2 up to 0.44) by a single tent-shaped linear combination of the corresponding five forward rates $y_t^{(\tau=12)}$, $f_t(12, 12)$, $f_t(24, 12)$, $f_t(36, 12)$, and $f_t(48, 12)$. Although the natural extension of this approach would imply putting all the available forward rates on the right hand side, this would lead to multicollinearity, as the Gurkaynak et al. (2007) data are based on a six-parameter functional form. Therefore, we follow Cochrane and Piazzesi (2008), and use only three forward rates, $y_t^{(\tau_{12})}$, $f_t(24, \tau_{12})$, and $f_t(48, \tau_{12})$, in order to eliminate problems of collinearity. The forecasts are obtained as follows. For each maturity τ , we estimate the following linear regression model:

$$y_t^{(\tau)} - y_{t-h}^{(\tau)} = \alpha_h + \beta_h y_{t-h}^{(12)} + \gamma f_{t-h}(24, 12) + \delta f_{t-h}(48, 12) + \varepsilon_t \quad (23)$$

Then, forecasts are computed as:

$$\hat{y}_{t+h}^{(\tau)} - y_t^{(\tau)} = \hat{\alpha}_h + \hat{\beta}_h y_t^{(\tau_{12})} + \hat{\gamma} f_t(24, 12) + \hat{\delta} f_t(48, 12) \quad (24)$$

3.6 Dynamic Nelson and Siegel model (Diebold-Li (2006), DL)

Diebold and Li (2006) proposed a dynamic version of the Nelson and Siegel (1987) exponential components framework. They obtained very good results in terms of out of sample yield curve forecasting. The starting point is the following yield curve interpolant:

$$y_t^{(\tau)} = \beta_1 + \beta_2 \frac{1 - e^{-\gamma\tau}}{\gamma\tau} + \beta_3 \left(\frac{1 - e^{-\gamma\tau}}{\gamma\tau} - e^{-\gamma\tau} \right), \quad (25)$$

where $\beta_1, \beta_2, \beta_3$ and γ are parameters. Diebold and Li (2006) interpreted equation (25) in a dynamic fashion as a latent factor model in which β_1, β_2 , and β_3 are time-varying level, slope, and curvature factors and the terms that multiply these factors are factor loadings:

$$y_t^{(\tau)} = \beta_{1t} + \beta_{2t} \frac{1 - e^{-\gamma\tau}}{\gamma\tau} + \beta_{3t} \left(\frac{1 - e^{-\gamma\tau}}{\gamma\tau} - e^{-\gamma\tau} \right). \quad (26)$$

The cross section of yields can be used to estimate via OLS² the coefficients $\beta_{1t}, \beta_{2t}, \beta_{3t}$ at each point in time. Repeating the OLS regression across sections at each point in time

²To do so we follow Diebold and Li (2006) and we fix the value of γ to 0.0609, which is the value that maximizes the loading on β_3 at exactly 30 months maturity. As stressed out by Diebold and Li this choice enhances numerical efficiency and robustness. An alternative approach could be to estimate the model with nonlinear least squares or Maximum Likelihood.

provides a time series of estimates of the three factors $\tilde{\beta}_{1t}, \tilde{\beta}_{2t}, \tilde{\beta}_{3t}$. Then, factors are assumed to follow a univariate autoregressive process and are forecasted as follows :

$$\hat{\beta}_{jt+h} = \hat{a}_0 + \hat{a}_1 \tilde{\beta}_{jt} \quad (27)$$

where \hat{a}_0 and \hat{a}_1 are the coefficients of a regression of $\tilde{\beta}_{jt}$ onto $\tilde{\beta}_{jt-h}$. Alternatively, one can assume a VAR structure for the factors, but such strategy typically produces a worse forecasting performance given that it increases the number of parameters to be estimated without having much to gain in terms of cross-variable interaction. For completeness we have also tried a VAR(1) specification as in Diebold et Li (2006). We label respectively $DL(AR)$ and $DL(VAR)$ the forecasts computed using a univariate AR and a VAR specification for the factors.

Once forecasts of the factors at time $t+h$ are available, the forecast of the yields can be retrieved simply by exploiting again the cross sectional dimension of the system:

$$\hat{y}_{t+h}^{(\tau)} = \hat{\beta}_{1t+h} + \hat{\beta}_{2t+h} \frac{1 - e^{-\gamma\tau}}{\gamma\tau} + \hat{\beta}_{3t+h} \left(\frac{1 - e^{-\gamma\tau}}{\gamma\tau} - e^{-\gamma\tau} \right). \quad (28)$$

The DL model described above has proved to be empirically successful but it is based on a purely statistical approach not grounded on finance theory. In a recent paper Christensen et al. (2007) have shown how the No-Arbitrage restrictions can be imposed onto the DL model. We have also tried this specification, and we found that although it is at least as good as the standard unrestricted DL specification for small sized systems, it is systematically outperformed in systems with a large cross section of yields.

3.7 Stock and Watson (2002a,b) (SW)

Stock and Watson (2002a,b) propose to overcome the curse of dimensionality problem arising in forecasting with large datasets by using a factor structure. The information contained in a set of predictors X_t is summarized by a set of K factors:

$$X_t = \Gamma F_t + u_t \quad (29)$$

where F_t is a K -dimensional multiple time series of factors and Γ a $N \times K$ matrix of loadings. The forecast for $\hat{y}_{t+h}^{(\tau)}$ given the predictors can be obtained through a two-step procedure, in which in the first step the set of predictors $\{X_t\}_{t=1}^T$ is used to estimate the factors $\{\hat{F}_t\}_{t=1}^T$ via principal components, and then the forecasts are obtained by projecting $\hat{y}_{t+h}^{(\tau)}$ onto \hat{F}_t and $y_t^{(\tau)}$. In the case at hand the set of predictors used in the first step comprises the yields of all

the remaining maturities. Once the factors have been derived via principal components, we estimate the following regression model:

$$y_t^{(\tau)} = \Theta_h \hat{F}_{t-h} + \beta_h y_{t-h}^{(\tau)} + \varepsilon_t, \quad (30)$$

and then the forecasts are obtained by projecting:

$$\hat{y}_{t+h}^{(\tau)} = \hat{\Theta}_h \hat{F}_t + \hat{\beta}_h y_t^{(\tau)} \quad (31)$$

Stock and Watson (2002a,b) develop theoretical results for this two-step procedure and show that under a set of moment and rank conditions the MSFE of the feasible forecast asymptotically approaches that of the optimal infeasible forecast (i.e. the one based on the knowledge of the factors) for N and T approaching infinity, see Bai and Ng (2006) for additional details.

3.8 Affine Term Structure Models (ATSM)

We consider the model proposed by Ang and Piazzesi (2003), which is a discrete-time version of the affine class introduced by Duffie and Kan (1996), where bond prices are exponential affine functions of underlying state variables. The assumption of no arbitrage (Harrison and Kreps, 1979) guarantees the existence of a risk neutral measure Q such that the price at time t of an asset V_t that does not pay any dividends at time $t+1$ satisfies $V_t = E_t^Q [\exp(-i_t)V_{t+1}]$, where the expectation is taken under the measure Q and i_t is the short term rate. The assumption of no arbitrage is equivalent to the assumption of the existence of the Radon-Nikodym derivative ξ_{t+1} , which allows to convert the risk neutral measure to the data generating measure: $E_t^Q [\exp(-i_t)V_{t+1}] = E_t [(\xi_{t+1}/\xi_t) \exp(-i_t)V_{t+1}]$. Assume ξ_{t+1} follows a log-normal process:

$$\xi_{t+1} = \xi_t \exp(-0.5\Lambda_t' \Lambda_t - \Lambda_t' \varepsilon_{t+1}). \quad (32)$$

The Λ_t are called market prices of risk and are an affine function of a vector of k factors F_t :

$$\Lambda_t = \Lambda_0 + \Lambda_1 F_t, \quad (33)$$

where Λ_0 is a k -dimensional vector and Λ_1 a $k \times k$ matrix. Note that by setting $\Lambda_1 = 0$ prices of risk are constant over time. Also the short term rate is assumed to be an affine function of F_t :

$$i_t = \delta_0 + \delta_1' F_t, \quad (34)$$

where δ_0 is a scalar and δ_1 a k -dimensional vector. We assume the factors follow a zero-mean stationary vector process:

$$F_t = \Psi F_{t-1} + \Omega \varepsilon_t, \quad (35)$$

where $\varepsilon_t \sim iidN(0, \Sigma_\varepsilon)$ with $\Sigma_\varepsilon = I$ with no loss of generality. The nominal pricing kernel is defined as:

$$m_{t+1} = \exp(-i_t)\xi_{t+1}/\xi_t = \exp(-\delta_0 - \delta'_1 F_t - 0.5\Lambda'_t \Lambda_t - \Lambda'_t \varepsilon_{t+1}), \quad (36)$$

where the second equality is obtained by using (34) and (32). The nominal pricing kernel prices all assets in the economy, so by letting $p_t^{(\tau)}$ denote the time t price of a τ -period zero coupon we have:

$$p_t^{(\tau+1)} = E_t(m_{t+1} p_{t+1}^{(\tau)}). \quad (37)$$

By using the above equations is possible to show that bond prices are an affine function of the state variables:

$$p_t^{(\tau)} = \exp(\bar{A}_\tau + \bar{B}'_\tau F_t), \quad (38)$$

where \bar{A}_τ and \bar{B}_τ are a scalar and a k -dimensional vector obeying to:

$$\begin{aligned} \bar{A}_{\tau+1} &= \bar{A}_\tau + \bar{B}'_\tau (-\Omega \Lambda_0) + 0.8'_\tau \Omega \Omega' \bar{B}_\tau - \delta_0 \\ \bar{B}'_{\tau+1} &= \bar{B}'_\tau (\Psi - \Omega \Lambda_1) - \delta'_1 \end{aligned}, \quad (39)$$

with $\bar{A}_1 = -\delta_0$ and $\bar{B}_1 = -\delta_1$. See Ang and Piazzesi (2003) for a derivation. The continuously compounded yield on a τ -period zero coupon bond is:

$$y_t^{(\tau)} = -\ln p_t^{(\tau)}/\tau = A_\tau + B'_\tau F_t, \quad (40)$$

with $A_\tau = -\bar{A}_\tau/\tau$ and $B_\tau = -\bar{B}_\tau/\tau$, so yields are also an affine function of the factors. Equations (35) and (40) can be interpreted as a transition and a measurement equation respectively, so the state space model for a vector of yields of q different maturities $\tau_1, \tau_2, \dots, \tau_q$ can be written as:

$$\begin{aligned} F_t &= \Psi F_{t-1} + \Omega \varepsilon_t, \\ Y_t &= A + B F_t + v_t, \end{aligned} \quad (41)$$

where $Y_t = (y_t^{(\tau_1)}, y_t^{(\tau_2)}, \dots, y_t^{(\tau_q)})'$ is a q -dimensional vector process collecting all the yields at hand, $A = (A_{\tau_1}, A_{\tau_2}, \dots, A_{\tau_q})'$ and $B = (B_{\tau_1}, B_{\tau_2}, \dots, B_{\tau_q})'$ are functions of the structural coefficients of the model according to equation (39), and where v_t is a vector of iid Gaussian measurement errors with variance Σ_v .

Following common practice we use three factors, which can be interpreted as the level, slope and curvature of the term structure. Given that scaling, shifting, or rotation of the factors provides observational equivalence, a normalization is required. Following Dai and Singleton (2000) we identify the factors by assuming factor mean equal to zero, a lower triangular structure for the matrix Ψ , and we set $\delta_1 = (1, 1, 0)'$. Given this identification

scheme the coefficient δ_0 equals the unconditional mean of the instantaneous rate, which can be approximated by the sample average of the 1-month yield. As for second order coefficients, we assume Ω and Σ_v to be diagonal, while we assume absence of correlation between ε_t and v_t , i.e. $E(\varepsilon_t v_t) = 0$.

We collect all the structural parameters to be estimated in the vector:

$$\vartheta = \{\Psi, \Omega, \Lambda_0, \Lambda_1, \Sigma_v\}. \quad (42)$$

In our application, we denote this model by $ATSM(tp)$, to emphasize that the model features time varying term premia. We also consider a specification which imposes constant risk premia, and it is obtained simply by setting $\Lambda_1 = 0$ in equation (33), and we label this specification $ATSM(cp)$. We estimate ϑ with the Expectation-Maximization (EM) algorithm, evaluating the likelihood at each iteration by means of the Kalman Filter. We maximize the likelihood of the VAR with no-arbitrage restrictions using the Broyden, Fletcher, Goldfarb, and Shanno (BFGS) algorithm with Brent line search. ATSMs are highly nonlinear in the parameters. Maximum likelihood estimation of such models involves a difficult optimization problem, as the likelihood features several local maxima and results are very sensitive to the initial conditions. A strategy that is often adopted for estimating such models is to start from several different initial conditions, find the maximum likelihood estimate for each of them, and finally pick the highest among the detected local maxima (see Duffee 2007). In our exercise we initialize our algorithm as follows. In the first estimation window we compute a maximum, then we draw 100 alternative starting points by randomizing around this maximum (drawing from a normal with variance derived from the Hessian at the maximum), maximize again, and check that none of the random initial points leads to a point with higher likelihood (i.e. a new maximum). If this is not the case, we take the new maximum and repeat the randomization until no points with higher likelihood are found. Then, for all the remaining estimation windows, we use the optimum obtained in the previous period $t - 1$ as initial condition for the maximization performed in period t .

4 Forecasting the Yield Curve

4.1 Data

We use the US Treasury zero coupon yield curve estimates by Gurkaynak et al. (2007). This dataset is publicly available on the website of the Federal Reserve Board, at the address <http://www.federalreserve.gov/pubs/feds/2006> and is updated regularly. Gurkaynak et al.

(2007) provide several details on the derivation of the yield curve from raw data, here we just note that the zero coupon yield curve estimates are obtained by using the extension by Svensson (1994) of the functional form that was initially proposed by Nelson and Siegel (1987). Such specification assumes that the forward rates are governed by six parameters.

We focus on monthly data going from January 1985 to December 2008. Although the dataset by Gurkaynak et al. (2007) goes further backward in time, we focus on the post-Volker period in order to avoid sample instability. We consider yields of the following 18 maturities (in months): 1, 3, 6, 9, 12, 15, 18, 21, 24, 30, 36, 48, 60, 72, 84, 96, 108, 120.

These choices provide us with a panel of 288 monthly observations on 18 different yields. Some descriptive statistics for the data at hand are contained in Table I, while a plot of the data is provided in Figure 1.

4.2 Forecasting Exercise

The forecasting exercise is performed in pseudo real time, i.e. we never use information which is not available at the time the forecast is made. For computing our results we use a rolling estimation window of 120 months (10 years). We produce forecasts for all the horizons up to 12-step ahead, but for brevity we present results only for the 1-, 3-, 6-, and 12- step ahead. Results for the remaining horizons are available upon request. The initial estimation window is 1985:1-1994:12, and the initial forecast window is 1995:1-1995:12. The last estimation window is 1998:1 to 2007:12, and the last forecast window is 2008:1 to 2008:12. The choice of a rolling scheme is suggested by two reasons. First, it is a natural way to avoid problems of instability, see e.g. Pesaran and Timmerman (2005). Second, having fixed the number of observations used to compute the forecasts and therefore the resulting time series of the forecast errors allows to use the Giacomini and White (2006) test for comparing forecast accuracy. Such a test is indeed valid provided that the size of the estimation window is fixed.

4.3 Statistical loss functions

Let $\hat{y}_{t+h}^{(\tau)}(M)$ denote the forecast of $y_{t+h}^{(\tau)}$ made by model M . The Mean Squared Forecast Error (MSFE) made by model M in forecasting the yield of maturity τ at horizon h is:

$$MSFE_{\tau,h}^M = \frac{1}{T_0} \sum \left(\hat{y}_{t+h}^{(\tau)}(M) - y_{t+h}^{(\tau)} \right)^2 \quad (43)$$

where the sum is computed over all the T_0 forecasts produced. We assess predictive accuracy in terms of Relative Mean Squared Forecast Error with respect to the Random Walk:

$$RMSFE_{\tau,h}^M = \frac{MSFE_{\tau,h}^M}{MSFE_{\tau,h}^{RW}}. \quad (44)$$

In order to give an idea of the absolute level of the errors, and to facilitate comparisons with other studies we shall also report the Root Mean Squared Forecast Error of the Random Walk forecast, i.e. $\sqrt{MSFE_{\tau,h}^{RW}}$. As it is customary in this literature we present results in terms of annualized percentage yields. For example, a root $MSFE$ of 0.31 signals an error of 31 basis points in predicting the annualized yield.

To assess whether the forecasts of two competing models are statistically significantly different we use the Giacomini and White (2006) test for forecast accuracy. This is a test of equal forecasting *method* accuracy and as such can handle forecasts based on both nested and non-nested models, regardless from the estimation procedures used in the derivation of the forecasts, including Bayesian and semi- and non-parametric estimation methods. The test is valid so long the number of observations used to estimate the forecasting models does not tend to infinity, which is the case in our example, as we use a rolling estimation window. We use the unconditional version of the test, which is based on the same statistic of Diebold and Mariano (1995).

4.4 Trading rules

Several papers have evaluated the forecasting performance of term structure models by looking at statistical measures such as the one discussed in the previous Section. Only a few papers (to our knowledge) have tried to use alternative loss functions to evaluate their performance. Cochrane and Piazzesi (2005) use a simple trading rule in order to evaluate the cumulative profits arising from using their proposed forward rate regression in out of sample forecasting, and Carriero and Giacomini (2009) have used a utility-based loss function in order to evaluate the degree of usefulness of the no arbitrage restrictions. In this paper we provide an assessment of the forecasting performance based on economic criteria. In particular we consider two alternative trading strategies and a portfolio utility loss function.

Note that in the forecasting exercise we use the models described in Section 2 to produce forecasts of the yields. However, the forecasts of the yield $\hat{y}_{t+1}^{(\tau)}$ provides a forecast of the holding period return $\hat{r}_{t+1}^{(\tau+1)}$ via a simple transformation:

$$\hat{r}_{t+1}^{(\tau+1)} = -\tau\hat{y}_{t+1}^{(\tau)} + (\tau + 1)y_t^{(\tau+1)}. \quad (45)$$

Also note that the forecast error made in forecasting the holding period return is proportional to that made in forecasting the yield of a given bond: $\hat{r}_{t+1}^{(\tau+1)} - r_{t+1}^{(\tau+1)} = -\tau(\hat{y}_{t+1}^{(\tau)} - \tau y_{t+1}^{(\tau)})$. This also implies that using the holding period return rather than the yields in the *statistical* loss function would just rescale the results.

The first trading rule is similar in spirit to that used by Carriero et al. (2009) for evaluating the performance of a portfolio of exchange rates. This rule simply suggests to buy or to short sell a bond of a given maturity depending on the sign of the forecasted variation in holding period returns. Profits at time $t + 1$ from this trading rule applied to the yield of maturity τ are given by:

$$\pi_{t,\tau}^{TR1} = \text{sign} \left(\left[\hat{r}_{t+1}^{(\tau)} - \hat{y}_{t+1}^{(1)} \right] \right) \times \left[r_{t+1}^{(\tau)} - y_{t+1}^{(1)} \right] \times 100 \quad (46)$$

As is clear, this rule uses the forecast of the excess one-period return with respect to the one-month yield $\left[\hat{r}_{t+1}^{(\tau)} - \hat{y}_{t+1}^{(1)} \right]$ to recommend the sign of a position which is subject to the ex post return $\left[r_{t+1}^{(\tau)} - y_{t+1}^{(1)} \right] \times 100$, where 100 is a notional capital. The profits of an equally weighted portfolio composed of returns of all the maturities at hand is simply given by $\sum_{\tau} w_{\tau} \pi_{t,\tau}^{TR1}$, where w_{τ} are equal weights. Cumulative profits from time t_0 to time t can be easily computed as $\Pi_t^{TR1} = \sum_{j=t_0}^t \sum_{\tau} w_{\tau} \pi_{j,\tau}^{TR1}$.

The second trading rule is the same as in the Cochrane and Piazzesi (2005, online Appendix). Profits at time $t + 1$ from this trading rule applied to the yield of maturity τ are given by:

$$\pi_{t,\tau}^{TR2} = \left(\left[\hat{r}_{t+1}^{(\tau)} - \hat{y}_{t+1}^{(1)} \right] \times 100 \right) \times \left[r_{t+1}^{(\tau)} - y_{t+1}^{(1)} \right] \times 100. \quad (47)$$

As is clear, this rule uses the forecast of the excess one-period return with respect to the one-month yield $\left[\hat{r}_{t+1}^{(\tau)} - \hat{y}_{t+1}^{(\tau_1)} \right]$ to recommend the size of a position which is subject to the ex post return $\left[r_{t+1}^{(\tau)} - y_{t+1}^{(\tau_1)} \right]$, where 100 is a notional capital. The recommended size of the investment is given by $\left[\hat{r}_{t+1}^{(\tau)} - \hat{y}_{t+1}^{(1)} \right] \times 100$, where 100 is the notional capital. The profits of an equally weighted portfolio composed of returns of all the maturities at hand is simply given by $\sum_{\tau} w_{\tau} \pi_{t,\tau}^{TR2}$, where w_{τ} are equal weights. Cumulative profits from time t_0 to time t can be easily computed as $\Pi_t^{TR2} = \sum_{j=t_0}^t \sum_{\tau} w_{\tau} \pi_{j,\tau}^{TR2}$.

The Sharpe ratios for a given maturity, for a generic trading rule TR used from t_0 to t can be computed as:

$$\text{Sharpe}_{\tau}^M = \frac{\bar{\pi}}{\sqrt{\frac{1}{t-t_0} \sum_{j=t_0}^t (\pi_{j,\tau}^{TR} - \bar{\pi})^2}}; \quad \bar{\pi} = \frac{1}{t-t_0} \sum_{j=t_0}^t \pi_{j,\tau}^{TR} \quad (48)$$

4.5 Optimal portfolio allocation

The trading strategies considered in the previous section provide a simple and straightforward criterion for the economic relevance of the forecasts produced by the various alternative models. Still, they focus on each yield separately and do not take explicitly into account the risk-return considerations. To address these latter issues, it is interesting to see how useful

the different forecasting models are, when used as a basis for optimal portfolio allocation. Following Carriero and Giacomini (2009) we consider a portfolio utility loss function which directly takes into account the risk return trade-off. Consider the asset allocation problem of an investor who is buying a portfolio of N assets in period t and then sells it in period $t+1$. In our application the relevant portfolio is a portfolio of bonds of maturities $\tau_1, \tau_2, \dots, \tau_N$ bought in period t and then sold in period $t+1$. The vector of expected returns on this investment is given by $E_t[\mathbf{r}_{t+1}] = E_t[(r_{t+1}^{(\tau_1+1)}, r_{t+1}^{(\tau_2+1)}, \dots, r_{t+1}^{(\tau_q+1)})']$. We further assume that the model does not specify conditional variance dynamics, so that the conditional variance of \mathbf{r}_{t+1} at time t simply equals the unconditional variance-covariance matrix of the N assets Σ , so that $Var_t[\mathbf{r}_{t+1}] = \Sigma$. We suppose that at each time t the forecaster constructs a portfolio by choosing the weights that minimize a quadratic utility loss function:

$$w^* = \arg \min_w \left\{ -w' E_t[\mathbf{r}_{t+1}] + \frac{\gamma}{2} w' \Sigma w \right\}, \quad (49)$$

where $E_t[\cdot]$ denotes the conditional mean at time t . The classical solution (Markowitz, 1952) to this problem is given by

$$w^* = a + B E_t[\mathbf{r}_{t+1}]; \quad (50)$$

$$a = \frac{\Sigma^{-1} \boldsymbol{\iota}}{\boldsymbol{\iota}' \Sigma^{-1} \boldsymbol{\iota}}; \quad (51)$$

$$B = \frac{1}{\gamma} \left(\Sigma^{-1} - \frac{\Sigma^{-1} \boldsymbol{\iota} \boldsymbol{\iota}' \Sigma^{-1}}{\boldsymbol{\iota}' \Sigma^{-1} \boldsymbol{\iota}} \right), \quad (52)$$

where $\boldsymbol{\iota}$ is a $N \times 1$ vector of ones and γ is a user-defined parameter related to the coefficient of relative risk aversion δ by the relationship $\frac{\gamma}{1-\gamma} = \delta$. Our empirical results are obtained by setting $\delta = 1$, so that $\gamma = .5$.

Each model at hand will provide a different forecast $\hat{\mathbf{r}}_{t+1} = E_t[\mathbf{r}_{t+1}]$, therefore providing a different set of optimal weights and a different ex-post utility loss at time t given by:

$$Loss_t = -w^*(\hat{\mathbf{r}}_{t+1})' \mathbf{r}_{t+1} + \frac{\gamma}{2} w^*(\hat{\mathbf{r}}_{t+1})' \Sigma w^*(\hat{\mathbf{r}}_{t+1}). \quad (53)$$

The average ex-post utility loss from t_0 to t is given by $\frac{1}{t-t_0} \sum_{j=t_0}^t Loss_j$. Note that the Giacomini and White (2006) test can be applied to any loss function which is itself a function of the forecasts, therefore we can use it to determine whether utility losses arising from two different models are statistically significantly different or not.

5 Results

5.1 Statistical Loss Function

Results in terms of Relative Mean Squared Forecast Errors ($RMSFE$, our statistical loss function) are displayed in Table II. The table is divided in four panels, each of them corresponding to a different forecast horizon, respectively 1, 3, 6, and 12 step ahead. Each column provides results for a different maturity. The first row in each panel contains the root $MSFE$ of the Random Walk (RW) forecasts. All the remaining rows report the $RMSFE$ of a given model vis a vis the RW , according to equation (44), so a figure smaller than one signals that the model under consideration outperforms the RW . The stars on the right of the cell entries signal the level at which the Giacomini and White (2006) test rejects the null of equal forecasting method accuracy (*, **, and *** stars mean respectively rejection at 10%, 5%, and 1% level). Note that for the 1-step ahead case the results for the AR and VAR with powering-up and direct approach produce by construction the same results and therefore they are grouped together. Several conclusions can be drawn from the table.

First, the RW confirms to be a very competitive benchmark in forecasting the term structure of bond yields. It systematically outperforms the AR models, which in general are a competitive benchmark in forecasting macroeconomic variables. The RW forecasts are generally more accurate than those of most of the competing models, especially at long horizons.

The $BVAR$ is the only method that is able to systematically outperform the RW at all maturities for short forecast horizons. The gains are in the range of 0% to 2% for the 1- and 12-step ahead horizon, and in the range of 2% to 4% for the 3-step and the 6-step ahead horizon. Such gains are not statistically significant, confirming the difficulty of beating the RW benchmark, but they are systematic for all maturities and forecast horizons.

Some models beat the $BVAR$ for a few selected maturities and forecast horizons but they perform much worse than the $BVAR$ in the remaining cases. Among these are the $DL(VAR)$, the $ATSM(tp)$ and the $VARs$, which all provide high gains for bonds of short maturities, especially at short forecast horizons, but their performance rapidly deteriorates when forecasting bonds of longer maturities.

Among linear models, it is worth noticing that the powering up produces overall better forecasts than the direct approach, both for ARs and for $VARs$. This finding suggests that the extent of mis-specification of linear models is limited. As expected the $VARs$ are often outperformed by simpler autoregressive specifications, which is likely due to overparametriza-

tion. Still it is very interesting to note that in some cases *VARs* do produce better forecasts than the simple *ARs*. This suggests that there is some valuable information in the cross section of yields, and the problem is to extract it efficiently, as this information probably gets lost due to the overparametrization of the model.

Turning to the forward rate regressions, their performance is somewhat mixed. In particular, both the *FB* and *CP* regression forecast quite well yields of short maturities but they have a hard time in beating the *RW* for yields of longer maturities. As for the internal ranking of the two, the *FB* regression seems to systematically outperform the *CP* regression in out of sample forecasting. We shall see how this pattern is reversed when one considers a different loss function.

Considering now the models based on the Nelson and Siegel interpolant, unexpectedly we find mixed results about the *DL* model depending on whether an *AR* or a *VAR* specification is assumed for the factors. In particular the model with a *VAR* specification, $DL(VAR)$, produces overall good forecasts, especially at short horizons and comparatively better than those based on the *AR* model, $DL(AR)$, for short and medium maturities and relatively short forecast horizon. The pattern is inverted for longer maturities and longer forecast horizons.

As for the *ATSM* models, the model with constant risk premium, $ATSM(crp)$, produces quite poor forecasts. The inclusion of a time varying premium improves dramatically forecast accuracy, although the $ATSM(tp)$ model can beat the *RW* only in few instances and overall is outperformed by alternative models. Also for *ATSMs* we shall see how results will change when a different loss function is considered.

5.2 Trading rules and portfolio utility loss.

We now consider the results in terms of our economic loss function, i.e. the two alternative trading strategies, and the portfolio utility loss. Table III presents the Sharpe Ratios (i.e. returns divided by standard deviation) obtained by using the trading strategies specified in (46) and (47) applied when the forecasts are obtained with our alternative methods. Note that as the strategies use only 1-step ahead forecasts, The *AR* and *VAR* models with powering-up or direct approach produce by construction the same result and are grouped together.

As is clear, the *BVAR* provides the highest returns and Sharpe ratios in most of the cases, followed by the $DL(AR)$ and $ATSM(tp)$ models. Interestingly, the $ATSM(tp)$ produce a rather high Sharpe ratio for both trading strategies, but importantly the presence of a time varying risk premium is essential to this result, as the $ATSM(crp)$ produces negative returns.

In order to assess possible time variation in the performance of the models, we proceed

along the lines of Cochrane and Piazzesi (2005, online Appendix) and we plot in Figure 4 the *cumulative* ex post profits on an equally weighted portfolio containing all the yields to maturity at hand. The plot shows the same ranking as Table III, namely the *BVAR*, *DL*, and *ATSM(tvp)* produce the best forecasts, and the inclusion of a time variation in the premia is essential for *ATSMs*. Also *CP* and *FB* regressions perform quite well. Finally it is interesting to note that for this trading strategy too the simple *VAR* produces rather good forecasts.

Finally, we consider the results for the Portfolio Utility Loss for portfolios whose weights are optimally chosen. We consider a portfolio of bonds of maturity 1, 3, 12, 24, 36, 60, and 120 months. Table IV provides the ex-post average utility loss obtained with each of the methods at hand. The model producing the lowest average (ex-post) utility loss is the *RW*, followed by the *BVAR*. The forward rate regressions *FB* and *CP* perform quite well but not as well as the *BVAR*. Overall, the methods based on multivariate models perform quite poorly, because they occasionally produce highly volatile forecasts. Finally, it is confirmed the importance of allowing for time variation in risk premia, as the loss implied by the *ATSM* doubles from 7.95 to 14.21 when constant risk premia are imposed on the model.

6 Robustness

In this section we check the temporal robustness of the results by performing the evaluation over two subsamples, and we conduct a simulation based on bootstrapped data that allows a better assessment of forecast uncertainty.

6.1 Sub-sample stability

In order to check for stability of our results in different samples, we run the forecasting exercise on two subsamples, each composed of seven years of forecasts. The first subsample goes from 1995:1 to 2000:12. Results on this subsample are displayed in Table V (for RMSFE), Table VI (for the trading rules) and Table VII (for the Portfolio utility loss).

Results in Table V are overall in line with those obtained on the whole sample³. The only

³In order to compare the results with those in Diebold and Li (2006) we also considered the period going from 1994:1 to 2000:12. To produce results for the year 1994 we used a shorter rolling window of 108 months rather than 120 as in our baseline application. Results for this subsample are available upon request and are in line with those obtained by Diebold and Li (2006) with the same sample and the Fama-Bliss dataset. In particular, the *DL* model with factors following an *AR* produces very good forecasts, and it outperforms the *RW* at horizons longer than 1-step ahead.

exception is the forecasting performance of the *AR* model, which works much better in this subsample. The *BVAR* is still systematically better than the *RW*.

Table VI contains results on the trading rules. Also in this subsample, the *BVAR* systematically provides the highest returns and Sharpe ratios. Also the *DL(AR)* provides good results, while the *DL(VAR)* produces negative returns. In this subsample the *CP* model works remarkably well, especially for the second trading strategy. This result is in line with Cochrane and Piazzesi (2005). Results for the *cumulative* ex post profits on an equally weighted portfolio containing all the yields and using the two trading strategies are displayed in Figure 5.

Table VII contains results on the portfolio utility loss and confirms the patterns found on the whole sample (see Table 4). The only noticeable differences are that in this subsample the forward rate regressions perform better than in the full sample, while the *ATSM* with constant risk premium performs relatively worse.

The second subsample contains forecasts for the period ranging from 2001:1 to 2008:12. Results on this subsample are displayed in Table VIII (for RMSFE), Table IX (for the trading rules) and Table X (for the Portfolio utility loss).

By looking at Table VIII it is clear that in this subsample all the models display a considerable reduction in forecast accuracy relative to the *RW*. The *AR*, which was competitive in the first subsample, becomes worse than the *RW* (the deterioration explains the overall performance of the *AR*). A similar result applies to the *DL* model, whose forecasting performance also deteriorates with respect to the first subsample. Also the forecasting performance of the *BVARs* deteriorates, but *BVARs* still produce some of the best forecasts, and in several instances are still able to outperform the *RW*. The remaining models, i.e. the *SW* model and the forward rate regressions have a consistently poor performance.

Table IX contains results on the trading rules. In this subsample the *BVAR* is no longer systematically the best model, even though it produces relatively high Sharpe ratios at least for medium and long maturities. The evidence in this subsample is instead more mixed, with the *RW*, the *BVAR*, the *FB* providing the highest Sharpe ratios. Results for the *cumulative* ex post profits on an equally weighted portfolio containing all the yields and using the two trading strategies are displayed in Figure 6. In this case the *ATSM* with time varying risk premium performs well when the second strategy is used.

Table X contains results on the portfolio utility loss. Overall, the reduced forecasting performance (based on the *MSFE*) in this subsample is confirmed: the loss from all the models increases, with the *BVAR* still producing the smallest possible utility loss closely

followed by the *RW* and the *CP*.

In summary, it is clear that in the second part of the sample there is an overall increase in unpredictability, i.e. it is even more difficult to beat a simple *RW* forecast. Most of the models producing good forecasts in the first part of the sample see their performance in the second part deteriorating.

6.2 Simulation with artificial data

To evaluate in more detail the performance and the robustness of the forecasting methods under analysis, we carry out a stochastic simulation in which we simulate and forecast a set of "artificial" term structures. Using several sets of artificial data allows to retrieve distributions of forecasts and forecast errors, which in turn can be used to assess the amount of variance implied by each competing model around the point forecasts. As we shall see, the simulation provides evidence that while most of the models feature an increase in forecast uncertainty with respect to the *RW*, our *BVARs* lead to a reduction in forecast uncertainty.

In doing a Monte Carlo simulation the researcher has to choose a Data Generating Process (DGP). Obviously the choice of a particular DGP over another would influence the results, advantaging one rather than the other model. Therefore, rather than concentrate on simulated data and an inevitably arbitrary data generation process, we carry out our simulation by bootstrapping the actual term structure dataset. The use of a real dataset as a basis for such a robustness analysis is referred to as a 'data based Monte Carlo method' and discussed further in, e.g., Ho and Sørensen (1996). Following this work, we create a set of artificial term structures by repeatedly bootstrapping the actual term structure.

In particular, we use the block bootstrapping algorithm described in Politis and Romano (1994). This algorithm is designed for block-bootstrapping from a set of stationary data, therefore to implement it we first difference the data on the yields and then we recover the bootstrapped level data by adding an initial condition and integrating the artificial data. We set the initial condition equal to the first actual data point, i.e. the term structure in 1985:1. To avoid generating inappropriate patterns for the artificial term structures, we eliminate all the draws involving negative values for any of the yields at any point in time.

We repeat this procedure 100 times, and then for each of the 100 repetitions we implement our forecasting exercise. We do not include in the experiment the *AFDNS* and the *ATSM* models, as these models are highly (especially the latter) nonlinear and their estimation via the *EM* algorithm makes the dimension of the problem unmanageable. This provides us with 100 time series of forecasts for each yield and forecast horizon, which can be used to

compute 100 *MSFEs* (one for each repetition) for each model. Finally, the *MSFEs* are averaged through the repetitions to get the average *MSFE* made by model M in predicting h -step ahead the yield of maturity τ . The square roots of these, denoted $avgMSFE_{\tau,h}^M$, are displayed in Table XI. Similarly, we can define the average of the Sharpe Ratio from a given trading strategy based on model M . These average Sharpe ratios are displayed in Table XII. Finally, we can similarly produce the average portfolio utility loss by averaging the loss in (53) over the 100 replications. Results for this case are in Table XIII.

Looking at the first row of Table XI, i.e. the Average Root *MSFE* of the *RW*, it is clear that the artificial data are relatively harder to predict with respect to the actual data (the Average Root *MSFE* of the *RW* is almost double with respect to the actual data). Interestingly all the models but the *BVARs* feature much higher average *MSFE* than the *RW*. As the bias for all the models is small⁴, this implies that the variance of the forecasts of all the models at hand tends to be higher than that of the *RW*. The only exception is the *BVAR*, whose forecast variability is reduced by the shrinkage towards the random walk forecasts. Second, comparing these results with those obtained with the actual data, only the performance of the *RW* and the *BVAR* appears to be robust to changes in the sample.

Similarly, Table XII shows that the *BVAR* is producing the best Sharpe ratios in the vast majority of cases. Finally, Table XIII shows that in the artificial samples the *RW*, *AR*, *BVAR*, *FB* and *CP* all produce similar results, while the *VAR* and all the factor based models' forecasts are rather poor.

7 Conclusions

In this paper we introduce a new statistical model for the entire term structure of interest rates and compare its forecasting performance with the current most promising alternatives, evaluating the forecasts in terms of both statistical and economic criteria.

Our proposal is a large Bayesian Vector Autoregression (BVAR), with time-variation in the amount of shrinking towards a set of univariate AR models. The underlying idea is that, while in general it is difficult to outperform univariate AR or Random Walk models, there are time periods with enhanced interaction among yields at different maturities, and capturing such an interaction only when it is indeed at work further enhances the forecasting performance.

We suggest an estimation procedure for the BVAR that, combined with the formulation of direct rather than iterated forecasts, totally eliminates the need of time-demanding simulation

⁴For brevity we do not report results on the biases, which are available upon request.

methods. Therefore, we can estimate the BVAR and forecast the entire term structure for different horizons in just a few seconds.

We then compare the forecasting performance of our BVAR with that of a large set of alternative models for the term structure of government bond yields. Specifically, we consider linear models (Random Walk, Autoregressive Models, Vector Autoregressions), factor models (Diebold-Li (2006), Christensen et al. (2007), Stock and Watson (2002a,b)), forward rate regression (Fama and Bliss (1987), Cochrane and Piazzesi (2005)), and Affine term structure Models (Ang and Piazzesi (2003)).

We extended the available empirical evidence on relative forecasting performance in three directions: i) For all the models considered we provide results for an homogeneous dataset, rather than using different maturities and sample periods. We use the novel dataset of Gurkaynak et al. (2007), which is publicly available and updated regularly. ii) While most of the existing evidence evaluates forecast accuracy only in terms of statistical measures such as Root Mean Squared Forecast Errors, we also evaluate forecasts on the basis of "economic" criteria. In particular we provide Sharpe Ratios arising from simple trading rules based on the alternative forecasts, and we use the alternative forecasts to perform optimal portfolio allocation. iii) We provide an extensive robustness check of the results. In particular, we provide results for two subsamples and we simulate and forecast a set of "artificial" term structures. Using several sets of artificial data allows to retrieve distributions of forecasts and forecast errors, which in turn can be used to assess the amount of variance implied by each competing model around the point forecasts.

Our main empirical results can be summarized as follows: i) our proposed BVAR approach produces forecasts systematically more accurate than the random walk forecasts, though the gains are small; ii) some models beat the BVAR for a few selected maturities and forecast horizons, but they perform much worse than the BVAR in the remaining cases; iii) predictive gains with respect to the random walk have decreased over time; iv) different loss functions (i.e., "statistical" vs "economic") lead to different ranking of specific models; v) modelling time variation in term premia is important and useful for forecasting.

In conclusion, we believe that this paper makes a simple but important theoretical contribution, the introduction of a large BVAR model with time-varying shrinkage, and a relevant empirical contribution, since this model seems to outperform common benchmarks and most competing specifications for the term structure of interest rates.

References

- [1] Almeida, C., Vicente, J., 2008. The role of no-arbitrage on forecasting: Lessons from a parametric term structure model. *Journal of Banking & Finance* 32(12) 2695-2705.
- [2] Ang, A. and M. Piazzesi, 2003. A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables, *Journal of Monetary Economics* 50, 4, 745-787.
- [3] Bai and Ng, 2006. Confidence Intervals for Diffusion Index Forecasts and Inference for Factor-Augmented Regressions. *Econometrica* 74(4), 1133-1150.
- [4] Banbura, Giannone and Reichlin, 2010. Large Bayesian Vector Autoregressions, *Journal of Applied Econometrics* 25(1), 71-92.
- [5] Bauwens L., Lubrano M., and Richard J.F., 1999. *Bayesian Inference in Dynamic Econometric Models*, Oxford University Press.
- [6] Carriero, A., Giacomini, R., 2009. How useful are no-arbitrage restrictions to forecast the term structure of interest rates? Mimeo.
- [7] Carriero, A., Kapetanios, G., and Marcellino M., 2009. Forecasting Exchange Rates with a Large Bayesian VAR. *International Journal of Forecasting*, 25(2), 400-417.
- [8] Carriero, A., Kapetanios, G., and Marcellino M., 2010. Forecasting Large Datasets with Bayesian Reduced Rank Multivariate Models. CEPR Discussion Papers, forthcoming in the *Journal of Applied Econometrics*.
- [9] Christensen, J.H.E., F.X. Diebold, and G. Rudebusch, 2007. The Affine Arbitrage-Free Class of Nelson-Siegel Term Structure Models. Manuscript, University of Pennsylvania and Federal Reserve Bank of San Francisco.
- [10] Cochrane, J. H., and M. Piazzesi, 2005. Bond Risk Premia. *American Economic Review* 95, 1, pp. 138-160.
- [11] Cox, J.C., Ingersoll, J.E., Ross, S.A., 1985. A theory of the term structure of interest rates. *Econometrica* 53, 385-407.
- [12] Dai, Q., Singleton, K., 2000. Specification analysis of affine term structure models. *Journal of Finance* 55, 1943-1978.

- [13] De Jong, F., 2000. Time series and cross section information in affine term structure models. *Journal of Business and Economic Statistics* 18, 300–314.
- [14] De Mol, Giannone, and Reichlin, 2008. Forecasting Using a Large Number of Predictors. Is Bayesian Regression a Valid Alternative to Principal Components? *Journal of Econometrics* 146(2), 318-328.
- [15] De Pooter, M., Ravazzolo F., van Dijk, D., 2007. Predicting the term structure of interest rates incorporating parameter uncertainty, model uncertainty and macroeconomic information. Tinbergen Institute Discussion Papers 07-028/4
- [16] Diebold, F. X., and Li, C., 2006. Forecasting the term structure of government bond yields. *Journal of Econometrics* 130, 337–364
- [17] Diebold, F. X and Mariano, R. 1995. Comparing Predictive Accuracy. *Journal of Business & Economic Statistics* 13(3), 253-63.
- [18] Diebold, F.X., Rudebusch, G.D. and Aruoba, B. 2006. The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach. *Journal of Econometrics* 131, 309-338.
- [19] Doan T., Litterman R., Sims C., 1984. Forecasting and Conditional Projection Using Realistic Prior Distributions. *Econometric Reviews* 3, 1-100.
- [20] Duffee, G., 2002. Term premia and interest rate forecasts in affine models. *Journal of Finance* 57, 405–443.
- [21] Duffie, D., Kan, R., 1996. A yield-factor model of interest rates. *Mathematical Finance* 6, 379–406.
- [22] Duffee, G., 2007. Forecasting with the term structure: the role of no-arbitrage. Mimeo.
- [23] Fama, E., Bliss, R., 1987. The information in long-maturity forward rates. *American Economic Review* 77, 680–692.
- [24] Favero, C.A., Niu L., Sala, L., 2007. Term Structure Forecasting: No-Arbitrage Restrictions vs Large Information Set. CEPR Discussion Papers 6206
- [25] Giacomini R., White H., 2006. Tests of Conditional Forecast Accuracy. *Econometrica* 74, 1545-1578.
- [26] Gurkaynak, Brian Sack, and Jonathan H. Wright, 2007. The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics* 54(8), 2291-2304.

- [27] Harrison, J.M., Kreps, D.M., 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 2, 381–408.
- [28] Ho, M. S., and B. E. Sørensen, 1996, Finding Cointegration Rank in High Dimensional Systems using the Johansen Test: An Illustration Using Data Based Monte Carlo Simulations, *Review of Economics and Statistics* 78, 726–732.
- [29] Kadiyala, K. R., and S. Karlsson 1997. Numerical Methods for Estimation and Inference in Bayesian VAR-Models. *Journal of Applied Econometrics* 12(2), 99–132.
- [30] Litterman R.B., 1986. Forecasting with Bayesian Vector Autoregressions-Five Years of Experience. *Journal of Business & Economic Statistics* 4, 25-38.
- [31] Marcellino, M., Stock, J. H., Watson, M.W., 2006. A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of Econometrics* 127(1-2), 499-526.
- [32] Markowitz, H. 1952, "Portfolio Selection". *The Journal of Finance*, 7(1), 77-91.
- [33] Moench, E., 2008. Forecasting the yield curve in a data-rich environment: A no-arbitrage factor-augmented VAR approach. *Journal of Econometrics* 146(1) 26-43.
- [34] Nelson, C.R., Siegel, A.F., 1987. Parsimonious modeling of yield curve. *Journal of Business* 60, 473–489.
- [35] Pesaran, M.H. and Timmermann A., 2005. Small sample properties of forecasts from autoregressive models under structural breaks, *Journal of Econometrics* 129(1-2), 183-217
- [36] Pesaran M.H, Pick A., Timmerman A. 2009. Variable Selection and Inference for Multi-period Forecasting Problems. Manuscript.
- [37] Politis, D.N., Romano, J.P., 1994. The Stationary Bootstrap. *Journal of the American Statistical Association* 89, No. 428, 1303-1313.
- [38] Sims, C. A., and T. Zha 1998. Bayesian Methods for Dynamic Multivariate Models. *International Economic Review* 39(4), 949–68.
- [39] Stock, J. H. and M. W. Watson 2002a. Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business and Economic Statistics* 20, 147-62.

- [40] Stock, J. H. and M. W. Watson 2002b. Forecasting Using Principal Components from a Large Number of Predictors. *Journal of the American Statistical Association* 97, 1167–1179.
- [41] Svensson, L.E.O., 1994. Estimating and Interpreting Forward Interest Rates: Sweden 1992-4. CEPR Discussion Papers 1051.
- [42] Vasicek, O., 1977. An equilibrium characterization of the term structure. *Journal of Financial Economics* 5, 177–188.
- [43] Zellner, A. (1973). *An Introduction to Bayesian Inference in Econometrics*. Wiley.

Table I. Descriptive Statistics

Maturity	Mean	StDev	1st Quartile	Median	3d Quartile	min	max
1	4.71	1.99	5.05	3.06	5.86	1.01	8.91
3	4.80	2.04	5.15	3.17	6.01	1.00	9.22
6	4.91	2.09	5.15	3.35	6.29	0.93	9.50
9	5.00	2.11	5.22	3.45	6.37	0.84	9.62
12	5.09	2.12	5.35	3.57	6.41	0.77	9.72
15	5.17	2.11	5.40	3.73	6.50	0.74	9.92
18	5.25	2.11	5.46	3.87	6.58	0.73	10.11
21	5.32	2.10	5.50	3.95	6.67	0.74	10.27
24	5.38	2.09	5.50	4.01	6.76	0.78	10.43
30	5.50	2.06	5.57	4.17	6.98	0.89	10.69
36	5.60	2.04	5.65	4.31	7.13	1.06	10.91
48	5.79	1.99	5.74	4.40	7.26	1.47	11.22
60	5.95	1.95	5.84	4.48	7.41	1.93	11.43
72	6.09	1.91	5.95	4.55	7.52	2.37	11.57
84	6.22	1.87	6.06	4.67	7.61	2.78	11.66
96	6.33	1.84	6.14	4.72	7.69	3.12	11.71
108	6.43	1.81	6.23	4.80	7.76	3.41	11.76
120	6.52	1.78	6.31	4.91	7.81	3.64	11.79

Descriptive statistics. Data are US Treasury zero coupon yield curve estimates by Gurkaynak et al (2006), publicly available at <http://www.federalreserve.gov/pubs/feds/2006>. Data are at monthly frequency, going from January 1985 to December 2008. The units are percentages on an annual basis. Data include the yields of the following maturities (in months): 1, 3, 6, 9, 12, 15, 18, 21, 24, 30, 36, 48, 60, 72, 84, 96, 108, 120.

Table III. Sharpe Ratios for alternative models and trading strategies

Panel A: Trading Strategy 1									
Maturity	1-month	3-month	6-month	9-month	12-month	15-month	18-month	21-month	24-month
RW	-	0.024	-0.007	0.058	0.108	0.131	0.128	0.125	0.120
AR (pu), AR(di)	-	0.016	0.027	-0.017	0.014	0.049	0.051	0.069	0.066
VAR (pu), VAR (di)	-	-0.015	0.079	0.122	0.115	0.064	0.070	0.094	0.060
BVAR	-	0.027	0.058	0.144	0.192	0.202	0.213	0.195	0.211
DL (AR)	-	0.122	0.013	-0.045	-0.031	0.016	0.051	0.102	0.100
DL (VAR)	-	0.149	0.045	0.031	0.001	0.065	0.044	0.101	0.116
SW	-	0.019	0.020	-0.003	-0.024	0.084	0.150	0.047	0.025
FB	-	0.118	0.171	0.196	0.174	0.130	0.125	0.123	0.072
CP	-	0.070	0.033	0.053	0.009	0.003	-0.015	-0.036	-0.048
ATSM (cp)	-	0.021	0.058	0.005	-0.073	-0.100	-0.138	-0.030	-0.011
ATSM (tvp)	-	0.189	0.105	0.000	-0.039	-0.012	0.004	0.046	0.041
Maturity	30-month	36-month	48-month	60-month	72-month	84-month	96-month	108-month	120-month
RW	0.090	0.094	0.108	0.107	0.122	0.134	0.145	0.133	0.144
AR (pu), AR(di)	0.068	0.063	0.048	0.128	0.080	0.045	0.023	0.018	0.030
VAR (pu), VAR (di)	0.044	0.003	-0.042	-0.079	-0.106	-0.109	-0.118	-0.138	-0.119
BVAR	0.199	0.190	0.181	0.175	0.171	0.168	0.165	0.162	0.159
DL (AR)	0.077	-0.013	-0.026	-0.064	-0.090	-0.011	0.017	0.156	0.182
DL (VAR)	0.068	0.024	-0.073	-0.043	-0.077	-0.053	-0.054	-0.028	0.038
SW	-0.034	-0.078	-0.057	0.047	0.084	0.051	-0.104	-0.064	-0.096
FB	0.011	0.045	0.042	0.025	0.064	0.100	0.103	0.103	0.037
CP	-0.007	-0.032	-0.033	-0.023	-0.005	-0.022	-0.041	-0.004	0.010
ATSM (cp)	0.006	0.034	0.038	0.028	0.007	-0.063	-0.029	-0.025	0.003
ATSM (tvp)	0.063	0.078	0.051	0.034	0.133	0.122	0.120	0.134	0.151
Panel B: Trading Strategy 2									
Maturity	1-month	3-month	6-month	9-month	12-month	15-month	18-month	21-month	24-month
RW	-	0.090	0.094	0.122	0.149	0.166	0.175	0.179	0.179
AR (pu), AR(di)	-	0.084	0.088	0.108	0.124	0.132	0.133	0.129	0.123
VAR (pu), VAR (di)	-	0.043	-0.001	0.000	0.007	0.015	0.022	0.028	0.033
BVAR	-	0.129	0.145	0.179	0.207	0.223	0.231	0.233	0.232
DL (AR)	-	0.076	0.062	0.052	0.049	0.053	0.063	0.072	0.073
DL (VAR)	-	0.105	0.093	0.090	0.096	0.096	0.069	0.042	0.022
SW	-	0.089	0.078	0.016	-0.009	0.015	0.086	0.071	0.048
FB	-	0.212	0.171	0.171	0.164	0.151	0.135	0.119	0.104
CP	-	0.125	0.147	0.116	0.087	0.064	0.047	0.035	0.028
ATSM (cp)	-	-0.025	0.012	-0.020	-0.056	-0.071	-0.062	-0.044	-0.026
ATSM (tvp)	-	0.162	0.092	0.072	0.072	0.086	0.102	0.117	0.128
Maturity	30-month	36-month	48-month	60-month	72-month	84-month	96-month	108-month	120-month
RW	0.175	0.169	0.157	0.148	0.142	0.136	0.132	0.129	0.126
AR (pu), AR(di)	0.109	0.097	0.079	0.068	0.062	0.058	0.055	0.049	0.041
VAR (pu), VAR (di)	0.041	0.048	0.055	0.057	0.055	0.051	0.045	0.037	0.029
BVAR	0.226	0.220	0.209	0.202	0.196	0.191	0.187	0.184	0.181
DL (AR)	0.055	0.031	0.000	-0.009	-0.005	0.014	0.063	0.139	0.133
DL (VAR)	-0.010	-0.039	-0.083	-0.099	-0.105	-0.105	-0.089	-0.036	0.007
SW	-0.066	-0.109	-0.103	-0.052	0.007	0.023	-0.127	-0.146	-0.128
FB	0.079	0.060	0.036	0.020	0.010	0.001	-0.005	-0.010	-0.014
CP	0.022	0.022	0.027	0.030	0.028	0.021	0.011	-0.001	-0.013
ATSM (cp)	-0.001	0.012	0.025	0.026	0.013	-0.051	-0.160	-0.111	-0.093
ATSM (tvp)	0.141	0.149	0.159	0.159	0.148	0.132	0.118	0.107	0.099

The Table contains the Sharpe Ratios obtained by using each of the models at hand and the trading strategies described in Section 4.4. The evaluation sample is 1995:1 to 2008:12. The competing models considered are Random Walk (RW), Linear models estimated with powering-up and direct approach (AR(pu), AR(di), VAR(pu), and VAR(di)), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (cp)) variation in the risk premium.

Table IV. Portfolio Utility Loss

	Average Loss	Significance
RW	0.5714	
AR	0.8805	***
VAR	1.5134	*
BVAR	0.5782	
DL (AR)	22.2288	***
DL (VAR)	22.1051	***
SW	535.542	***
FB	0.7963	
CP	0.6402	
ATSM (crp)	14.2197	***
ATSM (tvp)	7.9544	***

The Table contains the Utility Loss obtained by using each of the models at hand and the trading strategies described in Section 4.5. The evaluation sample is 1995:1 to 2008:12. The competing models considered are Random Walk (RW), Linear models (AR, VAR), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (crp)) variation in the risk premium. The stars on the right of the cell entries signal the level at which the Giacomini and White (2006) test rejects the null of equal forecasting accuracy against the random walk (*, **, and *** mean respectively rejection at 10%, 5%, and 1% level).

Table VI. Sharpe Ratios for alternative models and trading strategies, Sample 1995:1 to 2000:12

Panel A: Trading Strategy 1

Maturity	1-month	3-month	6-month	9-month	12-month	15-month	18-month	21-month	24-month
RW	-	0.089	0.062	0.130	0.182	0.221	0.210	0.193	0.183
AR (pu), AR(di)	-	0.132	0.114	0.020	0.038	0.086	0.079	0.092	0.081
VAR (pu), VAR (di)	-	0.034	0.139	0.147	0.135	0.125	0.133	0.178	0.138
BVAR	-	0.156	0.186	0.261	0.250	0.256	0.236	0.220	0.283
DL (AR)	-	0.101	0.022	0.003	-0.021	-0.012	-0.017	0.020	0.007
DL (VAR)	-	0.135	0.034	0.008	-0.023	-0.010	-0.078	0.065	0.042
SW	-	-0.041	-0.055	-0.140	-0.087	0.118	0.218	0.016	0.017
FB	-	0.265	0.290	0.303	0.242	0.214	0.190	0.170	0.147
CP	-	0.156	0.054	0.060	-0.027	-0.033	-0.008	0.000	-0.023
ATSM (cp)	-	-0.090	-0.012	-0.015	-0.040	-0.082	-0.137	-0.018	-0.012
ATSM (tvp)	-	0.263	0.137	0.093	0.019	0.075	0.071	0.082	0.039

Maturity	30-month	36-month	48-month	60-month	72-month	84-month	96-month	108-month	120-month
RW	0.120	0.124	0.069	0.045	0.093	0.093	0.094	0.096	0.093
AR (pu), AR(di)	0.054	0.018	-0.031	0.050	-0.032	-0.057	-0.049	-0.037	-0.041
VAR (pu), VAR (di)	0.071	0.011	0.023	-0.017	-0.065	-0.108	-0.122	-0.123	-0.083
BVAR	0.253	0.232	0.207	0.194	0.185	0.180	0.175	0.171	0.167
DL (AR)	-0.018	-0.076	-0.006	-0.066	-0.076	0.090	0.041	0.051	0.210
DL (VAR)	0.057	-0.048	-0.120	-0.077	-0.060	0.009	-0.070	-0.043	0.015
SW	-0.044	-0.043	-0.003	0.070	0.119	0.046	-0.107	-0.100	-0.153
FB	0.097	0.106	0.084	0.108	0.116	0.180	0.175	0.171	0.167
CP	0.017	-0.010	-0.025	-0.008	0.030	0.001	0.015	0.023	0.056
ATSM (cp)	0.061	0.119	0.141	0.140	0.079	-0.014	-0.033	0.001	0.059
ATSM (tvp)	0.075	0.080	0.118	0.073	0.159	0.123	0.115	0.105	0.102

Panel B: Trading Strategy 2

Maturity	1-month	3-month	6-month	9-month	12-month	15-month	18-month	21-month	24-month
RW	-	0.127	0.115	0.130	0.153	0.169	0.176	0.175	0.171
AR (pu), AR(di)	-	0.126	0.110	0.114	0.119	0.120	0.115	0.108	0.100
VAR (pu), VAR (di)	-	0.090	0.019	0.013	0.020	0.030	0.042	0.055	0.066
BVAR	-	0.185	0.193	0.226	0.256	0.272	0.276	0.272	0.265
DL (AR)	-	0.090	0.076	0.067	0.060	0.052	0.038	0.012	-0.022
DL (VAR)	-	0.129	0.122	0.122	0.126	0.112	0.013	-0.064	-0.100
SW	-	0.110	0.036	-0.062	-0.093	-0.045	0.186	0.096	0.060
FB	-	0.325	0.251	0.250	0.245	0.232	0.217	0.201	0.185
CP	-	0.192	0.246	0.201	0.155	0.124	0.104	0.093	0.088
ATSM (cp)	-	-0.088	-0.005	-0.006	-0.016	-0.012	0.006	0.027	0.048
ATSM (tvp)	-	0.210	0.134	0.117	0.123	0.141	0.163	0.184	0.200

Maturity	30-month	36-month	48-month	60-month	72-month	84-month	96-month	108-month	120-month
RW	0.156	0.141	0.117	0.100	0.087	0.078	0.072	0.068	0.067
AR (pu), AR(di)	0.087	0.077	0.060	0.041	0.019	-0.007	-0.037	-0.069	-0.098
VAR (pu), VAR (di)	0.085	0.098	0.112	0.117	0.117	0.114	0.109	0.101	0.093
BVAR	0.248	0.234	0.214	0.202	0.193	0.187	0.183	0.179	0.176
DL (AR)	-0.079	-0.103	-0.093	-0.056	-0.013	0.039	0.097	0.120	0.112
DL (VAR)	-0.136	-0.152	-0.149	-0.124	-0.097	-0.069	-0.032	0.017	0.046
SW	-0.112	-0.146	-0.110	-0.042	0.036	0.059	-0.103	-0.144	-0.145
FB	0.155	0.129	0.089	0.066	0.056	0.054	0.055	0.058	0.060
CP	0.089	0.097	0.113	0.119	0.118	0.110	0.098	0.084	0.068
ATSM (cp)	0.080	0.103	0.133	0.148	0.147	0.089	-0.141	-0.148	-0.135
ATSM (tvp)	0.219	0.228	0.226	0.199	0.152	0.097	0.046	0.003	-0.028

The Table contains the Sharpe Ratios obtained by using each of the models at hand and the trading strategies described in Section 4.4. The evaluation sample is 1995:1 to 2000:12. The competing models considered are Random Walk (RW), Linear models estimated with powering-up and direct approach (AR(pu), AR(di), VAR(pu), and VAR(di)), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (cp)) variation in the risk premium.

Table VII. Portfolio Utility Loss, Sample 1995:1-2000:12

	Average Loss	Significance
RW	0.4447	
AR	0.48	
VAR	1.8625	
BVAR	0.4322	
DL (AR)	15.1092	***
DL (VAR)	14.2607	***
SW	760.2786	***
FB	0.1047	
CP	0.2643	
ATSM (crp)	15.6363	***
ATSM (tvp)	4.505	***

The Table contains the Utility Loss obtained by using each of the models at hand and the trading strategies described in Section 4.5. The evaluation sample is 1995:1 to 2000:12. The competing models considered are Random Walk (RW), Linear models (AR, VAR), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (crp)) variation in the risk premium. The stars on the right of the cell entries signal the level at which the Giacomini and White (2006) test rejects the null of equal forecasting accuracy against the random walk (*, **, and *** mean respectively rejection at 10%, 5%, and 1% level).

Table IX. Sharpe Ratios for alternative models and trading strategies, Sample 2001:1 to 2008:12

Panel A: Trading Strategy 1

Maturity	1-month	3-month	6-month	9-month	12-month	15-month	18-month	21-month	24-month
RW	-	-0.093	-0.121	-0.043	0.012	0.025	0.035	0.051	0.053
AR (pu), AR(di)	-	-0.201	-0.114	-0.070	-0.017	0.004	0.017	0.042	0.048
VAR (pu), VAR (di)	-	-0.107	-0.017	0.088	0.088	-0.009	-0.004	0.002	-0.025
BVAR	-	-0.212	-0.148	-0.017	0.119	0.136	0.185	0.166	0.135
DL (AR)	-	0.169	-0.001	-0.115	-0.043	0.049	0.130	0.195	0.205
DL (VAR)	-	0.182	0.064	0.064	0.034	0.159	0.188	0.142	0.198
SW	-	0.133	0.146	0.195	0.057	0.043	0.072	0.081	0.034
FB	-	-0.135	-0.008	0.054	0.087	0.030	0.050	0.071	-0.010
CP	-	-0.083	-0.002	0.043	0.057	0.047	-0.023	-0.076	-0.076
ATSM (cp)	-	0.233	0.177	0.035	-0.117	-0.122	-0.138	-0.043	-0.010
ATSM (tvp)	-	0.069	0.055	-0.135	-0.115	-0.120	-0.075	0.005	0.043
Maturity	30-month	36-month	48-month	60-month	72-month	84-month	96-month	108-month	120-month
RW	0.058	0.063	0.149	0.173	0.152	0.177	0.198	0.172	0.196
AR (pu), AR(di)	0.082	0.112	0.131	0.209	0.197	0.150	0.097	0.073	0.103
VAR (pu), VAR (di)	0.015	-0.005	-0.110	-0.143	-0.149	-0.110	-0.114	-0.154	-0.156
BVAR	0.143	0.148	0.153	0.156	0.157	0.156	0.155	0.153	0.151
DL (AR)	0.180	0.052	-0.047	-0.062	-0.104	-0.114	-0.007	0.269	0.154
DL (VAR)	0.080	0.099	-0.025	-0.008	-0.094	-0.117	-0.037	-0.012	0.062
SW	-0.024	-0.115	-0.112	0.023	0.049	0.055	-0.101	-0.028	-0.039
FB	-0.081	-0.018	-0.002	-0.060	0.011	0.022	0.031	0.036	-0.095
CP	-0.032	-0.055	-0.041	-0.039	-0.041	-0.045	-0.098	-0.031	-0.036
ATSM (cp)	-0.052	-0.054	-0.066	-0.085	-0.067	-0.114	-0.026	-0.052	-0.055
ATSM (tvp)	0.049	0.076	-0.018	-0.006	0.107	0.121	0.125	0.163	0.202

Panel B: Trading Strategy 2

Maturity	1-month	3-month	6-month	9-month	12-month	15-month	18-month	21-month	24-month
RW	-	-0.020	0.063	0.111	0.143	0.167	0.183	0.193	0.200
AR (pu), AR(di)	-	-0.067	0.052	0.101	0.130	0.145	0.151	0.150	0.145
VAR (pu), VAR (di)	-	-0.202	-0.156	-0.086	-0.049	-0.035	-0.032	-0.033	-0.034
BVAR	-	-0.064	0.050	0.104	0.140	0.165	0.182	0.194	0.201
DL (AR)	-	0.126	0.066	0.029	0.035	0.076	0.131	0.176	0.198
DL (VAR)	-	0.127	0.041	0.021	0.042	0.081	0.114	0.135	0.148
SW	-	0.135	0.181	0.200	0.182	0.118	-0.005	0.024	0.028
FB	-	-0.097	0.003	0.025	0.029	0.030	0.031	0.031	0.031
CP	-	-0.128	-0.054	-0.008	0.004	0.000	-0.008	-0.017	-0.025
ATSM (cp)	-	0.116	0.036	-0.036	-0.100	-0.136	-0.127	-0.102	-0.080
ATSM (tvp)	-	0.082	-0.031	-0.066	-0.053	-0.023	0.005	0.027	0.043
Maturity	30-month	36-month	48-month	60-month	72-month	84-month	96-month	108-month	120-month
RW	0.204	0.204	0.199	0.193	0.188	0.183	0.179	0.175	0.172
AR (pu), AR(di)	0.131	0.117	0.097	0.091	0.097	0.115	0.141	0.169	0.183
VAR (pu), VAR (di)	-0.035	-0.034	-0.032	-0.033	-0.038	-0.046	-0.055	-0.065	-0.074
BVAR	0.208	0.209	0.208	0.204	0.201	0.197	0.194	0.190	0.186
DL (AR)	0.184	0.136	0.056	0.015	-0.001	0.002	0.037	0.202	0.170
DL (VAR)	0.153	0.126	0.008	-0.071	-0.112	-0.141	-0.145	-0.080	-0.021
SW	-0.016	-0.063	-0.095	-0.061	-0.016	-0.016	-0.152	-0.150	-0.120
FB	0.028	0.024	0.013	0.001	-0.013	-0.027	-0.040	-0.051	-0.060
CP	-0.034	-0.038	-0.041	-0.044	-0.049	-0.057	-0.067	-0.078	-0.089
ATSM (cp)	-0.052	-0.038	-0.029	-0.033	-0.051	-0.122	-0.181	-0.089	-0.065
ATSM (tvp)	0.064	0.079	0.108	0.136	0.155	0.166	0.170	0.172	0.173

The Table contains the Sharpe Ratios obtained by using each of the models at hand and the trading strategies described in Section 4.4. The evaluation sample is 2001:1 to 2008:12. The competing models considered are Random Walk (RW), Linear models estimated with powering-up and direct approach (AR(pu), AR(di), VAR(pu), and VAR(di)), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (cp)) variation in the risk premium.

Table X. Portfolio Utility Loss, Sample 2001:1-2008:12

	Average Loss	Significance
RW	0.9889	
AR	1.6235	***
VAR	1.2632	
BVAR	1.1034	
DL (AR)	38.2162	***
DL (VAR)	37.8461	***
SW	137.431	***
FB	2.3247	***
CP	1.1977	
ATSM (crp)	18.4412	***
ATSM (tvp)	13.8551	***

The Table contains the Utility Loss obtained by using each of the models at hand and the trading strategies described in Section 4.5. The evaluation sample is 2001:1 to 2008:12. The competing models considered are Random Walk (RW), Linear models (AR, VAR), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (cp)) variation in the risk premium. The stars on the right of the cell entries signal the level at which the Giacomini and White (2006) test rejects the null of equal forecasting accuracy against the random walk (*, **, and *** mean respectively rejection at 10%, 5%, and 1% level).

Table XII. Sharpe Ratios, averages over artificial samples

Panel A: Trading Strategy 1

Maturity	1-month	3-month	6-month	9-month	12-month	15-month	18-month	21-month	24-month
RW	-	1.33	1.09	0.88	0.73	0.62	0.56	0.50	0.46
AR	-	1.32	1.08	0.87	0.70	0.60	0.52	0.47	0.42
VAR	-	0.80	0.55	0.39	0.30	0.24	0.20	0.18	0.16
BVAR	-	1.29	1.04	0.84	0.69	0.60	0.53	0.48	0.45
DL (AR)	-	0.76	0.52	0.46	0.37	0.32	0.37	0.39	0.35
DL (VAR)	-	0.76	0.51	0.45	0.37	0.31	0.35	0.38	0.34
SW	-	1.06	0.91	0.75	0.60	0.48	0.41	0.37	0.33
FB	-	1.29	1.05	0.84	0.68	0.58	0.51	0.46	0.41
CP	-	1.29	1.00	0.77	0.61	0.51	0.44	0.38	0.34

Maturity	30-month	36-month	48-month	60-month	72-month	84-month	96-month	108-month	120-month
RW	0.40	0.35	0.28	0.24	0.22	0.20	0.18	0.17	0.16
AR	0.35	0.30	0.24	0.19	0.16	0.14	0.12	0.11	0.10
VAR	0.13	0.11	0.08	0.07	0.06	0.06	0.05	0.05	0.04
BVAR	0.39	0.35	0.30	0.27	0.25	0.24	0.22	0.21	0.20
DL (AR)	0.21	0.17	0.12	0.09	0.07	0.06	0.03	0.02	0.02
DL (VAR)	0.21	0.17	0.12	0.08	0.06	0.05	0.02	0.02	0.02
SW	0.28	0.24	0.18	0.14	0.12	0.10	0.09	0.08	0.07
FB	0.34	0.29	0.23	0.18	0.15	0.13	0.11	0.10	0.09
CP	0.28	0.24	0.18	0.15	0.13	0.11	0.10	0.09	0.08

Panel B: Trading Strategy 2

Maturity	1-month	3-month	6-month	9-month	12-month	15-month	18-month	21-month	24-month
RW	-	1.72	1.25	0.94	0.73	0.62	0.54	0.49	0.44
AR	-	1.71	1.23	0.93	0.72	0.59	0.52	0.46	0.41
VAR	-	1.58	1.12	0.82	0.63	0.52	0.44	0.38	0.34
BVAR	-	1.74	1.27	0.97	0.77	0.66	0.58	0.53	0.48
DL (AR)	-	1.10	0.60	0.47	0.36	0.31	0.36	0.40	0.35
DL (VAR)	-	1.10	0.59	0.47	0.36	0.31	0.35	0.39	0.34
SW	-	1.63	1.19	0.88	0.68	0.55	0.46	0.41	0.37
FB	-	1.70	1.23	0.93	0.72	0.60	0.53	0.46	0.42
CP	-	1.69	1.20	0.88	0.68	0.55	0.47	0.41	0.36

Maturity	30-month	36-month	48-month	60-month	72-month	84-month	96-month	108-month	120-month
RW	0.38	0.33	0.27	0.23	0.20	0.18	0.17	0.15	0.14
AR	0.35	0.30	0.23	0.18	0.16	0.14	0.12	0.11	0.10
VAR	0.28	0.24	0.18	0.15	0.12	0.10	0.09	0.09	0.07
BVAR	0.42	0.37	0.31	0.26	0.24	0.22	0.21	0.19	0.18
DL (AR)	0.22	0.18	0.13	0.09	0.07	0.05	0.01	0.00	0.01
DL (VAR)	0.22	0.18	0.13	0.09	0.07	0.04	0.01	0.00	0.01
SW	0.31	0.27	0.20	0.16	0.13	0.12	0.10	0.09	0.08
FB	0.35	0.31	0.24	0.20	0.17	0.14	0.13	0.11	0.10
CP	0.30	0.26	0.20	0.16	0.13	0.12	0.10	0.09	0.08

The Table contains the average Sharpe Ratios over 100 bootstrap replications, obtained by using each of the models at hand and the trading strategies described in Section 4.4, and the data based Monte Carlo method described in Section 6.2. The competing models considered are Random Walk (RW), Linear models (AR, VAR), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (cp)) variation in the risk premium. The stars on the right of the cell entries signal the level at which the Giacomini and White (2006) test rejects the null of equal forecasting accuracy (*, **, and *** mean respectively rejection at 10%, 5%, and 1% level).

Table XIII. Portfolio Utility Loss, artificial samples

	Average Loss
RW	-28.427
AR	-27.639
VAR	391.414
BVAR	-28.712
DL (AR)	482.525
DL (VAR)	481.527
SW	461.344
FB	-26.134
CP	-27.701

The Table contains the average Utility Loss over 100 bootstrap replications, obtained by using each of the models at hand and the trading strategies described in Section 4.5, and the data based Monte Carlo method described in Section 6.2. The competing models considered are Random Walk (RW), Linear models (AR, VAR), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (cp)) variation in the risk premium. The stars on the right of the cell entries signal the level at which the Giacomini and White (2006) test rejects the null of equal forecasting accuracy against the random walk (*, **, and *** mean respectively rejection at 10%, 5%, and 1% level).

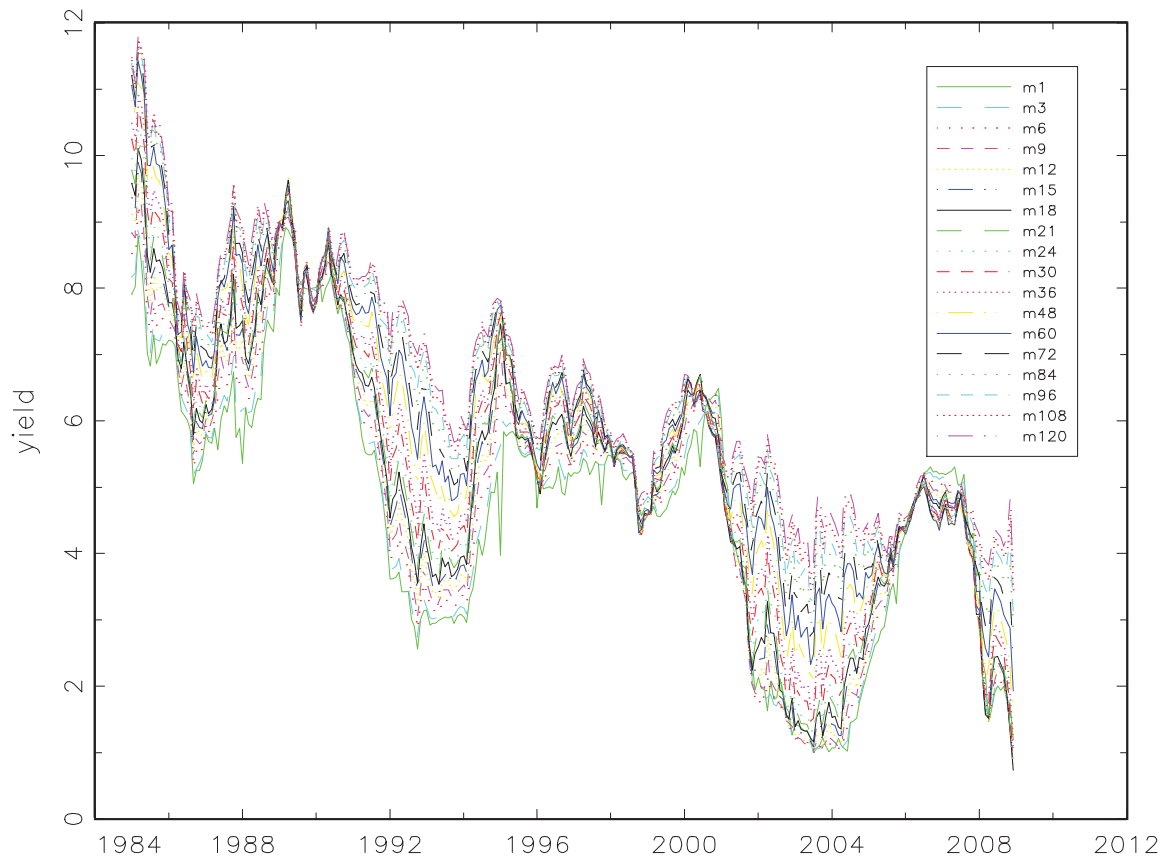


Figure 1: US Treasury zero coupon yield curve estimates by Gurkaynak et al (2006), publicly available at <http://www.federalreserve.gov/pubs/feds/2006>. Data are at monthly frequency, the scale is percentages on annual basis. Data include the yields of the following maturities (in months): 1, 3, 6, 9, 12, 15, 18, 21, 24, 30, 36, 48, 60, 72, 84, 96, 108, 120.

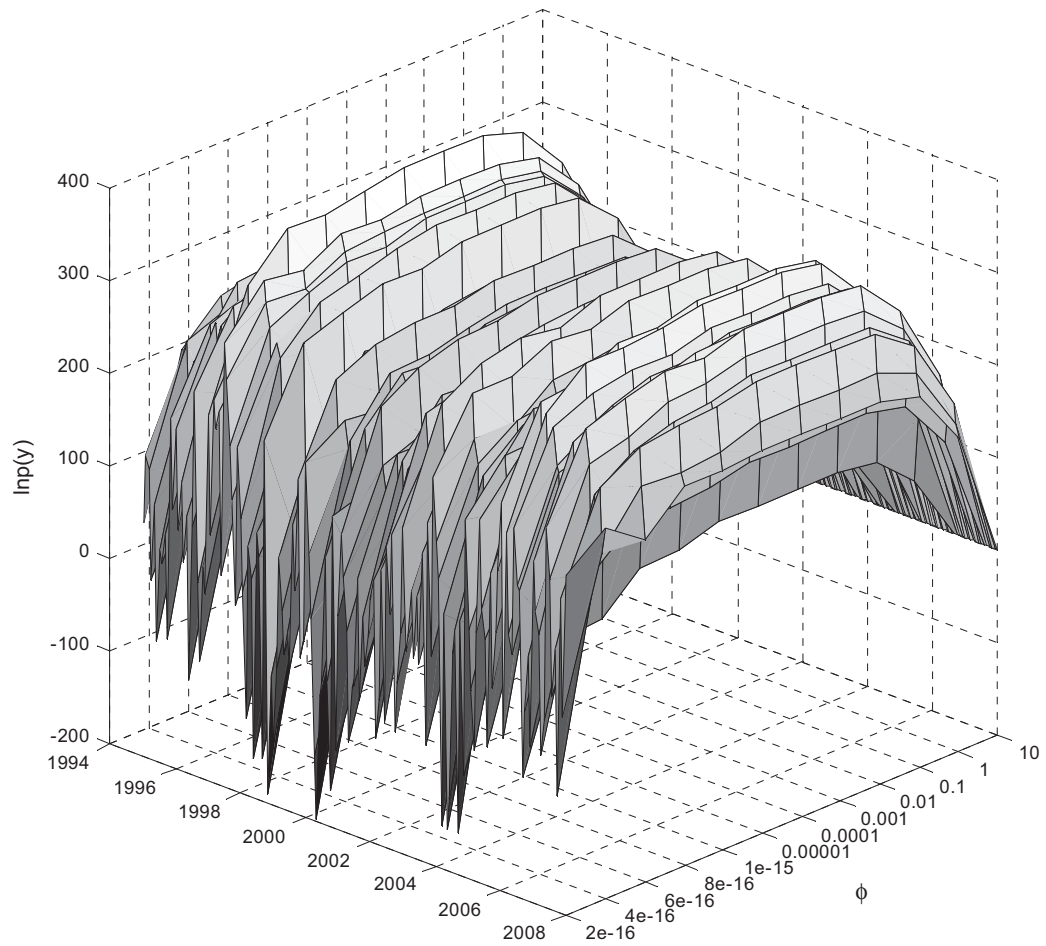


Figure 2: Marginal Data Density. The surface represents the value of the log-data density $\ln p(Y) = \ln \int p(Y|\Theta)p(\Theta)d\Theta$ for the case $h = 1$ (1-step ahead) as a function of the shrinkage parameter θ and time. To facilitate comparison across time the surface is rescaled so that it equals 0 in the point $\phi = 10$.

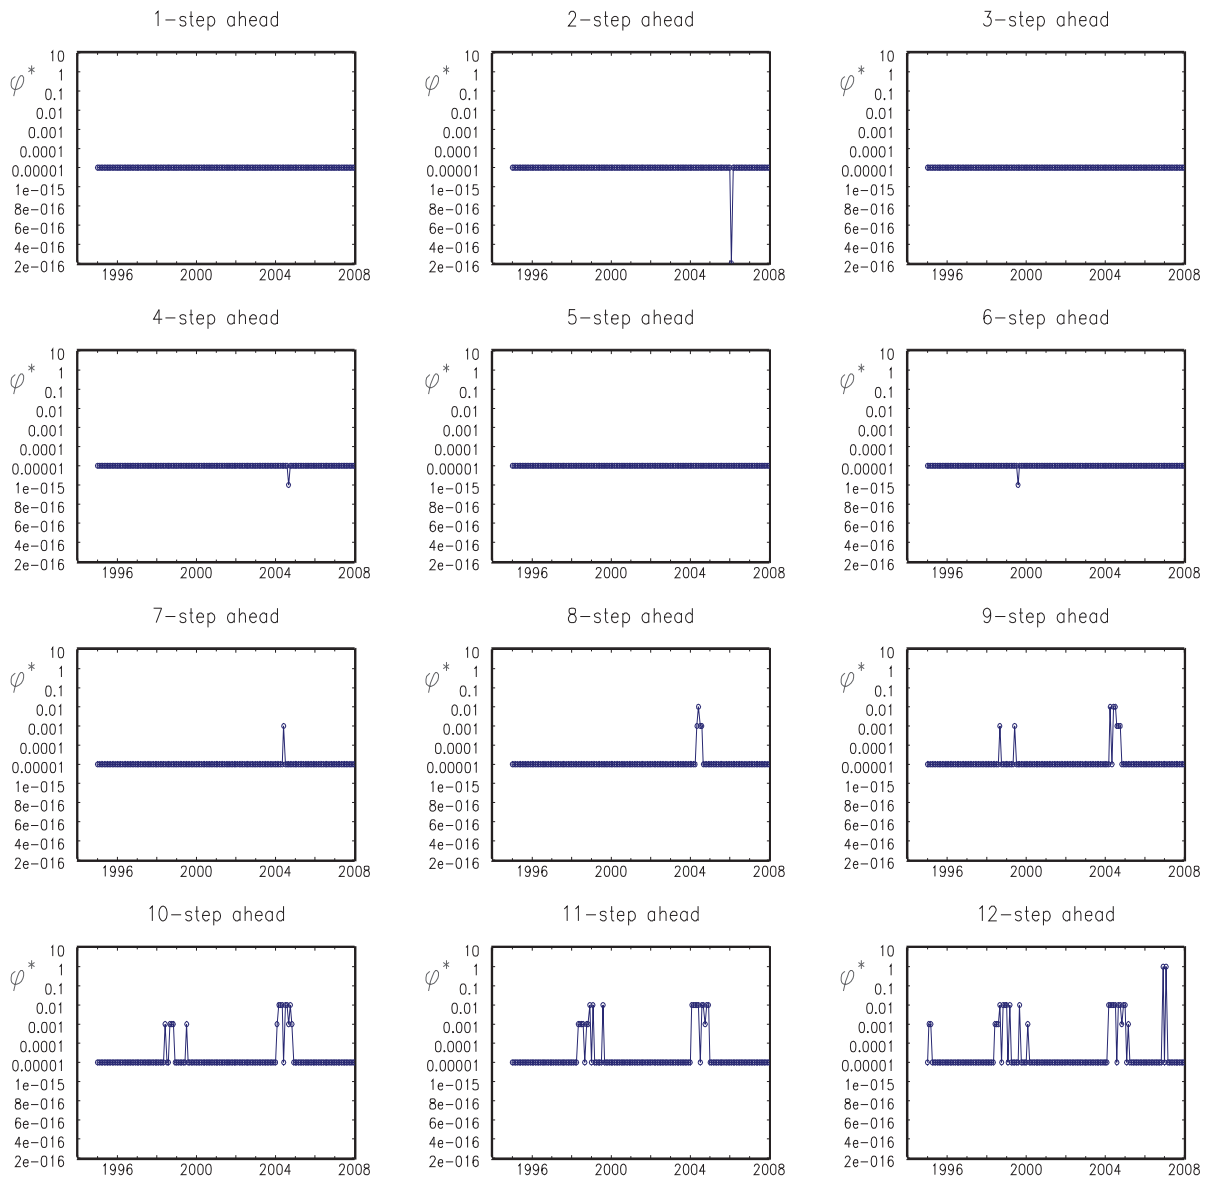


Figure 3: Time series path of the optimal shrinkage $\theta_t^* = \arg \max_{\theta} p(Y)$ for $h = 1, \dots, 12$.

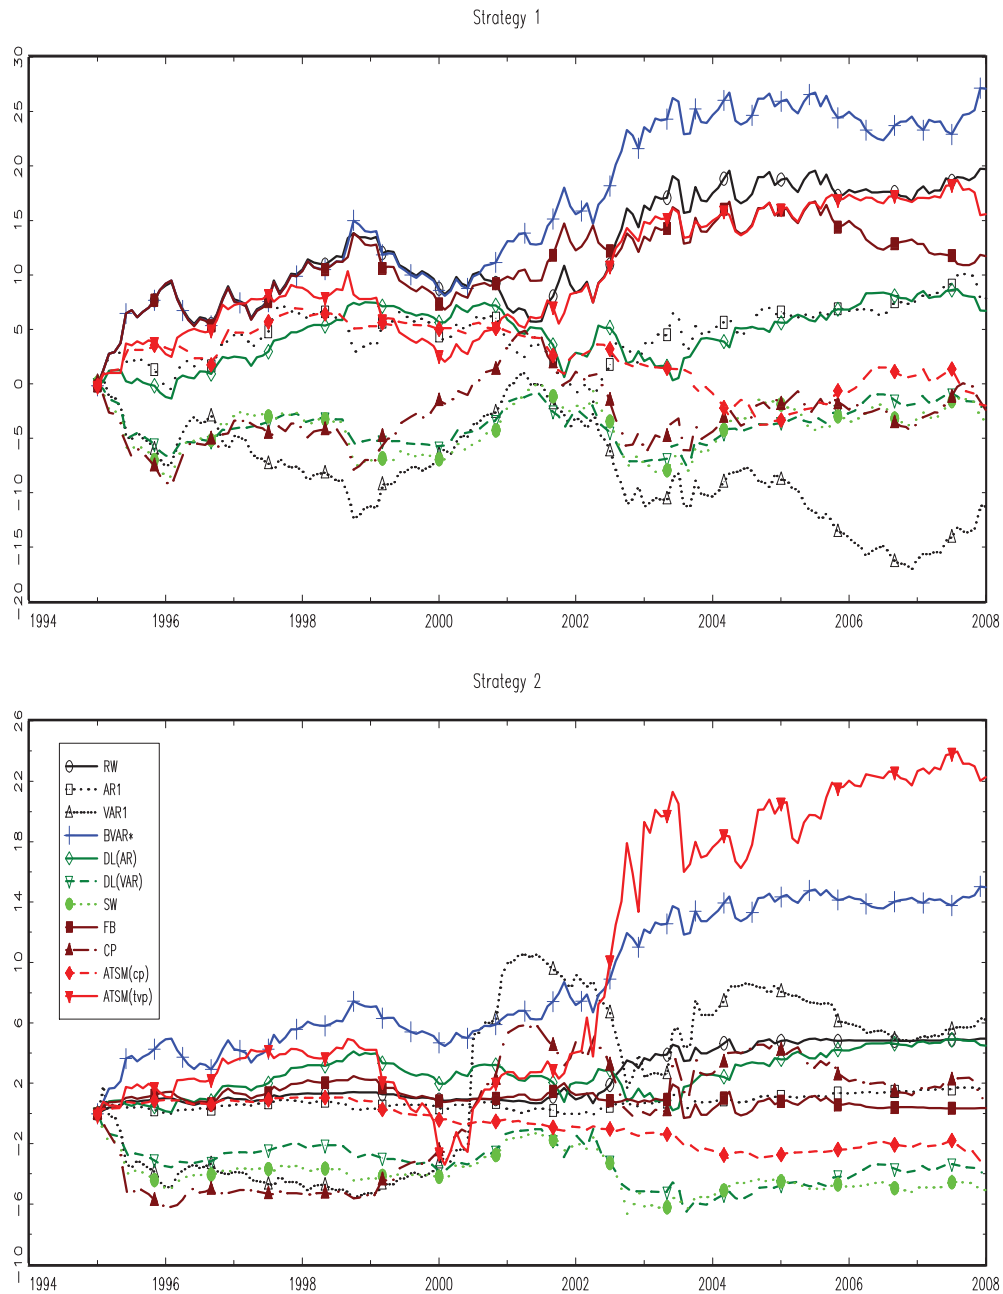


Figure 4: Cumulative ex-post profits on an equally weighted portfolio managed using the two simple trading rules described in Section 4.4. The first panel contains results based on trading strategy 1, i.e., a trading rule that simply suggests to buy or to short sell a bond of a given maturity depending on the sign of the forecasted variation in holding period returns. The second panel contains results based on the trading strategy 2, i.e., a trading rule that uses the forecast of the excess one-period return with respect to the one-month yield to recommend the *size* of a position which is subject to the ex post return. The competing models considered are Random Walk (RW), Linear models (AR and VAR), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tvp)) and without (ATSM (cp)) variation in the risk premium.

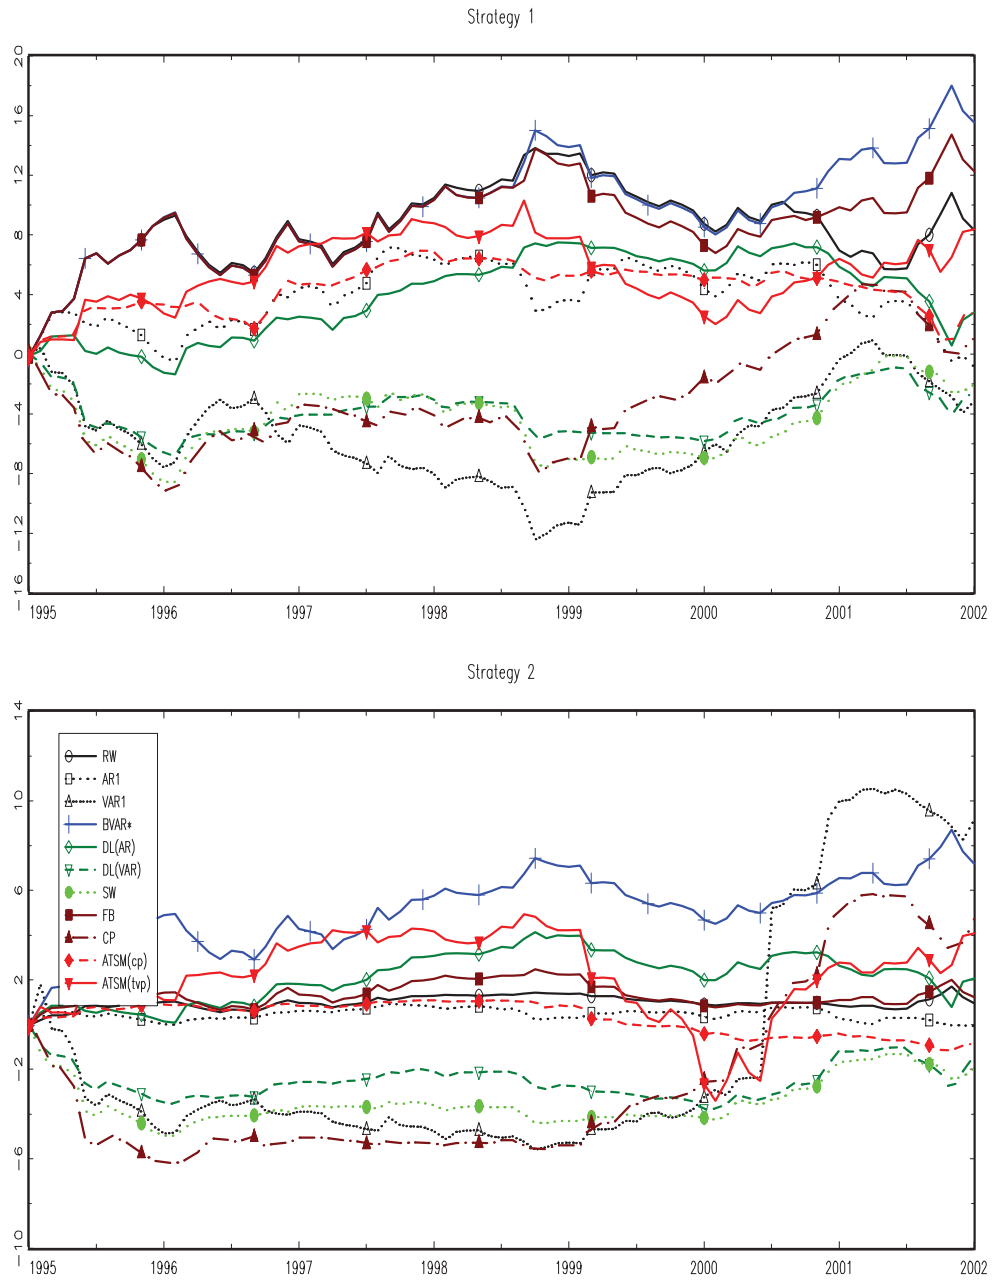


Figure 5: Cumulative ex-post profits on an equally weighted portfolio managed using the two simple trading rules described in Section 4.4. Results over the subsample 1995:1 to 2000:12. See notes to Figure 4 for additional details. The competing models considered are Random Walk (RW), Linear models (AR and VAR), Bayesian VAR (BVAR), Diebold-Li model based on AR and VAR factor structure (DL(AR) and DL(VAR)), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (ATSM (tp)) and without (ATSM (cp)) variation in the risk premium.

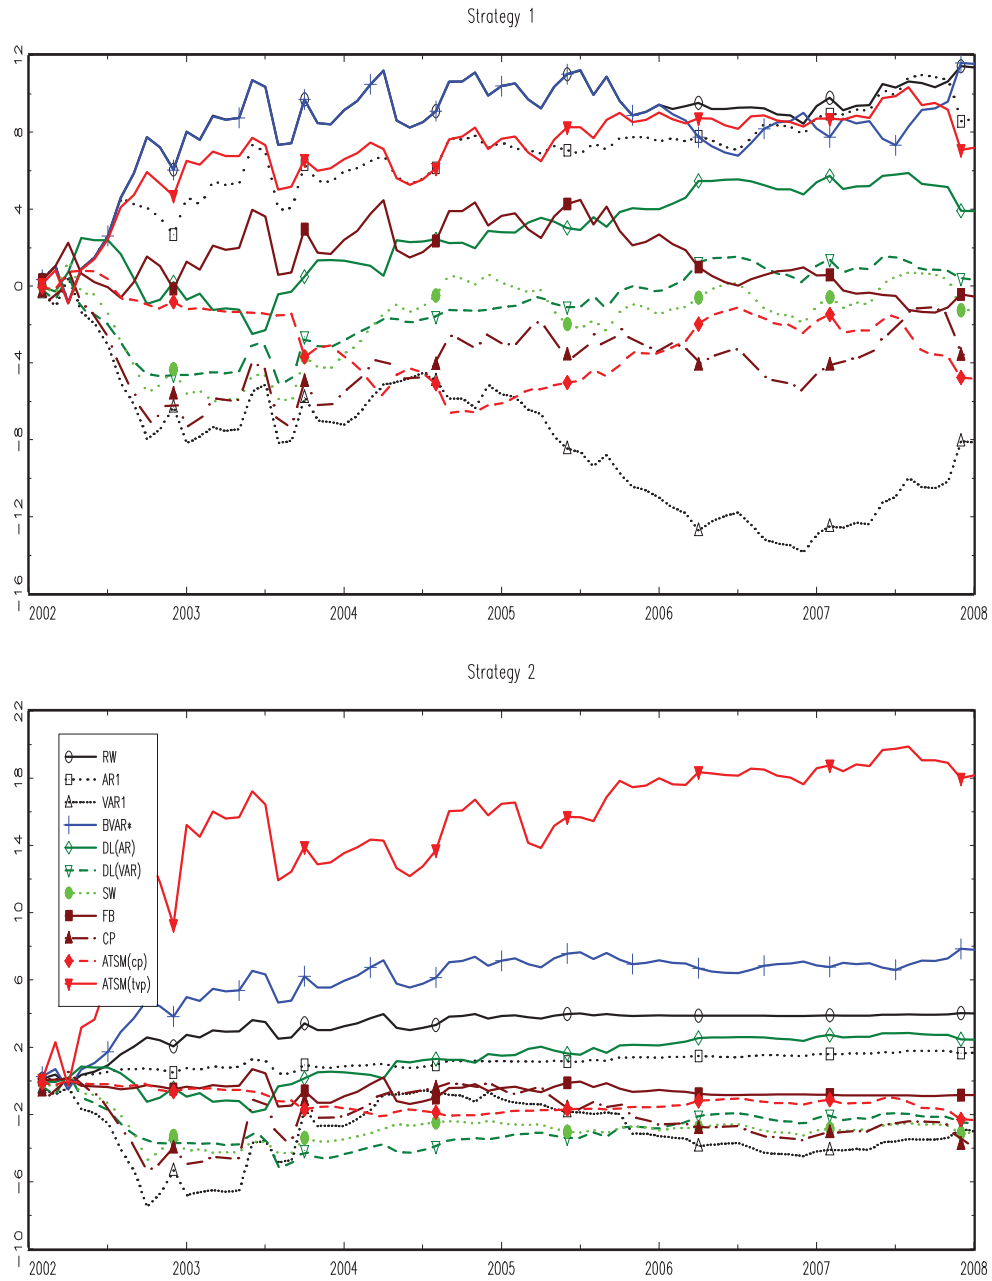


Figure 6: Cumulative ex-post profits on an equally weighted portfolio managed using the two simple trading rules described in Section 4.4. Results over the subsample 2001:1 to 2008:12. See notes to Figure 4 for additional details. The competing models considered are Random Walk (RW), Linear models (AR and VAR), Bayesian VAR (BVAR), Diebold-Li models based on AR and VAR structure of the factors (DL), Stock and Watson factor model (SW), forward rate regressions (FB and CP), and ATSM with (tvp) and without (cp) variation in the risk premium.

**This working paper has been produced by
the School of Economics and Finance at
Queen Mary, University of London**

**Copyright © 2010 Andrea Carriero, George Kapetanios
and Massimiliano Marcellino. All rights reserved**

**School of Economics and Finance
Queen Mary, University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 5096
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/papers/wp.htm**