



Forecasting Stock Returns under Economic Constraints

Davide Pettenuzzo, Economics Department, Brandeis University

Allan Timmermann, University of California, San Diego

Rossen Valkanov, University of California, San Diego

Working Paper Series

Forecasting Stock Returns under Economic Constraints*

Davide Pettenuzzo
Brandeis University[†]

Allan Timmermann
UCSD, CEPR, and CREATES[‡]

Rossen Valkanov
UCSD[§]

May 9, 2013

Abstract

We propose a new approach to imposing economic constraints on time-series forecasts of the equity premium. Economic constraints are used to modify the posterior distribution of the parameters of the predictive return regression in a way that better allows the model to learn from the data. We consider two types of constraints: Non-negative equity premia and bounds on the conditional Sharpe ratio, the latter of which incorporates time-varying volatility in the predictive regression framework. Empirically, we find that economic constraints systematically reduce uncertainty about model parameters, reduce the risk of selecting a poor forecasting model, and improve both statistical and economic measures of out-of-sample forecast performance. The Sharpe ratio constraint, in particular, results in considerable economic gains.

Key words: Economic constraints; Sharpe ratio, Equity premium predictions; Bayesian analysis

JEL classification: C11, C22, G11, G12

*We thank John Campbell, Wayne Ferson, Blake LeBaron, Lubos Pastor, Seth Pruitt, and Guofu Zhou for helpful comments and suggestions. Xia Meng provided excellent research assistance. We are grateful to John Campbell and David Rapach for making their codes available.

[†]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.

[§]University of California, San Diego, 9500 Gilman Drive, MC 0553, La Jolla CA 92093. Tel: (858) 534-0898. Email: rvalkanov@ucsd.edu.

1 Introduction

Equity premium forecasts play a central role in areas as diverse as asset pricing, portfolio allocation, and performance evaluation of investment managers.¹ Yet, despite more than twenty five years of research, it is commonly found that models that allow for time-varying return predictability produce worse out-of-sample forecasts than a simple benchmark that assumes a constant risk premium. This finding has led authors such as Bossaerts and Hillion (1999) and Welch and Goyal (2008) to question the economic value of ex-ante return forecasts that allow for time-varying expected returns.

Economically motivated constraints offer the potential to sharpen forecasts, particularly when the data is noisy and parameter uncertainty is a concern as in return prediction models. While economic constraints have previously been found to improve forecasts of asset returns, there is no broad consensus on how to impose such constraints. For example, Ang and Piazzesi (2003) impose no-arbitrage restrictions to identify the parameters in a term structure model; Campbell and Thompson (2008) truncate their equity premium forecasts at zero and also constrain the sign of the slope coefficients in return prediction models, while Pastor and Stambaugh (2009) and Pastor and Stambaugh (2012) use informative priors to ensure that the sign of the correlation between shocks to unexpected and expected returns is negative.

This paper proposes a new approach for incorporating economic information via inequality constraints on moments of the predictive distribution of the equity premium. We focus on two types of economic constraints. The first, which we label the equity premium constraint, follows the idea of Campbell and Thompson (2008) and constrains the conditional mean of the equity premium to be non-negative. It is difficult to imagine an equilibrium setting where risk-averse investors would hold stocks if their expected compensations were negative, and so this seems like a mild restriction. The second constraint imposes that the conditional Sharpe ratio has to lie between zero and a predetermined upper bound. The zero lower bound is identical to the equity premium (EP) constraint, whereas the upper bound rules out that the price of risk becomes too high. The Sharpe ratio of the market portfolio is extensively used in finance and, much like the equity premium, academics and investors can be expected to have strong priors about its magnitude.² Yet, Sharpe ratio (SR) constraints cast as inequality constraints on the predictive moments of the return distribution have not, to our knowledge, previously been

¹Papers on time-series predictability of stock returns include Campbell (1987), Campbell and Shiller (1988), Fama and French (1988), Fama and French (1989), Ferson and Harvey (1991), Keim and Stambaugh (1986) and Pesaran and Timmermann (1995). Examples of asset allocation studies under return predictability include Ait-Sahalia and Brandt (2001), Barberis (2000), Brennan et al. (1997), Campbell and Viceira (1999), Kandel and Stambaugh (1996) and Xia (2001). Avramov and Wermers (2006) and Ferson and Schadt (1996) consider mutual fund performance under time-varying investment opportunities.

²See Lettau and Wachter (2007) and Lettau and Wachter (2011) for recent examples of theoretical asset pricing models that rely on calibrations using the Sharpe ratio. For a good treatment of the Sharpe ratio and its theoretical and empirical links to asset pricing models, see Cochrane (2001) and Lettau and Ludvigson (2010).

explicitly explored in the return predictability literature.³

Other studies have considered bounds on the maximum Sharpe ratio in the context of cross-sectional pricing models, which is quite different from our focus here. MacKinlay (1995) introduces a bound on the maximum squared Sharpe ratio as a way to distinguish between risk- and non-risk explanations of deviations from the CAPM. MacKinlay and Pastor (2000) provide estimates of factor pricing models that condition on a given value of the Sharpe ratio. In a Bayesian setting this corresponds to investors having different degrees of confidence in the asset pricing model; setting the maximum Sharpe ratio at zero corresponds to dogmatic beliefs, while a very large Sharpe ratio corresponds to completely skeptical beliefs about the model.⁴

To incorporate economic information, we develop a Bayesian approach that lets us compute the predictive density of the equity premium subject to economic constraints. Importantly, the approach makes efficient use of the entire sequence of observations in computing the predictive density and also accounts for parameter uncertainty. Our approach builds on the conventional linear prediction model and so simplifies to the conventional model if the economic constraints are not binding in a particular sample.

The predictive moments of the return distribution get updated as new data arrive and so the inequality constraints give rise to dynamic learning effects. To see how this works, suppose that a new observation arrives that, under the previous parameter estimates, imply a negative conditional equity premium. Since this is ruled out, the economic constraints can force the posterior distribution of the parameter estimates to shift significantly – even in situations where the estimates of the standard linear model do not change at all. This effect turns out to be empirically important, particularly for “large” values of the predictor variables. Our empirical analysis finds that the posterior variance of the equity premium distribution—one measure of parameter estimation uncertainty—can be several times bigger for the unconstrained model compared with the constrained models, when evaluated at such large values of the predictor variables.

Our approach of incorporating economic constraints works very differently from that taken by previous studies such as Campbell and Thompson (2008). To highlight such differences, consider the constraint that the equity premium is non-negative. Campbell and Thompson (2008) impose this constraint by truncating the predicted equity premium at zero if the predicted

³Ross (2005) and Zhou (2010) consider constraints on the R^2 of the predictive return distribution. In practice, there will be a close relationship between constraints on the Sharpe ratio and constraints on the R^2 , see, e.g., Campbell and Thompson (2008) for investors with mean variance utility. Wachter and Warusawitharana (2009) also consider priors on the slope coefficient in the return equation which translate into priors about the predictive R^2 of the return equation. Shanken and Tamayo (2012) study return predictability by allowing for time-varying risk and specify a prior on the Sharpe ratio.

⁴Cochrane and Saa-Requejo (2000) also discuss upper bounds on Sharpe ratios and relate them to so-called good-deal bounds. Ross (1976) introduces asset pricing constraints by assuming that portfolios cannot have Sharpe ratios greater than twice the Sharpe ratio of the market portfolio.

value is negative. While this truncation approach can be viewed as a first approximation to imposing moment or parameter constraints, it does not make efficient use of the information in the theoretical constraints. In particular, this approach never learns from the information that comes from observing that the model implies negative forecasts of the equity premium and so the underlying model continues to repeat the same mistakes when faced with new data similar to previously observed data.

To illustrate the practical importance of these differences, in Figure 1 we plot equity premium forecasts (top panel) of monthly U.S. stock (excess) returns computed recursively over the period 1947-2010 and based on an intercept and the log dividend yield. The unconstrained forecasts and simple truncated forecasts are identical in all periods where the unconstrained posterior mean is non-negative—since the constraint is non-binding in this case—which holds in most periods up to 1990. As a result of the declining dividend yield from 1990 and onwards, the posterior means of the unconstrained excess return forecasts turn negative and so the constrained posterior density is shifted upwards, resulting in a higher mean for the remainder of the sample. Nevertheless, for the majority of periods in the sample, the unconstrained and truncated forecasts are identical.

In contrast, our approach constrains the equity premium forecast to be non-negative at each point in time. This implies that we have T constraints in a sample of T observations, rather than just a single constraint. Every time a new pair of observations on the predictor variable and returns becomes available, the non-negativity constraint on the conditional equity premium is used to rule out values of the parameter that are infeasible given the constraint. As shown in Figure 1, this makes a big difference in practice; the EP constraint makes an impact much earlier in the sample, as can be seen from the cumulative sum of squared forecast errors (shown in the bottom panel) measured relative to the prevailing mean model. In fact, the equity premium constrained forecasts steadily outperform the prevailing mean forecasts over the entire sample while, conversely, the unconstrained or truncated forecasts are subject to periods of large under-performance such as during the mid-nineties.

In addition to the conditional EP constraint, we also explore whether imposing a lower and an upper bound on the Sharpe ratio of the market portfolio provide further improvements. An upper bound on the Sharpe ratio is equivalent to a time-varying upper bound on the equity premium that is proportional to the market volatility. The implementation of such a constraint is non-trivial as it involves modeling the conditional volatility of the market portfolio in a predictive regression framework. We use a parsimonious parameterization that allows us to explore time-variation in the conditional first and second moments of returns. We find that the SR constraint increases the statistical and economic gains not only relative to the unconstrained case, but also relative to the EP constraint.

Attempts at producing improved forecasts of stock returns have spawned a huge literature that originated from studies by Campbell (1987), Campbell and Shiller (1988), Fama and French

(1988), Fama and French (1989), Ferson and Harvey (1991), and Keim and Stambaugh (1986) who provided convincing economic arguments and in-sample empirical results that some of the fluctuations in returns are predictable because of persistent time variation in expected returns. In-sample evidence for predictability is accumulating as various new variables have been suggested as predictors of excess returns (Pontiff and Schall (1998), Lamont (1998), Lettau and Ludvigson (2001), Polk et al. (2006), among others). Out-of-sample predictability evidence, however, has been much less conclusive. Recent studies by Paye and Timmermann (2006) and Lettau and Van Nieuwerburgh (2008) argue that predictability weakened or disappeared during the 1990s. Bossaerts and Hillion (1999), Goyal and Welch (2003), and Welch and Goyal (2008) provide an even sharper critique by arguing that predictability was largely an in-sample or *ex-post* phenomenon which disappears once the forecasting models are used to guide forecasts on new, out-of-sample, data. Rapach and Zhou (2012) provide an extensive review of this literature.

To evaluate our approach empirically, we consider the large set of predictor variables used by Welch and Goyal (2008). When implemented along the lines proposed in our paper, we find that for nearly all of the predictors and at both the monthly, quarterly and annual frequencies, both the equity premium (EP) and Sharpe ratio (SR) constraints lead to substantial improvements in the predictive accuracy of the equity premium forecasts. Across all variables, we find that when comparing the unconstrained to the equity constrained forecasts, the average out-of-sample R^2 improves from -0.53% to 0.19% at the monthly frequency, from -0.23% to 0.47% at the quarterly frequency and from -5.27% to 3.10% at the annual frequency. Similarly, comparing the unconstrained to the Sharpe ratio constrained forecasts, the out-of-sample R^2 improves from -0.53% to 0.18% at the monthly frequency, from -0.23% to 1.02% at the quarterly frequency and from -5.27% to 3.86% at the annual frequency. Hence, the improvement in predictive accuracy tends to get larger as the forecast horizon is extended and the effect of estimation error in a conventional unconstrained model gets stronger.

We also consider the economic value of using constrained forecasts in the portfolio allocation of a representative investor endowed with power utility. In the benchmark case with a coefficient of relative risk aversion of five, we compare the certainty equivalent return (CER) obtained from using a given predictor relative to the prevailing mean model. The comparison is conducted for the unconstrained as well as the EP-constrained and the SR-constrained cases at monthly, quarterly, and annual horizons, for the entire sample and a few subsamples. Here again, we find that the economic constraints lead to higher CER-values at all horizons and across practically all predictors (the one exception being the stock variance). Specifically, the EP constraint results in a higher CER (relative to the unconstrained case) of about 50 basis points per year, whereas for the SR-constrained models, the increase is about 100 basis points per year. Consistent with the predictive accuracy results, we generally find that the SR constraint produces higher CER improvements than the EP constraint, which suggests that there are economically important

interactions between the estimated mean and volatility. Robustness checks reveal that a higher (lower) risk aversion coefficient of 10 (2) reduces (increases) the spread in performance across models, as the investor’s willingness to exploit any predictability is inversely proportional to the risk aversion. Even in the most conservative (i.e., high risk aversion) case, we observe an increase in CER of 20 to 30 basis points for the EP-constrained case and 50 to 70 basis points for the SR constraints across predictors.

The previous results refer to univariate regression models with a single predictor variable. We also consider two ways to incorporate multivariate information. First, we consider equal-weighted forecast combinations. Consistent with Rapach et al. (2010), we find that simple forecast combinations improve on the average forecast performance, particularly for the unconstrained forecasts which are most adversely affected by parameter estimation error. Second, we consider a diffusion index approach that extracts common components from the cross-section of predictor variables followed by unconstrained or constrained equity premium predictions using these components. Empirically, the diffusion index approach produces better statistical and economic performance than the equal-weighted combination approach both across subsamples and in the full sample. Moreover, this approach works best for the economically constrained models. For example, at the quarterly horizon, the out-of-sample R^2 of the diffusion index is 0.42%, 3.02%, and 2.95% for the unconstrained, EP constrained, and SR constrained models, respectively, with associated CER-values of -0.04%, 0.53%, and 0.95% per annum.

The plan of the paper is as follows. Section 2 introduces our new methodology for efficiently incorporating theoretical constraints on the predictive moments of the equity premium distribution. Section 3 introduces the data and presents empirical results for both unconstrained and constrained prediction models using a range of predictor variables. Section 4 evaluates the economic value of imposing economic constraints on the forecasts. Section 5 presents an extension to incorporate multivariate information and conducts a range of robustness tests, and Section 6 concludes.

2 Methodology

This section describes how we estimate and forecast the equity premium subject to constraints motivated by economic theory. These constraints take the form of inequalities on the conditional equity premium or bounds on the conditional Sharpe ratio.

2.1 Economic Constraints on the Return Prediction Model

It is common practice in the literature on return predictability to assume that the conditional mean of stock returns, measured in excess of a risk-free rate, $r_{\tau+1}$, is a linear function of lagged

predictor variables, x_τ :

$$\begin{aligned} r_{\tau+1} &= \mu + \beta x_\tau + \varepsilon_{\tau+1}, \quad \tau = 1, \dots, t-1, \\ \varepsilon_{\tau+1} &\sim N(0, \sigma_\varepsilon^2). \end{aligned} \tag{1}$$

The linear model is simple to interpret and only requires estimating two mean parameters, μ and β , which can readily be accomplished by OLS.

Economic theory generally does not restrict the functional form of the mapping linking predictor variables, x_τ , to the conditional mean of excess returns, $r_{\tau+1}$, so the use of the linear specification in (1) should be viewed as an approximation. However, we argue that economically motivated constraints can be used to improve on this model. We next consider two such constraints.

2.1.1 Equity Premium Constraint

Under broad conditions the conditional equity risk premium can be expected to be positive.⁵ This reasoning implies a constraint on the predictive moments of the distribution of excess returns. In turn, this has implications for the estimated parameters of the return prediction model (1). Specifically, to efficiently exploit the information embedded in the constraint that the conditional equity premium is non-negative, the parameters μ and β should be estimated subject to the constraint $\mu + \beta x_\tau \geq 0$ at all points in time:

$$\mu + \beta x_\tau \geq 0 \quad \text{for } \tau = 1, \dots, t. \tag{2}$$

Although this constraint on the predictive moments of the equity premium is not directly a constraint on the model parameters, $\theta = (\mu, \beta, \sigma_\varepsilon^2)$, it clearly affects these parameters since they have to be consistent with (2). Moreover, because the conditional EP constraint has to hold at each point in time, the number of constraints grows in proportion with the length of the sample size. The seemingly simple EP constraint in (2) therefore potentially yields a very powerful way to pin down the parameters of the return forecasting model and obtain more precise estimates.

To see how the constraint in (2) works to restrict the $\mu - \beta$ parameter space, consider Figure 2. Panel (a) shows how different values of x , when its values are always negative (log dividend yield case), constrains the feasible set of μ and β . Panel (b) repeats this exercise when x only takes on positive values (T-bill case), whereas panel (c) illustrates the case with a predictor that can take on both negative and positive values (log dividend payout ratio case). These graphs illustrate that whenever a new observation of x arrives, both small and large values of this predictor can lead to new constraints on the set of feasible parameter values. Moreover, there will be T constraints on the parameters in a sample with T observations.

⁵For example, this rules out that stocks hedge against other risk factors affecting the performance of a market portfolio that comprises a broader set of asset classes.

Campbell and Thompson (2008) (CT, henceforth) were the first to argue in favor of imposing a non-negative EP constraint.⁶ They implement this idea by using a truncated forecast which is simply the largest of the unconstrained OLS forecast and zero:

$$\hat{r}_{t+1|t} = \max(0, \hat{\mu}_t + \hat{\beta}_t x_t), \quad (3)$$

where $\hat{\mu}_t$ and $\hat{\beta}_t$ are the OLS estimates from (1), i.e.,

$$(\hat{\mu}_t \quad \hat{\beta}_t)' = \left(\sum_{\tau=1}^{t-1} z_{\tau} z_{\tau}' \right)^{-1} \left(\sum_{\tau=1}^{t-1} z_{\tau} r_{\tau+1} \right), \quad (4)$$

and $z_{\tau} = (1 \ x_{\tau})'$. This truncation prevents the predicted equity premium from becoming negative, but the theoretical constraint is not used by CT to obtain improved estimates of μ and β in the manner reflected in Figure 2. Specifically, CT simply overrule the forecast if it is negative and do not impose on their parameters that $\hat{r}_{\tau+1|t} = \hat{\mu}_t + \hat{\beta}_t x_{\tau} \geq 0$ for $\tau = 1, \dots, t$. While potentially an improvement over the simple unconstrained model, this approach therefore does not make efficient use of the theoretical constraints in (2).

The linear-normal prediction model implies that the x -variables have unbounded support. We do not take this implication literally, and instead view this model as an approximation. We assume that investors only impose the EP constraint conditional on the data they have seen up to a given point in time, $\tau = 1, \dots, t$. This makes the length of the initial data sample important. Our implementation assumes a long (20-year) warm-up sample, which ensures that investors will have seen a wide range of values for x_{τ} before making their first prediction. It also ensures that new observations on the predictors within the historically observed range do not tighten the constraints. Conversely, observations on the predictors outside the historical range will trigger new learning dynamics, which we think is an attractive feature of our setup.

2.1.2 Sharpe Ratio Constraint

In this section, we explore a novel way of sharpening the forecasts of excess market returns, namely, by placing constraints on the conditional Sharpe ratio of the market portfolio. Such constraints might be motivated from an asset pricing perspective, as the Sharpe ratio is frequently used in the calibration and evaluation of structural asset pricing models.⁷ In US data, it is well-known that the Sharpe ratio is time-varying and countercyclical (Brandt (2010), Lettau and Ludvigson (2010)). More importantly, the empirical Sharpe ratio is quite a bit more

⁶Prior to this, some papers tested non-negativity of the equity premium. For example, Ost diek (1998) studies sign restrictions on the ex-ante equity premium and develops tests for whether this premium is non-negative using a conditional multiple inequality approach.

⁷See Cochrane (2001) for a textbook treatment of the Sharpe ratio's use in evaluating asset pricing models. Lettau and Ludvigson (2010) review whether some leading asset pricing models can replicate the stylized facts regarding the Sharpe ratio in the US. Lettau and Wachter (2007) and Lettau and Wachter (2011) use the Sharpe ratio in the calibration of their asset pricing model.

volatile than what the leading asset pricing models would suggest. This empirical fact has been labeled the “Sharpe ratio variability puzzle” by Lettau and Ludvigson (2010). Naturally, the Sharpe ratio is most often used for portfolio performance evaluation (see Brandt (2010) for a review article). Given all the theoretical and empirical work on this subject, most academics and practitioners are likely to have some priors about what constitutes a “reasonable” Sharpe ratio.

The conditional Sharpe ratio depends on both the conditional mean and volatility of the return distribution. Since time-variation in volatility is a well documented fact in empirical finance (see, e.g., Andersen et al. (2006)), we modify (1) as follows:

$$r_{\tau+1} = \mu + \beta x_{\tau} + \exp(h_{\tau+1}) u_{\tau+1}, \quad (5)$$

where $h_{\tau+1}$ denotes the (log of) return volatility at time $\tau + 1$ and $u_{\tau+1} \sim N(0, 1)$. Following common stochastic volatility models, log-volatility is assumed to evolve as a driftless random walk,

$$h_{\tau+1} = h_{\tau} + \xi_{\tau+1}, \quad (6)$$

where $\xi_{\tau+1} \sim N(0, \sigma_{\xi}^2)$ and u_{τ} and ξ_s are mutually independent for all τ and s .

Next, define the (annualized) conditional Sharpe ratio at time τ as

$$SR_{\tau+1|\tau} = \frac{\sqrt{H}(\mu + \beta x_{\tau})}{\exp(h_{\tau} + 0.5\sigma_{\xi}^2)}, \quad (7)$$

where H denotes the number of observations per year (i.e., $H = 12, 4$, and 1 with monthly, quarterly, and annual data, respectively). We assume that the conditional Sharpe ratio is bounded both from below and above at all points in time:

$$SR^l \leq SR_{\tau+1|\tau} \leq SR^u \quad \text{for } \tau = 1, \dots, t. \quad (8)$$

While (8) does not directly impose restrictions on the model parameters, $\theta = (\mu, \beta, \sigma_{\xi}^2)$ and the sequence of log return volatilities $h^t \equiv \{h_1, h_2, \dots, h_t\}$, it does so indirectly since not all parameter values are consistent with the SR constraint (8). Also, from (7) and (8), it is immediately clear that the SR constraint in effect imposes a time-varying upper bound on the equity premium which is proportional to the conditional volatility.

In the empirical implementation below, we set the lower bound at $SR^l = 0$, which is consistent with the EP constraint (2) augmented to account for time-varying volatility. Annualized values of $SR_{t+1|t}$ around 0.5 are seen as “normal” in the context of the market portfolio, given estimates of its mean and volatility (e.g., Cochrane (2001) and Brandt (2010)). Sharpe ratios higher than one are highly improbable for a non-leveraged market portfolio, so we accordingly

set $SR^u = 1$.⁸ By letting the constraint $[0, 1]$ be relatively wide, we accommodate the fact that Sharpe ratios are imprecisely estimated (Jobson and Korkie (1981)) and implicitly allow a large set of asset pricing models—consumption and non-consumption-based—to be consistent with it.⁹ We next explain how we estimate the econometric models and impose the constraints.

2.2 Priors

Theoretical constraints such as (2) and (8) are naturally interpreted as reflecting the forecaster’s prior beliefs on return predictability. Viewed in this way, they can best be imposed using Bayesian techniques and this is the approach followed here. Moreover, a major advantage of our Bayesian approach is that we obtain the full predictive densities of returns in a way that accounts for parameter estimation error. Such densities are vastly more informative than point forecasts of excess returns based on conventional plug-in least squares estimates.

We begin by describing the choices of priors, starting from the case where no constraints are imposed. Next, we show how to incorporate constraints on the predictive moments of the return distribution.

Following standard practice¹⁰, the priors for the parameters μ and β in (1) are assumed to be normal and independent of σ_ε^2 ,

$$\begin{bmatrix} \mu \\ \beta \end{bmatrix} \sim N(\underline{b}, \underline{V}), \quad (9)$$

where

$$\underline{b} = \begin{bmatrix} \bar{r}_t \\ 0 \end{bmatrix}, \quad \underline{V} = \begin{bmatrix} \psi^2 s_{r,t}^2 & 0 \\ 0 & \psi s_{r,t}^2 / s_{x,t}^2 \end{bmatrix}, \quad (10)$$

with data based moments

$$\begin{aligned} \bar{r}_t &= \frac{1}{t-1} \sum_{\tau=1}^{t-1} r_{\tau+1}, \quad s_{r,t}^2 = \frac{1}{t-2} \sum_{\tau=1}^{t-1} (r_{\tau+1} - \bar{r}_t)^2, \\ \bar{x}_t &= \frac{1}{t-1} \sum_{\tau=1}^{t-1} x_\tau, \quad s_{x,t}^2 = \frac{1}{t-2} \sum_{\tau=1}^{t-1} (x_\tau - \bar{x}_t)^2. \end{aligned}$$

Here $\underline{\psi}$ is a constant that controls the tightness of the prior, with $\underline{\psi} \rightarrow \infty$ corresponding to a diffuse prior on μ and β . Our benchmark analysis sets $\underline{\psi} = 2.5$, but we also consider alternative specifications with both lower and higher values of $\underline{\psi}$. The terms $s_{r,t}^2$ and $s_{r,t}^2 / s_{x,t}^2$ in the diagonal of the prior variance, \underline{V} , are scaling factors that are introduced to guarantee comparability of

⁸Setting the upper bound much higher than one, e.g., at 1.5, means that this bound does not bind very often and so the SR constraint becomes very similar to the EP constraint.

⁹Lettau and Ludvigson (2010) show that many of the leading consumption-based asset pricing models cannot generate the volatility that is observed in empirically estimated Sharpe ratios. Lettau and Wachter (2007) and Lettau and Wachter (2011) depart from the consumption-based asset pricing models to accommodate pricing kernels with higher conditional volatility which better fit the dynamic behavior of the Sharpe ratio.

¹⁰See for example Koop (2003), section 4.2.

the priors across different predictors and across different data frequencies.¹¹ Our choice of the prior mean vector \underline{b} reflects the “no predictability” view that the best predictor of stock 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.

Next, we specify a gamma prior for the error precision of the return innovation, σ_ε^{-2} :

$$\sigma_\varepsilon^{-2} \sim G(s_{r,t}^{-2}, \underline{v}_0(t-1)), \quad (11)$$

where \underline{v}_0 is a prior hyperparameter that controls the degree of informativeness of this prior, with $\underline{v}_0 \rightarrow 0$ corresponding to a diffuse prior on σ_ε^{-2} .¹² Our benchmark sets $\underline{v}_0 = 0.1$, which, loosely speaking, means that the prior weight is approximately 10% of the weight put on the data.

The SR constraint (8) requires specifying a joint prior for the sequence of log return volatilities, h^t , and the error precision, σ_ξ^{-2} . Writing $p(h^t, \sigma_\xi^{-2}) = p(h^t | \sigma_\xi^{-2}) p(\sigma_\xi^{-2})$, it follows from (6) that

$$p(h^t | \sigma_\xi^{-2}) = \prod_{\tau=1}^{t-1} p(h_{\tau+1} | h_\tau, \sigma_\xi^{-2}) p(h_1), \quad (12)$$

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

$$h_1 \sim N(\ln(s_{r,t}), \underline{k}_h), \quad (13)$$

$$\sigma_\xi^{-2} \sim G(1/\underline{k}_\xi, 1). \quad (14)$$

We set $\underline{k}_\xi = 0.01$ and choose the remaining hyperparameters in (13) and (14) to imply uninformative priors, allowing the data to determine the degree of time variation in the return volatility. Accordingly, we specify $\underline{k}_h = 10$, and set the degrees of freedom for σ_ξ^{-2} to 1.

We next describe how we impose the economic constraints on the model parameters. Starting with the EP constraint, we modify the priors on μ and β in (9) to

$$\begin{bmatrix} \mu \\ \beta \end{bmatrix} \sim N(\underline{b}, \underline{V}), \quad \mu, \beta \in A_t, \quad (15)$$

where A_t is a set such that

$$A_t = \{\mu + \beta x_\tau \geq 0, \tau = 1, \dots, t\}. \quad (16)$$

¹¹This approach is used routinely in macroeconomic Bayesian VAR models. See for example Kadiyala and Karlsson (1997) and Banbura et al. (2010).

¹²Following Koop (2003), we adopt the Gamma distribution parametrization of Poirier (1995). Nameley, if the continuous random variable Y has a Gamma distribution with mean $\mu > 0$ and degrees of freedom $v > 0$, we write $Y \sim G(\mu, v)$. Then, in this case, $E(Y) = \mu$ and $Var(Y) = 2\mu^2/v$.

Similarly, for the SR constraint, we modify the priors on μ, β to ensure

$$\begin{bmatrix} \mu \\ \beta \end{bmatrix} \sim N(\underline{b}, \underline{V}), \mu, \beta \in \tilde{A}_t, \quad (17)$$

where \tilde{A}_t is a set satisfying

$$\tilde{A}_t = \left\{ SR^l \leq SR_{\tau+1|\tau} \leq SR^u, \tau = 1, \dots, t \right\}, \quad (18)$$

and $SR_{\tau+1|\tau}$ is given in (7).

The Appendix provides details of how we estimate the parameters and compute forecasts for the unconstrained and constrained models. Section 5 discusses robustness of our results with respect to changes in the priors.

As a final point about the above analysis, we note that the boundaries of the constraints (2) and (8) are constants (0 , SR^l , and SR^u), motivated by economic considerations. However, one might view the boundaries themselves as being parameters with associated priors. In that case, our specification corresponds to having dogmatic priors on these specific parameters. This generalization might be less meaningful for constraints that are readily imposed by economic theory (such as the zero lower bound on the equity premium and Sharpe ratio) than for others (such as the upper bound on the Sharpe ratio). From an econometric perspective, updating priors about the boundary parameters is non-trivial. Given that the benefits of such a generalization are not clear, while the tractability and computational costs of imposing it are substantial, we conduct our empirical analysis by imposing constraints (2) and (8) as discussed above.

3 Empirical Results

This section presents the data and empirical results using the methods for incorporating economic constraints described in Section 2 to predict the equity premium.

3.1 Data

Our empirical analysis uses data on stock returns along with a set of seventeen predictor variables originally analyzed in Welch and Goyal (2008) and subsequently extended up to 2010 by the same authors. Stock returns are computed from the S&P500 index and include dividends. A short T-bill rate is subtracted from stock returns in order to capture excess returns. Data samples vary considerably across the individual predictor variables. To be able to compare results across the individual predictor variables, we use the longest common sample which is 1927-2010. In addition, we use the first 20 years of data as a training sample. For example, for the monthly data we initially estimate our regression models over the period January 1927–December 1946, and use the estimated coefficients to forecast excess returns for January 1947. We next include

January 1947 in the estimation sample, which thus becomes January 1927–January 1947, and use the corresponding estimates to predict excess returns for February 1947. We proceed in this recursive fashion until the last observation in the sample, thus producing a time series of one-step-ahead forecasts spanning the time period from January 1947 to December 2010.

The identity of the predictor variables, along with summary statistics, is provided in Table 1. Most variables fall into three broad categories, namely (i) valuation ratios capturing some measure of ‘fundamentals’ to market value such as the dividend price ratio, the dividend yield, the earnings-price ratio, the 10-year earnings-price ratio or the book-to-market ratio; (ii) measures of bond yields capturing level effects (the three-month T-bill rate and the yield on long term government bonds), slope effects (the term spread), and default risk effects (the default yield spread defined as the yield spread between BAA and AAA rated corporate bonds, and the default return spread defined as the difference between the yield on long-term corporate and government bonds); (iii) estimates of equity risk such as the long term return and stock variance (a volatility estimate based on daily squared returns). Finally, three corporate finance variables, namely the dividend payout ratio (the log of the dividend-earnings ratio), net equity expansion (the ratio of 12-month net issues by NYSE-listed stocks over the year-end market capitalization), percent equity issuing (the ratio of equity issuing activity as a fraction of total issuing activity) and a macroeconomic variable, inflation (the rate of change in the consumer price index), are considered.¹³

To make our results comparable to studies from the literature on return predictability such as Campbell and Thompson (2008) and Welch and Goyal (2008), we focus on univariate regressions with a single predictor variable. However, we also discuss in Section 5 how our approach can be extended to incorporate multivariate information. Finally, since there are too many variables to cover in detail, we focus our analysis on three predictors, namely the log dividend-price ratio, the T-bill rate, and the default yield spread, all of which have featured prominently in the literature on return predictability.

3.2 Coefficient Estimates and Predictive Densities

The economic constraints on the predictive moments of the return distribution affect the parameter estimates in a way that reflects the entire sequence of data points. This gives rise to parameter estimates that are very different from the standard, unconstrained ones typically applied in the literature on return predictability. To better understand the effect of the constraints, we begin by studying the posterior distribution of the parameter estimates.

Figure 3 plots the posterior density for the slope coefficient, β , in the equity premium equation (1) using either the log dividend-price ratio (top panel), the T-bill rate (middle), or the

¹³We follow Welch and Goyal (2008) and, for monthly and quarterly data, lag inflation an extra period to account for the delay in CPI releases.

default yield spread (bottom) as predictors. For all predictors, the posterior densities are displayed for the unconstrained case (solid line), the EP constraint (dark dash-dotted line), and the SR constraint (light dark-dotted line). In each case, the unconstrained posterior density for β is considerably wider than those of the constrained densities, suggesting that the economic constraints reduce parameter uncertainty. Moreover, whereas the unconstrained posterior densities are symmetric, the constrained ones are asymmetric in a direction that mostly reflects that the equity premium has to be non-negative. For example, for the log dividend price ratio, which is always negative, the EP constraint rules out large positive values of β , which could otherwise induce a negative equity premium. Conversely, the constrained posterior distributions rule out large negative values of β for variables that take on positive values such as the T-bill rate and the default yield spread. The upper bound on the Sharpe ratio also matters for the posterior distribution of β , however, which helps explain why for positive predictors such as the T-bill rate the posterior distribution of β under the SR constraint is shifted to the left compared with its distribution under the EP constraint.¹⁴

To evaluate the economic significance of the changes in the parameter estimates caused by the constraints, we next compare the ex-ante equity premium under the unconstrained and constrained models. To this end, Figures 4-6 show the predictive densities for the equity premium, computed as of the end of the sample (December 2010). To illustrate how expected returns depend on the value taken by the predictor, we show the predictive densities conditional on $x_T = \bar{x}$ as well as $x_T = \bar{x} \pm 2 \times SE(x)$, where \bar{x} and $SE(x)$ are the full-sample average and standard deviation of x , respectively.

First consider the results based on the log dividend-price ratio, $\log(D/P)$ (Figure 4). This predictor is always negative and the associated posterior estimates of β are centered on a positive value. Comparing the plots for the three values of x illustrates how the constraints work. When $\log(D/P)$ is set at its sample mean (top panel), the three posterior densities have comparable spreads, although the unconstrained model has a lower mean than the EP constrained and SR constrained models. Reducing the log dividend-price ratio to two standard errors below its mean (middle panel) results in a very different picture. The unconstrained posterior density for the equity premium is now much more dispersed and shifted far further to the right, whereas the two constrained forecasts have more probability mass to the right of zero with a tighter support. When $\log(D/P)$ is very low (bottom panel), the lower bounds imposed by the EP and SR constraints bind, thus preventing the probability mass from shifting to the left which otherwise happens mechanically in a linear model (as can be seen for the unconstrained forecast). This case is empirically relevant for the 1960s and 1990 to 2005 periods with abnormally low log

¹⁴Differences between the restricted densities do not always occur in the tail that one would expect. This happens because the upper constraint can be satisfied by simultaneously reducing large negative slope coefficients (as in the T-bill rate model) and shifting the density for the intercept, μ , to the right.

dividend price ratios. Conversely, when $\log(D/P)$ is very high (bottom panel), the constraints are less likely to bind, and so the three densities are more similar in shape, although once again the centers of the distributions clearly differ.

For the T-bill rate (Figure 5), we see similar mechanisms at work, although now with the opposite sign since the T-bill rate is always positive and the posterior estimates of β are centered on a negative value. This means that the lower constraints now bind when the T-bill rate is set at $\bar{x} + 2 \times SE(x)$ (bottom panel), once again leading to much tighter distributions under the EP and SR constraints than for the unconstrained case. Empirically, this occurred in the early 1980s, when the T-bill rate was particularly high. Finally, the model based on the default yield spread (Figure 6), shows less of an asymmetry across the three conditioning scenarios regarding the shape and spread of the conditional posterior density estimates of the equity premium.

These figures imply that the economic constraints tighten the predictive density for the equity premium in a manner that depends asymmetrically on whether the predictor variables take on large negative or positive values. Hence, how “informative” the bounds are, i.e., by how much they shift and tighten the posterior density, depends on the value taken by the predictor variable, x . We illustrate this effect in Figure 7 for the plots based on the T-bill rate.¹⁵ The top panel plots the posterior mean of the equity premium distribution as a function of the T-bill rate. The posterior mean declines linearly for the unconstrained model from a level near 1% per month for the lowest values of the T-bill range to a level near zero for the highest values.¹⁶ Under the SR and EP constrained models, the posterior mean is also reduced as the T-bill rate increases, but by far less than under the unconstrained model.

Turning to the uncertainty surrounding the predicted equity premium, the posterior variance of the equity premium distribution (bottom panel) is large and rises sharply under the unconstrained model as the T-bill rate moves far away from its sample average. In contrast, while the posterior variance of the constrained equity premium distributions does rise when the T-bill rate takes on very small or very large values, it does so at a far slower rate. For example, for very high values of the T-bill rate, the posterior variance of the equity premium under the unconstrained model is close to four times higher than under the constrained models.

3.3 Forecasts of Equity Premia

Using these insights into how economic constraints affect forecasts of equity premia, we next study the sequence of recursively generated out-of-sample equity premium forecasts. To this end, Figure 8 presents monthly values of the mean of the predictive distribution of the equity premium over the period 1947-2010. Economic constraints clearly make a substantial difference

¹⁵The plots for the log dividend-price ratio and the default yield spread are very similar and so are omitted here.

¹⁶Consistent with Figure 5, the T-bill rate varies between $\bar{x} - 2 \times SE(x)$ and $\bar{x} + 2 \times SE(x)$, with \bar{x} and $SE(x)$ denoting the full-sample average and standard deviation of the T-bill rate, respectively.

during most periods. For example, the unconstrained model forecasts based on the log-dividend price ratio (top panel) are lower and far more volatile than their constrained counterparts and turn negative for most of the period between 1990 and 2005. Interestingly, even though none of the recursive forecasts from the unconstrained model turn negative prior to 1960, the constrained forecasts are quite different prior to this period. This happens due to our requirement that the entire sequence of model-implied fitted equity premia be non-negative, and so clearly this condition is violated by the unconstrained model prior to 1960 (see Figure 1). The economic constraints lead to predicted equity premia whose differences from the unconstrained counterparts can last very long, e.g., from 1955 through to 1975 and again from around 1985 to the end of the sample.

Large and persistent differences in predicted mean returns are also found for the return model based on the T-bill rate (middle panel). For this model, negative values of the unconstrained forecasts occur most of the time between 1970 and 1985, whereas the constrained forecasts hover around small, but positive values throughout the sample. The SR constrained forecasts are smaller than the EP constrained forecasts for long periods of time, and both series are notably more stable than the unconstrained equity premium forecasts.

The unconstrained equity premium forecasts based on the model that uses the default yield spread as a predictor (bottom panel) only turn negative during the first few months of the sample and are otherwise quite similar to the mean forecasts from the EP constrained model which in turn are smaller than the SR constrained forecasts. These results are consistent with our earlier findings that the constraints tend to bind on fewer occasions for this predictor variable.

Figure 9 plots monthly volatility forecasts based on the stochastic volatility model (6). We only present results for a single predictor (the log dividend-price ratio) since results are very similar across different predictors. Volatility hovers around 5% per month, but spikes notably in 1975, after October 1987, and during the global financial crisis at the end of the sample.

Conditional Sharpe ratios are plotted in Figure 10. For the unconstrained model that assumes constant volatility, these plots essentially mirror the movements in expected returns in Figure 8. Conversely, the SR constrained forecasts fluctuate both because of variations in expected returns and in the conditional volatility. Interestingly, the spikes in the posterior mean forecasts seen around 1975—and again after 1980 under the default yield spread model—are not mirrored in the conditional Sharpe ratio plot for this variable (bottom panel), showing that there is an interaction between the estimated mean and volatility.

Figure 7 showed that the posterior volatility of the equity premium forecasts tends to be smaller under the two constrained models than under the unconstrained model. This has important consequences for the time-series of forecasts. To illustrate this, Figure 11 shows 95% posterior probability intervals for μ and β for the unconstrained and EP constrained models that

use the T-bill rate as a predictor.¹⁷ We focus on the period between 1965 and 1985 to better see the effect of specific events on the parameter estimation uncertainty. It is quite clear from these plots that the EP constraint reduces the uncertainty about β more than it does for μ . Moreover, the high T-bill rates during the Fed’s “Monetarist Experiment” from 1979-82 clearly reduce the width of the confidence interval for the constrained model, but not for the unconstrained model.

3.4 Out-of-Sample Predictive Performance

We next evaluate the predictive accuracy of the equity premium forecasts. As in Welch and Goyal (2008) and Campbell and Thompson (2008), the predictive performance of each model is measured relative to the prevailing mean model. The inputs to the analysis are the time series of predictive densities of excess returns obtained as described in Section 2. To simplify the exposition, let $\{r_{t+1}^j\}$, $j = 1, \dots, J$, denote draws from the predictive density of excess returns for the prevailing mean model, conditional on data known at time t . Further, let $\{r_{t+1,i}^j\}$, $j = 1, \dots, J$, be draws from the predictive density of excess returns for the model based on the i th predictor, again conditional on data known at time t . As further explained in the Appendix, for the unconstrained and EP constrained models, these draws are obtained by applying a Gibbs sampler to

$$p(r_{t+1}|\mathcal{D}^t) = \int_{\mu, \beta, \sigma_\varepsilon^{-2}} p(r_{t+1}|\mu, \beta, \sigma_\varepsilon^{-2}, \mathcal{D}^t) p(\mu, \beta, \sigma_\varepsilon^{-2}|\mathcal{D}^t) d\mu d\beta d\sigma_\varepsilon^{-2}, \quad (19)$$

where $\mathcal{D}^t = \{r_{\tau+1}, x_\tau\}_{\tau=1}^{t-1} \cup x_t$ is the information set at time t . Likewise, for the SR constrained model, return draws are based on the predictive density

$$\begin{aligned} p(r_{t+1}|\mathcal{D}^t) &= \int_{\mu, \beta, h^{t+1}, \sigma_\xi^{-2}} p(r_{t+1}|h_{t+1}, \mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t) \\ &\quad \times p(h_{t+1}|\mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t) \\ &\quad \times p(\mu, \beta, h^t, \sigma_\xi^{-2}|\mathcal{D}^t) d\mu d\beta dh^{t+1} d\sigma_\xi^{-2}, \end{aligned} \quad (20)$$

where h^{t+1} denotes the sequence of conditional variance states up to time $t + 1$.

To compare our results with conventional performance measures used in the literature (see, e.g., Welch and Goyal (2008), Campbell and Thompson (2008), and Rapach and Zhou (2012)), we compute the posterior mean from the densities in (19) or (20) to obtain point forecasts. Specifically, define time t forecast errors for the prevailing mean model and the model based on

¹⁷These posterior probability intervals (sometimes referred to as credible intervals) represent the probability that a parameter falls within a given region of the parameter space, given the observed data. So, for example, the (2.5, 97.5)% posterior probability interval represents the compact region of the parameter space for which there is a 2.5% probability that the parameter is higher than the region’s upper bound, and a 2.5% probability that it is lower than the region’s lower bound.

predictor i as

$$e_t = r_t - \frac{1}{J} \sum_{j=1}^J r_t^j, \quad t = \underline{t}, \dots, \bar{t}, \quad (21)$$

$$e_{t,i} = r_t - \frac{1}{J} \sum_{j=1}^J r_{t,i}^j, \quad t = \underline{t}, \dots, \bar{t}, \quad (22)$$

where \underline{t} and \bar{t} denote the beginning and the end of the forecast evaluation period, respectively.

The period- t difference in the cumulative sum of squared errors (SSE) between the prevailing mean and the i th predictor model is then equal to

$$\Delta CumSSE_t = \sum_{\tau=\underline{t}}^t e_{\tau}^2 - \sum_{\tau=\underline{t}}^t e_{\tau,i}^2, \quad (23)$$

while the out-of-sample R^2 is

$$R_{OoS,i}^2 = 1 - \frac{\sum_{\tau=\underline{t}}^{\bar{t}} e_{\tau,i}^2}{\sum_{\tau=\underline{t}}^{\bar{t}} e_{\tau}^2}. \quad (24)$$

Table 2 presents the out-of-sample R^2 for the unconstrained, EP constrained and SR constrained forecasts estimated on monthly data. Out of the 16 unconstrained forecast models, 12 produce negative R_{OoS}^2 . In contrast, the EP constrained monthly forecasts only generate a negative R_{OoS}^2 for three of the 16 variables whereas the SR constrained models generate a negative R_{OoS}^2 for six variables. Compared with the unconstrained forecasts, the EP and SR constrained forecasts lead to a higher R_{OoS}^2 for 14 out of 17 variables. This is also reflected in the average R_{OoS}^2 computed across the univariate prediction models which is -0.53% for the unconstrained models, 0.19% for the EP constrained model, and 0.18% for the SR constrained models. Notable improvements are seen for the models based on valuation ratios such as the dividend yield or earnings-price ratio.

Panels B and C in Table 2 show that the improvement in forecast performance resulting from imposing the economic constraints carries over to the two subsamples 1947-1978 and 1979-2010, obtained by splitting the forecast evaluation period in two halves. In the first subsample, the average improvement in the R_{OoS}^2 -values is between 0.60% and 0.70% (from -0.17% for the unconstrained to 0.44% and 0.55% for the EP and SR constrained models, respectively). It is a slightly better 0.70%-0.80% in the second subsample (from -0.80% for the average unconstrained model to 0.01% and -0.10% for the EP and SR constrained models).

For the quarterly models (Table 3), the benefits from imposing economic constraints on the equity premium forecasts get even bigger. At this frequency, we find that the EP and SR constrained forecasts generate a higher R_{OoS}^2 for 14 out of 15 predictors. Moreover, whereas the average R_{OoS}^2 is -2.33% for the unconstrained model, it is 0.47% and 1.02% for the EP and SR

constrained models, respectively. Again, notable improvements are seen for the models based on valuation ratios such as the dividend yield or earnings-price ratio. Improvements in the average R_{OoS}^2 due to imposing economic constraints again carry over to the two subsamples and exceed 2.2% in the first subsample (1947-1978) and 3.2% in the second subsample, although the latter reflects a clear deterioration in the performance of the unconstrained model during the period 1979-2010.

Turning to the annual results, Table 4 shows that 14 of the 16 unconstrained models generate a negative R_{OoS}^2 , the average R_{OoS}^2 being -5.27%. In contrast, all of the constrained forecasts generate a positive R_{OoS}^2 , in each case higher than that of the corresponding unconstrained model.¹⁸ Moreover, the average R_{OoS}^2 computed across the 16 prediction models tends to be quite high: 3.10% for the EP constrained models and 3.86% for the SR constrained models. Once again, imposing the constraints lowers the probability of very poor forecast performance. For example, the lowest R_{OoS}^2 -value of any unconstrained model is -16.2% in the annual data, versus 0.07% for the EP constrained model and 2.99% for the SR constrained models. Moreover, the improvements in the average R_{OoS}^2 -value carries over to the two subsamples and exceeds 6.7% in all cases.

Following Rapach et al. (2010), we use stars in tables 2-4 to indicate the statistical significance of pair-wise differences in the predictive accuracy between a given forecasting model and the benchmark model based on the Clark and West (2007) p -values.¹⁹ Economically constrained models appear to produce significantly better return forecasts than the unconstrained forecasts for most of the valuation ratios and many of the interest rate variables. Moreover, the results tend to get stronger at the quarterly and annual forecast horizons.

The results in tables 2-4 indicate that the superior performance of the constrained forecasts relative to the prevailing mean tends to strengthen as the forecast horizon grows from monthly via quarterly to annual, whereas the opposite happens for the unconstrained forecasts. Two effects are at play here. On the one hand, the power of the predictive signal tends to increase, the longer the forecast horizon. On the other hand, forecasts become more uncertain at the longer horizons as a result of the fewer data points available for estimation. For the unconstrained models, the second effect clearly dominates and so forecast performance tends to deteriorate as the horizon is extended. Conversely, the economic constraints provide an effective way to deal with parameter estimation error and so the performance of the constrained models improves as we move from the monthly to the annual horizon.

¹⁸The stochastic volatility model (6) is used to capture time-varying volatility at the monthly and quarterly horizons. At the annual horizon we found that there were too few observations to reliably identify the parameters of this model and ensure convergence of the parameter estimates. Instead we use a simple AR(1) specification for the realized variance to model the variance at the annual horizon.

¹⁹Such p -values should be interpreted with caution. In the spirit of Diebold (2012) they can be interpreted as a measure of the relative accuracy of the sequence of forecasts.

To help identify how the prediction models performed in specific periods, Figure 12 presents the time-series of $\Delta CumSSE$ for three of our models. For the model based on $\log(D/P)$ (top panel), the forecast performance of the unconstrained model deteriorates notably between 1995 and 2000 – a period during which this model generated large negative equity premium forecasts although average stock returns were positive. For the model based on the T-bill rate (middle panel), the unconstrained forecasts again tend to be less precise than their constrained counterparts, the main exception being an episode around 1974-75 during which the unconstrained model correctly predicted negative excess returns. Note also the consistently better forecast performance of the SR constrained forecasts compared with the EP constrained forecasts based on the T-bill rate. Finally, for the default yield premium model (bottom panel), the cumulative squared errors of the unconstrained forecasts are almost uniformly worse than those of the constrained forecasts.

In summary, economically motivated constraints on the equity premium predictions lead to substantially better forecast performance at the monthly, quarterly, and annual horizons. They also reduce the risk of selecting a bad forecast model which is important in situations, such as here, characterized by considerable model uncertainty.

4 Economic Performance and Portfolio Choice

So far we have compared the statistical performance of return forecasts generated by economically constrained prediction models to the performance of unconstrained models. We next evaluate the economic significance of these return forecasts by considering the optimal portfolio choice of an investor who uses the return forecasts. An advantage of our approach is that it accounts for parameter estimation error—a point whose importance has been emphasized by Barberis (2000). Moreover, our approach provides the full predictive density which means that we are not reduced to considering only mean-variance utility but can use utility functions such as power utility with better properties.

4.1 Framework

Consider the optimal asset allocations of a representative investor with utility function U . At time t , the investor solves the optimal asset allocation problem

$$\omega_t^* = \arg \max_{\omega_t} E \left[U(\omega_t, r_{t+1}) | \mathcal{D}^t \right], \quad (25)$$

with \mathcal{D}^t denoting all information available up to time t , and $t = \underline{t} - 1, \dots, \bar{t} - 1$. The investor is assumed to have power utility

$$U(\omega_t, r_{t+1}) = \frac{\left[(1 - \omega_t) \exp(r_t^f) + \omega_t \exp(r_t^f + r_{t+1}) \right]^{1-A}}{1 - A}. \quad (26)$$

Here r_t^f is the continuously compounded T-bill rate at time t , while A is the investor's coefficient of relative risk aversion. The t subscript on the portfolio weight reflects that the investor solves the portfolio optimization problem using only information available at time t .

Taking expectations in (25) with respect to the predictive density of r_t , we can rewrite (25) as

$$\omega_t^* = \arg \max_{\omega_t} \int U(\omega_t, r_{t+1}) p(r_{t+1} | \mathcal{D}^t) dr_{t+1}. \quad (27)$$

The integral in (27) can thus be approximated using the draws from the predictive densities as described in Section 2 and in the Appendix. Specifically, under the prevailing mean model, for suitably large values of J the solution to (27) can be approximated by

$$\hat{\omega}_t = \arg \max_{\omega_t} \frac{1}{J} \sum_{j=1}^J \left\{ \frac{\left[(1 - \omega_t) \exp(r_t^f) + \omega_t \exp(r_t^f + r_{t+1}^j) \right]^{1-A}}{1 - A} \right\}. \quad (28)$$

Similarly, the solution to the models with time-varying expected returns, (27), can be approximated by

$$\hat{\omega}_{t,i} = \arg \max_{\omega_t} \frac{1}{J} \sum_{j=1}^J \left\{ \frac{\left[(1 - \omega_t) \exp(r_t^f) + \omega_t \exp(r_t^f + r_{t+1,i}^j) \right]^{1-A}}{1 - A} \right\}, \quad (29)$$

where i indexes the predictor variable.

The sequence of portfolio weights $\{\hat{\omega}_t\}_{t=\bar{t}-1}^{\bar{t}-1}$ and $\{\hat{\omega}_{t,i}\}_{t=\bar{t}-1}^{\bar{t}-1}$ are used to compute the investor's realized utilities under the prevailing mean model and the model based on predictor i . In particular, let \widehat{W}_{t+1} and $\widehat{W}_{t+1,i}$ be the realized wealths at time $t+1$. \widehat{W}_{t+1} and $\widehat{W}_{t+1,i}$ are functions of time $t+1$ realized excess return, r_{t+1} , as well as the optimal allocations to stocks computed in (28) and (29):

$$\begin{aligned} \widehat{W}_{t+1} &= (1 - \hat{\omega}_t) \exp(r_t^f) + \hat{\omega}_t \exp(r_t^f + r_{t+1}), \\ \widehat{W}_{t+1,i} &= (1 - \hat{\omega}_{t,i}) \exp(r_t^f) + \hat{\omega}_{t,i} \exp(r_t^f + r_{t+1,i}). \end{aligned} \quad (30)$$

The certainty equivalent return for the model based on predictor i , CER_i , is defined as the value that equates the average realized utility of the prevailing mean model to the average realized utility of the model based on the i -th predictor, over the forecast evaluation sample:

$$CER_i = \left[\frac{\sum_{\tau=\bar{t}}^{\bar{t}} \widehat{U}_{\tau,i}}{\sum_{\tau=\bar{t}}^{\bar{t}} \widehat{U}_{\tau}} \right]^{\frac{1}{1-A}} - 1, \quad (31)$$

where \widehat{U}_{τ} and $\widehat{U}_{\tau,i}$ denote time τ realized utilities, i.e., $\widehat{U}_{\tau} = \widehat{W}_{\tau}^{1-A}/(1-A)$, $\widehat{U}_{\tau,i} = \widehat{W}_{\tau,i}^{1-A}/(1-A)$.

In addition to evaluating the economic values of the various models over the full forecast evaluation sample, we also study how the different models perform in real time. Specifically, we first calculate the single-period $CER_{t,i}$ as

$$CER_{t,i} = \left[\frac{\widehat{U}_{t,i}}{\widehat{U}_t} \right]^{\frac{1}{1-A}} - 1. \quad (32)$$

To parallel the cumulative SSE measures in (23), we also inspect the economic performance of the individual models by plotting the cumulative sum of CERs over time:²⁰

$$CumCER_{t,i} = \sum_{\tau=t}^t \log(1 + CER_{\tau,i}). \quad (33)$$

4.2 Empirical Results

Turning to the empirical asset allocation results, Table 5 reports annualized CER values for the monthly return regressions computed for an investor with power utility and a coefficient of relative risk aversion, $A = 5$. At the monthly horizon (Panel A), the average CER value, measured relative to the prevailing mean model, is -.12% for the unconstrained models, 0.28% for the EP-constrained models and 0.81% for the SR-constrained models. All but one of the EP-constrained models deliver higher CER values than their unconstrained counterparts, the exception being the stock variance. For the SR-constrained models, the CER values are higher than the corresponding benchmarks across all predictors.

At the quarterly horizon (Panel B), the constrained models retain their higher CER values relative to the unconstrained counterparts for all but one case. The average CER values, computed across all variables, is -0.14% for the unconstrained models, 0.29% for the EP constrained models and 0.32% for the SR constrained models. Finally, at the annual horizon (Panel C), the average CER value is -0.24% for the unconstrained models, 0.33% for the EP constrained models and 0.67% for the SR constrained models and the constrained models produce higher CER values than the unconstrained counterparts for every single predictor.

In Table 6, we show that the observed improvements in economic utility carry over to our two subsamples. There again, the constrained models do better than the unconstrained ones for the vast majority of cases. Interestingly, there is no evidence that the economic benefits from using economically constrained forecasts deteriorates over time. For example, over the subsample 1979-2010 the mean CER value for the annual model is -0.62% for the unconstrained model and 0.08% and 0.76% for the EP and SR constrained models – a bigger differential than in the earlier subsample 1947-78.

²⁰Since the CER is already defined relative to the prevailing mean model, we do not need to compute differential values here.

Using the log dividend-price ratio as a predictor, Figure 13 plots the sequence of stock portfolio weights along with the cumulative (continuously compounded) CER estimates computed according to Equation (33). Under the unconstrained model that uses the log dividend-price ratio the allocations to stocks are generally lower than those under the constrained models. The portfolio weights vary considerably over time under both the unconstrained and SR constrained models, but are much smoother under the EP constraint. Moreover, the cumulative CER values of the constrained models consistently lie above the CER estimates of the unconstrained model. At the end of the sample, the cumulative CER value of the unconstrained model is around -30%, whereas it exceeds 20% and 80% for the EP and SR constrained models, respectively. These numbers capture the cumulative risk-adjusted economic value of the economically constrained forecasts relative to the prevailing mean forecasts.

We conclude that there are economically large benefits from imposing economic constraints on the equity premium prediction models. The benefits appear to be present at monthly, quarterly, and annual horizons and are largest for the SR constraint which allows for time-varying volatility. Moreover, the benefits do not appear to be deteriorating over time.

5 Extensions and Robustness Analysis

This section extends our analysis to incorporate multivariate information. Moreover, we present a range of sensitivity analyses that shed light on the robustness of our findings.

5.1 Multivariate Results

So far, we have followed much of the finance literature on return predictability and focused on univariate prediction models. We next extend the analysis to a multivariate setting. With N predictor variables available, there are N different predictive densities. Using i to index the predictors as we have done above, we denote these predictive densities by $p(r_{t+1}|\mathcal{D}^t, M_i)$, $i = 1, \dots, N$, where M_i refers to model i . Instead of conditioning only on a single predictor, investors may want to take advantage of the information contained in all of the N predictors. We consider two different ways to combine the information contained in N different predictors.

The first approach relies on forecast combination methods. Specifically, we construct a combined predictive density as an equal-weighted average of the N predictive densities

$$p(r_{t+1}|\mathcal{D}^t) = \sum_{i=1}^N w_{i,t} \times p(r_{t+1}|\mathcal{D}^t, M_i), \quad (34)$$

where $w_{i,t} = 1/N$ for all i and $t = \underline{t}-1, \dots, \bar{t}-1$. We compute the equal weighted predictive density in (34) across all predictor variables applying this approach separately to the unconstrained, EP constrained, and SR constrained models. For point forecasts this approach was previously adopted by Rapach et al. (2010).

Our second approach relies on diffusion indexes. As shown by Stock and Watson (2006), diffusion indexes provide a convenient framework for extracting the key common drivers from a large number of potential predictors. Ludvigson and Ng (2007) and Neely et al. (2012) show that diffusion indexes can be used to improve equity premium forecasting.

The diffusion index approach assumes a common factor structure for the N potential predictors,

$$x_{i\tau} = \lambda_i' f_\tau + e_{i,\tau}, \quad \tau = 1, \dots, t-1 \quad (35)$$

where i indexes the predictor, λ is a $(q \times 1)$ vector of factor loadings ($q \ll N$), f_τ is a $(q \times 1)$ vector of latent factors containing the common components extracted from the N predictors, and $e_{i,\tau}$ is a zero-mean disturbance term. Following Rapach et al. (2010), we restrict our analysis to considering a single factor; the results do not appear to improve if we include two or more factors in the model.

We estimate the common factors using principal components, and use them as predictors for stock returns in the following equation:

$$r_{\tau+1} = \mu_{DI} + \beta_{DI} f_\tau + \varepsilon_{\tau+1}, \quad \tau = 1, \dots, t-1, \quad (36)$$

where β_{DI} is a $(q \times 1)$ vector of slope coefficients and $\varepsilon_{\tau+1} \sim N(0, \sigma_{\varepsilon, DI}^2)$. As for the univariate models in Section 2, we specify an independent normal-gamma prior for the parameters in (36), and use a Gibbs sampler for estimation. Next, draws from the corresponding predictive density are obtained as

$$\begin{aligned} p(r_{t+1} | \mathcal{D}^t) &= \int_{\mu_{DI}, \beta_{DI}, \sigma_{\varepsilon, DI}^{-2}} p(r_{t+1} | \mu_{DI}, \beta_{DI}, \sigma_{\varepsilon, DI}^{-2}, \mathcal{D}^t) \\ &\quad \times p(\mu_{DI}, \beta_{DI}, \sigma_{\varepsilon, DI}^{-2} | \mathcal{D}^t) d\mu_{DI} d\beta_{DI} d\sigma_{\varepsilon, DI}^{-2} \end{aligned} \quad (37)$$

where $p(\mu_{DI}, \beta_{DI}, \sigma_{\varepsilon, DI}^{-2} | \mathcal{D}^t)$ is the joint posterior density of all parameters in (36). We estimate the diffusion index model in (36) and derive the predictive density in (37) for the unconstrained, EP constrained, and SR constrained models.

5.1.1 Empirical Findings

Table 7 presents empirical results for the equal-weighted combination as well as for the diffusion index. First consider the statistical measures of forecast performance. In all cases these improve when compared to the average forecast performance computed across the individual models. At all three horizons, the equal weighted combination yields the largest improvement in R^2 -performance for the unconstrained models. For example, at the monthly horizon the equal-weighted combination of unconstrained forecasts generates a R^2 of 0.62% versus -0.53% as the average value of the individual models. We also see improvements for the constrained

models, but these tend to be smaller. Which equal-weighted combination is best depends on the frequency: At the monthly horizon, combining unconstrained forecasts seem to work best; at the quarterly horizon, combining the SR constrained forecasts produces the best performance, while at the annual horizon the three approaches perform comparably.

The performance improves even more in the case of the diffusion index which works particularly well under the economic constraints. In fact, the EP-constrained forecasts deliver higher R^2 -values than the equal-weighted unconstrained forecasts at all three horizons, and across both subsamples. The SR-constrained forecasts based on the diffusion index also perform very well.

Turning to the economic performance measures, again, these generally lead to higher CER values when compared against the average values produced by the individual univariate models. While the beneficial effect of using equal-weighted forecasts remains largest for the unconstrained forecasts, the resulting CER values are always smaller than those produced by the corresponding equal-weighted constrained forecasts with differences ranging from 0.13% to 0.31% for the EP-constrained forecasts, and from 0.39% to 0.72% for the SR-constrained forecasts. Moreover, the best results from using the diffusion index is obtained for the constrained forecasts with differences ranging from 0.51% to 0.71% for the EP-constrained forecasts and from 0.99% to 1.27% for the SR-constrained forecasts. In fact, the diffusion index approach works better than the equal-weighted combination for the EP-constrained and the SR-constrained cases at all horizons.

5.2 Performance in Recessions and Expansions

Table 8 shows results separately for recession and expansion periods as defined by the NBER indicator and applied to the monthly and quarterly data. This type of analysis has been proposed by authors such as Rapach et al. (2010) and Henkel et al. (2011). Consistent with the findings in these studies, the unconstrained return prediction models do better during recessions than during expansions. Interestingly, the converse holds for the constrained monthly forecasts which far outperform the unconstrained forecasts during expansions, but perform worse during recessions.

While this may seem surprising, this finding can be explained by the fact that the state variable being sorted on (recessions) is correlated with returns. Returns tend to be negative during recessions and positive during expansions and so models that impose non-negative predictions will almost by construction do relatively better during expansions and worse during recessions.

Despite this effect, at the quarterly horizon the EP constrained forecasts do almost as well as the unconstrained forecasts during recessions (and much better during expansions), while the SR constrained forecasts do significantly better than the unconstrained forecasts both during expansions and recessions.

5.3 Sensitivity to Risk Aversion

Our main analysis of the economic value of equity premium forecasts in Section 4 assumed a coefficient of relative risk aversion of $A = 5$. To explore the sensitivity of our results to this value, we also considered lower ($A = 2$) and higher ($A = 10$) values of this parameter. Results are shown in Table 9.

First consider the case with $A = 2$, i.e., lower risk aversion compared with the baseline case that has $A = 5$. At the monthly horizon, the average CER performance of the unconstrained performance models is reduced from -0.12% to -0.27%, while conversely the average CER values of the EP constrained models increases from 0.28% to 0.68% and the SR constrained models' average performance is essentially unchanged at 0.79%. At the quarterly horizon, the mean CER value of the unconstrained forecasts remains unchanged, while the mean CER values under the EP and SR constraints increase from around 0.3% to 0.8% and 1.0%, respectively. Similarly, the constrained models' mean CER values increase to 0.83% and 1.03%, respectively (previously 0.33% and 0.67%) at the annual horizon, whereas the unconstrained model's mean CER value, at -0.07%, remains negative (previously -0.14%). Lowering the coefficient of risk aversion from $A = 5$ to $A = 2$ thus has the effect of boosting the economic performance of the constrained models whose mean CER value exceeds that of the unconstrained model by more than 0.90% (EP constraints) and 1.06% (SR constraints).

Conversely, increasing the risk aversion from $A = 5$ to $A = 10$ reduces the spread in the performance of the different models, as an investor with such a high level of risk aversion refrains from taking large positions in equity even in the presence of strong evidence of return predictability. At this higher level of risk aversion, the constrained models continue to outperform the unconstrained ones, although the difference in average CER values is reduced to between 0.2% and 0.35% for the EP constrained models and to between 0.3% and 0.5% for the SR constrained models.

5.4 Sensitivity to Priors

We also test the robustness of our results to alternative prior assumptions, and perform a sensitivity analysis where we experiment with different values for some of the key prior hyperparameters. Given the similarities in the results obtained under the models based on the EP and SR constraints, as well as the more computationally demanding estimation algorithm required by models imposing the SR constraint, we focus our attention on the effectiveness of the equity premium constraint as the priors change, and explore the effect of altering the hyperparameter ψ and \underline{v}_0 in (10) and (11). As discussed in section 2.2, the hyperparameter ψ plays the role of a scaling factor controlling the tightness of the priors for μ and β , and our benchmark models set $\psi = 2.5$. As sensitivities, we experiment with $\psi = 1.25$ and $\psi = 5$, which imply more

concentrated prior distributions (in the case of $\psi = 1.25$) or more dispersed prior distributions (in the case of $\psi = 5$) for μ and β . Similarly, the prior parameter \underline{v}_0 controls the tightness of the prior for σ_ε^{-2} , and our benchmark models set $\underline{v}_0 = 0.1$, which corresponds to a hypothetical prior sample size as large as 10% of the actual sample.²¹ As a sensitivity, we experiment with $\underline{v}_0 = 0.5$ and $\underline{v}_0 = 0.05$, which imply, respectively, an hypothetical prior sample as large as 50% of the underlying estimation sample (in the case of $\underline{v}_0 = 0.5$), or as large as 5% of the underlying estimation sample (in the case of $\underline{v}_0 = 0.05$).

Table 10 summarize the (relative) statistical and economic performances of both unconstrained and equity premium constrained models under two alternative prior choices, over the period 1947-2010. A comparison with tables 2, 3, 4, and 5 reveal that the key results derived under the benchmark priors remain largely unaffected when the priors change, with perhaps the only noticeable difference being related to the Clark and West (2007) p -values in table 10 for the more dispersed prior choice, which is a direct consequence of the slightly wider predictive densities resulting from this prior.

5.5 Other predictor variables

So far we focused our empirical analysis exclusively on the predictor variables considered by Welch and Goyal (2008). One additional predictor variable that has recently garnered considerable interest is the single factor extracted from the cross-section of book-to-market ratios by Kelly and Pruitt (2012). For comparison, we compute the out-of-sample R^2 obtained from this predictor for the same sample period as that used here, 1947-2012.²² For the unconstrained, EP and SR constrained models, we obtain R^2 -values of 0.16, -0.16, and 0.01, respectively. This suggests, first, that the economic constraints do not improve the forecasts from a linear regression model based on this particular predictor variable. Second, however, it is interesting to note that many of our models produce higher out-of-sample R^2 -values than those obtained for this new predictor variable.

6 Conclusion

Economic arguments can be used to constrain the predictive moments of the equity premium distribution. We develop a new methodology for imposing constraints that rule out negative equity premia and bound the conditional Sharpe ratio from above and below. Our approach efficiently exploits the information in such constraints in a way that incorporates the entire

²¹Under conjugate priors, the information contained in the priors can be viewed as “fictitious sample information” in that it is combined with the sample in exactly the same way that additional sample information would be combined. The only difference is that the prior information is observed in the mind of the researcher, not in the real world (see Koop (2003)).

²²We are grateful to Seth Pruitt for making the data available to us.

sequence of data points, while accounting for parameter uncertainty.

When evaluated empirically, we find that a key effect of the economic constraints is to reduce the impact of “large” values of the predictor variables on the expected equity premium. Conventional linear models tend to generate noisy forecasts following large variations in the predictor variables. In contrast, economic constraints effectively shrink the predicted equity premium and reduce the effect of outliers since they are more likely to bind for “large” values of the predictor variables. This gives rise to dynamic learning effects which can shift the entire predictive distribution for the equity premium even in situations where no such change is observed under a conventional unconstrained model.

Imposing economic constraints on the equity premium forecast improves the predictive accuracy for nearly all of the prediction models we consider. Moreover, the benefits from the economic constraints seem to improve, the longer the forecast horizon. In turn, when used to select portfolio weights, the constrained forecasts are found to yield higher certainty equivalent returns than their unconstrained counterparts.

References

- 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.
- 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.
- Avramov, D. and R. Wermers (2006). Investing in mutual funds when returns are predictable. *Journal of Financial Economics* 81(2), 339 – 377.
- Aït-Sahalia, Y. and M. W. Brandt (2001). Variable selection for portfolio choice. *The Journal of Finance* 56(4), pp. 1297–1351.
- Banbura, M., D. Giannone, and L. Reichlin (2010). Large bayesian vector auto regressions. *Journal of Applied Econometrics* 25(1), 71–92.
- Barberis, N. (2000). Investing for the long run when returns are predictable. *The Journal of Finance* 55(1), 225 – 264.
- Bossaerts, P. and P. Hillion (1999). Implementing statistical criteria to select return forecasting models: what do we learn? *Review of Financial Studies* 12(2), 405–428.
- Brandt, M. W. (2010). Portfolio choice problems. *Handbook of Financial Econometrics* 1, 269 – 336.

- Brennan, M. J., E. S. Schwartz, and R. Lagnado (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control* 21(8&AŠ9), 1377 – 1403.
- Campbell, J. and R. Shiller (1988). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1(3), 195–228.
- Campbell, J. Y. (1987). Stock returns and the term structure. *Journal of Financial Economics* 18(2), 373 – 399.
- 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.
- Campbell, J. Y. and L. M. Viceira (1999). Consumption and portfolio decisions when expected returns are time varying. *The Quarterly Journal of Economics* 114(2), 433–495.
- Carter, C. K. and R. Kohn (1994). On gibbs sampling for state space models. *Biometrika* 81(3), pp. 541–553.
- 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. (2001). *Asset pricing*. Princeton University Press.
- Cochrane, J. H. and J. Saa-Requejo (2000). Beyond arbitrage: Good-deal asset price bounds in incomplete markets. *Journal of Political Economy* 108(1), 79–119.
- Diebold, F. X. (2012). Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of diebold-mariano tests. Working paper, University of Pennsylvania.
- 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.
- Fama, E. F. and K. R. French (1988). Dividend yields and expected stock returns. *Journal of Financial Economics* 22(1), 3 – 25.
- Fama, E. F. and K. R. French (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25(1), 23 – 49.
- Ferson, W. E. and C. R. Harvey (1991). The variation of economic risk premiums. *Journal of Political Economy* 99(2), pp. 385–415.
- Ferson, W. E. and R. W. Schadt (1996). Measuring fund strategy and performance in changing economic conditions. *The Journal of Finance* 51(2), pp. 425–461.

- Goyal, A. and I. Welch (2003). Predicting the equity premium with dividend ratios. *Management Science* 49(5), pp. 639–654.
- Henkel, S. J., J. S. Martin, and F. Nardari (2011). Time-varying short-horizon predictability. *Journal of Financial Economics* 99(3), 560 – 580.
- Jobson, J. D. and B. M. Korkie (1981). Performance hypothesis testing with the sharpe and treynor measures. *The Journal of Finance* 36(4), pp. 889–908.
- Kadiyala, K. R. and S. Karlsson (1997). Numerical methods for estimation and inference in bayesian var-models. *Journal of Applied Econometrics* 12(2), 99–132.
- 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.
- Keim, D. B. and R. F. Stambaugh (1986). Predicting returns in the stock and bond markets. *Journal of Financial Economics* 17(2), 357 – 390.
- Kelly, B. T. and S. Pruitt (2012). Market expectations in the cross section of present values. *Journal of Finance* (forthcoming).
- 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.
- Koop, G. (2003). *Bayesian Econometrics*. John Wiley & Sons, Ltd.
- Lamont, O. (1998). Earnings and expected returns. *The Journal of Finance* 53(5), pp. 1563–1587.
- Lettau, M. and S. Ludvigson (2001). Consumption, aggregate wealth, and expected stock returns. *The Journal of Finance* 56(3), 815–849.
- Lettau, M. and S. C. Ludvigson (2010). Measuring and modeling variation in the risk-return trade-off. Volume 1 of *Handbook of Financial Econometrics*, pp. 617 – 690. Elsevier.
- Lettau, M. and S. Van Nieuwerburgh (2008). Reconciling the return predictability evidence. *Review of Financial Studies* 21(4), 1607–1652.
- Lettau, M. and J. A. Wachter (2007). Why is long-horizon equity less risky? a duration-based explanation of the value premium. *The Journal of Finance* 62(1), 55–92.
- Lettau, M. and J. A. Wachter (2011). The term structures of equity and interest rates. *Journal of Financial Economics* 101(1), 90 – 113.

- Ludvigson, S. C. and S. Ng (2007). The empirical risk-return relation: A factor analysis approach. *Journal of Financial Economics* 83(1), 171 – 222.
- MacKinlay, A. (1995). Multifactor models do not explain deviations from the capm. *Journal of Financial Economics* 38(1), 3 – 28.
- MacKinlay, A. C. and L. Pastor (2000). Asset pricing models: Implications for expected returns and portfolio selection. *The Review of Financial Studies* 13(4), pp. 883–916.
- Neely, C., D. E. Rapach, J. Tu, and G. Zhou (2012). Forecasting the equity risk premium: The role of technical indicators. Federal Reserve Bank of St. Louis working paper.
- Ostdiek, B. (1998). The world ex ante risk premium: an empirical investigation. *Journal of International Money and Finance* 17(6), 967 – 999.
- Pastor, L. and R. F. Stambaugh (2009). Predictive systems: Living with imperfect predictors. *The Journal of Finance* 64(4), 1583–1628.
- Pastor, L. and R. F. Stambaugh (2012). Are stocks really less volatile in the long run? *The Journal of Finance* 67(2), 431–478.
- Paye, B. S. and A. Timmermann (2006). Instability of return prediction models. *Journal of Empirical Finance* 13(3), 274 – 315.
- Pesaran, M. H. and A. Timmermann (1995). Predictability of stock returns: Robustness and economic significance. *The Journal of Finance* 50(4), 1201–1228.
- Poirier, D. J. (1995). *Intermediate Statistics and Econometrics: A Comparative Approach*. MIT Press.
- Polk, C., S. Thompson, and T. Vuolteenaho (2006). Cross-sectional forecasts of the equity premium. *Journal of Financial Economics* 81(1), 101 – 141.
- Pontiff, J. and L. D. Schall (1998). Book-to-market ratios as predictors of market returns. *Journal of Financial Economics* 49(2), 141 – 160.
- 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.
- Rapach, D. E. and G. Zhou (2012). Forecasting stock returns. Working paper - In preparation for the Handbook of Economic Forecasting, volume 2.

- Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13(3), 341 – 360.
- Ross, S. A. (2005). *Neoclassical finance*. Princeton University Press.
- Shanken, J. and A. Tamayo (2012). Payout yield, risk, and mispricing: A bayesian analysis. *Journal of Financial Economics* 105(1), 131 – 152.
- Stock, J. H. and M. W. Watson (2006). Forecasting with many predictors. Volume 1 of *Handbook of Economic Forecasting*, pp. 515 – 554. Elsevier.
- Wachter, J. A. and M. Warusawitharana (2009). Predictable returns and asset allocation: Should a skeptical investor time the market? *Journal of Econometrics* 148(2), 162 – 178.
- 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.
- Xia, Y. (2001). Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation. *The Journal of Finance* 56(1), 205–246.
- Zhou, G. (2010). How much stock return predictability can we expect from an asset pricing model? *Economics Letters* 108(2), 184 – 186.

A Appendix

This appendix explains how we obtain parameter estimates for the models described in Section 2, and how we use these to generate predictive densities for excess returns. We begin by discussing the unconstrained model in (1), and next turn to the models that incorporate the constraints in (2) and (8).

A.1 No constraints

In the unconstrained case, the goal is to obtain draws from the joint posterior distribution $p(\mu, \beta, \sigma_\varepsilon^{-2} | \mathcal{D}^t)$ with \mathcal{D}^t denoting all information available up to time t . Combining the priors in (9)-(11) with the likelihood function yields the following conditional posteriors:

$$\left[\begin{array}{c} \mu \\ \beta \end{array} \right] \Big| \sigma_\varepsilon^{-2}, D^t \sim N(\bar{b}, \bar{V}), \quad (\text{A-1})$$

and

$$\sigma_\varepsilon^{-2} | \mu, \beta, D^t \sim G(\bar{s}^{-2}, \bar{v}), \quad (\text{A-2})$$

where

$$\begin{aligned}\bar{V} &= \left[\underline{V}^{-1} + \sigma_\varepsilon^{-2} \sum_{\tau=1}^{t-1} z_\tau z_\tau' \right]^{-1}, \\ \bar{b} &= \bar{V} \left[\underline{V}^{-1} \underline{b} + \sigma_\varepsilon^{-2} \sum_{\tau=1}^{t-1} z_\tau r_{\tau+1} \right], \\ \bar{v} &= \underline{v}_0 + (t-1),\end{aligned}\tag{A-3}$$

and

$$\bar{s}^2 = \frac{\sum_{\tau=1}^{t-1} (r_{\tau+1} - \mu - \beta x_\tau)^2 + (s_{r,t}^2 \times \underline{v}_0 (t-1))}{\bar{v}}.\tag{A-4}$$

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

$$p(r_{t+1} | \mathcal{D}^t) = \int_{\mu, \beta, \sigma_\varepsilon^{-2}} p(r_{t+1} | \mu, \beta, \sigma_\varepsilon^{-2}, \mathcal{D}^t) p(\mu, \beta, \sigma_\varepsilon^{-2} | \mathcal{D}^t) d\mu d\beta d\sigma_\varepsilon^{-2}.\tag{A-5}$$

A.2 Equity premium constraints

The approach used for the unconstrained model also works when EP constraints are imposed subject to simply introducing an accept-reject step in (A-1):

$$\begin{aligned}\begin{bmatrix} \mu \\ \beta \end{bmatrix} \Big| \sigma_\varepsilon^{-2}, D^t &\sim N(\bar{b}, \bar{V}) \times \mu, \beta \in A_t, \\ \sigma_\varepsilon^{-2} \Big| \mu, \beta, D^t &\sim G(\bar{s}^{-2}, \bar{v}),\end{aligned}\tag{A-6}$$

with A_t defined in Equation (16), and, once again, using Equation (A-5).

A.3 Conditional Sharpe Ratio Constraints

To obtain draws from the joint posterior distribution $p(\mu, \beta, h^t, \sigma_\xi^{-2} | \mathcal{D}^t)$ under the SR constraints, we use the Gibbs sampler to draw recursively from the following three conditional distributions:

1. $p(h^t | \mu, \beta, \sigma_\xi^{-2}, \mathcal{D}^t)$.
2. $p(\mu, \beta | h^t, \sigma_\xi^{-2}, \mathcal{D}^t)$.
3. $p(\sigma_\xi^{-2} | \mu, \beta, h^t, \mathcal{D}^t)$.

We simulate from each of these blocks as follows. Starting with $p(h^t | \mu, \beta, \sigma_\xi^{-2}, \mathcal{D}^t)$, we employ the algorithm of Kim et al. (1998). Define $r_{\tau+1}^* = r_{\tau+1} - \mu - \beta x_\tau$ and note that $r_{\tau+1}^*$ is observable conditional on μ, β . Next, rewrite (5) as

$$r_{\tau+1}^* = \exp(h_{\tau+1}) u_{\tau+1}. \quad (\text{A-7})$$

Squaring and taking logs on both sides of (A-7) yields a new state space system that replaces (5)-(6) with

$$r_{\tau+1}^{**} = 2h_{\tau+1} + u_{\tau+1}^{**}, \quad (\text{A-8})$$

$$h_{\tau+1} = h_\tau + \xi_{\tau+1}, \quad (\text{A-9})$$

where $r_{\tau+1}^{**} = \ln(r_{\tau+1}^2)$, and $u_{\tau+1}^{**} = \ln(u_{\tau+1}^2)$, with u_τ^{**} independent of ξ_s for all τ and s . Since $u_{t+1}^{**} \sim \ln(\chi_1^2)$, we cannot resort to standard Kalman recursions and simulation algorithms such as those in Carter and Kohn (1994) or Durbin and Koopman (2002). To obviate this problem, Kim et al. (1998) employ a data augmentation approach and introduce a new state variable $s_{\tau+1}$, $\tau = 1, \dots, t-1$, turning their focus on drawing from $p(h^t | \mu, \beta, \sigma_\xi^{-2}, s^t, \mathcal{D}^t)$ instead of $p(h^t | \mu, \beta, \sigma_\xi^{-2}, \mathcal{D}^t)$.²³ The introduction of the state variable $s_{\tau+1}$ allows us to rewrite the linear non-Gaussian state space representation in (A-8)-(A-9) as a linear Gaussian state space model, making use of the following approximation,

$$\varepsilon_{\tau+1}^{**} \approx \sum_{j=1}^7 q_j f_N(m_j - 1.2704, v_j^2), \quad (\text{A-10})$$

where m_j , v_j^2 , and q_j , $j = 1, 2, \dots, 7$, are constants specified in Kim et al. (1998) and thus need not be estimated. In turn, (A-10) implies

$$\varepsilon_{\tau+1}^{**} | s_{\tau+1} = j \sim N(m_j - 1.2704, v_j^2), \quad (\text{A-11})$$

where each state has probability

$$\Pr(s_{\tau+1} = j) = q_j. \quad (\text{A-12})$$

Conditional on s^t , we can rewrite the nonlinear state space system as follows:

$$\begin{aligned} r_{\tau+1}^{**} &= 2h_{\tau+1} + e_{\tau+1}, \\ h_{\tau+1} &= h_\tau + \xi_{\tau+1}, \end{aligned} \quad (\text{A-13})$$

where $e_{\tau+1} \sim N(m_j - 1.2704, v_j^2)$ with probability q_j . 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 .

²³Here $s^t = \{s_1, s_2, \dots, s_t\}$ denotes the history of the state variable s up to time t .

Finally, conditional on the whole sequence h^t , draws for the sequence of states s^t can easily be obtained, noting that

$$\Pr(s_{\tau+1} = j | r_{\tau+1}^{**}, h_{\tau+1}) = \frac{f_N(r_{\tau+1}^{**} | 2h_{\tau+1} - m_j + 1.2704, v_j^2)}{\sum_{l=1}^7 f_N(r_{\tau+1}^{**} | 2h_{\tau+1} - m_l + 1.2704, v_l^2)}. \quad (\text{A-14})$$

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

$$\begin{bmatrix} \mu \\ \beta \end{bmatrix} \Big| h^t, \sigma_\xi^{-2}, \mathcal{D}^t \sim N(\bar{b}, \bar{V}) \times \mu, \beta \in \tilde{A}_t, \quad (\text{A-15})$$

where \tilde{A}_t is defined in (18), with

$$\begin{aligned} \bar{V} &= \left\{ \underline{V}^{-1} + \sum_{\tau=1}^{t-1} \frac{1}{\exp(h_{\tau+1})^2} z_\tau z_\tau' \right\}^{-1}, \\ \bar{b} &= \bar{V} \left\{ \underline{V}^{-1} \underline{b} + \sum_{\tau=1}^{t-1} \frac{1}{\exp(h_{\tau+1})^2} z_\tau r_{\tau+1} \right\}. \end{aligned}$$

In turn, the posterior distribution for $p(\sigma_\xi^{-2} | \mu, \beta, h^t, \mathcal{D}^t)$ is readily available as,

$$\sigma_\xi^{-2} | \mu, \beta, h^t, \mathcal{D}^t \sim G \left(\left[\frac{k_\xi + \sum_{\tau=2}^{t-1} (h_{\tau+1} - h_\tau)^2}{t-1} \right]^{-1}, t-1 \right). \quad (\text{A-16})$$

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

$$\begin{aligned} p(r_{t+1} | \mathcal{D}^t) &= \int_{\mu, \beta, h^{t+1}, \sigma_\xi^{-2}} p(r_{t+1} | h_{t+1}, \mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t) \\ &\quad \times p(h_{t+1} | \mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t) \\ &\quad \times p(\mu, \beta, h^t, \sigma_\xi^{-2} | \mathcal{D}^t) d\mu d\beta dh^{t+1} d\sigma_\xi^{-2}. \end{aligned} \quad (\text{A-17})$$

The first term in the integral above, $p(r_{t+1} | h_{t+1}, \mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t)$, represents the period $t+1$ predictive density of excess returns, treating model parameters as if they were known with certainty, and so is straightforward to calculate. The second term in the integral, $p(h_{t+1} | \mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t)$, reflects how period $t+1$ volatility may drift away from h_t over time. Finally, the last term in the integral, $p(\mu, \beta, h^t, \sigma_\xi^{-2} | \mathcal{D}^t)$, measures parameter uncertainty in the sample.

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

1. Simulating from $p(\mu, \beta, h^t, \sigma_\xi^{-2} | \mathcal{D}^t)$: draws from $p(\mu, \beta, h^t, \sigma_\xi^{-2} | \mathcal{D}^t)$ are obtained from the Gibbs sampling algorithm described above;

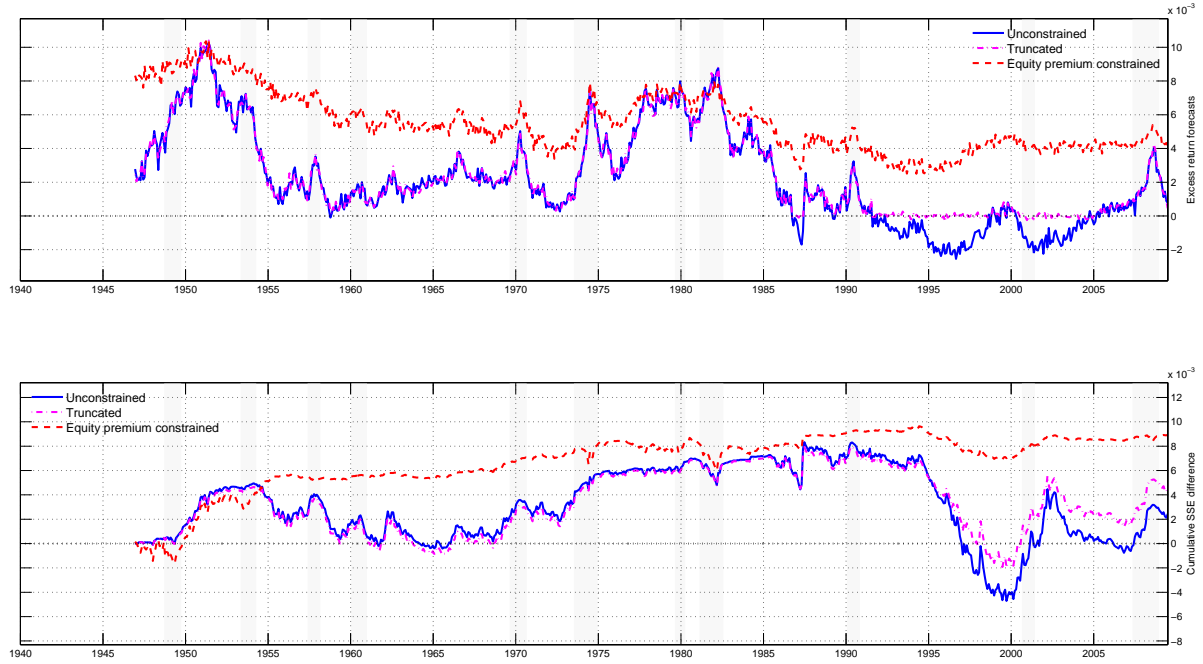
2. Simulating from $p(h_{t+1}|\mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t)$: having processed data up to time t , the next step is to simulate the future volatility, h_{t+1} . For a given h_t and σ_ξ^{-2} , note that μ and β and the history of volatilities up to t become redundant, i.e., $p(h_{t+1}|\mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t) = p(h_{t+1}|h_t, \sigma_\xi^{-2}, \mathcal{D}^t)$. Note also that (6) along with the distributional assumptions made with regards to ξ_{t+1} imply that

$$h_{t+1}|h_t, \sigma_\xi^{-2}, \mathcal{D}^t \sim N(h_t, \sigma_\xi^2). \quad (\text{A-18})$$

3. Simulating from $p(r_{t+1}|h_{t+1}, \mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t)$: For a given h_{t+1} , μ , and β , note that h^t and σ_ξ^{-2} become redundant, i.e., $p(r_{t+1}|h_{t+1}, \mu, \beta, h^t, \sigma_\xi^{-2}, \mathcal{D}^t) = p(r_{t+1}|h_{t+1}, \mu, \beta, \mathcal{D}^t)$. Then use the fact that

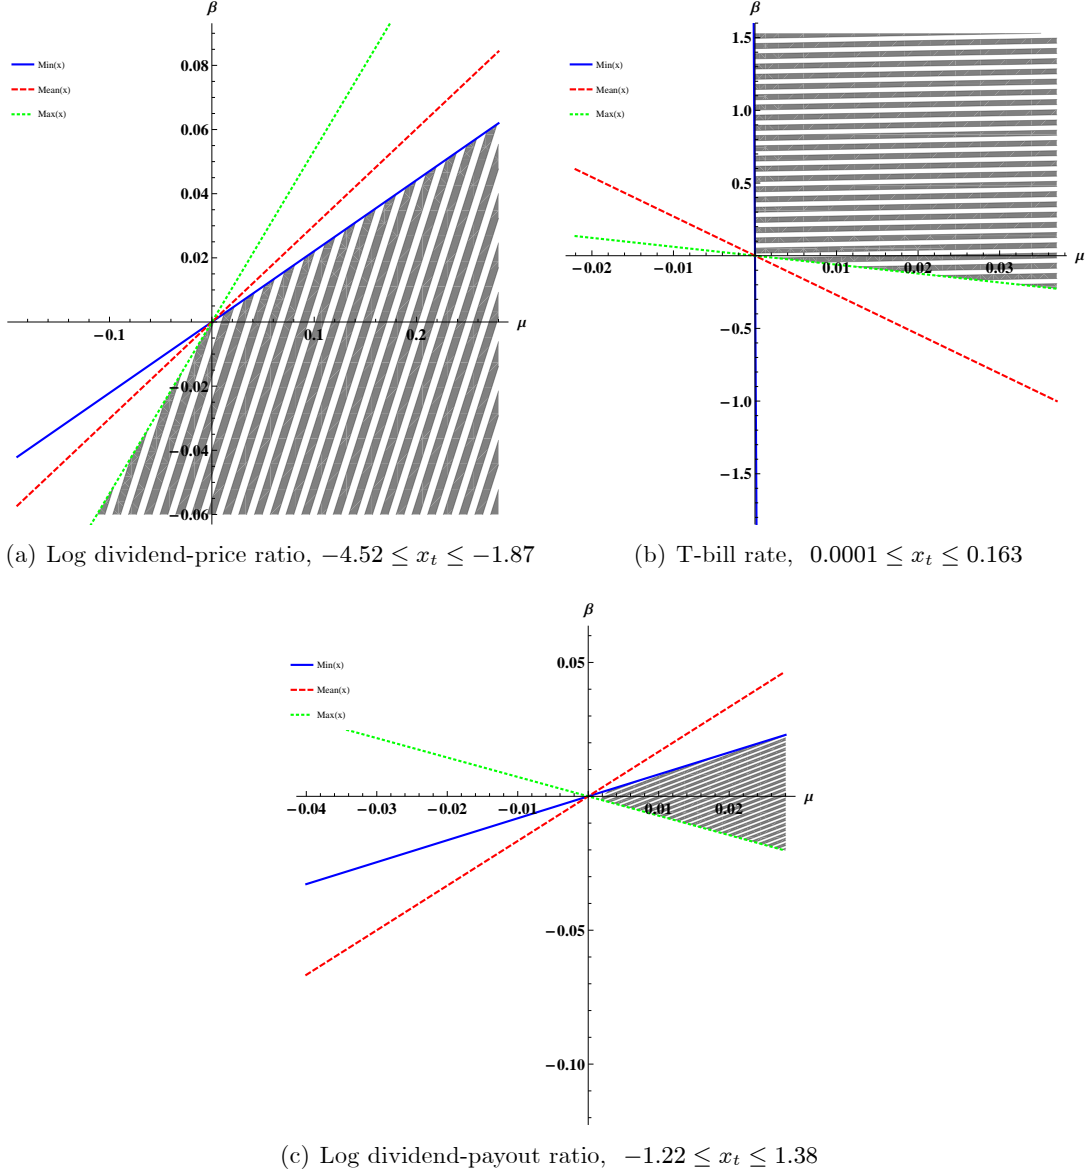
$$r_{t+1}|h_{t+1}, \mu, \beta, \mathcal{D}^t \sim N(\mu + \beta x_t, \exp(h_{t+1})). \quad (\text{A-19})$$

Figure 1: Out-of-sample equity premium forecasts and forecast performance under constrained and unconstrained models



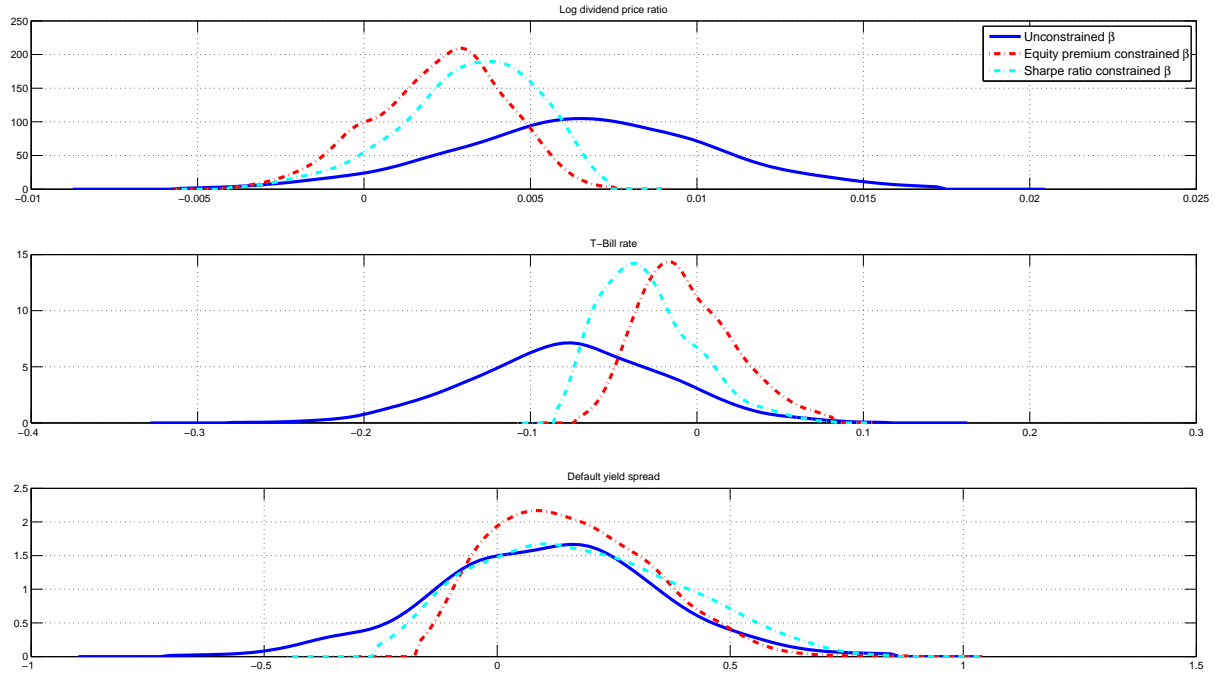
Each month we regress stock returns measured in excess of the T-bill rate (r_{t+1}) on an intercept and the log dividend-price ratio, $r_{t+1} = \mu + \beta \log(D_t/P_t) + \varepsilon_{t+1}$. Recursive estimates from this model are then used to generate one-step-ahead out-of-sample forecasts of excess returns, and the process is repeated up to the end of the sample. Unconstrained forecasts, $\hat{r}_{t+1|t}$, are based on recursive estimates using uninformative priors. Truncated forecasts use $\text{Max}(0, \hat{r}_{t+1|t})$, while the equity premium constrained forecasts use our Bayesian methodology. The top window plots the sequence of recursive out-of-sample forecasts, using 20 years of data as a warm-up period. The bottom window plots the cumulative sum of squared forecast errors (SSE_t) of the prevailing mean forecasts (SSE_t^{PM}) relative to the forecasts based on the log dividend-price ratio, $SSE_t^{PM} - SSE_t^{dp}$.

Figure 2: Equity premium constraint



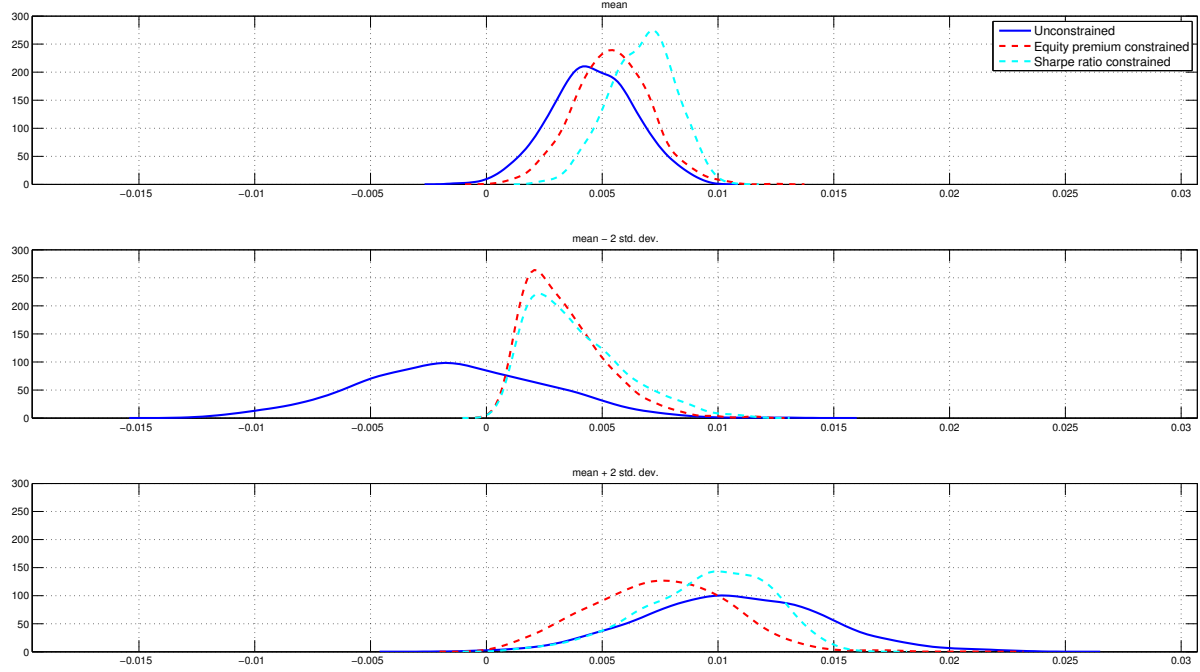
This figure shows the effect on the parameters μ and β from imposing the equity premium constraint $\mu + \beta x_t \geq 0$, as indicated by the shaded area.

Figure 3: Slope coefficient of predictors under constrained and unconstrained models



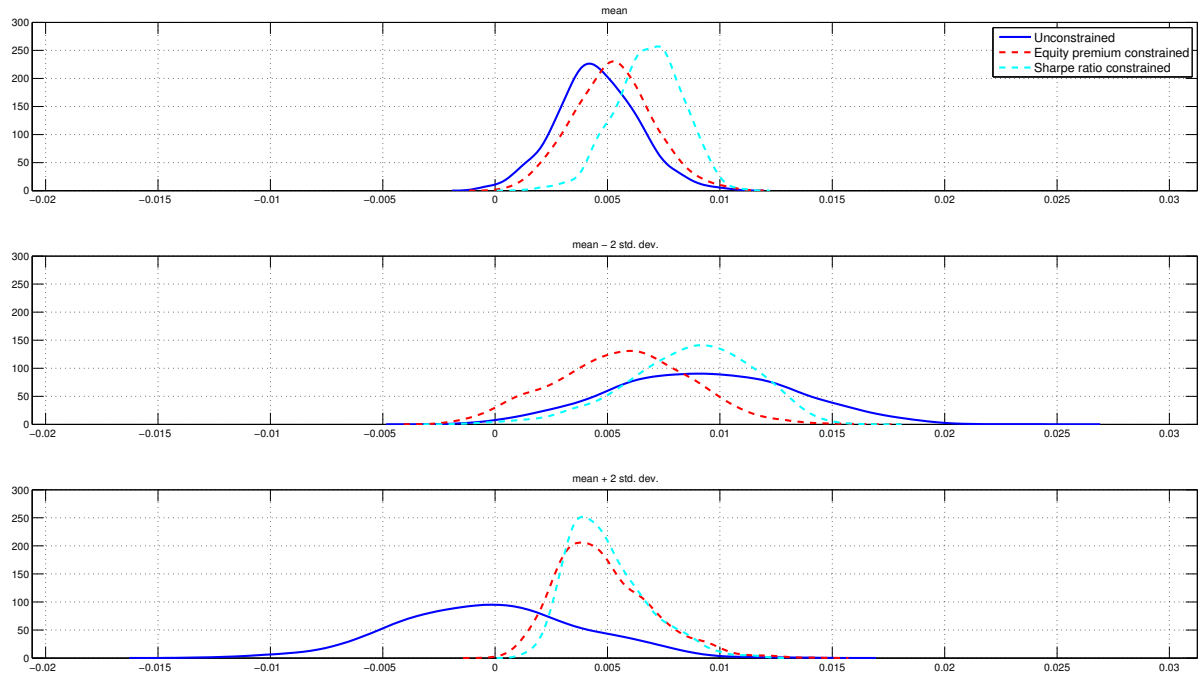
This figure shows the posterior density of the slope coefficient, β , from a regression of monthly excess returns (r_{t+1}) on an intercept and a lagged predictor variable, x_t : $r_{t+1} = \mu + \beta x_t + \varepsilon_{t+1}$. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T-1$ and information set \mathcal{D}^t , while the Sharpe ratio constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T-1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. The posterior density estimates shown here are based on the full sample at the end of 2010. Panels A, B, and C use the log dividend-price ratio, T-bill rate, and the default yield spread as predictors, respectively.

Figure 4: Posterior density of the equity premium under constrained and unconstrained models (Dividend-price ratio)



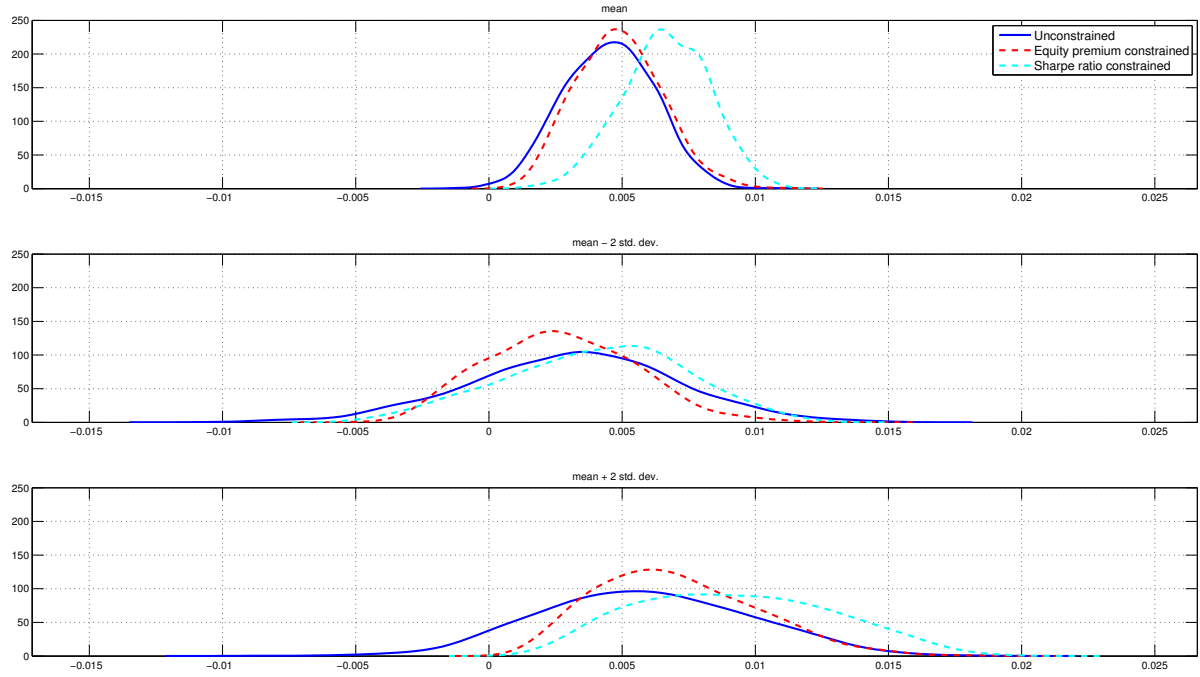
This figure shows the posterior density of the equity premium as a function of the log dividend-price ratio, $\mu + \beta x_T$, where x_T is set at the sample mean of the log dividend-price ratio \bar{x} (top window), $\bar{x} - 2st.dev(x)$ (middle window), and $\bar{x} + 2st.dev(x)$ (bottom window). The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. All posterior density estimates are based on the full data sample as of end-2010.

Figure 5: Posterior density of the equity premium under constrained and unconstrained models (T-bill rate)



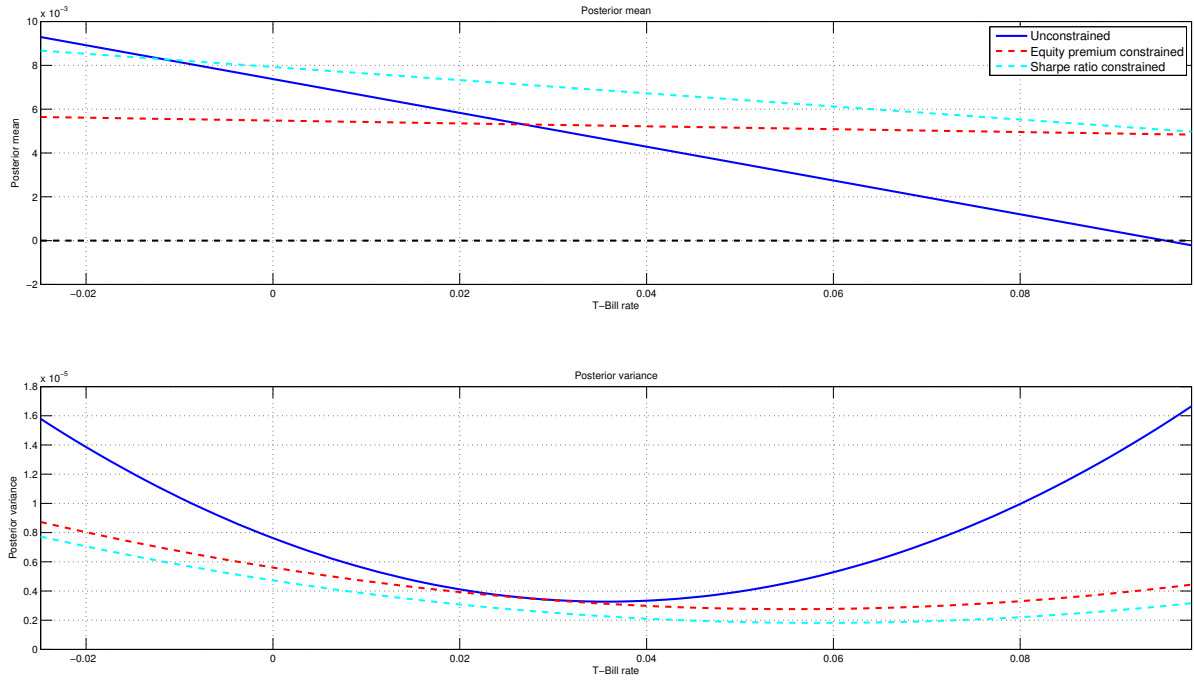
This figure shows the posterior density of the equity premium as a function of the T-bill rate, $\mu + \beta x_T$, where x_T is set at the sample mean of the T-bill rate \bar{x} (top window), $\bar{x} - 2st.dev(x)$ (middle window), and $\bar{x} + 2st.dev(x)$ (bottom window). The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. All posterior density estimates are based on the full data sample as of end-2010.

Figure 6: Posterior density of the equity premium under constrained and unconstrained models (default spread)



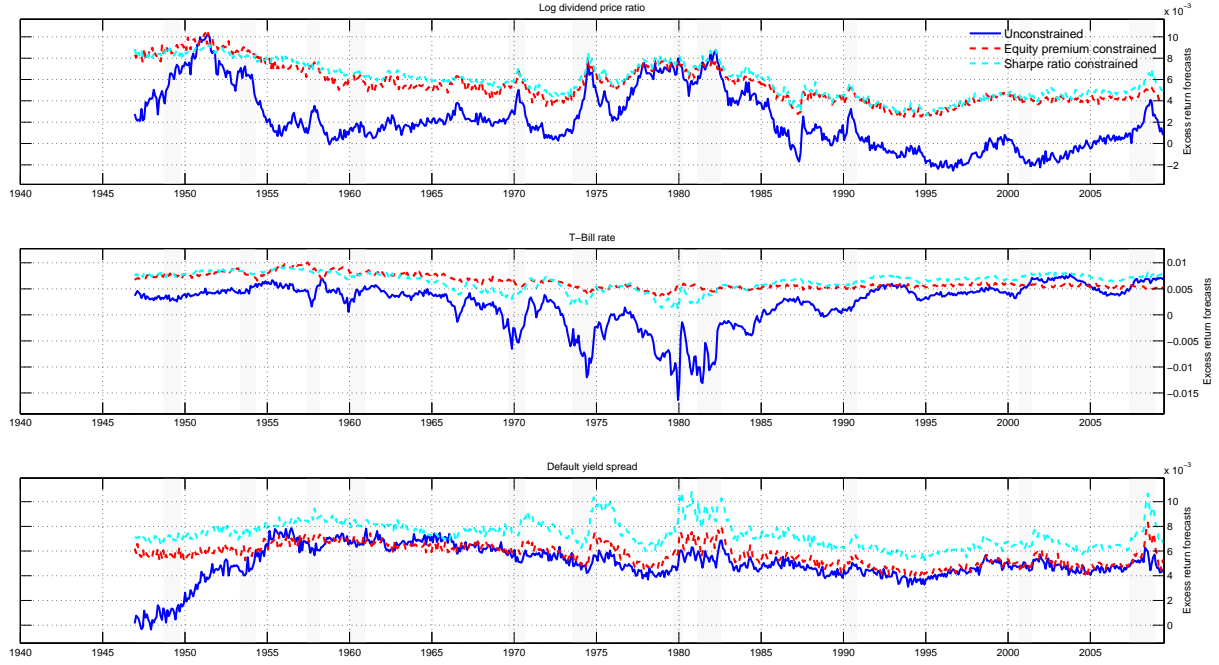
This figure shows the posterior density of the equity premium as a function of the default yield spread, $\mu + \beta x_T$, where x_T is set at the sample mean of the default yield spread \bar{x} (top window), $\bar{x} - 2st.dev(x)$ (middle window), and $\bar{x} + 2st.dev(x)$ (bottom window). The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. All posterior density estimates are based on the full data sample as of end-2010.

Figure 7: Posterior mean and variance of the equity premium as a function of the value of the predictor variable



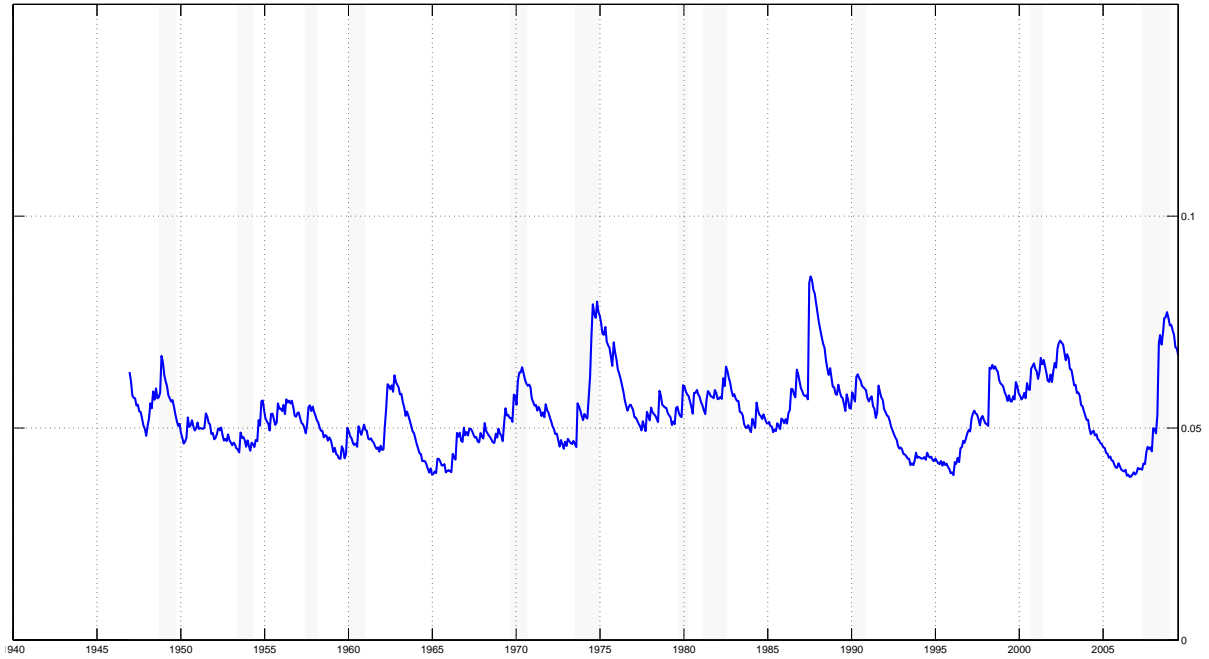
This figure plots the posterior mean (top window) and posterior variance (bottom window) of the equity premium as a function of the T-bill rate under unconstrained, equity premium constrained and Sharpe ratio constrained models, respectively. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. All posterior density estimates are based on the full data sample as of end-2010.

Figure 8: Out-of-sample equity premium forecasts under unconstrained and constrained models



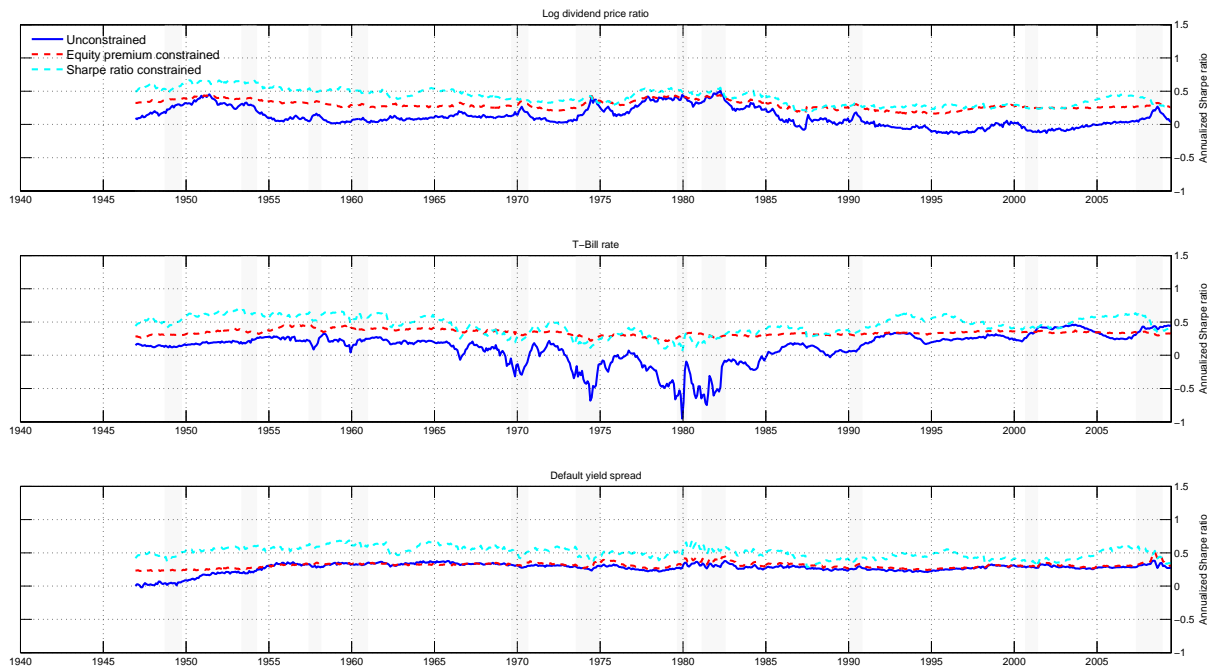
Each month we regress excess returns (r_{t+1}) on an intercept and a lagged predictor, x_t : $r_{t+1} = \mu + \beta x_t + \varepsilon_{t+1}$. Estimates from this model are then used to generate recursive one-step out-of-sample forecasts of excess returns and the process is repeated up to the end of the sample in 2010. Unconstrained forecasts, $\hat{r}_{t+1|t}$, are based on recursive least-squares estimates and uninformative priors. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. The three windows use different predictor variables, namely, the log dividend-price ratio (top window), the T-bill rate (middle), and the default yield spread (bottom).

Figure 9: Volatility forecasts



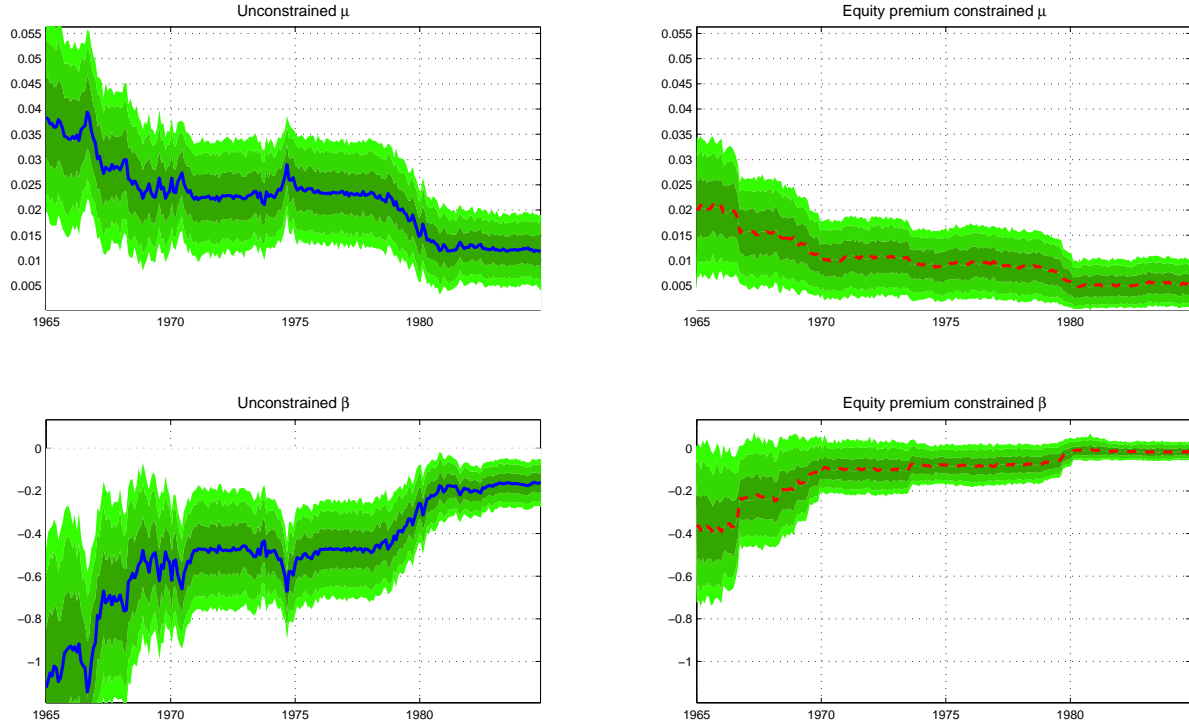
This figure shows the one-step-ahead recursive conditional volatility forecasts computed from the predictive return distribution based on the stochastic volatility model that uses the log dividend-price ratio as predictor, $r_{t+1} = \mu + \beta \log(D_t/P_t) + \exp(h_{t+1})u_{t+1}$, $h_{t+1} = h_t + \xi_{t+1}$.

Figure 10: Conditional Sharpe ratios



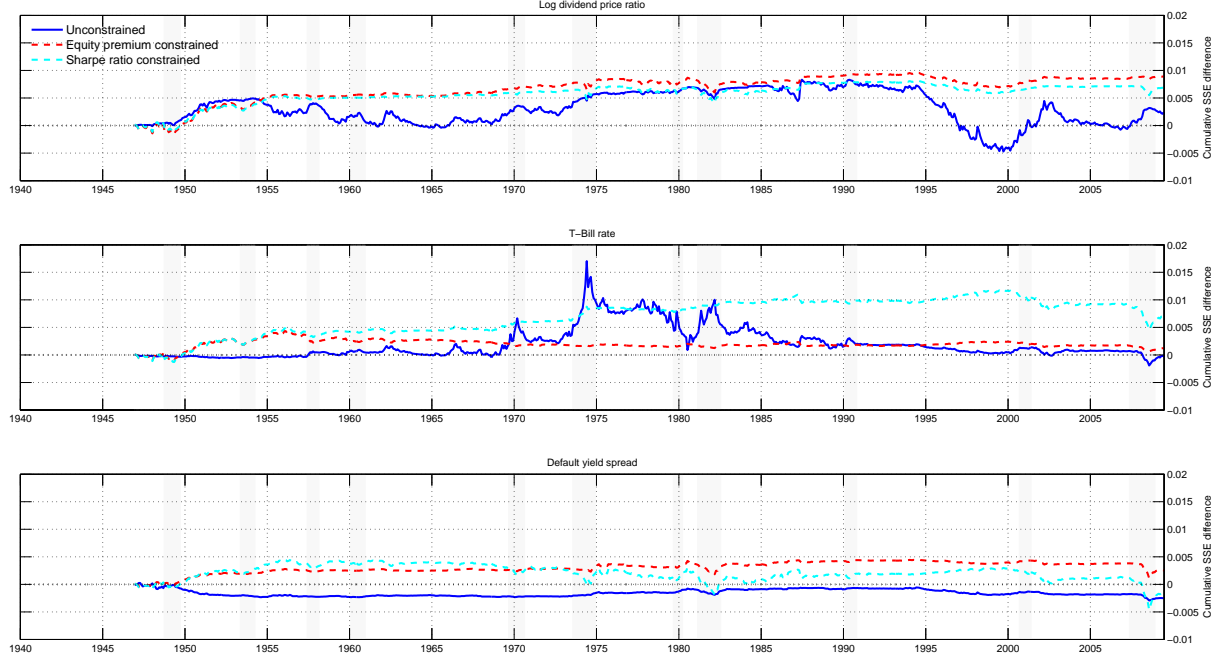
This figure shows the time series of conditional Sharpe ratios computed from the predictive density of excess returns based on the unconstrained, equity premium constrained, and Sharpe ratio constrained models. The three windows use different predictor variables, namely, the log dividend-price ratio (top window), the T-bill rate (middle), and the default yield spread (bottom).

Figure 11: Posterior probability intervals for the parameters of the return prediction model



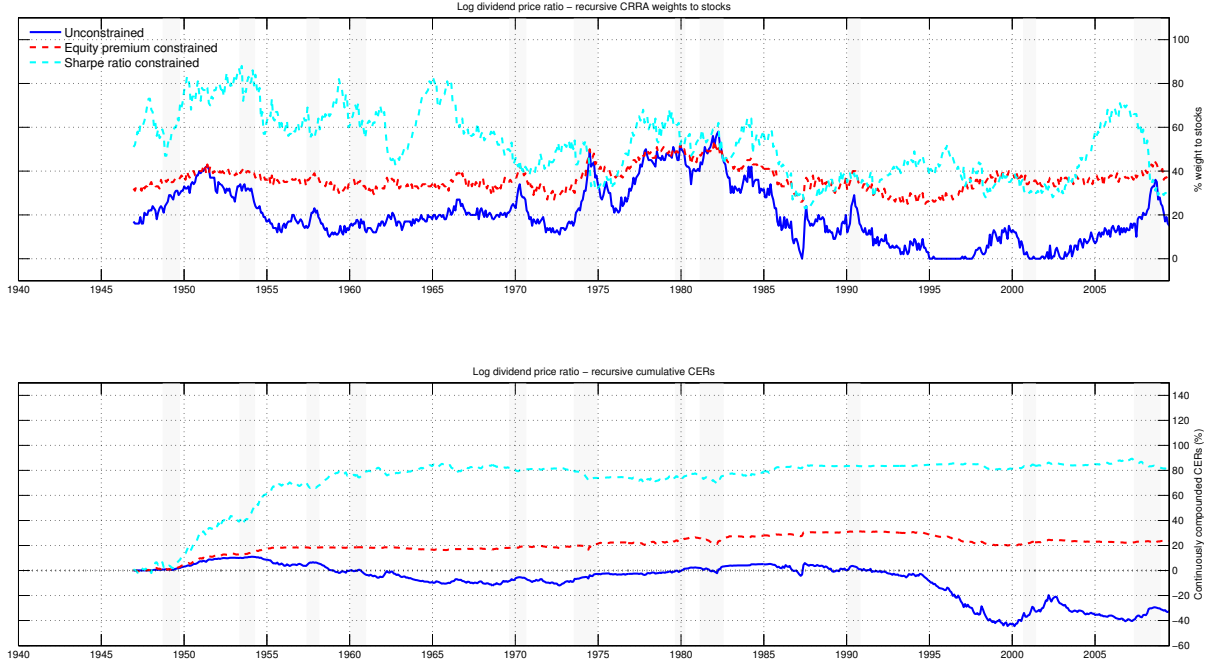
Each month we update the posterior density of the parameters μ, β of the return prediction model $r_{t+1} = \mu + \beta x_t + \varepsilon_{t+1}$, where x_t is the T-bill rate and r_{t+1} is the return on the S&P 500 index, measured in excess of the T-bill rate. We then compute posterior probability intervals for these parameter estimates. Windows to the left report the (2.5, 97.5) percentile posterior probability intervals for the parameters of the unconstrained model, while windows to the right show results for the equity-premium constrained model. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | D^t) d\mu d\beta > 0$ for all $t = 1, \dots, T - 1$ and information set \mathcal{D}^t . All density estimates are updated recursively through time.

Figure 12: Forecast performance: cumulative sum of squared forecast error differentials



This figure shows the sum of squared forecast errors of the prevailing mean model minus the sum of squared forecast errors of a forecast model with time-varying predictors. Each month we estimate the parameters of the forecast models recursively and generate one-step-ahead forecasts of excess returns which are in turn used to compute out-of-sample forecast errors. This procedure uses univariate forecast models based on the log dividend-price ratio (top window), the T-bill rate (middle window), or the default yield spread (bottom window) or a simple prevailing mean model which is our benchmark. We then plot the cumulative sum of squared forecast errors (SSE_t) of the prevailing mean forecasts (SSE_t^{PM}) relative to the univariate forecasts, $SSE_t^{PM} - SSE_t$. Values above zero indicate that a univariate forecast model generates better performance than the prevailing mean benchmark, while negative values suggest the opposite. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T-1$ and information set \mathcal{D}^t , while the Sharpe ratio constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T-1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate from a stochastic volatility model.

Figure 13: Portfolio allocation and economic value of forecasts



This figure plots the percentage allocation to stocks and the resulting cumulative certainty equivalent returns measured relative to the prevailing mean model. Each month we compute the optimal allocation to stocks and T-bills based on the predictive density of excess returns. The investor is assumed to have power utility with a coefficient of relative risk aversion of five and the weight on stocks is constrained to lie in the interval $[0, 0.99]$. The top window shows the recursively computed optimal weight on stocks, while the bottom window shows the cumulative certainty equivalent return measured relative to the prevailing mean model. The equity premium constrained model imposes the constraint that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T-1$ and information set \mathcal{D}^t , while the Sharpe ratio constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T-1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. The predictor used is the log dividend-price ratio.

Table 1: Summary statistics

Monthly				
Variables	Mean	Std. deviation	Skewness	Kurthosis
Excess returns	0.005	0.056	-0.405	10.603
Log dividend price ratio	-3.329	0.452	-0.403	3.044
Log dividend yield	-3.324	0.450	-0.435	3.030
Log earning price ratio	-2.720	0.426	-0.708	5.659
Log smooth earning price ratio	-2.912	0.376	-0.002	3.559
Log dividend-payout ratio	-0.609	0.325	1.616	9.452
Book-to-market ratio	0.589	0.267	0.671	4.456
T-Bill rate	0.037	0.031	1.025	4.246
Long-term yield	0.053	0.028	0.991	3.407
Long-term return	0.005	0.024	0.618	8.259
Term spread	0.016	0.013	-0.218	3.128
Default yield spread	0.011	0.007	2.382	11.049
Default return spread	0.000	0.013	-0.302	11.490
Stock variance	0.003	0.005	5.875	48.302
Net equity expansion	0.019	0.024	1.468	10.638
Inflation	0.002	0.005	-0.069	6.535
Log total net payout yield	-2.137	0.224	-1.268	6.213
Quarterly				
Variables	Mean	Std. deviation	Skewness	Kurthosis
Excess returns	0.014	0.108	0.201	11.087
Log dividend price ratio	-3.328	0.456	-0.372	3.077
Log dividend yield	-3.314	0.450	-0.471	3.037
Log earning price ratio	-2.719	0.432	-0.777	5.932
Log smooth earning price ratio	-2.906	0.378	0.028	3.654
Log dividend-payout ratio	-0.609	0.332	1.702	9.919
Book-to-market ratio	0.594	0.268	0.745	4.905
T-Bill rate	0.037	0.031	1.040	4.313
Long-term yield	0.053	0.028	1.008	3.484
Long-term return	0.014	0.045	1.067	7.369
Term spread	0.016	0.013	-0.260	3.285
Default yield spread	0.011	0.007	2.390	11.007
Default return spread	0.001	0.021	0.355	16.437
Stock variance	0.008	0.013	4.523	28.492
Net equity expansion	0.019	0.025	1.416	10.179
Inflation	0.007	0.013	-0.383	5.341
Annual				
Variables	Mean	Std. deviation	Skewness	Kurthosis
Excess returns	0.053	0.202	-0.904	4.104
Log dividend price ratio	-3.337	0.464	-0.415	2.873
Log dividend yield	-3.286	0.444	-0.732	3.115
Log earning price ratio	-2.722	0.420	-0.339	3.672
Log smooth earning price ratio	-2.895	0.377	-0.097	3.078
Log dividend-payout ratio	-0.615	0.319	1.068	5.664
Book-to-market ratio	0.585	0.263	0.506	3.285
T-Bill rate	0.037	0.031	1.028	4.388
Long-term yield	0.053	0.028	0.914	3.168
Long-term return	0.058	0.096	1.035	4.591
Term spread	0.016	0.014	-0.453	3.925
Default yield spread	0.012	0.008	2.278	9.532
Default return spread	0.004	0.043	-0.134	7.845
Stock variance	0.030	0.040	2.906	12.014
Net equity expansion	0.019	0.026	2.498	15.603
Inflation	0.030	0.038	-0.343	5.876
Percent equity issuing	0.194	0.110	1.733	8.368

This table reports summary statistics for excess returns, computed as returns on the S&P500 portfolio minus the T-bill rate, and for the predictor variables used in the study. The sample period is 1927-2010.

Table 2: Out-of-sample forecast performance (monthly)

	Panel A: full sample (1947-2010)			Panel B: 1947-1978			Panel C: 1979-2010			
Variables	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	
Log dividend price ratio	0.10%	*	0.64%	***	0.49%	**	1.01%	**	1.03%	**
Log dividend yield	-0.16%	**	0.76%	***	0.54%	**	0.97%	***	1.05%	**
Log earning price ratio	-1.38%	*	0.25%	*	0.30%	*	-1.17%	*	0.43%	*
Log smooth earning price ratio	-0.79%	*	0.48%	**	0.47%	**	-0.14%	*	0.81%	**
Log dividend-payout ratio	-1.57%		-0.12%		0.03%		-2.64%		0.23%	
Book-to-market ratio	-1.40%		0.02%		0.04%		-0.42%		0.08%	
T-Bill rate	0.01%	*	0.10%	**	0.54%	**	1.50%	**	1.40%	***
Long-term yield	-0.94%	*	0.31%	**	0.55%	***	-0.47%	*	0.78%	**
Long-term return	-0.74%		0.19%	**	0.53%	**	-1.90%	*	0.31%	**
Term spread	0.06%		0.05%	*	0.24%	*	0.41%		0.14%	*
Default yield spread	-0.19%		0.18%		-0.08%		-0.24%		0.53%	
Default return spread	-0.24%		0.15%		-0.30%		-0.97%		0.22%	**
Stock variance	0.24%		-0.17%		-0.29%		-0.07%	*	-0.42%	
Net equity expansion	-0.79%		-0.01%		-0.08%		0.08%		0.36%	*
Inflation	-0.13%		0.07%		-0.12%		0.05%		0.15%	
Log total net payout yield	-0.49%	*	0.18%		-0.01%		1.44%	***	0.45%	*
									0.25%	
									0.43%	
									0.08%	
									0.36%	
									0.03%	
									0.14%	
									0.20%	
									0.02%	
									0.00%	
									0.17%	
									0.31%	
									0.23%	
									0.12%	
									0.03%	
									0.02%	
									-0.11%	
									-0.04%	
									-0.11%	
									-0.02%	
									-0.02%	
									-0.31%	
									-0.07%	
									0.00%	
									-0.32%	
									-0.57%	
									-0.48%	
									-0.13%	
									-0.47%	
									-0.16%	
									-0.20%	
									0.02%	
									0.00%	
									-0.17%	

This table reports the out-of-sample R^2 of unconstrained and constrained univariate prediction models for the monthly excess return, r_{t+1} , measured relative to the prevailing mean model:

$$R_{OoS}^2 = 1 - \frac{\sum_{\tau=t-1}^{\bar{t}-1} (\tau_{t+1} - \hat{\tau}_{t+1|t})^2}{\sum_{\tau=t-1}^{\bar{t}-1} (\tau_{t+1} - \bar{\tau}_{t+1|t})^2},$$

where $\hat{\tau}_{t+1|t}$ is the mean of the predictive return distribution based on a regression of monthly excess returns on an intercept and a lagged predictor variable, x_t : $\tau_{t+1} = \mu + \beta x_t + \varepsilon_{t+1}$. $\bar{\tau}_{t+1|t}$ is the forecast from the prevailing mean model which assumes that $\beta = 0$. The equity premium constrained model imposes that $\hat{\tau}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio (SR) constraint imposes that $0 \leq \hat{\tau}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. Panel A shows results for the full sample (1947-2010), while panels B and C report results for the subsamples 1947-1978 and 1979-2010. Bold figures highlight instances where the constrained R_{OoS}^2 is higher than its unconstrained counterpart. We measure statistical significance relative to the prevailing mean model using the Clark-West (2007) test statistic. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 3: Out-of-sample forecast performance (quarterly)

Variables	Panel A: full sample (1947-2010)			Panel B: 1947-1978			Panel C: 1979-2010		
	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint
Log dividend price ratio	-0.16% **	2.11% ***	1.68% ***	3.22% ***	4.27% ***	3.07% ***	-3.10% ***	0.24% ***	0.47% ***
Log dividend yield	0.75% **	2.04% ***	1.88% ***	4.19% ***	4.04% ***	3.38% ***	-2.22% ***	0.32% ***	0.57% ***
Log earning price ratio	-6.33% *	0.70% **	0.92% **	4.65% **	0.78% *	1.60% **	-7.78% **	0.33% **	0.33% **
Log smooth earning price ratio	-3.58% **	1.57% **	1.60% **	-0.80% **	3.03% **	2.96% **	-5.98% **	0.30% **	0.41% **
Log dividend-payout ratio	-3.69% **	0.04% **	0.81% **	-5.16% **	-0.04% **	1.75% **	-2.42% **	0.11% **	0.00% **
Book-to-market ratio	-6.24% **	-0.15% **	0.96% **	-1.64% **	0.04% **	1.64% **	-10.23% **	-0.32% **	0.37% **
T-Bill rate	-0.53% ***	0.01% ***	1.87% ***	2.40% ***	0.34% ***	3.87% ***	-3.06% ***	-0.27% ***	0.14% ***
Long-term yield	-2.70% **	0.61% **	1.55% **	-1.82% **	1.51% **	3.41% **	-3.45% **	-0.16% **	-0.06% **
Long-term return	-0.64% **	0.44% **	0.66% **	-0.41% **	0.68% **	1.46% **	-0.84% **	0.24% **	-0.03% **
Term spread	0.11% **	0.04% **	1.21% **	1.06% **	0.39% **	2.17% **	-0.71% **	-0.25% **	0.38% **
Default yield spread	-0.82% **	0.50% **	0.42% **	-0.94% **	1.79% **	1.49% **	-0.72% **	-0.62% **	-0.51% **
Default return spread	-6.14% **	-1.22% **	0.77% **	-11.02% **	-1.33% **	0.83% **	-1.92% **	-1.12% **	0.72% **
Stock variance	-0.14% **	0.12% **	-0.26% **	0.01% **	1.66% **	1.53% **	-0.26% **	-1.22% **	-1.81% **
Net equity expansion	-4.91% **	-0.19% **	0.29% **	-0.47% **	0.28% **	0.87% **	-8.75% **	-0.60% **	-0.20% **
Inflation	0.01% **	0.40% **	0.95% **	-0.26% **	0.81% **	1.63% **	0.25% **	0.05% **	0.35% **

This table reports the out-of-sample R^2 of unconstrained and constrained univariate prediction models for the quarterly excess return, r_{t+1} , measured relative to the prevailing mean model:

$$R_{OOS}^2 = 1 - \frac{\sum_{\tau=\underline{t}-1}^{\bar{t}-1} (r_{t+1} - \hat{r}_{t+1|t})^2}{\sum_{\tau=\underline{t}-1}^{\bar{t}-1} (\tau_{t+1} - \bar{\tau}_{t+1|t})^2},$$

where $\hat{r}_{t+1|t}$ is the posterior mean of the predictive return distribution based on a regression of quarterly excess returns on an intercept and a lagged predictor variable, x_t : $r_{t+1} = \mu + \beta x_t + \varepsilon_{t+1}$. $\bar{r}_{t+1|t}$ is the forecast from the prevailing mean model which assumes that $\beta = 0$. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T-1$ and information set \mathcal{D}^t , while the Sharpe ratio (SR) constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T-1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. Panel A shows results for the full sample (1947-2010), while panels B and C report results for the subsamples 1947-1978 and 1979-2010. Bold figures highlight instances where the constrained R_{OOS}^2 is higher than its unconstrained counterpart. We measure statistical significance relative to the prevailing mean model using the Clark-West (2007) test statistic. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 4: Out-of-sample forecast performance (annual)

Variables	Panel A: full sample (1947-2010)			Panel B: 1947-1978			Panel C: 1979-2010		
	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint
Log dividend price ratio	1.98% *	5.92% ***	5.35% ***	9.34% **	11.98% ***	10.21% ***	-5.69% ***	-0.38% ***	0.30% ***
Log dividend yield	-8.54% *	6.49% **	4.29% **	-2.46% **	15.86% ***	10.01% ***	-14.86% ***	-3.27% ***	-1.66% ***
Log earning price ratio	1.66% *	5.34% **	4.84% ***	10.99% **	8.87% **	8.10% ***	-8.05% ***	1.67% ***	1.45% ***
Log smooth earning price ratio	-5.78% **	7.16% ***	6.46% ***	3.46% **	14.27% ***	11.74% ***	-15.39% ***	-0.24% ***	0.98% ***
Log dividend-payout ratio	-4.00%	2.01%	3.90% ***	-4.61%	3.87% *	7.59% ***	-3.37% ***	0.08% ***	0.05% ***
Book-to-market ratio	-8.52%	3.00%	4.70% **	2.46%	6.27% *	8.57% ***	-19.96% ***	-0.41% ***	0.67% ***
T-Bill rate	-8.95%	2.50%	3.67% **	-4.81%	6.18% **	7.47% ***	-13.27% ***	-1.33% ***	-0.28% ***
Long-term yield	-8.31%	2.27%	4.12% ***	-3.86%	5.33% **	7.89% ***	-12.94% ***	-0.91% ***	0.20% ***
Long-term return	-5.67%	3.24% **	4.83% ***	1.46%	9.57% **	8.64% ***	-13.08% ***	-3.34% ***	0.87% ***
Term spread	-5.39%	2.40%	3.80% ***	-8.25%	5.47% **	6.82% ***	-2.41% ***	-0.80% ***	0.66% ***
Default yield spread	-1.43%	2.86% **	3.33% **	-2.34%	5.01% **	5.94% ***	-0.49% ***	0.62% ***	0.61% ***
Default return spread	-3.35%	1.38%	3.42% ***	4.67%	4.42% *	6.49% ***	-11.70% ***	-1.78% ***	0.22% ***
Stock variance	-1.09%	2.47% **	3.31% **	-1.16%	4.28% **	5.67% **	-1.02% **	0.58% **	0.85% **
Net equity expansion	-16.23%	0.07%	3.15% ***	-2.64%	2.20%	6.63% ***	-30.37% ***	-2.15% ***	-0.47% ***
Inflation	-6.29%	0.80%	3.16% ***	-10.19%	1.96%	6.09% ***	-2.24% ***	-0.40% ***	0.11% ***
Percent equity issuing	-4.40%	1.66%	3.41% ***	10.69%	4.81% *	6.77% ***	-20.11% ***	-1.61% ***	-0.08% ***

This table reports the out-of-sample R^2 of unconstrained and constrained univariate prediction models for the annual excess return, r_{t+1} , measured relative to the prevailing mean model:

$$R_{OoS}^2 = 1 - \frac{\sum_{t=\bar{T}-1}^{\bar{T}-1} (r_{t+1} - \hat{r}_{t+1|t})^2}{\sum_{t=\bar{T}-1}^{\bar{T}-1} (r_{t+1} - \bar{r}_{t+1|t})^2},$$

where $\hat{r}_{t+1|t}$ is the posterior mean of the predictive return distribution based on a regression of annual excess returns on an intercept and a lagged predictor variable, x_t : $r_{t+1} = \mu + \beta x_t + \varepsilon_{t+1}$. $\bar{r}_{t+1|t}$ is the forecast from the prevailing mean model which assumes that $\beta = 0$. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T-1$ and information set \mathcal{D}^t , while the Sharpe ratio (SR) constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T-1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from the realized volatility. Panel A shows results for the full sample (1947-2010), while panels B and C report results for the subsamples 1947-1978 and 1979-2010. Bold figures highlight instances where the constrained R_{OoS}^2 is higher than its unconstrained counterpart. We measure statistical significance relative to the prevailing mean model using the Clark-West (2007) test statistic. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 5: Economic performance of portfolios based on out-of-sample return forecasts

Variables	Panel A: Monthly			Panel B: Quarterly			Panel C: Annual		
	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint
Log dividend price ratio	-0.28%	0.32%	0.74%	-0.18%	0.38%	0.50%	-0.11%	0.39%	0.92%
Log dividend yield	-0.26%	0.41%	0.76%	-0.08%	0.39%	0.62%	-0.04%	0.50%	0.90%
Log earning price ratio	0.13%	0.49%	0.84%	0.18%	0.54%	0.33%	0.18%	0.47%	0.79%
Log smooth earning price ratio	-0.29%	0.38%	0.81%	-0.16%	0.44%	0.36%	0.03%	0.58%	1.03%
Log dividend-payout ratio	0.20%	0.37%	0.70%	0.15%	0.35%	0.24%	-0.09%	0.31%	0.63%
Book-to-market ratio	-0.76%	0.28%	0.63%	-0.54%	0.38%	0.19%	-0.27%	0.44%	0.53%
T-Bill rate	-0.15%	0.25%	1.06%	-0.13%	0.24%	0.85%	-0.37%	0.27%	0.63%
Long-term yield	-0.26%	0.25%	1.00%	-0.22%	0.26%	0.70%	-0.25%	0.26%	0.76%
Long-term return	-0.03%	0.35%	1.41%	-0.03%	0.33%	0.01%	0.22%	0.44%	0.75%
Term spread	0.21%	0.29%	0.96%	0.15%	0.28%	0.52%	-0.07%	0.32%	0.56%
Default yield spread	-0.11%	0.10%	0.56%	-0.19%	0.12%	0.00%	-0.04%	0.24%	0.49%
Default return spread	-0.06%	0.31%	0.65%	-0.74%	0.03%	0.30%	-0.46%	0.22%	0.53%
Stock variance	0.17%	-0.14%	0.61%	0.00%	0.00%	-0.20%	-0.02%	0.21%	0.55%
Net equity expansion	-0.05%	0.25%	0.77%	-0.26%	0.29%	0.04%	-1.75%	0.19%	0.50%
Inflation	-0.09%	0.26%	0.72%	0.02%	0.31%	0.39%	-0.16%	0.27%	0.54%
Log total net payout yield	-0.30%	0.30%	0.68%						
Percent equity issuing							-0.57%	0.25%	0.58%

This table reports certainty equivalent values for portfolio decisions based on recursive out-of-sample forecasts of excess returns. Each period an investor with power utility and coefficient of relative risk aversion of five selects stocks and T-bills based on his predictive density. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int(\mu + \beta x_t)p(\mu, \beta | \mathcal{D}^t)d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio (SR) constraint imposes that $0 \leq \hat{r}_{t+1|t}/\hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. Certainty equivalent values are annualized and are measured relative to the prevailing mean model which assumes a constant equity premium. Bold figures highlight instances where the constrained annualized certainty equivalent values are higher than their unconstrained counterparts. All results are based on the sample period 1947-2010.

Table 6: Economic performance measures: subsamples

Variables	1947-1978											
	Panel A: Monthly				Panel B: Quarterly				Panel C: Annual			
	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	No constraint	Equity premium constraint	SR constraint	No constraint
Log dividend price ratio	0.06%	0.57%	1.41%	0.22%	0.67%	0.82%	0.22%	0.69%	0.22%	0.69%	1.11%	0.22%
Log dividend yield	0.16%	0.68%	1.46%	0.30%	0.66%	1.05%	0.69%	1.00%	0.69%	1.00%	1.21%	0.69%
Log earning price ratio	0.19%	0.54%	1.16%	0.30%	0.61%	0.16%	0.33%	0.59%	0.33%	0.59%	0.64%	0.33%
Log smooth earning price ratio	-0.36%	0.52%	1.47%	-0.14%	0.64%	0.70%	0.13%	0.82%	0.13%	0.82%	1.22%	0.13%
Log dividend-payout ratio	0.18%	0.51%	1.28%	0.20%	0.51%	0.36%	-0.02%	0.52%	-0.02%	0.52%	0.58%	-0.02%
Book-to-market ratio	-0.48%	0.48%	1.01%	-0.35%	0.58%	0.10%	0.04%	0.59%	0.04%	0.59%	0.23%	0.04%
T-Bill rate	0.30%	0.43%	1.88%	0.19%	0.39%	1.38%	-0.29%	0.52%	-0.29%	0.52%	0.64%	-0.29%
Long-term yield	0.04%	0.44%	1.79%	0.01%	0.43%	1.17%	0.06%	0.49%	0.06%	0.49%	0.71%	0.06%
Long-term return	-0.38%	0.47%	1.94%	-0.08%	0.44%	-0.14%	0.40%	0.71%	0.40%	0.71%	0.59%	0.40%
Term spread	0.33%	0.46%	1.50%	0.25%	0.47%	0.70%	-0.30%	0.55%	-0.30%	0.55%	0.33%	-0.30%
Default yield spread	-0.04%	0.27%	1.09%	-0.15%	0.30%	0.04%	-0.02%	0.36%	-0.02%	0.36%	0.27%	-0.02%
Default return spread	-0.31%	0.49%	0.89%	-1.09%	0.24%	-0.11%	0.20%	0.51%	0.20%	0.51%	0.40%	0.20%
Stock variance	0.01%	0.15%	1.11%	0.05%	0.28%	0.04%	0.03%	0.35%	0.03%	0.35%	0.34%	0.03%
Net equity expansion	0.19%	0.43%	1.22%	0.10%	0.48%	-0.02%	0.06%	0.51%	0.06%	0.51%	0.45%	0.06%
Inflation	-0.03%	0.38%	1.29%	-0.03%	0.44%	0.27%	-0.24%	0.46%	-0.24%	0.46%	0.27%	-0.24%
Log total net payout yield	0.26%	0.48%	1.18%									
Percent equity issuing							0.86%	0.53%	0.86%	0.53%	0.45%	0.86%
Variables	1979-2010											
	Panel A: Monthly				Panel B: Quarterly				Panel C: Annual			
	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	No constraint	Equity premium constraint	SR constraint	No constraint
Log dividend price ratio	-0.61%	0.07%	0.07%	-0.58%	0.08%	0.17%	-0.47%	0.07%	-0.47%	0.07%	0.72%	-0.47%
Log dividend yield	-0.69%	0.14%	0.06%	-0.45%	0.12%	0.18%	-0.81%	-0.02%	-0.81%	-0.02%	0.57%	-0.81%
Log earning price ratio	0.08%	0.44%	0.51%	0.05%	0.47%	0.52%	0.02%	0.35%	0.02%	0.35%	0.97%	0.02%
Log smooth earning price ratio	-0.22%	0.25%	0.15%	-0.19%	0.23%	0.02%	-0.08%	0.32%	-0.08%	0.32%	0.83%	-0.08%
Log dividend-payout ratio	0.21%	0.23%	0.12%	0.10%	0.19%	0.11%	-0.16%	0.07%	-0.16%	0.07%	0.68%	-0.16%
Book-to-market ratio	-1.04%	0.08%	0.24%	-0.73%	0.18%	0.28%	-0.60%	0.28%	-0.60%	0.28%	0.86%	-0.60%
T-Bill rate	-0.61%	0.06%	0.25%	-0.45%	0.09%	0.31%	-0.46%	0.00%	-0.46%	0.00%	0.63%	-0.46%
Long-term yield	-0.56%	0.06%	0.21%	-0.47%	0.10%	0.22%	-0.58%	0.00%	-0.58%	0.00%	0.80%	-0.58%
Long-term return	0.33%	0.23%	0.87%	0.02%	0.22%	0.16%	0.02%	0.15%	0.02%	0.15%	0.92%	0.02%
Term spread	0.09%	0.11%	0.42%	0.05%	0.10%	0.33%	0.19%	0.07%	0.19%	0.07%	0.82%	0.19%
Default yield spread	-0.17%	-0.06%	0.03%	-0.22%	-0.07%	-0.03%	-0.05%	0.11%	-0.05%	0.11%	0.74%	-0.05%
Default return spread	0.19%	0.13%	0.40%	-0.38%	-0.19%	0.72%	-1.15%	-0.10%	-1.15%	-0.10%	0.67%	-1.15%
Stock variance	0.33%	-0.44%	0.12%	-0.05%	-0.28%	-0.43%	-0.08%	0.05%	-0.08%	0.05%	0.79%	-0.08%
Net equity expansion	-0.29%	0.06%	0.31%	-0.63%	0.10%	0.11%	-3.55%	-0.15%	-3.55%	-0.15%	0.56%	-3.55%
Inflation	-0.16%	0.15%	0.16%	0.07%	0.17%	0.51%	-0.08%	0.06%	-0.08%	0.06%	0.84%	-0.08%
Log total net payout yield	-0.86%	0.11%	0.17%									
Percent equity issuing							-2.02%	-0.04%	-2.02%	-0.04%	0.72%	-2.02%

This table reports certainty equivalent values for portfolio decisions based on recursive out-of-sample forecasts of excess returns. Each period an investor with power utility and coefficient of relative risk aversion of five selects stocks and T-bills based on his predictive density. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio (SR) constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. Certainty equivalent values are annualized and are measured relative to the prevailing mean model which assumes a constant equity premium. Bold figures highlight instances where the constrained annualized certainty equivalent values are higher than their unconstrained counterparts.

Table 7: Out-of-sample forecast and economic performance: Multivariate methods

Method	Monthly			Quarterly			Annual		
	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint
Panel A: Forecast Performance (out of sample R-squared)									
Full sample									
Equal weighted combination	0.62% ***	0.25% *	0.31% *	1.19% **	0.77% *	1.32% ***	4.48% **	4.13% **	4.45% ***
Diffusion index	0.23% **	0.74% ***	0.76% ***	0.42% **	3.02% ***	2.95% ***	-1.53% ***	7.57% ***	6.33% ***
1947-1978									
Equal weighted combination	1.09% ***	0.53% *	0.65% *	2.57% **	1.59% *	2.38% ***	9.26% ***	8.06% **	7.99% ***
Diffusion index	1.62% ***	1.46% ***	1.53% ***	3.09% **	4.99% ***	4.93% ***	1.71% ***	13.68% ***	11.44% ***
1979-2010									
Equal weighted combination	0.26% -0.83%	0.06% 0.22%	0.09% 0.19%	-0.02% -1.89%	0.07% 1.32%	0.40% 1.24%	-0.50% -4.89%	0.04% 1.21%	0.76% 1.00%
Panel B: Economic Performance (Certainty equivalent return)									
Full sample									
Equal weighted combination	0.14% -0.07%	0.27% 0.44%	0.86% 1.02%	0.02% -0.04%	0.33% 0.53%	0.41% 0.95%	0.07% -0.25%	0.37% 0.46%	0.69% 1.02%

Panel A reports the out-of-sample R^2 of unconstrained and constrained multivariate prediction models for the excess return, r_{t+1} , measured relative to the prevailing mean model:

$$R_{OoS}^2 = 1 - \frac{\sum_{\tau=\underline{t}-1}^{\bar{t}-1} (r_{t+1} - \hat{r}_{t+1|t})^2}{\sum_{\tau=\underline{t}-1}^{\bar{t}-1} (r_{t+1} - \bar{r}_{t+1|t})^2},$$

where $\hat{r}_{t+1|t}$ is the posterior mean of the predictive return distribution based on either an equal-weighted forecast combination or a diffusion index approach. The equity premium constrained model imposes that $\hat{r}_{t+1|t} > 0$, for $t = 1, \dots, T-1$ and information set \mathcal{D}^t , while the Sharpe ratio (SR) constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T-1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. Panel B reports certainty equivalent values for portfolio decisions based on recursive out-of-sample forecasts of excess returns. Each period an investor with power utility and coefficient of relative risk aversion of five selects stocks and T-bills based on his predictive density. Certainty equivalent values are annualized and are measured relative to the prevailing mean model which assumes a constant equity premium. Bold figures highlight instances where the constrained R_{OoS}^2 (Panel A) or annualized certainty equivalent returns (Panel B) are higher than their unconstrained counterparts. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 8: Out-of-sample forecast performance: Recessions and expansions

Monthly										
Variables	Panel A: Expansions					Panel B: Recessions				
	No constraint	Equity premium constraint	SR constraint			No constraint	Equity premium constraint	SR constraint		
Log dividend price ratio	-0.93%	0.72% ***	0.70% ***			2.39% ***	0.47%		0.02%	
Log dividend yield	-1.77%	0.77% ***	0.67% ***			3.44% ***	0.75% *		0.27%	
Log earning price ratio	-0.96% **	0.85% ***	0.75% ***			-2.32%	-1.07%		-0.72%	
Log smooth earning price ratio	-2.25%	0.79% ***	0.88% ***			2.45% **	-0.21%		-0.43%	
Log dividend-payout ratio	-0.50%	0.34% *	0.41% **			-3.95%	-1.13%		-0.83%	
Book-to-market ratio	-2.08%	0.38% *	0.46% **			0.12%	-0.79%		-0.91%	
T-Bill rate	-0.77%	0.30% *	0.78% ***			1.75% *	-0.36%		-0.01%	
Long-term yield	-1.81%	0.50% **	0.79% ***			0.98%	-0.09%		0.04%	
Long-term return	-1.35%	0.40% **	-0.17%			0.61%	-0.29%		2.07% **	
Term spread	-0.37%	0.28% *	0.49% **			1.01% *	-0.48%		-0.31%	
Default yield spread	-0.28%	0.30% **	0.41% *			0.02%	-0.07%		-1.18%	
Default return spread	-0.16%	0.46% **	0.42% **			-0.41%	-0.52%		-1.90%	
Stock variance	-0.08%	0.03%	0.24%			0.95%	-0.62%		-1.47%	
Net equity expansion	0.49% **	0.37% **	0.48% **			-3.66%	-0.86%		-1.32%	
Inflation	-0.10%	0.39% **	0.59% **			-0.20%	-0.63%		-1.70%	
Log total net payout yield	-1.21%	0.52% **	0.48% **			1.12%	-0.55%		-1.13%	

Quarterly										
Variables	Panel A: Expansions					Panel B: Recessions				
	No constraint	Equity premium constraint	SR constraint			No constraint	Equity premium constraint	SR constraint		
Log dividend price ratio	-4.35%	2.81% ***	1.81% ***			6.42% ***	1.00%		1.46% *	
Log dividend yield	-2.54%	2.66% ***	1.91% ***			5.94% ***	1.08%		1.82% **	
Log earning price ratio	-5.82% *	2.76% ***	1.63% ***			-7.13%	-2.53%		-0.20%	
Log smooth earning price ratio	-9.43%	2.93% ***	2.29% ***			5.63% **	-0.57%		0.50%	
Log dividend-payout ratio	-1.24%	1.21% **	1.35% **			-7.56%	-1.80%		-0.02%	
Book-to-market ratio	-10.20%	1.25% **	1.46% **			-0.02%	-2.36%		0.17%	
T-Bill rate	-1.61%	1.04% **	2.31% ***			1.17%	-1.61%		1.19%	
Long-term yield	-4.71%	1.36% **	1.87% ***			0.47%	-0.56%		1.05%	
Long-term return	0.70%	1.73% **	1.83% **			-2.74%	-1.59%		-1.17%	
Term spread	-0.51%	0.96% *	1.40% *			1.09%	-1.41%		0.90%	
Default yield spread	-1.25%	1.14% **	1.43% **			-0.16%	-0.50%		-1.18%	
Default return spread	-4.07%	1.00% **	1.33% **			-9.40%	-4.70%		-0.11%	
Stock variance	-0.09%	1.31% ***	1.59% **			-0.21%	-1.76%		-3.16%	
Net equity expansion	-0.30% **	1.28% **	1.43% **			-12.17%	-2.50%		-1.49%	
Inflation	-0.41%	1.36% **	1.27% **			0.68%	-1.11%		0.44%	

This table reports the out-of-sample R^2 of unconstrained and constrained univariate prediction models for the monthly excess return, r_{t+1} , measured relative to the prevailing mean model:

$$R_{OoS}^2 = 1 - \frac{\sum_{\tau=\underline{t}-1}^{\bar{t}-1} (r_{t+1} - \hat{r}_{t+1|t})^2}{\sum_{\tau=\underline{t}-1}^{\bar{t}-1} (r_{t+1} - \bar{r}_{t+1|t})^2},$$

where $\hat{r}_{t+1|t}$ is the posterior mean of the predictive return distribution based on a regression of monthly excess returns on an intercept and a lagged predictor variable, x_t : $r_{t+1} = \mu + \beta x_t + \varepsilon_{t+1}$. $\bar{r}_{t+1|t}$ is the forecast from the prevailing mean model which assumes that $\beta = 0$. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T-1$ and information set \mathcal{D}^t , while the Sharpe ratio (SR) constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T-1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. Panels A and B show monthly results for expansions and recessions, while panels C and D show quarterly results for expansions and recessions. Bold figures highlight instances where the constrained R_{OoS}^2 is higher than its unconstrained counterpart. * significance at 10% level; ** significance at 5% level; *** significance at 1% level.

Table 9: Effect of risk aversion on economic performance measures

A = 2											
Variables	Panel A: Monthly			Panel B: Quarterly			Panel C: Annual				
	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint		
Log dividend price ratio	-0.89%	0.68%	0.78%	-0.47%	0.90%	1.02%	-0.20%	0.83%	1.09%		
Log dividend yield	-1.06%	0.66%	0.77%	-0.22%	0.92%	1.00%	-0.51%	0.90%	1.09%		
Log earning price ratio	0.15%	1.04%	0.75%	0.32%	1.41%	0.96%	0.36%	1.09%	1.06%		
Log smooth earning price ratio	-1.11%	0.72%	0.78%	-0.85%	0.96%	1.14%	-0.33%	1.04%	1.08%		
Log dividend-payout ratio	0.70%	0.83%	0.64%	0.79%	0.86%	0.93%	0.15%	0.76%	1.00%		
Book-to-market ratio	-1.36%	0.73%	0.73%	-1.05%	1.13%	1.06%	-0.64%	1.13%	1.24%		
T-Bill rate	-0.29%	0.64%	0.89%	-0.24%	0.59%	1.03%	-0.92%	0.71%	1.01%		
Long-term yield	-0.69%	0.61%	0.72%	-0.54%	0.63%	0.94%	-0.71%	0.65%	1.02%		
Long-term return	-0.12%	0.77%	1.31%	-0.11%	0.82%	0.67%	0.38%	0.87%	1.07%		
Term spread	0.67%	0.76%	0.82%	0.85%	0.71%	1.02%	-0.07%	0.84%	1.03%		
Default yield spread	-0.17%	0.33%	0.83%	-0.37%	0.37%	1.08%	-0.11%	0.53%	1.02%		
Default return spread	0.01%	0.76%	0.83%	-1.17%	0.68%	1.01%	0.10%	0.80%	1.02%		
Stock variance	0.34%	0.26%	0.74%	0.00%	0.36%	1.13%	-0.09%	0.48%	1.00%		
Net equity expansion	0.45%	0.72%	0.77%	0.71%	0.91%	0.99%	0.52%	0.96%	0.93%		
Inflation	-0.12%	0.69%	0.59%	0.03%	0.73%	0.97%	-0.01%	0.83%	0.98%		
Log total net payout yield	-0.83%	0.73%	0.69%				1.02%	0.85%	0.91%		
Percent equity issuing											

A = 10											
Variables	Panel A: Monthly			Panel B: Quarterly			Panel C: Annual				
	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint	No constraint	Equity premium constraint	SR constraint		
Log dividend price ratio	-0.13%	0.18%	0.39%	-0.10%	0.19%	0.27%	-0.09%	0.19%	0.49%		
Log dividend yield	-0.13%	0.20%	0.41%	-0.06%	0.19%	0.34%	-0.04%	0.27%	0.47%		
Log earning price ratio	0.08%	0.27%	0.44%	0.09%	0.26%	0.19%	0.04%	0.24%	0.44%		
Log smooth earning price ratio	-0.15%	0.20%	0.43%	-0.10%	0.23%	0.21%	-0.05%	0.29%	0.55%		
Log dividend-payout ratio	0.12%	0.19%	0.36%	0.09%	0.18%	0.16%	-0.03%	0.17%	0.32%		
Book-to-market ratio	-0.36%	0.16%	0.33%	-0.25%	0.20%	0.12%	-0.16%	0.24%	0.29%		
T-Bill rate	-0.08%	0.13%	0.55%	-0.08%	0.12%	0.44%	-0.20%	0.15%	0.35%		
Long-term yield	-0.13%	0.13%	0.52%	-0.14%	0.13%	0.36%	-0.13%	0.14%	0.40%		
Long-term return	0.00%	0.19%	0.74%	-0.03%	0.17%	-0.02%	0.08%	0.24%	0.37%		
Term spread	0.12%	0.16%	0.49%	0.06%	0.15%	0.29%	-0.07%	0.18%	0.32%		
Default yield spread	-0.04%	0.06%	0.29%	-0.10%	0.05%	0.03%	-0.02%	0.13%	0.27%		
Default return spread	-0.07%	0.16%	0.28%	-0.38%	0.02%	0.21%	-0.25%	0.12%	0.28%		
Stock variance	0.10%	-0.08%	0.33%	-0.02%	-0.01%	-0.09%	-0.04%	0.11%	0.29%		
Net equity expansion	-0.12%	0.13%	0.40%	-0.36%	0.15%	0.07%	-1.49%	0.11%	0.24%		
Inflation	-0.04%	0.13%	0.37%	0.00%	0.14%	0.23%	-0.09%	0.15%	0.26%		
Log total net payout yield	-0.15%	0.15%	0.34%				-0.25%	0.13%	0.31%		
Percent equity issuing											

This table reports certainty equivalent values for portfolio decisions based on recursive out-of-sample forecasts of excess returns. Each period an investor with power utility and coefficient of relative risk aversion of two (top panels) or ten (bottom panels) selects stocks and T-bills based on his predictive density. The equity premium constrained model imposes that $\hat{r}_{t+1|t} = \int (\mu + \beta x_t) p(\mu, \beta | \mathcal{D}^t) d\mu d\beta > 0$, for $t = 1, \dots, T - 1$ and information set \mathcal{D}^t , while the Sharpe ratio (SR) constraint imposes that $0 \leq \hat{r}_{t+1|t} / \hat{\sigma}_{t+1|t} \leq 1$, for $t = 1, \dots, T - 1$, where $\hat{\sigma}_{t+1|t}$ is the posterior volatility estimate obtained from a stochastic volatility model. Certainty equivalent values are annualized and are measured relative to the prevailing mean model which assumes a constant equity premium. Bold figures highlight instances where the constrained annualized certainty equivalent values are higher than their unconstrained counterparts.

