JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 56, No. 4, June 2021, pp. 1228–1258
© THE AUTHOR(S), 2020. PUBLISHED BY CAMBRIDGE UNIVERSITY PRESS ON BEHALF OF THE MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON
doi:10.1017/S002210902000054X

When Are Stocks Less Volatile in the Long Run?

Eric Jondeau
Swiss Finance Institute and University of Lausanne
eric.jondeau@unil.ch

Qunzi Zhang
Shandong University School of Economics
qunzi.zhang@sdu.edu.cn

Xiaoneng Zhu
Shanghai University of Finance and Economics School of Finance and Shanghai Institute of International Finance and Economics
xiaonengz@gmail.com (corresponding author)

Abstract

Pástor and Stambaugh (2012) find that from a forward-looking perspective, stocks are more volatile in the long run than they are in the short run. We demonstrate that when the nonnegative equity premium (NEP) condition is imposed on predictive regressions, stocks are in fact less volatile in the long run, even after taking estimation risk and uncertainties into account. The reason is that the NEP provides an additional parameter identification condition and prior information for future returns. Combined with the mean reversion of stock returns, this condition substantially reduces uncertainty on future returns and leads to lower long-run predictive variance.

I. Introduction

A prevailing view in finance is that stocks are less volatile in the long run than they are in the short run (Siegel (1994), Campbell and Viceira (2002), (2005)). An important implication of this evidence is that a long-term investor should hold more stocks and fewer bonds than a short-term investor (Barberis (2000), Campbell, Chan and Viceira (2003), Campbell and Viceira (2005)). However, recent work by Pástor and Stambaugh (2012) points out that this view holds true when one considers the true variance of the return process (assuming the model's parameters are known) but does not hold for the predictive variance, which matters for an investor's portfolio decisions. In a predictive system that recognizes the imperfection of predictors and takes all parameter uncertainty into account, Pástor and Stambaugh (2012) demonstrate that from a forward-looking perspective, stocks

Zhang acknowledges financial support from the Natural Science Foundation of China (No. 71903112, 71873079), National Social Science Fund (No. 18ZDA078), China Postdoctoral Science Foundation (No. 2020T130380), and Shandong University Research Fund (No. IFYT19001). Zhu acknowledges financial support from the Natural Science Foundation of China (No. 71473281), Shanghai University of Finance and Economics (SUFE) (No. 2018110698), and Shanghai Institute of International Finance and Economics (SIIFE) (No. 2018110262). The order of authors' names is alphabetical.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

are actually more volatile over long horizons than they are over short horizons. Johannes, Korteweg and Polson (2014) obtain similar results using several alternative specifications. In particular, they find that in a model with a time-varying regression coefficient (constant variance–drifting coefficient, or CV-DC, model), predictive variance also increases with the investment horizon.

In this article, we provide evidence that stocks are in fact less volatile in the long run, even from an investor's perspective, when the nonnegativity of the equity premium is taken into account. A basic tenet in finance is that risk-averse investors must be compensated for bearing systematic risks, which implies that such investors expect a nonnegative equity premium (NEP). Recent studies find that the forecasting accuracy of stock returns is substantially improved when the NEP condition is imposed on predictive models (Campbell and Thompson (2008), Pettenuzzo, Timmermann, and Valkanov (2014), Penasse (2016)). In this spirit, we incorporate the NEP condition into two standard predictive models (the predictive system of Pástor and Stambaugh (2012) and the CV-DC model of Johannes et al. (2014)) and find that the predictive variance is indeed lower in the long run. This result has important implications for long-horizon portfolio managers, such as pension funds, endowment funds, and sovereign wealth funds.

It is worth noting that our result is not inconsistent with Pástor and Stambaugh (2012) and Johannes et al. (2014). In our sample, we also find that predictive variance effectively increases with the investment horizon when the predictive models do not impose the NEP condition. Investors find that stocks are less volatile in the long run only when they expect the market risk premium to be nonnegative and thus incorporate the NEP condition into the predictive models. Our result thus enriches the understanding of predictive variance by showing that investors' beliefs on the NEP condition determine whether the predictive variance is higher or lower in the long run from an investor's perspective.

To estimate the predictive system and the CV-DC model, we use a Bayesian approach that takes into account the uncertainty surrounding the parameters (or estimation risk). Investors face estimation risk when they make financial decisions because the parameters driving the data-generating process are estimated with sampling error, even in a long historical sample. More importantly, Bayesian estimation provides us with a natural way to impose the NEP condition. Specifically, we use the Metropolis–Hastings algorithm to impose the NEP condition on the predictive models, following the approach proposed by O'Donnell, Rambaldi, and Doran (2001), and O'Donnell and Coelli (2005). This algorithm ensures that investors learn from all observations spanning the whole information set to sharpen their stock return forecasts and reduce estimation risk and other uncertainties plaguing the estimation.

To provide insights on the necessity of the NEP condition for predictive models, we employ the Bayes factor to compare alternative model specifications. This analysis provides evidence supporting the predictive specifications that impose the NEP condition. We then demonstrate that the predictive models imposing the NEP condition have a higher explanatory power than the unrestricted models and improve the accuracy of stock return forecasts.

Imposing the NEP condition has important implications for the level of predictive variance over investment horizons. In a predictive system that incorporates the

imperfection of predictors and the negative correlation between unexpected returns and innovations in expected returns, Pástor and Stambaugh (2012) show that predictive variance varies over investment horizons according to 5 components: independent and identically distributed (IID) uncertainty, mean reversion in stock returns, uncertainty about current expected returns, uncertainty about future expected returns, and estimation risk. In the predictive models without the NEP condition, the estimation risk and the uncertainty about future expected returns drive the increase in predictive variance for long investment horizons. In contrast, our empirical analysis demonstrates that imposing the NEP condition on the predictive models substantially reduces the estimation risk and the uncertainty about future expected returns. Consequently, the mean reversion in stock returns dominates the other components and thus leads to a decrease in predictive variance for long horizons.

The intuition behind the reduction in estimation risk and uncertainty about future expected returns is the following. First, we interpret the NEP condition as an identification condition that provides additional information not contained in the data to estimate the parameters of the predictive models (e.g., Pástor (2000), Avramov, Cederburg, and Lučivjanská (2018)). Thus, the NEP condition reduces estimation risk. In addition, this condition provides valuable, theoretically grounded prior information about the equity premium. As a result, it also reduces the uncertainty about future expected returns (e.g., Campbell and Thompson (2008), Pettenuzzo et al. (2014)). Recently, Avramov et al. (2018) have shown that investors with long-run-risk priors perceive weaker mean reversion and riskier equities. Our article is complementary to their study. Whereas they argue that using prior information based on the long-run-risk model may change the shape of the term structure of predictive variance, we find that the NEP condition affects stock variance and implies that stocks are less volatile in the long run.

Finally, we provide evidence on the economic value delivered by the NEP condition for a buy-and-hold investor. We demonstrate that an investor using the unrestricted version of the predictive models would allocate less to stocks when the horizon increases, whereas an investor using the restricted version would allocate more to stocks. In addition, the investor using the restrictive model would invest substantially more in stocks than the investor using the unrestricted model. The Sharpe ratio and certainty equivalent return delivered by the portfolio based on the restricted model are much higher than their counterparts based on the unrestricted model.

Clearly, the result that the NEP condition helps reduce the uncertainty in the long run depends on the model under consideration. In particular, it may not apply to some models, such as stochastic volatility models. In Section VII, we use the long-run-risk model of Bansal and Yaron (2004) to simulate an equity premium that is known to be positive and highly persistent. Based on these simulated data, we estimate the predictive models with and without the NEP condition and measure the out-of-sample predictive performance of these models. Our analysis shows that out of sample, imposing the NEP condition allows investors to generate a higher risk-adjusted performance.

Relevant literature. An important implication of stock return predictability is that the term structure of risk, measured by the annualized variance across investment horizons, should decrease with the horizon. This implication has been

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

analyzed by Siegel (1994) for the sample variance and by Campbell and Viceira (2002), (2005) for the variance conditional on a vector autoregressive model, and they find that the term structure of risk is indeed decreasing. Pástor and Stambaugh (2012) argue that the relevant measure of risk is predictive variance, which involves additional sources of risk and uncertainty. They find that the term structure of predictive variance is in fact increasing with the horizon. Johannes et al. (2014) perform a similar analysis in the context of models with time-varying regression coefficients or time-varying volatility and report similar conclusions. From these two articles, we draw the predictive models that we use to forecast the future variance.

Our article is also related to the strand of research on how to improve the predictability of stock returns. Goyal and Welch (2008) provide evidence that most standard predictors have poor in-sample and out-of-sample performance, notably because of the instability of regression parameters. Several articles have investigated new predictors, for instance, those by Hong and Yogo (2012), Li, Ng, and Swaminathan (2013), van Binsbergen, Hueskes, Koijen, and Vrugt (2013), and Golež (2014). Other articles have addressed the issue of parameter instability using more sophisticated specifications, such as those by Lettau and van Nieuwerburgh (2008), Ferreira and Santa-Clara (2011), and Henkel, Martin, and Nardari (2011).[1] Several authors find that imposing restrictions on the data-generating process may improve the predictability of returns. In particular, Campbell and Thompson (2008), van Binsbergen and Koijen (2010), Pettenuzzo et al. (2014), and Penasse (2016) show that imposing the nonnegativity of the equity premium improves stock return forecasts. For instance, Penasse (2016) considers three types of restrictions in predictive regressions, including the NEP condition. He finds that learning from prior information about the NEP substantially increases the strength of return predictability. In addition, the NEP condition reduces estimation risks. From this strand of research, we draw the NEP condition, which we impose on the predictive system of Pástor and Stambaugh (2009) and Pástor and Stambaugh (2012) and the CV-DC model of Johannes et al. (2014).

Another important field of research investigates the effect of parameter and model uncertainty on stock return predictability. Parameter uncertainty in predictive regressions is addressed using Bayesian techniques by Stambaugh (1999), Barberis (2000), Pástor and Stambaugh (2009), (2012), Shanken and Tamayo (2012), Hoevenaars, Molenaar, Schotman, and Steenkamp (2014), and Johannes et al. (2014). To cope with parameter uncertainty, Avramov (2002) and Cremers (2002) also use Bayesian techniques to select truly predictive variables. Last, Bayesian techniques are used to explore the effect of stock return predictability in asset allocation, particularly in a multiperiod perspective (e.g., Kandel and Stambaugh (1996), Barberis (2000), Pástor (2000), Xia (2001), Avramov (2004), Brandt, Goyal, Santa-Clara, and Stroud (2005), and Wachter and Warusawitharana (2009)).

[1] Another avenue followed by researchers consists of combining forecasts produced by individual predictors. Rapach, Strauss, and Zhou (2010) find that this approach delivers statistically and economically significant out-of-sample gains.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

The remainder of the article is organized as follows: Section II presents the predictive models and discusses the Bayesian estimation of these models. Section III reports the estimates of parameters and the specification testing results. Section IV analyzes the implications of the NEP condition for predictive variance and, in particular, its properties at long horizons. Section V investigates the economic value of imposing the NEP condition for a buy-and-hold investor. Section VI conducts robustness checks. In Section VII, we simulate a long-run-risk model that imposes the nonnegativity of the equity premium and compare the out-of-sample performance of the unrestricted and restricted predictive models in predicting future stock returns. Section VIII concludes.

II. Predictive Models

This section presents the predictive system of Pástor and Stambaugh (2012) and the CV-DC model of Johannes et al. (2014). Both models imply that from a forward-looking perspective, stocks should be more volatile in the long run. In contrast, a typical linear predictive regression or a predictive vector autoregression (VAR) implies a constant predictive variance over investment horizons.

A. The Predictive System

Standard predictive regressions take predictors as a perfect proxy for the expected return, regardless of whether parameter uncertainty is taken into account. However, investors may realize that predictors are in fact only imperfect proxies. To investigate this issue, Pástor and Stambaugh (2012) develop a predictive system that captures the imperfection of the predictors.

The predictive system allows a predictor, x_t, to be imperfect; that is, the conditional expected return $\mu_t = E_t(r_{t+1})$ is not an exact linear function of x_t. The information contained in the imperfect predictor is supplemented by lags of the return and the predictor. The predictive system is a restricted VAR for (r_{t+1}, x_{t+1}, μ_{t+1}):

$$r_{t+1} = \mu_t + u_{t+1}, \tag{1}$$

$$x_{t+1} = (1-\gamma)E_x + \gamma x_t + v_{t+1}, \tag{2}$$

$$\mu_{t+1} = (1-\beta)E_r + \beta\mu_t + w_{t+1}, \tag{3}$$

where E_x and E_r denote the unconditional mean of the predictor and the unconditional mean return, respectively. Residuals in the predictive system are assumed to be IID over time, as follows:

$$\begin{bmatrix} u_{t+1} \\ v_{t+1} \\ w_{t+1} \end{bmatrix} \sim N(0,\Omega) \quad \text{where} \quad \Omega = \begin{bmatrix} \sigma_u^2 & \sigma_{uv} & \sigma_{uw} \\ \sigma_{vu} & \sigma_v^2 & \sigma_{vw} \\ \sigma_{wu} & \sigma_{wv} & \sigma_w^2 \end{bmatrix}. \tag{4}$$

The predictive system is assumed to be stationary, with $0<\beta<1$ and $0<\gamma<1$. The Bayesian estimation method of the predictive system is described in the Appendix of Pástor and Stambaugh (2009).[2]

Because r_t and x_t are observed but μ_t is not, the predictive system in equations (1)–(3) is a state-space model. Equation (1) defines the unobserved conditional expected return, or equity premium, μ_t. Equation (2) is a standard assumption in the predictability literature, which describes the dynamics of the predictive variable. Equation (3) captures the dynamics of the conditional expected return.

Because the predictive system is a state-space model, the Kalman filter can be used to filter out the expected return. As shown by Pástor and Stambaugh (2009), the conditional expected return can be written as the unconditional expected return plus a linear combination of past return forecast errors and innovations in the predictive variable. Specifically, we define the return forecast error in period $t+1$ as $\varepsilon_{t+1}=r_{t+1}-E(r_{t+1}|Z_t)$, where $Z_t=\{z_i\}_{i=1}^{t}$ is the conditional information set, including the history of the return and the predictor observed up to time t. Then, the conditional expected return is given by

$$E(r_{t+1}|Z_t)=E_r+\sum_{s=0}^{\infty}(\lambda_s\varepsilon_{t-s}+\phi_s v_{t-s}), \tag{5}$$

where in the steady state, $\lambda_s=m\beta^s$ and $\phi_s=n\beta^s$, with m and n being two functions of the parameters in equations (1)–(4).[3] In this predictive system, it is straightforward to show that $E(\mu_t|Z_t)=E(r_{t+1}|Z_t)$.[4]

A widely accepted belief about the expected return is that the unexpected return is negatively correlated with the innovation in expected return. The intuition is that asset prices tend to fall when the discount rate rises. This negative correlation is confirmed by the large body of evidence that an unexpected increase in the discount rate is accompanied by an unanticipated negative return (e.g., Shiller (1981), Campbell (1991), Siegel (1994), Pástor and Stambaugh (2009), (2012), and van Binsbergen and Koijen (2010)). For example, Campbell (1991), Campbell and Ammer (1993), and van Binsbergen and Koijen (2010) consistently find a negative correlation, ranging between -0.71 and -0.95. Because a negative correlation is intuitively appealing and empirically well documented, investors might have strong priors about this large negative correlation. In their Bayesian framework, Pástor and Stambaugh (2009) incorporate this prior into the predictive system. They show that with informative prior beliefs about the negative correlation, the predictive system produces much more accurate

[2]The Supplementary Material provides a more detailed description of the Bayesian estimation method proposed by Pástor and Stambaugh (2009).

[3]There is no simple analytical expression for m and n, but their expression is given in equation (A20) in Appendix A of Pástor and Stambaugh (2009).

[4]Equation (5) also indicates that the conditional expected return depends on the full history of the return and the predictor. This result stems from the fact that the predictive system allows the predictor x_t to be imperfect: Whereas the standard predictive regression implies $\mu_t=a+bx_t$, the predictive system allows $\mu_t\neq a+bx_t$. When the predictor is imperfect, its predictive power is supplemented by lags of the return and the predictor. The predictive system captures this type of imperfection in a parsimonious way.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

estimates of expected returns. An important feature of the predictive system is that the effect of predictor imperfections on the estimate of the expected return depends on the correlation between the unexpected return and the innovation in expected return. We follow Pástor and Stambaugh and also impose a negative-correlation prior on the estimation.

In the predictive system, μ_t is a first-order autoregressive (AR(1)) process, given by equation (3). According to the Wold representation, this AR(1) process is rewritten as an infinite moving average (MA) process:

$$\mu_t = E_r + \sum_{i=0}^{\infty} \beta^i w_{t-i}. \tag{6}$$

Using equations (1) and (6), the predictive system implies that the k-period-ahead return is

$$\begin{aligned} r_{t+k} &= E_r + \sum_{i=0}^{\infty} \beta^i w_{t+k-1-i} + u_{t+k} \\ &= \left(1-\beta^{k-1}\right)E_r + \beta^{k-1}\mu_t + \sum_{i=1}^{k-1} \beta^{k-1-i} w_{t+i} + u_{t+k}. \end{aligned} \tag{7}$$

We now define the cumulative excess stock return between t and $t+k$ as $R_{t:t+k} = \sum_{i=1}^{k} r_{t+i}$. The predictive system implies the following:

$$R_{t:t+k} = kE_r + \frac{1-\beta^k}{1-\beta}(\mu_t - E_r) + \sum_{i=1}^{k-1} \frac{1-\beta^{k-i}}{1-\beta} w_{t+i} + \sum_{i=1}^{k} u_{t+i}. \tag{8}$$

Taking parameter uncertainty into account, the conditional predictive variance of $R_{t:t+k}$ is decomposed as

$$\mathrm{var}(R_{t:t+k}|Z_t) = E(\mathrm{var}(R_{t:t+k}|\Theta, Z_t)|Z_t) + \mathrm{var}(E(R_{t:t+k}|\Theta, Z_t)|Z_t), \tag{9}$$

where Θ denotes the vector of parameters. The first term is the expectation of the conditional variance of the k-period return. The second term is the variance of the expected cumulative return, which reflects that this true value is unknown. Pástor and Stambaugh (2012) provide the solution for this predictive variance. They show that the conditional predictive variance includes 5 components:[5]

$$\begin{aligned} \mathrm{var}(R_{t:t+k}|Z_t) = {} & E\left(k\sigma_u^2|Z_t\right) + E\left(2k\sigma_u^2 \overline{d}\rho_{uw}A(k)|Z_t\right) + E\left(k\sigma_u^2\overline{d}^2 B(k)|Z_t\right) \\ & + E\left(\left(\frac{1-\beta^k}{1-\beta}\right)^2 q_t|Z_t\right) + \mathrm{var}\left(kE_r + \frac{1-\beta^k}{1-\beta}(b_t - E_r)|Z_t\right), \end{aligned} \tag{10}$$

[5]This calculation of the predictive variance is based on the assumption that the parameters of the predictive system remain constant over the sample period. As Pástor and Stambaugh (2012) stress, this assumption simplifies the analysis.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

where

$$A(k) = 1 + \frac{1}{k}\left(-1 - \beta\frac{1-\beta^{k-1}}{1-\beta}\right)$$
$$B(k) = 1 + \frac{1}{k}\left(-1 - 2\beta\frac{1-\beta^{k-1}}{1-\beta} + \beta^2\frac{1-\beta^{2(k-1)}}{1-\beta^2}\right)$$
$$\overline{d} = \left(\frac{1+\beta}{1-\beta}\,\frac{R^2}{1-R^2}\right)^{1/2}$$
$$b_t = E(\mu_t|\Theta, Z_t)$$
$$q_t = Var(\mu_t|\Theta, Z_t),$$

and R^2 is the ratio of the variance of μ_t to the variance of r_{t+1}. The first three terms on the right-hand side of equation (10) correspond to $E(\text{var}(R_{t:t+k}|\Theta, Z_t)|Z_t)$, and the last two terms correspond to $\text{var}(E(R_{t:t+k}|\Theta, Z_t)|Z_t)$. These terms can be interpreted as follows: The first term captures the well-known feature of IID returns that the variance of the cumulative k-period return increases linearly with k; the second term reflects the mean reversion in returns arising from the likely negative correlation between the unexpected return and the innovations in expected return ($\rho_{uw} < 0$), which contributes negatively to long-horizon variance; the third term captures the effect of uncertainty about future values of μ_t, which contributes positively to the long-horizon variance; the fourth term reflects the uncertainty about current μ_t, or predictor imperfection; and the last term captures the effect of parameter uncertainty.

B. The CV-DC Model

Johannes et al. (2014) consider various specifications to predict stock returns. Whereas some specifications imply that predictive variance decreases with the investment horizon, the specification that delivers the most striking increasing term structure of predictive variance is the CV-DC model. We therefore consider this model in this study.

In essence, the CV-DC model is a VAR predictive system with time-varying regression coefficients. Compared with a typical predictive regression, which assumes a stable relationship between predictors and expected returns, this specification allows for the time-varying predictability of stock returns. This specification is consistent with theoretical arguments (Menzly, Santos, and Veronesi (2004), Santos and Veronesi (2006)) and supported by empirical evidence reported in previous studies (e.g., Lettau and van Nieuwerburgh (2008), Henkel et al. (2011), Dangl and Halling (2012), Pettenuzzo et al. (2014), and Zhu (2015)).

In the CV-DC model, the excess return, r_{t+1}, is related to the lagged predictor x_t:

$$r_{t+1} = \alpha + \beta_0 x_t + \beta_{t+1} x_t + \varepsilon_{t+1}, \tag{11}$$

$$x_{t+1} = \beta_x x_t + \varepsilon_{t+1}, \tag{12}$$

$$\beta_{t+1} = b\beta_t + \eta_{t+1}, \tag{13}$$

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

where $\beta_x < 1$, $\text{cov}(\varepsilon_{t+1}, \varepsilon_{t+1}) = \Omega = \sigma_{\varepsilon\varepsilon}$, and $\text{var}(\eta_{t+1}) = \sigma_\eta^2$. The CV-DC model is also a state-space model: Equations (11) and (12) are the observation equations, and equation (13) is a state equation that describes the dynamics of time-varying coefficients and captures the relationship between the predictor and the expected return.

III. Estimation and Specifications

A. Estimation Method

Pástor and Stambaugh (2009) present details of their Markov-chain Monte Carlo (MCMC) algorithm used to estimate the predictive system. To summarize, in each step of Gibbs sampling, they first draw the model parameters $\Theta = (E_x, E_r, \gamma, \beta, \Omega)$ conditional on the current draw of $\{\mu_t\}$. Then, the technique of forward filtering and backward sampling proposed by Carter and Kohn (1994) is used to draw the time series of $\{\mu_t\}$ conditional on the current draw of Θ. We develop a similar Bayesian approach for the estimation of the CV-DC model. Specifically, we employ the Gibbs sampling technique to obtain the posterior distributions for the model parameters.

In our context, we need to impose the NEP condition on the predictive system and the CV-DC model. A popular method uses an accept–reject algorithm when imposing the NEP condition on predictive models (e.g., Pettenuzzo et al. (2014), Avramov et al. (2018)). For each parameter vector generated by the Gibbs sampler, the NEP condition is evaluated at each point in time; the parameter vector is included in the sample if the NEP condition is not violated, and is rejected otherwise. A disadvantage of this accept–reject algorithm is that it is not efficient (e.g., O'Donnell and Coelli (2005)). In addition, it may be more restrictive than necessary. As an alternative, a few studies propose to use the Metropolis–Hastings algorithm to impose restrictions on a model (O'Donnell et al. (2001), O'Donnell and Coelli (2005)). We adopt a similar strategy and use the Metropolis–Hastings algorithm to estimate the predictive system and the CV-DC model under the NEP condition.

The Gibbs sampling method and the Metropolis–Hasting algorithm that we use in the article are described in Section 1 of the Supplementary Material.

B. Data

The empirical analysis in this article uses postwar data. The stock market index is the end-of-month value of the Standard & Poor's (S&P) 500 index from the Center for Research in Security Prices (CRSP). Returns are continuously compounded, including dividends. In calculating excess returns, we use the 30-day Treasury bill rate, which also comes from CRSP.

We use the dividend yield as the predictor for stock returns.[6] Dividends are defined as 12-month moving sums of dividends paid by firms belonging to the S&P

[6]Pástor and Stambaugh (2009), (2012) consider different predictors at annual and quarterly frequencies but mainly focus on the dividend yield because it exhibits the strongest relation to the expected return. Johannes et al. (2014) only consider the dividend yield as a predictor.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

500 index. In so doing, we circumvent the seasonality pattern in the dividend distribution. Finally, the dividend yield is defined as the value of dividends divided by the value of the index at the end of the month.[7]

The data set consists of monthly observations spanning 751 months from June 1952 to Dec. 2014. The choice of monthly data is motivated by prior literature, which employs monthly data to examine stock return predictability and optimal portfolio choices (e.g., Kandel and Stambaugh (1996), Stambaugh (1999), Barberis (2000), and Johannes et al. (2014)).

Table 1 reports the summary statistics of the excess stock market return and the dividend yield. The table reveals the low persistence in the excess stock return and the high persistence in the dividend yield.

C. Model Specification Test

The first question that we address is whether the NEP condition should be imposed on the predictive system and the CV-DC model. To answer this question, we conduct a comparison of the models with and without the NEP condition. Specifically, we employ the Bayes factor, a widely used statistic in the Bayesian literature. Kass and Raftery (1995) present a detailed discussion of the Bayes factor. Consider two models, M_1 and M_2, conditional on observed data Y_T and prior odds $p(M_1)/p(M_2)$. The Bayes factor compares these models based on the posterior odds:

$$\frac{p(M_1|Y_T)}{p(M_2|Y_T)} = \frac{p(Y_T|M_1)}{p(Y_T|M_2)} \times \frac{p(M_1)}{p(M_2)}. \tag{14}$$

A typical input for prior odds is $p(M_1) = p(M_2) = 0.5$. In this setting, the Bayes factor provides a summary of the evidence given by the data in favor of one model as opposed to another. An advantage of the Bayes factor is that it includes a

TABLE 1

Summary Statistics of Excess Stock Returns and Dividend Yields

Table 1 reports summary statistics for excess stock returns and log dividend yields using monthly data. *AR(1)* and *AR(12)* denote the 1st- and 12th-order autocorrelation of the process. *Cross-correlation* denotes the contemporaneous correlation between the excess stock return and the log dividend yield. The sample period is June 1952 and Dec. 2014.

	Excess Stock Return	Log Dividend Yield
Mean	0.0058	0.0317
Standard deviation	0.0420	0.0119
Maximum	0.1619	0.0637
Minimum	−0.2209	0.0108
AR(1)	0.0503	0.9882
AR(12)	0.0491	0.8448
Cross-correlation	−0.0544	

[7]Our finding of a lower predictive variance in the long run also holds for other predictors, such as the term spread, the cay factor (Lettau and Ludvigson (2001)), and the bond yield.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

penalty for incorporating too great a model structure.[8] Let M_1 denote the specification with the NEP condition and M_2 the specification without the condition. The empirical test leads to a Bayes factor equal to 7.6 for the predictive system and 13.5 for the CV-DC model. These results provide substantial to strong support for the specifications with the NEP condition.

D. Estimation Results

A likely property of the unexpected return is its negative correlation with the innovation in expected return. In the estimation of the predictive system, Pástor and Stambaugh (2009), (2012) use a negative prior. To gain some intuition on this correlation, we run the linear regression

$$r_{t+1} = a + bx_t + e_{t+1} \tag{15}$$

to predict the expected return based on the dividend yield. The estimated slope coefficient $\widehat{b}$ is equal to 0.21, with a t-value of 1.68 and an adjusted R^2 of 0.3%. Let $\widehat{v}_t$ denote the innovation in the dividend yield from an AR(1) regression. The empirical correlation between the unexpected return and the innovation in expected return is $\text{corr}(\widehat{e}_t, \widehat{v}_t) = -0.92$. This estimate is consistent with prior literature (e.g., Campbell (1991), Campbell and Ammer (1993), Pástor and Stambaugh (2009), and van Binsbergen and Koijen (2010)).

Table 2 presents the parameter estimates of the unrestricted and restricted versions of the predictive system (Panel A). Figure 1 plots the posterior densities of the parameters of the unrestricted predictive system (dotted lines). Graph A displays the posterior distribution of E_r. We estimate a mean value equal to 0.43%, with a 90% confidence interval of [–0.0003; 0.0082]. Graph B displays the posterior density of the autoregressive coefficient of the expected returns, β. The mean value

TABLE 2

Parameter Estimates of the Predictive Models

Table 2 reports parameter estimates for the unrestricted and restricted predictive system (Panel A) and the constant variance–drifting coefficient (CV-DC) model (Panel B), respectively. *, **, and *** denote statistical significance at the 10%, 5%, and 1% confidence levels, respectively. The sample period is June 1952–Dec. 2014.

Panel A. Predictive System

Parameters	Unrestricted	Restricted
E_r	0.43%***	0.89%***
β	0.88***	0.74***
γ	0.98***	0.96***
E_x	0.03%*	0.02%
σ_u	4.01%***	3.74%***
σ_v	1.37%***	1.41%***
σ_w	0.87%***	0.82%***
ρ_{uv}	0.43***	0.62***
ρ_{uw}	−0.19**	−0.21**
ρ_{vw}	0.62***	0.84***

Panel B. CV-DC Model

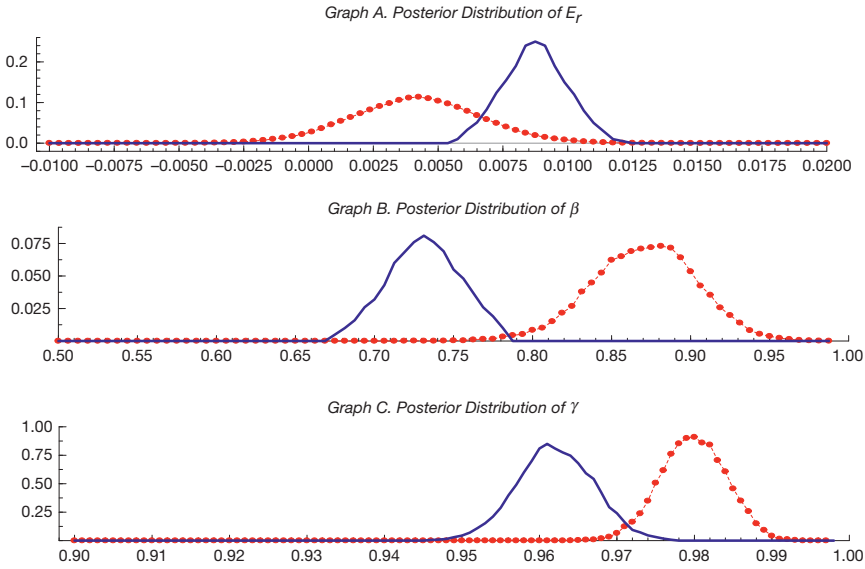
Parameters	Unrestricted	Restricted
α	0.10%*	0.12%*
β_0	0.24**	0.26**
β_x	0.96***	0.95***
b	0.78***	0.91***
σ_ε	3.85%***	3.72%***
σ_ξ	1.28%***	1.22%***
σ_η	2.05%***	1.76%***
$\rho_{\varepsilon\xi}$	0.17*	0.15*

[8] A typical interpretation of the Bayes factor in empirical analysis is as follows (Jeffreys (1961)): 1–3.2 is barely worth mentioning, 3.2–10 is substantial, 10–100 is strong, and greater than 100 is decisive.

FIGURE 1

Posterior Distribution of Coefficients in the Predictive System

Figure 1 plots the posterior distributions of the key coefficients in the unrestricted predictive system (light gray dotted lines) and in the predictive system with the nonnegative equity premium (NEP) condition (dark gray solid lines). Graphs A, B, and C plot the posterior distributions of E_r, β, and γ in equations (2) and (3), respectively. The predictor is the dividend yield. The sample period for estimating the predictive models is June 1952–Dec. 2014.



is equal to 0.89, suggesting that expected returns are somewhat persistent, consistent with Pástor and Stambaugh (2009) and van Binsbergen and Koijen (2010). The persistence of the dividend yield, γ (Graph C), has a mean value equal to 0.98, with a 90% confidence interval of [0:972; 0:988].

The solid lines in Figure 1 plot the posterior densities of the parameters of the predictive system when the NEP condition is imposed. Compared with the unrestricted estimation, the restricted distribution of E_r shifts substantially to the right, with a mean value equal to 0.89% and a confidence interval of [0:0062; 0:0107]. If we look at the posterior distribution of the slope coefficient β, the mean value becomes 0.74, indicating that expected returns are less persistent in the restricted predictive system. The persistence of the dividend yield has a mean value equal to 0.96. In sum, these results demonstrate that the NEP condition significantly alters the parameter estimates, suggesting that the NEP condition is binding.

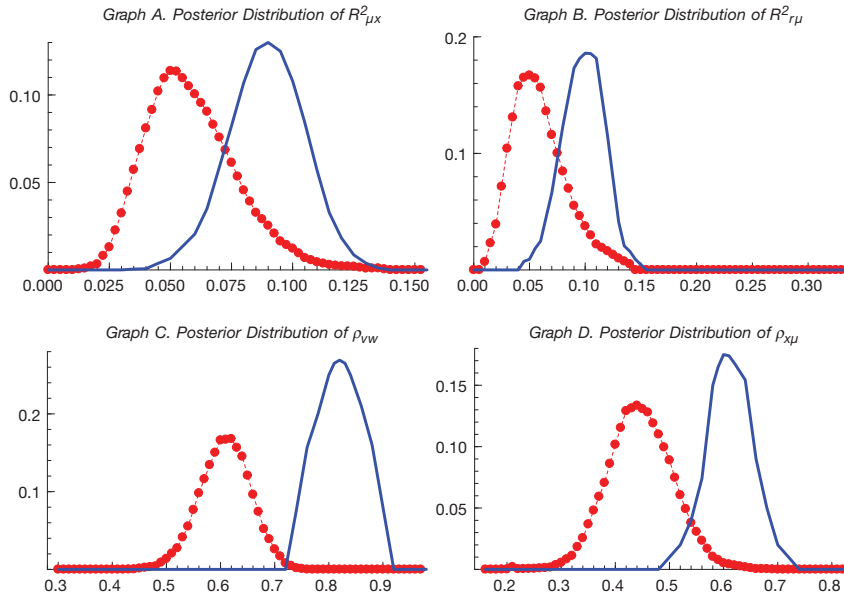
To provide insights on the relative performance of the unrestricted and restricted predictive models in predicting stock returns, we analyze several statistics. First, we compute the R^2 of the regression of the expected return μ_t on the dividend yield x_t (denoted by $R^2_{\mu x}$), which measures how much of the variance of the expected return is explained by the dividend yield. We also compute the predictive R^2 of the regression of the return r_{t+1} on the expected return μ_t (denoted by $R^2_{r\mu}$), which measures how much of the variance of the future return is explained by the current expected return. Figure 2 displays the posterior distribution of $R^2_{\mu x}$ and $R^2_{r\mu}$ (Graphs A and B, respectively) obtained with the unrestricted and restricted models.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 2

Posterior Distribution of Additional Statistics

Graph A of Figure 2 plots the posterior distribution of the fraction of the variation in the expected return μ_t that is explained by the predictor x_t. Graph B plots the posterior distribution of the predictive R^2 from the regression of returns r_{t+1} on expected returns μ_t. Panel C plots the posterior distribution of the conditional correlation ρ_{vw} between x_t and μ_t. Panel D plots the posterior distribution of the unconditional correlation $\rho_{x\mu}$ between x_t and μ_t. In each panel, the light gray dotted lines and dark gray solid lines represent the unrestricted predictive system and the predictive system with the nonnegative equity premium (NEP) condition, respectively. The predictor is the dividend yield. The sample period for estimating the predictive models is June 1952–Dec. 2014.



As the figure shows, imposing the NEP condition tends to increase both posterior R^2 statistics compared with the unrestricted case. This result indicates that the restricted model performs better in capturing the dynamics of the expected return with the dividend yield ($R^2_{\mu x}$) and in predicting the future stock return ($R^2_{r\mu}$).

We also report, in Graphs C and D of Figure 2, the posterior distributions of the conditional and unconditional correlations between the dividend yield x_t and the expected return μ_t: ρ_{vw} corresponds to the correlation between the dividend innovation v_t and the innovation in expected return w_t; $\rho_{x\mu}$ corresponds to the correlation between the dividend x_t and the expected return μ_t. We find that the NEP condition makes the dividend yield more highly correlated with the expected return, both conditionally and unconditionally. This result indicates that imposing the NEP condition improves the predictive ability of the dividend yield.

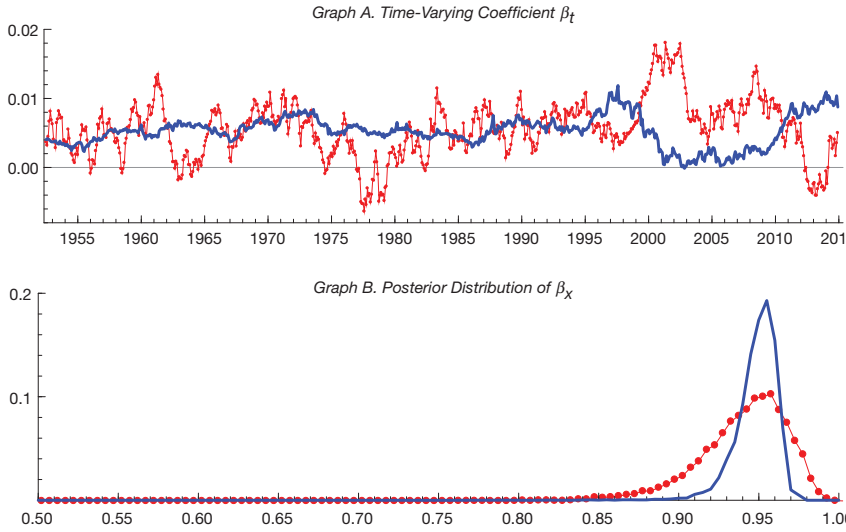
Turning to the CV-DC model, Panel B of Table 2 presents the parameter estimates of the unrestricted and restricted versions of the model. The time-varying predictability coefficient β_{t+1} in equation (11) is found to be persistent ($b = 0.78$). The table also reveals that the dividend yield process is found to be highly persistent ($\beta_x = 0.96$). To evaluate the impact of the NEP condition on the parameter estimates, we compare the evolution of the predictability coefficient β_{t+1} in the unrestricted and restricted models. Graph A of Figure 3 plots the evolution of the posterior

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 3

Predictability Coefficient Estimates in the CV-DC Model

Graph A of Figure 3 plots the time series of the posterior mean of the predictability coefficient for the unrestricted constant variance–drifting coefficient (CV-DC) model (dotted line) and the CV-DC model with the nonnegative equity premium (NEP) condition (solid line). It corresponds to the loading on the predictor ($\beta_0 + \beta_t$) in equation (11). Graph B plots the posterior distribution of β_x for the unrestricted CV-DC model (dotted line) and the CV-DC model with the NEP condition (solid line). The predictor is the dividend yield. The sample period for estimating the predictive models is June 1952–Dec. 2014.



mean of β_t. The difference is substantial between the parameter estimates obtained with the specifications with and without the NEP condition. Graph B presents the posterior distributions of β_x under both specifications. Imposing the NEP condition produces a narrower distribution of β_x and therefore helps reduce the uncertainty in the CV-DC model. To summarize, our empirical evidence clearly suggests that the NEP condition is binding.

To gain some intuition on how the NEP condition alters stock return forecasts, Figure 4 plots the time series of the equity premium obtained with the predictive models. Graph A displays the premium estimated with the (unrestricted and restricted) predictive system, and Graph B displays the premium estimated with the (unrestricted and restricted) CV-DC model. The figure clearly illustrates how imposing this condition affects the equity premium: For both models, the NEP condition is binding during some stressful episodes, in particular in the post–Internet bubble recession (2000–2003) and during the financial crisis (2007–2010). We also find that the correlation with the realized return is higher for the premium obtained with the restricted model. Indeed, the correlation between the equity premium and the realized return is equal to 0.22 for the unrestricted predictive system and 0.35 for the restricted predictive system. Similarly, the correlation between the equity premium and the realized return is equal to 0.16 for the unrestricted CV-DC model and 0.25 for the restricted CV-DC model.

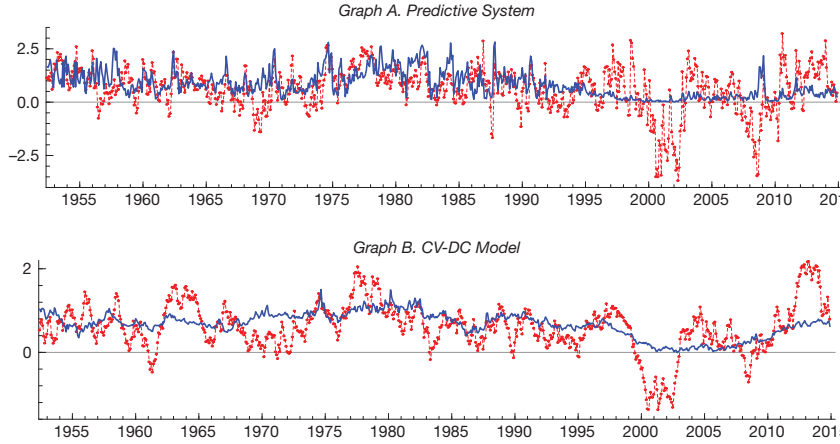
Even though the Bayes-factor analysis reveals that the NEP condition is an informative structure of predictive models, a more straightforward way to

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 4

Equity-Premium Forecasts

Figure 4 plots the time series of the 1-month-ahead equity-premium forecasts implied by the unrestricted model (light gray dotted line) and the model with the nonnegative equity premium (NEP) condition (dark gray solid line). Graph A corresponds to the predictive system, and Graph B corresponds to the constant variance–drifting coefficient (CV-DC) model. The sample period for estimating the predictive models is June 1952–Dec. 2014.



demonstrate the importance of this condition is to evaluate its ability to predict the future variance. To this end, we now proceed in two steps. First, we demonstrate that with the NEP condition, the predictive variance of stock returns is lower in the long run than it is in the short run. Second, we evaluate the economic value of the NEP condition and show that it delivers systematic value to a buy-and-hold investor.

IV. The Term Structure of Predictive Variance

This section investigates the slope of the term structure of predictive variance, that is, whether stocks are less volatile in the long run when investors expect an NEP. It is noteworthy that in the linear predictive regression of equation (15), predictive variance is constant over time horizons. Therefore, imposing the NEP condition on this predictive regression still generates a flat curve of predictive variance. In contrast, for predictive models, Pástor and Stambaugh (2012) find that stocks are actually more volatile in the long run, when the imperfection of predictors and the estimation risk are taken into account. Johannes et al. (2014) document a similar pattern based on the CV-DC model. Along this line of literature, we evaluate each of the 5 components of the multiperiod predictive variance in equation (10) conditional on the information available at time t. In our empirical setting, the information set consists of the complete history of returns and dividend yields.[9]

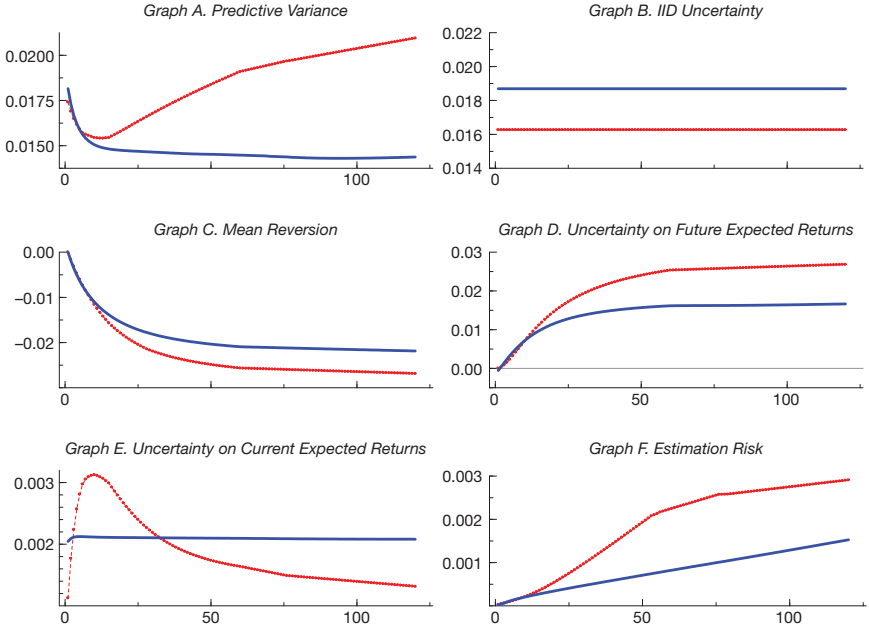
[9] In this analysis, the NEP condition affects the term structure of predictive variance through its effect on the estimation of the model's parameters. Therefore, for a given set of parameters, the term structure is essentially unaffected by the information set Z_t and does not depend on time. At the end of this section, we consider the case of a recursive estimation of the models and illustrate how recursively estimating model parameters affects the slope of the term structure over time.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 5

Term Structure of Annualized Predictive Variance and Its Components: Predictive System

Graph A of Figure 5 plots the annualized predictive variance of the distribution of multiperiod returns for horizons between 1 and 120 months. Graphs B–F plot the 5 components of the predictive variance: the independent and identically distributed (IID) uncertainty, the mean reversion, the uncertainty on future expected returns, the uncertainty on current expected returns, and the estimation risk, respectively. The light gray dotted (dark gray solid) lines are associated with the unrestricted (restricted) predictive system.



We start with the predictive system. Figure 5 displays the annualized predictive variance and its 5 components for the horizons from $k = 1$ to 120 months. Graph A reveals that the unrestricted predictive system implies an increasing term structure of predictive variance, confirming the robustness of the findings of Pástor and Stambaugh (2012). However, when the NEP condition is imposed, we obtain a decreasing term structure of predictive variance.

To identify the drivers of this twist in the term structure, Graphs B–F of Figure 5 plot the predictive variance components attributable to the IID uncertainty, mean reversion, uncertainty about future expected returns, uncertainty about current expected returns, and estimation risk, respectively. The twist of the term structure of predictive variance is due mainly to the decrease in the uncertainty about future expected returns (Graph D). This uncertainty increases with the horizon for both the unrestricted and restricted predictive systems. However, the increase is almost halved when the NEP condition is imposed, resulting in a much flatter curve. The decrease in the estimation risk for long horizons also contributes to a lower predictive variance in the long run (Graph F). In contrast, the contribution of the mean reversion is slightly increased when the NEP is imposed, which partly compensates for these effects, but it is not sufficient to reverse the decrease in the predictive variance relative to the unrestricted model.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

The NEP condition also implies small increases in the contribution of the IID uncertainty and the uncertainty about current expected returns (Graphs B and E of Figure 5, respectively). However, these effects are stable over investment horizons or are small in magnitude and do not significantly alter the shape of the term structure of predictive variance.

Figure 6 plots the term structure of predictive variance and its components for the CV-DC model with and without the NEP condition. We achieve similar conclusions. The NEP condition substantially reduces the predictive-variance component attributable to the uncertainty about future expected returns and the estimation risk, whereas the other components have essentially no impact on the term structure of predictive variance.

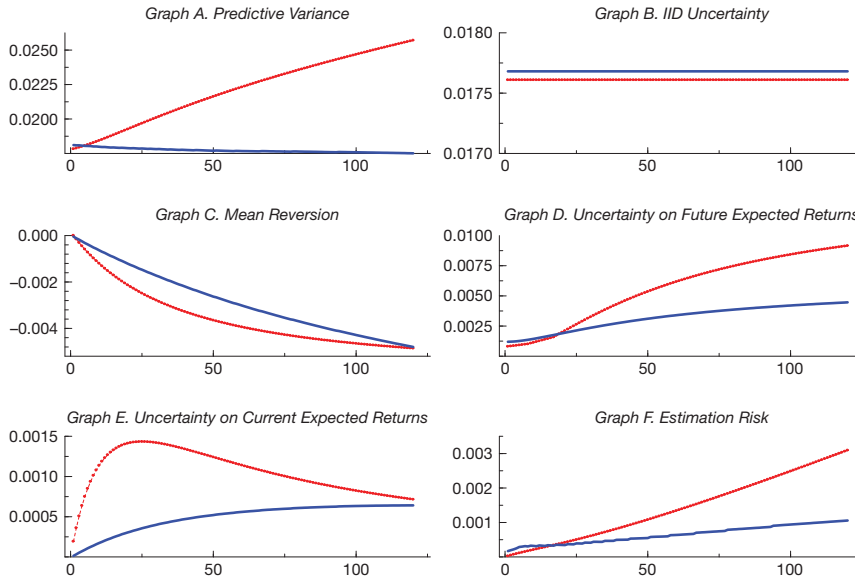
Overall, our findings suggest that from a forward-looking perspective, stocks are less volatile in the long run when investors expect an NEP. Our results demonstrate that the decrease in the term structure of predictive variance is driven mostly by the reduction in the uncertainty about future expected returns and the estimation risk. We now consider the driving forces of this reduction.

A first force is that the NEP condition sharpens the return forecasts and reduces the uncertainty about future expected returns (e.g., Campbell and Thompson (2008), Pettenuzzo et al. (2014)). To illustrate this point, Figure 7 plots the probability of a negative equity premium at each forecast horizon obtained with the unrestricted predictive system and the CV-DC model. As the figure reveals, the

FIGURE 6

Term Structure of Annualized Predictive Variance and Its Components: CV-DC Model

Graph A of Figure 6 plots the annualized predictive variance of the distribution of multiperiod returns for horizons between 1 and 120 months. Graphs B–F plot the 5 components of the predictive variance: the independent and identically distributed (IID) uncertainty, the mean reversion, the uncertainty on future expected returns, the uncertainty on current expected returns, and the estimation risk, respectively. The light gray dotted (dark gray solid) lines are associated with the unrestricted (restricted) constant variance–drifting coefficient (CV-DC) model.

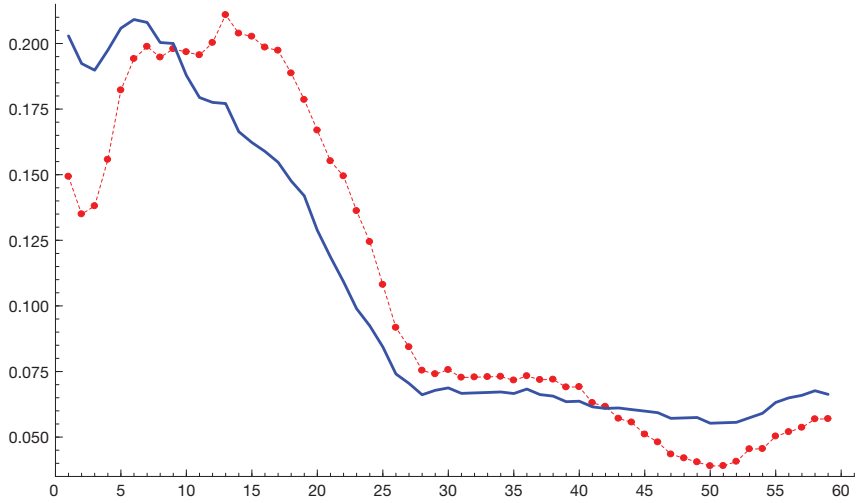


https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 7

Probability of a Negative Equity Premium

Figure 7 plots the probability that the equity premium is negative at each forecast horizon for investors who use the unrestricted predictive system (dark gray solid line) and the unrestricted constant variance–drifting coefficient (CV-DC) model (light gray dotted line) to predict stock returns. Forecast horizons cover the range from 1 to 60 months.



probability of a negative equity premium is as high as 20% for short horizons. This probability is still substantial, although lower (close to 5%), at a 5-year forecasting horizon, suggesting that imposing the NEP condition has the potential to significantly reduce the uncertainty about future stock returns.

Second, the NEP condition constitutes an identification condition that provides additional information that is not driven by the data to estimate the parameters of the predictive models (e.g., Pástor (2000), Avramov et al. (2018)). It thus reduces the estimation risk for a given horizon relative to the unrestricted model. It is important to note that the uncertainty on the future expected return and the estimation risk still increase with the horizon when the NEP condition is imposed. However, these increases are much less pronounced and no longer compensate for the decreasing effects of the mean reversion on the predictive variance. Therefore, with the restricted predictive system, mean reversion dominates, and predictive variance is lower in the long run than it is in the short run.[10]

As illustrated in Figure 4, the NEP condition is generally binding during the recent crisis episodes (2000–2003 and 2007–2010). For this reason, we expect the

[10]Precisely measuring what components contribute the most to the downward slope of the volatility term structure in the restricted model (compared with the upward-sloping curve in the unrestricted model) is a difficult task because the effects are highly nonlinear and interdependent. We used an interval method to measure how much of the difference between predictive variance in the unrestricted and restricted predictive systems is due to the difference in parameter estimates and to the lower standard errors of the estimates in the restricted model (due to the NEP condition). We find that roughly 38%–45% of the predictive variance difference is due to the difference in the standard errors of the parameter estimates for the predictive system (29%–35% in the CV-DC model). In Section 2 of the Supplementary Material, we provide details on the implementation of the interval method.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

term structure of predictive variance to be more affected by the NEP condition in the recent period. To investigate this point, we recursively estimate the predictive models with the NEP condition, using a 10-year rolling window, and compute the evolution of the term structure of predictive variance over time. In Figure 8, we display the 1-month-ahead and 10-year-ahead predictive variances for the restricted predictive system and the CV-DC model, respectively. These horizons correspond to the first and last points of the term structure in Graph A of Figures 5 and 6.

Figure 8 reveals that in the first part of the sample, stocks are less volatile in the long run than they are in the short run but with a relatively limited difference. In contrast, from 2000 on, when the NEP condition is more binding, we observe a substantial increase in the short-run predictive variance relative to the long-run predictive variance. This pattern is driven by the following forces: The IID uncertainty increases in the recent period, so that the predictive variance increases at all horizons, and the mean reversion is more pronounced, so that the long-run variance eventually exhibits a lower level.

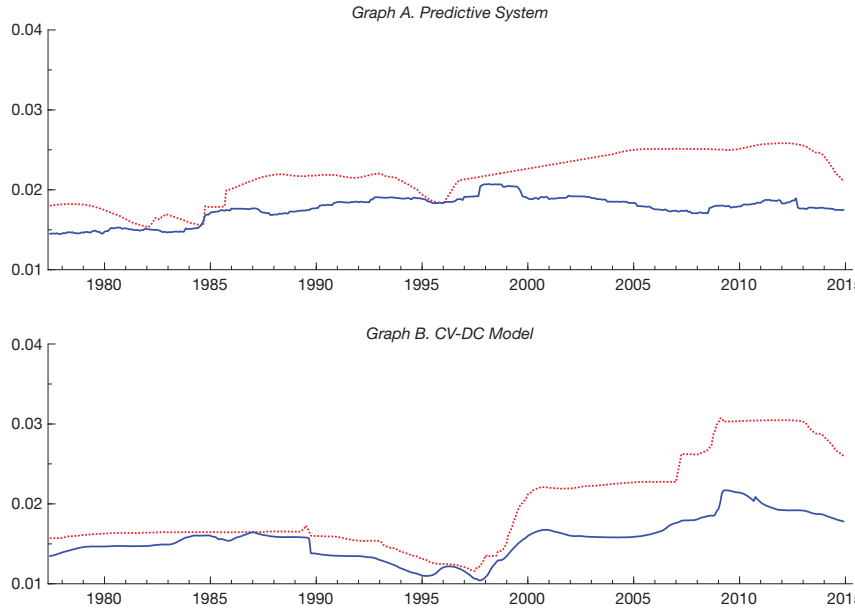
V. The Importance of the NEP Condition

In this section, we present the framework for exploring how the NEP condition affects the optimal portfolio choice. This framework is based on that of Barberis (2000) and Kandel and Stambaugh (1996). After the presentation of the framework

FIGURE 8

1-Month-Ahead and 10-Year-Ahead Predictive Variance

Figure 8 plots the temporal evolution of the 1-month-ahead and 10-year-ahead predictive variances, using a 10-year recursive estimation of the restricted predictive system (Graph A) and the restricted constant variance–drifting coefficient (CV-DC) model (Graph B), respectively. Each plot displays the 1-month predictive variance (light gray dotted line) and the 10-year predictive variance (dark gray solid line).



https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

and the discussion of the optimal weights, we compare the performance of the portfolios allocated using stock return forecasts based on the predictive model with and without imposing the NEP condition.

A. The Framework

We consider a risk-averse investor who allocates funds to a riskless Treasury bill and a risky stock market index for a k-period investment horizon. The investor believes that her portfolio decision has no effect on the probability distribution of asset returns. Let $W_t = 1$ denote the investor's current wealth and ω_t the proportion of the wealth invested in the stock market index at time t. Assuming a buy-and-hold strategy, the investor's wealth evolves from t to $t+k$ as follows:

$$\begin{aligned} W_{t+k} &= W_t \prod_{i=1}^{k} \left[(1-\omega_t) \exp\left(r^f_{t+i}\right) + \omega_t \exp\left(r^f_{t+i} + r_{t+i}\right) \right] \\ &= W_t \left[(1-\omega_t) \exp\left(R^f_{t:t+k}\right) + \omega_t \exp\left(R^f_{t:t+k} + R_{t:t+k}\right) \right], \end{aligned} \tag{16}$$

where r^f_{t+i} is the risk-free interest rate for the period between time $t+i-1$ and $t+i$, and $R^f_{t:t+k} = \sum_{i=1}^{k} r^f_{t+i}$ is the cumulative risk-free rate between t and $t+k$. Finally, the investor's preference over her terminal wealth is described by a power utility function with constant relative risk-aversion parameter A:

$$v(W_{t+k}) = \frac{W^{1-A}_{t+k}}{1-A}. \tag{17}$$

The investor chooses ω_t to maximize the expected utility of the end-of-horizon wealth, so the buy-and-hold investor's problem is to solve the following:

$$\max_{\{\omega\}} \; E_t(v(W_{t+k})), \tag{18}$$

where E_t denotes the fact that the allocation decision is made based on the information set Z_t, which contains the history of the stock return and the dividend yield up to time t.

B. Optimal Portfolio Weights

A difficulty in computing the expected utility in equation (18) is that parameters Θ are not known with certainty. A Bayesian investor takes this uncertainty into account through the posterior distribution $p(\Theta|Z_t)$, given the observations Z_t. By integrating over the posterior distribution, we obtain the predictive distribution $p(R_{t:t+k}|Z_t) = \int p(R_{t:t+k}|Z_t,\Theta)p(\Theta|Z_t)d\Theta$, which we can use to solve the investor's problem:

$$\max_{\{\omega\}} \int v(W_{t+k}) p(R_{t:t+k}|Z_t) dR_{t:t+k}. \tag{19}$$

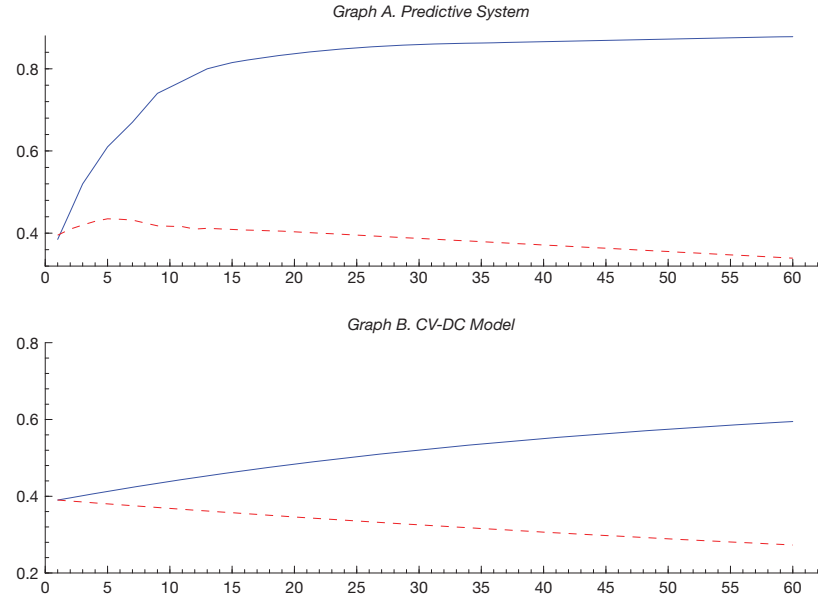
We solve the problem in equation (19) using numerical simulations, as proposed by Barberis (2000) and Brandt et al. (2005). We calculate the integral in the

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 9

Optimal Allocation of a Buy-and-Hold Investor

Figure 9 plots the optimal allocation to stocks of a buy-and-hold investor for investment horizons ranging from 1 to 60 months. The two lines on each graph correspond to the unrestricted (light gray dotted line) and restricted (dark gray solid line) predictive models that the investor would use to predict the equity premium. The risk-aversion parameter is $A = 4$.



problem in equation (19) for values of ω between ω_L and ω_U and report the weight ω that maximizes the expected utility.[11] Following Barberis (2000), we simulate S = 1,000,000 samples of Θ from the posterior distribution $p(\Theta|Z_t)$. For each simulated draw Θ^s, we do ensure that the model satisfies the NEP condition. Then, for each Θ^s, we simulate one draw from the return distribution $p(R_{t:t+k}|Z_t, \Theta^s)$. This process gives a sample of 1,000,000 draws from the predictive distribution $p(R_{t:t+k}|Z_t)$. We then solve the problem in equation (19) for ω by replacing the integral by the sum $(1/S)\sum_{s=1}^{S} v(W_{t+k})p(R_{t:t+k}|Z_t, \Theta^s)dR_{t:t+k}$.

Figure 9 illustrates the optimal allocation of the buy-and-hold investor. We set the risk-aversion parameter to $A = 4$ in our empirical analysis. Within each graph, each line shows the percentage allocated to stocks for investment horizons ranging from 1 to 60 months. The two lines in each panel correspond to the allocation obtained with the unrestricted and restricted predictive models used by the investor to predict the equity premium. With the unrestricted versions of the predictive system and the CV-DC model, the allocation to stocks decreases over the horizon. With a horizon increasing from 1 month to 5 years, the allocation to stocks would decrease from 40% to 35% with the predictive system and from 40% to 28% with the CV-DC model. The reason for this decrease is that the predictive variance of stock returns under these models rises quickly as the time horizon increases, which

[11] Throughout the article, we restrict the allocation to the interval $-2 \leq \omega \leq 3$, as in Johannes et al. (2014).

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

leads to a lower allocation to risky stocks. In contrast, once the NEP condition is imposed on predictive regressions, predictive variance decreases with the horizon, and the optimal allocation to stocks is higher for a long-horizon investor than it is for a short-horizon investor. It increases from 40% to 84% with the predictive system and from 40% to 65% with the CV-DC model, respectively.

Another implication of these patterns for predictive variance is that the investor using the restricted model invests substantially more in stocks than the investor using the unrestricted version. For the 2-year investment horizon, the allocation to stocks is equal to 82% and 40% for an investor using the restricted and unrestricted versions of the predictive system, respectively. The allocation is equal to 50% and 35% for the investors using the restricted and unrestricted versions of the CV-DC model.

C. Economic Values

Given that statistical return predictability does not necessarily imply economic value (e.g., Leitch and Tanner (1991), Della Corte, Sarno, and Thornton (2012)), we assess the economic value of the NEP condition by comparing the utility gains accrued to an investor who uses the predictive models with the NEP condition to choose the optimal portfolio relative to the unrestricted predictive models. When the investment horizon is longer than 1 year, our setting allows investors to rebalance their portfolios every year, as in Barberis (2000). Following Johannes et al. (2014), we bound portfolio weights at $\omega_L = -2$ and $\omega_U = 3$, although the portfolio weights for the predictive system and the CV-DC model rarely hit the lower and upper bounds. Our empirical analysis considers investment horizons k of 1 and 2 years ($k = 12$ and 24).

In each period, the investor determines the optimal weight that maximizes the expected utility. The investor holds the portfolio for some period of time, updates posterior distributions, and chooses portfolio weights. We measure the economic value of imposing the NEP condition using the Sharpe ratio (SR) and the certainty equivalent return (CER), which are defined as follows: For an allocation in month t for k months, the ex post portfolio excess return is calculated at the end of month $t+k$ as $R_{p,t:t+k} = (1-\omega_t)\exp(R^f_{t:t+k}) + \omega_t \exp(R^f_{t:t+k} + R_{t:t+k})$. Iterating this process for all dates from $t = 1$ to $(T-k)$, we obtain a time series of ex post excess returns for the optimal portfolio. Denoting the sample mean by $\overline{R}_p(k)$ and the sample variance of the portfolio return by $\sigma_p(k)^2$, the Sharpe ratio, measured as $\mathrm{SR}(k) = \overline{R}_p(k)/\sigma_p(k)$, is the risk-adjusted performance of the strategy. The CER is the risk-free return that an investor would consider equivalent to investing in the strategy. It is measured as follows:

$$1 + \mathrm{CER} = \left(\frac{1}{T} \sum_{t}^{T} \left(1 + R_{p,t}\right)^{1-A} \right)^{1/(1-A)}.$$

To test whether the SR of the strategy is equal to the SR of the strategy based on the historical mean of the market return (computed up to time t), denoted by SR_0, we follow the approach of DeMiguel, Garlappi, and Uppal (2009). We proceed in a similar way to test whether the CER of the strategy is equal to the CER of the

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

TABLE 3

Performance of the Optimal Portfolios

Table 3 summarizes the performance of the optimal portfolios when the dividend yield is used as the predictor for the 1- and 2-year investment horizons. Investors have power utility with risk-aversion parameter A and allocate their funds between a risky stock index and a riskless 1-period Treasury bill. Statistics are the annualized Sharpe ratio and the certainty equivalent return (CER) (i.e., the annualized risk-free return that gives the investor the same utility as the risky portfolio strategy). The models are the unrestricted and restricted predictive systems (Panel A) and the unrestricted and restricted constant variance–drifting coefficient (CV-DC) model (Panel B). The asymptotic distribution for the test of the null hypothesis that the Sharpe ratio of the model is equal to the Sharpe ratio based on the historical mean is given in DeMiguel et al. (2009). *, **, and *** denote significance at the 10%, 5%, and 1% confidence levels, respectively.

	Sharpe Ratio				CER			
	$A=4$		$A=6$		$A=4$		$A=6$	
	1 Year	2 Years	1 Year	2 Years	1 Year	2 Years	1 Year	2 Years
Panel A. Predictive System								
Unrestricted system	0.080	0.100*	0.107	0.125*	4.98*	4.39*	5.21**	5.27**
Restricted system	0.163**	0.173**	0.187**	0.201**	6.90**	6.45**	6.85**	6.51*
Panel B. CV-DC Model								
Unrestricted model	0.090*	0.083	0.073	0.076	3.96*	3.79	3.90	3.62
Restricted model	0.163**	0.142**	0.121*	0.128*	5.87**	6.37**	6.18*	6.40**

strategy based on the historical mean of the market excess return, denoted by CER_0.[12]

Table 3 reports the annualized SR and CER of portfolio returns for each model, risk aversion, and investment horizons. The results are based on monthly observations from 1952 to 2014, with the dividend yield as the predictor of U.S. stock returns. For a risk-aversion parameter equal to $A = 4$, the SR values of the unrestricted predictive system are equal to 0.08 and 0.1 at the 1-year and 2-year investment horizons, respectively. The gain relative to the strategy based on the historical mean is insignificantly different from 0. In contrast, the predictive system with the NEP condition delivers SR values equal to 0.16 and 0.17, respectively. In words, the SR values obtained with the restricted model are statistically higher than the value obtained with the historical mean and approximately twice as high as the values obtained with the unrestricted model. Similar results are obtained with the CV-DC model. Turning to the CER, we obtain similar findings: The CER values of the unrestricted predictive system are equal to 4.98% and 4.39% at the 1-year and 2-year investment horizons, respectively. With the NEP condition, the model delivers CER values equal to 6.9% and 6.45%, respectively. All these results confirm that the NEP condition plays an important role in optimal portfolio choice and delivers economic value to investors.

Figure 8 has revealed that the recent period, when the NEP condition is more binding, is associated with a high short-run predictive variance relative to the long-run predictive variance. From an allocation perspective, Figure 9 demonstrates that because stocks are relatively less risky in the long run, a long-horizon investor will typically invest more to stocks than a short-horizon investor will. We now illustrate in Figure 10 how the performance of investment strategies based on the restricted models is affected in the recent period. We plot the evolution of the

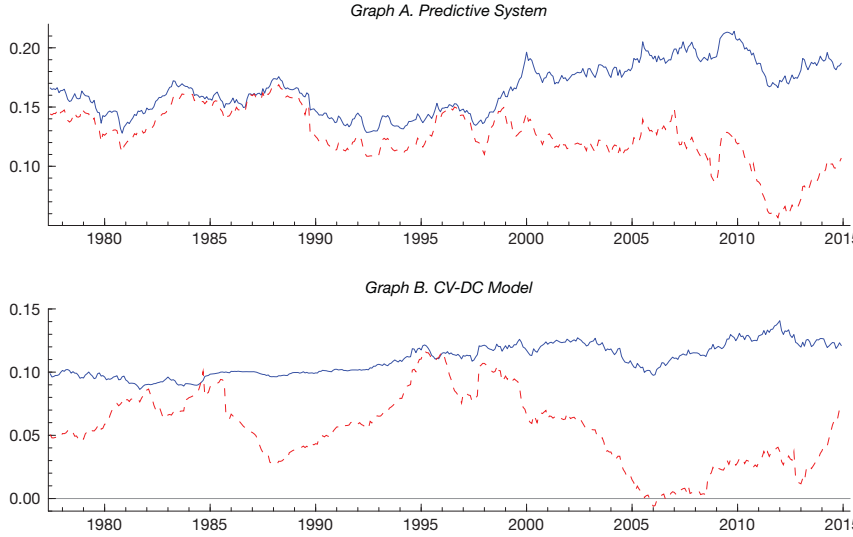
[12]To test the null hypothesis that $SR = SR_0$, we use the statistic given in footnote 16 of DeMiguel et al. (2009). To test the null hypothesis that $CE = CER_0$, we use the statistic for the test of equal CER given in their footnote 18.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 10

Sharpe Ratio for 1-Month and 10-Year Investment Horizons

Figure 10 plots the temporal evolution of the Sharpe-ratio values for the 1-month and 10-year investment horizons, using a 10-year recursive estimation of the restricted predictive system (Graph A) and the restricted constant variance–drifting coefficient (CV-DC) model (Graph B), respectively. Each plot displays the Sharpe ratio for the 1-month (light gray dotted line) and 10-year (dark gray solid line) investment horizons.



SR of the 1-month and 10-year investment strategies (computed using a 10-year recursive estimation) for the predictive system and the CV-DC model (Graphs A and B, respectively). The figure reveals that in the recent period, the SR of a long-horizon investor is substantially higher than the SR of a short-horizon investor because of the overweighting of stocks. These results suggest that periods with low long-run predictive variance relative to the short-run predictive variance are associated with better risk-adjusted performance of long-horizon investment strategies.

VI. Robustness Checks

In this section, we investigate the robustness of our main results by performing the same analysis with the alternative samples used by Pástor and Stambaugh (2012) and Johannes et al. (2014). First, as in Pástor and Stambaugh, we use a long sample of annual data starting in 1802 to estimate the unrestricted and reinitially compiled by Siegel (1992), (1994). The return is the annual real log return on the U.S. equity market, and the predictor is the dividend yield on U.S. equity. In our Bayesian empirical setting, the information set therefore consists of the 215-year history of returns and dividend yields.

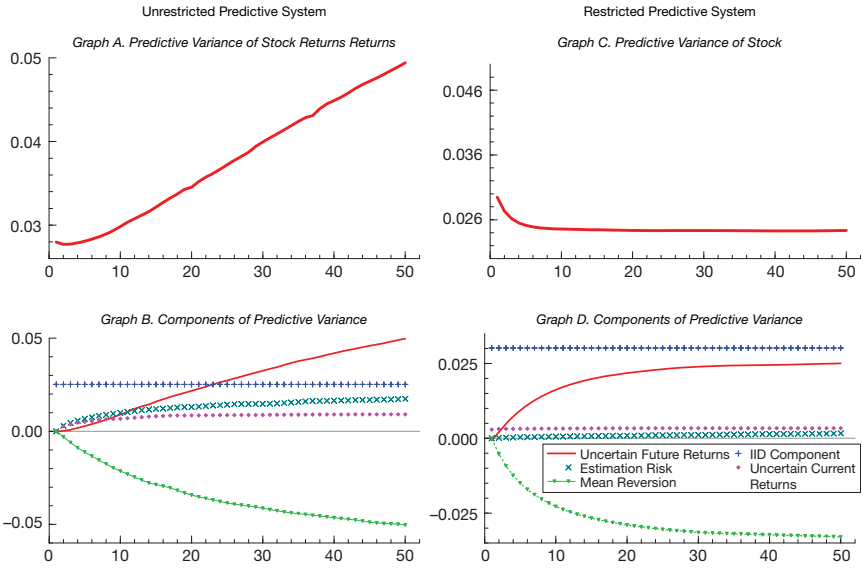
Figure 11 plots the predictive variance implied by the unrestricted predictive system and its 5 components for horizons from 1 to 50 years (Graphs A and B). The predictive variance increases substantially with the investment horizon: The annualized variance exceeds the 1-year variance by approximately 45% at a

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 11

Term Structure of Annualized Predictive Variance and Its Components: Predictive System

Graphs A and C (B and D) of Figure 11 plot the annualized predictive variance of the multiperiod returns for the unrestricted (restricted) predictive systems. Investment horizons range from 1 to 50 years. All quantities are divided by k, the number of periods in the return horizon. The results are obtained by estimating the predictive system on annual real U.S. stock market returns between 1802 and 2016.



30-year horizon and by 80% at a 50-year horizon. Our analysis confirms the findings of Pástor and Stambaugh (2012), which suggest that predictive variance is an increasing function of the investment horizon, after taking uncertainties into account.

We now evaluate the extent to which the NEP condition changes the shape of the term structure of predictive variance. For this purpose, we estimate the restricted predictive system based on the 215-year sample and calculate the term structure of predictive variance (Graphs C and D of Figure 11). It is evident from the figure that as opposed to the evidence from the unrestricted model, the annualized variance now decreases with the investment horizon: The predictive variance for a 30-year (50-year) horizon is only 19% (21%) of the 1-year variance. These results confirm, in a long sample, that stocks are less volatile in the long run when investors expect the equity premium to be nonnegative.

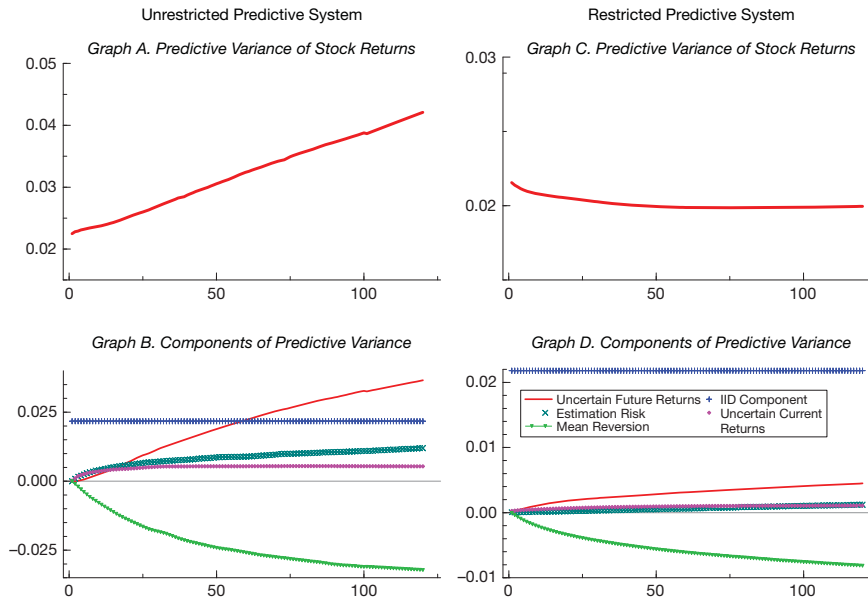
Second, we use the same alternative data set as in Johannes et al. (2014), covering monthly data starting in 1927, to estimate the CV-DC model. The data consist of monthly log-excess returns from the value-weighted New York Stock Exchange (NYSE)/American Stock Exchange (AMEX)/National Association of Securities Dealers Automated Quotations (NASDAQ) index (including dividends) minus the 1-month Treasury bill rate from Ibbotson and Associates, extended to 2016. Graphs A and B of Figure 12 display the predictive variance generated by the unrestricted model. Consistent with the evidence reported by Pástor and Stambaugh (2012) and Johannes et al., the annualized variance exceeds the 1-year

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

FIGURE 12

Term Structure of Annualized Predictive Variance and Its Components: CV-DC Model

Graphs A and C (B and D) of Figure 12 plot the annualized predictive variance of the multiperiod returns for the unrestricted (restricted) constant variance–drifting coefficient (CV-DC) models. Investment horizons range from 1 to 120 months. All quantities are divided by k, the number of periods in the return horizon. The results are obtained by estimating the predictive system on monthly U.S. stock market returns between 1927 and 2016.



variance by approximately 43% at a 60-month horizon and 87% at a 120-month horizon. However, the figure also reveals that when the NEP condition is imposed, the predictive variance decreases with the investment horizon, as predicted by the traditional view (e.g., Siegel (1994), Campbell and Viceira (2002), (2005)) (Graphs C and D). The predictive variance is reduced by approximately 10% for long horizons.

VII. Out-of-Sample Prediction in a Simulated Experiment

In this section, we conduct an out-of-sample exercise based on simulated data to show that the restricted model does a better job of predicting mean returns out of sample when the equity premium is indeed nonnegative. We use the long-run-risk model of Bansal and Yaron (2004) as the data-generating process to simulate an equity premium that is known to be nonnegative and highly persistent. In the long-run-risk model, the main cause of stock price variability relative to dividends is the predictable and persistent variation in consumption growth, which creates similar variation in dividend growth. This finding implies that stock returns are predictable by dividend yield. These two properties of the long-run-risk model make it a good data-generating process to evaluate the usefulness of the NEP condition by examining the out-of-sample forecasting performance of the unrestricted and restricted models. Specifically, using simulated data, we recursively estimate the unrestricted

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

and restricted predictive systems and compare their relative performance in predicting stock returns.[13] In Section 3 of the Supplementary Material, we provide a brief description of the long-run-risk model of Bansal and Yaron (2004) and describe how we calibrate and simulate the model. To save space, the Supplementary Material also reports summary statistics on the parameter estimates of the (unrestricted and restricted) predictive systems and CV-DC models based on these simulated data. We note that the parameter estimates are in the ballpark of the estimates reported for our data (Table 2). In addition, as expected, the unconditional equity premium increases when the NEP condition is imposed on the model.

We use these simulated data to investigate the out-of-sample forecasting performance of the unrestricted and restricted predictive models. Following Goyal and Welch (2008) and Rapach et al. (2010), we predict the market return using a sequence of expanding windows. For the first window, we use the first 376 observations (half of the sample). Then, for the sample ending in month $t = s_0, \cdots, T-1$, we recursively estimate the unrestricted and restricted predictive models to obtain a return forecast ($\widehat{r}_t$). The empirical benchmark is the historical average based on stock returns up to time t ($\overline{r}_t$). We use the out-of-sample R^2 statistic, R^2_{OOS}, suggested by Campbell and Thompson (2008), to compare the forecasts:

$$R^2_{\text{OOS}} = 1 - \frac{\sum_{t=s_0}^{T-1} (r_{t+1} - \widehat{r}_t)^2}{\sum_{t=s_0}^{T-1} (r_{t+1} - \overline{r}_t)^2}.$$

The R^2_{OOS} statistic measures the reduction in the mean-squared prediction error for the predictive models. As in Section V, we consider the allocation to the riskless asset and the risky stock market index, based on the return forecast (with a 1-period horizon). Additional measures of performance are, as before, the SR and the CER. We also consider the same levels of risk aversion ($A = 4$ and 6).

We repeat the out-of-sample exercise for 1,000 samples and report the average R^2_{OOS} statistic, the SR, and the CER for the restricted and unrestricted predictive systems and CV-DC models in Table 4. Whereas the unrestricted predictive system delivers an R^2_{OOS} statistic of 1.03%, the restricted predictive system generates an R^2_{OOS} statistic of 2.08%. Therefore, imposing the NEP condition on the predictive system and the CV-DC model helps predict returns when the data are drawn from the long-run-risk model. This result confirms that even out of sample, this condition plays a key role in prediction. This role is confirmed in the allocation experiment: The SR and CER values generated by the restricted model are always higher than the values generated by the unrestricted model. The SR values obtained with the restricted models are at least twice as large as the values obtained with the unrestricted models, the same order of magnitude that we find using actual data in Section V.C. Using the method suggested by DeMiguel et al. (2009), we test

[13] In sample, the restricted predictive system should generate better performance than the unrestricted system because it takes the NEP condition into account. Out of sample, there is a trade-off between the additional flexibility offered by the unrestricted model and the additional structure imposed by the NEP condition.

TABLE 4

Out-of-Sample Forecast Performance

Table 4 reports the out-of-sample forecasting performance of the restricted and unrestricted predictive system (Panels A and B) and the constant variance–drifting coefficient (CV-DC) model (Panels C and D). The reported statistics are the out-of-sample R^2 (R^2_{OOS}), the Sharpe ratio, and the certainty equivalent return (CER) (for risk-aversion parameters, $A=4$ and 6). The significance of the out-of-sample R^2 is based on the Clark and West (2007) mean-squared prediction error (MSPE)-adjusted test. The asymptotic distribution for the test of the null hypothesis that the Sharpe ratio of the model is equal to the Sharpe ratio based on the historical mean is given in DeMiguel et al. (2009). *, **, and *** denote significance at the 10%, 5%, and 1% confidence levels, respectively.

	R^2_{OOS}	Sharpe Ratio		CER	
		$A=4$	$A=6$	$A=4$	$A=6$
Panel A. Unrestricted Predictive System					
Full sample	1.03%**	0.20*	0.18*	3.17*	3.45*
Expansions	0.68%*	0.13	0.12	2.65	3.07
Recessions	1.57%***	0.37*	0.33*	4.48**	5.16**
Panel B. Restricted Predictive System					
Full sample	2.08%***	0.52*	0.47*	3.35**	3.68**
Expansions	1.18%**	0.28	0.29	2.66*	3.10*
Recessions	3.35%***	0.62*	0.56*	4.78***	5.35***
Panel C. Unrestricted CV-DC Model					
Full sample	0.87%**	0.18*	0.17*	2.89*	3.14*
Expansions	0.58%	0.11*	0.12*	2.33	2.57**
Recessions	1.37%***	0.31**	0.29**	4.15**	4.39**
Panel D. Restricted CV-DC Model					
Full sample	1.68%***	0.41*	0.38*	3.36*	3.75**
Expansions	1.26%**	0.29*	0.29*	2.69	3.17*
Recessions	2.55%***	0.57**	0.52**	4.85**	5.66**

whether the predictive models deliver the same SR values as the historical mean. Our test rejects the null hypothesis, implying that the restricted model generates a higher SR than its unrestricted counterpart. Turning to the restricted and unrestricted CV-DC models, we find a similar pattern. Overall, the out-of-sample forecasting exercise based on the simulated data confirms that the NEP condition fundamentally helps predict returns and generate higher risk-adjusted performance.

To better understand when the NEP condition is the most useful, we split the samples between expansion and recession periods. We define expansion (recession) periods as periods corresponding to positive (negative) consumption growth. We find that the predictive power is stronger during the recession period. During recessions, investors are more reluctant to take on risk. Because investors are more risk averse, the equity premium is higher, and the return predictability is stronger, consistent with Dangl and Halling (2012) and Zhu (2015), among others.

VIII. Conclusion

In this study, we impose the NEP condition on the predictive system and the CV-DC model to forecast stock returns. We find that this condition significantly increases the predictive power of predictive models. The Bayes-factor analysis further indicates that the NEP condition is a necessary structure of predictive models. The NEP condition also plays a fundamental role in the term structure of predictive variance. Indeed, we find that the term structure of predictive variance is increasing in the unrestricted model but decreasing in the restricted model. We interpret this result by decomposing the predictive variance into its components:

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

We find that the mean reversion of stock returns always leads to a decrease in the term structure but that this effect is more than compensated for in the unrestricted models by the uncertainty about future expected returns and the estimation risk, so that stocks appear more volatile in the long run. However, when the NEP condition is imposed, the mean reversion dominates, and stocks are less volatile in the long run than they are in the short run.

We also demonstrate that imposing the NEP condition to forecast future stock returns has a substantial impact on the optimal portfolio choice of a long-horizon investor. The gain of using economic variables for forecasting stock returns is often reduced when several aspects of uncertainty are taken into account, such as IID uncertainty, estimation risk, uncertainty on current and future expected returns, and mean reversion. Our results show that imposing the NEP condition on the predictive models allows a long-term investor to obtain more accurate predictions of stock returns. As a result, the long-term investor invests more in stocks than an investor using the unrestricted model for forecasting and eventually benefits from a higher SR and CER.

Supplementary Material

To view supplementary material for this article, please visit http://dx.doi.org/10.1017/S002210902000054X.

References

Avramov, D. "Stock Return Predictability and Model Uncertainty." *Journal of Financial Economics*, 64 (2002), 423–458.

Avramov, D. "Stock Return Predictability and Asset Pricing Models." *Review of Financial Studies*, 17 (2004), 699–738.

Avramov, D.; S. Cederburg; and K. Lučivjanská. "Are Stocks Riskier over the Long Run? Taking Cues from Economic Theory." *Review of Financial Studies*, 31 (2018), 556–594.

Bansal, R., and A. Yaron. "Risks for the Long Run." *Journal of Finance*, 59 (2004), 1481–1509.

Barberis, N. "Investing for the Long Run When Returns Are Predictable." *Journal of Finance*, 55 (2000), 225–264.

van Binsbergenvan, J.; W. Hueskes; R. Koijen; and E. Vrugt. "Equity Yields." *Journal of Financial Economics*, 110 (2013), 503–519.

van Binsbergenvan, J., and R. Koijen. "Predictive Regressions: A Present-Value Approach." *Journal of Finance*, 65 (2010), 1439–1471.

Brandt, M.; A. Goyal; P. Santa-Clara; and J. Stroud. "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning about Return Predictability." *Review of Financial Studies*, 18 (2005), 831–873.

Campbell, J. "A Variance Decomposition for Stock Returns." *Economic Journal*, 101 (1991), 157–179.

Campbell, J., and J. Ammer. "What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns." *Journal of Finance*, 48 (1993), 3–37.

Campbell, J.; Y. Chan; and L. Viceira. "A Multivariate Model of Strategic Asset Allocation." *Journal of Financial Economics*, 67 (2003), 41–80.

Campbell, J., and S. Thompson. "Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average?" *Review of Financial Studies*, 21 (2008), 1509–1531.

Campbell, J., and L. Viceira. *Strategic Asset Allocation*, Oxford, UK: Oxford University Press (2002).

Campbell, J., and L. Viceira. "The Term Structure of the Risk-Return Tradeoff." *Financial Analysts Journal*, 61 (2005), 34–44.

Carter, C., and R. Kohn. "On Gibbs Sampling for State Space Models." *Biometrika*, 81 (1994), 541–553.

Clark, T., and K. West. "Approximately Normal Tests for Equal Predictive Accuracy in Nested Models." *Journal of Econometrics*, 138 (2007), 85–110.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

Cremers, K. "Stock Return Predictability: A Bayesian Model Selection Perspective." *Review of Financial Studies*, 15 (2002), 1223–1249.
Dangl, T., and T. Halling. "Predictive Regressions with Time-Varying Coefficients." *Journal of Financial Economics*, 106 (2012), 157–181.
Della Corte, P.; L. Sarno; and D. Thornton. "The Expectations Hypothesis of the Term Structure of Very Short-Term Rates: Statistical Tests and Economic Value." *Journal of Financial Economics*, 89 (2012), 158–174.
DeMiguel, V.; L. Garlappi; and R. Uppal. "Optimal versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?" *Review of Financial Studies*, 22 (2009), 889–908.
Ferreira, M., and P. Santa-Clara. "Forecasting Stock Market Returns: The Sum of the Parts Is More Than the Whole." *Journal of Financial Economics*, 100 (2011), 514–537.
Golež, B. "Expected Returns and Dividend Growth Rates Implied in Derivative Markets." *Review of Financial Studies*, 27 (2014), 790–822.
Goyal, A., and I. Welch. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies*, 21 (2008), 1455–1508.
Henkel, S.; J. Martin; and F. Nardari. "Time-Varying Short-Horizon Predictability." *Journal of Financial Economics*, 99 (2011), 560–580.
Hoevenaars, R.; R. Molenaar; P. C. Schotman; and T. Steenkamp. "Strategic Asset Allocation for Long-Term Investors: Parameter Uncertainty and Prior Information." *Journal of Applied Econometrics*, 29 (2014), 353–376.
Hong, H., and M. Yogo. "What Does Futures Market Interest Tell Us about the Macroeconomy and Asset Prices?" *Journal of Financial Economics*, 105 (2012), 473–490.
Jeffreys, H. *The Theory of Probability*, Oxford, UK: Oxford University Press (1961).
Johannes, M.; A. Korteweg; and N. Polson. "Sequential Learning, Predictability, and Optimal Portfolio Returns." *Journal of Finance*, 69 (2014), 611–644.
Kandel, S., and R. Stambaugh. "On the Predictability of Stock Returns: An Asset Allocation Perspective." *Journal of Finance*, 51 (1996), 385–424.
Kass, R., and A. Raftery. "Bayes Factors." *Journal of American Statistical Association*, 90 (1995), 773–795.
Leitch, G., and J. Tanner. "Economic Forecast Evaluation: Profits versus the Conventional Measures." *American Economic Review*, 81 (1991), 580–590.
Lettau, M., and S. Ludvigson. "Consumption, Aggregate Wealth, and Expected Stock Returns." *Journal of Finance*, 56 (2001), 815–849.
Lettau, M., and S. van Nieuwerburgh. "Reconciling the Return Predictability Evidence." *Review of Financial Studies*, 21 (2008), 1607–1652.
Li, Y.; D. Ng; and B. Swaminathan. "Predicting Market Returns Using Aggregate Implied Cost of Capital." *Journal of Financial Economics*, 110 (2013), 419–436.
Menzly, L.; T. Santos; and P. Veronesi. "Understanding Predictability." *Journal of Political Economy*, 112 (2004), 1–47.
O'Donnell, C., and T. Coelli. "A Bayesian Approach to Imposing Curvature on Distance Functions." *Journal of Econometrics*, 126 (2005), 493–523.
O'Donnell, C.; A. Rambaldi; and H. Doran. "Estimating Economic Relationship Subject to Firm- and Time-Varying Equality and Inequality Constraints." *Journal of Applied Econometrics*, 16 (2001), 709–726.
Pástor, L. "Portfolio Selection and Asset Pricing Models." *Journal of Finance*, 55 (2000), 179–223.
Pástor, L., and R. Stambaugh. "Predictive Systems: Living with Imperfect Predictors." *Journal of Finance*, 64 (2009), 1583–1627.
Pástor, L., and R. Stambaugh. "Are Stocks Really Less Volatile in the Long Run?" *Journal of Finance*, 67 (2012), 431–477.
Penasse, J. "Return Predictability: Learning from the Cross-Section." Working Paper, available at https://ssrn.com/abstract=2575725 (2016).
Pettenuzzo, D.; A. Timmermann; and R. Valkanov. "Forecasting Stock Returns under Economic Constraints." *Journal of Financial Economics*, 114 (2014), 517–553.
Rapach, D.; J. Strauss; and G. Zhou. "Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy." *Review of Financial Studies*, 23 (2010), 821–862.
Santos, T., and P. Veronesi. "Labor Income and Predictable Stock Returns." *Review of Financial Studies*, 19 (2006), 1–44.
Shanken, J., and A. Tamayo. "Payout Yield, Risk, and Mispricing: A Bayesian Analysis." *Journal of Financial Economics*, 105 (2012), 131–152.
Shiller, R. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review*, 71 (1981), 421–436.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press

Siegel, J. "The Real Rate of Interest from 1800–1990: A Study of the U.S. and the U.K." *Journal of Monetary Economics*, 29 (1992), 227–252.
Siegel, J. *Stocks for the Long Run*, New York, NY: McGraw-Hill (1994).
Stambaugh, R. F. "Predictive Regressions." *Journal of Financial Economics*, 54 (1999), 375–421.
Wachter, J., and M. Warusawitharana. "Predictable Returns and Asset Allocation: Should a Skeptical Investor Time the Market?" *Journal of Econometrics*, 148 (2009), 162–178.
Xia, Y. "Learning About Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation." *Journal of Finance*, 56 (2001), 205–246.
Zhu, X. "Tug-of-War: Time-Varying Predictability of Stock Returns and Dividend Growth." *Review of Finance*, 19 (2015), 2359–2399.

https://doi.org/10.1017/S002210902000054X Published online by Cambridge University Press