Bond Return Predictability: Economic Value and Links to the Macroeconomy

Davide Pettenuzzo, Economics Department, Brandeis University

KAntonio Gargano, University of Melbourne

Allan Timmermann, University of California, San Diego

Working Paper Series

2016 | 75R

Bond Return Predictability: Economic Value and Links to the Macroeconomy*

Antonio Gargano†
University of Melbourne

Davide Pettenuzzo‡
Brandeis University

Allan Timmermann§
University of California San Diego

July 28, 2016

Abstract

Studies of bond return predictability find a puzzling disparity between strong statistical evidence of return predictability and the failure to convert return forecasts into economic gains. We show that resolving this puzzle requires accounting for important features of bond return models such as time varying parameters, volatility dynamics, and unspanned macro factors. A three-factor model comprising the Fama and Bliss (1987) forward spread, the Cochrane and Piazzesi (2005) combination of forward rates and the Ludvigson and Ng (2009) macro factor generates notable gains in out-of-sample forecast accuracy compared with a model based on the expectations hypothesis. Such gains in predictive accuracy translate into higher risk-adjusted portfolio returns after accounting for estimation error and model uncertainty. Consistent with models featuring unspanned macro factors, our forecasts of future bond excess returns are strongly negatively correlated with survey forecasts of short rates.

Key words: bond returns; yield curve; macro factors; stochastic volatility; time-varying parameters; unspanned macro risk factors.

JEL codes: G11, G12, G17

*We thank Pierluigi Balduzzi, Alessandro Beber, Carlos Carvalho, Jens Hilscher, Blake LeBaron, Spencer Martin and seminar participants at USC, University of Michigan, Central Bank of Belgium, Bank of Canada, Carleton University, Imperial College, ESSEC Paris, Boston Fed, SoFiE Meeting, McCombs School of Business, University of Connecticut, University of Illinois Urbana-Champaign, and Econometric Society Australasian Meeting (ESAM) for comments on the paper.

†University of Melbourne, Building 110, Room 11.042, 198 Berkeley Street, Melbourne, 3010. Email: antonio.gargano@unimelb.edu.au

‡Brandeis University, Sachar International Center, 415 South St, Waltham, MA, Tel: (781) 736-2834. Email: dpettenu@brandeis.edu

§University of California, San Diego, 9500 Gilman Drive, MC 0553, La Jolla CA 92093. Tel: (858) 534-0894. Email: atimmerm@ucsd.edu

1 Introduction

Treasury bonds play an important role in many investors' portfolios so an understanding of the risk and return dynamics for this asset class is of central economic importance.[1] Some studies document significant in-sample predictability of Treasury bond excess returns for 2-5 year maturities by means of variables such as forward spreads (Fama and Bliss, 1987), yield spreads (Campbell and Shiller, 1991), a linear combination of forward rates (Cochrane and Piazzesi, 2005) and factors extracted from a cross-section of macroeconomic variables (Ludvigson and Ng, 2009).

While empirical studies provide statistical evidence in support of bond return predictability, there is so far little evidence that such predictability could have been used in real time to improve investors' economic utility. Thornton and Valente (2012) find that forward spread predictors, when used to guide the investment decisions of an investor with mean-variance preferences, do not lead to higher out-of-sample Sharpe ratios or higher economic utility compared with investments based on a no-predictability expectations hypothesis (EH) model. Sarno et al. (2016) reach a similar conclusion.

To address this puzzling contradiction between the statistical and economic evidence on bond return predictability, we propose a new empirical modeling approach that generalizes the existing literature in economically insightful ways. Modeling bond return dynamics requires adding several features absent from the regression models used in the existing literature. First, bond prices, and thus bond returns, are sensitive to monetary policy and inflation prospects, both of which are known to shift over time.[2] This suggests that it is important to adopt a framework that accounts for time varying parameters and even for the possibility that the forecasting model may shift over time, requiring that we allow for model uncertainty. Second, uncertainty about inflation prospects changes over time and the volatility of bond yields has also undergone shifts–most notably during the Fed's monetarist experiment from 1979-1982–underscoring the need to allow for time varying volatility.[3] Third, risk-averse bond investors are concerned not only with the most likely outcomes but also with the degree of uncertainty surrounding future bond returns, indicating the need to model the full probability distribution of bond returns.

The literature on bond return predictability has noted the importance of parameter estimation error, model instability, and model uncertainty. However, no study on bond return

[1] According to the Securities Industry and Financial Markets Association, the size of the U.S. Treasury bond market was $11.9 trillion in 2013Q4. This is almost 30% of the entire U.S. bond market which includes corporate debt, mortgage and municipal bonds, money market instruments, agency and asset-backed securities.

[2] Stock and Watson (1999) and Cogley and Sargent (2002) find strong evidence of time variation in a Phillips curve model for U.S. inflation.

[3] Sims and Zha (2006) and Cogley et al. (2010) find that it is important to account for time varying volatility when modeling the dynamics of U.S. macroeconomic variables.

predictability has so far addressed how these considerations, jointly, impact the results. To accomplish this, we propose a novel Bayesian approach that brings several advantages to inference about the return prediction models and to their use in portfolio allocation analysis.

Our approach allows us, first, to integrate out uncertainty about the unknown parameters and to evaluate the effect of estimation error on the results. Even after observing 50 years of monthly observations, the coefficients of the return prediction models are surrounded by considerable uncertainty and so accounting for estimation error turns out to be important. Indeed, we find many cases with strong improvements in forecasting performance as a result of incorporating estimation error.[4]

Second, our approach produces predictive densities of bond excess returns. This allows us to analyze the economic value of bond return predictability from the perspective of an investor with power utility. Thornton and Valente (2012) are limited to considering mean-variance utility since they model only the first two moments of bond returns.[5]

Third, we allow for time varying (stochastic) volatility in the bond excess return model. Bond market volatility spiked during the monetarist experiment from 1979 to 1982, but we find clear advantages from allowing for stochastic volatility beyond this episode, particularly for bonds with shorter maturities. Stochastic volatility models do not, in general, lead to better point forecasts of bond returns but they produce far better density forecasts which, when used by a risk averse investor, generate better economic performance. In addition to lowering risk during periods with spikes in volatility, the stochastic volatility models imply that investors load more heavily on risky bonds during times with relatively low interest rate volatility such as during the 1990s.

A fourth advantage of our approach is that it allows for time variation in the regression parameters. We find evidence that the slope coefficients on both the yield spreads and the macrofactors vary considerably during our sample and that accounting for time varying parameters generally leads to more accurate forecasts and better economic performance. For example, whereas a univariate constant-coefficient model based on the Cochrane and Piazzesi (2005) forward rate factor does not improve forecasting performance relative to a constant expected return model, allowing the parameters of this model to vary over time leads to significantly better results.

Fifth, we address model uncertainty and model instability through forecast combination methods. Model uncertainty is important in our analysis which considers a variety of univariate

[4] Altavilla et al. (2014) find that an exponential tilting approach helps improve the accuracy of out-of-sample forecasts of bond yields. While their approach is not Bayesian, their tilting approach also attenuates the effect of estimation error on the model estimates.

[5] Sarno et al. (2016) use an approximate solution to compute optimal portfolio weights under power utility. They do not find evidence of economically exploitable return predictability but also do not consider parameter uncertainty.

and multivariate models as well as different model specifications. We consider equal-weighted averages of predictive densities, Bayesian model averaging, as well as combinations based on the optimal pooling method of Geweke and Amisano (2011). The latter forms a portfolio of the individual prediction models using weights that reflect the models' posterior probabilities. Models that are more strongly supported by the data get a larger weight in this average, so our combinations accommodate shifts in the relative forecasting performance of different models. The model combination results are generally better than the results for the individual models and thus suggest that model uncertainty–and model instability–can be effectively addressed through combination methods.[6]

Our empirical analysis uses the daily treasury yield data from Gurkaynak et al. (2007) to construct monthly excess returns for bond maturities between two and five years over the period 1962-2011. While previous studies have focused on the annual holding period, focusing on the higher frequency affords several advantages. Most obviously, it expands the number of non-overlapping observations, a point of considerable importance given the impact of parameter estimation error. Moreover, it allows us to identify short-lived dynamics in both first and second moments of bond returns which are missed by models of annual returns. This is an important consideration during events such as the global financial crisis of 2007-09 and around turning points of the economic cycle.

We conduct our analysis in the context of a three-variable model that includes the Fama-Bliss forward spread, the Cochrane-Piazzesi linear combination of forward rates, and a macro factor constructed using the methodology of Ludvigson and Ng (2009). Each variable is weighted according to its ability to improve on the predictive power of the bond return equation. Since forecasting studies have found that simpler models often do well in out-of-sample experiments, we also consider simpler univariate models.[7]

To assess the statistical evidence on bond return predictability, we use our models to generate out-of-sample forecasts over the period 1990-2011. Our return forecasts are based on recursively updated parameter estimates and use only historically available information, thus allowing us to assess how valuable the forecasts would have been to investors in real time. Compared to the benchmark EH model that assumes no return predictability, we find that many of the return predictability models generate significantly positive out-of-sample R^2 values. Interestingly, the Bayesian return prediction models generally perform better than the least squares counterparts so far explored in the literature.

Turning to the economic value of such out-of-sample forecasts, we next consider the portfolio

[6]Using an iterated combination approach, Lin et al. (2016) uncovers statistical and economic predictability in *corporate* bond returns

[7]Other studies considering macroeconomic determinants of the term structure of interest rates include Ang and Piazzesi (2003), Ang et al. (2007), Bikbov and Chernov (2010), Dewachter et al. (2014), Duffee (2011) and Joslin et al. (2014).

choice between a risk-free Treasury bill versus a bond with 2-5 years maturity for an investor with power utility. We find that the best return prediction models that account for volatility dynamics and changing parameters deliver sizeable gains in certainty equivalent returns relative to an EH model that assumes no predictability of bond returns, particularly in the absence of tight constraints on the portfolio weights.

There are several reasons why our findings differ from studies such as Thornton and Valente (2012) and Sarno et al. (2016) which argue that the statistical evidence on bond return predictability does not translate into economic gains. Allowing for stochastic volatility and time varying parameters, while accounting for parameter estimation error, leads to important gains in economic performance for many models.[8] Our results on forecast combinations also emphasize the importance of accounting for model uncertainty and the ability to capture changes in the performance of individual prediction models.

To interpret the economic sources of our findings on bond return predictability, we analyze the extent to which such predictability is concentrated in certain economic states and whether it is correlated with variables expected to be key drivers of time varying bond risk premia. We find that bond return predictability is stronger in recessions than during expansions, consistent with similar findings for stock returns by Henkel et al. (2011) and Dangl and Halling (2012). Using data from survey expectations we find that, consistent with a risk-premium story, our bond excess return forecasts are strongly negatively correlated with economic growth prospects (thus being higher during recessions) and strongly positively correlated with inflation uncertainty.

Our finding that the macro factor of Ludvigson and Ng (2009) possesses considerable predictive power over bond excess returns implies that information embedded in the yield curve does not subsume information contained in such macro variables. We address possible explanations of this finding, including the unspanned risk factor models of Joslin et al. (2014) and Duffee (2011) which suggest that macro variables move forecasts of future bond excess returns and forecasts of future short rates by the same magnitude but in opposite directions. We find support for this explanation as our bond excess return forecasts are strongly negatively correlated with survey forecasts of future short rates.

The outline of the paper is as follows. Section 2 describes the construction of the bond data, including bond returns, forward rates and the predictor variables. Section 3 sets up the prediction models and introduces our Bayesian estimation approach. Section 4 presents both full-sample and out-of-sample empirical results on bond return predictability. Section 5 assesses the economic value of bond return predictability for a risk averse investor when this investor uses the bond return predictions to form a portfolio of risky bonds and a risk-free asset. Section 6 analyzes economic sources of bond return predictability such as recession risk, time variations

[8] Thornton and Valente (2012) use a rolling window to update their parameter estimates but do not have a formal model that predicts future volatility or parameter values.

in inflation uncertainty, and the presence of unspanned risk factors. Section 7 presents model combination results and Section 8 concludes.

2 Data

This section describes how we construct our monthly series of bond returns and introduces the predictor variables used in the bond return models.

2.1 Returns and Forward Rates

Previous studies on bond return predictability such as Cochrane and Piazzesi (2005), Ludvigson and Ng (2009) and Thornton and Valente (2012) use overlapping 12-month returns data. This overlap induces strong serial correlation in the regression residuals. To handle this issue, we reconstruct the yield curve at the daily frequency starting from the parameters estimated by Gurkaynak et al. (2007), who rely on methods developed in Nelson and Siegel (1987) and Svensson (1994). Specifically, the time t zero coupon log yield on a bond maturing in n years, $y_t^{(n)}$, gets computed as[9]

$$\begin{aligned} y_t^{(n)} &= \beta_{0t} + \beta_{1t}\frac{1-\exp\left(-\frac{n}{\tau_1}\right)}{\frac{n}{\tau_1}} + \beta_{2t}\left[\frac{1-\exp\left(-\frac{n}{\tau_1}\right)}{\frac{n}{\tau_1}} - \exp\left(-\frac{n}{\tau_1}\right)\right] \\ &\quad + \beta_{3t}\left[\frac{1-\exp\left(-\frac{n}{\tau_2}\right)}{\frac{n}{\tau_2}} - \exp\left(-\frac{n}{\tau_2}\right)\right]. \end{aligned} \tag{1}$$

The parameters $(\beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2)$ are provided by Gurkaynak et al. (2007), who report daily estimates of the yield curve from June 1961 onward for the entire maturity range spanned by outstanding Treasury securities. We consider maturities ranging from 12 to 60 months and, in what follows, focus on the last day of each month's estimated log yields.[10]

Denote the frequency at which returns are computed by h, so $h = 1, 3$ for the monthly and quarterly frequencies, respectively. Also, let n be the bond maturity in years. For $n > h/12$ we compute returns and excess returns, relative to the $h-$period T-bill rate[11]

$$r_{t+h/12}^{(n)} = p_{t+h/12}^{(n-h/12)} - p_t^{(n)} = ny_t^{(n)} - (n-h/12)y_{t+h/12}^{(n-h/12)}, \tag{2}$$

$$rx_{t+h/12}^{(n)} = r_{t+h/12}^{(n)} - y_t^{h/12}(h/12). \tag{3}$$

[9] The third term was excluded from the calculations prior to January 1, 1980.

[10] The data is available at http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html. Because of idiosyncrasies at the very short end of the yield curve, we do not compute yields for maturities less than twelve months. For estimation purposes, the Gurkaynak et al. (2007) curve drops all bills and coupon bearing securities with a remaining time to maturity less than 6 months, while downweighting securities that are close to this window. The coefficients of the yield curve are estimated using daily cross-sections and thus avoid introducing look-ahead biases in the estimated yields.

[11] The formulas assume that the yields have been annualized, so we multiply $y_t^{(h/12)}$ by $h/12$.

Here $p_t^{(n)}$ is the logarithm of the time t price of a bond with n periods to maturity. Similarly, forward rates are computed as[12]

$$f_t^{(n-h/12,n)} = p_t^{(n-h/12)} - p_t^{(n)} = ny_t^{(n)} - (n-h/12)y_t^{(n-h/12)}. \tag{4}$$

2.2 Data Summary

Our bond excess return data span the period 1962:01-2011:12. We focus our analysis on the monthly holding period which offers several advantages over the annual returns data which have been the focus of most studies in the literature on bond return predictability. Most obviously, using monthly rather than annual data provides a sizeable increase in the number of (non-overlapping) data points available for model estimation. This is important in light of the low power of the return prediction models. Second, some of the most dramatic swings in bond prices occur over short periods of time lasting less than a year–e.g., the effect of the bankruptcy of Lehman Brothers on September 15, 2008–and are easily missed by models focusing on the annual holding period. This point is also important for the analysis of how return predictability is linked to recessions versus expansions; bond returns recorded at the annual horizon easily overlook important variations around turning points of the economic cycle.

Figure 1 plots monthly bond returns for the 2, 3, 4, and 5-year maturities, computed in excess of the 1-month T-bill rate. All four series are notably more volatile during 1979-82 and the volatility clearly increases with the maturity of the bonds. Panel A.1 in Table 1 presents summary statistics for the four monthly excess return series. Returns on the shortest maturities are right-skewed and fat-tailed, more so than the longer maturities. This observation suggests that it is inappropriate to use models that assume a normal distribution for bond returns.

Because the data used in our study differ from those used in most existing studies, it is worth highlighting the main differences and showing how they affect our data. Whereas most studies such as Cochrane and Piazzesi (2005) use overlapping 12-month bond returns computed in excess of the yield on a 12-month Treasury bill, we use monthly (non-overlapping) bond returns computed in excess of the 1-month T-bill rate.

Panels A.2 and A.3 in Table 1 provide summary statistics on the more conventional overlapping 12-month returns constructed either from our monthly data (Panel A.2) or as in Cochrane and Piazzesi (2005) (Panel A.3), using the Fama-Bliss CRSP files. The two series have very similar means which in turn are lower than the mean excess return on the monthly series in Panel A.1 due to the lower mean of the risk-free rate (1-month T-bill) used in Panel A.1 compared to the mean of the 12-month T-bill rate used in Panels A.2 and A.3. Comparing the monthly series in Panel A.1 to the 12-month series in Panels A.2 and A.3, we see that the kurtosis in the monthly excess returns data is much greater than in the 12-month data, while the serial

[12] For $n = h/12$, $f_t^{(n,n)} = ny_t^{(n)}$ and $y_t^{(n-h/12)} = y_t^{(0)}$ equals zero because $P_t^{(0)} = 1$ and its logarithm is zero.

correlation is much stronger in the 12-month series. Both findings are unsurprising because of the overlap in the 12-month return data which creates a smoother return series.

2.3 Predictor variables

Our empirical strategy entails regressing bond excess returns on a range of the most prominent predictors proposed in the literature on bond return predictability. Specifically, we consider forward spreads as proposed by Fama and Bliss (1987), a linear combination of forward rates as proposed by Cochrane and Piazzesi (2005), and a linear combination of macro factors, as proposed by Ludvigson and Ng (2009).

To motivate our use of these three predictor variables, note that the n-period bond yield is related to expected future short yields and expected future excess returns (Duffee, 2013):

$$y_t^{(n)} = \frac{1}{n}\left(\sum_{j=0}^{n-1} E[y_{t+j}^{(1)}|\boldsymbol{z}_t]\right) + \frac{1}{n}\left(\sum_{j=0}^{n-1} E_t[rx_{t+j+1}^{(n-j)}|\boldsymbol{z}_t]\right), \tag{5}$$

where $rx_{t+j+1}^{(n-j)}$ is the excess return in period $t+j+1$ on a bond with $n-j$ periods to maturity and $E[.|\boldsymbol{z}_t]$ denotes the conditional expectation given market information at time t, $\boldsymbol{z}_t$. Equation (5) suggests that current yields or, equivalently, forward spreads should have predictive power over future bond excess returns and so motivates our use of these variables in the excess return regressions.

The use of non-yield predictors is more contentious. In fact, if the vector of conditioning information variables, $\boldsymbol{z}_t$, is of sufficiently low dimension, we can invert (5) to get $\boldsymbol{z}_t = g(\boldsymbol{y}_t)$. In this case information in the current yield curve subsumes all other predictors of future excess returns and so macro variables should be irrelevant when added to the prediction model. The unspanned risk factor models of Joslin et al. (2014) and Duffee (2011) offer an explanation for why macro variables help predict bond excess returns over and above information contained in the yield curve. These models suggest that the effect of additional state variable on expected future short rates and expected future bond excess returns cancel out in Equation (5). Such cancellations imply that the additional state variables do not show up in bond yields although they can have predictive power over bond excess returns.

Our predictor variables are computed as follows. The Fama-Bliss (FB) forward spreads are given by

$$fs_t^{(n,h)} = f_t^{(n-h/12,n)} - y_t^{(h/12)}(h/12). \tag{6}$$

The Cochrane-Piazzesi (CP) factor is given as a linear combination of forward rates

$$CP_t^h = \hat{\boldsymbol{\gamma}}^{h\prime}\boldsymbol{f}_t^{(\boldsymbol{n}-h/12,\boldsymbol{n})}, \tag{7}$$

where

$$\boldsymbol{f}_t^{(\boldsymbol{n}-h/12,\boldsymbol{n})} = \left[f_t^{(n_1-h/12,n_1)}, f_t^{(n_2-h/12,n_2)}, ..., f_t^{(n_k-h/12,n_k)}\right].$$

Here $\boldsymbol{n} = [1, 2, 3, 4, 5]$ denotes the vector of maturities measured in years. As in Cochrane and Piazzesi (2005), the coefficient vector $\hat{\boldsymbol{\gamma}}$ is estimated from

$$\frac{1}{4}\sum_{n=2}^{5} rx_{t+h/12}^{(n)} = \gamma_0^h + \gamma_1^h f_t^{(1-1/12,1)} + \gamma_2^h f_t^{(2-1/12,2)} + \gamma_3^h f_t^{(3-1/12,3)} + \gamma_4^h f_t^{(4-1/12,4)} + \gamma_5^h f_t^{(5-1/12,5)} + \overline{\varepsilon}_{t+h/12}. \tag{8}$$

Ludvigson and Ng (2009) propose to use macro factors to predict bond returns. Suppose we observe a $T \times M$ panel of macroeconomic variables $\{x_{i,t}\}$ generated by a factor model

$$x_{i,t} = \kappa_i g_t + \epsilon_{i,t}, \tag{9}$$

where g_t is an $s \times 1$ vector of common factors and $s << M$. The unobserved common factor, g_t is replaced by an estimate, $\hat{g}_t$, obtained using principal components analysis. Following Ludvigson and Ng (2009), we build a single linear combination from a subset of the first eight estimated principal components, $\widehat{\boldsymbol{G}}_t = [\hat{g}_{1,t}, \hat{g}_{1,t}^3, \hat{g}_{3,t}, \hat{g}_{4,t}, \hat{g}_{8,t}]$ to obtain the LN factor[13]

$$LN_t^h = \hat{\boldsymbol{\lambda}}^{h\prime} \widehat{\boldsymbol{G}}_t, \tag{10}$$

where $\hat{\boldsymbol{\lambda}}$ is obtained from the projection

$$\frac{1}{4}\sum_{n=2}^{5} rx_{t+h/12}^{(n)} = \lambda_0^h + \lambda_1^h \hat{g}_{1,t} + \lambda_2^h \hat{g}_{1,t}^3 + \lambda_3^h \hat{g}_{3,t} + \lambda_4^h \hat{g}_{4,t} + \lambda_5^h \hat{g}_{8,t} + \overline{\eta}_{t+h/12}. \tag{11}$$

Panel B in Table 1 presents summary statistics for the Fama-Bliss forward spreads along with the CP and LN factors. The Fama-Bliss forward spreads are strongly positively autocorrelated with first-order autocorrelation coefficients around 0.90. The CP and LN factors are far less autocorrelated with first-order autocorrelations of 0.67 and 0.41, respectively.

Panel C shows that the Fama-Bliss spreads are strongly positively correlated. In turn, these spreads are positively correlated with the CP factor, with correlations around 0.5, but are uncorrelated with the LN factor. The LN factor captures a largely orthogonal component in relation to the other predictors. For example, its correlation with CP is only 0.18. It is also less persistent than the FB and CP factors.

3 Return Prediction Models and Estimation Methods

We next introduce the return prediction models and describe the estimation methods used in the paper.

[13] Ludvigson and Ng (2009) selected this particular combination of factors using the Schwarz information criterion.

3.1 Model specifications

Our analysis considers the three predictor variables described in the previous section. Specifically, we consider three univariate models, each of which includes one of these three factors, along with a multivariate model that includes all three predictors. This produces a total of four models:

1. Fama-Bliss (FB) univariate

$$rx_{t+h/12}^{(n)} = \beta_0 + \beta_1 fs_t^{(n,h)} + \varepsilon_{t+h/12}. \tag{12}$$

2. Cochrane-Piazzesi (CP) univariate

$$rx_{t+h/12}^{(n)} = \beta_0 + \beta_1 CP_t^h + \varepsilon_{t+h/12}. \tag{13}$$

3. Ludvigson-Ng (LN) univariate

$$rx_{t+h/12}^{(n)} = \beta_0 + \beta_1 LN_t^h + \varepsilon_{t+h/12}. \tag{14}$$

4. Fama-Bliss, Cochrane-Piazzesi and Ludvigson-Ng predictors (FB-CP-LN)

$$rx_{t+h/12}^{(n)} = \beta_0 + \beta_1 fs_t^{(n,h)} + \beta_2 CP_t^h + \beta_3 LN_t^h + \varepsilon_{t+h/12}. \tag{15}$$

These models are in turn compared against the Expectation Hypothesis benchmark

$$rx_{t+h/12}^{(n)} = \beta_0 + \varepsilon_{t+h/12}, \tag{16}$$

that assumes no predictability. In each case $n \in \{2, 3, 4, 5\}$.

We consider four classes of models: (i) constant coefficient models with constant volatility; (ii) constant coefficient models with stochastic volatility; (iii) time varying parameter models with constant volatility; and (iv) time varying parameter models with stochastic volatility.

The constant coefficient, constant volatility model serves as a natural starting point for the out-of-sample analysis. There is no guarantee that the more complicated models with stochastic volatility and time varying regression coefficients produce better out-of-sample forecasts since their parameters may be imprecisely estimated.

To estimate the models we adopt a Bayesian approach that offers several advantages over the conventional estimation methods adopted by previous studies on bond return predictability. First, imprecisely estimated parameters is a big issue in the return predictability literature and so it is important to account for parameter uncertainty as is explicitly done by the Bayesian approach. Second, portfolio allocation analysis requires estimating not only the conditional

mean, but also the conditional variance (under mean-variance preferences) or the full predictive density (under power utility) of returns. This is accomplished by our method which generates the (posterior) predictive return distribution. Third, our approach also allows us to handle model uncertainty (and model instability) by combining forecasting models.

We next describe our estimation approach for each of the four classes of models. To ease the notation, for the remainder of the paper we drop the notation $t + h/12$ and replace $h/12$ with 1, with the understanding that the definition of a period depends on the data frequency.

3.2 Constant Coefficients and Constant Volatility Model

The linear model projects bond excess returns $rx_{\tau+1}^{(n)}$ on a set of lagged predictors, $\boldsymbol{x}_{\tau}^{(n)}$:

$$\begin{aligned} rx_{\tau+1}^{(n)} &= \mu + \boldsymbol{\beta}'\boldsymbol{x}_{\tau}^{(n)} + \varepsilon_{\tau+1}, \quad \tau = 1, ..., t-1, \\ \varepsilon_{\tau+1} &\sim \mathcal{N}(0, \sigma_{\varepsilon}^2). \end{aligned} \tag{17}$$

Ordinary least squares (OLS) estimation of this model is straightforward and so is not further explained. However, we also consider Bayesian estimation so we briefly describe how the prior and likelihood are specified. Following standard practice, the priors for the parameters μ and $\boldsymbol{\beta}$ in (17) are assumed to be normal and independent of σ_{ε}^2

$$\begin{bmatrix} \mu \\ \boldsymbol{\beta} \end{bmatrix} \sim \mathcal{N}\left(\underline{\boldsymbol{b}}, \underline{\boldsymbol{V}}\right), \tag{18}$$

where

$$\underline{\boldsymbol{b}} = \begin{bmatrix} \overline{rx}_t^{(n)} \\ \mathbf{0} \end{bmatrix}, \ \underline{\boldsymbol{V}} = \underline{\psi}^2 \left[\left(s_{rx,t}^{(n)}\right)^2 \left(\sum_{\tau=1}^{t-1} \boldsymbol{x}_{\tau}^{(n)} \boldsymbol{x}_{\tau}^{(n)\prime} \right)^{-1} \right], \tag{19}$$

and $\overline{rx}_t^{(n)}$ and $\left(s_{rx,t}^{(n)}\right)^2$ are data-based moments:

$$\begin{aligned} \overline{rx}_t^{(n)} &= \frac{1}{t-1} \sum_{\tau=1}^{t-1} rx_{\tau+1}^{(n)}, \\ \left(s_{rx,t}^{(n)}\right)^2 &= \frac{1}{t-2} \sum_{\tau=1}^{t-1} \left(rx_{\tau+1}^{(n)} - \overline{rx}_t^{(n)}\right)^2. \end{aligned}$$

Our choice of the prior mean vector $\underline{\boldsymbol{b}}$ reflects the "no predictability" view that the best predictor of bond excess returns is the average of past returns. We therefore center the prior intercept on the prevailing mean of historical excess returns, while the prior slope coefficient is centered on zero.

It is common to base the priors of the hyperparameters on sample estimates, see Stock and Watson (2006) and Efron (2010). Our analysis can thus be viewed as an empirical Bayes approach rather than a more traditional Bayesian approach that fixes the prior distribution

before any data are observed. We find that, at least for a reasonable range of values, the choice of priors has modest impact on our results. In (19), $\underline{\psi}$ is a constant that controls the tightness of the prior, with $\underline{\psi} \to \infty$ corresponding to a diffuse prior on μ and $\boldsymbol{\beta}$. Our benchmark analysis sets $\underline{\psi} = n/2$. This choice means that the prior becomes looser for the longer bond maturities for which fundamentals-based information is likely to be more important. It also means that the posterior parameter estimates are shrunk more towards their priors for the shortest maturities which are most strongly affected by estimation error.

We assume a standard gamma prior for the error precision of the return innovation, $\sigma_{\varepsilon}^{-2}$:

$$\sigma_{\varepsilon}^{-2} \sim \mathcal{G}\left(\left(s_{rx,t}^{(n)}\right)^{-2}, \underline{v}_0 (t-1)\right), \tag{20}$$

where $\underline{v}_0$ is a prior hyperparameter that controls how informative the prior is with $\underline{v}_0 \to 0$ corresponding to a diffuse prior on $\sigma_{\varepsilon}^{-2}$. Our baseline analysis sets $\underline{v}_0 = 2/n$, again letting the priors be more diffuse, the longer the bond maturity.

3.3 Stochastic Volatility Model

A large literature has found strong empirical evidence of time varying return volatility, see Andersen et al. (2006). We accommodate such effects through a simple stochastic volatility (SV) model:

$$rx_{\tau+1}^{(n)} = \mu + \boldsymbol{\beta}' \boldsymbol{x}_{\tau}^{(n)} + \exp(h_{\tau+1})\, u_{\tau+1}, \tag{21}$$

where $h_{\tau+1}$ denotes the (log of) bond return volatility at time $\tau+1$ and $u_{\tau+1} \sim \mathcal{N}(0,1)$. The log-volatility $h_{\tau+1}$ is assumed to follow a stationary and mean reverting process:

$$h_{\tau+1} = \lambda_0 + \lambda_1 h_{\tau} + \xi_{\tau+1}, \tag{22}$$

where $\xi_{\tau+1} \sim \mathcal{N}\left(0, \sigma_{\xi}^2\right)$, $|\lambda_1| < 1$, and u_{τ} and ξ_s are mutually independent for all τ and s. Appendix A explains how we estimate the SV model and set the priors.

3.4 Time varying Parameter Model

Studies such as Thornton and Valente (2012) find considerable evidence of instability in the parameters of bond return prediction models. The following time varying parameter (TVP) model allows the regression coefficients in (17) to change over time:

$$\begin{aligned} rx_{\tau+1}^{(n)} &= (\mu + \mu_{\tau}) + (\boldsymbol{\beta} + \boldsymbol{\beta}_{\tau})' \boldsymbol{x}_{\tau}^{(n)} + \varepsilon_{\tau+1}, \quad \tau = 1, ..., t-1, \\ \varepsilon_{\tau+1} &\sim \mathcal{N}(0, \sigma_{\varepsilon}^2). \end{aligned} \tag{23}$$

The intercept and slope parameters $\boldsymbol{\theta}_{\tau} = \left(\mu_{\tau}, \boldsymbol{\beta}_{\tau}'\right)'$ are assumed to follow a zero-mean, stationary process

$$\boldsymbol{\theta}_{\tau+1} = diag\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}\right) \boldsymbol{\theta}_{\tau} + \boldsymbol{\eta}_{\tau+1}, \tag{24}$$

where $\boldsymbol{\theta}_1 = \mathbf{0}$, $\boldsymbol{\eta}_{\tau+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$, and the elements in $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$ are restricted to lie between -1 and 1. In addition, ε_τ and $\boldsymbol{\eta}_s$ are mutually independent for all τ and s.[14] The key parameter is $\boldsymbol{Q}$ which determines how rapidly the parameters $\boldsymbol{\theta}$ are allowed to change over time. We set the priors to ensure that the parameters are allowed to change only gradually. Again Appendix A provides details on how we estimate the model and set the priors.

3.5 Time varying Parameter, Stochastic Volatility Model

Finally, we consider a general model that admits both time varying parameters and stochastic volatility (TVP-SV):

$$rx_{\tau+1}^{(n)} = (\mu + \mu_\tau) + (\boldsymbol{\beta} + \boldsymbol{\beta}_\tau)' \boldsymbol{x}_\tau^{(n)} + \exp(h_{\tau+1})\, u_{\tau+1}, \tag{25}$$

with

$$\boldsymbol{\theta}_{\tau+1} = diag(\boldsymbol{\gamma}_{\boldsymbol{\theta}})\, \boldsymbol{\theta}_\tau + \boldsymbol{\eta}_{\tau+1}, \tag{26}$$

where again $\boldsymbol{\theta}_\tau = (\mu_\tau, \boldsymbol{\beta}_\tau')'$, and

$$h_{\tau+1} = \lambda_0 + \lambda_1 h_\tau + \xi_{\tau+1}. \tag{27}$$

We assume that $u_{\tau+1} \sim \mathcal{N}(0,1)$, $\boldsymbol{\eta}_{\tau+1} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{Q})$, $\xi_{\tau+1} \sim \mathcal{N}\left(0, \sigma_\xi^2\right)$ and u_τ, $\boldsymbol{\eta}_s$ and ξ_l are mutually independent for all τ, s, and l. Again we refer to Appendix A for further details on this model.

The models are estimated by Gibbs sampling methods. This allows us to generate draws of excess returns, $rx_{t+1}^{(n)}$, in a way that only conditions on a given model and the data at hand. This is convenient when computing bond return forecasts and determining the optimal bond holdings.

4 Empirical Results

This section describes our empirical results. For comparison with the existing literature, and to convey results on the importance of different features of the models, we first report results based on full-sample estimates. This is followed by an out-of-sample analysis of both the statistical and economic evidence on return predictability.

4.1 Full-sample Estimates

For comparison with extant results, Table 2 presents full-sample (1962:01-2011:12) least squares estimates for the bond return prediction models with constant parameters. While no investors could have based their historical portfolio choices on these estimates, such results are important

[14] This is equivalent to writing $rx_{\tau+1}^{(n)} = \widetilde{\mu}_\tau + \widetilde{\boldsymbol{\beta}}_\tau' \boldsymbol{x}_\tau^{(n)} + \varepsilon_{\tau+1}$, where $\widetilde{\boldsymbol{\theta}}_1 \equiv \left(\widetilde{\mu}_1, \widetilde{\boldsymbol{\beta}}_1'\right)$ is left unrestricted.

for our understanding of how the various models work. The slope coefficients for the univariate models increase monotonically in the maturity of the bonds. With the exception of the coefficients on the CP factor in the multivariate model, the coefficients are significant across all maturities and forecasting models.[15]

Table 2 shows R^2 values around 1-2% for the model that uses FB as a predictor, 2.5% for the model that uses the CP factor and around 5% for the model based on the LN factor. These values, which increase to 7-8% for the multivariate model, are notably smaller than those conventionally reported for the overlapping 12-month horizon. For comparison, at the one-year horizon we obtain R^2 values of 10-11%, 17-24%, and 14-19% for the FB, CP, and LN models, respectively.[16]

The extent of time variations in the parameters of the three-factor FB-CP-LN model is displayed in Figure 2. All coefficients are notably volatile around 1980 and the coefficients continue to fluctuate throughout the sample, particularly for the CP regressor.

An advantage of our approach is its ability to deal with parameter estimation error. To get a sense of the importance of this issue, Figure 3 plots full-sample posterior densities of the regression coefficients for the three-factor model that uses the FB, CP and LN factors as predictors. The spread of the densities in this figure shows the considerable uncertainty surrounding the parameter estimates even at the end of the sample. As expected, parameter uncertainty is greatest for the TVP and TVP-SV models which allow for the greatest amount of flexibility–clearly this comes at the cost of less precisely estimated parameters. The SV model generates more precise estimates than the constant volatility benchmark, reflecting the ability of the SV model to reduce the weight on observations in highly volatile periods.

The effect of such parameter uncertainty on the predictive density of bond excess returns is depicted in Figure 4. This figure evaluates the univariate LN model at the mean of this predictor, plus or minus two times its standard deviation. The TVP and TVP-SV models imply a greater dispersion for bond returns and their densities shift further out in the tails as the predictor variable moves away from its mean. The four models clearly imply very different probability distributions for bond returns and so have very different implications when used by investors to form portfolios.

Figure 5 plots the time series of the posterior means and volatilities of bond excess returns for the FB-CP-LN model. Mean excess returns (top panel) vary substantially during the sample,

[15] As emphasized by Cochrane and Piazzesi (2005), care has to be exercised when evaluating the statistical significance of these results due to the highly persistent FB and CP regressors. Wei and Wright (2013) find that conventional tests applied to bond excess return regressions that use yield spreads or yields as predictors are subject to considerable finite-sample distortions. However, their reverse regression approach confirms that, even after accounting for such biases, bond excess returns still appear to be predictable.

[16] These values are a bit lower than those reported in the literature but are consistent with the range of results reported by Duffee (2013). The weaker evidence reflects our use of an extended sample along with a tendency for the regression coefficients to decline towards zero at the end of the sample.

peaking during the early eighties, and again during 2008. Stochastic volatility effects (bottom panel) also appear to be empirically important. The conditional volatility is very high during 1979-1982, while subsequent spells with above-average volatility are more muted and short-lived. Interestingly, there are relatively long spells with below-average conditional volatility such as during the late nineties and mid-2000s.

4.2 Calculation of out-of-sample Forecasts

To gauge the real-time value of the bond return prediction models, following Ludvigson and Ng (2009) and Thornton and Valente (2012), we next conduct an out-of-sample forecasting experiment.[17] This experiment relies on information up to period t to compute return forecasts for period $t+1$ and uses an expanding estimation window. Notably, when constructing the CP and LN factors we also restrict our information to end at time t. Hence, we re-estimate each period the principal components and the regression coefficients in equations (8) and (11).

We use 1962:01-1989:12 as our initial warm-up estimation sample and 1990:01-2011:12 as the forecast evaluation period. As before, we set $n = 2, 3, 4, 5$ and so predict 2, 3, 4, and 5-year bond returns in excess of the one-month T-bill rate.

The predictive accuracy of the bond excess return forecasts is measured relative to recursively updated forecasts from the expectations hypothesis (EH) model (16) that projects excess returns on a constant. Specifically, at each point in time we obtain draws from the predictive densities of the benchmark model and the models with time varying predictors. For a given bond maturity, n, we denote draws from the predictive density of the EH model, given the information set at time t, $\mathcal{D}^t = \{rx_{\tau+1}^{(n)}\}_{\tau=1}^{t-1}$, by $\left\{rx_{t+1}^{(n),j}\right\}, j = 1, ..., J$. Similarly, draws from the predictive density of any of the other models (labeled model i) given $\mathcal{D}^t = \{rx_{\tau+1}^{(n)}, \boldsymbol{x}_\tau^{(n)}\}_{\tau=1}^{t-1} \cup \boldsymbol{x}_t^{(n)}$ are denoted $\left\{rx_{t+1,i}^{(n),j}\right\}, j = 1, ..., J$.[18]

For the constant parameter, constant volatility model, return draws are obtained by applying a Gibbs sampler to

$$p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right) = \int p\left(rx_{t+1}^{(n)}\middle|\,\mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \mathcal{D}^t\right) p\left(\mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}\middle|\,\mathcal{D}^t\right) d\mu d\boldsymbol{\beta} d\sigma_\varepsilon^{-2}. \tag{28}$$

[17] Out-of-sample analysis also provides a way to guard against overfitting. Duffee (2010) shows that in-sample overfitting can generate unrealistically high Sharpe ratios.

[18] We run the Gibbs sampling algorithms recursively for all time periods betweeen 1990:01 and 2011:12. At each point in time, we retain 1,000 draws from the Gibbs samplers after a burn-in period of 500 iterations. For the TVP, SV, and TVP-SV models we run the Gibbs samplers five times longer while at the same time thinning the chains by keeping only one in every five draws, thus effectively eliminating any autocorrelation left in the draws. Additional details on these algorithms are presented in Appendix A.

Return draws for the most general TVP-SV model are obtained from the predictive density[19]

$$\begin{aligned} p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right) &= \int p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},h_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right) \\ &\quad\times p\left(\boldsymbol{\theta}_{t+1},h_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right) \\ &\quad\times p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\middle|\,\mathcal{D}^t\right) d\mu d\boldsymbol{\beta} d\boldsymbol{\theta}^{t+1} d\boldsymbol{\gamma_\theta} d\boldsymbol{Q} dh^{t+1} d\lambda_0 d\lambda_1 d\sigma_\xi^{-2}. \end{aligned} \tag{29}$$

where $h^{t+1}=(h_1,...,h_{t+1})$ and $\boldsymbol{\theta}^{t+1}=(\boldsymbol{\theta}_1,...,\boldsymbol{\theta}_{t+1})$ denote the sequence of conditional variance states and time varying regression parameters up to time $t+1$, respectively. Draws from the SV and TVP models are obtained as special cases of (29). All Bayesian models integrate out uncertainty about the parameters.

4.3 Forecasting Performance

Although our models generate a full predictive distribution for bond returns it is insightful to also report results based on conventional point forecasts. To obtain point forecasts we first compute the posterior mean from the densities in (28) and (29). We denote these by $\overline{rx}_{t,EH}^{(n)}=\frac{1}{J}\sum_{j=1}^{J}rx_t^{(n),j}$ and $\overline{rx}_{t,i}^{(n)}=\frac{1}{J}\sum_{j=1}^{J}rx_{t,i}^{(n),j}$, for the EH and alternative models, respectively. Using such point forecasts, we obtain the corresponding forecast errors as $e_{t,EH}^{(n)}=rx_t^{(n)}-\overline{rx}_{t,EH}^{(n)}$ and $e_{t,i}^{(n)}=rx_t^{(n)}-\overline{rx}_{t,i}^{(n)}$, $t=\underline{t},...,\bar{t}$, where $\underline{t}=1990:01$ and $\bar{t}=2011:12$ denote the beginning and end of the forecast evaluation period.

Following Campbell and Thompson (2008), we compute the out-of-sample R^2 of model i relative to the EH model as

$$R_{OoS,i}^{(n)2}=1-\frac{\sum_{\tau=\underline{t}}^{\bar{t}}e_{\tau,i}^{(n)2}}{\sum_{\tau=\underline{t}}^{\bar{t}}e_{\tau,EH}^{(n)2}}. \tag{30}$$

Positive values of this statistic suggest evidence of time varying return predictability.

Table 3 reports R_{OoS}^2 values for the OLS, linear, SV, TVP and TVP-SV models across the four bond maturities. For the two-year maturity we find little evidence that models estimated by OLS are able to improve on the predictive accuracy of the EH model. Conversely, all models estimated using our Bayesian approach generate significantly more accurate forecasts at either the 10% or 1% significance levels, using the test for equal predictive accuracy suggested by Clark and West (2007). Similar results are obtained for the SV and TVP models which generate R_{OoS}^2 values close to or above 5% for the models that include the LN predictor. Interestingly, the best results are obtained for the TVP-SV model which notably outperforms the other univariate models based on the FB or CP variables.

While the OLS models fare considerably better for the 3-5 year bond maturities, the ability of the linear Bayesian model to generate accurate forecasts does not appear to depend as

[19]For each draw retained from the Gibbs sampler, we produce 100 draws from the corresponding predictive densities.

strongly on the maturity. Moreover, the Bayesian approach performs notably better than its OLS counterpart, particularly for the multivariate model which requires estimating more parameters which increases the importance of parameter estimation error.

Comparing results across predictor variables, the univariate CP model performs worst with only the TVP and TVP-SV specifications suggesting some (weak) evidence of improvements in the predictive accuracy over the EH benchmark. This suggests that the CP predictor has some predictive power over bond excess returns but that its coefficient is unstable over time, consistent with the third plot in Figure 2. Conversely, the FB and, in particular, the LN predictor, add considerable improvements in out-of-sample predictive performance.

Ranking the different specifications across bond maturities and predictor variables, we find that the TVP-SV models produce the best out-of-sample forecasts in around half of all cases with the SV model a distant second best. These results suggest that the more sophisticated models that allow for time varying parameters and time varying volatility manage to produce better out-of-sample forecasts than simple constant parameter, constant volatility models. Even in cases where the TVP-SV model is not the best specification, it still performs nearly as well as the best model. In contrast, there are instances where the other models are highly inferior to the TVP-SV model.

To identify which periods the models perform best, following Welch and Goyal (2008), we use the out-of-sample forecast errors to compute the difference in the cumulative sum of squared errors (SSE) for the EH model versus the ith model:

$$\Delta CumSSE_{t,i}^{(n)} = \sum_{\tau=\underline{t}}^{t} \left(e_{\tau,EH}^{(n)}\right)^2 - \sum_{\tau=\underline{t}}^{t} \left(e_{\tau,i}^{(n)}\right)^2. \tag{31}$$

Positive and increasing values of $\Delta CumSSE_t$ suggest that the model with time varying return predictability generates more accurate point forecasts than the EH benchmark.

Figure 6 plots $\Delta CumSSE_t$ for the three univariate models and the three-factor model, assuming a two-year bond maturity. These plots show periods during which the various models perform well relative to the EH model–periods where the lines are increasing and above zero–and periods where the models underperform against this benchmark–periods with decreasing graphs. The univariate FB model performs quite poorly due to spells of poor performance in 1994, 2000, and 2008, while the CP model underperforms between 1993 and 2006. In contrast, except for a few isolated months in 2002, 2008 and 2009, the LN model consistently beats the EH benchmark up to 2010, at which point its performance flattens against the EH model. A similar performance is seen for the multivariate model.

The predictive accuracy measures in (30) and (31) ignore information on the full probability distribution of returns. To evaluate the accuracy of the density forecasts obtained in (28) and (29), we compute the implied predictive likelihoods. The predictive likelihood, or score, is

commonly viewed as the broadest measure of accuracy of density forecasts, see, e.g., Geweke and Amisano (2010). At each point in time t, the log predictive score is obtained by taking the natural log of the predictive densities (28)–(29) evaluated at the observed bond excess return, $rx_t^{(n)}$, denoted by $LS_{t,EH}$ and $LS_{t,i}$ for the EH and alternative models, respectively.

Table 4 reports the average log-score differential for each of our models, again measured relative to the EH benchmark.[20] The results show that the linear model performs significantly better than the EH benchmark across all specifications for the four- and five-year bond maturities and for the models containing the LN predictor for the two- and three-year bond maturities. The TVP model produces similar, if modest, improvements in the log-score of the EH model. Notably stronger results are obtained for the SV and TVP-SV models which dominate the EH benchmark across the board.

Figure 7 supplements Table 4 by showing the cumulative log score (LS) differentials between the EH model and the ith model, computed analogously to (31) as

$$\Delta CumLS_{t,i} = \sum_{\tau=\underline{t}}^{t} \left[LS_{\tau,i} - LS_{\tau} \right]. \tag{32}$$

The dominant performance of the density forecasts generated by the SV and TVP-SV models is clear from these plots. In contrast, the linear and TVP models offer only modest improvements over the EH benchmark by this measure.

4.4 Robustness to Choice of Priors

Choice of priors can always be debated in Bayesian analysis, so we conduct a sensitivity analysis with regard to two of the priors, namely $\underline{\psi}$ and $\underline{v}_0$, which together control how informative the baseline priors are. Our first experiment sets $\underline{\psi} = 5$ and $\underline{v}_0 = 1/5$. This choice corresponds to using more diffuse priors than in the baseline scenario. Compared with the baseline prior, this prior produces worse results (lower out-of-sample R^2 values) for the two shortest maturities ($n = 2, 3$), but stronger results for the longest maturities ($n = 4, 5$).

Our second experiment sets $\underline{\psi} = 0.5, \underline{v}_0 = 5$, corresponding to tighter priors. Under these priors, the results improve for the shorter bond maturities but get weaker at the longest maturities. In both cases, the conclusion that the best prediction models dominate the EH benchmark continues to hold even for such large shifts in priors. Detailed results are available on request.

[20] To test if the differences in forecast accuracy are significant, we follow Clark and Ravazzolo (2015) and apply the Diebold and Mariano (1995) t-test for equality of the average log-scores based on the statistic $\overline{LS}_i = \frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} (LS_{\tau.i} - LS_{\tau,EH})$. The p-values for this statistic are based on t-statistics computed with a serial correlation-robust variance, using the pre-whitened quadratic spectral estimator of Andrews and Monahan (1992). Monte Carlo evidence in Clark and McCracken (2011) indicates that, with nested models, the Diebold-Mariano test compared against normal critical values can be viewed as a somewhat conservative test for equal predictive accuracy in finite samples. Since all models considered here nest the EH benchmark, we report p-values based on one-sided tests, taking the nested EH benchmark as the null and the nesting model as the alternative.

5 Economic Value of Return Forecasts

So far our analysis concentrated on statistical measures of predictive accuracy. It is important to evaluate the extent to which the apparent gains in predictive accuracy translate into better investment performance. In fact, for an investor with mean-variance preferences, Thornton and Valente (2012) find that improvements in the statistical accuracy of bond return forecasts do not imply improved portfolio performance.

5.1 Bond Holdings

We consider the asset allocation decisions of an investor that selects the weight, $\omega_t^{(n)}$, on a risky bond with n periods to maturity versus a one-month T-bill that pays the riskfree rate, $\tilde{y}_t = y_t^{(1/12)}$. The investor has power utility and coefficient of relative risk aversion A:

$$U\left(\omega_t^{(n)}, rx_{t+1}^{(n)}\right) = \frac{\left[\left(1-\omega_t^{(n)}\right)\exp\left(\tilde{y}_t\right) + \omega_t^{(n)}\exp\left(\tilde{y}_t + rx_{t+1}^{(n)}\right)\right]^{1-A}}{1-A}, \quad A > 0. \tag{33}$$

Using all information at time t, $\mathcal{D}^t$, to evaluate the predictive density of $rx_{t+1}^{(n)}$, the investor solves the optimal asset allocation problem

$$\omega_t^{(n)*} = \arg\max_{\omega_t^{(n)}} \int U\left(\omega_t^{(n)}, rx_{t+1}^{(n)}\right) p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right) drx_{t+1}^{(n)}. \tag{34}$$

The integral in (34) can be approximated by generating a large number of draws, $rx_{t+1,i}^{(n),j}$, $j = 1, .., J$, from the predictive densities specified in (28) and (29). For each of the candidate models, i, we approximate the solution to (34) by

$$\widehat{\omega}_{t,i}^{(n)} = \arg\max_{\omega_{t,i}^{(n)}} \frac{1}{J}\sum_{j=1}^{J}\left\{\frac{\left[\left(1-\omega_{t,i}^{(n)}\right)\exp\left(\tilde{y}_t\right) + \omega_{t,i}^{(n)}\exp\left(\tilde{y}_t + rx_{t+1,i}^{(n),j}\right)\right]^{1-A}}{1-A}\right\}. \tag{35}$$

The resulting sequences of portfolio weights $\left\{\widehat{\omega}_{t,EH}^{(n)}\right\}$ and $\left\{\widehat{\omega}_{t,i}^{(n)}\right\}$ are used to compute realized utilities. For each model, i, we convert these into certainty equivalent returns (CER), i.e., values that equate the average utility of the EH model with the average utility of any of the alternative models.

We consider two different sets of assumptions about the portfolio weights. The first scenario restricts the weights on the risky bonds to the interval $[0, 0.99]$ to ensure that the expected utility is finite even with an unbounded return distribution; see Geweke (2001) and Kandel and Stambaugh (1996) for a discussion of this point. The second scenario instead restricts the weights to lie between -2 and 3 but bounds the monthly return distribution between -100% and 100%. This follows the procedure in Johannes et al. (2014) who argue that this approach ensures that the expected utility exists.

In both cases we set the coefficient of relative risk aversion to $A = 10$, a value higher than normally considered. Our choice reflects the high Sharpe ratios observed for the bond portfolios during our sample—see Table 1. For lower values of A, this causes the weights on the risky bonds to almost always hit the upper bound (0.99) of the first scenario for both the EH and time varying predictability models and so does not allow us to differentiate between these models. However, provided that we do not impose too tight limits on the bond portfolio weights (Scenario 2), informative results can still be obtained for lower values of A, e.g., $A = 5$.

5.2 Empirical Results

Table 5 shows annualized CER values computed relative to the EH model so positive values indicate that the time varying predictability models perform better than the EH model. For the scenario with the weights constrained to $[0, 0.99]$ (Panels A-D), the CER values generally increase with the bond maturity. In 11 of 16 cases, the highest CER values are found for the three-factor TVP-SV models. For these models the CER values increase from around 0.24% ($n = 2$) to 0.82% ($n = 3$) and 1.2%-1.6% for the longest bond maturity. To test if the annualized CER values are statistically greater than zero we use a Diebold-Mariano test.[21] With the notable exception of the two-year bond maturity, most of the CER values for the SV and TVP-SV models are significantly higher than those generated by the EH benchmark.

Turning to the case with portfolio weights constrained to lie in the [-2, 3] interval (Panels E-H in Table 5), the CER values generally increase substantially, notably for the two-year bond maturity. For example, for the TVP-SV model with three predictors the CER increases from 0.24% to 2.95% (two-year bond) and from 1.55% to 1.96% (five-year maturity). In general, all models improve strongly when loosening the constraint imposed on the portfolio weights in Panels A-D.

Comparing results across the different specifications, the SV and TVP-SV models generally perform best. Their improvements over the benchmark model are particularly large under the less constrained weights where we find CER gains above 200 basis points per year for the models that include the LN predictor

Figure 8 plots cumulative CER values, computed relative to the EH benchmark, for the three-factor model assuming $\hat{\omega}_t^{(n)} \in [-2, 3]$. These graphs parallel the cumulated sum of squared error difference plots in (31), the key difference being that they show the cumulated risk-adjusted

[21] Specifically, we estimate the regression $u_{t+1.i}^{(n)} - u_{t+1,EH}^{(n)} = \alpha^{(n)} + \epsilon_{t+1}$ where

$$u_{t+1,i}^{(n)} = \frac{1}{1-A}\left[\left(1-\omega_{t,i}^{(n)}\right)\exp\left(\tilde{y}_t\right) + \omega_{t,i}^{(n)}\exp\left(\tilde{y}_t + rx_{t+1}^{(n)}\right)\right]^{1-A},$$

and

$$u_{t+1,EH}^{(n)} = \frac{1}{1-A}\left[\left(1-\omega_{t,EH}^{(n)}\right)\exp\left(\tilde{y}_t\right) + \omega_{t,EH}^{(n)}\exp\left(\tilde{y}_t + rx_{t+1}^{(n)}\right)\right]^{1-A},$$

and test if $\alpha^{(n)}$ equals zero.

gains from using a particular model instead of the EH model. Interestingly, the models' best performance is concentrated from 1990 to 1993 and from 2001 onwards. Across all bond maturities The cumulated CER value at the end of the sample exceeds 50 percent for all models.

The CER values reported in Table 5 ignore transaction costs. However, when we allow for transaction costs we continue to see sizeable gains over the EH benchmark. For example, assuming a one-way transaction cost of 10 basis points, the CER value for the strategy that uses the TVP-SV model to predict bond excess returns is reduced from 2.95% to 1.80% for the two-year bond and from 1.96% to 1.48% for the five-year bond.

We conclude from these results that there is strong statistical and economic evidence that the returns on 2-5 year bonds can be predicted using predictor variables proposed in the literature. Moreover, the best performing models do not assume constant parameters but allow for time varying mean and volatility dynamics.

5.3 Comparison with Other Studies

Our results are very different from those reported by Thornton and Valente (2012). These authors find that statistical evidence of out-of-sample return predictability fails to translate into an ability for investors to use return forecasts in a way that generates higher out-of-sample average utility than forecasts from the EH model. Notably, we find that incorporating time varying parameters and stochastic volatility in many cases improves bond portfolio performance. Instead, Thornton and Valente (2012) find that the Sharpe ratios of their bond portfolios decrease when accounting for such effects through rolling window estimation.

Besides differences in modeling approaches, a reason for such differences is the focus of Thornton and Valente (2012) on 12-month bond returns, whereas we use monthly bond returns. To address the importance of the return horizon, we repeat the out-of-sample analysis using non-overlapping quarterly and annual returns data. Compared with the monthly results, the quarterly and annual R^2_{OoS} values decline somewhat. At the quarterly horizon the univariate FB and LN models, along with the bivariate FB-LN model, continue to perform well across the four bond maturities. The LN and FB-LN models also perform well at the annual horizon, particularly for the bonds with shorter maturities ($n = 2, 3$). The associated CER values continue to be positive and, in most cases, significant at the quarterly horizon, but are substantially smaller at the annual horizon. These findings indicate a fast moving predictable component in bond returns that is missed when using longer return horizons and so help explain the difference between our results and those of Thornton and Valente (2012) and Dewachter et al. (2014).

The setup of Sarno et al. (2016) is closest to that adopted here as they also consider results for one-month returns and still obtain negative economic values from using their time varying bond return forecasts compared with the EH model. Such differences in results reflect (i) different modeling assumptions: Sarno et al. (2016) compute expected excess returns in the context

of an affine term structure model and also do not consider stochastic volatility or time varying parameters; (ii) different predictor variables: Sarno et al. (2016) use latent state variables extracted from their term structure model to predict bond excess returns; and, (iii) different estimation methodologies: Sarno et al. (2016) do not follow the same Bayesian methodology that we use here and thus ignore parameter uncertainty.

Duffee (2013) expresses concerns related to data mining when interpreting results for macro predictors whose effects are not underpinned by theory. A particular concern is that the strong results for the LN factor are sample specific. The sample used by Ludvigson and Ng ends in 2003:12. One way to address this concern is by inspecting the performance of the three-factor model in the subsequent sample, i.e., from 2004:01 to 2011:12. Figure 7 and Figure 8 show that the prediction models continue to generate higher CER values and log-density scores than the EH benchmark after 2003, although Figure 5 suggests weaker R^2 performance for the two- and three-year bond maturities. These results show that the predictive power of the LN factor is not limited to the original sample used to construct this variable.

6 Economic Drivers of Bond Return Predictability

To address the economic sources of our results, this section first analyzes variations in bond return predictability across the economic cycle. Next, we explore whether our bond return forecasts are correlated with drivers of time varying bond risk premia. Finally, we discuss whether unspanned risk factors help explain the predictive power of the LN macro factor.

6.1 Cyclical Variations in Bond Return Predictability

Recent studies such as Rapach et al. (2010), Henkel et al. (2011) and Dangl and Halling (2012) report that predictability of stock returns is concentrated in economic recessions and is largely absent during expansions. This finding is important since it suggests that return predictability is linked to cyclical variations and that time varying risk premia may be important drivers of expected returns.

To see if bond return predictability varies over the economic cycle, we split the data into recession and expansion periods using the NBER recession indicator which equals one in recessions and zero in expansions. Table 6 uses full-sample parameter estimates, but computes R^2 values separately for the recession and expansion samples. We use full-sample information because there are only three recessions in our out-of-sample period, 1990-2011.

Table 6 shows that the R^2 values are generally higher during recessions than in expansions. This finding is consistent with the findings for stock market returns as indicated by the earlier references. Moreover, it is robust across model specifications and predictor variables, the only exception being the univariate FB model for which return predictability actually is stronger

during expansions. Conversely, note that the R^2 values are particularly high in recessions for the TVP models that include the LN variable.

To test if the differences in R^2 values are statistically significant, we conduct a simple bootstrap test that exploits the monotonic relationship between the mean squared prediction error (MSE) of the forecasting model, measured relative to that of the EH model, and the R^2 measure in (30). Specifically, we test the null that the predictive accuracy of a given prediction model (measured relative to the EH benchmark) is the same across recessions and expansions, against the one-sided alternative that the relative MSE is higher in expansions,

$$\begin{aligned} H_0: &\quad E[\underbrace{e^2_{EH,0} - e^2_{i,0}}_{\Delta_0}] = E[\underbrace{e^2_{EH,1} - e^2_{i,1}}_{\Delta_1}] \\ H_1: &\quad E[e^2_{EH,0} - e^2_{i,0}] < E[e^2_{EH,1} - e^2_{i,1}]. \end{aligned} \tag{36}$$

Here e_{EH} and e_i are the forecast errors under the EH and model i, respectively, and the subscript refers to expansions (0) and recessions (1). By computing a particular model's MSE relative to the MSE of the EH model in the same state we control for differences in bond return variances in recessions versus expansions. Our test uses a bootstrap based on the frequency with which $\Delta_0 - \Delta_1$ is smaller than 10,000 counterparts bootstrapped under the null of $\Delta_0 = \Delta_1$.[22]

Outcomes from this test are indicated by stars in the recession columns of Table 6. With the notable exception of the univariate FB model, we find that not only is the fit of the bond return prediction models better in recessions than in expansions, but this difference is highly statistically significant in most cases.

Large differences between bond return predictability in recessions and expansions are also observed in the out-of-sample period 1990-2011. However, in this case we do not have a large enough number of recessions for the test of equal predictive power to have sufficient power to reject the null hypothesis in (36).[23]

We also compute results that split the sample into recessions and expansions using the unemployment gap recession indicator of Stock and Watson (2010).[24] This indicator is computable in real time and so is arguably more relevant than the NBER indicator which gets released with

[22] The p-value for the test is computed as follows: i) impose the null of equal-predictability across states i.e., compute $\hat{\Delta}_0 = \Delta_0 - \hat{\mu}(\Delta_0)$ and $\hat{\Delta}_1 = \Delta_1 - \hat{\mu}(\Delta_1)$; ii) estimate the distribution under the null by using an i.i.d. bootstrap, to generate B bootstrap samples from $\hat{\Delta}_0$ and $\hat{\Delta}_1$ and for each of these compute $J^b = \mu(\hat{\Delta}^b_0) - \mu(\hat{\Delta}^b_1)$; iii) compute p-values as $p_{val} = \frac{1}{B}\sum_{b=1}^{B} 1[J > J^b]$ where $J = \mu(\Delta_0) - \mu(\Delta_1)$ is based on the data.

[23] Engsted et al. (2013) find that bond return predictability is stronger during expansions than during recessions, concluding that return predictability displays opposite patterns in the bond and stock markets. However, they use returns on a 20-year Treasury bond obtained from Ibbotson International. As we have seen, bond return predictability strongly depends on the bond maturity and so this is likely to explain the difference between their results and ours.

[24] This measure is based on the difference between the current unemployment rate and a three-year moving average of past unemployment rates.

a considerable lag. Using this alternative measure of recessions we continue to find that return predictability tends to be stronger in recessions than in expansions.

6.2 Time varying Risk Premia

Asset pricing models such as Campbell and Cochrane (1999) suggest that the Sharpe ratio on risky assets should be higher during recessions due to higher consumption volatility and a lower surplus consumption ratio. To see if this implication is consistent with our models, Table 7 reports Sharpe ratios for the bond portfolios computed separately for recession and expansion periods. Following authors such as Henkel et al. (2011) these results are again based on the full sample to ensure enough observations in recessions. With exception of the univariate FB regressions, the Sharpe ratios are substantially higher during recessions than in expansions.

To further analyze risk-based explanations of bond return predictability, we explore two important sources of bond risk premia. First, some asset pricing models suggest that investors' risk aversion should vary countercyclically, being higher around recessions and lower in expansions. To the extent that our forecasts of bond excess returns reflect time varying risk premia, we should therefore expect a negative correlation between economic growth and bond return forecasts. Second, inflation risk is likely to be an important determinant of bond return dynamics, see, e.g., Wright (2011) and Abrahams et al. (2013) and we should expect to find a positive correlation between inflation uncertainty and bond return forecasts.

Table 8 tests this implication. The table reports the contemporaneous correlations between forecasts of two-year bond excess returns and current real GDP growth (Panel A), inflation (Panel B), real GDP growth uncertainty (Panel C) and inflation uncertainty (Panel D). Real GDP growth is computed as $\Delta \log(GDP_{t+1})$, where GDP_{t+1} is the real gross domestic product, while inflation is computed as $\Delta \log(CPI_{t+1})$, where CPI is the consumer price index for all urban consumers. We measure real GDP growth and inflation uncertainty using the cross-sectional dispersion (the difference between the 75th percentile and the 25th percentile) in real GDP and CPI one-quarter-ahead forecasts, respectively, as reported by the Survey of Professional Forecasters (SPF). An advantage of this measure is that it affords a model-free approach.

The correlations between out-of-sample bond excess return forecasts, on the one hand, and current GDP growth or inflation, on the other, are negative and, in most cases, highly significant. Thus lower economic growth and reduced inflation appear to be associated with expectations of higher bond excess returns.

Turning to the uncertainty measures, we find a strongly positive and, in most cases, highly significant correlation between uncertainty about economic growth and future inflation, on the one hand, and expected bond excess returns on the other. This is consistent with our bond excess return forecasts being driven, at least in part, by time varying inflation risk premia.

Correlations are particularly strong for the models that include the LN macro factor which can be expected to be particularly sensitive to the economic cycle.

6.3 Unspanned Macro Factors

Many studies use only information in the yield curve to predict bond excess returns so our finding that the LN macro factor improves such forecasts may seem puzzling. However, as discussed by Duffee (2013), a possible explanation is that the macro variables are hidden or unspanned risk factors which do not show up in the yield curve because their effect on expected future bond excess returns and expected future short rates work in opposite directions and so tend to cancel out in (5). To see if this possibility holds up, we use SPF survey forecasts of the future 3-month T-bill rate as a proxy for the first term on the right hand side in (5). Consistent with the unspanned risk factor story, Table 9 shows that our forecasts of bond excess returns and survey expectations of future treasury yields are strongly negatively correlated with three-factor (FB-CP-LN) R^2 values ranging from 0.14 to 0.24 for $n = 2$ years and from 0.40 to 0.44 for $n = 5$ years.

The final part of our analysis estimates bond risk premia by fitting an affine term structure model to the cross section of bond yields. In Appendix B we explain how we use the approach of Joslin et al. (2011) and Wright (2011) to fit term structure models with unspanned macro risks to compute bond risk premia. We use the resulting risk premium estimates to regress, for a given maturity n, the corresponding mean bond excess returns, $\overline{rx}_t$, on a constant and the corresponding risk premium estimates $\overline{rp}_t^{(n)}$,

$$\overline{rx}_t^{(n)} = \mu + \beta \overline{rp}_t^{(n)} + u_t, \tag{37}$$

where $\overline{rx}_t^{(n)}$ denotes the predicted bond excess-return and $\overline{rp}_t^{(n)}$ denotes the risk premium estimate. Table 10 reports the estimated coefficient β along with its t-statistics for the FB-CP-LN model. For all specifications, the estimated β coefficient has the right sign (positive) and it is statistically significant for the SV and TVP-SV models fitted to the two shortest bond maturities.

7 Model Combinations

In addition to parameter uncertainty, investors face model uncertainty along with the possibility that the best model may change over time, i.e., model instability. This raises the question whether, in real time, investors could have selected forecasting models that would have generated accurate forecasts. Model uncertainty would not be a concern if all prediction models produced improvements over the EH benchmark. However, as we have seen in the empirical analysis, there is a great deal of heterogeneity across the models' predictive performance. To address this issue,

we turn to model combination. Model combinations form portfolios of individual prediction models with weights reflecting the models' historical performance. The better a model's fit relative to its complexity, the larger its weight. Similar to diversification benefits obtained for asset portfolios, model combination tends to stabilize forecasts relative to forecasts generated by individual return prediction models.

Recent model combination approaches such as Bayesian model averaging and the optimal prediction pool of Geweke and Amisano (2011) allow the weights on individual forecasting models to reflect their predictive accuracy. Such combination schemes can therefore accommodate time variations in the relative performance of different models. This matters if the importance of features such as time varying parameters and stochastic volatility dynamics changes over time.

A final reason for our interest in model combinations is that studies on predictability of stock returns such as Rapach et al. (2010), Dangl and Halling (2012), and Pettenuzzo et al. (2014) find that combinations improve on the average performance of the individual models. This result has only been established for stock returns, however.

To see if it carries over to bond returns, we consider three different combination schemes applied to all possible models obtained by combining the FB, CP and LN predictors, estimated using the linear, SV, TVP and TVP-SV approaches.

7.1 Combination Schemes

We begin by considering the equal-weighted pool (EW) which weighs each of the N models, M_i, equally

$$p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right)=\frac{1}{N}\sum_{i=1}^{N}p\left(rx_{t+1}^{(n)}\middle|\,M_i,\mathcal{D}^t\right), \tag{38}$$

where $\left\{p\left(rx_{t+1}^{(n)}\middle|\,M_i,\mathcal{D}^t\right)\right\}_{i=1}^{N}$ denotes the predictive densities specified in (28) and (29). This approach does not allow the weights on different models to change over time as a result of differences in predictive accuracy.

We also consider Bayesian model averaging (BMA) weights:

$$p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right)=\sum_{i=1}^{N}\Pr\left(M_i|\,\mathcal{D}^t\right)p\left(rx_{t+1}^{(n)}\middle|\,M_i,\mathcal{D}^t\right). \tag{39}$$

Here $\Pr\left(M_i|\,\mathcal{D}^t\right)$ denotes the posterior probability of model i, relative to all models under consideration, computed using information available at time t, $\mathcal{D}^t$. This is given by

$$\Pr\left(M_i|\,\mathcal{D}^t\right)=\frac{\Pr\left(\mathcal{D}^t\middle|\,M_i\right)\Pr\left(M_i\right)}{\sum_{j=1}^{N}\Pr\left(\mathcal{D}^t|\,M_j\right)\Pr\left(M_j\right)}. \tag{40}$$

$\Pr\left(\mathcal{D}^t \middle| M_i\right)$ and $\Pr\left(M_i\right)$ denote the marginal likelihood and prior probability for model i, respectively. We assume that all models are equally likely a priori and so set $\Pr\left(M_i\right) = 1/N$.[25]

A limitation of the BMA approach is that it assumes that the true prediction model is contained in the set of models under consideration. One approach that does not require this assumption is the optimal predictive pool (OW) proposed by Geweke and Amisano (2011). This approach again computes a weighted average of the predictive densities:

$$p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right) = \sum_{i=1}^{N} w_{t,i}^* \times p\left(rx_{t+1}^{(n)} \middle| M_i, \mathcal{D}^t\right). \tag{41}$$

The $(N \times 1)$ vector of model weights $\boldsymbol{w}_t^* = \left[w_{t,1}^*, ..., w_{t,N}^*\right]$ is determined by recursively solving the following maximization problem

$$\boldsymbol{w}_t^* = \arg\max_{\boldsymbol{w}_t} \sum_{\tau=1}^{t-1} \log\left[\sum_{i=1}^{N} w_{ti} \times S_{\tau+1,i}\right], \tag{42}$$

where $S_{\tau+1,i} = \exp\left(LS_{\tau+1,i}\right)$ is the recursively computed log-score for model i at time $\tau+1$, and $\boldsymbol{w}_t^* \in [0,1]^N$. As $t \to \infty$ the weights in (42) minimize the Kullback-Leibler distance between the combined predictive density and the data generating process, see Hall and Mitchell (2007).

By recursively updating the combination weights in (39) and (42), these combination methods accommodate changes in the relative performance of the different models. This is empirically important as we shall see.

7.2 Empirical Findings

Table 11 presents statistical and economic measures of out-of-sample forecasting performance for the three combination schemes. The optimal prediction pool generates R^2_{OoS} values at or above 5% regardless of the bond maturity while the R^2_{OoS} values range between 3 and 5% for the EW combination scheme (Panel A). The BMA combination scheme performs better than the EW combination but worse than the optimal pool. In all cases, the forecast combinations perform better–considerably so in the case of the optimal pool–than what one would expect from simply selecting a model at random.

The predictive likelihood tests shown in Panel B of Table 11 strongly reject the null of equal predictive accuracy relative to the EH model. Finally, the CER values with constrained weights (Panel C) are quite similar for the three combination schemes, rising from about 0.2% for $n = 2$ to 1%–1.5% for $n = 5$. These CER values are similar to those achieved by the best of the individual models reported in Table 5 and, for bond maturities of three years or longer, are

[25] We follow Geweke and Amisano (2010) and compute the marginal likelihoods by cumulating the predictive log scores of each model over time after conditioning on the initial warm-up estimation sample $\Pr\left(\{rx_{\tau+1}^{(n)}\}_{\tau=1}^{t-1} \middle| M_i\right) = \exp\left(\sum_{\tau=\underline{t}}^{t} LS_{\tau,i}\right)$.

significantly higher than those obtained by an EH investor, suggesting that model combination can be used to effectively deal with model uncertainty. The CER values for the case with the [-2,3] weights (Panel D) are very high, falling between 1.43% and 3.27%.

Figure 9 shows the evolution over time in the weights on the individual predictors in the optimal prediction pool. The FB and LN predictors are assigned close to full weights throughout the sample, whereas the CP predictor only gets assigned modest weight during short spells of time.

8 Conclusion

We analyze predictability of excess returns on Treasury bonds with maturities ranging from two through five years. As predictors we use the forward spread variable of Fama and Bliss (1987), the Cochrane and Piazzesi (2005) combination of forward rates, and the Ludvigson and Ng (2009) macro factors. Our analysis allows for time varying regression parameters and stochastic volatility dynamics and accounts for both parameter estimation error and model uncertainty. Using a flexible setup turns out to be important as we find significant statistical and economic gains over the constant coefficient, constant volatility models generally adopted in the existing literature.

Our findings suggest that there is evidence of both statistically and economically significant predictability in bond excess returns. This contrasts with the findings of Thornton and Valente (2012) who conclude that the statistical evidence on bond return predictability fails to translate into economic return predictability. We find that such differences can be attributed to the importance of modeling stochastic volatility and time-varying parameters in the monthly bond excess return series that we study and to rigorously accounting for estimation error.

We link the evidence on return predictability to the economic cycle, finding that the degree of return predictability is significantly higher during recessions. Our bond return forecasts are strongly positively correlated with inflation uncertainty and negatively correlated with economic growth, suggesting that time varying risk premia are an important driver of the results. Consistent with unspanned risk factor models, forecasts of bond excess returns that incorporate information on macro variables are strongly negatively correlated with survey forecasts of future short term yields.

References

Abrahams, M., T. Adrian, R. K. Crump, and E. Moench (2013). Decomposing real and nominal yield curves. Federal Reserve of New York Staff Report No. 570, New York, NY.

Altavilla, C., R. Giacomini, and R. Costantini (2014). Bond returns and market expectations. *Journal of Financial Econometrics 12*(4), 708–729.

Andersen, T. G., T. Bollerslev, P. F. Christoffersen, and F. X. Diebold (2006). Volatility and correlation forecasting. Volume 1 of *Handbook of Economic Forecasting*, pp. 777 – 878. Elsevier.

Andrews, D. W. K. and J. C. Monahan (1992). An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica 60*(4), pp. 953–966.

Ang, A., S. Dong, and M. Piazzesi (2007). No-arbitrage Taylor rules. NBER Working paper 13448.

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.

Bikbov, R. and M. Chernov (2010). No-arbitrage macroeconomic determinants of the yield curve. *Journal of Econometrics 159*(1), 166 – 182.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), pp. 205–251.

Campbell, J. Y. and R. J. Shiller (1991). Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies 58*(3), 495–514.

Campbell, J. Y. and S. B. Thompson (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies 21*(4), 1509–1531.

Carter, C. K. and R. Kohn (1994). On gibbs sampling for state space models. *Biometrika 81*(3), pp. 541–553.

Clark, T. E. and M. McCracken (2011). Testing for unconditional predictive ability. In M. Clements and D. Hendry (Eds.), *Oxford Handbook of Economic Forecasting*. Oxford University Press: Oxford.

Clark, T. E. and F. Ravazzolo (2015). Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *Journal of Applied Econometrics 30*(4), 551–575.

Clark, T. E. and K. D. West (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics 138*(1), 291 – 311.

Cochrane, J. H. and M. Piazzesi (2005). Bond risk premia. *American Economic Review 95*(1), 138–160.

Cogley, T., G. E. Primiceri, and T. J. Sargent (2010). Inflation-gap persistence in the us. *American Economic Journal: Macroeconomics 2*, 43–69.

Cogley, T. and T. J. Sargent (2002). Evolving post-world war ii u.s. inflation dynamics. In B. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001*. Cambridge, U.S.: MIT Press.

Dangl, T. and M. Halling (2012). Predictive regressions with time-varying coefficients. *Journal of Financial Economics 106*(1), 157 – 181.

Del Negro, M. and G. E. Primiceri (2015). Time varying structural vector autoregressions and monetary policy: A corrigendum. *The Review of Economic Studies*, forthcoming.

Dewachter, H., L. Iania, and M. Lyrio (2014). Information in the yield curve: A macro-finance approach. *Journal of Applied Econometrics 29*(1), 42–64.

Diebold, F. X. and R. S. Mariano (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics 13*(3), 253–263.

Duffee, G. (2010). Sharpe ratios in term structure models. Wokring paper, John Hopkins University.

Duffee, G. (2013). Forecasting interest rates. In G. Elliot and A. Timmermann (Eds.), *Handbook of Economic Forecasting, volume 2*, pp. 385 – 426. North Holland.

Duffee, G. R. (2011). Information in (and not in) the term structure. *Review of Financial Studies 24*(9), 2895–2934.

Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika 89*(3), pp. 603–615.

Efron, B. (2010). *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge University Press.

Engsted, T., S. Moller, and M. Sander (2013). Bond return predictability in expansions and recessions. CREATES research paper 2013-13.

Fama, E. F. and R. R. Bliss (1987). The information in long-maturity forward rates. *The American Economic Review 77*(4), pp. 680–692.

Geweke, J. (2001). A note on some limitations of crra utility. *Economics Letters 71*(3), 341 – 345.

Geweke, J. and G. Amisano (2010). Comparing and evaluating bayesian predictive distributions of asset returns. *International Journal of Forecasting 26*(2), 216 – 230.

Geweke, J. and G. Amisano (2011). Optimal prediction pools. *Journal of Econometrics 164*(1), 130 – 141.

Gurkaynak, R. S., B. Sack, and J. H. Wright (2007). The u.s. treasury yield curve: 1961 to the present. *Journal of Monetary Economics 54*(8), 2291 – 2304.

Hall, S. G. and J. Mitchell (2007). Combining density forecasts. *International Journal of Forecasting 23*(1), 1 – 13.

Henkel, S. J., J. S. Martin, and F. Nardari (2011). Time-varying short-horizon predictability. *Journal of Financial Economics 99*(3), 560 – 580.

Johannes, M., A. Korteweg, and N. Polson (2014). Sequential learning, predictive regressions, and optimal portfolio returns. *Journal of Finance 69*(2), 611–644.

Joslin, S., M. Priebsch, and K. J. Singleton (2014). Risk premiums in dynamic term structure models with unspanned macro risks. *The Journal of Finance 69*(3), 1197–1233.

Joslin, S., K. J. Singleton, and H. Zhu (2011). A new perspective on gaussian dynamic term structure models. *Review of Financial Studies 24*(3), 926–970.

Kandel, S. and R. F. Stambaugh (1996). On the predictability of stock returns: An asset-allocation perspective. *The Journal of Finance 51*(2), pp. 385–424.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: Likelihood inference and comparison with arch models. *The Review of Economic Studies 65*(3), 361–393.

Lin, H., C. Wu, and G. Zhou (2016). Forecasting corporate bond returns with a large set of predictors: An iterated combination approach. *Working Paper*.

Ludvigson, S. C. and S. Ng (2009). Macro factors in bond risk premia. *Review of Financial Studies 22*(12), 5027–5067.

Nelson, C. R. and A. F. Siegel (1987). Parsimonious modeling of yield curves. *The Journal of Business 60*(4), pp. 473–489.

Pettenuzzo, D., A. Timmermann, and R. Valkanov (2014). Forecasting stock returns under economic constraints. *Journal of Financial Economics 114*(3), 517–553.

Rapach, D. E., J. K. Strauss, and G. Zhou (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies 23*(2), 821–862.

Sarno, L., P. Schneider, and C. Wagner (2016). The economic value of predicting bond risk premia. *Journal of Empirical Finance 37*, 247 – 267.

Sims, C. A. and T. Zha (2006). Were there regime switches in u.s. monetary policy? *American Economic Review 96*(1), 54–81.

Stock, J. H. and M. W. Watson (1999). Forecasting inflation. *Journal of Monetary Economics 44*(2), 293 – 335.

Stock, J. H. and M. W. Watson (2006). Forecasting with many predictors. Volume 1 of *Handbook of Economic Forecasting*, pp. 515 – 554. Elsevier.

Stock, J. H. and M. W. Watson (2010). Modeling inflation after the crisis. Presented at the FRB Kansas city symposium on macroeconomic policy: post-crisis and risks ahead, Jackson Hole, Wyoming.

Svensson, L. E. O. (1994). Estimating and interpreting forward interest rates: Sweden 1992-1994. IMF Working Paper No. 94/114.

Thornton, D. L. and G. Valente (2012). Out-of-sample predictions of bond excess returns and forward rates: An asset allocation perspective. *Review of Financial Studies 25*(10), 3141–3168.

Wei, M. and J. H. Wright (2013). Reverse regressions and long-horizon forecasting. *Journal of Applied Econometrics 28*(3), 353–371.

Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies 21*(4), 1455–1508.

Wright, J. H. (2011). Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review 101*(4), 1514–34.

Appendix A Bayesian estimation and predictions

This appendix explains how we obtain parameter estimates for the models described in Section 3 and shows how we use these to generate predictive densities for bond excess returns. We begin by discussing the linear regression model in (17), then turn to the SV model in (21)-(22), the TVP model in (23)-(24), and the general TVP-SV model in (25)-(27).

A.1 Constant coefficient, constant volatility model

The goal for the simple linear regression model is to obtain draws from the joint posterior distribution $p\left(\mu, \boldsymbol{\beta}, \sigma_{\varepsilon}^{-2} \middle| \mathcal{D}^t\right)$, where D^t denotes all information available up to time t. Combining the priors in (18)-(20) with the likelihood function yields the following posteriors:

$$\left[\begin{array}{c} \mu \\ \boldsymbol{\beta} \end{array}\right] \middle| \sigma_{\varepsilon}^{-2}, \mathcal{D}^t \sim \mathcal{N}\left(\overline{\boldsymbol{b}}, \overline{\boldsymbol{V}}\right), \tag{A-1}$$

and

$$\sigma_{\varepsilon}^{-2} \middle| \mu, \boldsymbol{\beta}, \mathcal{D}^t \sim \mathcal{G}\left(\overline{s}^{-2}, \overline{v}\right), \tag{A-2}$$

where

$$\begin{aligned} \overline{\boldsymbol{V}} &= \left[\underline{\boldsymbol{V}}^{-1} + \sigma_{\varepsilon}^{-2} \sum_{\tau=1}^{t-1} \boldsymbol{x}_{\tau}^{(n)} \boldsymbol{x}_{\tau}^{(n)\prime}\right]^{-1}, \\ \overline{\boldsymbol{b}} &= \overline{\boldsymbol{V}}\left[\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{b}} + \sigma_{\varepsilon}^{-2} \sum_{\tau=1}^{t-1} \boldsymbol{x}_{\tau}^{(n)} rx_{\tau+1}^{(n)}\right], \\ \overline{v} &= (1+\underline{v}_0)(t-1). \end{aligned} \tag{A-3}$$

and

$$\overline{s}^2 = \frac{\sum_{\tau=1}^{t-1}\left(rx_{\tau+1}^{(n)} - \mu - \boldsymbol{\beta}'\boldsymbol{x}_{\tau}^{(n)}\right)^2 + \left(\left(s_{rx,t}^{(n)}\right)^2 \times \underline{v}_0 (t-1)\right)}{\overline{v}}. \tag{A-4}$$

Gibbs sampling can be used to iterate back and forth between (A-1) and (A-2), yielding a series of draws for the parameter vector $\left(\mu, \boldsymbol{\beta}, \sigma_{\varepsilon}^{-2}\right)$. Draws from the predictive density $p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right)$ can then be obtained by noting that

$$p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right) = \int p\left(rx_{t+1}^{(n)} \middle| \mu, \boldsymbol{\beta}, \sigma_{\varepsilon}^{-2}, \mathcal{D}^t\right) p\left(\mu, \boldsymbol{\beta}, \sigma_{\varepsilon}^{-2} \middle| \mathcal{D}^t\right) d\mu d\boldsymbol{\beta} d\sigma_{\varepsilon}^{-2}. \tag{A-5}$$

A.2 Stochastic Volatility model[26]

The SV model requires specifying a joint prior for the sequence of log return volatilities, h^t, the parameters λ_0 and λ_1, and the error precision, σ_{ξ}^{-2}. Writing $p\left(h^t, \lambda_0, \lambda_1, \sigma_{\xi}^{-2}\right) =$

[26] See Pettenuzzo et al. (2014) for a description of a similar algorithm where the priors are modified to impose economic constraints on the model parameters.

$p\left(h^t\middle|\,\lambda_0,\lambda_1,\sigma_\xi^{-2}\right)p\left(\lambda_0,\lambda_1\right)p\left(\sigma_\xi^{-2}\right)$, it follows from (22) that

$$p\left(h^t\middle|\,\lambda_0,\lambda_1,\sigma_\xi^{-2}\right)=\prod_{\tau=1}^{t-1}p\left(h_{\tau+1}\middle|\,h_\tau,\lambda_0,\lambda_1,\sigma_\xi^{-2}\right)p\left(h_1\right), \tag{A-6}$$

with $h_{\tau+1}|\,h_\tau,\lambda_0,\lambda_1,\sigma_\xi^{-2}\sim\mathcal{N}\left(\lambda_0+\lambda_1 h_\tau,\sigma_\xi^2\right)$. Thus, to complete the prior elicitation for $p\left(h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\right)$, we only need to specify priors for h_1, the initial log volatility, λ_0, λ_1, and σ_ξ^{-2}. We choose these from the normal-gamma family as follows:

$$h_1\sim\mathcal{N}\left(\ln\left(s_{rx,t}^{(n)}\right),\underline{k}_h\right), \tag{A-7}$$

$$\begin{bmatrix}\lambda_0\\ \lambda_1\end{bmatrix}\sim\mathcal{N}\left(\begin{bmatrix}\underline{m}_{\lambda_0}\\ \underline{m}_{\lambda_1}\end{bmatrix},\begin{bmatrix}\underline{V}_{\lambda_0} & 0\\ 0 & \underline{V}_{\lambda_1}\end{bmatrix}\right),\quad \lambda_1\in(-1,1), \tag{A-8}$$

and

$$\sigma_\xi^{-2}\sim\mathcal{G}\left(1/\underline{k}_\xi,\underline{\upsilon}_\xi\left(t-1\right)\right). \tag{A-9}$$

We set $\underline{k}_\xi=0.01$ and set the remaining hyperparameters in (A-7) and (A-9) at $\underline{k}_h=10$ and $\underline{\upsilon}_\xi=1$ to imply uninformative priors, thus allowing the data to determine the degree of time variation in the return volatility. Following Clark and Ravazzolo (2015) we set the hyperparameters to $\underline{m}_{\lambda_0}=0$, $\underline{m}_{\lambda_1}=0.9$, $\underline{V}_{\lambda_0}=0.25$, and $\underline{V}_{\lambda_0}=1.0e^{-4}$. This corresponds to setting the prior mean and standard deviation of the intercept to 0 and 0.5, respectively, and represents uninformative priors on the intercept of the log volatility specification and a prior mean of the AR(1) coefficient, λ_1, of 0.9 with a standard deviation of 0.01. This is a more informative prior that matches persistent dynamics in the log volatility process.

To obtain draws from the joint posterior distribution $p\left(\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\middle|\,\mathcal{D}^t\right)$ under the SV model, we use the Gibbs sampler to draw recursively from the following four conditional distributions:

1. $p\left(h^t\middle|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$.
2. $p\left(\mu,\boldsymbol{\beta}\middle|\,h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$.
3. $p\left(\lambda_0,\lambda_1\middle|\,\mu,\boldsymbol{\beta},h^t,\sigma_\xi^{-2},\mathcal{D}^t\right)$.
4. $p\left(\sigma_\xi^{-2}\middle|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\mathcal{D}^t\right)$.

We simulate from each of these blocks as follows. Starting with $p\left(h^t\middle|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$, we employ the algorithm of Kim et al. (1998).[27] Define $rx_{\tau+1}^{(n)*}=rx_{\tau+1}^{(n)}-\mu-\beta'\boldsymbol{x}_\tau^{(n)}$ and note that $r_{\tau+1}^{(n)*}$ is observable conditional on μ,β. Next, rewrite (21) as

$$rx_{\tau+1}^{(n)*}=\exp\left(h_{\tau+1}\right)u_{\tau+1}. \tag{A-10}$$

[27] We apply the correction to the ordering of steps detailed in Del Negro and Primiceri (2015).

Squaring and taking logs on both sides of (A-10) yields a new state space system that replaces (21)-(22) with

$$rx_{\tau+1}^{(n)**} = 2h_{\tau+1} + u_{\tau+1}^{**}, \tag{A-11}$$

$$h_{\tau+1} = \lambda_0 + \lambda_1 h_\tau + \xi_{\tau+1}, \tag{A-12}$$

where $rx_{\tau+1}^{(n)**} = \ln\left[\left(rx_{\tau+1}^{(n)*}\right)^2\right]$, and $u_{\tau+1}^{**} = \ln\left(u_{\tau+1}^2\right)$, with u_τ^{**} independent of ξ_s for all τ and s. Since $u_{t+1}^{**} \sim \ln\left(\chi_1^2\right)$, we cannot resort to standard Kalman recursions and simulation algorithms such as those in Carter and Kohn (1994) or Durbin and Koopman (2002). To get around this problem, Kim et al. (1998) employ a data augmentation approach and introduce a new state variable $s_{\tau+1}$, $\tau = 1, .., t-1$, turning their focus to drawing from $p\left(h^t \middle| \mu, \boldsymbol{\beta}, \lambda_0, \lambda_1, \sigma_\xi^{-2}, s^t, \mathcal{D}^t\right)$ instead of $p\left(h^t \middle| \mu, \boldsymbol{\beta}, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right)$, where $s^t = \{s_2, ..., s_t\}$ denotes the history up to time t of the new state variable s.

The introduction of the state variable $s_{\tau+1}$ allows us to rewrite the linear non-Gaussian state space representation in (A-11)-(A-12) as a linear Gaussian state space model, making use of the following approximation,

$$u_{\tau+1}^{**} \approx \sum_{j=1}^{7} q_j \mathcal{N}\left(m_j - 1.2704, v_j^2\right), \tag{A-13}$$

where m_j, v_j^2, and q_j, $j = 1, 2, ..., 7$, are constants specified in Kim et al. (1998) and thus need not be estimated. In turn, (A-13) implies

$$u_{\tau+1}^{**} \middle| \, s_{\tau+1} = j \sim \mathcal{N}\left(m_j - 1.2704, v_j^2\right), \tag{A-14}$$

where each state has probability

$$\Pr\left(s_{\tau+1} = j\right) = q_j. \tag{A-15}$$

Draws for the sequence of states s^t can be easily obtained, noting that each of its elements can be independently drawn from the discrete density defined by

$$\Pr\left(s_{\tau+1} = j \middle| \, \mu, \boldsymbol{\beta}, \lambda_0, \lambda_1, \sigma_\xi^{-2}, h^t, \mathcal{D}^t\right) = \frac{q_i f_{\mathcal{N}}\left(rx_{\tau+1}^{(n)**} \middle| \, 2h_{\tau+1} + m_j - 1.2704, v_j^2\right)}{\sum_{l=1}^{7} q_l f_{\mathcal{N}}\left(rx_{\tau+1}^{(n)**} \middle| \, 2h_{\tau+1} + m_l - 1.2704, v_l^2\right)}. \tag{A-16}$$

for $\tau = 1, ..., t-1$ and $j = 1, ..., 7$, and where $f_{\mathcal{N}}$ denotes the kernel of a normal density. Next, conditional on s^t, we can rewrite the nonlinear state space system as follows:

$$\begin{aligned} rx_{\tau+1}^{(n)**} &= 2h_{\tau+1} + e_{\tau+1}, \\ h_{\tau+1} &= \lambda_0 + \lambda_1 h_\tau + \xi_{\tau+1}, \end{aligned} \tag{A-17}$$

where $e_{\tau+1} \sim \mathcal{N}\left(m_j - 1.2704, v_j^2\right)$ with probability $\Pr\left(s_{\tau+1} = j \middle| \, \mu, \boldsymbol{\beta}, \lambda_0, \lambda_1, \sigma_\xi^{-2}, h^t, \mathcal{D}^t\right)$. For this linear Gaussian state space system, we can use the algorithm of Carter and Kohn (1994) to draw the whole sequence of stochastic volatilities, h^t.

Moving on to $p\left(\mu,\boldsymbol{\beta}|\, h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$, conditional on h^t it is straightforward to draw μ and $\boldsymbol{\beta}$ and apply standard results. Specifically,

$$\left[\begin{array}{c}\mu\\ \boldsymbol{\beta}\end{array}\right]\Big|\, h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\sim\mathcal{N}\left(\overline{\boldsymbol{b}},\overline{\boldsymbol{V}}\right), \tag{A-18}$$

with

$$\begin{aligned}\overline{\boldsymbol{V}} &= \left\{\underline{\boldsymbol{V}}^{-1}+\sum_{\tau=1}^{t-1}\frac{1}{\exp\left(h_{\tau+1}\right)^2}\boldsymbol{x}_\tau^{(n)}\boldsymbol{x}_\tau^{(n)\prime}\right\}^{-1},\\ \overline{\boldsymbol{b}} &= \overline{\boldsymbol{V}}\left\{\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{b}}+\sum_{\tau=1}^{t-1}\frac{1}{\exp\left(h_{\tau+1}\right)^2}\boldsymbol{x}_\tau^{(n)}rx_{\tau+1}^{(n)}\right\}.\end{aligned}$$

Next, the distribution $p\left(\lambda_0,\lambda_1|\,\mu,\boldsymbol{\beta},h^t,\sigma_\xi^{-2},\mathcal{D}^t\right)$ takes the form

$$\lambda_0,\lambda_1|\,\mu,\boldsymbol{\beta},h^t,\sigma_\xi^{-2},\mathcal{D}^t\sim\mathcal{N}\left(\left[\begin{array}{c}\overline{m}_{\lambda_0}\\ \overline{m}_{\lambda 1}\end{array}\right],\overline{\mathbf{V}}_\lambda\right)\times\lambda_1\in(-1,1),$$

where

$$\overline{\mathbf{V}}_\lambda=\left\{\left[\begin{array}{cc}\underline{V}_{\lambda_0}^{-1}&0\\ 0&\underline{V}_{\lambda_1}^{-1}\end{array}\right]+\sigma_\xi^{-2}\sum_{\tau=1}^{t-1}\left[\begin{array}{c}1\\ h_\tau\end{array}\right][1,h_\tau]\right\}^{-1}, \tag{A-19}$$

and

$$\left[\begin{array}{c}\overline{m}_{\lambda_0}\\ \overline{m}_{\lambda 1}\end{array}\right]=\overline{\mathbf{V}}_\lambda\left\{\left[\begin{array}{cc}\underline{V}_{\lambda_0}^{-1}&0\\ 0&\underline{V}_{\lambda_1}^{-1}\end{array}\right]\left[\begin{array}{c}\underline{m}_{\lambda_0}\\ \underline{m}_{\lambda_1}\end{array}\right]+\sigma_\xi^{-2}\sum_{\tau=1}^{t-1}\left[\begin{array}{c}1\\ h_\tau\end{array}\right]h_{\tau+1}\right\}. \tag{A-20}$$

Finally, the posterior distribution for $p\left(\sigma_\xi^{-2}\Big|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\mathcal{D}^t\right)$ is readily available using

$$\sigma_\xi^{-2}\Big|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\mathcal{D}^t\sim\mathcal{G}\left(\left[\frac{\underline{k}_\xi\underline{v}_\xi(t-1)+\sum_{\tau=1}^{t-1}\left(h_{\tau+1}-\lambda_0-\lambda_1h_\tau\right)^2}{\left(1+\underline{v}_\xi\right)(t-1)}\right]^{-1},\left(1+\underline{v}_\xi\right)(t-1)\right). \tag{A-21}$$

Draws from the predictive density $p\left(rx_{t+1}^{(n)}\Big|\,\mathcal{D}^t\right)$ can be obtained by noting that

$$\begin{aligned}p\left(rx_{t+1}^{(n)}\Big|\,\mathcal{D}^t\right) &= \int p\left(rx_{t+1}^{(n)}\Big|\,h_{t+1},\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)\\ &\quad\times p\left(h_{t+1}|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)\\ &\quad\times p\left(\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\Big|\,\mathcal{D}^t\right)d\mu d\boldsymbol{\beta}dh^{t+1}d\lambda_0d\lambda_1d\sigma_\xi^{-2}.\end{aligned} \tag{A-22}$$

The first term in the integral above, $p\left(rx_{t+1}^{(n)}\Big|\,h_{t+1},\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$, represents the period $t+1$ predictive density of bond excess returns, treating model parameters as if they were known with certainty, and so is straightforward to calculate. The second term in the integral, $p\left(h_{t+1}|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$, reflects how period $t+1$ volatility may drift away from h_t over

time. Finally, the last term in the integral, $p\left(\mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2} \middle| \mathcal{D}^t\right)$, measures parameter uncertainty in the sample.

To obtain draws for $p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right)$, we proceed in three steps:

1. Simulate from $p\left(\mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2} \middle| \mathcal{D}^t\right)$: draws from $p\left(\mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2} \middle| \mathcal{D}^t\right)$ are obtained from the Gibbs sampling algorithm described above.

2. Simulate from $p\left(h_{t+1} | \mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right)$: having processed data up to time t, the next step is to simulate the future volatility, h_{t+1}. For a given h_t and σ_ξ^{-2}, note that μ and β and the history of volatilities up to t become redundant, i.e., $p\left(h_{t+1} | \mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right) = p\left(h_{t+1} | h_t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right)$. Note also that (22) along with the distributional assumptions made on $\xi_{\tau+1}$ imply that

$$h_{t+1} | h_t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t \sim \mathcal{N}\left(\lambda_0 + \lambda_1 h_t, \sigma_\xi^2\right). \tag{A-23}$$

3. Simulate from $p\left(rx_{t+1}^{(n)} \middle| h_{t+1}, \mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right)$: For a given h_{t+1}, μ, and β, note that h^t, λ_0, λ_1, and σ_ξ^{-2} become redundant, i.e., $p\left(rx_{t+1}^{(n)} \middle| h_{t+1}, \mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right) = p\left(rx_{t+1}^{(n)} \middle| h_{t+1}, \mu, \boldsymbol{\beta}, \mathcal{D}^t\right)$. Then use the fact that

$$rx_{t+1}^{(n)} \middle| h_{t+1}, \mu, \boldsymbol{\beta}, \mathcal{D}^t \sim \mathcal{N}\left(\mu + \boldsymbol{\beta}' \boldsymbol{x}_\tau^{(n)}, \exp\left(h_{t+1}\right)\right). \tag{A-24}$$

A.3 Time varying Parameter Model

In addition to specifying prior distributions and hyperparameters for $[\mu, \boldsymbol{\beta}]'$ and σ_ε^2, the TVP model in (23)-(24) requires eliciting a joint prior for the sequence of time varying parameters $\boldsymbol{\theta}^t = \{\boldsymbol{\theta}_2, ..., \boldsymbol{\theta}_t\}$, the parameter vector $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$, and the variance covariance matrix $\boldsymbol{Q}$. For $[\mu, \boldsymbol{\beta}]'$ and σ_ε^2, we follow the same prior choices made for the linear model:

$$\begin{bmatrix} \mu \\ \boldsymbol{\beta} \end{bmatrix} \sim \mathcal{N}\left(\underline{\boldsymbol{b}}, \underline{\boldsymbol{V}}\right), \tag{A-25}$$

and

$$\sigma_\varepsilon^{-2} \sim \mathcal{G}\left(\left(s_{rx,t}^{(n)}\right)^{-2}, \underline{v}_0 (t-1)\right). \tag{A-26}$$

Turning to $\boldsymbol{\theta}^t$, $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$, and $\boldsymbol{Q}$, we first write $p\left(\boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right) = p\left(\boldsymbol{\theta}^t \middle| \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right) p\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}\right) p\left(\boldsymbol{Q}\right)$, and note that (24) along with the assumption that $\boldsymbol{\theta}_1 = \mathbf{0}$ implies

$$p\left(\boldsymbol{\theta}^t \middle| \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right) = \prod_{\tau=1}^{t-1} p\left(\boldsymbol{\theta}_{\tau+1} | \boldsymbol{\theta}_\tau, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right), \tag{A-27}$$

with $\boldsymbol{\theta}_{\tau+1} | \boldsymbol{\theta}_\tau, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q} \sim \mathcal{N}\left(diag\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}\right) \boldsymbol{\theta}_\tau, \boldsymbol{Q}\right)$. Thus, to complete the prior elicitation for $p\left(\boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right)$ we need to specify priors for $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$ and $\boldsymbol{Q}$.

We choose an Inverted Wishart distribution for $\mathbf{Q}$:

$$\mathbf{Q} \sim \mathcal{IW}\left(\underline{\mathbf{Q}}, \underline{v}_{\mathbf{Q}}\left(t_0 - 1\right)\right), \tag{A-28}$$

with

$$\underline{\mathbf{Q}} = \underline{k}_{\mathbf{Q}} \underline{v}_{\mathbf{Q}}\left(t_0 - 1\right) \underline{\mathbf{V}}. \tag{A-29}$$

$\underline{k}_{\mathbf{Q}}$ controls the degree of variation in the time-varying regression coefficients $\boldsymbol{\theta}_\tau$, with larger values of $\underline{k}_{\mathbf{Q}}$ implying greater variation in $\boldsymbol{\theta}_\tau$. Our analysis sets $\underline{k}_{\mathbf{Q}} = \left(\underline{\psi}/100\right)^2$ and $\underline{v}_{\mathbf{Q}} = 10$. These are more informative priors than the earlier choices and limit the changes to the regression coefficients to be $\underline{\psi}/100$ on average.

We specify the elements of $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$ to be a priori independent of each other with generic element $\gamma^i_{\boldsymbol{\theta}}$

$$\gamma^i_{\boldsymbol{\theta}} \sim \mathcal{N}\left(\underline{m}_{\gamma_{\boldsymbol{\theta}}}, \underline{V}_{\gamma_{\boldsymbol{\theta}}}\right), \quad \gamma^i_{\boldsymbol{\theta}} \in (-1, 1), \quad i = 1, ..., k. \tag{A-30}$$

where $\underline{m}_{\gamma_{\boldsymbol{\theta}}} = 0.8$, and $\underline{V}_{\gamma_{\boldsymbol{\theta}}} = 1.0e^{-6}$, implying relatively high autocorrelations.

To obtain draws from the joint posterior distribution $p\left(\mu, \boldsymbol{\beta}, \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q} \middle| \mathcal{D}^t\right)$ under the TVP model we use the Gibbs sampler to draw recursively from the following conditional distributions:

1. $p\left(\boldsymbol{\theta}^t \middle| \mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t\right)$.
2. $p\left(\mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2} \middle| \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t\right)$.
3. $p\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}} \middle| \mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \boldsymbol{\theta}^t, \boldsymbol{Q}, \mathcal{D}^t\right)$.
4. $p\left(\boldsymbol{Q} \middle| \mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \mathcal{D}^t\right)$.

We simulate from each of these blocks as follows. Starting with θ^t, we focus on $p\left(\boldsymbol{\theta}^t \middle| \mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t\right)$. Define $rx^{(n)*}_{\tau+1} = rx^{(n)}_{\tau+1} - \mu - \beta' \boldsymbol{x}^{(n)}_\tau$ and rewrite (23) as follows:

$$rx^{(n)*}_{\tau+1} = \mu_\tau - \boldsymbol{\beta}'_\tau \boldsymbol{x}^{(n)}_\tau + \varepsilon_{\tau+1} \tag{A-31}$$

Knowledge of μ and $\boldsymbol{\beta}$ makes $rx^{(n)*}_{\tau+1}$ observable, and reduces (23) to the measurement equation of a standard linear Gaussian state space model with homoskedastic errors. Thus, the sequence of time varying parameters $\boldsymbol{\theta}^t$ can be drawn from (A-31) using the algorithm of Carter and Kohn (1994).

Moving on to $p\left(\mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2} \middle| \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t\right)$, conditional on $\boldsymbol{\theta}^t$ it is straightforward to draw $\mu, \boldsymbol{\beta}$, and σ_ε^{-2} by applying standard results. Specifically,

$$\begin{bmatrix} \mu \\ \boldsymbol{\beta} \end{bmatrix} \Bigg| \sigma_\varepsilon^{-2}, \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t \sim \mathcal{N}\left(\overline{\boldsymbol{b}}, \overline{\boldsymbol{V}}\right), \tag{A-32}$$

and

$$\sigma_\varepsilon^{-2} \middle| \mu, \boldsymbol{\beta}, \boldsymbol{\theta}^t, \boldsymbol{Q}, \mathcal{D}^t \sim \mathcal{G}\left(\overline{s}^{-2}, \overline{v}\right), \tag{A-33}$$

where

$$\begin{aligned}\overline{\boldsymbol{V}} &= \left[\underline{\boldsymbol{V}}^{-1}+\sigma_{\varepsilon}^{-2}\sum_{\tau=1}^{t-1}\boldsymbol{x}_{\tau}^{(n)}\boldsymbol{x}_{\tau}^{(n)\prime}\right]^{-1}, \\ \overline{\boldsymbol{b}} &= \overline{\boldsymbol{V}}\left[\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{b}}+\sigma_{\varepsilon}^{-2}\sum_{\tau=1}^{t-1}\boldsymbol{x}_{\tau}^{(n)}\left(rx_{\tau+1}^{(n)}-\mu_{\tau}-\boldsymbol{\beta}_{\tau}'\boldsymbol{x}_{\tau}^{(n)}\right)\right], \end{aligned} \tag{A-34}$$

$$\overline{s}^{2}=\frac{\sum_{\tau=1}^{t-1}\left(rx_{\tau+1}^{(n)*}-\mu_{\tau}-\boldsymbol{\beta}_{\tau}'\boldsymbol{x}_{\tau}^{(n)}\right)^{2}+\left(\left(s_{rx,t}^{(n)}\right)^{2}\times\underline{v}_{0}\left(t-1\right)\right)}{\overline{v}}, \tag{A-35}$$

and $\overline{v}=\left(1+\underline{v}_{0}\right)\left(t-1\right)$.

Next, obtaining draws from $p\left(\boldsymbol{\gamma_{\theta}}|\,\mu,\boldsymbol{\beta},\sigma_{\varepsilon}^{-2},\boldsymbol{\theta}^{t},\boldsymbol{Q},\mathcal{D}^{t}\right)$ is straightforward. The $i-$th element $\gamma_{\boldsymbol{\theta}}^{i}$ is drawn from the following distribution

$$\gamma_{\boldsymbol{\theta}}^{i}|\,\mu,\boldsymbol{\beta},\sigma_{\varepsilon}^{-2},\boldsymbol{\theta}^{t},\boldsymbol{Q},\mathcal{D}^{t}\sim\mathcal{N}\left(\overline{m}_{\gamma_{\boldsymbol{\theta}}}^{i},\overline{V}_{\gamma_{\boldsymbol{\theta}}}^{i}\right)\times\gamma_{\boldsymbol{\theta}}^{i}\in(-1,1) \tag{A-36}$$

where

$$\begin{aligned}\overline{V}_{\gamma_{\boldsymbol{\theta}}}^{i} &= \left[\underline{V}_{\gamma_{\boldsymbol{\theta}}}^{-1}+\mathbf{Q}^{ii}\sum_{\tau=1}^{t-1}\left(\theta_{\tau}^{i}\right)^{2}\right]^{-1}, \\ \overline{m}_{\gamma_{\boldsymbol{\theta}}}^{i} &= \overline{V}_{\gamma_{\boldsymbol{\theta}}}^{i}\left[\underline{V}_{\gamma_{\boldsymbol{\theta}}}^{-1}\underline{m}_{\gamma_{\boldsymbol{\theta}}}+\mathbf{Q}^{ii}\sum_{\tau=1}^{t-1}\theta_{\tau}^{i}\theta_{\tau+1}^{i}\right], \end{aligned} \tag{A-37}$$

and $\mathbf{Q}^{ii}$ is the $i-$th diagonal element of $\mathbf{Q}^{-1}$.

As for $p\left(\boldsymbol{Q}|\,\mu,\boldsymbol{\beta},\sigma_{\varepsilon}^{-2},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\mathcal{D}^{t}\right)$, we have that

$$\boldsymbol{Q}|\,\mu,\boldsymbol{\beta},\sigma_{\varepsilon}^{-2},\boldsymbol{\theta}^{t},\mathcal{D}^{t}\sim\mathcal{IW}\left(\overline{\mathbf{Q}},\overline{v}_{\mathbf{Q}}\right), \tag{A-38}$$

where

$$\overline{\mathbf{Q}}=\underline{\mathbf{Q}}+\sum_{\tau=1}^{t-1}\left(\boldsymbol{\theta}_{\tau+1}-diag\left(\boldsymbol{\gamma_{\theta}}\right)\boldsymbol{\theta}_{\tau}\right)\left(\boldsymbol{\theta}_{\tau+1}-diag\left(\boldsymbol{\gamma_{\theta}}\right)\boldsymbol{\theta}_{\tau}\right)'. \tag{A-39}$$

and $\overline{v}_{\mathbf{Q}}=\left(1+\underline{v}_{\mathbf{Q}}\right)\left(t-1\right)$.

Finally, draws from the predictive density $p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^{t}\right)$ can be obtained by noting than

$$\begin{aligned}p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^{t}\right) &= \int p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\boldsymbol{Q},\sigma_{\varepsilon}^{-2},\mathcal{D}^{t}\right) \\ &\quad\times p\left(\boldsymbol{\theta}_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\boldsymbol{Q},\sigma_{\varepsilon}^{-2},\mathcal{D}^{t}\right) \\ &\quad\times p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\boldsymbol{Q},\sigma_{\varepsilon}^{-2}\middle|\,\mathcal{D}^{t}\right)d\mu d\boldsymbol{\beta}d\boldsymbol{\theta}^{t+1}d\boldsymbol{\gamma_{\theta}}d\boldsymbol{Q}d\sigma_{\varepsilon}^{-2}. \end{aligned} \tag{A-40}$$

The first term in the integral above, $p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\boldsymbol{Q},\sigma_{\varepsilon}^{-2},\mathcal{D}^{t}\right)$, represents the period $t+1$ predictive density of bond excess returns, treating model parameters as if they were

known with certainty, and so is straightforward to calculate. The second term in the integral, $p\left(\boldsymbol{\theta}_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right)$, reflects that the regression parameters may drift away from $\boldsymbol{\theta}_t$ over time. Finally, the last term in the integral, $p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2}\middle|\,\mathcal{D}^t\right)$, measures parameter uncertainty.

To obtain draws for $p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right)$, we proceed in three steps:

1. Simulate from $p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2}\middle|\,\mathcal{D}^t\right)$: draws from $p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2}\middle|\,\mathcal{D}^t\right)$ are obtained from the Gibbs sampling algorithm described above;

2. Simulate from $p\left(\boldsymbol{\theta}_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right)$: For a given θ_t and Q, note that μ, β, σ_ε^{-2}, and the history of regression parameters up to t become redundant, i.e., $p\left(\boldsymbol{\theta}_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right) = p\left(\boldsymbol{\theta}_{t+1}|\,\boldsymbol{\theta}_t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\mathcal{D}^t\right)$. Note also that (24), along with the distributional assumptions made with regards to $\eta_{\tau+1}$, imply that

$$\boldsymbol{\theta}_{t+1}|\,\boldsymbol{\theta}_t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\mathcal{D}^t \sim \mathcal{N}\left(diag\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}\right)\boldsymbol{\theta}_t,\boldsymbol{Q}\right). \tag{A-41}$$

3. Simulate from $p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right)$: For a given $\boldsymbol{\theta}_{t+1}$, μ, $\boldsymbol{\beta}$, and σ_ε^{-2}, $\boldsymbol{\theta}^t$, $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$, and $\boldsymbol{Q}$ become redundant so $p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right) = p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\sigma_\varepsilon^{-2},\mathcal{D}^t\right)$. Then use the fact that

$$rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\sigma_\varepsilon^{-2},\mathcal{D}^t \sim \mathcal{N}\left(\left(\mu+\mu_t\right)+\left(\boldsymbol{\beta}+\boldsymbol{\beta}_t\right)'\boldsymbol{x}_\tau^{(n)},\sigma_\varepsilon^2\right). \tag{A-42}$$

A.4 Time varying Parameter, Stochastic Volatility Model

Our priors for the TVP-SV model combine the earlier choices for the TVP and SV models, i.e., (A-25) and (A-26) for the regression parameters, (A-7) and (A-9) for the SV component, and (A-28) and (A-29) for the TVP component.

To obtain draws from the joint posterior distribution $p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\middle|\,\mathcal{D}^t\right)$ under the TVP-SV model, we use the Gibbs sampler to draw recursively from the following seven conditional distributions:

1. $p\left(\boldsymbol{\theta}^t\middle|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\varepsilon^{-2},\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\mathcal{D}^t\right)$.
2. $p\left(\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\varepsilon^{-2}\middle|\,\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\mathcal{D}^t\right)$.
3. $p\left(h^t\middle|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$.
4. $p\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}|\,\mu,\boldsymbol{\beta},\sigma_\varepsilon^{-2},\boldsymbol{\theta}^t,\boldsymbol{Q},\mathcal{D}^t\right)$.
5. $p\left(\boldsymbol{Q}|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\varepsilon^{-2},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\mathcal{D}^t\right)$.
6. $p\left(\lambda_0,\lambda_1|\,\mu,\boldsymbol{\beta},h^t,\sigma_\xi^{-2},\mathcal{D}^t\right)$.

7. $p\left(\sigma_\xi^{-2}\middle|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\mathcal{D}^t\right)$.

With minor modifications, these steps are similar to the steps described in the TVP and SV sections above. Draws from the predictive density $p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right)$ can be obtained from

$$\begin{aligned} p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right) &= \int p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},h_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right) \\ &\quad\times p\left(\boldsymbol{\theta}_{t+1},h_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right) \\ &\quad\times p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\middle|\,\mathcal{D}^t\right) d\mu d\boldsymbol{\beta} d\boldsymbol{\theta}^{t+1} d\boldsymbol{\gamma_\theta} d\boldsymbol{Q} dh^{t+1} d\lambda_0 d\lambda_1 d\sigma_\xi^{-2}. \end{aligned} \tag{A-43}$$

and following the steps described in the SV and TVP sections above.

Appendix B Estimation of bond risk premia

In this appendix we describe how we estimate bond risk premia using a dynamic Gaussian affine term structure model. Yields are first collected in a vector, Y_t, which contains rates for J different maturities. The risk factors that determine the yields are denoted by Z_t; these include both macro factors (M_t) and yield factors ($F_t^{\mathcal{L}}$) extracted as the first $\mathcal{L}$ principal components of yields, i.e., $Z_t = \left(M_t', F_t^{\mathcal{L}\prime}\right)$. The macro factors can be unspanned, i.e., they are allowed to predict future yields without having additional explanatory power for current yields beyond the yield factors.

We begin by assuming that the continuously compounded nominal one-period interest rate r_t depends on the yield factors, but not on the macro factors,

$$r_t = \delta_0 + \delta_{\mathcal{F}}' F_t^{\mathcal{L}} + 0_{\mathcal{M}}' M_t. \tag{B-1}$$

Next, the risk factors are assumed to follow a Gaussian vector autoregression (VAR) under the risk neutral probability measure:

$$F_t^{\mathcal{L}} = \mu^{\mathbb{Q}} + \phi^{\mathbb{Q}} F_{t-1}^{\mathcal{L}} + 0_{\mathcal{L}\times\mathcal{M}}' M_{t-1} + \Sigma\varepsilon_t^{\mathbb{Q}}, \tag{B-2}$$

where $\varepsilon_t^{\mathbb{Q}} \sim \mathcal{N}(0, I)$. Finally, the evolution in Z_t under the physical measure takes the form:

$$Z_t = \mu + \phi Z_{t-1} + \Sigma\varepsilon_t, \tag{B-3}$$

where $\varepsilon_t \sim \mathcal{N}(0, I)$. Under these assumptions bond prices are exponentially affine in the yield factors and do not depend on the macro factors,

$$P_t^{(n)} = \exp\left(\mathcal{A}_n + \mathcal{B}_n F_t^{\mathcal{L}} + 0_{\mathcal{M}\times\mathcal{M}}' M_t\right), \tag{B-4}$$

where $\mathcal{A} = \mathcal{A}\left(\mu^{\mathbb{Q}}, \phi^{\mathbb{Q}}, \delta_0, \delta_{\mathcal{F}}, \Sigma\right)$ and $\mathcal{B} = \mathcal{B}\left(\phi^{\mathbb{Q}}, \delta_{\mathcal{F}}\right)$ are affine loadings which are given by the following recursions:

$$\begin{aligned} \mathcal{A}_{n+1} &= \mathcal{A}_n + \left(\mu^{\mathbb{Q}}\right)' \mathcal{B}_n + \frac{1}{2}\mathcal{B}_n' \Sigma\Sigma' \mathcal{B}_n - \delta_0 \\ \mathcal{B}_{n+1} &= \left(\phi^{\mathbb{Q}}\right)' \mathcal{B}_n - \delta_{\mathcal{F}} \end{aligned} \tag{B-5}$$

with starting values $\mathcal{A}_0 = 0$ and $\mathcal{B}_0 = 0$.

Using these coefficients, the model-implied yields are obtained as

$$y_t^{(n)} = -\frac{1}{n}\log\left(P_t^{(n)}\right) = A_j + B_n F_t^{\mathcal{L}}, \tag{B-6}$$

where

$$A_n = -\frac{1}{n}\mathcal{A}_n, \qquad B_n = -\frac{1}{n}\mathcal{B}_n. \tag{B-7}$$

Similarly, the risk neutral price of a n-period bond, $\widetilde{P}_t^{(n)}$, and its implied yield, $\widetilde{y}_t^{(n)}$, can be calculated as

$$\widetilde{P}_t^{(n)} = \exp\left(\widetilde{\mathcal{A}}_n + \widetilde{\mathcal{B}}_n Z_t\right), \tag{B-8}$$

and

$$\widetilde{y}_t^{(n)} = \widetilde{A}_n + \widetilde{B}_n Z_t, \tag{B-9}$$

where

$$\widetilde{A}_n = -\frac{1}{n}\widetilde{\mathcal{A}}_n, \qquad \widetilde{B}_n = -\frac{1}{n}\widetilde{\mathcal{B}}_n. \tag{B-10}$$

and where $\widetilde{\mathcal{A}}_n$ and $\widetilde{\mathcal{B}}_n$ are given by the following recursions:

$$\begin{aligned} \widetilde{\mathcal{A}}_{n+1} &= \widetilde{\mathcal{A}}_n + \mu' \widetilde{\mathcal{B}}_n + \frac{1}{2}\widetilde{\mathcal{B}}_n' \Sigma\Sigma' \widetilde{\mathcal{B}}_n - \delta_0 \\ \widetilde{\mathcal{B}}_{n+1} &= \phi' \widetilde{\mathcal{B}}_j - \delta_{\mathcal{F}} \end{aligned} \tag{B-11}$$

initialized at $\widetilde{\mathcal{A}}_0 = 0$ and $\widetilde{\mathcal{B}}_0 = 0$.

We follow Joslin et al. (2011) and Wright (2011) and include in $F_t^{\mathcal{L}}$ the first three principal components of zero-coupon bond yields using maturities ranging from three months to ten years. As macro factors, M_t, we use exponentially weighted moving averages of monthly inflation and IPI growth. The data used for estimation are monthly yields on zero-coupon bonds, inflation, and IPI growth, from 1982:01 to 2011:12. Our estimation approach also follows closely Joslin et al. (2011). Using the resulting estimates of the model parameters, we compute the risk premium at all maturities as the difference between the yields computed under the risk neutral measure, $\mathbb{Q}$, and the yields calculated under the physical measure, i.e.,

$$rp_t^{(n)} = y_t^{(n)} - \widetilde{y}_t^{(n)}. \tag{B-12}$$

Figure 1. **Bond excess returns**

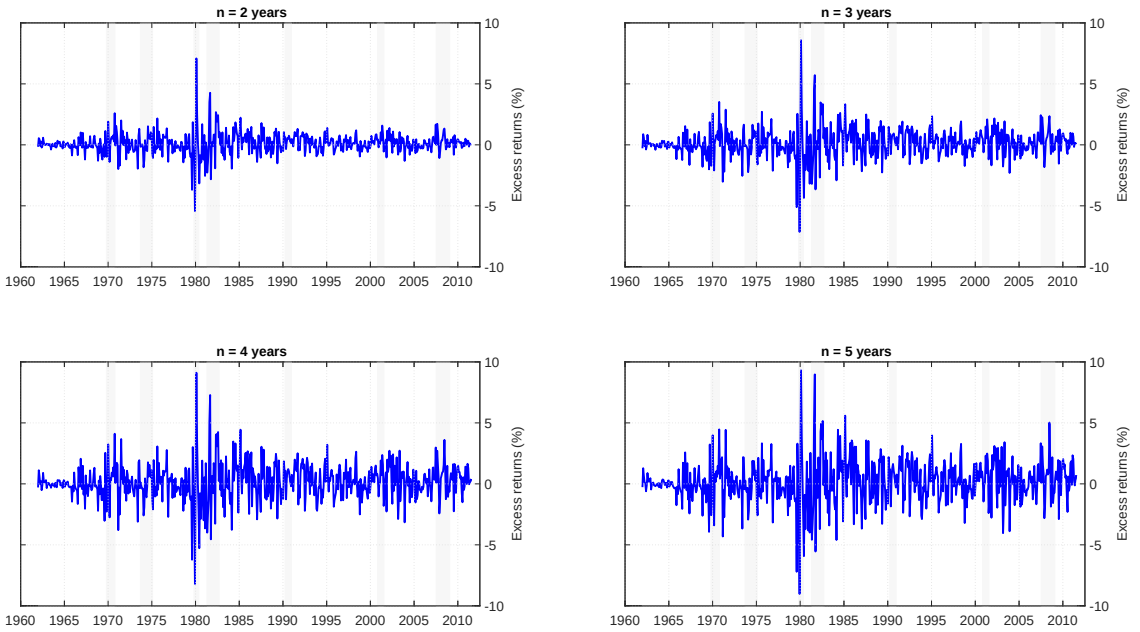

This figure shows time series of monthly bond excess returns (in percentage terms) for maturities (n) ranging from 2 years through 5 years. Monthly bond excess returns, $rx_{t+1/12}^{(n)}$, are computed from monthly yields, $y_t^{(n)}$, and are expressed in deviations from the 1-month T-bill rate, $rx_{t+1/12}^{(n)} = r_{t+1/12}^{(n)} - (1/12)y_t^{1/12}$, with $r_{t+1/12}^{(n)} = ny_t^{(n)} - (n - 1/12)y_{t+1/12}^{n-1/12}$. The sample ranges from 1962:01 to 2011:12.

Figure 2. **Parameter estimates for bond return forecasting model**

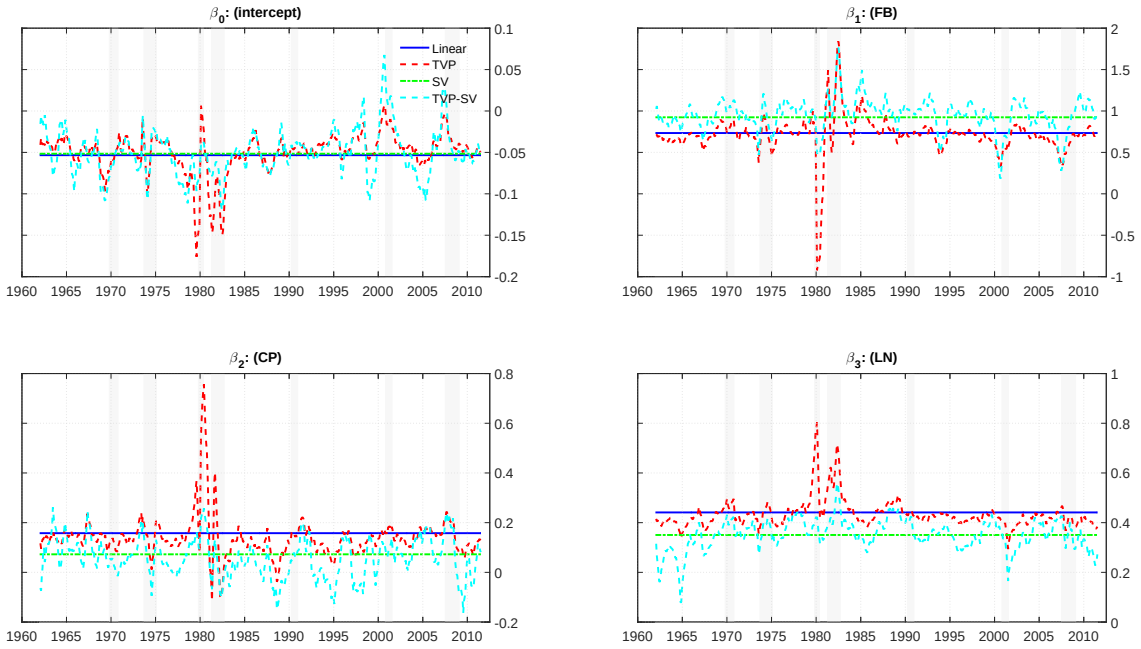

This figure displays parameter estimates for the FB-CP-LN model used to forecast monthly 3-year bond excess returns using as predictors the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) variables. The blue solid line represents the linear, constant coefficient model (Linear); the red dashed line tracks the parameter estimates for the time-varying parameter model (TVP); the green dashed-dotted line depicts the parameters for the stochastic volatility model (SV), while the dotted light-blue line shows estimates for the time-varying parameter, stochastic volatility (TVP-SV) model. The top left panel plots estimates of the intercept and the top right panel displays the coefficients on the FB predictor. The bottom left and right panels plot the coefficients on the CP and LN factors, respectively. The sample ranges from 1962:01 to 2011:12 and the parameter estimates are based on full-sample information.

Figure 3. **Posterior densities for model parameters**

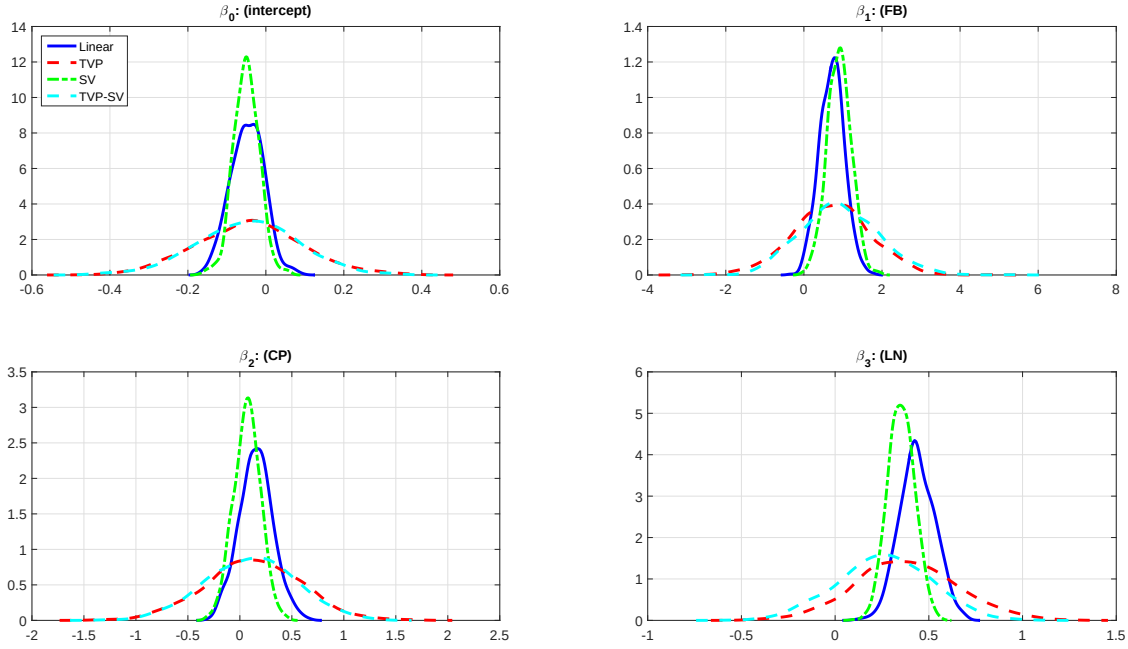

This figure displays posterior densities for the coefficients of the FB-CP-LN return model fitted to 3-year Treasury bonds, using as predictors the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) factors. The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line shows the parameter posterior density for the time-varying parameter (TVP) model; the green dashed-dotted line represents the stochastic volatility (SV) model, while the dotted light-blue line shows the posterior density for the time-varying parameter, stochastic volatility (TVP-SV) model. The first panel shows densities for the intercept. The second panel shows densities for the coefficient on the FB predictor. The third and fourth panels show densities for the coefficients on the CP and LN factors, respectively. The posterior density estimates shown here are based on their values as of 2011:12.

Figure 4. **Posterior densities for bond returns**

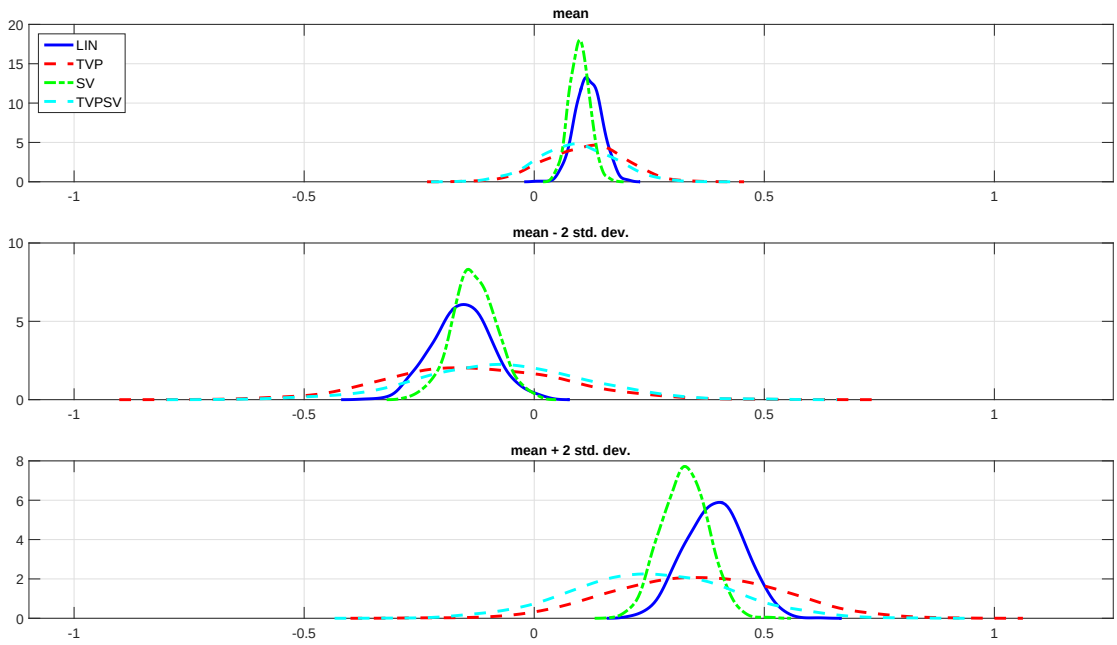

This figure shows the posterior density for excess returns on a three-year Treasury bond using the univariate Ludvigson-Ng (LN) state variable as a predictor. The LN variable is set at its sample mean $\overline{LN}$ (top panel), $\overline{LN} - 2stdev\,(LN)$ (middle panel), and $\overline{LN} + 2stdev\,(LN)$ (bottom panel). The blue solid line represents the linear, constant coefficient (Linear) model. the red dashed line tracks densities for the time-varying parameter (TVP) model. The green dashed-dotted line represents the stochastic volatility (SV) model, and the dotted light-blue line refers to the time varying parameter, stochastic volatility (TVP-SV) model. All posterior density estimates are based on the full data sample at the end of 2011.

Figure 5. **Conditional mean and volatility estimates for bond excess returns**

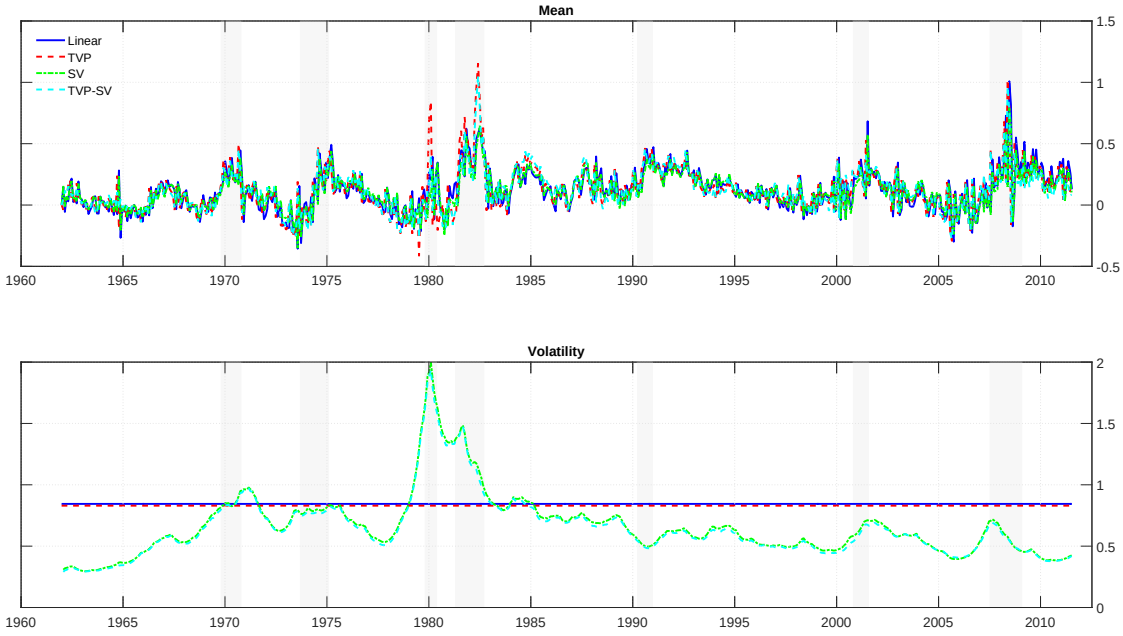

The top panel shows time-series of expected bond excess returns obtained from a range of models used to forecast monthly returns on a three-year Tresury bond using as predictors the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) factors. The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line tracks the time-varying parameter (TVP) model; the green dashed-dotted line depicts the stochastic volatility (SV) model, while the dotted light-blue line displays values for the time varying parameter, stochastic volatility (TVP-SV) model. The bottom panel displays volatility estimates for the FB-CP-LN models. The sample ranges from 1962:01 to 2011:12 and the estimates are based on full-sample information.

Figure 6. **Cumulative sum of squared forecast error differentials**

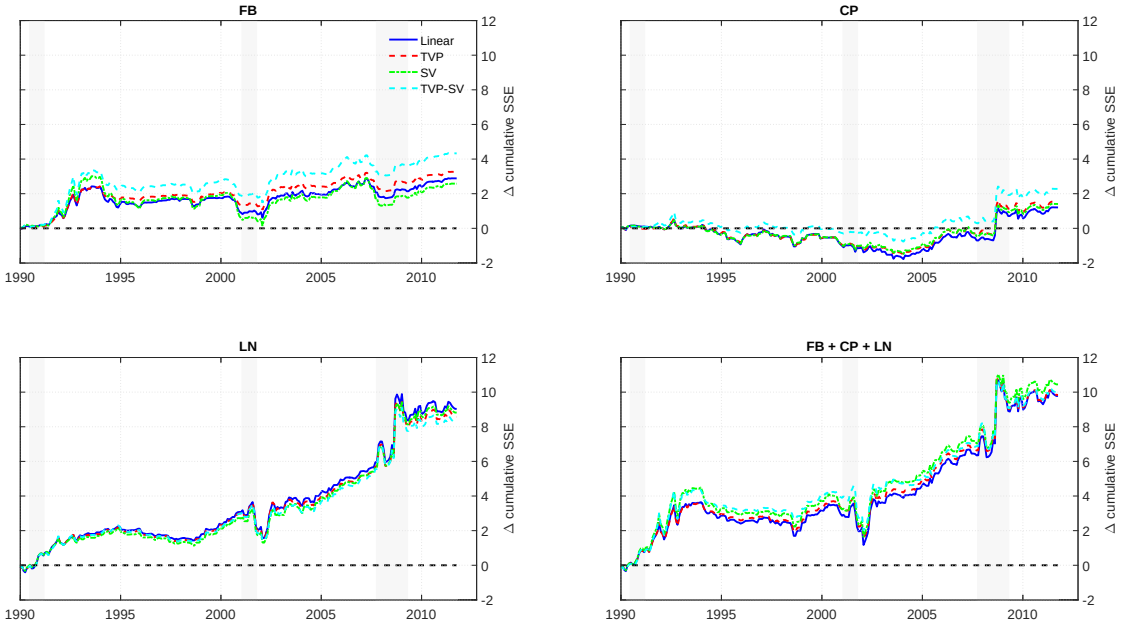

This figure shows the recursively calculated sum of squared forecast errors for the expectations hypothesis (EH) model minus the sum of squared forecast errors for a forecasting model with time-varying expected returns for a bond with a two year maturity, ($n = 2$). Each month we recursively estimate the parameters of the forecasting models and generate one-step-ahead forecasts of bond excess returns which are in turn used to compute out-of-sample forecasts. This procedure is applied to the EH model, which is our benchmark, as well as to forecasting models based on the Fama-Bliss (FB) predictor (1st window), the Cochrane-Piazzesi (CP) factor (2nd window), the Ludvigson-Ng (LN) factor (3rd window), and a multivariate model with all three predictors included (4th window). We then plot the cumulative sum of squared forecast errors (SSE_t) of the EH forecasts (SSE_t^{EH}) minus the corresponding value from the model with time-varying mean, $SSE_t^{EH} - SSE_t$. Values above zero indicate that a forecasting model with time-varying predictors produces more accurate forecasts than the EH benchmark, while negative values suggest the opposite. The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line tracks the time-varying parameter (TVP) model; the green dashed-dotted line represents the stochastic volatility (SV) model, while the dotted light-blue line refers to the time-varying parameter, stochastic volatility (TVP-SV) model. The out-of-sample period is 1990:01 - 2011:12.

Figure 7. **Cumulative sum of log-score differentials**

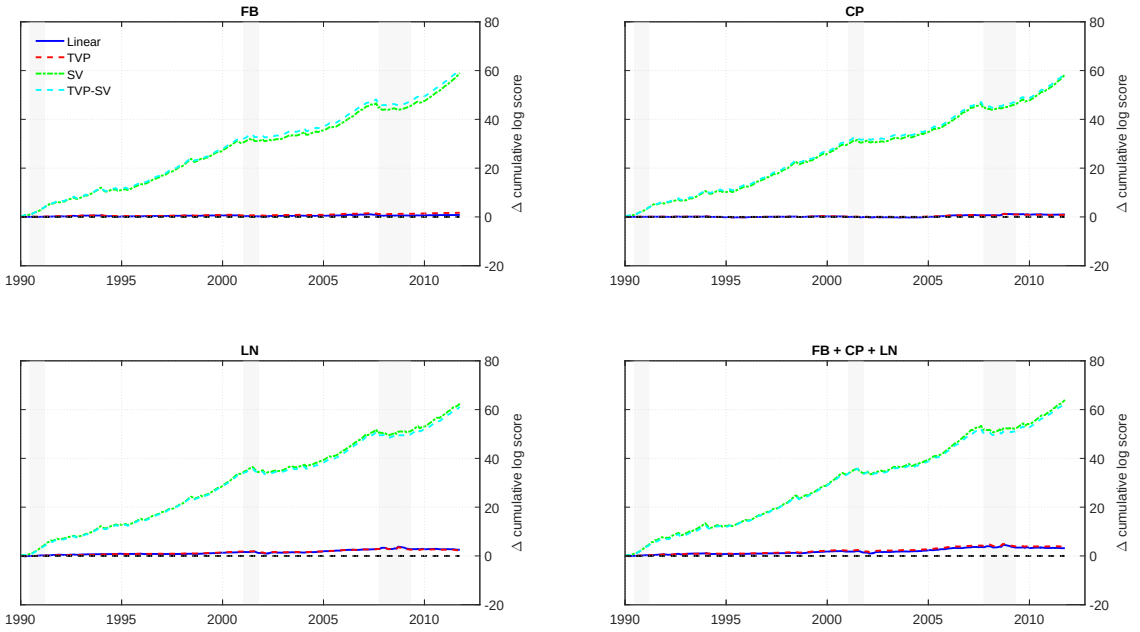

This figure shows the recursively calculated sum of log predictive scores from forecasting models with time-varying predictors minus the corresponding sum of log predictive scores for the EH model, using a 2-year Treasury bond. Each month we recursively estimate the parameters of the forecasting models and generate one-step-ahead density forecasts of bond excess returns which are in turn used to compute log-predictive scores. This procedure is applied to the benchmark EH model as well as to forecasting models based on the Fama-Bliss (FB) predictor (1st window), the Cochrane-Piazzesi (CP) factor (2nd window), the Ludvigson-Ng (LN) factor (3rd window), and a multivariate FB-CP-LN model (4th window). We then plot the cumulative sum of log predictive scores (LS_t) for the models with time-varying predictors minus the cumulative sum of log-predictive scores of the EH model, $LS_t - LS_t^{EH}$. Values above zero indicate that the time-varying mean model generates more accurate forecasts than the EH benchmark, while negative values suggest the opposite. The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line tracks the time-varying parameter (TVP) model; the green dashed-dotted line represents the stochastic volatility (SV) model, while the dotted light-blue line shows the time-varying parameter, stochastic volatility (TVP-SV) model. The out-of-sample period is 1990:01 - 2011:12.

Figure 8. **Economic value of out-of-sample bond return forecasts**

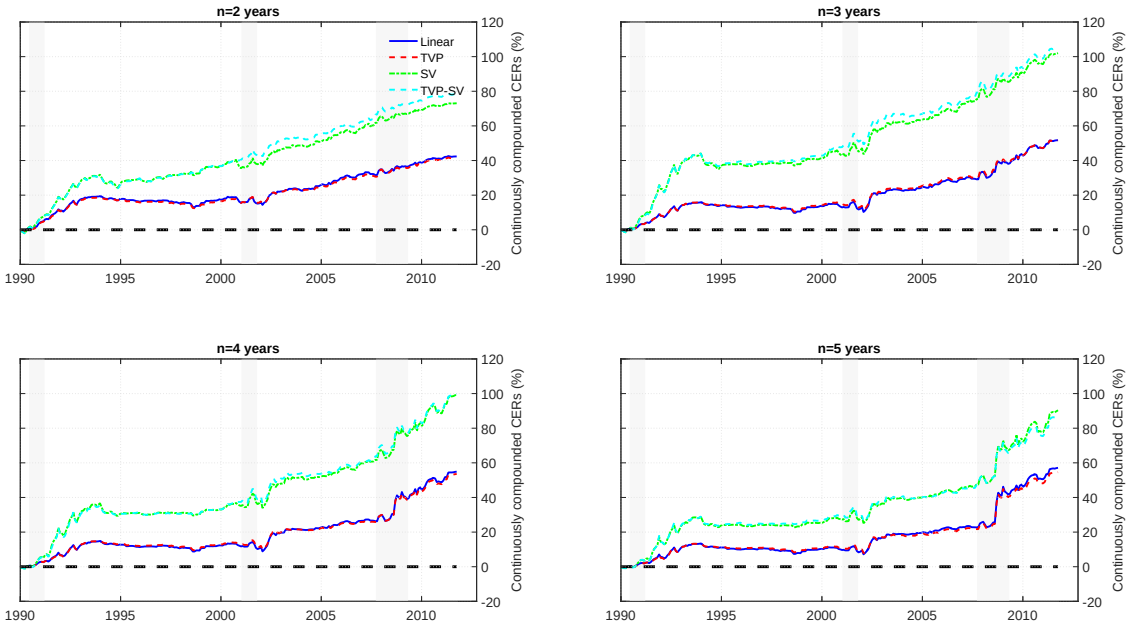

This figure plots cumulative certainty equivalent returns for the three-factor FB-CP-LN forecasting model that uses the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) factors as predictors, measured relative to the expectations hypothesis (EH) model. Each month we compute the optimal allocation to bonds and T-bills based on the predictive densities of bond excess returns. The investor is assumed to have power utility with a coefficient of relative risk aversion of ten and the weight on bonds is constrained to lie in the interval $[-200\%, 300\%]$. Each panel displays a different bond maturity, ranging from 2 years (1st panel) to 5 years (4th panel) The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line tracks the time-varying parameter (TVP) model; the green dashed-dotted line represents the stochastic volatility (SV) model, while the dotted light-blue line shows results for the time-varying parameter, stochastic volatility (TVP-SV) model. The out-of-sample period is 1990:01 - 2011:12.

Figure 9. **Weights on individual predictors in the optimal prediction pool**

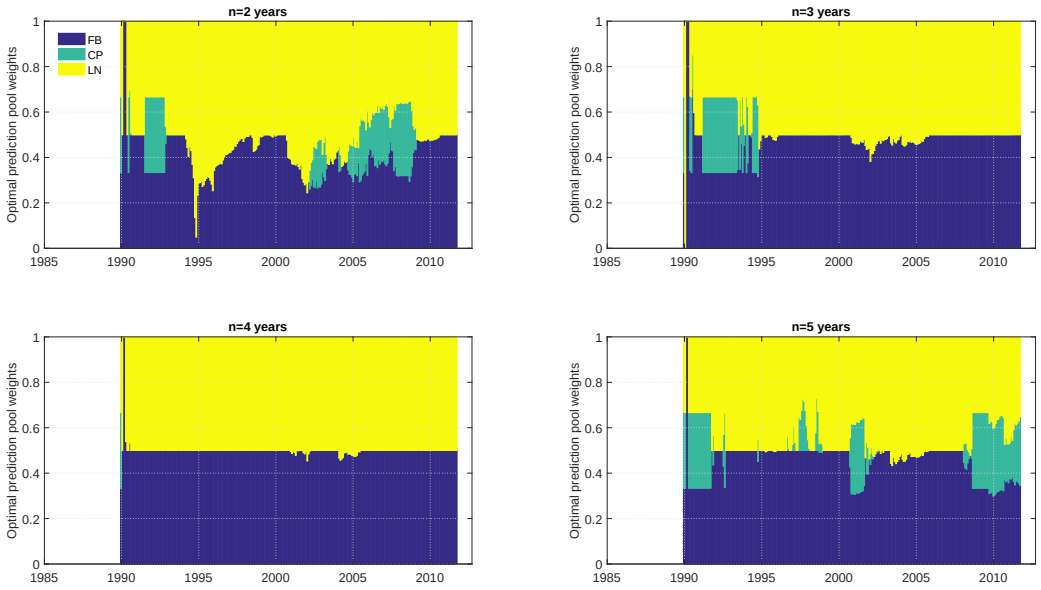

This figure plots the total weights on the individual predictors in the optimal predictive pool. At each point in time t, the weights are computed as $\mathcal{A}'\mathbf{w}_t$, where $\mathcal{A}$ is a 28×3 matrix representing all forecasting models by their unique combinations of zeros and ones and $\mathbf{w}_t$ is a 28×1 vector of period t optimal weights in the predictive pool, obtained in real time by solving the minimization problem

$$\mathbf{w}_t^* = \arg\max_{\mathbf{w}} \sum_{\tau=1}^{t-1} \log \left[\sum_{i=1}^{N} w_i \times S_{\tau+1,i} \right]$$

subject to $\mathbf{w}_t^*$ belonging to the $N-$dimensional unit simplex and N is the number of forecasting models. $S_{\tau+1,i}$ denotes the time $\tau+1$ log score for model i, $S_{\tau+1,i} = \exp(LS_{\tau+1,i})$. Blue bars show the combination weights associated with the Fama-Bliss (FB) factor; green bars show the weight assigned to the Cochrane-Piazzesi (CP) factor, and red bars show the weights assigned to the Ludvigson-Ng (LN) factor. Bond maturities range from 2 years (top left panel) to 5 years (bottom right panel).

Table 1. **Summary Statistics.**

	2 years	3 years	4 years	5 years
Panel A.1: One-month excess returns				
mean	1.4147	1.7316	1.9868	2.1941
mean (gross)	6.442	6.753	7.011	7.214
st.dev.	2.9711	4.1555	5.2174	6.2252
skew	0.4995	0.2079	0.0566	0.0149
kurt	14.8625	10.6482	7.9003	6.5797
AC(1)	0.1692	0.1533	0.1384	0.1235
Panel A.2: 12-month overlapping excess returns				
mean	0.5149	0.8601	1.1176	1.3177
mean (gross)	6.3511	6.6892	6.9353	7.1220
st.dev.	1.7527	3.1811	4.4103	5.5139
skew	-0.1014	-0.0618	-0.0196	0.0157
kurt	3.7801	3.5919	3.5172	3.5311
AC(1)	0.9313	0.9321	0.9322	0.9312
Panel A.3: 12-month overlapping excess returns (Cochrane-Piazzesi)				
mean	0.4768	0.8662	1.1435	1.2197
mean(gross)	6.2300	6.6193	6.8966	6.9729
st.dev.	1.7691	3.2438	4.5064	5.5381
skew	0.0636	-0.0256	0.0208	0.0157
kurt	3.4734	3.5507	3.4963	3.4792
AC(1)	0.9312	0.9336	0.9325	0.9234

Panel B: Predictors						
	Fama Bliss				CP	LN
	2-years	3-years	4-years	5-years		
mean	0.1078	0.1287	0.1451	0.1584	0.1533	0.1533
st.dev.	0.0996	0.1155	0.1278	0.1376	0.2126	0.3123
skew	-0.0716	-0.2379	-0.2234	-0.1434	0.7852	0.8604
kurt	3.7558	3.3782	3.0409	2.7548	5.2884	5.7283
AC(1)	0.8787	0.8995	0.9127	0.9227	0.6740	0.4109

Panel C: Correlation Matrix						
	FB-2	FB-3	FB-4	FB-5	CP	LN
FB-2	1.000	0.973	0.926	0.879	0.460	-0.087
FB-3		1.000	0.987	0.961	0.472	-0.049
FB-4			1.000	0.993	0.490	-0.010
FB-5				1.000	0.500	0.022
CP					1.000	0.184
LN						1.000

This table reports summary statistics for monthly bond excess returns and the predictor variables used in our study. Panels A.1-A.3 report the mean, standard deviation, skewness, kurtosis and first-order autocorrelation (AC(1)) of bond excess returns for 2 to 5-year bond maturities. Panel A.1 is based on monthly returns computed in excess of a one-month T-bill rate while Panels A.2 and Panel A.3 are based on 12-month overlapping returns, computed in excess of a 12-month T-bill rate. Gross returns do not subtract the risk-free rate. In Panels A.1 and A.2 returns are constructed using daily treasury yield data from Gurkaynak et al. (2007) while in Panel A.3 returns are constructed as in Cochrane and Piazzesi (2005) using the Fama-Bliss CRSP files. The row $mean(gross)$ reports the sample mean of gross bond returns (without subtracting the risk-free). Panel B reports the same summary statistics for the predictors: the Fama-Bliss (FB) forward spreads (2, 3, 4, and 5 years), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) factors. Panel C reports the correlation matrix for the predictors. The sample period is 1962-2011.

Table 2. **Full-sample OLS estimates**

	FB	CP	LN	FB+CP+LN
		2 years		
β_{FB}	1.1648***			1.1221***
β_{CP}		0.6549***		0.2317
β_{LN}			0.6713***	0.6736***
R^2	0.0166	0.0247	0.0581	0.0822
		3 years		
β_{FB}	1.3742***			1.2060***
β_{CP}		0.8785***		0.3299
β_{LN}			0.9068***	0.8877***
R^2	0.0158	0.0226	0.0541	0.0743
		4 years		
β_{FB}	1.6661***			1.3500***
β_{CP}		1.1054***		0.4202
β_{LN}			1.1149***	1.0679***
R^2	0.0183	0.0227	0.0518	0.0719
		5 years		
β_{FB}	1.9726***			1.4877***
β_{CP}		1.3613***		0.5502*
β_{LN}			1.3070***	1.2241***
R^2	0.0211	0.0243	0.0499	0.0712

This table reports OLS estimates of the slope coefficients for four linear models based on inclusion or exclusion of the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) predictor computed from a projection of the time series of cross-sectional averages of the 2, 3, 4, 5 bond excess returns on the 1, 2, 3, 4 and 5 year forward rates, and the Ludvigson-Ng (LN) predictor computed from a projection of the time-series of cross-sectional averages of the 2, 3, 4, 5 bond excess returns on five principal components obtained from a large panel of macroeconomic variables. Columns (1)-(3) report results for the univariate models, column (4) for the multivariate model that includes all three predictors. The last row in each panel reports the adjusted R^2. Stars indicate statistical significance based on Newey-West standard errors. ***: significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table 3. **Out-of-sample forecasting performance: R^2 values**

	Panel A: 2 years					Panel B: 3 years				
Model	OLS	LIN	SV	TVP	TVPSV	OLS	LIN	SV	TVP	TVPSV
FB	0.22%*	0.96%*	0.31%*	1.79%**	**2.48%****	1.77%**	1.46%**	1.31%**	1.66%**	**2.20%****
CP	-1.58%	0.52%*	0.67%*	0.70%*	**1.82%****	-0.45%	0.61%	0.71%	0.78%	**1.16%***
LN	-0.27%	4.94%***	5.51%***	4.88%***	**5.52%****	2.60%***	**4.59%****	4.48%***	4.35%***	4.20%***
$FB+CP+LN$	-3.06%	4.72%***	5.75%***	5.05%***	**6.47%****	0.43%***	4.97%***	**5.31%****	5.00%***	5.07%***
	Panel C: 4 years					Panel D: 5 years				
Model	OLS	LIN	SV	TVP	TVPSV	OLS	LIN	SV	TVP	TVPSV
FB	**2.39%****	1.64%***	1.72%**	1.64%**	2.08%***	**2.61%****	1.67%**	1.75%**	1.61%**	1.81%**
CP	0.25%	0.69%	0.63%	**0.94%***	0.87%*	0.73%	0.80%*	0.65%	0.79%*	**0.95%***
LN	3.77%***	**4.02%****	3.67%***	3.58%***	3.52%***	**4.32%****	3.43%***	3.26%***	3.07%**	3.15%**
$FB+CP+LN$	1.89%***	4.78%***	**4.93%****	4.56%***	4.67%***	2.51%***	4.45%***	**4.53%****	4.25%***	4.26%***

This table reports out-of-sample R^2 values for four prediction models based on the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) predictors fitted to monthly bond excess returns, rx_{t+1}, measured relative to the one-month T-bill rate. The R^2_{OoS} is measured relative to the EH model: $R^2_{OoS} = 1 - \frac{\sum_{\tau=\underline{t}-1}^{\bar{t}-1}(r_{xt+1}-\hat{r}_{xt+1|t})^2}{\sum_{\tau=\underline{t}-1}^{\bar{t}-1}(r_{xt+1}-\bar{r}_{xt+1|t})^2}$ where $\hat{r}_{xt+1|t}$ is the conditional mean of bond returns based on a regression of monthly excess returns on an intercept and lagged predictor variable(s), x_t: $r_{xt+1} = \mu + \beta' x_t + \varepsilon_{t+1}$. $\bar{r}_{t+1|t}$ is the forecast from the EH model which assumes that the βs are zero. We report results for five specifications: (i) ordinary least squares (OLS), (ii) a linear specification with constant coefficients and constant volatility (LIN), (iii) a model that allows for stochastic volatility (SV), (iv) a model that allows for time-varying coefficients (TVP) and (v) a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). The out-of-sample period starts in January 1990 and ends in 2011:12. We measure statistical significance relative to the expectation hypothesis model using the Clark and West (2007) test statistic. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 4. **Out-of-sample forecasting performance: predictive likelihood**

	Panel A: 2 years				Panel B: 3 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.003	0.223***	0.006***	**0.230***	0.003	0.130***	0.006***	**0.132***
CP	0.004	0.221***	0.003	**0.223***	0.004	**0.129***	0.004	0.128***
LN	0.010**	**0.237***	0.009**	0.232***	0.011**	**0.141***	0.011**	0.138***
$FB + CP + LN$	0.012**	**0.242***	0.015***	0.238***	0.013**	**0.146***	0.016***	0.141***
	Panel C: 4 years				Panel D: 5 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.007***	0.090***	0.006**	**0.090***	0.006**	0.066***	0.006**	**0.067***
CP	0.006**	0.086***	0.005**	**0.086***	0.005**	0.062***	0.006**	**0.066***
LN	0.013**	**0.099***	0.011**	0.098***	0.013**	**0.075***	0.012**	0.073***
$FB + CP + LN$	0.017***	**0.104***	0.017**	0.101***	0.016**	**0.079***	0.017**	0.078***

This table reports the log predictive score for four forecasting models that allow for time-varying predictors relative to the log-predictive score computed under the expectation hypothesis (EH) model. The four forecasting models use the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. Positive values of the test statistic indicate that the model with time-varying predictors generates more precise forecasts than the EH benchmark. We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). The results are based on out-of-sample estimates over the sample period 1990 - 2011. ***: significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table 5. **Out-of-sample economic performance of bond portfolios**

	Weights constrained between 0% and 99%							
	Panel A: 2 years				Panel B: 3 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.23%	-0.25%	-0.15%	**-0.02%**	0.06%	0.21%	0.08%	**0.36%****
CP	-0.20%	0.03%	-0.15%	**0.07%**	-0.15%	0.37%**	-0.09%	**0.45%*****
LN	**0.11%**	0.07%	0.09%	0.09%	0.58%***	**0.70%*****	0.57%***	0.62%***
$FB+CP+LN$	0.05%	0.15%	0.04%	**0.24%****	0.49%**	0.72%***	0.50%**	**0.82%*****
	Panel C: 4 years				Panel D: 5 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.46%*	0.71%**	0.50%**	**0.81%*****	0.67%**	**1.02%****	0.62%**	1.01%**
CP	0.04%	0.46%**	0.12%	**0.64%*****	0.18%	0.55%**	0.18%	**0.69%*****
LN	0.87%***	**1.23%*****	0.78%***	1.21%***	0.86%***	**1.38%*****	0.78%***	1.36%***
$FB+CP+LN$	0.94%***	1.19%***	0.94%***	**1.25%*****	1.07%***	1.54%***	1.06%***	**1.55%*****
	Weights constrained between -200% and 300%							
	Panel E: 2 years				Panel F: 3 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.25%	0.99%*	0.39%*	**1.78%*****	0.50%**	1.30%**	0.53%**	**1.72%*****
CP	0.14%	1.47%***	0.17%	**1.85%*****	0.22%	1.15%**	0.25%	**1.28%****
LN	1.68%***	**2.84%*****	1.62%***	2.83%***	1.61%***	**2.98%*****	1.53%***	2.94%***
$FB+CP+LN$	1.69%***	2.73%***	1.65%***	**2.95%*****	1.79%***	3.27%***	1.78%***	**3.34%*****
	Panel G: 4 years				Panel H: 5 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.63%**	1.33%**	0.61%**	**1.47%****	0.69%**	**1.26%****	0.65%**	1.24%**
CP	0.24%	0.83%**	0.32%	**0.91%****	0.28%	0.70%**	0.27%	**0.88%****
LN	1.47%***	**2.35%*****	1.31%***	2.29%***	1.32%**	**2.02%*****	1.15%**	1.78%**
$FB+CP+LN$	1.73%***	**2.78%*****	1.65%***	2.66%**	1.66%**	**2.27%****	1.49%**	1.96%*

This table reports annualized certainty equivalent return values for portfolio decisions based on recursive out-of-sample forecasts of bond excess returns. Each period an investor with power utility and coefficient of relative risk aversion of 10 selects 2, 3, 4, or 5-year bond and 1-month T-bills based on the predictive density implied by a given model. The four forecasting models use the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model with both time varying coefficients and stochastic volatility ($TVPSV$). Statistical significance is based on a one-sided Diebold-Mariano test applied to the out-of-sample period 1990-2011. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 6. **Bond return predictability in expansions and recessions**

Model	LIN		SV		TVP		TVPSV	
	Exp	Rec	Exp	Rec	Exp	Rec	Exp	Rec
				Panel A: 2 years				
FB	2.53%	0.44%	3.01%	0.16%	6.76%	**8.11%****	7.96%	3.76%
CP	1.85%	**2.91%****	1.93%	**2.62%***	6.08%	5.61%**	6.47%	4.39%**
LN	2.11%	**9.30%*****	2.26%	**7.73%*****	4.94%	**15.46%*****	6.05%	**10.66%*****
$FB+CP+LN$	5.33%	**10.69%*****	5.51%	**8.66%*****	12.30%	**23.46%*****	12.74%	**14.93%*****
				Panel B: 3 years				
FB	2.17%	0.06%	2.74%	-0.28%	4.56%	**4.74%****	5.32%	1.83%
CP	1.36%	**2.70%****	1.42%	**2.80%****	3.83%	**4.33%****	4.03%	3.82%**
LN	2.09%	**7.93%*****	2.12%	**7.29%*****	3.77%	**11.89%*****	4.16%	**9.32%*****
$FB+CP+LN$	4.50%	**8.73%*****	4.63%	**7.85%*****	8.58%	**17.21%*****	8.74%	**11.83%*****
				Panel C: 4 years				
FB	1.99%	0.18%	2.45%	-0.05%	3.65%	3.11%*	4.20%	1.41%
CP	1.12%	**2.70%****	1.20%	**2.83%****	2.87%	**4.08%****	2.94%	**3.79%****
LN	1.89%	**7.03%*****	1.93%	**6.93%*****	3.06%	**10.11%*****	3.28%	**8.60%*****
$FB+CP+LN$	3.91%	**7.88%*****	4.02%	**7.54%*****	6.83%	**14.04%*****	6.84%	**10.69%*****
				Panel D: 5 years				
FB	1.88%	0.42%	2.23%	0.25%	3.11%	2.33%	3.55%	1.24%
CP	1.03%	**2.78%****	1.08%	**2.85%****	2.37%	**3.88%****	2.44%	**3.70%****
LN	1.65%	**6.32%*****	1.69%	**6.19%*****	2.52%	**8.85%*****	2.74%	**8.05%*****
$FB+CP+LN$	3.41%	**7.30%*****	3.49%	**7.10%*****	5.59%	**12.04%*****	5.62%	**9.87%*****

This table reports the R^2 from regressions of bond excess returns on the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. We report results separately for expansions (Exp) and recessions (Rec) as defined by the NBER recession index. Results are shown for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). The R^2 in expansions is computed as $R^2_{i,0} = 1 - \frac{e_{i,0}{}' e_{i,0}}{e_{EH,0}{}' e_{EH,0}}$ where $e_{i,0}$ and $e_{EH,0}$ denote the vectors of residuals of the alternative and the benchmark model, respectively, during expansions. Similarly, the R^2 in recessions only uses the vector of residuals in recessions: $R^2_{i,1} = 1 - \frac{e_{i,1}{}' e_{i,1}}{e_{EH,1}{}' e_{EH,1}}$. We test whether the R^2 is higher in recessions than in expansions using a bootstrap methodology. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 7. **Sharpe ratios in expansions and recessions**

Model	LIN		SV		TVP		TVPSV	
	Exp	Rec	Exp	Rec	Exp	Rec	Exp	Rec
				Panel A: 2 years				
FB	0.49%	0.43%	0.49%	0.35%	0.48%	**0.56%**	0.50%	0.48%
CP	0.46%	**0.63%**	0.46%	**0.54%**	0.44%	**0.72%**	0.47%	**0.65%**
LN	0.38%	**1.12%**	0.35%	**0.90%**	0.37%	**1.22%**	0.37%	**1.00%**
$FB+CP+LN$	0.38%	**1.14%**	0.42%	**0.91%**	0.37%	**1.33%**	0.42%	**1.05%**
				Panel B: 3 years				
FB	0.43%	0.38%	0.48%	0.36%	0.43%	**0.44%**	0.49%	0.42%
CP	0.40%	**0.53%**	0.47%	**0.56%**	0.39%	**0.59%**	0.47%	**0.61%**
LN	0.35%	**0.93%**	0.37%	**0.93%**	0.34%	**1.00%**	0.38%	**1.00%**
$FB+CP+LN$	0.34%	**0.92%**	0.42%	**0.94%**	0.34%	**1.04%**	0.43%	**1.02%**
				Panel C: 4 years				
FB	0.39%	0.35%	0.49%	0.38%	0.39%	0.38%	0.48%	0.40%
CP	0.37%	**0.48%**	0.47%	**0.55%**	0.37%	**0.52%**	0.48%	**0.61%**
LN	0.32%	**0.80%**	0.39%	**0.93%**	0.32%	**0.85%**	0.39%	**0.98%**
$FB+CP+LN$	0.32%	**0.78%**	0.43%	**0.93%**	0.31%	**0.88%**	0.43%	**0.99%**
				Panel D: 5 years				
FB	0.36%	0.32%	0.48%	0.38%	0.35%	0.34%	0.48%	0.41%
CP	0.34%	**0.44%**	0.47%	**0.55%**	0.34%	**0.47%**	0.47%	**0.58%**
LN	0.31%	**0.71%**	0.40%	**0.89%**	0.30%	**0.75%**	0.40%	**0.96%**
$FB+CP+LN$	0.30%	**0.69%**	0.43%	**0.89%**	0.31%	**0.77%**	0.43%	**0.97%**

This table reports the annualized Sharpe ratio computed from conditional mean and conditional volatility estimates implied by regressions of bond excess returns on the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. We report results separately for expansions (Exp) and recessions (Rec) as defined by the NBER recession index. Results are shown for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$).

Table 8. **Correlations between expected bond excess returns and economic variables**

	Panel A: GDP				Panel B: Inflation			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.14	0.15	0.08	0.06	-0.11	-0.09	-0.13	-0.11
CP	-0.34***	-0.32***	-0.39***	-0.41***	-0.31***	-0.32***	-0.35***	-0.40***
LN	-0.61***	-0.61***	-0.62***	-0.63***	-0.38***	-0.36***	-0.38***	-0.36***
$FB + CP + LN$	-0.49***	-0.41***	-0.51***	-0.47***	-0.36***	-0.31***	-0.36***	-0.33***
	Panel C: GDP Uncertainty				Panel D: Inflation Uncertainty			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.29***	0.27***	0.32***	0.31***	0.01	-0.02	0.04	0.02
CP	0.45***	0.44***	0.46***	0.46***	0.39***	0.37***	0.40***	0.35***
LN	0.51***	0.52***	0.50***	0.49***	0.49***	0.47***	0.48***	0.45***
$FB + CP + LN$	0.57***	0.57***	0.56***	0.54***	0.45***	0.39***	0.45***	0.38***

This table reports the contemporaneous correlations between out-of-sample forecasts of excess returns on a two-year Treasury bond and real GDP growth (Panel A), inflation (Panel B), real GDP growth uncertainty (Panel C) and inflation uncertainty (Panel D). Real GDP growth is computed as $\Delta log(GDP_{t+1})$ where GDP_{t+1} is the real gross domestic product (GDPMC1 Fred mnemonic). Inflation is computed as $\Delta log(CPI_{t+1})$ where CPI is the consumer price index for all urban consumers (CPIAUCSL Fred mnemonic). Real GDP growth uncertainty is the cross-sectional dispersion (the difference between the 75th percentile and the 25th percentile) for real GDP forecasts from the Philadelphia Fed Survey of Professional Forecasters. Inflation uncertainty is the cross-sectional dispersion (the difference between the 75th percentile and the 25th percentile) for CPI forecasts from the Philadelphia Fed Survey of Professional Forecasters. The bond return prediction models use the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). Finally, we test whether the correlation coefficients are statistically different from zero. All results are based on the out-of-sample period 1990-2011. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 9. **Expected bond excess returns and survey forecasts of the 3-Month T-bill rate.**

	Panel A: 2 years							
	Slope coefficient				R^2			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.04***	-0.04***	-0.03***	-0.04***	0.20	0.15	0.20	0.14
CP	-0.05***	-0.05***	-0.05***	-0.04***	0.30	0.30	0.30	0.24
LN	-0.06***	-0.05***	-0.06***	-0.04***	0.15	0.12	0.14	0.09
$FB+CP+LN$	-0.10***	-0.08***	-0.09***	-0.07***	0.24	0.20	0.22	0.14
	Panel B: 3 years							
	Slope coefficient				R^2			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.06***	-0.08***	-0.06***	-0.07***	0.45	0.36	0.45	0.35
CP	-0.06***	-0.06***	-0.06***	-0.06***	0.38	0.39	0.37	0.35
LN	-0.08***	-0.07***	-0.08***	-0.07***	0.21	0.19	0.20	0.16
$FB+CP+LN$	-0.13***	-0.12***	-0.12***	-0.11***	0.33	0.32	0.31	0.27
	Panel C: 4 years							
	Slope coefficient				R^2			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.09***	-0.10***	-0.09***	-0.10***	0.59	0.52	0.59	0.49
CP	-0.07***	-0.07***	-0.07***	-0.07***	0.45	0.45	0.44	0.43
LN	-0.09***	-0.08***	-0.09***	-0.08***	0.26	0.25	0.24	0.21
$FB+CP+LN$	-0.15***	-0.15***	-0.15***	-0.15***	0.40	0.39	0.38	0.35
	Panel D: 5 years							
	Slope coefficient				R^2			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.11***	-0.13***	-0.11***	-0.12***	0.69	0.62	0.67	0.58
CP	-0.08***	-0.08***	-0.08***	-0.08***	0.50	0.50	0.48	0.48
LN	-0.10***	-0.09***	-0.10***	-0.10***	0.30	0.28	0.28	0.26
$FB+CP+LN$	-0.17***	-0.17***	-0.17***	-0.17***	0.44	0.44	0.42	0.40

This table reports the R^2 and OLS estimates of slope coefficients from a regression of the predicted bond excess return on consensus forecasts of the 3-month T-bill rate reported by the Philadelphia Fed Survey of Professional Forecasters. The bond return prediction models use the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. Panels A-D display results for 2-5 year bond maturities. We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). All results are based on the sample 1990-2011. Stars indicate statistical significance based on Newey-West standard errors. ***: significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table 10. **Risk-premium regression**

	Panel A: 2 years				Panel B: 3 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
β	0.18	0.52	0.21	0.60	0.12	0.35	0.13	0.37
t-stat	0.78	2.45	0.95	2.95	0.58	1.77	0.67	1.85
	Panel C: 4 years				Panel D: 5 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
β	0.12	0.28	0.11	0.27	0.13	0.23	0.11	0.21
t-stat	0.63	1.52	0.56	1.42	0.73	1.33	0.59	1.16

This table reports OLS estimates of the slope coefficients (and the relative t-stats) from the following regression

$$\overline{rx}_t = \mu + \beta \overline{rp}_t + u_t,$$

where $\overline{rx}_t$ denotes the predicted bond excess returns and $\overline{rp}_t$ denotes the risk premium estimates, obtained from a term structure model with unspanned macro risks, based on the approach of Joslin et. a. (2011) and Wright (2011). We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). All results are based on the sample 1990-2011. Stars indicate statistical significance based on Newey-West standard errors. ***: significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table 11. **Economic and statistical performance of forecast combinations**

Method	2 years	3 years	4 years	5 years
	Panel A: Out-of-sample R^2			
OW	**5.30%*****	**4.99%*****	**5.13%*****	**5.07%*****
EW	**5.35%*****	**4.12%*****	**3.59%*****	**3.11%*****
BMA	**5.51%*****	**4.66%*****	**4.17%*****	**3.67%*****
	Panel B: Predictive Likelihood			
OW	**0.24*****	**0.15*****	**0.10*****	**0.07*****
EW	**0.14*****	**0.09*****	**0.06*****	**0.05*****
BMA	**0.24*****	**0.14*****	**0.09*****	**0.07*****
	Panel C: CER (Weights between 0% and 99%)			
OW	**0.18%**	**0.77%*****	**1.23%*****	**1.55%*****
EW	**0.17%***	**0.60%*****	**1.02%*****	**1.05%*****
BMA	**0.22%***	**0.80%*****	**1.22%*****	**1.48%*****
	Panel D: CER (Weights between -200% and 300%)			
OW	**2.57%*****	**3.27%*****	**2.88%*****	**1.99%*****
EW	**1.98%*****	**1.83%*****	**1.67%*****	**1.43%*****
BMA	**2.80%*****	**2.95%*****	**2.27%*****	**1.83%*****

This table reports out-of-sample results for the optimal predictive pool (OW) of Geweke and Amisano (2011), an equal-weighted (EW) model combination scheme, and Bayesian Model Averaging (BMA) applied to 28 forecasting models that models based on all possible combinations of the CP, FB and LN factors estimated using linear, SV, TVP and TVPSV methods. In each case the models and combination weights are estimated recursively using only data up to the point of the forecast. The R^2 values in Panel A use the out-of-sample R^2 measure proposed by Campbell and Thompson (2008). The predictive likelihood in Panel B is the value of the test for equal accuracy of the predictive density log-scores proposed by Clark and Ravazzolo (2014). CER values in Panels C and D are the annualized certainty equivalent returns derived for an investor with power utility and a coefficient of relative risk aversion of 10 who uses the posterior predictive density implied by the forecast combination. The forecast evaluation sample is 1990:01-2011:12. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Bond Return Predictability: Economic Value and Links to the Macroeconomy*

Antonio Gargano†
University of Melbourne

Davide Pettenuzzo‡
Brandeis University

Allan Timmermann§
University of California San Diego

July 28, 2016

Abstract

Studies of bond return predictability find a puzzling disparity between strong statistical evidence of return predictability and the failure to convert return forecasts into economic gains. We show that resolving this puzzle requires accounting for important features of bond return models such as time varying parameters, volatility dynamics, and unspanned macro factors. A three-factor model comprising the Fama and Bliss (1987) forward spread, the Cochrane and Piazzesi (2005) combination of forward rates and the Ludvigson and Ng (2009) macro factor generates notable gains in out-of-sample forecast accuracy compared with a model based on the expectations hypothesis. Such gains in predictive accuracy translate into higher risk-adjusted portfolio returns after accounting for estimation error and model uncertainty. Consistent with models featuring unspanned macro factors, our forecasts of future bond excess returns are strongly negatively correlated with survey forecasts of short rates.

Key words: bond returns; yield curve; macro factors; stochastic volatility; time-varying parameters; unspanned macro risk factors.

JEL codes: G11, G12, G17

*We thank Pierluigi Balduzzi, Alessandro Beber, Carlos Carvalho, Jens Hilscher, Blake LeBaron, Spencer Martin and seminar participants at USC, University of Michigan, Central Bank of Belgium, Bank of Canada, Carleton University, Imperial College, ESSEC Paris, Boston Fed, SoFiE Meeting, McCombs School of Business, University of Connecticut, University of Illinois Urbana-Champaign, and Econometric Society Australasian Meeting (ESAM) for comments on the paper.

†University of Melbourne, Building 110, Room 11.042, 198 Berkeley Street, Melbourne, 3010. Email: antonio.gargano@unimelb.edu.au

‡Brandeis University, Sachar International Center, 415 South St, Waltham, MA, Tel: (781) 736-2834. Email: dpettenu@brandeis.edu

§University of California, San Diego, 9500 Gilman Drive, MC 0553, La Jolla CA 92093. Tel: (858) 534-0894. Email: atimmerm@ucsd.edu

1 Introduction

Treasury bonds play an important role in many investors' portfolios so an understanding of the risk and return dynamics for this asset class is of central economic importance.[1] Some studies document significant in-sample predictability of Treasury bond excess returns for 2-5 year maturities by means of variables such as forward spreads (Fama and Bliss, 1987), yield spreads (Campbell and Shiller, 1991), a linear combination of forward rates (Cochrane and Piazzesi, 2005) and factors extracted from a cross-section of macroeconomic variables (Ludvigson and Ng, 2009).

While empirical studies provide statistical evidence in support of bond return predictability, there is so far little evidence that such predictability could have been used in real time to improve investors' economic utility. Thornton and Valente (2012) find that forward spread predictors, when used to guide the investment decisions of an investor with mean-variance preferences, do not lead to higher out-of-sample Sharpe ratios or higher economic utility compared with investments based on a no-predictability expectations hypothesis (EH) model. Sarno et al. (2016) reach a similar conclusion.

To address this puzzling contradiction between the statistical and economic evidence on bond return predictability, we propose a new empirical modeling approach that generalizes the existing literature in economically insightful ways. Modeling bond return dynamics requires adding several features absent from the regression models used in the existing literature. First, bond prices, and thus bond returns, are sensitive to monetary policy and inflation prospects, both of which are known to shift over time.[2] This suggests that it is important to adopt a framework that accounts for time varying parameters and even for the possibility that the forecasting model may shift over time, requiring that we allow for model uncertainty. Second, uncertainty about inflation prospects changes over time and the volatility of bond yields has also undergone shifts–most notably during the Fed's monetarist experiment from 1979-1982–underscoring the need to allow for time varying volatility.[3] Third, risk-averse bond investors are concerned not only with the most likely outcomes but also with the degree of uncertainty surrounding future bond returns, indicating the need to model the full probability distribution of bond returns.

The literature on bond return predictability has noted the importance of parameter estimation error, model instability, and model uncertainty. However, no study on bond return

[1] According to the Securities Industry and Financial Markets Association, the size of the U.S. Treasury bond market was $11.9 trillion in 2013Q4. This is almost 30% of the entire U.S. bond market which includes corporate debt, mortgage and municipal bonds, money market instruments, agency and asset-backed securities.

[2] Stock and Watson (1999) and Cogley and Sargent (2002) find strong evidence of time variation in a Phillips curve model for U.S. inflation.

[3] Sims and Zha (2006) and Cogley et al. (2010) find that it is important to account for time varying volatility when modeling the dynamics of U.S. macroeconomic variables.

predictability has so far addressed how these considerations, jointly, impact the results. To accomplish this, we propose a novel Bayesian approach that brings several advantages to inference about the return prediction models and to their use in portfolio allocation analysis.

Our approach allows us, first, to integrate out uncertainty about the unknown parameters and to evaluate the effect of estimation error on the results. Even after observing 50 years of monthly observations, the coefficients of the return prediction models are surrounded by considerable uncertainty and so accounting for estimation error turns out to be important. Indeed, we find many cases with strong improvements in forecasting performance as a result of incorporating estimation error.[4]

Second, our approach produces predictive densities of bond excess returns. This allows us to analyze the economic value of bond return predictability from the perspective of an investor with power utility. Thornton and Valente (2012) are limited to considering mean-variance utility since they model only the first two moments of bond returns.[5]

Third, we allow for time varying (stochastic) volatility in the bond excess return model. Bond market volatility spiked during the monetarist experiment from 1979 to 1982, but we find clear advantages from allowing for stochastic volatility beyond this episode, particularly for bonds with shorter maturities. Stochastic volatility models do not, in general, lead to better point forecasts of bond returns but they produce far better density forecasts which, when used by a risk averse investor, generate better economic performance. In addition to lowering risk during periods with spikes in volatility, the stochastic volatility models imply that investors load more heavily on risky bonds during times with relatively low interest rate volatility such as during the 1990s.

A fourth advantage of our approach is that it allows for time variation in the regression parameters. We find evidence that the slope coefficients on both the yield spreads and the macrofactors vary considerably during our sample and that accounting for time varying parameters generally leads to more accurate forecasts and better economic performance. For example, whereas a univariate constant-coefficient model based on the Cochrane and Piazzesi (2005) forward rate factor does not improve forecasting performance relative to a constant expected return model, allowing the parameters of this model to vary over time leads to significantly better results.

Fifth, we address model uncertainty and model instability through forecast combination methods. Model uncertainty is important in our analysis which considers a variety of univariate

[4] Altavilla et al. (2014) find that an exponential tilting approach helps improve the accuracy of out-of-sample forecasts of bond yields. While their approach is not Bayesian, their tilting approach also attenuates the effect of estimation error on the model estimates.

[5] Sarno et al. (2016) use an approximate solution to compute optimal portfolio weights under power utility. They do not find evidence of economically exploitable return predictability but also do not consider parameter uncertainty.

and multivariate models as well as different model specifications. We consider equal-weighted averages of predictive densities, Bayesian model averaging, as well as combinations based on the optimal pooling method of Geweke and Amisano (2011). The latter forms a portfolio of the individual prediction models using weights that reflect the models' posterior probabilities. Models that are more strongly supported by the data get a larger weight in this average, so our combinations accommodate shifts in the relative forecasting performance of different models. The model combination results are generally better than the results for the individual models and thus suggest that model uncertainty–and model instability–can be effectively addressed through combination methods.[6]

Our empirical analysis uses the daily treasury yield data from Gurkaynak et al. (2007) to construct monthly excess returns for bond maturities between two and five years over the period 1962-2011. While previous studies have focused on the annual holding period, focusing on the higher frequency affords several advantages. Most obviously, it expands the number of non-overlapping observations, a point of considerable importance given the impact of parameter estimation error. Moreover, it allows us to identify short-lived dynamics in both first and second moments of bond returns which are missed by models of annual returns. This is an important consideration during events such as the global financial crisis of 2007-09 and around turning points of the economic cycle.

We conduct our analysis in the context of a three-variable model that includes the Fama-Bliss forward spread, the Cochrane-Piazzesi linear combination of forward rates, and a macro factor constructed using the methodology of Ludvigson and Ng (2009). Each variable is weighted according to its ability to improve on the predictive power of the bond return equation. Since forecasting studies have found that simpler models often do well in out-of-sample experiments, we also consider simpler univariate models.[7]

To assess the statistical evidence on bond return predictability, we use our models to generate out-of-sample forecasts over the period 1990-2011. Our return forecasts are based on recursively updated parameter estimates and use only historically available information, thus allowing us to assess how valuable the forecasts would have been to investors in real time. Compared to the benchmark EH model that assumes no return predictability, we find that many of the return predictability models generate significantly positive out-of-sample R^2 values. Interestingly, the Bayesian return prediction models generally perform better than the least squares counterparts so far explored in the literature.

Turning to the economic value of such out-of-sample forecasts, we next consider the portfolio

[6]Using an iterated combination approach, Lin et al. (2016) uncovers statistical and economic predictability in *corporate* bond returns

[7]Other studies considering macroeconomic determinants of the term structure of interest rates include Ang and Piazzesi (2003), Ang et al. (2007), Bikbov and Chernov (2010), Dewachter et al. (2014), Duffee (2011) and Joslin et al. (2014).

choice between a risk-free Treasury bill versus a bond with 2-5 years maturity for an investor with power utility. We find that the best return prediction models that account for volatility dynamics and changing parameters deliver sizeable gains in certainty equivalent returns relative to an EH model that assumes no predictability of bond returns, particularly in the absence of tight constraints on the portfolio weights.

There are several reasons why our findings differ from studies such as Thornton and Valente (2012) and Sarno et al. (2016) which argue that the statistical evidence on bond return predictability does not translate into economic gains. Allowing for stochastic volatility and time varying parameters, while accounting for parameter estimation error, leads to important gains in economic performance for many models.[8] Our results on forecast combinations also emphasize the importance of accounting for model uncertainty and the ability to capture changes in the performance of individual prediction models.

To interpret the economic sources of our findings on bond return predictability, we analyze the extent to which such predictability is concentrated in certain economic states and whether it is correlated with variables expected to be key drivers of time varying bond risk premia. We find that bond return predictability is stronger in recessions than during expansions, consistent with similar findings for stock returns by Henkel et al. (2011) and Dangl and Halling (2012). Using data from survey expectations we find that, consistent with a risk-premium story, our bond excess return forecasts are strongly negatively correlated with economic growth prospects (thus being higher during recessions) and strongly positively correlated with inflation uncertainty.

Our finding that the macro factor of Ludvigson and Ng (2009) possesses considerable predictive power over bond excess returns implies that information embedded in the yield curve does not subsume information contained in such macro variables. We address possible explanations of this finding, including the unspanned risk factor models of Joslin et al. (2014) and Duffee (2011) which suggest that macro variables move forecasts of future bond excess returns and forecasts of future short rates by the same magnitude but in opposite directions. We find support for this explanation as our bond excess return forecasts are strongly negatively correlated with survey forecasts of future short rates.

The outline of the paper is as follows. Section 2 describes the construction of the bond data, including bond returns, forward rates and the predictor variables. Section 3 sets up the prediction models and introduces our Bayesian estimation approach. Section 4 presents both full-sample and out-of-sample empirical results on bond return predictability. Section 5 assesses the economic value of bond return predictability for a risk averse investor when this investor uses the bond return predictions to form a portfolio of risky bonds and a risk-free asset. Section 6 analyzes economic sources of bond return predictability such as recession risk, time variations

[8] Thornton and Valente (2012) use a rolling window to update their parameter estimates but do not have a formal model that predicts future volatility or parameter values.

in inflation uncertainty, and the presence of unspanned risk factors. Section 7 presents model combination results and Section 8 concludes.

2 Data

This section describes how we construct our monthly series of bond returns and introduces the predictor variables used in the bond return models.

2.1 Returns and Forward Rates

Previous studies on bond return predictability such as Cochrane and Piazzesi (2005), Ludvigson and Ng (2009) and Thornton and Valente (2012) use overlapping 12-month returns data. This overlap induces strong serial correlation in the regression residuals. To handle this issue, we reconstruct the yield curve at the daily frequency starting from the parameters estimated by Gurkaynak et al. (2007), who rely on methods developed in Nelson and Siegel (1987) and Svensson (1994). Specifically, the time t zero coupon log yield on a bond maturing in n years, $y_t^{(n)}$, gets computed as[9]

$$\begin{aligned} y_t^{(n)} &= \beta_{0t} + \beta_{1t}\frac{1-\exp\left(-\frac{n}{\tau_1}\right)}{\frac{n}{\tau_1}} + \beta_{2t}\left[\frac{1-\exp\left(-\frac{n}{\tau_1}\right)}{\frac{n}{\tau_1}} - \exp\left(-\frac{n}{\tau_1}\right)\right] \\ &\quad + \beta_{3t}\left[\frac{1-\exp\left(-\frac{n}{\tau_2}\right)}{\frac{n}{\tau_2}} - \exp\left(-\frac{n}{\tau_2}\right)\right]. \end{aligned} \tag{1}$$

The parameters $(\beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2)$ are provided by Gurkaynak et al. (2007), who report daily estimates of the yield curve from June 1961 onward for the entire maturity range spanned by outstanding Treasury securities. We consider maturities ranging from 12 to 60 months and, in what follows, focus on the last day of each month's estimated log yields.[10]

Denote the frequency at which returns are computed by h, so $h = 1, 3$ for the monthly and quarterly frequencies, respectively. Also, let n be the bond maturity in years. For $n > h/12$ we compute returns and excess returns, relative to the $h-$period T-bill rate[11]

$$r_{t+h/12}^{(n)} = p_{t+h/12}^{(n-h/12)} - p_t^{(n)} = ny_t^{(n)} - (n-h/12)y_{t+h/12}^{(n-h/12)}, \tag{2}$$

$$rx_{t+h/12}^{(n)} = r_{t+h/12}^{(n)} - y_t^{h/12}(h/12). \tag{3}$$

[9] The third term was excluded from the calculations prior to January 1, 1980.

[10] The data is available at http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html. Because of idiosyncrasies at the very short end of the yield curve, we do not compute yields for maturities less than twelve months. For estimation purposes, the Gurkaynak et al. (2007) curve drops all bills and coupon bearing securities with a remaining time to maturity less than 6 months, while downweighting securities that are close to this window. The coefficients of the yield curve are estimated using daily cross-sections and thus avoid introducing look-ahead biases in the estimated yields.

[11] The formulas assume that the yields have been annualized, so we multiply $y_t^{(h/12)}$ by $h/12$.

Here $p_t^{(n)}$ is the logarithm of the time t price of a bond with n periods to maturity. Similarly, forward rates are computed as[12]

$$f_t^{(n-h/12,n)} = p_t^{(n-h/12)} - p_t^{(n)} = ny_t^{(n)} - (n-h/12)y_t^{(n-h/12)}. \tag{4}$$

2.2 Data Summary

Our bond excess return data span the period 1962:01-2011:12. We focus our analysis on the monthly holding period which offers several advantages over the annual returns data which have been the focus of most studies in the literature on bond return predictability. Most obviously, using monthly rather than annual data provides a sizeable increase in the number of (non-overlapping) data points available for model estimation. This is important in light of the low power of the return prediction models. Second, some of the most dramatic swings in bond prices occur over short periods of time lasting less than a year–e.g., the effect of the bankruptcy of Lehman Brothers on September 15, 2008–and are easily missed by models focusing on the annual holding period. This point is also important for the analysis of how return predictability is linked to recessions versus expansions; bond returns recorded at the annual horizon easily overlook important variations around turning points of the economic cycle.

Figure 1 plots monthly bond returns for the 2, 3, 4, and 5-year maturities, computed in excess of the 1-month T-bill rate. All four series are notably more volatile during 1979-82 and the volatility clearly increases with the maturity of the bonds. Panel A.1 in Table 1 presents summary statistics for the four monthly excess return series. Returns on the shortest maturities are right-skewed and fat-tailed, more so than the longer maturities. This observation suggests that it is inappropriate to use models that assume a normal distribution for bond returns.

Because the data used in our study differ from those used in most existing studies, it is worth highlighting the main differences and showing how they affect our data. Whereas most studies such as Cochrane and Piazzesi (2005) use overlapping 12-month bond returns computed in excess of the yield on a 12-month Treasury bill, we use monthly (non-overlapping) bond returns computed in excess of the 1-month T-bill rate.

Panels A.2 and A.3 in Table 1 provide summary statistics on the more conventional overlapping 12-month returns constructed either from our monthly data (Panel A.2) or as in Cochrane and Piazzesi (2005) (Panel A.3), using the Fama-Bliss CRSP files. The two series have very similar means which in turn are lower than the mean excess return on the monthly series in Panel A.1 due to the lower mean of the risk-free rate (1-month T-bill) used in Panel A.1 compared to the mean of the 12-month T-bill rate used in Panels A.2 and A.3. Comparing the monthly series in Panel A.1 to the 12-month series in Panels A.2 and A.3, we see that the kurtosis in the monthly excess returns data is much greater than in the 12-month data, while the serial

[12] For $n = h/12$, $f_t^{(n,n)} = ny_t^{(n)}$ and $y_t^{(n-h/12)} = y_t^{(0)}$ equals zero because $P_t^{(0)} = 1$ and its logarithm is zero.

correlation is much stronger in the 12-month series. Both findings are unsurprising because of the overlap in the 12-month return data which creates a smoother return series.

2.3 Predictor variables

Our empirical strategy entails regressing bond excess returns on a range of the most prominent predictors proposed in the literature on bond return predictability. Specifically, we consider forward spreads as proposed by Fama and Bliss (1987), a linear combination of forward rates as proposed by Cochrane and Piazzesi (2005), and a linear combination of macro factors, as proposed by Ludvigson and Ng (2009).

To motivate our use of these three predictor variables, note that the n-period bond yield is related to expected future short yields and expected future excess returns (Duffee, 2013):

$$y_t^{(n)} = \frac{1}{n}\left(\sum_{j=0}^{n-1} E[y_{t+j}^{(1)}|\boldsymbol{z}_t]\right) + \frac{1}{n}\left(\sum_{j=0}^{n-1} E_t[rx_{t+j+1}^{(n-j)}|\boldsymbol{z}_t]\right), \tag{5}$$

where $rx_{t+j+1}^{(n-j)}$ is the excess return in period $t+j+1$ on a bond with $n-j$ periods to maturity and $E[.|\boldsymbol{z}_t]$ denotes the conditional expectation given market information at time t, $\boldsymbol{z}_t$. Equation (5) suggests that current yields or, equivalently, forward spreads should have predictive power over future bond excess returns and so motivates our use of these variables in the excess return regressions.

The use of non-yield predictors is more contentious. In fact, if the vector of conditioning information variables, $\boldsymbol{z}_t$, is of sufficiently low dimension, we can invert (5) to get $\boldsymbol{z}_t = g(\boldsymbol{y}_t)$. In this case information in the current yield curve subsumes all other predictors of future excess returns and so macro variables should be irrelevant when added to the prediction model. The unspanned risk factor models of Joslin et al. (2014) and Duffee (2011) offer an explanation for why macro variables help predict bond excess returns over and above information contained in the yield curve. These models suggest that the effect of additional state variable on expected future short rates and expected future bond excess returns cancel out in Equation (5). Such cancellations imply that the additional state variables do not show up in bond yields although they can have predictive power over bond excess returns.

Our predictor variables are computed as follows. The Fama-Bliss (FB) forward spreads are given by

$$fs_t^{(n,h)} = f_t^{(n-h/12,n)} - y_t^{(h/12)}(h/12). \tag{6}$$

The Cochrane-Piazzesi (CP) factor is given as a linear combination of forward rates

$$CP_t^h = \hat{\boldsymbol{\gamma}}^{h\prime}\boldsymbol{f}_t^{(\boldsymbol{n}-h/12,\boldsymbol{n})}, \tag{7}$$

where

$$\boldsymbol{f}_t^{(\boldsymbol{n}-h/12,\boldsymbol{n})} = \left[f_t^{(n_1-h/12,n_1)}, f_t^{(n_2-h/12,n_2)}, ..., f_t^{(n_k-h/12,n_k)}\right].$$

Here $\boldsymbol{n} = [1, 2, 3, 4, 5]$ denotes the vector of maturities measured in years. As in Cochrane and Piazzesi (2005), the coefficient vector $\hat{\boldsymbol{\gamma}}$ is estimated from

$$\frac{1}{4}\sum_{n=2}^{5} rx_{t+h/12}^{(n)} = \gamma_0^h + \gamma_1^h f_t^{(1-1/12,1)} + \gamma_2^h f_t^{(2-1/12,2)} + \gamma_3^h f_t^{(3-1/12,3)} + \gamma_4^h f_t^{(4-1/12,4)} + \gamma_5^h f_t^{(5-1/12,5)} + \overline{\varepsilon}_{t+h/12}. \tag{8}$$

Ludvigson and Ng (2009) propose to use macro factors to predict bond returns. Suppose we observe a $T \times M$ panel of macroeconomic variables $\{x_{i,t}\}$ generated by a factor model

$$x_{i,t} = \kappa_i g_t + \epsilon_{i,t}, \tag{9}$$

where g_t is an $s \times 1$ vector of common factors and $s << M$. The unobserved common factor, g_t is replaced by an estimate, $\hat{g}_t$, obtained using principal components analysis. Following Ludvigson and Ng (2009), we build a single linear combination from a subset of the first eight estimated principal components, $\widehat{\boldsymbol{G}}_t = [\hat{g}_{1,t}, \hat{g}_{1,t}^3, \hat{g}_{3,t}, \hat{g}_{4,t}, \hat{g}_{8,t}]$ to obtain the LN factor[13]

$$LN_t^h = \hat{\boldsymbol{\lambda}}^{h\prime} \widehat{\boldsymbol{G}}_t, \tag{10}$$

where $\hat{\boldsymbol{\lambda}}$ is obtained from the projection

$$\frac{1}{4}\sum_{n=2}^{5} rx_{t+h/12}^{(n)} = \lambda_0^h + \lambda_1^h \hat{g}_{1,t} + \lambda_2^h \hat{g}_{1,t}^3 + \lambda_3^h \hat{g}_{3,t} + \lambda_4^h \hat{g}_{4,t} + \lambda_5^h \hat{g}_{8,t} + \overline{\eta}_{t+h/12}. \tag{11}$$

Panel B in Table 1 presents summary statistics for the Fama-Bliss forward spreads along with the CP and LN factors. The Fama-Bliss forward spreads are strongly positively autocorrelated with first-order autocorrelation coefficients around 0.90. The CP and LN factors are far less autocorrelated with first-order autocorrelations of 0.67 and 0.41, respectively.

Panel C shows that the Fama-Bliss spreads are strongly positively correlated. In turn, these spreads are positively correlated with the CP factor, with correlations around 0.5, but are uncorrelated with the LN factor. The LN factor captures a largely orthogonal component in relation to the other predictors. For example, its correlation with CP is only 0.18. It is also less persistent than the FB and CP factors.

3 Return Prediction Models and Estimation Methods

We next introduce the return prediction models and describe the estimation methods used in the paper.

[13] Ludvigson and Ng (2009) selected this particular combination of factors using the Schwarz information criterion.

3.1 Model specifications

Our analysis considers the three predictor variables described in the previous section. Specifically, we consider three univariate models, each of which includes one of these three factors, along with a multivariate model that includes all three predictors. This produces a total of four models:

1. Fama-Bliss (FB) univariate

$$rx_{t+h/12}^{(n)} = \beta_0 + \beta_1 fs_t^{(n,h)} + \varepsilon_{t+h/12}. \tag{12}$$

2. Cochrane-Piazzesi (CP) univariate

$$rx_{t+h/12}^{(n)} = \beta_0 + \beta_1 CP_t^h + \varepsilon_{t+h/12}. \tag{13}$$

3. Ludvigson-Ng (LN) univariate

$$rx_{t+h/12}^{(n)} = \beta_0 + \beta_1 LN_t^h + \varepsilon_{t+h/12}. \tag{14}$$

4. Fama-Bliss, Cochrane-Piazzesi and Ludvigson-Ng predictors (FB-CP-LN)

$$rx_{t+h/12}^{(n)} = \beta_0 + \beta_1 fs_t^{(n,h)} + \beta_2 CP_t^h + \beta_3 LN_t^h + \varepsilon_{t+h/12}. \tag{15}$$

These models are in turn compared against the Expectation Hypothesis benchmark

$$rx_{t+h/12}^{(n)} = \beta_0 + \varepsilon_{t+h/12}, \tag{16}$$

that assumes no predictability. In each case $n \in \{2, 3, 4, 5\}$.

We consider four classes of models: (i) constant coefficient models with constant volatility; (ii) constant coefficient models with stochastic volatility; (iii) time varying parameter models with constant volatility; and (iv) time varying parameter models with stochastic volatility.

The constant coefficient, constant volatility model serves as a natural starting point for the out-of-sample analysis. There is no guarantee that the more complicated models with stochastic volatility and time varying regression coefficients produce better out-of-sample forecasts since their parameters may be imprecisely estimated.

To estimate the models we adopt a Bayesian approach that offers several advantages over the conventional estimation methods adopted by previous studies on bond return predictability. First, imprecisely estimated parameters is a big issue in the return predictability literature and so it is important to account for parameter uncertainty as is explicitly done by the Bayesian approach. Second, portfolio allocation analysis requires estimating not only the conditional

mean, but also the conditional variance (under mean-variance preferences) or the full predictive density (under power utility) of returns. This is accomplished by our method which generates the (posterior) predictive return distribution. Third, our approach also allows us to handle model uncertainty (and model instability) by combining forecasting models.

We next describe our estimation approach for each of the four classes of models. To ease the notation, for the remainder of the paper we drop the notation $t + h/12$ and replace $h/12$ with 1, with the understanding that the definition of a period depends on the data frequency.

3.2 Constant Coefficients and Constant Volatility Model

The linear model projects bond excess returns $rx_{\tau+1}^{(n)}$ on a set of lagged predictors, $\boldsymbol{x}_{\tau}^{(n)}$:

$$\begin{aligned} rx_{\tau+1}^{(n)} &= \mu + \boldsymbol{\beta}'\boldsymbol{x}_{\tau}^{(n)} + \varepsilon_{\tau+1}, \quad \tau = 1, ..., t-1, \\ \varepsilon_{\tau+1} &\sim \mathcal{N}(0, \sigma_{\varepsilon}^2). \end{aligned} \tag{17}$$

Ordinary least squares (OLS) estimation of this model is straightforward and so is not further explained. However, we also consider Bayesian estimation so we briefly describe how the prior and likelihood are specified. Following standard practice, the priors for the parameters μ and $\boldsymbol{\beta}$ in (17) are assumed to be normal and independent of σ_{ε}^2

$$\begin{bmatrix} \mu \\ \boldsymbol{\beta} \end{bmatrix} \sim \mathcal{N}\left(\underline{\boldsymbol{b}}, \underline{\boldsymbol{V}}\right), \tag{18}$$

where

$$\underline{\boldsymbol{b}} = \begin{bmatrix} \overline{rx}_t^{(n)} \\ \mathbf{0} \end{bmatrix}, \ \underline{\boldsymbol{V}} = \underline{\psi}^2 \left[\left(s_{rx,t}^{(n)}\right)^2 \left(\sum_{\tau=1}^{t-1} \boldsymbol{x}_{\tau}^{(n)} \boldsymbol{x}_{\tau}^{(n)\prime} \right)^{-1} \right], \tag{19}$$

and $\overline{rx}_t^{(n)}$ and $\left(s_{rx,t}^{(n)}\right)^2$ are data-based moments:

$$\begin{aligned} \overline{rx}_t^{(n)} &= \frac{1}{t-1} \sum_{\tau=1}^{t-1} rx_{\tau+1}^{(n)}, \\ \left(s_{rx,t}^{(n)}\right)^2 &= \frac{1}{t-2} \sum_{\tau=1}^{t-1} \left(rx_{\tau+1}^{(n)} - \overline{rx}_t^{(n)}\right)^2. \end{aligned}$$

Our choice of the prior mean vector $\underline{\boldsymbol{b}}$ reflects the "no predictability" view that the best predictor of bond excess returns is the average of past returns. We therefore center the prior intercept on the prevailing mean of historical excess returns, while the prior slope coefficient is centered on zero.

It is common to base the priors of the hyperparameters on sample estimates, see Stock and Watson (2006) and Efron (2010). Our analysis can thus be viewed as an empirical Bayes approach rather than a more traditional Bayesian approach that fixes the prior distribution

before any data are observed. We find that, at least for a reasonable range of values, the choice of priors has modest impact on our results. In (19), $\underline{\psi}$ is a constant that controls the tightness of the prior, with $\underline{\psi} \to \infty$ corresponding to a diffuse prior on μ and $\boldsymbol{\beta}$. Our benchmark analysis sets $\underline{\psi} = n/2$. This choice means that the prior becomes looser for the longer bond maturities for which fundamentals-based information is likely to be more important. It also means that the posterior parameter estimates are shrunk more towards their priors for the shortest maturities which are most strongly affected by estimation error.

We assume a standard gamma prior for the error precision of the return innovation, $\sigma_{\varepsilon}^{-2}$:

$$\sigma_{\varepsilon}^{-2} \sim \mathcal{G}\left(\left(s_{rx,t}^{(n)}\right)^{-2}, \underline{v}_0 (t-1)\right), \tag{20}$$

where $\underline{v}_0$ is a prior hyperparameter that controls how informative the prior is with $\underline{v}_0 \to 0$ corresponding to a diffuse prior on $\sigma_{\varepsilon}^{-2}$. Our baseline analysis sets $\underline{v}_0 = 2/n$, again letting the priors be more diffuse, the longer the bond maturity.

3.3 Stochastic Volatility Model

A large literature has found strong empirical evidence of time varying return volatility, see Andersen et al. (2006). We accommodate such effects through a simple stochastic volatility (SV) model:

$$rx_{\tau+1}^{(n)} = \mu + \boldsymbol{\beta}' \boldsymbol{x}_{\tau}^{(n)} + \exp(h_{\tau+1})\, u_{\tau+1}, \tag{21}$$

where $h_{\tau+1}$ denotes the (log of) bond return volatility at time $\tau+1$ and $u_{\tau+1} \sim \mathcal{N}(0,1)$. The log-volatility $h_{\tau+1}$ is assumed to follow a stationary and mean reverting process:

$$h_{\tau+1} = \lambda_0 + \lambda_1 h_{\tau} + \xi_{\tau+1}, \tag{22}$$

where $\xi_{\tau+1} \sim \mathcal{N}\left(0, \sigma_{\xi}^2\right)$, $|\lambda_1| < 1$, and u_{τ} and ξ_s are mutually independent for all τ and s. Appendix A explains how we estimate the SV model and set the priors.

3.4 Time varying Parameter Model

Studies such as Thornton and Valente (2012) find considerable evidence of instability in the parameters of bond return prediction models. The following time varying parameter (TVP) model allows the regression coefficients in (17) to change over time:

$$\begin{aligned} rx_{\tau+1}^{(n)} &= (\mu + \mu_{\tau}) + (\boldsymbol{\beta} + \boldsymbol{\beta}_{\tau})' \boldsymbol{x}_{\tau}^{(n)} + \varepsilon_{\tau+1}, \quad \tau = 1, ..., t-1, \\ \varepsilon_{\tau+1} &\sim \mathcal{N}(0, \sigma_{\varepsilon}^2). \end{aligned} \tag{23}$$

The intercept and slope parameters $\boldsymbol{\theta}_{\tau} = \left(\mu_{\tau}, \boldsymbol{\beta}_{\tau}'\right)'$ are assumed to follow a zero-mean, stationary process

$$\boldsymbol{\theta}_{\tau+1} = diag\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}\right) \boldsymbol{\theta}_{\tau} + \boldsymbol{\eta}_{\tau+1}, \tag{24}$$

where $\boldsymbol{\theta}_1 = \mathbf{0}$, $\boldsymbol{\eta}_{\tau+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$, and the elements in $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$ are restricted to lie between -1 and 1. In addition, ε_τ and $\boldsymbol{\eta}_s$ are mutually independent for all τ and s.[14] The key parameter is $\boldsymbol{Q}$ which determines how rapidly the parameters $\boldsymbol{\theta}$ are allowed to change over time. We set the priors to ensure that the parameters are allowed to change only gradually. Again Appendix A provides details on how we estimate the model and set the priors.

3.5 Time varying Parameter, Stochastic Volatility Model

Finally, we consider a general model that admits both time varying parameters and stochastic volatility (TVP-SV):

$$rx_{\tau+1}^{(n)} = (\mu + \mu_\tau) + (\boldsymbol{\beta} + \boldsymbol{\beta}_\tau)' \boldsymbol{x}_\tau^{(n)} + \exp(h_{\tau+1})\, u_{\tau+1}, \tag{25}$$

with

$$\boldsymbol{\theta}_{\tau+1} = diag(\boldsymbol{\gamma}_{\boldsymbol{\theta}})\, \boldsymbol{\theta}_\tau + \boldsymbol{\eta}_{\tau+1}, \tag{26}$$

where again $\boldsymbol{\theta}_\tau = (\mu_\tau, \boldsymbol{\beta}_\tau')'$, and

$$h_{\tau+1} = \lambda_0 + \lambda_1 h_\tau + \xi_{\tau+1}. \tag{27}$$

We assume that $u_{\tau+1} \sim \mathcal{N}(0,1)$, $\boldsymbol{\eta}_{\tau+1} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{Q})$, $\xi_{\tau+1} \sim \mathcal{N}\left(0, \sigma_\xi^2\right)$ and u_τ, $\boldsymbol{\eta}_s$ and ξ_l are mutually independent for all τ, s, and l. Again we refer to Appendix A for further details on this model.

The models are estimated by Gibbs sampling methods. This allows us to generate draws of excess returns, $rx_{t+1}^{(n)}$, in a way that only conditions on a given model and the data at hand. This is convenient when computing bond return forecasts and determining the optimal bond holdings.

4 Empirical Results

This section describes our empirical results. For comparison with the existing literature, and to convey results on the importance of different features of the models, we first report results based on full-sample estimates. This is followed by an out-of-sample analysis of both the statistical and economic evidence on return predictability.

4.1 Full-sample Estimates

For comparison with extant results, Table 2 presents full-sample (1962:01-2011:12) least squares estimates for the bond return prediction models with constant parameters. While no investors could have based their historical portfolio choices on these estimates, such results are important

[14] This is equivalent to writing $rx_{\tau+1}^{(n)} = \widetilde{\mu}_\tau + \widetilde{\boldsymbol{\beta}}_\tau' \boldsymbol{x}_\tau^{(n)} + \varepsilon_{\tau+1}$, where $\widetilde{\boldsymbol{\theta}}_1 \equiv \left(\widetilde{\mu}_1, \widetilde{\boldsymbol{\beta}}_1'\right)$ is left unrestricted.

for our understanding of how the various models work. The slope coefficients for the univariate models increase monotonically in the maturity of the bonds. With the exception of the coefficients on the CP factor in the multivariate model, the coefficients are significant across all maturities and forecasting models.[15]

Table 2 shows R^2 values around 1-2% for the model that uses FB as a predictor, 2.5% for the model that uses the CP factor and around 5% for the model based on the LN factor. These values, which increase to 7-8% for the multivariate model, are notably smaller than those conventionally reported for the overlapping 12-month horizon. For comparison, at the one-year horizon we obtain R^2 values of 10-11%, 17-24%, and 14-19% for the FB, CP, and LN models, respectively.[16]

The extent of time variations in the parameters of the three-factor FB-CP-LN model is displayed in Figure 2. All coefficients are notably volatile around 1980 and the coefficients continue to fluctuate throughout the sample, particularly for the CP regressor.

An advantage of our approach is its ability to deal with parameter estimation error. To get a sense of the importance of this issue, Figure 3 plots full-sample posterior densities of the regression coefficients for the three-factor model that uses the FB, CP and LN factors as predictors. The spread of the densities in this figure shows the considerable uncertainty surrounding the parameter estimates even at the end of the sample. As expected, parameter uncertainty is greatest for the TVP and TVP-SV models which allow for the greatest amount of flexibility–clearly this comes at the cost of less precisely estimated parameters. The SV model generates more precise estimates than the constant volatility benchmark, reflecting the ability of the SV model to reduce the weight on observations in highly volatile periods.

The effect of such parameter uncertainty on the predictive density of bond excess returns is depicted in Figure 4. This figure evaluates the univariate LN model at the mean of this predictor, plus or minus two times its standard deviation. The TVP and TVP-SV models imply a greater dispersion for bond returns and their densities shift further out in the tails as the predictor variable moves away from its mean. The four models clearly imply very different probability distributions for bond returns and so have very different implications when used by investors to form portfolios.

Figure 5 plots the time series of the posterior means and volatilities of bond excess returns for the FB-CP-LN model. Mean excess returns (top panel) vary substantially during the sample,

[15] As emphasized by Cochrane and Piazzesi (2005), care has to be exercised when evaluating the statistical significance of these results due to the highly persistent FB and CP regressors. Wei and Wright (2013) find that conventional tests applied to bond excess return regressions that use yield spreads or yields as predictors are subject to considerable finite-sample distortions. However, their reverse regression approach confirms that, even after accounting for such biases, bond excess returns still appear to be predictable.

[16] These values are a bit lower than those reported in the literature but are consistent with the range of results reported by Duffee (2013). The weaker evidence reflects our use of an extended sample along with a tendency for the regression coefficients to decline towards zero at the end of the sample.

peaking during the early eighties, and again during 2008. Stochastic volatility effects (bottom panel) also appear to be empirically important. The conditional volatility is very high during 1979-1982, while subsequent spells with above-average volatility are more muted and short-lived. Interestingly, there are relatively long spells with below-average conditional volatility such as during the late nineties and mid-2000s.

4.2 Calculation of out-of-sample Forecasts

To gauge the real-time value of the bond return prediction models, following Ludvigson and Ng (2009) and Thornton and Valente (2012), we next conduct an out-of-sample forecasting experiment.[17] This experiment relies on information up to period t to compute return forecasts for period $t+1$ and uses an expanding estimation window. Notably, when constructing the CP and LN factors we also restrict our information to end at time t. Hence, we re-estimate each period the principal components and the regression coefficients in equations (8) and (11).

We use 1962:01-1989:12 as our initial warm-up estimation sample and 1990:01-2011:12 as the forecast evaluation period. As before, we set $n = 2, 3, 4, 5$ and so predict 2, 3, 4, and 5-year bond returns in excess of the one-month T-bill rate.

The predictive accuracy of the bond excess return forecasts is measured relative to recursively updated forecasts from the expectations hypothesis (EH) model (16) that projects excess returns on a constant. Specifically, at each point in time we obtain draws from the predictive densities of the benchmark model and the models with time varying predictors. For a given bond maturity, n, we denote draws from the predictive density of the EH model, given the information set at time t, $\mathcal{D}^t = \{rx_{\tau+1}^{(n)}\}_{\tau=1}^{t-1}$, by $\left\{rx_{t+1}^{(n),j}\right\}, j = 1, ..., J$. Similarly, draws from the predictive density of any of the other models (labeled model i) given $\mathcal{D}^t = \{rx_{\tau+1}^{(n)}, \boldsymbol{x}_\tau^{(n)}\}_{\tau=1}^{t-1} \cup \boldsymbol{x}_t^{(n)}$ are denoted $\left\{rx_{t+1,i}^{(n),j}\right\}, j = 1, ..., J$.[18]

For the constant parameter, constant volatility model, return draws are obtained by applying a Gibbs sampler to

$$p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right) = \int p\left(rx_{t+1}^{(n)}\middle|\,\mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \mathcal{D}^t\right) p\left(\mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}\middle|\,\mathcal{D}^t\right) d\mu d\boldsymbol{\beta} d\sigma_\varepsilon^{-2}. \tag{28}$$

[17] Out-of-sample analysis also provides a way to guard against overfitting. Duffee (2010) shows that in-sample overfitting can generate unrealistically high Sharpe ratios.

[18] We run the Gibbs sampling algorithms recursively for all time periods betweeen 1990:01 and 2011:12. At each point in time, we retain 1,000 draws from the Gibbs samplers after a burn-in period of 500 iterations. For the TVP, SV, and TVP-SV models we run the Gibbs samplers five times longer while at the same time thinning the chains by keeping only one in every five draws, thus effectively eliminating any autocorrelation left in the draws. Additional details on these algorithms are presented in Appendix A.

Return draws for the most general TVP-SV model are obtained from the predictive density[19]

$$\begin{aligned} p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right) &= \int p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},h_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right) \\ &\quad\times p\left(\boldsymbol{\theta}_{t+1},h_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right) \\ &\quad\times p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\middle|\,\mathcal{D}^t\right) d\mu d\boldsymbol{\beta} d\boldsymbol{\theta}^{t+1} d\boldsymbol{\gamma_\theta} d\boldsymbol{Q} dh^{t+1} d\lambda_0 d\lambda_1 d\sigma_\xi^{-2}. \end{aligned} \tag{29}$$

where $h^{t+1}=(h_1,...,h_{t+1})$ and $\boldsymbol{\theta}^{t+1}=(\boldsymbol{\theta}_1,...,\boldsymbol{\theta}_{t+1})$ denote the sequence of conditional variance states and time varying regression parameters up to time $t+1$, respectively. Draws from the SV and TVP models are obtained as special cases of (29). All Bayesian models integrate out uncertainty about the parameters.

4.3 Forecasting Performance

Although our models generate a full predictive distribution for bond returns it is insightful to also report results based on conventional point forecasts. To obtain point forecasts we first compute the posterior mean from the densities in (28) and (29). We denote these by $\overline{rx}_{t,EH}^{(n)}=\frac{1}{J}\sum_{j=1}^{J}rx_t^{(n),j}$ and $\overline{rx}_{t,i}^{(n)}=\frac{1}{J}\sum_{j=1}^{J}rx_{t,i}^{(n),j}$, for the EH and alternative models, respectively. Using such point forecasts, we obtain the corresponding forecast errors as $e_{t,EH}^{(n)}=rx_t^{(n)}-\overline{rx}_{t,EH}^{(n)}$ and $e_{t,i}^{(n)}=rx_t^{(n)}-\overline{rx}_{t,i}^{(n)}$, $t=\underline{t},...,\bar{t}$, where $\underline{t}=1990:01$ and $\bar{t}=2011:12$ denote the beginning and end of the forecast evaluation period.

Following Campbell and Thompson (2008), we compute the out-of-sample R^2 of model i relative to the EH model as

$$R_{OoS,i}^{(n)2}=1-\frac{\sum_{\tau=\underline{t}}^{\bar{t}}e_{\tau,i}^{(n)2}}{\sum_{\tau=\underline{t}}^{\bar{t}}e_{\tau,EH}^{(n)2}}. \tag{30}$$

Positive values of this statistic suggest evidence of time varying return predictability.

Table 3 reports R_{OoS}^2 values for the OLS, linear, SV, TVP and TVP-SV models across the four bond maturities. For the two-year maturity we find little evidence that models estimated by OLS are able to improve on the predictive accuracy of the EH model. Conversely, all models estimated using our Bayesian approach generate significantly more accurate forecasts at either the 10% or 1% significance levels, using the test for equal predictive accuracy suggested by Clark and West (2007). Similar results are obtained for the SV and TVP models which generate R_{OoS}^2 values close to or above 5% for the models that include the LN predictor. Interestingly, the best results are obtained for the TVP-SV model which notably outperforms the other univariate models based on the FB or CP variables.

While the OLS models fare considerably better for the 3-5 year bond maturities, the ability of the linear Bayesian model to generate accurate forecasts does not appear to depend as

[19]For each draw retained from the Gibbs sampler, we produce 100 draws from the corresponding predictive densities.

strongly on the maturity. Moreover, the Bayesian approach performs notably better than its OLS counterpart, particularly for the multivariate model which requires estimating more parameters which increases the importance of parameter estimation error.

Comparing results across predictor variables, the univariate CP model performs worst with only the TVP and TVP-SV specifications suggesting some (weak) evidence of improvements in the predictive accuracy over the EH benchmark. This suggests that the CP predictor has some predictive power over bond excess returns but that its coefficient is unstable over time, consistent with the third plot in Figure 2. Conversely, the FB and, in particular, the LN predictor, add considerable improvements in out-of-sample predictive performance.

Ranking the different specifications across bond maturities and predictor variables, we find that the TVP-SV models produce the best out-of-sample forecasts in around half of all cases with the SV model a distant second best. These results suggest that the more sophisticated models that allow for time varying parameters and time varying volatility manage to produce better out-of-sample forecasts than simple constant parameter, constant volatility models. Even in cases where the TVP-SV model is not the best specification, it still performs nearly as well as the best model. In contrast, there are instances where the other models are highly inferior to the TVP-SV model.

To identify which periods the models perform best, following Welch and Goyal (2008), we use the out-of-sample forecast errors to compute the difference in the cumulative sum of squared errors (SSE) for the EH model versus the ith model:

$$\Delta CumSSE_{t,i}^{(n)} = \sum_{\tau=\underline{t}}^{t} \left(e_{\tau,EH}^{(n)}\right)^2 - \sum_{\tau=\underline{t}}^{t} \left(e_{\tau,i}^{(n)}\right)^2. \tag{31}$$

Positive and increasing values of $\Delta CumSSE_t$ suggest that the model with time varying return predictability generates more accurate point forecasts than the EH benchmark.

Figure 6 plots $\Delta CumSSE_t$ for the three univariate models and the three-factor model, assuming a two-year bond maturity. These plots show periods during which the various models perform well relative to the EH model–periods where the lines are increasing and above zero–and periods where the models underperform against this benchmark–periods with decreasing graphs. The univariate FB model performs quite poorly due to spells of poor performance in 1994, 2000, and 2008, while the CP model underperforms between 1993 and 2006. In contrast, except for a few isolated months in 2002, 2008 and 2009, the LN model consistently beats the EH benchmark up to 2010, at which point its performance flattens against the EH model. A similar performance is seen for the multivariate model.

The predictive accuracy measures in (30) and (31) ignore information on the full probability distribution of returns. To evaluate the accuracy of the density forecasts obtained in (28) and (29), we compute the implied predictive likelihoods. The predictive likelihood, or score, is

commonly viewed as the broadest measure of accuracy of density forecasts, see, e.g., Geweke and Amisano (2010). At each point in time t, the log predictive score is obtained by taking the natural log of the predictive densities (28)–(29) evaluated at the observed bond excess return, $rx_t^{(n)}$, denoted by $LS_{t,EH}$ and $LS_{t,i}$ for the EH and alternative models, respectively.

Table 4 reports the average log-score differential for each of our models, again measured relative to the EH benchmark.[20] The results show that the linear model performs significantly better than the EH benchmark across all specifications for the four- and five-year bond maturities and for the models containing the LN predictor for the two- and three-year bond maturities. The TVP model produces similar, if modest, improvements in the log-score of the EH model. Notably stronger results are obtained for the SV and TVP-SV models which dominate the EH benchmark across the board.

Figure 7 supplements Table 4 by showing the cumulative log score (LS) differentials between the EH model and the ith model, computed analogously to (31) as

$$\Delta CumLS_{t,i} = \sum_{\tau=\underline{t}}^{t} \left[LS_{\tau,i} - LS_{\tau} \right]. \tag{32}$$

The dominant performance of the density forecasts generated by the SV and TVP-SV models is clear from these plots. In contrast, the linear and TVP models offer only modest improvements over the EH benchmark by this measure.

4.4 Robustness to Choice of Priors

Choice of priors can always be debated in Bayesian analysis, so we conduct a sensitivity analysis with regard to two of the priors, namely $\underline{\psi}$ and $\underline{v}_0$, which together control how informative the baseline priors are. Our first experiment sets $\underline{\psi} = 5$ and $\underline{v}_0 = 1/5$. This choice corresponds to using more diffuse priors than in the baseline scenario. Compared with the baseline prior, this prior produces worse results (lower out-of-sample R^2 values) for the two shortest maturities ($n = 2, 3$), but stronger results for the longest maturities ($n = 4, 5$).

Our second experiment sets $\underline{\psi} = 0.5, \underline{v}_0 = 5$, corresponding to tighter priors. Under these priors, the results improve for the shorter bond maturities but get weaker at the longest maturities. In both cases, the conclusion that the best prediction models dominate the EH benchmark continues to hold even for such large shifts in priors. Detailed results are available on request.

[20] To test if the differences in forecast accuracy are significant, we follow Clark and Ravazzolo (2015) and apply the Diebold and Mariano (1995) t-test for equality of the average log-scores based on the statistic $\overline{LS}_i = \frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} (LS_{\tau.i} - LS_{\tau,EH})$. The p-values for this statistic are based on t-statistics computed with a serial correlation-robust variance, using the pre-whitened quadratic spectral estimator of Andrews and Monahan (1992). Monte Carlo evidence in Clark and McCracken (2011) indicates that, with nested models, the Diebold-Mariano test compared against normal critical values can be viewed as a somewhat conservative test for equal predictive accuracy in finite samples. Since all models considered here nest the EH benchmark, we report p-values based on one-sided tests, taking the nested EH benchmark as the null and the nesting model as the alternative.

5 Economic Value of Return Forecasts

So far our analysis concentrated on statistical measures of predictive accuracy. It is important to evaluate the extent to which the apparent gains in predictive accuracy translate into better investment performance. In fact, for an investor with mean-variance preferences, Thornton and Valente (2012) find that improvements in the statistical accuracy of bond return forecasts do not imply improved portfolio performance.

5.1 Bond Holdings

We consider the asset allocation decisions of an investor that selects the weight, $\omega_t^{(n)}$, on a risky bond with n periods to maturity versus a one-month T-bill that pays the riskfree rate, $\tilde{y}_t = y_t^{(1/12)}$. The investor has power utility and coefficient of relative risk aversion A:

$$U\left(\omega_t^{(n)}, rx_{t+1}^{(n)}\right) = \frac{\left[\left(1-\omega_t^{(n)}\right)\exp\left(\tilde{y}_t\right) + \omega_t^{(n)}\exp\left(\tilde{y}_t + rx_{t+1}^{(n)}\right)\right]^{1-A}}{1-A}, \quad A > 0. \tag{33}$$

Using all information at time t, $\mathcal{D}^t$, to evaluate the predictive density of $rx_{t+1}^{(n)}$, the investor solves the optimal asset allocation problem

$$\omega_t^{(n)*} = \arg\max_{\omega_t^{(n)}} \int U\left(\omega_t^{(n)}, rx_{t+1}^{(n)}\right) p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right) drx_{t+1}^{(n)}. \tag{34}$$

The integral in (34) can be approximated by generating a large number of draws, $rx_{t+1,i}^{(n),j}$, $j = 1, .., J$, from the predictive densities specified in (28) and (29). For each of the candidate models, i, we approximate the solution to (34) by

$$\widehat{\omega}_{t,i}^{(n)} = \arg\max_{\omega_{t,i}^{(n)}} \frac{1}{J}\sum_{j=1}^{J}\left\{\frac{\left[\left(1-\omega_{t,i}^{(n)}\right)\exp\left(\tilde{y}_t\right) + \omega_{t,i}^{(n)}\exp\left(\tilde{y}_t + rx_{t+1,i}^{(n),j}\right)\right]^{1-A}}{1-A}\right\}. \tag{35}$$

The resulting sequences of portfolio weights $\left\{\widehat{\omega}_{t,EH}^{(n)}\right\}$ and $\left\{\widehat{\omega}_{t,i}^{(n)}\right\}$ are used to compute realized utilities. For each model, i, we convert these into certainty equivalent returns (CER), i.e., values that equate the average utility of the EH model with the average utility of any of the alternative models.

We consider two different sets of assumptions about the portfolio weights. The first scenario restricts the weights on the risky bonds to the interval $[0, 0.99]$ to ensure that the expected utility is finite even with an unbounded return distribution; see Geweke (2001) and Kandel and Stambaugh (1996) for a discussion of this point. The second scenario instead restricts the weights to lie between -2 and 3 but bounds the monthly return distribution between -100% and 100%. This follows the procedure in Johannes et al. (2014) who argue that this approach ensures that the expected utility exists.

In both cases we set the coefficient of relative risk aversion to $A = 10$, a value higher than normally considered. Our choice reflects the high Sharpe ratios observed for the bond portfolios during our sample—see Table 1. For lower values of A, this causes the weights on the risky bonds to almost always hit the upper bound (0.99) of the first scenario for both the EH and time varying predictability models and so does not allow us to differentiate between these models. However, provided that we do not impose too tight limits on the bond portfolio weights (Scenario 2), informative results can still be obtained for lower values of A, e.g., $A = 5$.

5.2 Empirical Results

Table 5 shows annualized CER values computed relative to the EH model so positive values indicate that the time varying predictability models perform better than the EH model. For the scenario with the weights constrained to $[0, 0.99]$ (Panels A-D), the CER values generally increase with the bond maturity. In 11 of 16 cases, the highest CER values are found for the three-factor TVP-SV models. For these models the CER values increase from around 0.24% ($n = 2$) to 0.82% ($n = 3$) and 1.2%-1.6% for the longest bond maturity. To test if the annualized CER values are statistically greater than zero we use a Diebold-Mariano test.[21] With the notable exception of the two-year bond maturity, most of the CER values for the SV and TVP-SV models are significantly higher than those generated by the EH benchmark.

Turning to the case with portfolio weights constrained to lie in the [-2, 3] interval (Panels E-H in Table 5), the CER values generally increase substantially, notably for the two-year bond maturity. For example, for the TVP-SV model with three predictors the CER increases from 0.24% to 2.95% (two-year bond) and from 1.55% to 1.96% (five-year maturity). In general, all models improve strongly when loosening the constraint imposed on the portfolio weights in Panels A-D.

Comparing results across the different specifications, the SV and TVP-SV models generally perform best. Their improvements over the benchmark model are particularly large under the less constrained weights where we find CER gains above 200 basis points per year for the models that include the LN predictor

Figure 8 plots cumulative CER values, computed relative to the EH benchmark, for the three-factor model assuming $\hat{\omega}_t^{(n)} \in [-2, 3]$. These graphs parallel the cumulated sum of squared error difference plots in (31), the key difference being that they show the cumulated risk-adjusted

[21] Specifically, we estimate the regression $u_{t+1.i}^{(n)} - u_{t+1,EH}^{(n)} = \alpha^{(n)} + \epsilon_{t+1}$ where

$$u_{t+1,i}^{(n)} = \frac{1}{1-A}\left[\left(1-\omega_{t,i}^{(n)}\right)\exp\left(\tilde{y}_t\right) + \omega_{t,i}^{(n)}\exp\left(\tilde{y}_t + rx_{t+1}^{(n)}\right)\right]^{1-A},$$

and

$$u_{t+1,EH}^{(n)} = \frac{1}{1-A}\left[\left(1-\omega_{t,EH}^{(n)}\right)\exp\left(\tilde{y}_t\right) + \omega_{t,EH}^{(n)}\exp\left(\tilde{y}_t + rx_{t+1}^{(n)}\right)\right]^{1-A},$$

and test if $\alpha^{(n)}$ equals zero.

gains from using a particular model instead of the EH model. Interestingly, the models' best performance is concentrated from 1990 to 1993 and from 2001 onwards. Across all bond maturities The cumulated CER value at the end of the sample exceeds 50 percent for all models.

The CER values reported in Table 5 ignore transaction costs. However, when we allow for transaction costs we continue to see sizeable gains over the EH benchmark. For example, assuming a one-way transaction cost of 10 basis points, the CER value for the strategy that uses the TVP-SV model to predict bond excess returns is reduced from 2.95% to 1.80% for the two-year bond and from 1.96% to 1.48% for the five-year bond.

We conclude from these results that there is strong statistical and economic evidence that the returns on 2-5 year bonds can be predicted using predictor variables proposed in the literature. Moreover, the best performing models do not assume constant parameters but allow for time varying mean and volatility dynamics.

5.3 Comparison with Other Studies

Our results are very different from those reported by Thornton and Valente (2012). These authors find that statistical evidence of out-of-sample return predictability fails to translate into an ability for investors to use return forecasts in a way that generates higher out-of-sample average utility than forecasts from the EH model. Notably, we find that incorporating time varying parameters and stochastic volatility in many cases improves bond portfolio performance. Instead, Thornton and Valente (2012) find that the Sharpe ratios of their bond portfolios decrease when accounting for such effects through rolling window estimation.

Besides differences in modeling approaches, a reason for such differences is the focus of Thornton and Valente (2012) on 12-month bond returns, whereas we use monthly bond returns. To address the importance of the return horizon, we repeat the out-of-sample analysis using non-overlapping quarterly and annual returns data. Compared with the monthly results, the quarterly and annual R^2_{OoS} values decline somewhat. At the quarterly horizon the univariate FB and LN models, along with the bivariate FB-LN model, continue to perform well across the four bond maturities. The LN and FB-LN models also perform well at the annual horizon, particularly for the bonds with shorter maturities ($n = 2, 3$). The associated CER values continue to be positive and, in most cases, significant at the quarterly horizon, but are substantially smaller at the annual horizon. These findings indicate a fast moving predictable component in bond returns that is missed when using longer return horizons and so help explain the difference between our results and those of Thornton and Valente (2012) and Dewachter et al. (2014).

The setup of Sarno et al. (2016) is closest to that adopted here as they also consider results for one-month returns and still obtain negative economic values from using their time varying bond return forecasts compared with the EH model. Such differences in results reflect (i) different modeling assumptions: Sarno et al. (2016) compute expected excess returns in the context

of an affine term structure model and also do not consider stochastic volatility or time varying parameters; (ii) different predictor variables: Sarno et al. (2016) use latent state variables extracted from their term structure model to predict bond excess returns; and, (iii) different estimation methodologies: Sarno et al. (2016) do not follow the same Bayesian methodology that we use here and thus ignore parameter uncertainty.

Duffee (2013) expresses concerns related to data mining when interpreting results for macro predictors whose effects are not underpinned by theory. A particular concern is that the strong results for the LN factor are sample specific. The sample used by Ludvigson and Ng ends in 2003:12. One way to address this concern is by inspecting the performance of the three-factor model in the subsequent sample, i.e., from 2004:01 to 2011:12. Figure 7 and Figure 8 show that the prediction models continue to generate higher CER values and log-density scores than the EH benchmark after 2003, although Figure 5 suggests weaker R^2 performance for the two- and three-year bond maturities. These results show that the predictive power of the LN factor is not limited to the original sample used to construct this variable.

6 Economic Drivers of Bond Return Predictability

To address the economic sources of our results, this section first analyzes variations in bond return predictability across the economic cycle. Next, we explore whether our bond return forecasts are correlated with drivers of time varying bond risk premia. Finally, we discuss whether unspanned risk factors help explain the predictive power of the LN macro factor.

6.1 Cyclical Variations in Bond Return Predictability

Recent studies such as Rapach et al. (2010), Henkel et al. (2011) and Dangl and Halling (2012) report that predictability of stock returns is concentrated in economic recessions and is largely absent during expansions. This finding is important since it suggests that return predictability is linked to cyclical variations and that time varying risk premia may be important drivers of expected returns.

To see if bond return predictability varies over the economic cycle, we split the data into recession and expansion periods using the NBER recession indicator which equals one in recessions and zero in expansions. Table 6 uses full-sample parameter estimates, but computes R^2 values separately for the recession and expansion samples. We use full-sample information because there are only three recessions in our out-of-sample period, 1990-2011.

Table 6 shows that the R^2 values are generally higher during recessions than in expansions. This finding is consistent with the findings for stock market returns as indicated by the earlier references. Moreover, it is robust across model specifications and predictor variables, the only exception being the univariate FB model for which return predictability actually is stronger

during expansions. Conversely, note that the R^2 values are particularly high in recessions for the TVP models that include the LN variable.

To test if the differences in R^2 values are statistically significant, we conduct a simple bootstrap test that exploits the monotonic relationship between the mean squared prediction error (MSE) of the forecasting model, measured relative to that of the EH model, and the R^2 measure in (30). Specifically, we test the null that the predictive accuracy of a given prediction model (measured relative to the EH benchmark) is the same across recessions and expansions, against the one-sided alternative that the relative MSE is higher in expansions,

$$\begin{aligned} H_0: &\quad E[\underbrace{e^2_{EH,0} - e^2_{i,0}}_{\Delta_0}] = E[\underbrace{e^2_{EH,1} - e^2_{i,1}}_{\Delta_1}] \\ H_1: &\quad E[e^2_{EH,0} - e^2_{i,0}] < E[e^2_{EH,1} - e^2_{i,1}]. \end{aligned} \tag{36}$$

Here e_{EH} and e_i are the forecast errors under the EH and model i, respectively, and the subscript refers to expansions (0) and recessions (1). By computing a particular model's MSE relative to the MSE of the EH model in the same state we control for differences in bond return variances in recessions versus expansions. Our test uses a bootstrap based on the frequency with which $\Delta_0 - \Delta_1$ is smaller than 10,000 counterparts bootstrapped under the null of $\Delta_0 = \Delta_1$.[22]

Outcomes from this test are indicated by stars in the recession columns of Table 6. With the notable exception of the univariate FB model, we find that not only is the fit of the bond return prediction models better in recessions than in expansions, but this difference is highly statistically significant in most cases.

Large differences between bond return predictability in recessions and expansions are also observed in the out-of-sample period 1990-2011. However, in this case we do not have a large enough number of recessions for the test of equal predictive power to have sufficient power to reject the null hypothesis in (36).[23]

We also compute results that split the sample into recessions and expansions using the unemployment gap recession indicator of Stock and Watson (2010).[24] This indicator is computable in real time and so is arguably more relevant than the NBER indicator which gets released with

[22] The p-value for the test is computed as follows: i) impose the null of equal-predictability across states i.e., compute $\hat{\Delta}_0 = \Delta_0 - \hat{\mu}(\Delta_0)$ and $\hat{\Delta}_1 = \Delta_1 - \hat{\mu}(\Delta_1)$; ii) estimate the distribution under the null by using an i.i.d. bootstrap, to generate B bootstrap samples from $\hat{\Delta}_0$ and $\hat{\Delta}_1$ and for each of these compute $J^b = \mu(\hat{\Delta}^b_0) - \mu(\hat{\Delta}^b_1)$; iii) compute p-values as $p_{val} = \frac{1}{B}\sum_{b=1}^{B} 1[J > J^b]$ where $J = \mu(\Delta_0) - \mu(\Delta_1)$ is based on the data.

[23] Engsted et al. (2013) find that bond return predictability is stronger during expansions than during recessions, concluding that return predictability displays opposite patterns in the bond and stock markets. However, they use returns on a 20-year Treasury bond obtained from Ibbotson International. As we have seen, bond return predictability strongly depends on the bond maturity and so this is likely to explain the difference between their results and ours.

[24] This measure is based on the difference between the current unemployment rate and a three-year moving average of past unemployment rates.

a considerable lag. Using this alternative measure of recessions we continue to find that return predictability tends to be stronger in recessions than in expansions.

6.2 Time varying Risk Premia

Asset pricing models such as Campbell and Cochrane (1999) suggest that the Sharpe ratio on risky assets should be higher during recessions due to higher consumption volatility and a lower surplus consumption ratio. To see if this implication is consistent with our models, Table 7 reports Sharpe ratios for the bond portfolios computed separately for recession and expansion periods. Following authors such as Henkel et al. (2011) these results are again based on the full sample to ensure enough observations in recessions. With exception of the univariate FB regressions, the Sharpe ratios are substantially higher during recessions than in expansions.

To further analyze risk-based explanations of bond return predictability, we explore two important sources of bond risk premia. First, some asset pricing models suggest that investors' risk aversion should vary countercyclically, being higher around recessions and lower in expansions. To the extent that our forecasts of bond excess returns reflect time varying risk premia, we should therefore expect a negative correlation between economic growth and bond return forecasts. Second, inflation risk is likely to be an important determinant of bond return dynamics, see, e.g., Wright (2011) and Abrahams et al. (2013) and we should expect to find a positive correlation between inflation uncertainty and bond return forecasts.

Table 8 tests this implication. The table reports the contemporaneous correlations between forecasts of two-year bond excess returns and current real GDP growth (Panel A), inflation (Panel B), real GDP growth uncertainty (Panel C) and inflation uncertainty (Panel D). Real GDP growth is computed as $\Delta \log(GDP_{t+1})$, where GDP_{t+1} is the real gross domestic product, while inflation is computed as $\Delta \log(CPI_{t+1})$, where CPI is the consumer price index for all urban consumers. We measure real GDP growth and inflation uncertainty using the cross-sectional dispersion (the difference between the 75th percentile and the 25th percentile) in real GDP and CPI one-quarter-ahead forecasts, respectively, as reported by the Survey of Professional Forecasters (SPF). An advantage of this measure is that it affords a model-free approach.

The correlations between out-of-sample bond excess return forecasts, on the one hand, and current GDP growth or inflation, on the other, are negative and, in most cases, highly significant. Thus lower economic growth and reduced inflation appear to be associated with expectations of higher bond excess returns.

Turning to the uncertainty measures, we find a strongly positive and, in most cases, highly significant correlation between uncertainty about economic growth and future inflation, on the one hand, and expected bond excess returns on the other. This is consistent with our bond excess return forecasts being driven, at least in part, by time varying inflation risk premia.

Correlations are particularly strong for the models that include the LN macro factor which can be expected to be particularly sensitive to the economic cycle.

6.3 Unspanned Macro Factors

Many studies use only information in the yield curve to predict bond excess returns so our finding that the LN macro factor improves such forecasts may seem puzzling. However, as discussed by Duffee (2013), a possible explanation is that the macro variables are hidden or unspanned risk factors which do not show up in the yield curve because their effect on expected future bond excess returns and expected future short rates work in opposite directions and so tend to cancel out in (5). To see if this possibility holds up, we use SPF survey forecasts of the future 3-month T-bill rate as a proxy for the first term on the right hand side in (5). Consistent with the unspanned risk factor story, Table 9 shows that our forecasts of bond excess returns and survey expectations of future treasury yields are strongly negatively correlated with three-factor (FB-CP-LN) R^2 values ranging from 0.14 to 0.24 for $n = 2$ years and from 0.40 to 0.44 for $n = 5$ years.

The final part of our analysis estimates bond risk premia by fitting an affine term structure model to the cross section of bond yields. In Appendix B we explain how we use the approach of Joslin et al. (2011) and Wright (2011) to fit term structure models with unspanned macro risks to compute bond risk premia. We use the resulting risk premium estimates to regress, for a given maturity n, the corresponding mean bond excess returns, $\overline{rx}_t$, on a constant and the corresponding risk premium estimates $\overline{rp}_t^{(n)}$,

$$\overline{rx}_t^{(n)} = \mu + \beta \overline{rp}_t^{(n)} + u_t, \tag{37}$$

where $\overline{rx}_t^{(n)}$ denotes the predicted bond excess-return and $\overline{rp}_t^{(n)}$ denotes the risk premium estimate. Table 10 reports the estimated coefficient β along with its t-statistics for the FB-CP-LN model. For all specifications, the estimated β coefficient has the right sign (positive) and it is statistically significant for the SV and TVP-SV models fitted to the two shortest bond maturities.

7 Model Combinations

In addition to parameter uncertainty, investors face model uncertainty along with the possibility that the best model may change over time, i.e., model instability. This raises the question whether, in real time, investors could have selected forecasting models that would have generated accurate forecasts. Model uncertainty would not be a concern if all prediction models produced improvements over the EH benchmark. However, as we have seen in the empirical analysis, there is a great deal of heterogeneity across the models' predictive performance. To address this issue,

we turn to model combination. Model combinations form portfolios of individual prediction models with weights reflecting the models' historical performance. The better a model's fit relative to its complexity, the larger its weight. Similar to diversification benefits obtained for asset portfolios, model combination tends to stabilize forecasts relative to forecasts generated by individual return prediction models.

Recent model combination approaches such as Bayesian model averaging and the optimal prediction pool of Geweke and Amisano (2011) allow the weights on individual forecasting models to reflect their predictive accuracy. Such combination schemes can therefore accommodate time variations in the relative performance of different models. This matters if the importance of features such as time varying parameters and stochastic volatility dynamics changes over time.

A final reason for our interest in model combinations is that studies on predictability of stock returns such as Rapach et al. (2010), Dangl and Halling (2012), and Pettenuzzo et al. (2014) find that combinations improve on the average performance of the individual models. This result has only been established for stock returns, however.

To see if it carries over to bond returns, we consider three different combination schemes applied to all possible models obtained by combining the FB, CP and LN predictors, estimated using the linear, SV, TVP and TVP-SV approaches.

7.1 Combination Schemes

We begin by considering the equal-weighted pool (EW) which weighs each of the N models, M_i, equally

$$p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right)=\frac{1}{N}\sum_{i=1}^{N}p\left(rx_{t+1}^{(n)}\middle|\,M_i,\mathcal{D}^t\right), \tag{38}$$

where $\left\{p\left(rx_{t+1}^{(n)}\middle|\,M_i,\mathcal{D}^t\right)\right\}_{i=1}^{N}$ denotes the predictive densities specified in (28) and (29). This approach does not allow the weights on different models to change over time as a result of differences in predictive accuracy.

We also consider Bayesian model averaging (BMA) weights:

$$p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right)=\sum_{i=1}^{N}\Pr\left(M_i|\,\mathcal{D}^t\right)p\left(rx_{t+1}^{(n)}\middle|\,M_i,\mathcal{D}^t\right). \tag{39}$$

Here $\Pr\left(M_i|\,\mathcal{D}^t\right)$ denotes the posterior probability of model i, relative to all models under consideration, computed using information available at time t, $\mathcal{D}^t$. This is given by

$$\Pr\left(M_i|\,\mathcal{D}^t\right)=\frac{\Pr\left(\mathcal{D}^t\middle|\,M_i\right)\Pr\left(M_i\right)}{\sum_{j=1}^{N}\Pr\left(\mathcal{D}^t|\,M_j\right)\Pr\left(M_j\right)}. \tag{40}$$

$\Pr\left(\mathcal{D}^t \middle| M_i\right)$ and $\Pr\left(M_i\right)$ denote the marginal likelihood and prior probability for model i, respectively. We assume that all models are equally likely a priori and so set $\Pr\left(M_i\right) = 1/N$.[25]

A limitation of the BMA approach is that it assumes that the true prediction model is contained in the set of models under consideration. One approach that does not require this assumption is the optimal predictive pool (OW) proposed by Geweke and Amisano (2011). This approach again computes a weighted average of the predictive densities:

$$p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right) = \sum_{i=1}^{N} w_{t,i}^* \times p\left(rx_{t+1}^{(n)} \middle| M_i, \mathcal{D}^t\right). \tag{41}$$

The $(N \times 1)$ vector of model weights $\boldsymbol{w}_t^* = \left[w_{t,1}^*, ..., w_{t,N}^*\right]$ is determined by recursively solving the following maximization problem

$$\boldsymbol{w}_t^* = \arg\max_{\boldsymbol{w}_t} \sum_{\tau=1}^{t-1} \log\left[\sum_{i=1}^{N} w_{ti} \times S_{\tau+1,i}\right], \tag{42}$$

where $S_{\tau+1,i} = \exp\left(LS_{\tau+1,i}\right)$ is the recursively computed log-score for model i at time $\tau+1$, and $\boldsymbol{w}_t^* \in [0,1]^N$. As $t \to \infty$ the weights in (42) minimize the Kullback-Leibler distance between the combined predictive density and the data generating process, see Hall and Mitchell (2007).

By recursively updating the combination weights in (39) and (42), these combination methods accommodate changes in the relative performance of the different models. This is empirically important as we shall see.

7.2 Empirical Findings

Table 11 presents statistical and economic measures of out-of-sample forecasting performance for the three combination schemes. The optimal prediction pool generates R^2_{OoS} values at or above 5% regardless of the bond maturity while the R^2_{OoS} values range between 3 and 5% for the EW combination scheme (Panel A). The BMA combination scheme performs better than the EW combination but worse than the optimal pool. In all cases, the forecast combinations perform better–considerably so in the case of the optimal pool–than what one would expect from simply selecting a model at random.

The predictive likelihood tests shown in Panel B of Table 11 strongly reject the null of equal predictive accuracy relative to the EH model. Finally, the CER values with constrained weights (Panel C) are quite similar for the three combination schemes, rising from about 0.2% for $n = 2$ to 1%–1.5% for $n = 5$. These CER values are similar to those achieved by the best of the individual models reported in Table 5 and, for bond maturities of three years or longer, are

[25] We follow Geweke and Amisano (2010) and compute the marginal likelihoods by cumulating the predictive log scores of each model over time after conditioning on the initial warm-up estimation sample $\Pr\left(\{rx_{\tau+1}^{(n)}\}_{\tau=1}^{t-1} \middle| M_i\right) = \exp\left(\sum_{\tau=\underline{t}}^{t} LS_{\tau,i}\right)$.

significantly higher than those obtained by an EH investor, suggesting that model combination can be used to effectively deal with model uncertainty. The CER values for the case with the [-2,3] weights (Panel D) are very high, falling between 1.43% and 3.27%.

Figure 9 shows the evolution over time in the weights on the individual predictors in the optimal prediction pool. The FB and LN predictors are assigned close to full weights throughout the sample, whereas the CP predictor only gets assigned modest weight during short spells of time.

8 Conclusion

We analyze predictability of excess returns on Treasury bonds with maturities ranging from two through five years. As predictors we use the forward spread variable of Fama and Bliss (1987), the Cochrane and Piazzesi (2005) combination of forward rates, and the Ludvigson and Ng (2009) macro factors. Our analysis allows for time varying regression parameters and stochastic volatility dynamics and accounts for both parameter estimation error and model uncertainty. Using a flexible setup turns out to be important as we find significant statistical and economic gains over the constant coefficient, constant volatility models generally adopted in the existing literature.

Our findings suggest that there is evidence of both statistically and economically significant predictability in bond excess returns. This contrasts with the findings of Thornton and Valente (2012) who conclude that the statistical evidence on bond return predictability fails to translate into economic return predictability. We find that such differences can be attributed to the importance of modeling stochastic volatility and time-varying parameters in the monthly bond excess return series that we study and to rigorously accounting for estimation error.

We link the evidence on return predictability to the economic cycle, finding that the degree of return predictability is significantly higher during recessions. Our bond return forecasts are strongly positively correlated with inflation uncertainty and negatively correlated with economic growth, suggesting that time varying risk premia are an important driver of the results. Consistent with unspanned risk factor models, forecasts of bond excess returns that incorporate information on macro variables are strongly negatively correlated with survey forecasts of future short term yields.

References

Abrahams, M., T. Adrian, R. K. Crump, and E. Moench (2013). Decomposing real and nominal yield curves. Federal Reserve of New York Staff Report No. 570, New York, NY.

Altavilla, C., R. Giacomini, and R. Costantini (2014). Bond returns and market expectations. *Journal of Financial Econometrics 12*(4), 708–729.

Andersen, T. G., T. Bollerslev, P. F. Christoffersen, and F. X. Diebold (2006). Volatility and correlation forecasting. Volume 1 of *Handbook of Economic Forecasting*, pp. 777 – 878. Elsevier.

Andrews, D. W. K. and J. C. Monahan (1992). An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica 60*(4), pp. 953–966.

Ang, A., S. Dong, and M. Piazzesi (2007). No-arbitrage Taylor rules. NBER Working paper 13448.

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.

Bikbov, R. and M. Chernov (2010). No-arbitrage macroeconomic determinants of the yield curve. *Journal of Econometrics 159*(1), 166 – 182.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), pp. 205–251.

Campbell, J. Y. and R. J. Shiller (1991). Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies 58*(3), 495–514.

Campbell, J. Y. and S. B. Thompson (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies 21*(4), 1509–1531.

Carter, C. K. and R. Kohn (1994). On gibbs sampling for state space models. *Biometrika 81*(3), pp. 541–553.

Clark, T. E. and M. McCracken (2011). Testing for unconditional predictive ability. In M. Clements and D. Hendry (Eds.), *Oxford Handbook of Economic Forecasting*. Oxford University Press: Oxford.

Clark, T. E. and F. Ravazzolo (2015). Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *Journal of Applied Econometrics 30*(4), 551–575.

Clark, T. E. and K. D. West (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics 138*(1), 291 – 311.

Cochrane, J. H. and M. Piazzesi (2005). Bond risk premia. *American Economic Review 95*(1), 138–160.

Cogley, T., G. E. Primiceri, and T. J. Sargent (2010). Inflation-gap persistence in the us. *American Economic Journal: Macroeconomics 2*, 43–69.

Cogley, T. and T. J. Sargent (2002). Evolving post-world war ii u.s. inflation dynamics. In B. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001*. Cambridge, U.S.: MIT Press.

Dangl, T. and M. Halling (2012). Predictive regressions with time-varying coefficients. *Journal of Financial Economics 106*(1), 157 – 181.

Del Negro, M. and G. E. Primiceri (2015). Time varying structural vector autoregressions and monetary policy: A corrigendum. *The Review of Economic Studies*, forthcoming.

Dewachter, H., L. Iania, and M. Lyrio (2014). Information in the yield curve: A macro-finance approach. *Journal of Applied Econometrics 29*(1), 42–64.

Diebold, F. X. and R. S. Mariano (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics 13*(3), 253–263.

Duffee, G. (2010). Sharpe ratios in term structure models. Wokring paper, John Hopkins University.

Duffee, G. (2013). Forecasting interest rates. In G. Elliot and A. Timmermann (Eds.), *Handbook of Economic Forecasting, volume 2*, pp. 385 – 426. North Holland.

Duffee, G. R. (2011). Information in (and not in) the term structure. *Review of Financial Studies 24*(9), 2895–2934.

Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika 89*(3), pp. 603–615.

Efron, B. (2010). *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge University Press.

Engsted, T., S. Moller, and M. Sander (2013). Bond return predictability in expansions and recessions. CREATES research paper 2013-13.

Fama, E. F. and R. R. Bliss (1987). The information in long-maturity forward rates. *The American Economic Review 77*(4), pp. 680–692.

Geweke, J. (2001). A note on some limitations of crra utility. *Economics Letters 71*(3), 341 – 345.

Geweke, J. and G. Amisano (2010). Comparing and evaluating bayesian predictive distributions of asset returns. *International Journal of Forecasting 26*(2), 216 – 230.

Geweke, J. and G. Amisano (2011). Optimal prediction pools. *Journal of Econometrics 164*(1), 130 – 141.

Gurkaynak, R. S., B. Sack, and J. H. Wright (2007). The u.s. treasury yield curve: 1961 to the present. *Journal of Monetary Economics 54*(8), 2291 – 2304.

Hall, S. G. and J. Mitchell (2007). Combining density forecasts. *International Journal of Forecasting 23*(1), 1 – 13.

Henkel, S. J., J. S. Martin, and F. Nardari (2011). Time-varying short-horizon predictability. *Journal of Financial Economics 99*(3), 560 – 580.

Johannes, M., A. Korteweg, and N. Polson (2014). Sequential learning, predictive regressions, and optimal portfolio returns. *Journal of Finance 69*(2), 611–644.

Joslin, S., M. Priebsch, and K. J. Singleton (2014). Risk premiums in dynamic term structure models with unspanned macro risks. *The Journal of Finance 69*(3), 1197–1233.

Joslin, S., K. J. Singleton, and H. Zhu (2011). A new perspective on gaussian dynamic term structure models. *Review of Financial Studies 24*(3), 926–970.

Kandel, S. and R. F. Stambaugh (1996). On the predictability of stock returns: An asset-allocation perspective. *The Journal of Finance 51*(2), pp. 385–424.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: Likelihood inference and comparison with arch models. *The Review of Economic Studies 65*(3), 361–393.

Lin, H., C. Wu, and G. Zhou (2016). Forecasting corporate bond returns with a large set of predictors: An iterated combination approach. *Working Paper*.

Ludvigson, S. C. and S. Ng (2009). Macro factors in bond risk premia. *Review of Financial Studies 22*(12), 5027–5067.

Nelson, C. R. and A. F. Siegel (1987). Parsimonious modeling of yield curves. *The Journal of Business 60*(4), pp. 473–489.

Pettenuzzo, D., A. Timmermann, and R. Valkanov (2014). Forecasting stock returns under economic constraints. *Journal of Financial Economics 114*(3), 517–553.

Rapach, D. E., J. K. Strauss, and G. Zhou (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies 23*(2), 821–862.

Sarno, L., P. Schneider, and C. Wagner (2016). The economic value of predicting bond risk premia. *Journal of Empirical Finance 37*, 247 – 267.

Sims, C. A. and T. Zha (2006). Were there regime switches in u.s. monetary policy? *American Economic Review 96*(1), 54–81.

Stock, J. H. and M. W. Watson (1999). Forecasting inflation. *Journal of Monetary Economics 44*(2), 293 – 335.

Stock, J. H. and M. W. Watson (2006). Forecasting with many predictors. Volume 1 of *Handbook of Economic Forecasting*, pp. 515 – 554. Elsevier.

Stock, J. H. and M. W. Watson (2010). Modeling inflation after the crisis. Presented at the FRB Kansas city symposium on macroeconomic policy: post-crisis and risks ahead, Jackson Hole, Wyoming.

Svensson, L. E. O. (1994). Estimating and interpreting forward interest rates: Sweden 1992-1994. IMF Working Paper No. 94/114.

Thornton, D. L. and G. Valente (2012). Out-of-sample predictions of bond excess returns and forward rates: An asset allocation perspective. *Review of Financial Studies 25*(10), 3141–3168.

Wei, M. and J. H. Wright (2013). Reverse regressions and long-horizon forecasting. *Journal of Applied Econometrics 28*(3), 353–371.

Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies 21*(4), 1455–1508.

Wright, J. H. (2011). Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review 101*(4), 1514–34.

Appendix A Bayesian estimation and predictions

This appendix explains how we obtain parameter estimates for the models described in Section 3 and shows how we use these to generate predictive densities for bond excess returns. We begin by discussing the linear regression model in (17), then turn to the SV model in (21)-(22), the TVP model in (23)-(24), and the general TVP-SV model in (25)-(27).

A.1 Constant coefficient, constant volatility model

The goal for the simple linear regression model is to obtain draws from the joint posterior distribution $p\left(\mu, \boldsymbol{\beta}, \sigma_{\varepsilon}^{-2} \middle| \mathcal{D}^t\right)$, where D^t denotes all information available up to time t. Combining the priors in (18)-(20) with the likelihood function yields the following posteriors:

$$\left[\begin{array}{c} \mu \\ \boldsymbol{\beta} \end{array}\right] \middle| \sigma_{\varepsilon}^{-2}, \mathcal{D}^t \sim \mathcal{N}\left(\overline{\boldsymbol{b}}, \overline{\boldsymbol{V}}\right), \tag{A-1}$$

and

$$\sigma_{\varepsilon}^{-2} \middle| \mu, \boldsymbol{\beta}, \mathcal{D}^t \sim \mathcal{G}\left(\overline{s}^{-2}, \overline{v}\right), \tag{A-2}$$

where

$$\begin{aligned} \overline{\boldsymbol{V}} &= \left[\underline{\boldsymbol{V}}^{-1} + \sigma_{\varepsilon}^{-2} \sum_{\tau=1}^{t-1} \boldsymbol{x}_{\tau}^{(n)} \boldsymbol{x}_{\tau}^{(n)\prime}\right]^{-1}, \\ \overline{\boldsymbol{b}} &= \overline{\boldsymbol{V}}\left[\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{b}} + \sigma_{\varepsilon}^{-2} \sum_{\tau=1}^{t-1} \boldsymbol{x}_{\tau}^{(n)} rx_{\tau+1}^{(n)}\right], \\ \overline{v} &= (1+\underline{v}_0)(t-1). \end{aligned} \tag{A-3}$$

and

$$\overline{s}^2 = \frac{\sum_{\tau=1}^{t-1}\left(rx_{\tau+1}^{(n)} - \mu - \boldsymbol{\beta}'\boldsymbol{x}_{\tau}^{(n)}\right)^2 + \left(\left(s_{rx,t}^{(n)}\right)^2 \times \underline{v}_0 (t-1)\right)}{\overline{v}}. \tag{A-4}$$

Gibbs sampling can be used to iterate back and forth between (A-1) and (A-2), yielding a series of draws for the parameter vector $\left(\mu, \boldsymbol{\beta}, \sigma_{\varepsilon}^{-2}\right)$. Draws from the predictive density $p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right)$ can then be obtained by noting that

$$p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right) = \int p\left(rx_{t+1}^{(n)} \middle| \mu, \boldsymbol{\beta}, \sigma_{\varepsilon}^{-2}, \mathcal{D}^t\right) p\left(\mu, \boldsymbol{\beta}, \sigma_{\varepsilon}^{-2} \middle| \mathcal{D}^t\right) d\mu d\boldsymbol{\beta} d\sigma_{\varepsilon}^{-2}. \tag{A-5}$$

A.2 Stochastic Volatility model[26]

The SV model requires specifying a joint prior for the sequence of log return volatilities, h^t, the parameters λ_0 and λ_1, and the error precision, σ_{ξ}^{-2}. Writing $p\left(h^t, \lambda_0, \lambda_1, \sigma_{\xi}^{-2}\right) =$

[26] See Pettenuzzo et al. (2014) for a description of a similar algorithm where the priors are modified to impose economic constraints on the model parameters.

$p\left(h^t\middle|\,\lambda_0,\lambda_1,\sigma_\xi^{-2}\right)p\left(\lambda_0,\lambda_1\right)p\left(\sigma_\xi^{-2}\right)$, it follows from (22) that

$$p\left(h^t\middle|\,\lambda_0,\lambda_1,\sigma_\xi^{-2}\right)=\prod_{\tau=1}^{t-1}p\left(h_{\tau+1}\middle|\,h_\tau,\lambda_0,\lambda_1,\sigma_\xi^{-2}\right)p\left(h_1\right), \tag{A-6}$$

with $h_{\tau+1}|\,h_\tau,\lambda_0,\lambda_1,\sigma_\xi^{-2}\sim\mathcal{N}\left(\lambda_0+\lambda_1 h_\tau,\sigma_\xi^2\right)$. Thus, to complete the prior elicitation for $p\left(h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\right)$, we only need to specify priors for h_1, the initial log volatility, λ_0, λ_1, and σ_ξ^{-2}. We choose these from the normal-gamma family as follows:

$$h_1\sim\mathcal{N}\left(\ln\left(s_{rx,t}^{(n)}\right),\underline{k}_h\right), \tag{A-7}$$

$$\begin{bmatrix}\lambda_0\\ \lambda_1\end{bmatrix}\sim\mathcal{N}\left(\begin{bmatrix}\underline{m}_{\lambda_0}\\ \underline{m}_{\lambda_1}\end{bmatrix},\begin{bmatrix}\underline{V}_{\lambda_0} & 0\\ 0 & \underline{V}_{\lambda_1}\end{bmatrix}\right),\quad \lambda_1\in(-1,1), \tag{A-8}$$

and

$$\sigma_\xi^{-2}\sim\mathcal{G}\left(1/\underline{k}_\xi,\underline{\upsilon}_\xi\left(t-1\right)\right). \tag{A-9}$$

We set $\underline{k}_\xi=0.01$ and set the remaining hyperparameters in (A-7) and (A-9) at $\underline{k}_h=10$ and $\underline{\upsilon}_\xi=1$ to imply uninformative priors, thus allowing the data to determine the degree of time variation in the return volatility. Following Clark and Ravazzolo (2015) we set the hyperparameters to $\underline{m}_{\lambda_0}=0$, $\underline{m}_{\lambda_1}=0.9$, $\underline{V}_{\lambda_0}=0.25$, and $\underline{V}_{\lambda_0}=1.0e^{-4}$. This corresponds to setting the prior mean and standard deviation of the intercept to 0 and 0.5, respectively, and represents uninformative priors on the intercept of the log volatility specification and a prior mean of the AR(1) coefficient, λ_1, of 0.9 with a standard deviation of 0.01. This is a more informative prior that matches persistent dynamics in the log volatility process.

To obtain draws from the joint posterior distribution $p\left(\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\middle|\,\mathcal{D}^t\right)$ under the SV model, we use the Gibbs sampler to draw recursively from the following four conditional distributions:

1. $p\left(h^t\middle|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$.
2. $p\left(\mu,\boldsymbol{\beta}\middle|\,h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$.
3. $p\left(\lambda_0,\lambda_1\middle|\,\mu,\boldsymbol{\beta},h^t,\sigma_\xi^{-2},\mathcal{D}^t\right)$.
4. $p\left(\sigma_\xi^{-2}\middle|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\mathcal{D}^t\right)$.

We simulate from each of these blocks as follows. Starting with $p\left(h^t\middle|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$, we employ the algorithm of Kim et al. (1998).[27] Define $rx_{\tau+1}^{(n)*}=rx_{\tau+1}^{(n)}-\mu-\beta'\boldsymbol{x}_\tau^{(n)}$ and note that $r_{\tau+1}^{(n)*}$ is observable conditional on μ,β. Next, rewrite (21) as

$$rx_{\tau+1}^{(n)*}=\exp\left(h_{\tau+1}\right)u_{\tau+1}. \tag{A-10}$$

[27] We apply the correction to the ordering of steps detailed in Del Negro and Primiceri (2015).

Squaring and taking logs on both sides of (A-10) yields a new state space system that replaces (21)-(22) with

$$rx_{\tau+1}^{(n)**} = 2h_{\tau+1} + u_{\tau+1}^{**}, \tag{A-11}$$

$$h_{\tau+1} = \lambda_0 + \lambda_1 h_\tau + \xi_{\tau+1}, \tag{A-12}$$

where $rx_{\tau+1}^{(n)**} = \ln\left[\left(rx_{\tau+1}^{(n)*}\right)^2\right]$, and $u_{\tau+1}^{**} = \ln\left(u_{\tau+1}^2\right)$, with u_τ^{**} independent of ξ_s for all τ and s. Since $u_{t+1}^{**} \sim \ln\left(\chi_1^2\right)$, we cannot resort to standard Kalman recursions and simulation algorithms such as those in Carter and Kohn (1994) or Durbin and Koopman (2002). To get around this problem, Kim et al. (1998) employ a data augmentation approach and introduce a new state variable $s_{\tau+1}$, $\tau = 1, .., t-1$, turning their focus to drawing from $p\left(h^t \middle| \mu, \boldsymbol{\beta}, \lambda_0, \lambda_1, \sigma_\xi^{-2}, s^t, \mathcal{D}^t\right)$ instead of $p\left(h^t \middle| \mu, \boldsymbol{\beta}, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right)$, where $s^t = \{s_2, ..., s_t\}$ denotes the history up to time t of the new state variable s.

The introduction of the state variable $s_{\tau+1}$ allows us to rewrite the linear non-Gaussian state space representation in (A-11)-(A-12) as a linear Gaussian state space model, making use of the following approximation,

$$u_{\tau+1}^{**} \approx \sum_{j=1}^{7} q_j \mathcal{N}\left(m_j - 1.2704, v_j^2\right), \tag{A-13}$$

where m_j, v_j^2, and q_j, $j = 1, 2, ..., 7$, are constants specified in Kim et al. (1998) and thus need not be estimated. In turn, (A-13) implies

$$u_{\tau+1}^{**} \middle| \, s_{\tau+1} = j \sim \mathcal{N}\left(m_j - 1.2704, v_j^2\right), \tag{A-14}$$

where each state has probability

$$\Pr\left(s_{\tau+1} = j\right) = q_j. \tag{A-15}$$

Draws for the sequence of states s^t can be easily obtained, noting that each of its elements can be independently drawn from the discrete density defined by

$$\Pr\left(s_{\tau+1} = j \middle| \, \mu, \boldsymbol{\beta}, \lambda_0, \lambda_1, \sigma_\xi^{-2}, h^t, \mathcal{D}^t\right) = \frac{q_i f_{\mathcal{N}}\left(rx_{\tau+1}^{(n)**} \middle| \, 2h_{\tau+1} + m_j - 1.2704, v_j^2\right)}{\sum_{l=1}^{7} q_l f_{\mathcal{N}}\left(rx_{\tau+1}^{(n)**} \middle| \, 2h_{\tau+1} + m_l - 1.2704, v_l^2\right)}. \tag{A-16}$$

for $\tau = 1, ..., t-1$ and $j = 1, ..., 7$, and where $f_{\mathcal{N}}$ denotes the kernel of a normal density. Next, conditional on s^t, we can rewrite the nonlinear state space system as follows:

$$\begin{aligned} rx_{\tau+1}^{(n)**} &= 2h_{\tau+1} + e_{\tau+1}, \\ h_{\tau+1} &= \lambda_0 + \lambda_1 h_\tau + \xi_{\tau+1}, \end{aligned} \tag{A-17}$$

where $e_{\tau+1} \sim \mathcal{N}\left(m_j - 1.2704, v_j^2\right)$ with probability $\Pr\left(s_{\tau+1} = j \middle| \, \mu, \boldsymbol{\beta}, \lambda_0, \lambda_1, \sigma_\xi^{-2}, h^t, \mathcal{D}^t\right)$. For this linear Gaussian state space system, we can use the algorithm of Carter and Kohn (1994) to draw the whole sequence of stochastic volatilities, h^t.

Moving on to $p\left(\mu,\boldsymbol{\beta}|\, h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$, conditional on h^t it is straightforward to draw μ and $\boldsymbol{\beta}$ and apply standard results. Specifically,

$$\left[\begin{array}{c}\mu\\ \boldsymbol{\beta}\end{array}\right]\Big|\, h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\sim\mathcal{N}\left(\overline{\boldsymbol{b}},\overline{\boldsymbol{V}}\right), \tag{A-18}$$

with

$$\begin{aligned}\overline{\boldsymbol{V}} &= \left\{\underline{\boldsymbol{V}}^{-1}+\sum_{\tau=1}^{t-1}\frac{1}{\exp\left(h_{\tau+1}\right)^2}\boldsymbol{x}_\tau^{(n)}\boldsymbol{x}_\tau^{(n)\prime}\right\}^{-1},\\ \overline{\boldsymbol{b}} &= \overline{\boldsymbol{V}}\left\{\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{b}}+\sum_{\tau=1}^{t-1}\frac{1}{\exp\left(h_{\tau+1}\right)^2}\boldsymbol{x}_\tau^{(n)}rx_{\tau+1}^{(n)}\right\}.\end{aligned}$$

Next, the distribution $p\left(\lambda_0,\lambda_1|\,\mu,\boldsymbol{\beta},h^t,\sigma_\xi^{-2},\mathcal{D}^t\right)$ takes the form

$$\lambda_0,\lambda_1|\,\mu,\boldsymbol{\beta},h^t,\sigma_\xi^{-2},\mathcal{D}^t\sim\mathcal{N}\left(\left[\begin{array}{c}\overline{m}_{\lambda_0}\\ \overline{m}_{\lambda 1}\end{array}\right],\overline{\mathbf{V}}_\lambda\right)\times\lambda_1\in(-1,1),$$

where

$$\overline{\mathbf{V}}_\lambda=\left\{\left[\begin{array}{cc}\underline{V}_{\lambda_0}^{-1}&0\\ 0&\underline{V}_{\lambda_1}^{-1}\end{array}\right]+\sigma_\xi^{-2}\sum_{\tau=1}^{t-1}\left[\begin{array}{c}1\\ h_\tau\end{array}\right][1,h_\tau]\right\}^{-1}, \tag{A-19}$$

and

$$\left[\begin{array}{c}\overline{m}_{\lambda_0}\\ \overline{m}_{\lambda 1}\end{array}\right]=\overline{\mathbf{V}}_\lambda\left\{\left[\begin{array}{cc}\underline{V}_{\lambda_0}^{-1}&0\\ 0&\underline{V}_{\lambda_1}^{-1}\end{array}\right]\left[\begin{array}{c}\underline{m}_{\lambda_0}\\ \underline{m}_{\lambda_1}\end{array}\right]+\sigma_\xi^{-2}\sum_{\tau=1}^{t-1}\left[\begin{array}{c}1\\ h_\tau\end{array}\right]h_{\tau+1}\right\}. \tag{A-20}$$

Finally, the posterior distribution for $p\left(\sigma_\xi^{-2}\Big|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\mathcal{D}^t\right)$ is readily available using

$$\sigma_\xi^{-2}\Big|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\mathcal{D}^t\sim\mathcal{G}\left(\left[\frac{\underline{k}_\xi\underline{v}_\xi(t-1)+\sum_{\tau=1}^{t-1}\left(h_{\tau+1}-\lambda_0-\lambda_1h_\tau\right)^2}{\left(1+\underline{v}_\xi\right)(t-1)}\right]^{-1},\left(1+\underline{v}_\xi\right)(t-1)\right). \tag{A-21}$$

Draws from the predictive density $p\left(rx_{t+1}^{(n)}\Big|\,\mathcal{D}^t\right)$ can be obtained by noting that

$$\begin{aligned}p\left(rx_{t+1}^{(n)}\Big|\,\mathcal{D}^t\right) &= \int p\left(rx_{t+1}^{(n)}\Big|\,h_{t+1},\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)\\ &\quad\times p\left(h_{t+1}|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)\\ &\quad\times p\left(\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\Big|\,\mathcal{D}^t\right)d\mu d\boldsymbol{\beta}dh^{t+1}d\lambda_0d\lambda_1d\sigma_\xi^{-2}.\end{aligned} \tag{A-22}$$

The first term in the integral above, $p\left(rx_{t+1}^{(n)}\Big|\,h_{t+1},\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$, represents the period $t+1$ predictive density of bond excess returns, treating model parameters as if they were known with certainty, and so is straightforward to calculate. The second term in the integral, $p\left(h_{t+1}|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$, reflects how period $t+1$ volatility may drift away from h_t over

time. Finally, the last term in the integral, $p\left(\mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2} \middle| \mathcal{D}^t\right)$, measures parameter uncertainty in the sample.

To obtain draws for $p\left(rx_{t+1}^{(n)} \middle| \mathcal{D}^t\right)$, we proceed in three steps:

1. Simulate from $p\left(\mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2} \middle| \mathcal{D}^t\right)$: draws from $p\left(\mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2} \middle| \mathcal{D}^t\right)$ are obtained from the Gibbs sampling algorithm described above.

2. Simulate from $p\left(h_{t+1} | \mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right)$: having processed data up to time t, the next step is to simulate the future volatility, h_{t+1}. For a given h_t and σ_ξ^{-2}, note that μ and β and the history of volatilities up to t become redundant, i.e., $p\left(h_{t+1} | \mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right) = p\left(h_{t+1} | h_t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right)$. Note also that (22) along with the distributional assumptions made on $\xi_{\tau+1}$ imply that

$$h_{t+1} | h_t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t \sim \mathcal{N}\left(\lambda_0 + \lambda_1 h_t, \sigma_\xi^2\right). \tag{A-23}$$

3. Simulate from $p\left(rx_{t+1}^{(n)} \middle| h_{t+1}, \mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right)$: For a given h_{t+1}, μ, and β, note that h^t, λ_0, λ_1, and σ_ξ^{-2} become redundant, i.e., $p\left(rx_{t+1}^{(n)} \middle| h_{t+1}, \mu, \boldsymbol{\beta}, h^t, \lambda_0, \lambda_1, \sigma_\xi^{-2}, \mathcal{D}^t\right) = p\left(rx_{t+1}^{(n)} \middle| h_{t+1}, \mu, \boldsymbol{\beta}, \mathcal{D}^t\right)$. Then use the fact that

$$rx_{t+1}^{(n)} \middle| h_{t+1}, \mu, \boldsymbol{\beta}, \mathcal{D}^t \sim \mathcal{N}\left(\mu + \boldsymbol{\beta}' \boldsymbol{x}_\tau^{(n)}, \exp\left(h_{t+1}\right)\right). \tag{A-24}$$

A.3 Time varying Parameter Model

In addition to specifying prior distributions and hyperparameters for $[\mu, \boldsymbol{\beta}]'$ and σ_ε^2, the TVP model in (23)-(24) requires eliciting a joint prior for the sequence of time varying parameters $\boldsymbol{\theta}^t = \{\boldsymbol{\theta}_2, ..., \boldsymbol{\theta}_t\}$, the parameter vector $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$, and the variance covariance matrix $\boldsymbol{Q}$. For $[\mu, \boldsymbol{\beta}]'$ and σ_ε^2, we follow the same prior choices made for the linear model:

$$\begin{bmatrix} \mu \\ \boldsymbol{\beta} \end{bmatrix} \sim \mathcal{N}\left(\underline{\boldsymbol{b}}, \underline{\boldsymbol{V}}\right), \tag{A-25}$$

and

$$\sigma_\varepsilon^{-2} \sim \mathcal{G}\left(\left(s_{rx,t}^{(n)}\right)^{-2}, \underline{v}_0 (t-1)\right). \tag{A-26}$$

Turning to $\boldsymbol{\theta}^t$, $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$, and $\boldsymbol{Q}$, we first write $p\left(\boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right) = p\left(\boldsymbol{\theta}^t \middle| \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right) p\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}\right) p\left(\boldsymbol{Q}\right)$, and note that (24) along with the assumption that $\boldsymbol{\theta}_1 = \mathbf{0}$ implies

$$p\left(\boldsymbol{\theta}^t \middle| \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right) = \prod_{\tau=1}^{t-1} p\left(\boldsymbol{\theta}_{\tau+1} | \boldsymbol{\theta}_\tau, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right), \tag{A-27}$$

with $\boldsymbol{\theta}_{\tau+1} | \boldsymbol{\theta}_\tau, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q} \sim \mathcal{N}\left(diag\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}\right) \boldsymbol{\theta}_\tau, \boldsymbol{Q}\right)$. Thus, to complete the prior elicitation for $p\left(\boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}\right)$ we need to specify priors for $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$ and $\boldsymbol{Q}$.

We choose an Inverted Wishart distribution for $\mathbf{Q}$:

$$\mathbf{Q} \sim \mathcal{IW}\left(\underline{\mathbf{Q}}, \underline{v}_{\mathbf{Q}}\left(t_0 - 1\right)\right), \tag{A-28}$$

with

$$\underline{\mathbf{Q}} = \underline{k}_{\mathbf{Q}} \underline{v}_{\mathbf{Q}}\left(t_0 - 1\right) \underline{\mathbf{V}}. \tag{A-29}$$

$\underline{k}_{\mathbf{Q}}$ controls the degree of variation in the time-varying regression coefficients $\boldsymbol{\theta}_\tau$, with larger values of $\underline{k}_{\mathbf{Q}}$ implying greater variation in $\boldsymbol{\theta}_\tau$. Our analysis sets $\underline{k}_{\mathbf{Q}} = \left(\underline{\psi}/100\right)^2$ and $\underline{v}_{\mathbf{Q}} = 10$. These are more informative priors than the earlier choices and limit the changes to the regression coefficients to be $\underline{\psi}/100$ on average.

We specify the elements of $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$ to be a priori independent of each other with generic element $\gamma^i_{\boldsymbol{\theta}}$

$$\gamma^i_{\boldsymbol{\theta}} \sim \mathcal{N}\left(\underline{m}_{\gamma_{\boldsymbol{\theta}}}, \underline{V}_{\gamma_{\boldsymbol{\theta}}}\right), \quad \gamma^i_{\boldsymbol{\theta}} \in (-1, 1), \quad i = 1, ..., k. \tag{A-30}$$

where $\underline{m}_{\gamma_{\boldsymbol{\theta}}} = 0.8$, and $\underline{V}_{\gamma_{\boldsymbol{\theta}}} = 1.0e^{-6}$, implying relatively high autocorrelations.

To obtain draws from the joint posterior distribution $p\left(\mu, \boldsymbol{\beta}, \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q} \middle| \mathcal{D}^t\right)$ under the TVP model we use the Gibbs sampler to draw recursively from the following conditional distributions:

1. $p\left(\boldsymbol{\theta}^t \middle| \mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t\right)$.
2. $p\left(\mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2} \middle| \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t\right)$.
3. $p\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}} \middle| \mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \boldsymbol{\theta}^t, \boldsymbol{Q}, \mathcal{D}^t\right)$.
4. $p\left(\boldsymbol{Q} \middle| \mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \mathcal{D}^t\right)$.

We simulate from each of these blocks as follows. Starting with θ^t, we focus on $p\left(\boldsymbol{\theta}^t \middle| \mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2}, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t\right)$. Define $rx^{(n)*}_{\tau+1} = rx^{(n)}_{\tau+1} - \mu - \beta' \boldsymbol{x}^{(n)}_\tau$ and rewrite (23) as follows:

$$rx^{(n)*}_{\tau+1} = \mu_\tau - \boldsymbol{\beta}'_\tau \boldsymbol{x}^{(n)}_\tau + \varepsilon_{\tau+1} \tag{A-31}$$

Knowledge of μ and $\boldsymbol{\beta}$ makes $rx^{(n)*}_{\tau+1}$ observable, and reduces (23) to the measurement equation of a standard linear Gaussian state space model with homoskedastic errors. Thus, the sequence of time varying parameters $\boldsymbol{\theta}^t$ can be drawn from (A-31) using the algorithm of Carter and Kohn (1994).

Moving on to $p\left(\mu, \boldsymbol{\beta}, \sigma_\varepsilon^{-2} \middle| \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t\right)$, conditional on $\boldsymbol{\theta}^t$ it is straightforward to draw $\mu, \boldsymbol{\beta}$, and σ_ε^{-2} by applying standard results. Specifically,

$$\begin{bmatrix} \mu \\ \boldsymbol{\beta} \end{bmatrix} \Bigg| \sigma_\varepsilon^{-2}, \boldsymbol{\theta}^t, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{\gamma}_{\boldsymbol{\theta}}, \boldsymbol{Q}, \mathcal{D}^t \sim \mathcal{N}\left(\overline{\boldsymbol{b}}, \overline{\boldsymbol{V}}\right), \tag{A-32}$$

and

$$\sigma_\varepsilon^{-2} \middle| \mu, \boldsymbol{\beta}, \boldsymbol{\theta}^t, \boldsymbol{Q}, \mathcal{D}^t \sim \mathcal{G}\left(\overline{s}^{-2}, \overline{v}\right), \tag{A-33}$$

where

$$\begin{aligned}\overline{\boldsymbol{V}} &= \left[\underline{\boldsymbol{V}}^{-1}+\sigma_{\varepsilon}^{-2}\sum_{\tau=1}^{t-1}\boldsymbol{x}_{\tau}^{(n)}\boldsymbol{x}_{\tau}^{(n)\prime}\right]^{-1}, \\ \overline{\boldsymbol{b}} &= \overline{\boldsymbol{V}}\left[\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{b}}+\sigma_{\varepsilon}^{-2}\sum_{\tau=1}^{t-1}\boldsymbol{x}_{\tau}^{(n)}\left(rx_{\tau+1}^{(n)}-\mu_{\tau}-\boldsymbol{\beta}_{\tau}'\boldsymbol{x}_{\tau}^{(n)}\right)\right], \end{aligned} \tag{A-34}$$

$$\overline{s}^{2}=\frac{\sum_{\tau=1}^{t-1}\left(rx_{\tau+1}^{(n)*}-\mu_{\tau}-\boldsymbol{\beta}_{\tau}'\boldsymbol{x}_{\tau}^{(n)}\right)^{2}+\left(\left(s_{rx,t}^{(n)}\right)^{2}\times\underline{v}_{0}\left(t-1\right)\right)}{\overline{v}}, \tag{A-35}$$

and $\overline{v}=\left(1+\underline{v}_{0}\right)\left(t-1\right)$.

Next, obtaining draws from $p\left(\boldsymbol{\gamma_{\theta}}|\,\mu,\boldsymbol{\beta},\sigma_{\varepsilon}^{-2},\boldsymbol{\theta}^{t},\boldsymbol{Q},\mathcal{D}^{t}\right)$ is straightforward. The $i-$th element $\gamma_{\boldsymbol{\theta}}^{i}$ is drawn from the following distribution

$$\gamma_{\boldsymbol{\theta}}^{i}|\,\mu,\boldsymbol{\beta},\sigma_{\varepsilon}^{-2},\boldsymbol{\theta}^{t},\boldsymbol{Q},\mathcal{D}^{t}\sim\mathcal{N}\left(\overline{m}_{\gamma_{\boldsymbol{\theta}}}^{i},\overline{V}_{\gamma_{\boldsymbol{\theta}}}^{i}\right)\times\gamma_{\boldsymbol{\theta}}^{i}\in(-1,1) \tag{A-36}$$

where

$$\begin{aligned}\overline{V}_{\gamma_{\boldsymbol{\theta}}}^{i} &= \left[\underline{V}_{\gamma_{\boldsymbol{\theta}}}^{-1}+\mathbf{Q}^{ii}\sum_{\tau=1}^{t-1}\left(\theta_{\tau}^{i}\right)^{2}\right]^{-1}, \\ \overline{m}_{\gamma_{\boldsymbol{\theta}}}^{i} &= \overline{V}_{\gamma_{\boldsymbol{\theta}}}^{i}\left[\underline{V}_{\gamma_{\boldsymbol{\theta}}}^{-1}\underline{m}_{\gamma_{\boldsymbol{\theta}}}+\mathbf{Q}^{ii}\sum_{\tau=1}^{t-1}\theta_{\tau}^{i}\theta_{\tau+1}^{i}\right], \end{aligned} \tag{A-37}$$

and $\mathbf{Q}^{ii}$ is the $i-$th diagonal element of $\mathbf{Q}^{-1}$.

As for $p\left(\boldsymbol{Q}|\,\mu,\boldsymbol{\beta},\sigma_{\varepsilon}^{-2},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\mathcal{D}^{t}\right)$, we have that

$$\boldsymbol{Q}|\,\mu,\boldsymbol{\beta},\sigma_{\varepsilon}^{-2},\boldsymbol{\theta}^{t},\mathcal{D}^{t}\sim\mathcal{IW}\left(\overline{\mathbf{Q}},\overline{v}_{\mathbf{Q}}\right), \tag{A-38}$$

where

$$\overline{\mathbf{Q}}=\underline{\mathbf{Q}}+\sum_{\tau=1}^{t-1}\left(\boldsymbol{\theta}_{\tau+1}-diag\left(\boldsymbol{\gamma_{\theta}}\right)\boldsymbol{\theta}_{\tau}\right)\left(\boldsymbol{\theta}_{\tau+1}-diag\left(\boldsymbol{\gamma_{\theta}}\right)\boldsymbol{\theta}_{\tau}\right)'. \tag{A-39}$$

and $\overline{v}_{\mathbf{Q}}=\left(1+\underline{v}_{\mathbf{Q}}\right)\left(t-1\right)$.

Finally, draws from the predictive density $p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^{t}\right)$ can be obtained by noting than

$$\begin{aligned}p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^{t}\right) &= \int p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\boldsymbol{Q},\sigma_{\varepsilon}^{-2},\mathcal{D}^{t}\right) \\ &\quad\times p\left(\boldsymbol{\theta}_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\boldsymbol{Q},\sigma_{\varepsilon}^{-2},\mathcal{D}^{t}\right) \\ &\quad\times p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\boldsymbol{Q},\sigma_{\varepsilon}^{-2}\middle|\,\mathcal{D}^{t}\right)d\mu d\boldsymbol{\beta}d\boldsymbol{\theta}^{t+1}d\boldsymbol{\gamma_{\theta}}d\boldsymbol{Q}d\sigma_{\varepsilon}^{-2}. \end{aligned} \tag{A-40}$$

The first term in the integral above, $p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^{t},\boldsymbol{\gamma_{\theta}},\boldsymbol{Q},\sigma_{\varepsilon}^{-2},\mathcal{D}^{t}\right)$, represents the period $t+1$ predictive density of bond excess returns, treating model parameters as if they were

known with certainty, and so is straightforward to calculate. The second term in the integral, $p\left(\boldsymbol{\theta}_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right)$, reflects that the regression parameters may drift away from $\boldsymbol{\theta}_t$ over time. Finally, the last term in the integral, $p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2}\middle|\,\mathcal{D}^t\right)$, measures parameter uncertainty.

To obtain draws for $p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right)$, we proceed in three steps:

1. Simulate from $p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2}\middle|\,\mathcal{D}^t\right)$: draws from $p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2}\middle|\,\mathcal{D}^t\right)$ are obtained from the Gibbs sampling algorithm described above;

2. Simulate from $p\left(\boldsymbol{\theta}_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right)$: For a given θ_t and Q, note that μ, β, σ_ε^{-2}, and the history of regression parameters up to t become redundant, i.e., $p\left(\boldsymbol{\theta}_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right) = p\left(\boldsymbol{\theta}_{t+1}|\,\boldsymbol{\theta}_t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\mathcal{D}^t\right)$. Note also that (24), along with the distributional assumptions made with regards to $\eta_{\tau+1}$, imply that

$$\boldsymbol{\theta}_{t+1}|\,\boldsymbol{\theta}_t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\mathcal{D}^t \sim \mathcal{N}\left(diag\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}\right)\boldsymbol{\theta}_t,\boldsymbol{Q}\right). \tag{A-41}$$

3. Simulate from $p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right)$: For a given $\boldsymbol{\theta}_{t+1}$, μ, $\boldsymbol{\beta}$, and σ_ε^{-2}, $\boldsymbol{\theta}^t$, $\boldsymbol{\gamma}_{\boldsymbol{\theta}}$, and $\boldsymbol{Q}$ become redundant so $p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\sigma_\varepsilon^{-2},\mathcal{D}^t\right) = p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\sigma_\varepsilon^{-2},\mathcal{D}^t\right)$. Then use the fact that

$$rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},\mu,\boldsymbol{\beta},\sigma_\varepsilon^{-2},\mathcal{D}^t \sim \mathcal{N}\left(\left(\mu+\mu_t\right)+\left(\boldsymbol{\beta}+\boldsymbol{\beta}_t\right)'\boldsymbol{x}_\tau^{(n)},\sigma_\varepsilon^2\right). \tag{A-42}$$

A.4 Time varying Parameter, Stochastic Volatility Model

Our priors for the TVP-SV model combine the earlier choices for the TVP and SV models, i.e., (A-25) and (A-26) for the regression parameters, (A-7) and (A-9) for the SV component, and (A-28) and (A-29) for the TVP component.

To obtain draws from the joint posterior distribution $p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\middle|\,\mathcal{D}^t\right)$ under the TVP-SV model, we use the Gibbs sampler to draw recursively from the following seven conditional distributions:

1. $p\left(\boldsymbol{\theta}^t\middle|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\varepsilon^{-2},\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\mathcal{D}^t\right)$.
2. $p\left(\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\varepsilon^{-2}\middle|\,\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\boldsymbol{Q},\mathcal{D}^t\right)$.
3. $p\left(h^t\middle|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right)$.
4. $p\left(\boldsymbol{\gamma}_{\boldsymbol{\theta}}|\,\mu,\boldsymbol{\beta},\sigma_\varepsilon^{-2},\boldsymbol{\theta}^t,\boldsymbol{Q},\mathcal{D}^t\right)$.
5. $p\left(\boldsymbol{Q}|\,\mu,\boldsymbol{\beta},\lambda_0,\lambda_1,\sigma_\varepsilon^{-2},\boldsymbol{\theta}^t,\boldsymbol{\gamma}_{\boldsymbol{\theta}},\mathcal{D}^t\right)$.
6. $p\left(\lambda_0,\lambda_1|\,\mu,\boldsymbol{\beta},h^t,\sigma_\xi^{-2},\mathcal{D}^t\right)$.

7. $p\left(\sigma_\xi^{-2}\middle|\,\mu,\boldsymbol{\beta},h^t,\lambda_0,\lambda_1,\mathcal{D}^t\right)$.

With minor modifications, these steps are similar to the steps described in the TVP and SV sections above. Draws from the predictive density $p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right)$ can be obtained from

$$\begin{aligned} p\left(rx_{t+1}^{(n)}\middle|\,\mathcal{D}^t\right) &= \int p\left(rx_{t+1}^{(n)}\middle|\,\boldsymbol{\theta}_{t+1},h_{t+1},\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right) \\ &\quad\times p\left(\boldsymbol{\theta}_{t+1},h_{t+1}|\,\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2},\mathcal{D}^t\right) \\ &\quad\times p\left(\mu,\boldsymbol{\beta},\boldsymbol{\theta}^t,\boldsymbol{\gamma_\theta},\boldsymbol{Q},h^t,\lambda_0,\lambda_1,\sigma_\xi^{-2}\middle|\,\mathcal{D}^t\right) d\mu d\boldsymbol{\beta} d\boldsymbol{\theta}^{t+1} d\boldsymbol{\gamma_\theta} d\boldsymbol{Q} dh^{t+1} d\lambda_0 d\lambda_1 d\sigma_\xi^{-2}. \end{aligned} \tag{A-43}$$

and following the steps described in the SV and TVP sections above.

Appendix B Estimation of bond risk premia

In this appendix we describe how we estimate bond risk premia using a dynamic Gaussian affine term structure model. Yields are first collected in a vector, Y_t, which contains rates for J different maturities. The risk factors that determine the yields are denoted by Z_t; these include both macro factors (M_t) and yield factors ($F_t^{\mathcal{L}}$) extracted as the first $\mathcal{L}$ principal components of yields, i.e., $Z_t = \left(M_t', F_t^{\mathcal{L}\prime}\right)$. The macro factors can be unspanned, i.e., they are allowed to predict future yields without having additional explanatory power for current yields beyond the yield factors.

We begin by assuming that the continuously compounded nominal one-period interest rate r_t depends on the yield factors, but not on the macro factors,

$$r_t = \delta_0 + \delta_{\mathcal{F}}' F_t^{\mathcal{L}} + 0_{\mathcal{M}}' M_t. \tag{B-1}$$

Next, the risk factors are assumed to follow a Gaussian vector autoregression (VAR) under the risk neutral probability measure:

$$F_t^{\mathcal{L}} = \mu^{\mathbb{Q}} + \phi^{\mathbb{Q}} F_{t-1}^{\mathcal{L}} + 0_{\mathcal{L}\times\mathcal{M}}' M_{t-1} + \Sigma\varepsilon_t^{\mathbb{Q}}, \tag{B-2}$$

where $\varepsilon_t^{\mathbb{Q}} \sim \mathcal{N}(0, I)$. Finally, the evolution in Z_t under the physical measure takes the form:

$$Z_t = \mu + \phi Z_{t-1} + \Sigma\varepsilon_t, \tag{B-3}$$

where $\varepsilon_t \sim \mathcal{N}(0, I)$. Under these assumptions bond prices are exponentially affine in the yield factors and do not depend on the macro factors,

$$P_t^{(n)} = \exp\left(\mathcal{A}_n + \mathcal{B}_n F_t^{\mathcal{L}} + 0_{\mathcal{M}\times\mathcal{M}}' M_t\right), \tag{B-4}$$

where $\mathcal{A} = \mathcal{A}\left(\mu^{\mathbb{Q}}, \phi^{\mathbb{Q}}, \delta_0, \delta_{\mathcal{F}}, \Sigma\right)$ and $\mathcal{B} = \mathcal{B}\left(\phi^{\mathbb{Q}}, \delta_{\mathcal{F}}\right)$ are affine loadings which are given by the following recursions:

$$\begin{aligned} \mathcal{A}_{n+1} &= \mathcal{A}_n + \left(\mu^{\mathbb{Q}}\right)' \mathcal{B}_n + \frac{1}{2}\mathcal{B}_n' \Sigma\Sigma' \mathcal{B}_n - \delta_0 \\ \mathcal{B}_{n+1} &= \left(\phi^{\mathbb{Q}}\right)' \mathcal{B}_n - \delta_{\mathcal{F}} \end{aligned} \tag{B-5}$$

with starting values $\mathcal{A}_0 = 0$ and $\mathcal{B}_0 = 0$.

Using these coefficients, the model-implied yields are obtained as

$$y_t^{(n)} = -\frac{1}{n}\log\left(P_t^{(n)}\right) = A_j + B_n F_t^{\mathcal{L}}, \tag{B-6}$$

where

$$A_n = -\frac{1}{n}\mathcal{A}_n, \qquad B_n = -\frac{1}{n}\mathcal{B}_n. \tag{B-7}$$

Similarly, the risk neutral price of a n-period bond, $\widetilde{P}_t^{(n)}$, and its implied yield, $\widetilde{y}_t^{(n)}$, can be calculated as

$$\widetilde{P}_t^{(n)} = \exp\left(\widetilde{\mathcal{A}}_n + \widetilde{\mathcal{B}}_n Z_t\right), \tag{B-8}$$

and

$$\widetilde{y}_t^{(n)} = \widetilde{A}_n + \widetilde{B}_n Z_t, \tag{B-9}$$

where

$$\widetilde{A}_n = -\frac{1}{n}\widetilde{\mathcal{A}}_n, \qquad \widetilde{B}_n = -\frac{1}{n}\widetilde{\mathcal{B}}_n. \tag{B-10}$$

and where $\widetilde{\mathcal{A}}_n$ and $\widetilde{\mathcal{B}}_n$ are given by the following recursions:

$$\begin{aligned} \widetilde{\mathcal{A}}_{n+1} &= \widetilde{\mathcal{A}}_n + \mu' \widetilde{\mathcal{B}}_n + \frac{1}{2}\widetilde{\mathcal{B}}_n' \Sigma\Sigma' \widetilde{\mathcal{B}}_n - \delta_0 \\ \widetilde{\mathcal{B}}_{n+1} &= \phi' \widetilde{\mathcal{B}}_j - \delta_{\mathcal{F}} \end{aligned} \tag{B-11}$$

initialized at $\widetilde{\mathcal{A}}_0 = 0$ and $\widetilde{\mathcal{B}}_0 = 0$.

We follow Joslin et al. (2011) and Wright (2011) and include in $F_t^{\mathcal{L}}$ the first three principal components of zero-coupon bond yields using maturities ranging from three months to ten years. As macro factors, M_t, we use exponentially weighted moving averages of monthly inflation and IPI growth. The data used for estimation are monthly yields on zero-coupon bonds, inflation, and IPI growth, from 1982:01 to 2011:12. Our estimation approach also follows closely Joslin et al. (2011). Using the resulting estimates of the model parameters, we compute the risk premium at all maturities as the difference between the yields computed under the risk neutral measure, $\mathbb{Q}$, and the yields calculated under the physical measure, i.e.,

$$rp_t^{(n)} = y_t^{(n)} - \widetilde{y}_t^{(n)}. \tag{B-12}$$

Figure 1. **Bond excess returns**

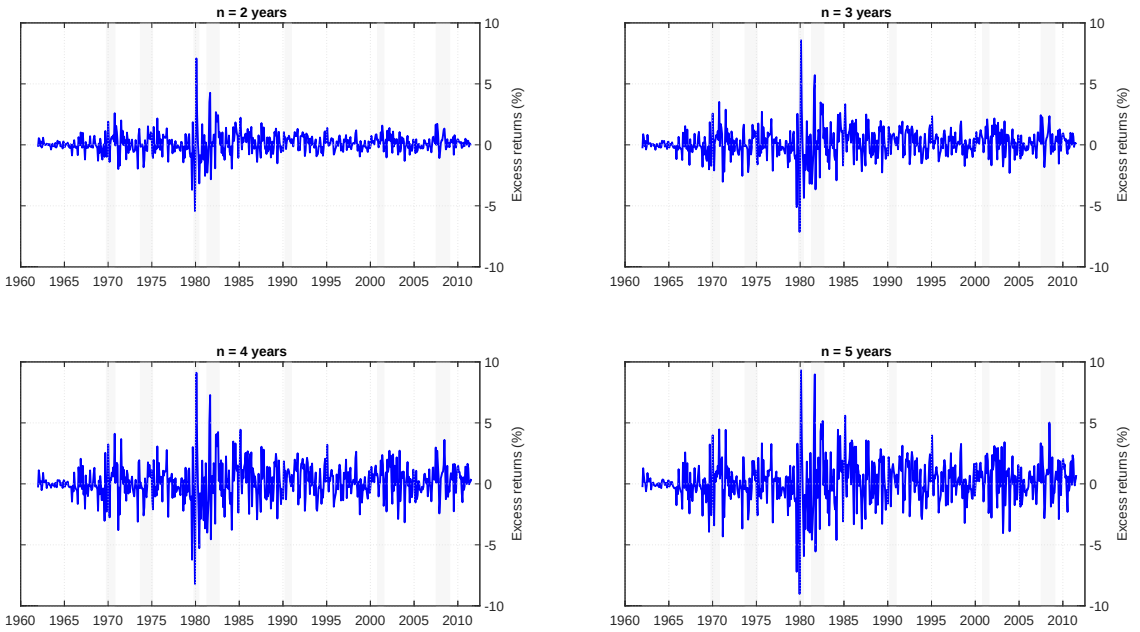

This figure shows time series of monthly bond excess returns (in percentage terms) for maturities (n) ranging from 2 years through 5 years. Monthly bond excess returns, $rx_{t+1/12}^{(n)}$, are computed from monthly yields, $y_t^{(n)}$, and are expressed in deviations from the 1-month T-bill rate, $rx_{t+1/12}^{(n)} = r_{t+1/12}^{(n)} - (1/12)y_t^{1/12}$, with $r_{t+1/12}^{(n)} = ny_t^{(n)} - (n - 1/12)y_{t+1/12}^{n-1/12}$. The sample ranges from 1962:01 to 2011:12.

Figure 2. **Parameter estimates for bond return forecasting model**

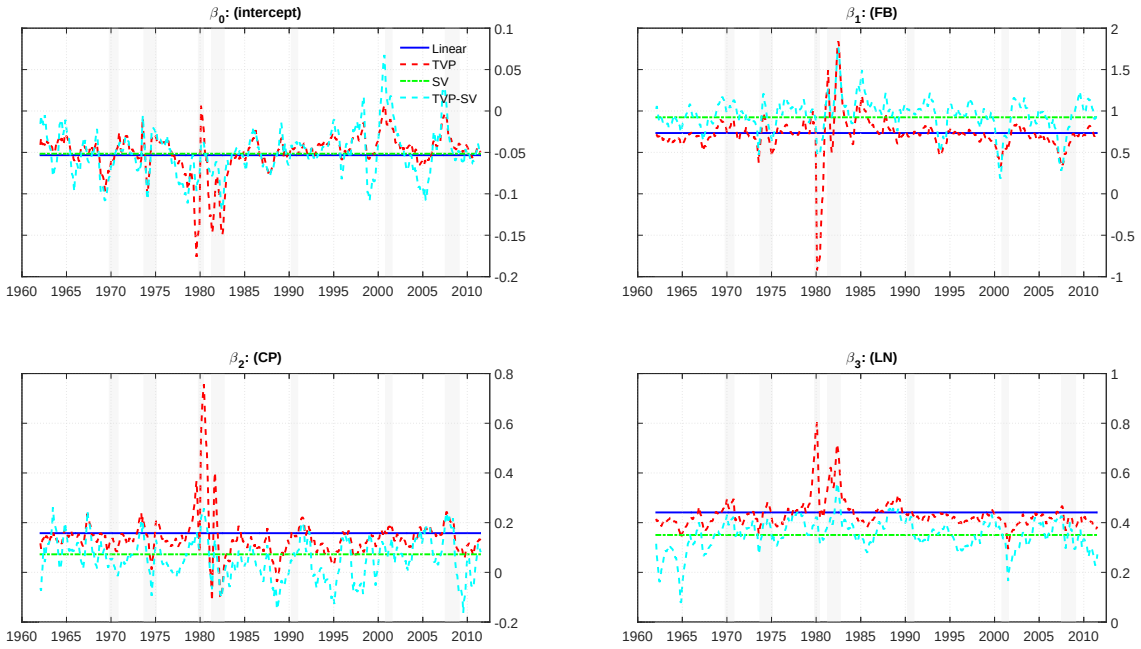

This figure displays parameter estimates for the FB-CP-LN model used to forecast monthly 3-year bond excess returns using as predictors the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) variables. The blue solid line represents the linear, constant coefficient model (Linear); the red dashed line tracks the parameter estimates for the time-varying parameter model (TVP); the green dashed-dotted line depicts the parameters for the stochastic volatility model (SV), while the dotted light-blue line shows estimates for the time-varying parameter, stochastic volatility (TVP-SV) model. The top left panel plots estimates of the intercept and the top right panel displays the coefficients on the FB predictor. The bottom left and right panels plot the coefficients on the CP and LN factors, respectively. The sample ranges from 1962:01 to 2011:12 and the parameter estimates are based on full-sample information.

Figure 3. **Posterior densities for model parameters**

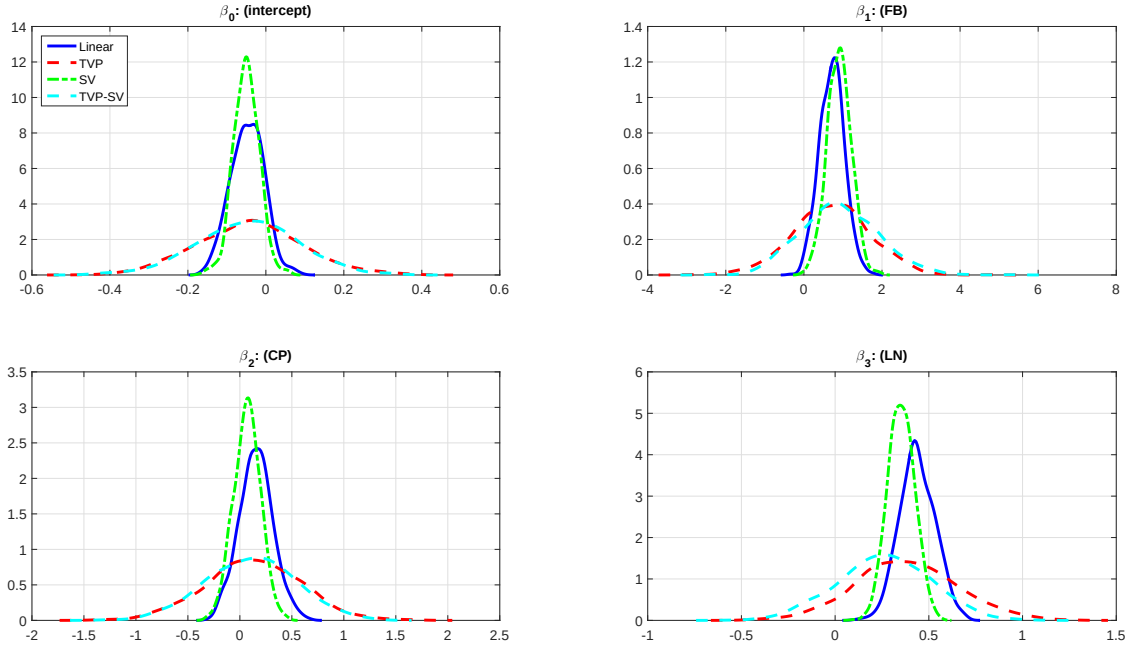

This figure displays posterior densities for the coefficients of the FB-CP-LN return model fitted to 3-year Treasury bonds, using as predictors the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) factors. The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line shows the parameter posterior density for the time-varying parameter (TVP) model; the green dashed-dotted line represents the stochastic volatility (SV) model, while the dotted light-blue line shows the posterior density for the time-varying parameter, stochastic volatility (TVP-SV) model. The first panel shows densities for the intercept. The second panel shows densities for the coefficient on the FB predictor. The third and fourth panels show densities for the coefficients on the CP and LN factors, respectively. The posterior density estimates shown here are based on their values as of 2011:12.

Figure 4. **Posterior densities for bond returns**

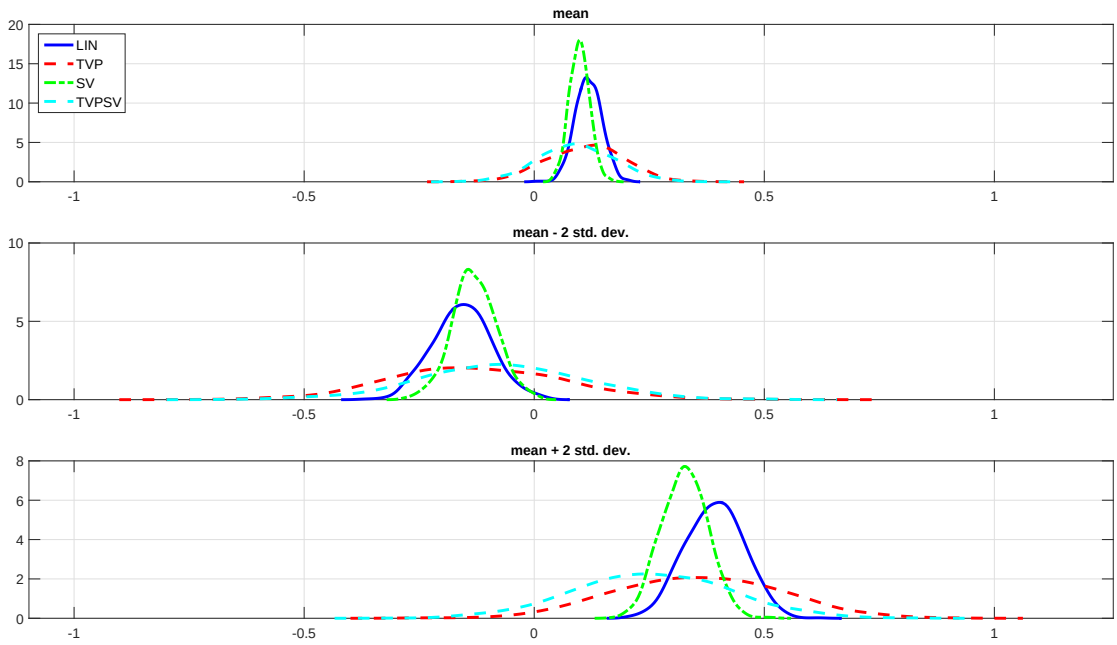

This figure shows the posterior density for excess returns on a three-year Treasury bond using the univariate Ludvigson-Ng (LN) state variable as a predictor. The LN variable is set at its sample mean $\overline{LN}$ (top panel), $\overline{LN} - 2stdev\,(LN)$ (middle panel), and $\overline{LN} + 2stdev\,(LN)$ (bottom panel). The blue solid line represents the linear, constant coefficient (Linear) model. the red dashed line tracks densities for the time-varying parameter (TVP) model. The green dashed-dotted line represents the stochastic volatility (SV) model, and the dotted light-blue line refers to the time varying parameter, stochastic volatility (TVP-SV) model. All posterior density estimates are based on the full data sample at the end of 2011.

Figure 5. **Conditional mean and volatility estimates for bond excess returns**

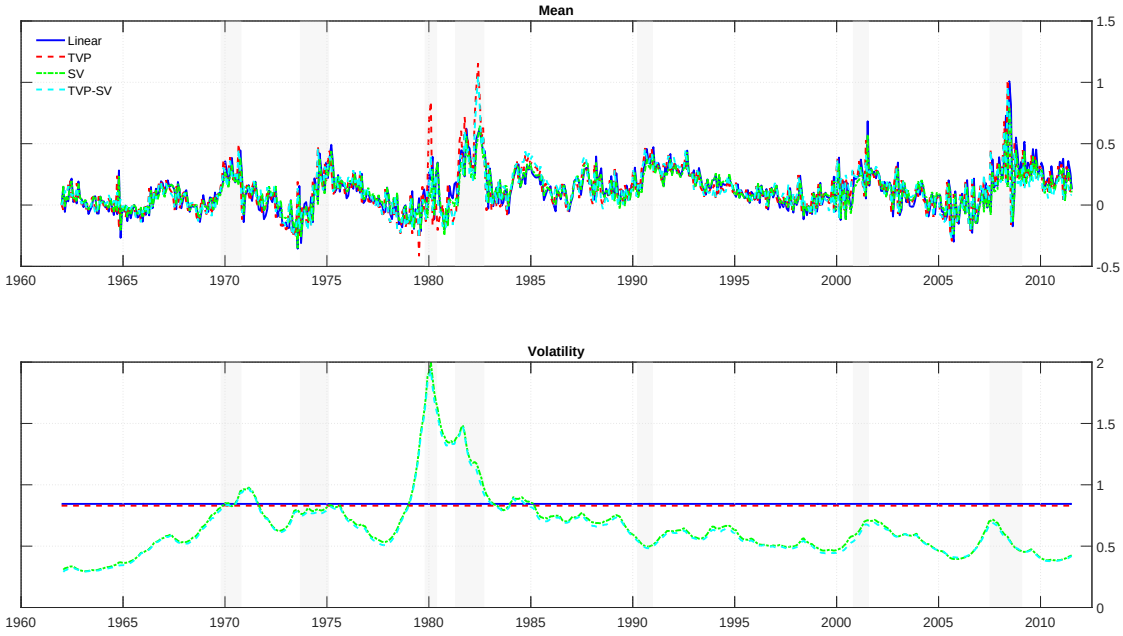

The top panel shows time-series of expected bond excess returns obtained from a range of models used to forecast monthly returns on a three-year Tresury bond using as predictors the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) factors. The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line tracks the time-varying parameter (TVP) model; the green dashed-dotted line depicts the stochastic volatility (SV) model, while the dotted light-blue line displays values for the time varying parameter, stochastic volatility (TVP-SV) model. The bottom panel displays volatility estimates for the FB-CP-LN models. The sample ranges from 1962:01 to 2011:12 and the estimates are based on full-sample information.

Figure 6. **Cumulative sum of squared forecast error differentials**

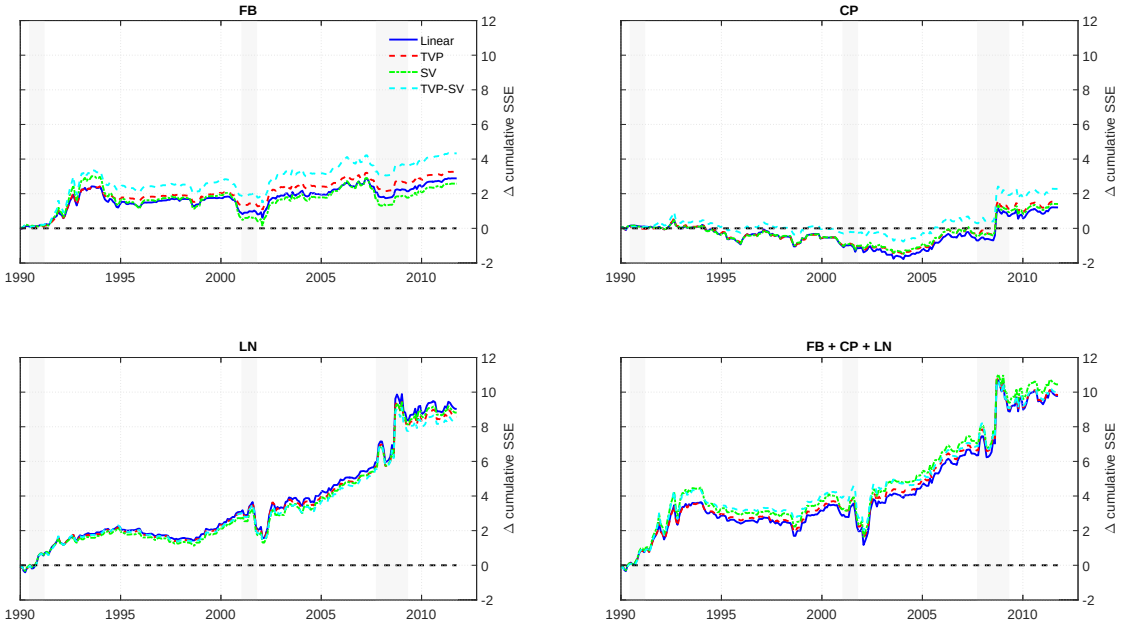

This figure shows the recursively calculated sum of squared forecast errors for the expectations hypothesis (EH) model minus the sum of squared forecast errors for a forecasting model with time-varying expected returns for a bond with a two year maturity, ($n = 2$). Each month we recursively estimate the parameters of the forecasting models and generate one-step-ahead forecasts of bond excess returns which are in turn used to compute out-of-sample forecasts. This procedure is applied to the EH model, which is our benchmark, as well as to forecasting models based on the Fama-Bliss (FB) predictor (1st window), the Cochrane-Piazzesi (CP) factor (2nd window), the Ludvigson-Ng (LN) factor (3rd window), and a multivariate model with all three predictors included (4th window). We then plot the cumulative sum of squared forecast errors (SSE_t) of the EH forecasts (SSE_t^{EH}) minus the corresponding value from the model with time-varying mean, $SSE_t^{EH} - SSE_t$. Values above zero indicate that a forecasting model with time-varying predictors produces more accurate forecasts than the EH benchmark, while negative values suggest the opposite. The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line tracks the time-varying parameter (TVP) model; the green dashed-dotted line represents the stochastic volatility (SV) model, while the dotted light-blue line refers to the time-varying parameter, stochastic volatility (TVP-SV) model. The out-of-sample period is 1990:01 - 2011:12.

Figure 7. **Cumulative sum of log-score differentials**

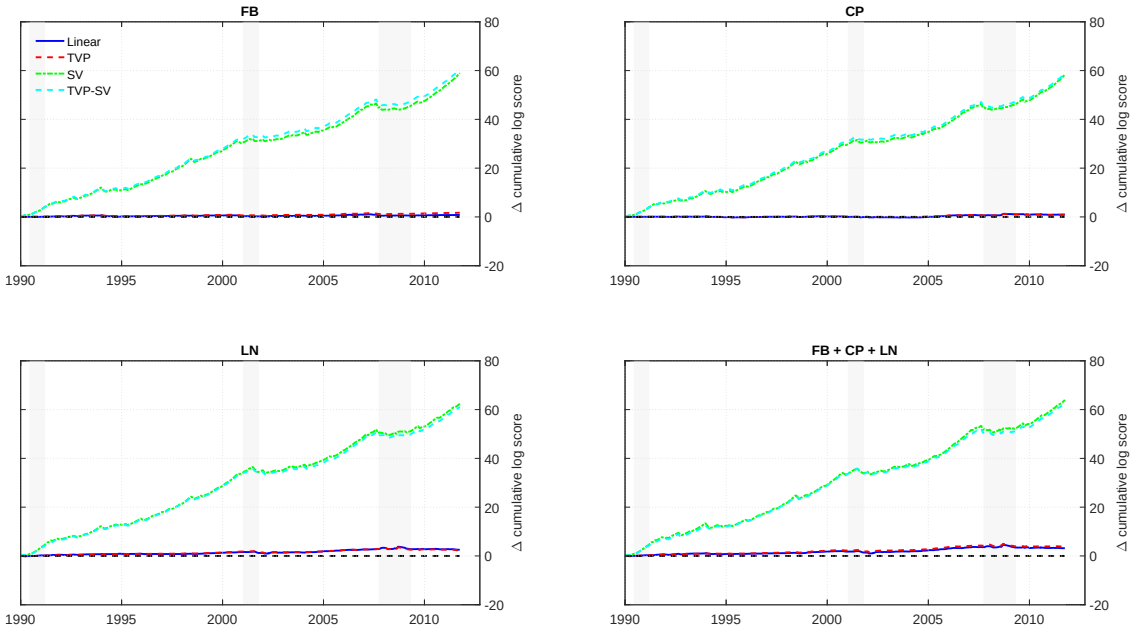

This figure shows the recursively calculated sum of log predictive scores from forecasting models with time-varying predictors minus the corresponding sum of log predictive scores for the EH model, using a 2-year Treasury bond. Each month we recursively estimate the parameters of the forecasting models and generate one-step-ahead density forecasts of bond excess returns which are in turn used to compute log-predictive scores. This procedure is applied to the benchmark EH model as well as to forecasting models based on the Fama-Bliss (FB) predictor (1st window), the Cochrane-Piazzesi (CP) factor (2nd window), the Ludvigson-Ng (LN) factor (3rd window), and a multivariate FB-CP-LN model (4th window). We then plot the cumulative sum of log predictive scores (LS_t) for the models with time-varying predictors minus the cumulative sum of log-predictive scores of the EH model, $LS_t - LS_t^{EH}$. Values above zero indicate that the time-varying mean model generates more accurate forecasts than the EH benchmark, while negative values suggest the opposite. The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line tracks the time-varying parameter (TVP) model; the green dashed-dotted line represents the stochastic volatility (SV) model, while the dotted light-blue line shows the time-varying parameter, stochastic volatility (TVP-SV) model. The out-of-sample period is 1990:01 - 2011:12.

Figure 8. **Economic value of out-of-sample bond return forecasts**

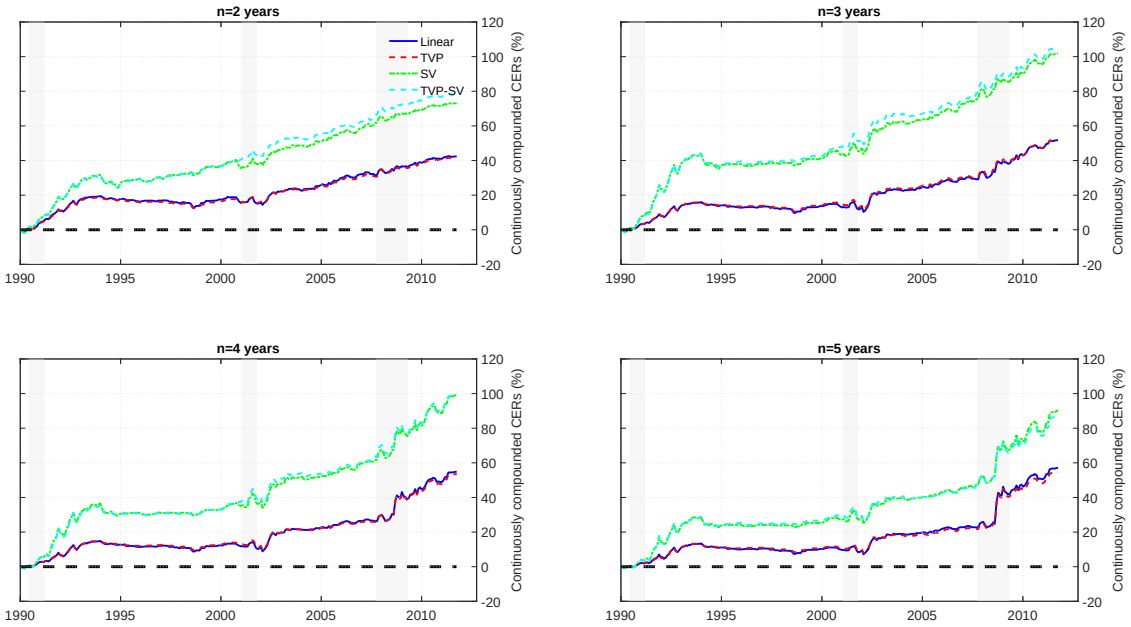

This figure plots cumulative certainty equivalent returns for the three-factor FB-CP-LN forecasting model that uses the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) factors as predictors, measured relative to the expectations hypothesis (EH) model. Each month we compute the optimal allocation to bonds and T-bills based on the predictive densities of bond excess returns. The investor is assumed to have power utility with a coefficient of relative risk aversion of ten and the weight on bonds is constrained to lie in the interval $[-200\%, 300\%]$. Each panel displays a different bond maturity, ranging from 2 years (1st panel) to 5 years (4th panel) The blue solid line represents the linear, constant coefficient (Linear) model; the red dashed line tracks the time-varying parameter (TVP) model; the green dashed-dotted line represents the stochastic volatility (SV) model, while the dotted light-blue line shows results for the time-varying parameter, stochastic volatility (TVP-SV) model. The out-of-sample period is 1990:01 - 2011:12.

Figure 9. **Weights on individual predictors in the optimal prediction pool**

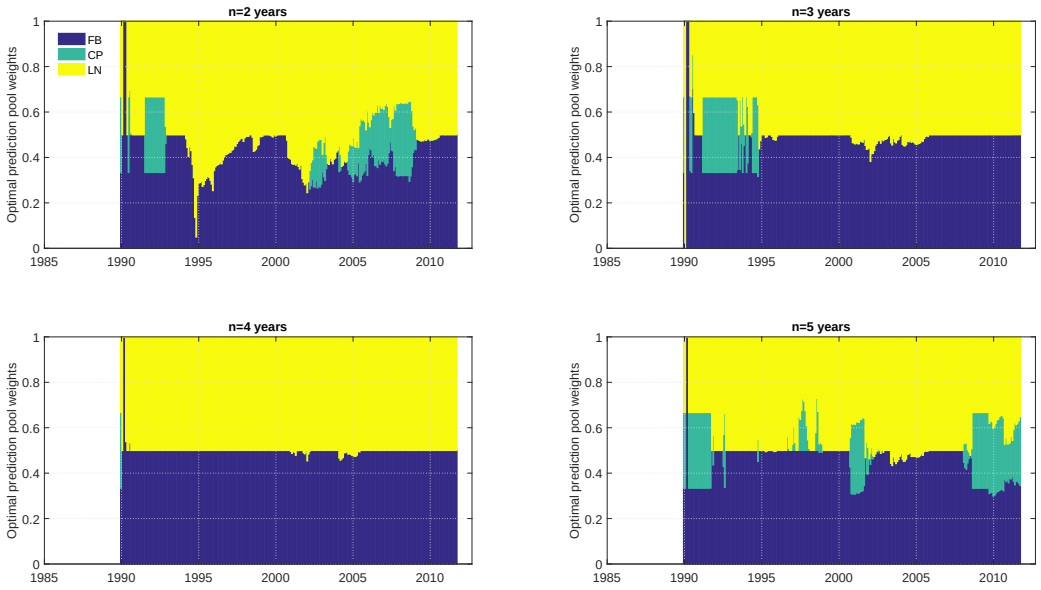

This figure plots the total weights on the individual predictors in the optimal predictive pool. At each point in time t, the weights are computed as $\mathcal{A}'\mathbf{w}_t$, where $\mathcal{A}$ is a 28×3 matrix representing all forecasting models by their unique combinations of zeros and ones and $\mathbf{w}_t$ is a 28×1 vector of period t optimal weights in the predictive pool, obtained in real time by solving the minimization problem

$$\mathbf{w}_t^* = \arg\max_{\mathbf{w}} \sum_{\tau=1}^{t-1} \log \left[\sum_{i=1}^{N} w_i \times S_{\tau+1,i} \right]$$

subject to $\mathbf{w}_t^*$ belonging to the $N-$dimensional unit simplex and N is the number of forecasting models. $S_{\tau+1,i}$ denotes the time $\tau+1$ log score for model i, $S_{\tau+1,i} = \exp(LS_{\tau+1,i})$. Blue bars show the combination weights associated with the Fama-Bliss (FB) factor; green bars show the weight assigned to the Cochrane-Piazzesi (CP) factor, and red bars show the weights assigned to the Ludvigson-Ng (LN) factor. Bond maturities range from 2 years (top left panel) to 5 years (bottom right panel).

Table 1. **Summary Statistics.**

	2 years	3 years	4 years	5 years
Panel A.1: One-month excess returns				
mean	1.4147	1.7316	1.9868	2.1941
mean (gross)	6.442	6.753	7.011	7.214
st.dev.	2.9711	4.1555	5.2174	6.2252
skew	0.4995	0.2079	0.0566	0.0149
kurt	14.8625	10.6482	7.9003	6.5797
AC(1)	0.1692	0.1533	0.1384	0.1235
Panel A.2: 12-month overlapping excess returns				
mean	0.5149	0.8601	1.1176	1.3177
mean (gross)	6.3511	6.6892	6.9353	7.1220
st.dev.	1.7527	3.1811	4.4103	5.5139
skew	-0.1014	-0.0618	-0.0196	0.0157
kurt	3.7801	3.5919	3.5172	3.5311
AC(1)	0.9313	0.9321	0.9322	0.9312
Panel A.3: 12-month overlapping excess returns (Cochrane-Piazzesi)				
mean	0.4768	0.8662	1.1435	1.2197
mean(gross)	6.2300	6.6193	6.8966	6.9729
st.dev.	1.7691	3.2438	4.5064	5.5381
skew	0.0636	-0.0256	0.0208	0.0157
kurt	3.4734	3.5507	3.4963	3.4792
AC(1)	0.9312	0.9336	0.9325	0.9234

Panel B: Predictors						
	Fama Bliss				CP	LN
	2-years	3-years	4-years	5-years		
mean	0.1078	0.1287	0.1451	0.1584	0.1533	0.1533
st.dev.	0.0996	0.1155	0.1278	0.1376	0.2126	0.3123
skew	-0.0716	-0.2379	-0.2234	-0.1434	0.7852	0.8604
kurt	3.7558	3.3782	3.0409	2.7548	5.2884	5.7283
AC(1)	0.8787	0.8995	0.9127	0.9227	0.6740	0.4109

Panel C: Correlation Matrix						
	FB-2	FB-3	FB-4	FB-5	CP	LN
FB-2	1.000	0.973	0.926	0.879	0.460	-0.087
FB-3		1.000	0.987	0.961	0.472	-0.049
FB-4			1.000	0.993	0.490	-0.010
FB-5				1.000	0.500	0.022
CP					1.000	0.184
LN						1.000

This table reports summary statistics for monthly bond excess returns and the predictor variables used in our study. Panels A.1-A.3 report the mean, standard deviation, skewness, kurtosis and first-order autocorrelation (AC(1)) of bond excess returns for 2 to 5-year bond maturities. Panel A.1 is based on monthly returns computed in excess of a one-month T-bill rate while Panels A.2 and Panel A.3 are based on 12-month overlapping returns, computed in excess of a 12-month T-bill rate. Gross returns do not subtract the risk-free rate. In Panels A.1 and A.2 returns are constructed using daily treasury yield data from Gurkaynak et al. (2007) while in Panel A.3 returns are constructed as in Cochrane and Piazzesi (2005) using the Fama-Bliss CRSP files. The row $mean(gross)$ reports the sample mean of gross bond returns (without subtracting the risk-free). Panel B reports the same summary statistics for the predictors: the Fama-Bliss (FB) forward spreads (2, 3, 4, and 5 years), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) factors. Panel C reports the correlation matrix for the predictors. The sample period is 1962-2011.

Table 2. **Full-sample OLS estimates**

	FB	CP	LN	FB+CP+LN
		2 years		
β_{FB}	1.1648***			1.1221***
β_{CP}		0.6549***		0.2317
β_{LN}			0.6713***	0.6736***
R^2	0.0166	0.0247	0.0581	0.0822
		3 years		
β_{FB}	1.3742***			1.2060***
β_{CP}		0.8785***		0.3299
β_{LN}			0.9068***	0.8877***
R^2	0.0158	0.0226	0.0541	0.0743
		4 years		
β_{FB}	1.6661***			1.3500***
β_{CP}		1.1054***		0.4202
β_{LN}			1.1149***	1.0679***
R^2	0.0183	0.0227	0.0518	0.0719
		5 years		
β_{FB}	1.9726***			1.4877***
β_{CP}		1.3613***		0.5502*
β_{LN}			1.3070***	1.2241***
R^2	0.0211	0.0243	0.0499	0.0712

This table reports OLS estimates of the slope coefficients for four linear models based on inclusion or exclusion of the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) predictor computed from a projection of the time series of cross-sectional averages of the 2, 3, 4, 5 bond excess returns on the 1, 2, 3, 4 and 5 year forward rates, and the Ludvigson-Ng (LN) predictor computed from a projection of the time-series of cross-sectional averages of the 2, 3, 4, 5 bond excess returns on five principal components obtained from a large panel of macroeconomic variables. Columns (1)-(3) report results for the univariate models, column (4) for the multivariate model that includes all three predictors. The last row in each panel reports the adjusted R^2. Stars indicate statistical significance based on Newey-West standard errors. ***: significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table 3. **Out-of-sample forecasting performance: R^2 values**

	Panel A: 2 years					Panel B: 3 years				
Model	OLS	LIN	SV	TVP	TVPSV	OLS	LIN	SV	TVP	TVPSV
FB	0.22%*	0.96%*	0.31%*	1.79%**	**2.48%****	1.77%**	1.46%**	1.31%**	1.66%**	**2.20%****
CP	-1.58%	0.52%*	0.67%*	0.70%*	**1.82%****	-0.45%	0.61%	0.71%	0.78%	**1.16%***
LN	-0.27%	4.94%***	5.51%***	4.88%***	**5.52%****	2.60%***	**4.59%****	4.48%***	4.35%***	4.20%***
$FB+CP+LN$	-3.06%	4.72%***	5.75%***	5.05%***	**6.47%****	0.43%***	4.97%***	**5.31%****	5.00%***	5.07%***
	Panel C: 4 years					Panel D: 5 years				
Model	OLS	LIN	SV	TVP	TVPSV	OLS	LIN	SV	TVP	TVPSV
FB	**2.39%****	1.64%***	1.72%**	1.64%**	2.08%***	**2.61%****	1.67%**	1.75%**	1.61%**	1.81%**
CP	0.25%	0.69%	0.63%	**0.94%***	0.87%*	0.73%	0.80%*	0.65%	0.79%*	**0.95%***
LN	3.77%***	**4.02%****	3.67%***	3.58%***	3.52%***	**4.32%****	3.43%***	3.26%***	3.07%**	3.15%**
$FB+CP+LN$	1.89%***	4.78%***	**4.93%****	4.56%***	4.67%***	2.51%***	4.45%***	**4.53%****	4.25%***	4.26%***

This table reports out-of-sample R^2 values for four prediction models based on the Fama-Bliss (FB), Cochrane-Piazzesi (CP), and Ludvigson-Ng (LN) predictors fitted to monthly bond excess returns, rx_{t+1}, measured relative to the one-month T-bill rate. The R^2_{OoS} is measured relative to the EH model: $R^2_{OoS} = 1 - \frac{\sum_{\tau=\underline{t}-1}^{\bar{t}-1}(r_{xt+1}-\hat{r}_{xt+1|t})^2}{\sum_{\tau=\underline{t}-1}^{\bar{t}-1}(r_{xt+1}-\bar{r}_{xt+1|t})^2}$ where $\hat{r}_{xt+1|t}$ is the conditional mean of bond returns based on a regression of monthly excess returns on an intercept and lagged predictor variable(s), x_t: $r_{xt+1} = \mu + \beta' x_t + \varepsilon_{t+1}$. $\bar{r}_{t+1|t}$ is the forecast from the EH model which assumes that the βs are zero. We report results for five specifications: (i) ordinary least squares (OLS), (ii) a linear specification with constant coefficients and constant volatility (LIN), (iii) a model that allows for stochastic volatility (SV), (iv) a model that allows for time-varying coefficients (TVP) and (v) a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). The out-of-sample period starts in January 1990 and ends in 2011:12. We measure statistical significance relative to the expectation hypothesis model using the Clark and West (2007) test statistic. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 4. **Out-of-sample forecasting performance: predictive likelihood**

	Panel A: 2 years				Panel B: 3 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.003	0.223***	0.006***	**0.230***	0.003	0.130***	0.006***	**0.132***
CP	0.004	0.221***	0.003	**0.223***	0.004	**0.129***	0.004	0.128***
LN	0.010**	**0.237***	0.009**	0.232***	0.011**	**0.141***	0.011**	0.138***
$FB + CP + LN$	0.012**	**0.242***	0.015***	0.238***	0.013**	**0.146***	0.016***	0.141***
	Panel C: 4 years				Panel D: 5 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.007***	0.090***	0.006**	**0.090***	0.006**	0.066***	0.006**	**0.067***
CP	0.006**	0.086***	0.005**	**0.086***	0.005**	0.062***	0.006**	**0.066***
LN	0.013**	**0.099***	0.011**	0.098***	0.013**	**0.075***	0.012**	0.073***
$FB + CP + LN$	0.017***	**0.104***	0.017**	0.101***	0.016**	**0.079***	0.017**	0.078***

This table reports the log predictive score for four forecasting models that allow for time-varying predictors relative to the log-predictive score computed under the expectation hypothesis (EH) model. The four forecasting models use the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. Positive values of the test statistic indicate that the model with time-varying predictors generates more precise forecasts than the EH benchmark. We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). The results are based on out-of-sample estimates over the sample period 1990 - 2011. ***: significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table 5. **Out-of-sample economic performance of bond portfolios**

	Weights constrained between 0% and 99%							
	Panel A: 2 years				Panel B: 3 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.23%	-0.25%	-0.15%	**-0.02%**	0.06%	0.21%	0.08%	**0.36%****
CP	-0.20%	0.03%	-0.15%	**0.07%**	-0.15%	0.37%**	-0.09%	**0.45%*****
LN	**0.11%**	0.07%	0.09%	0.09%	0.58%***	**0.70%*****	0.57%***	0.62%***
$FB+CP+LN$	0.05%	0.15%	0.04%	**0.24%****	0.49%**	0.72%***	0.50%**	**0.82%*****
	Panel C: 4 years				Panel D: 5 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.46%*	0.71%**	0.50%**	**0.81%*****	0.67%**	**1.02%****	0.62%**	1.01%**
CP	0.04%	0.46%**	0.12%	**0.64%*****	0.18%	0.55%**	0.18%	**0.69%*****
LN	0.87%***	**1.23%*****	0.78%***	1.21%***	0.86%***	**1.38%*****	0.78%***	1.36%***
$FB+CP+LN$	0.94%***	1.19%***	0.94%***	**1.25%*****	1.07%***	1.54%***	1.06%***	**1.55%*****
	Weights constrained between -200% and 300%							
	Panel E: 2 years				Panel F: 3 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.25%	0.99%*	0.39%*	**1.78%*****	0.50%**	1.30%**	0.53%**	**1.72%*****
CP	0.14%	1.47%***	0.17%	**1.85%*****	0.22%	1.15%**	0.25%	**1.28%****
LN	1.68%***	**2.84%*****	1.62%***	2.83%***	1.61%***	**2.98%*****	1.53%***	2.94%***
$FB+CP+LN$	1.69%***	2.73%***	1.65%***	**2.95%*****	1.79%***	3.27%***	1.78%***	**3.34%*****
	Panel G: 4 years				Panel H: 5 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.63%**	1.33%**	0.61%**	**1.47%****	0.69%**	**1.26%****	0.65%**	1.24%**
CP	0.24%	0.83%**	0.32%	**0.91%****	0.28%	0.70%**	0.27%	**0.88%****
LN	1.47%***	**2.35%*****	1.31%***	2.29%***	1.32%**	**2.02%*****	1.15%**	1.78%**
$FB+CP+LN$	1.73%***	**2.78%*****	1.65%***	2.66%**	1.66%**	**2.27%****	1.49%**	1.96%*

This table reports annualized certainty equivalent return values for portfolio decisions based on recursive out-of-sample forecasts of bond excess returns. Each period an investor with power utility and coefficient of relative risk aversion of 10 selects 2, 3, 4, or 5-year bond and 1-month T-bills based on the predictive density implied by a given model. The four forecasting models use the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model with both time varying coefficients and stochastic volatility ($TVPSV$). Statistical significance is based on a one-sided Diebold-Mariano test applied to the out-of-sample period 1990-2011. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 6. **Bond return predictability in expansions and recessions**

Model	LIN		SV		TVP		TVPSV	
	Exp	Rec	Exp	Rec	Exp	Rec	Exp	Rec
				Panel A: 2 years				
FB	2.53%	0.44%	3.01%	0.16%	6.76%	**8.11%****	7.96%	3.76%
CP	1.85%	**2.91%****	1.93%	**2.62%***	6.08%	5.61%**	6.47%	4.39%**
LN	2.11%	**9.30%*****	2.26%	**7.73%*****	4.94%	**15.46%*****	6.05%	**10.66%*****
$FB+CP+LN$	5.33%	**10.69%*****	5.51%	**8.66%*****	12.30%	**23.46%*****	12.74%	**14.93%*****
				Panel B: 3 years				
FB	2.17%	0.06%	2.74%	-0.28%	4.56%	**4.74%****	5.32%	1.83%
CP	1.36%	**2.70%****	1.42%	**2.80%****	3.83%	**4.33%****	4.03%	3.82%**
LN	2.09%	**7.93%*****	2.12%	**7.29%*****	3.77%	**11.89%*****	4.16%	**9.32%*****
$FB+CP+LN$	4.50%	**8.73%*****	4.63%	**7.85%*****	8.58%	**17.21%*****	8.74%	**11.83%*****
				Panel C: 4 years				
FB	1.99%	0.18%	2.45%	-0.05%	3.65%	3.11%*	4.20%	1.41%
CP	1.12%	**2.70%****	1.20%	**2.83%****	2.87%	**4.08%****	2.94%	**3.79%****
LN	1.89%	**7.03%*****	1.93%	**6.93%*****	3.06%	**10.11%*****	3.28%	**8.60%*****
$FB+CP+LN$	3.91%	**7.88%*****	4.02%	**7.54%*****	6.83%	**14.04%*****	6.84%	**10.69%*****
				Panel D: 5 years				
FB	1.88%	0.42%	2.23%	0.25%	3.11%	2.33%	3.55%	1.24%
CP	1.03%	**2.78%****	1.08%	**2.85%****	2.37%	**3.88%****	2.44%	**3.70%****
LN	1.65%	**6.32%*****	1.69%	**6.19%*****	2.52%	**8.85%*****	2.74%	**8.05%*****
$FB+CP+LN$	3.41%	**7.30%*****	3.49%	**7.10%*****	5.59%	**12.04%*****	5.62%	**9.87%*****

This table reports the R^2 from regressions of bond excess returns on the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. We report results separately for expansions (Exp) and recessions (Rec) as defined by the NBER recession index. Results are shown for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). The R^2 in expansions is computed as $R^2_{i,0} = 1 - \frac{e_{i,0}{}' e_{i,0}}{e_{EH,0}{}' e_{EH,0}}$ where $e_{i,0}$ and $e_{EH,0}$ denote the vectors of residuals of the alternative and the benchmark model, respectively, during expansions. Similarly, the R^2 in recessions only uses the vector of residuals in recessions: $R^2_{i,1} = 1 - \frac{e_{i,1}{}' e_{i,1}}{e_{EH,1}{}' e_{EH,1}}$. We test whether the R^2 is higher in recessions than in expansions using a bootstrap methodology. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 7. **Sharpe ratios in expansions and recessions**

Model	LIN		SV		TVP		TVPSV	
	Exp	Rec	Exp	Rec	Exp	Rec	Exp	Rec
				Panel A: 2 years				
FB	0.49%	0.43%	0.49%	0.35%	0.48%	**0.56%**	0.50%	0.48%
CP	0.46%	**0.63%**	0.46%	**0.54%**	0.44%	**0.72%**	0.47%	**0.65%**
LN	0.38%	**1.12%**	0.35%	**0.90%**	0.37%	**1.22%**	0.37%	**1.00%**
$FB+CP+LN$	0.38%	**1.14%**	0.42%	**0.91%**	0.37%	**1.33%**	0.42%	**1.05%**
				Panel B: 3 years				
FB	0.43%	0.38%	0.48%	0.36%	0.43%	**0.44%**	0.49%	0.42%
CP	0.40%	**0.53%**	0.47%	**0.56%**	0.39%	**0.59%**	0.47%	**0.61%**
LN	0.35%	**0.93%**	0.37%	**0.93%**	0.34%	**1.00%**	0.38%	**1.00%**
$FB+CP+LN$	0.34%	**0.92%**	0.42%	**0.94%**	0.34%	**1.04%**	0.43%	**1.02%**
				Panel C: 4 years				
FB	0.39%	0.35%	0.49%	0.38%	0.39%	0.38%	0.48%	0.40%
CP	0.37%	**0.48%**	0.47%	**0.55%**	0.37%	**0.52%**	0.48%	**0.61%**
LN	0.32%	**0.80%**	0.39%	**0.93%**	0.32%	**0.85%**	0.39%	**0.98%**
$FB+CP+LN$	0.32%	**0.78%**	0.43%	**0.93%**	0.31%	**0.88%**	0.43%	**0.99%**
				Panel D: 5 years				
FB	0.36%	0.32%	0.48%	0.38%	0.35%	0.34%	0.48%	0.41%
CP	0.34%	**0.44%**	0.47%	**0.55%**	0.34%	**0.47%**	0.47%	**0.58%**
LN	0.31%	**0.71%**	0.40%	**0.89%**	0.30%	**0.75%**	0.40%	**0.96%**
$FB+CP+LN$	0.30%	**0.69%**	0.43%	**0.89%**	0.31%	**0.77%**	0.43%	**0.97%**

This table reports the annualized Sharpe ratio computed from conditional mean and conditional volatility estimates implied by regressions of bond excess returns on the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. We report results separately for expansions (Exp) and recessions (Rec) as defined by the NBER recession index. Results are shown for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$).

Table 8. **Correlations between expected bond excess returns and economic variables**

	Panel A: GDP				Panel B: Inflation			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.14	0.15	0.08	0.06	-0.11	-0.09	-0.13	-0.11
CP	-0.34***	-0.32***	-0.39***	-0.41***	-0.31***	-0.32***	-0.35***	-0.40***
LN	-0.61***	-0.61***	-0.62***	-0.63***	-0.38***	-0.36***	-0.38***	-0.36***
$FB + CP + LN$	-0.49***	-0.41***	-0.51***	-0.47***	-0.36***	-0.31***	-0.36***	-0.33***
	Panel C: GDP Uncertainty				Panel D: Inflation Uncertainty			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	0.29***	0.27***	0.32***	0.31***	0.01	-0.02	0.04	0.02
CP	0.45***	0.44***	0.46***	0.46***	0.39***	0.37***	0.40***	0.35***
LN	0.51***	0.52***	0.50***	0.49***	0.49***	0.47***	0.48***	0.45***
$FB + CP + LN$	0.57***	0.57***	0.56***	0.54***	0.45***	0.39***	0.45***	0.38***

This table reports the contemporaneous correlations between out-of-sample forecasts of excess returns on a two-year Treasury bond and real GDP growth (Panel A), inflation (Panel B), real GDP growth uncertainty (Panel C) and inflation uncertainty (Panel D). Real GDP growth is computed as $\Delta log(GDP_{t+1})$ where GDP_{t+1} is the real gross domestic product (GDPMC1 Fred mnemonic). Inflation is computed as $\Delta log(CPI_{t+1})$ where CPI is the consumer price index for all urban consumers (CPIAUCSL Fred mnemonic). Real GDP growth uncertainty is the cross-sectional dispersion (the difference between the 75th percentile and the 25th percentile) for real GDP forecasts from the Philadelphia Fed Survey of Professional Forecasters. Inflation uncertainty is the cross-sectional dispersion (the difference between the 75th percentile and the 25th percentile) for CPI forecasts from the Philadelphia Fed Survey of Professional Forecasters. The bond return prediction models use the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). Finally, we test whether the correlation coefficients are statistically different from zero. All results are based on the out-of-sample period 1990-2011. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 9. **Expected bond excess returns and survey forecasts of the 3-Month T-bill rate.**

	Panel A: 2 years							
	Slope coefficient				R^2			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.04***	-0.04***	-0.03***	-0.04***	0.20	0.15	0.20	0.14
CP	-0.05***	-0.05***	-0.05***	-0.04***	0.30	0.30	0.30	0.24
LN	-0.06***	-0.05***	-0.06***	-0.04***	0.15	0.12	0.14	0.09
$FB+CP+LN$	-0.10***	-0.08***	-0.09***	-0.07***	0.24	0.20	0.22	0.14
	Panel B: 3 years							
	Slope coefficient				R^2			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.06***	-0.08***	-0.06***	-0.07***	0.45	0.36	0.45	0.35
CP	-0.06***	-0.06***	-0.06***	-0.06***	0.38	0.39	0.37	0.35
LN	-0.08***	-0.07***	-0.08***	-0.07***	0.21	0.19	0.20	0.16
$FB+CP+LN$	-0.13***	-0.12***	-0.12***	-0.11***	0.33	0.32	0.31	0.27
	Panel C: 4 years							
	Slope coefficient				R^2			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.09***	-0.10***	-0.09***	-0.10***	0.59	0.52	0.59	0.49
CP	-0.07***	-0.07***	-0.07***	-0.07***	0.45	0.45	0.44	0.43
LN	-0.09***	-0.08***	-0.09***	-0.08***	0.26	0.25	0.24	0.21
$FB+CP+LN$	-0.15***	-0.15***	-0.15***	-0.15***	0.40	0.39	0.38	0.35
	Panel D: 5 years							
	Slope coefficient				R^2			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
FB	-0.11***	-0.13***	-0.11***	-0.12***	0.69	0.62	0.67	0.58
CP	-0.08***	-0.08***	-0.08***	-0.08***	0.50	0.50	0.48	0.48
LN	-0.10***	-0.09***	-0.10***	-0.10***	0.30	0.28	0.28	0.26
$FB+CP+LN$	-0.17***	-0.17***	-0.17***	-0.17***	0.44	0.44	0.42	0.40

This table reports the R^2 and OLS estimates of slope coefficients from a regression of the predicted bond excess return on consensus forecasts of the 3-month T-bill rate reported by the Philadelphia Fed Survey of Professional Forecasters. The bond return prediction models use the Fama-Bliss (FB) forward spread predictor, the Cochrane-Piazzesi (CP) combination of forward rates, the Ludvigson-Ng (LN) macro factor, and the combination of these. Panels A-D display results for 2-5 year bond maturities. We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). All results are based on the sample 1990-2011. Stars indicate statistical significance based on Newey-West standard errors. ***: significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table 10. **Risk-premium regression**

	Panel A: 2 years				Panel B: 3 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
β	0.18	0.52	0.21	0.60	0.12	0.35	0.13	0.37
t-stat	0.78	2.45	0.95	2.95	0.58	1.77	0.67	1.85
	Panel C: 4 years				Panel D: 5 years			
Model	LIN	SV	TVP	TVPSV	LIN	SV	TVP	TVPSV
β	0.12	0.28	0.11	0.27	0.13	0.23	0.11	0.21
t-stat	0.63	1.52	0.56	1.42	0.73	1.33	0.59	1.16

This table reports OLS estimates of the slope coefficients (and the relative t-stats) from the following regression

$$\overline{rx}_t = \mu + \beta \overline{rp}_t + u_t,$$

where $\overline{rx}_t$ denotes the predicted bond excess returns and $\overline{rp}_t$ denotes the risk premium estimates, obtained from a term structure model with unspanned macro risks, based on the approach of Joslin et. a. (2011) and Wright (2011). We report results for a linear specification with constant coefficients and constant volatility (LIN), a model that allows for stochastic volatility (SV), a model that allows for time-varying coefficients (TVP) and a model that allows for both time-varying coefficients and stochastic volatility ($TVPSV$). All results are based on the sample 1990-2011. Stars indicate statistical significance based on Newey-West standard errors. ***: significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table 11. **Economic and statistical performance of forecast combinations**

Method	2 years	3 years	4 years	5 years
	Panel A: Out-of-sample R^2			
OW	**5.30%*****	**4.99%*****	**5.13%*****	**5.07%*****
EW	**5.35%*****	**4.12%*****	**3.59%*****	**3.11%*****
BMA	**5.51%*****	**4.66%*****	**4.17%*****	**3.67%*****
	Panel B: Predictive Likelihood			
OW	**0.24*****	**0.15*****	**0.10*****	**0.07*****
EW	**0.14*****	**0.09*****	**0.06*****	**0.05*****
BMA	**0.24*****	**0.14*****	**0.09*****	**0.07*****
	Panel C: CER (Weights between 0% and 99%)			
OW	**0.18%**	**0.77%*****	**1.23%*****	**1.55%*****
EW	**0.17%***	**0.60%*****	**1.02%*****	**1.05%*****
BMA	**0.22%***	**0.80%*****	**1.22%*****	**1.48%*****
	Panel D: CER (Weights between -200% and 300%)			
OW	**2.57%*****	**3.27%*****	**2.88%*****	**1.99%*****
EW	**1.98%*****	**1.83%*****	**1.67%*****	**1.43%*****
BMA	**2.80%*****	**2.95%*****	**2.27%*****	**1.83%*****

This table reports out-of-sample results for the optimal predictive pool (OW) of Geweke and Amisano (2011), an equal-weighted (EW) model combination scheme, and Bayesian Model Averaging (BMA) applied to 28 forecasting models that models based on all possible combinations of the CP, FB and LN factors estimated using linear, SV, TVP and TVPSV methods. In each case the models and combination weights are estimated recursively using only data up to the point of the forecast. The R^2 values in Panel A use the out-of-sample R^2 measure proposed by Campbell and Thompson (2008). The predictive likelihood in Panel B is the value of the test for equal accuracy of the predictive density log-scores proposed by Clark and Ravazzolo (2014). CER values in Panels C and D are the annualized certainty equivalent returns derived for an investor with power utility and a coefficient of relative risk aversion of 10 who uses the posterior predictive density implied by the forecast combination. The forecast evaluation sample is 1990:01-2011:12. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.