On the Economic Value of Stock Market Return Predictors

Scott Cederburg[*] Travis L. Johnson[†] Michael S. O'Doherty[‡]

November 17, 2018

Abstract

Kandel and Stambaugh (1996) demonstrate that forecasting variables with weak statistical support in predictive return regressions can exert considerable economic influence on portfolio decisions. Using a Bayesian vector autoregression framework with stochastic volatility in market returns and predictor variables, we assess the economic value of return predictability and reach a complementary conclusion. Statistically strong predictors can be economically unimportant if they tend to take extreme values in high-volatility periods, have low persistence, and/or follow distributions with fat tails. Several popular predictors exhibit these properties such that their impressive statistical results do not translate into large economic gains for investors.

JEL classifications: G10, G11, G12

[*]Eller College of Management, University of Arizona. Email: `cederburg@email.arizona.edu`.
[†]McCombs School of Business, University of Texas at Austin. Email: `travis.johnson@mccombs.utexas.edu`.
[‡]Trulaske College of Business, University of Missouri. Email: `odohertym@missouri.edu`.

1 Introduction

Although investors and academics have long studied the issue of whether stock market returns are predictable, many historically important predictor candidates generate only weak evidence in favor of predictability with marginal statistical significance and low in-sample predictive R^2s. In a seminal paper, Kandel and Stambaugh (1996) demonstrate that even this relatively weak statistical evidence of predictability can still produce substantive economic effects. In particular, Bayesian investors who consider return predictability evidence display large shifts in their optimal allocation to stocks over time and accrue substantial utility gains. More recently, researchers have uncovered new predictors that achieve strong statistical significance in return forecasting relative to their historical peers. Extrapolating the conclusions from Kandel and Stambaugh (1996), these new variables may promise very large economic gains for investors. Several of these predictor variables do, however, display extreme time-series properties and strong relations with stock market volatility that could cast doubt on their benefits to investors.

We find that the strength of statistical evidence for a forecasting variable does not fully describe its potential value to investors, particularly among newer predictors. Variables that tend to take extreme values in high-volatility periods, have low persistence, and/or follow distributions with fat tails may provide relatively little value even if they successfully forecast market returns in predictive regressions. Our findings complement those of Kandel and Stambaugh (1996) by showing that, while statistically weak evidence can be economically important, the converse also holds true as strong statistical performance may be accompanied by low economic value.

Our conclusions are drawn from an approach that builds on Kandel and Stambaugh's (1996) seminal framework. We consider predictability from the perspective of Bayesian investors who learn about the dynamics of the market return and the predictor from a vector autoregression model. Our model has three primary design features. First, we consider the effects of time-varying volatility on the value of market return predictability. We specifically model returns and predictors to have stochastic volatility, and we show several interesting interactions between volatility and predictability. Second, we investigate the economic value of predictability for multi-period investors. We demonstrate that longer investment horizons impact the perceived value of return predictability evidence, particularly in settings with stochastic volatility and low persistence in the predictor. Third, we implement an in-sample design that preserves the time-series properties of the predictor variables and their relations to market volatility. Using this model, we study a broad set of 25 predictor variables, 14 from Goyal and Welch (2008) and 11 from more recent publications in top finance journals. This large set of predictors provides variation in predictor characteristics, and it allows us to make generalizations about economic value across different predictor types.

Our initial tests on the dividend-price ratio provides a good illustration of Kandel and Stambaugh's (1996) point that weak statistical evidence can translate into important economic effects. The ordinary least squares (OLS) slope coefficient on this variable in a monthly forecasting regression has a p-value of just 0.30, and the regression R^2 value is only 0.002. When the one-month, constant-volatility investor studied by Kandel and Stambaugh (1996) forms beliefs based on information from the dividend-price ratio, she varies her allocation to stocks between 4% and 98% during the 1927-2017 period and achieves a non-trivial utility gain of 0.24% per year in certainty equivalent return (CER).

Within the one-month, constant-volatility framework, we demonstrate that many of the 25 predictors provide substantial economic value to investors. Among the 14 Goyal-Welch variables, most of the CER gains are relatively modest, ranging from 0.03% per year for the dividend-earnings ratio to 0.69% per year for net equity expansion. The 11 new predictors tend to produce large economic benefits, however, with a range of annual CER gain from 0.55% for the nearness to Dow historical high variable of Li and Yu (2012) to 4.37% for the variance risk premium of Bollerslev, Tauchen, and Zhou (2009). Further, we find that OLS R^2 provides a very good fit to CER gain across predictors, such that this statistical metric provides a good indication of in-sample economic value in the one-month, constant-volatility setting.

We proceed to study return predictability in the context of Bayesian investors who account for time variation in stock market volatility. Stochastic volatility has two primary effects on return predictability evidence with a one-month horizon. First, a purely statistical effect changes inferences about the statistical strength of a predictor. Specifically, an investor who believes in stochastic volatility weights observations in the market return predictability regression based on their precision, such that the investor effectively downweights (upweights) information from high-volatility (low-volatility) periods and more efficiently learns about the predictability relation. Inferences about the strength of a predictor can thus differ across the constant-volatility and stochastic-volatility settings. Second, stochastic volatility in market returns is also important for evaluating the economic value of a predictor during the portfolio optimization stage. If a given predictor tends to make extreme return forecasts that coincide with periods of pronounced market volatility, the investor will tend to moderate her bets on stocks and assign a lower economic value to the information in the predictor. We find evidence that both of these effects are at work across the predictors.

An investor forming beliefs about market returns while incorporating stochastic volatility produces different inferences about the value of predictors relative to the constant-volatility investor. Considering time-varying volatility increases the utility gain from some predictors, as is the case for the Treasury bill yield which has a CER gain of 1.18% per year with stochastic volatility versus only 0.35% with constant volatility. Alternatively, some predictors are less valuable under the stochastic-volatility framework, such as the partial least squares aggregated book-to-market ratio of Kelly and Pruitt (2013) with a CER gain of only 0.53% in this setting versus 3.38% in the constant-volatility model. Overall, considering stochastic volatility increases perceived economic value for 11 of the 25 predictors (including seven of the 14 Goyal-Welch predictors and four of the 11 new predictors), whereas benefits fall for the remaining 14 variables. The most pronounced positive differences occur for variables from Goyal and Welch (2008) related to interest rate levels. Several of the new predictors, in contrast, tend to take extreme values during high market volatility periods which leads to lower utility gains. Looking across all 25 variables, the strong relation between standard measures of in-sample statistical significance and CER gains observable in the constant-volatility results is much less pronounced in the stochastic-volatility framework.

We also consider the effect of horizon on the value of return predictability evidence. The dynamics of predictors and expected return forecasts are important considerations for those who do not adjust their market positions at a high frequency. The utility gain for multi-period investors from considering predictability evidence is likely to be lower when a given predictor variable is less persistent because expected return quickly reverts to its mean as horizon increases. In addition to this direct channel, we show an effect on the perceived risk of stocks that arises from uncertainty about future expected return and is most pronounced when a predictor has low persistence and stochastic volatility. Intuitively, the set of low-persistence, high-variance predictors can produce extreme changes in expected market return over short time frames, such that placing a large bet on a forecasted market return runs the risk that the conditional mean return will quickly shift to oppose the bet. Past literature (e.g., Pástor and Stambaugh (2012) and Avramov, Cederburg, and Lučivjanská (2018)) shows the long-horizon effects of uncertainty about future expected return on risk from an investor's perspective, and we add to this literature by demonstrating substantial short-term effects for low-persistence predictors.

To analyze these effects, we study three-month Bayesian investors in either a constant-volatility or a stochastic-volatility framework. We demonstrate that CER gains are often substantially lessened by increasing the horizon from one month to just three months. For example, the CER gains in the constant-volatility setup for the variance risk premium are 4.37% per year with a one-month horizon and 1.00% per year with a three-month horizon. In our sample, each predictor with monthly autocorrelation below 0.90 exhibits substantially less economic value at a three-month horizon compared with a one-month horizon.

In sum, we find important interactions between many return predictors and realistic aspects of market returns relating to stochastic volatility and holding period. While these mechanisms affect each predictor differently, we note several tendencies across predictor variables. First, a broad assessment across predictors indicates that OLS R^2 is a good indication of economic value in the one-month, constant-volatility case but is less meaningful in other, perhaps more realistic, frameworks. In this sense, we find evidence that some predictors are statistically strong but economically weak whereas other variables have weak statistical evidence with relatively high economic value. Second, incorporating stochastic volatility tends to be more favorable to the traditional Goyal-Welch variables compared with the new predictors. Third, the set of predictors that are not highly persistent produce substantially lower utility gains even for multi-period investors with relatively short three-month horizons. Fourth, although we are cautious of making sharp comparisons of CER gains across predictors because they are calculated using different predictive return distributions, there appears to be some tendency for CER gains to be lower in each alternative setting compared with the one-month, constant-volatility framework.

Our investigation of return predictability considers the economic value of the in-sample statistical evidence while considering realistic features of returns and predictors like stochastic volatility. We note, however, that the design of our study is not particularly well-suited for weighing in on the debate about whether a given variable is truly related to expected return. A variable could be a statistically strong predictor of stock market returns, which provides valuable information about how investors price assets in a rational expectations framework, while being economically weak in our setting due to portfolio-choice considerations of Bayesian investors who do not condition on unobserved parameters. As such, our results about the economic value of various predictor variables complement the statistical evidence found in other studies to provide a more complete picture of stock market return predictability using a given variable.

Our study fits within an important literature that investigates return predictability through the lens of Bayesian investors. Kandel and Stambaugh (1996), Stambaugh (1999), Barberis (2000), Cremers (2003), and Wachter and Warusawitharana (2009, 2015) study allocations between stocks and a risk-free security based on market return predictability, as in our study. Avramov (2002, 2004) and Tu and Zhou (2010) investigate portfolio choice with multiple assets that have predictable returns. Related to these studies, Pástor and Stambaugh (2012), Avramov, Cederburg, and Lučivjanská (2018), and Carvalho, Lopes, and McCulloch (2018) consider long-horizon predictive return variance from the perspective of Bayesian investors in settings with predictable market returns.

Three papers that study return predictability from the perspective of Bayesian investors who consider time-varying market volatility are more closely related to ours.[1] Shanken and Tamayo (2012) model returns with time-varying volatility and with expected return as a function of time-varying volatility and payout yield. Johannes, Korteweg, and Polson (2014) examine the question of whether real-time investors can benefit from information in payout yield and stochastic volatility. While these studies also share the feature of incorporating stochastic volatility while considering return predictability, our focus is different. We consider a broad set of predictors to understand the economic value of predictability in an in-sample setting, and we find many of the most interesting effects in predictors that are not closely related to payout yield. Finally, Pettenuzzo, Timmermann, and Valkanov (2014) introduce stochastic volatility in market returns in the context of specifying a constraint on the conditional market Sharpe ratio while estimating the predictive return regression.

The rest of the paper is organized as follows. Section 2 introduces the Bayesian investor's problem, the models for stock market returns, and our estimation procedures. Section 3 discusses the data. Section 4 presents our main results on evaluating market return predictors in frameworks that incorporate stochastic volatility and alternative investment horizons. Section 5 concludes.

2 Methodology

This section develops our approach to investigating the economic importance of the characteristics of stock market return predictors. Section 2.1 introduces the Bayesian investor's problem. Section 2.2 describes models for returns that either do not or do incorporate return predictability and stochastic volatility. Section 2.3 discusses estimation, and Section 2.4 lays out our approach to measuring the economic value of the information contained in a given return predictor.

2.1 Bayesian Investor

We consider a Bayesian investor who chooses optimal allocations between the stock market and a risk-free security. The investor has power utility over wealth,

$$U(W_{T+1}) = \frac{W_{T+1}^{1-\gamma}}{1-\gamma}, \tag{1}$$

[1] Fleming, Kirby, and Ostdiek (2001) and Fleming, Kirby, and Ostdiek (2003) consider the impact of time-varying volatility on portfolio choice in the absence of predictability.

where γ is the coefficient of relative risk aversion. Wealth at time $T+1$ is given by

$$W_{T+1} = W_T \left(R_{f,T+1} + \omega_T R_{T+1} \right), \quad (2)$$

where $R_{f,T+1}$ is the risk-free rate, R_{T+1} is the stock market return in excess of the risk-free rate, and ω_T is a portfolio allocation to stocks that is chosen at time T. In the base case, we consider an investor with $\gamma = 5$ and a one-month investment horizon. We consider investors with longer horizons in Section 4.2. Results corresponding to investors with $\gamma = 2$ or $\gamma = 8$ are available in the Appendix, and inferences about return predictors are similar to the $\gamma = 5$ base case.

Investor i's beliefs about stock market return dynamics are based on a model, $\mathcal{M}_i$. The investor maximizes expected utility by choosing an optimal allocation to stocks,

$$\max_{\omega_T} E[U(W_{T+1})|\mathcal{M}_i, D_T], \quad (3)$$

where the conditional expectation is taken with respect to the predictive distribution of excess stock market returns,

$$p(R_{T+1}|\mathcal{M}_i, D_T) = \int p(R_{T+1}|\mathcal{M}_i, \theta, D_T) p(\theta|\mathcal{M}_i, D_T) d\theta, \quad (4)$$

in which θ is the set of parameters in model $\mathcal{M}_i$, D_T denotes the time series of returns and state variables in model $\mathcal{M}_i$, and $p(\theta|\mathcal{M}_i, D_T)$ is the posterior distribution of θ.

The predictive distribution of excess returns in equation (4) accounts for uncertainty about the parameters in the return process, such that the conditional expectation in equation (3) integrates over this uncertainty. This feature of the asset allocation problem for a Bayesian investor was introduced by Klein and Bawa (1976), and it has the effect that parameter uncertainty increases the riskiness of stocks from the investor's perspective relative to an environment with known parameters. To illustrate this point, note that the predictive variance of the excess return is given by

$$Var(R_{T+1}|\mathcal{M}_i, D_T) = E[Var(R_{T+1}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T] + Var[E(R_{T+1}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T]. \quad (5)$$

The first term is the expected return variance and the second term captures the effect on predictive return variance of uncertainty about the conditional expected market return that exists because the investor does not know the set of parameters θ. The predictive return distribution has higher variance when uncertainty about expected return is considered, and integrating over parameter uncertainty also affects higher moments by producing fat tails in the predictive distribution. As a result, a risk-averse Bayesian investor moderates her investment position relative to an investor with the same preferences who conditions on parameter point estimates.

2.2 Return Process

We study the implications of stock return predictability, stochastic volatility in stock returns, and the interaction of these two effects for Bayesian investors' utility. As such, we specify four alternative models that either do not or do carry these features: (i) no predictability with constant volatility (NP-CV), (ii) predictability with constant volatility (P-CV), (iii) no predictability with stochastic volatility (NP-SV), and (iv) predictability with stochastic volatility (P-SV). Given a candidate predictor variable x_t, the processes for the excess stock market return and the state variable are given by

$$r_{t+1} = \alpha + \beta x_t + \epsilon^r_{t+1}, \tag{6}$$

$$x_{t+1} = \alpha_x + \beta_x x_t + \epsilon^x_{t+1}, \tag{7}$$

where r_{t+1} is the log excess return. The models with no predictability (NP-CV and NP-SV) have the restriction $\beta = 0$. Following much of the return predictability literature, the expected log excess return is specified as a linear function of x_t in models P-CV and P-SV. Finally, the predictor variable follows a stationary AR(1) process such that $-1 < \beta_x < 1$.

The error terms in equations (6) and (7) are conditionally normally distributed, but these conditional distributions differ across the constant-volatility models and the stochastic-volatility models. The errors for the constant-volatility models (NP-CV and P-CV) are distributed bivariate normal,

$$\begin{bmatrix} \epsilon^r_{t+1} \\ \epsilon^x_{t+1} \end{bmatrix} \sim N(0, \Sigma), \qquad \Sigma_{t+1} = \begin{bmatrix} \sigma^2_r & \sigma_{rx} \\ \sigma_{rx} & \sigma^2_x \end{bmatrix}. \tag{8}$$

The errors for the stochastic-volatility models (NP-SV and P-SV) follow the specification of Primiceri (2005). In particular, ϵ^r_{t+1} and ϵ^x_{t+1} are conditionally normally distributed,

$$\begin{bmatrix} \epsilon^r_{t+1} \\ \epsilon^x_{t+1} \end{bmatrix} \sim N(0, \Sigma_{t+1}), \qquad \Sigma_{t+1} = \begin{bmatrix} \sigma^2_{r,t+1} & \sigma_{rx,t+1} \\ \sigma_{rx,t+1} & \sigma^2_{x,t+1} \end{bmatrix}. \tag{9}$$

The conditional covariance matrix Σ_{t+1} can be decomposed as

$$\begin{bmatrix} 1 & 0 \\ a_{t+1} & 1 \end{bmatrix} \Sigma_{t+1} \begin{bmatrix} 1 & a_{t+1} \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} \sigma^2_{r,t+1} & 0 \\ 0 & \tilde{\sigma}^2_{x,t+1} \end{bmatrix}, \tag{10}$$

and the processes for the log standard deviations are

$$\log(\sigma_{r,t+1}) = \log(\sigma_{r,t}) + \eta^r_{t+1}, \tag{11}$$

$$\log(\tilde{\sigma}_{x,t+1}) = \log(\tilde{\sigma}_{x,t}) + \eta^v_{t+1}, \tag{12}$$

where

$$\begin{bmatrix} \eta^r_{t+1} \\ \eta^v_{t+1} \end{bmatrix} \sim N(0, \Omega). \tag{13}$$

Finally, the a_{t+1} process is a random walk with normally distributed errors. This specification for Σ_{t+1} allows for time variation in the conditional volatilities of the return and the state variable as well as time variation in the contemporaneous correlation between the errors.

2.3 Estimation

Each of the four models introduced in Section 2.2 is a restricted Bayesian vector autoregression (BVAR), and we use methods from the literature to estimate each model. The NP-CV and P-CV models are restricted BVARs with a constant covariance matrix Σ. We specify conjugate normal-inverse-Wishart priors with prior parameters chosen to produce diffuse prior distributions. The NP-SV and P-SV models are similar to the specification of the BVAR with time-varying parameters and stochastic volatility in Primiceri (2005), but with constant parameters and parameter restrictions in equations (6) and (7). We broadly follow Primiceri's (2005) estimation approach, with the exception of the estimation of regression parameters in equations (6) and (7).[2] For these parameters, the setup mirrors the constant-volatility models so we specify the same normal prior distribution as in the NP-CV and P-CV cases. In each model, we condition on the initial predictor variable, x_0, such that our posteriors are derived based on conditional likelihoods. As such, our approach does not account for potential Stambaugh (1999) bias in the predictive regression coefficient. Johnson (2018) shows that Stambaugh (1999) bias is small for most of the predictors we consider, and we do not believe that this feature of our research design is likely to affect the spirit of our findings about the

[2]Our estimation procedure incorporates the update to Primiceri (2005) that is shown by Del Negro and Primiceri (2015).

interactions between predictability and stochastic volatility.

We estimate the BVARs using Markov chain Monte Carlo (MCMC) approaches. We use the full time series of data, D_T, to estimate model parameters, such that our study is best viewed as an examination of in-sample return predictability evidence. For each combination of model and predictor, we run the MCMC chain for 110,000 draws and discard the first 10,000 as a burn-in period. As such, we produce 100,000 draws from the posterior distribution for each model and predictor combination. A full description of the MCMC chain and prior parameters is available in the Appendix.

2.4 Economic Value of a Predictor

We examine Bayesian investors' certainty equivalent returns (CERs) to quantify the economic importance of return predictability and stochastic volatility. In particular, we evaluate the impact of information from a predictor variable by comparing the CER for the optimal policy from a model that includes the predictor with the CER that corresponds to the policy that would be optimal under an otherwise similar model without predictability. For example, to discern the value of return predictability in a setting with stochastic volatility, we compare the CER for the optimal P-SV model with the CER that the P-SV investor assigns to the optimal policy for the NP-SV model. In these cases, expected utility is taken with respect to the predictive return distribution from the P-SV model, and the two CERs under comparison are based on the optimal portfolio weights for the P-SV and NP-SV models to isolate the economic effect of the return predictability signal when returns have stochastic volatility. Formally, investor i forms beliefs about the predictive return distribution in equation (4) using model $\mathcal{M}_i$. Investor i's CER with the optimal policy under model $\mathcal{M}_i$ denoted as $\omega^*_{i,T}$,

$$CER_i = [(1-\gamma)E[U(W_T(R_{f,T+1} + \omega^*_{i,T} R_{T+1}))|\mathcal{M}_i, D_T]]^{\frac{1}{1-\gamma}}, \tag{14}$$

can be compared with investor i's CER from adopting the optimal policy $\omega^*_{j,T}$ from an alternative model $\mathcal{M}_j$,

$$CER_{i,j} = [(1-\gamma)E[U(W_T(R_{f,T+1} + \omega^*_{j,T} R_{T+1}))|\mathcal{M}_i, D_T]]^{\frac{1}{1-\gamma}}. \tag{15}$$

The CER difference, $\Delta CER_{i,j} = CER_i - CER_{i,j}$, reflects the economic magnitude of the difference between the optimal policies under models i and j from investor i's perspective. This method of measuring the economic value of information is used by Kandel and Stambaugh (1996), Pástor and Stambaugh (2000), and Avramov (2004), among others.

We are particularly interested in time-series properties of the predictor variables and the corresponding effects on the economic value of the predictors. Several of the predictor variables exhibit extreme values and display interesting relations to stock market volatility. As such, our analysis is designed to preserve much of the structure of the time series from the observed sample period when we evaluate a given predictor. In particular, we examine CERs based on predictive return distributions for each month $t + 1$ of the sample period that are generated based on the observed level of the predictor variable x_t and posterior draws of the parameters. This approach preserves the observed time-series properties of x_t. Further, with this design the models with stochastic volatility maintain the time-series relation between x_t and Σ_{t+1} that is estimated from the data, which is important for our goal of investigating the interactions between return predictability and stochastic volatility.

Within this general structure, our approach to estimating CER differences has four steps:

1. We estimate a model using the full time series of data to produce 100,000 draws from the posterior distribution as described in Section 2.3.

2. We produce 100,000 draws from the conditional predictive distribution of stock market excess returns in equation (4) for each month in the sample period. Specifically, to draw returns in month $t + 1$, we condition on the value of the predictor variable from month t, x_t. We then draw a random log excess return from the model in equation (6) conditional on the parameters from a posterior draw. In the models with stochastic volatility (NP-SV and P-SV), the posterior draw also includes a full time series of stochastic volatility of returns, such that we condition on a draw of $\sigma^2_{r,t+1}$. Finally, we convert the draws of log excess returns to excess returns to produce a set of 100,000 draws from the predictive distribution, which integrates over parameter uncertainty.

3. Given draws from a predictive return distribution, we find the portfolio weight that maximizes average utility across the draws.[3] We assume that the risk-free security pays a constant rate of 2% per year. Because the conditional predictive distribution varies by month for all models but NP-CV, we produce an optimal portfolio weight for each combination of model and month except for the NP-CV model which has a constant optimal weight.

[3]Given that we do not constrain the investment in stocks between zero and one, there is some positive probability of losses exceeding 100% of wealth for any model-month combination that produces an optimal weight outside of these bounds. We assume that the worst possible investment return that can be realized is -90%, which solves the theoretical and numerical problems associated with zero or negative wealth and could be motivated by the real-world existence of bankruptcy laws that limit actual harm in the event of large losses and liabilities.

4. We calculate the CER of the optimal weight and the CERs of weights that are optimal under other models. To arrive at a single CER value for a time series of weights that spans the sample period, we calculate the average utility across all months that is achieved by following the policy and find the CER that provides equivalent utility. Finally, we annualize by multiplying monthly CER differences by 12.

The design of our study reflects our focus on in-sample return predictability evidence. Specifically, our goal is to examine the correspondence between the strength of in-sample statistical evidence and the economic value of a predictor while taking into account realistic features of returns and investors. Using the predictive return distributions from the sample period while conditioning on the full sample for estimation produces results for Bayesian investors that are directly comparable to in-sample statistical evidence from the predictive regressions that are common in the literature. We purposefully abstract away from the potential difficulties in translating in-sample performance into out-of-sample gains to maintain this direct comparison.

3 Data

Our empirical tests focus on forecasting log excess stock market returns using a variety of predictor variables. Our proxy for the market portfolio is the Center for Research in Securities Prices (CRSP) value-weighted index. We collect monthly time-series data on the excess market return and the risk-free rate from Kenneth French's website.[4] The log excess market return is the log return on the CRSP index less the log return on the risk-free asset.

We consider a wide range of forecasting variables from prior literature. We examine stock return predictability at a monthly horizon, so we restrict the sample to predictors that are available at a monthly frequency. We also require that each predictor variable has data availability through December 2017. We provide a brief overview of the predictors below. Full details on variable definitions, data sources, and construction methods are available in the Appendix.

We start with the 14 monthly predictor variables from Goyal and Welch (2008). This set of predictors includes the dividend-price ratio (DP), the dividend yield (DY), the earnings-price ratio (EP), the dividend-earnings ratio (DE), stock market variance ($SVAR$), the book-to-market ratio (BM), net equity expansion ($NTIS$), the Treasury bill yield (TBL), the long-term Treasury bond yield (LTY), the long-term Treasury bond return (LTR), the term spread (TMS), the default yield

[4] See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/. We thank Kenneth French for making these data available.

spread (DFY), the default return spread (DFR), and inflation ($INFL$). This group of forecasting variables is widely used in the literature on stock market return predictability.

We augment the Goyal and Welch (2008) predictors with a second group of 11 forecasting variables introduced in more recent literature. We specifically search articles appearing in top finance journals subsequent to the publication of Goyal and Welch (2008) and identify predictors with data availability at a monthly frequency. This group includes Kelly and Pruitt's (2013) partial least squares aggregated book-to-market ratio (PLS), Cooper and Priestly's (2009) output gap (GAP), Jones and Tuzel's (2013) new orders-to-shipments of durable goods (NOS), Li and Yu's (2012) nearness to Dow historical high (DOW), Kelly and Jiang's (2014) tail risk ($TAIL$), Pollett and Wilson's (2010) average correlation (COR), Rapach, Ringgenberg, and Zhou's (2016) short interest index (SII), Huang and Kilic's (2018) gold-to-platinum price ratio (GP), Driesprong, Jacobsen, and Maat's (2008) oil price change (OIL), Bollerslev, Tauchen, and Zhou's (2009) variance risk premium (VRP), and Bollerslev, Todorov, and Xu's (2015) left jump tail variation (LJV).

Table I reports summary statistics for the 25 predictor variables. Panel A presents the sample period start date, mean, standard deviation, skewness, kurtosis, and monthly autocorrelation coefficient for each of the Goyal and Welch (2008) predictors, and Panel B shows the corresponding statistics for the new predictors. The times series for the Goyal and Welch (2008) variables in Panel A and the PLS variable in Panel B all begin in January 1927. Data for the other new predictors cover shorter sample periods.

Table I highlights two properties of the predictors that are relevant for our subsequent analysis. First, the empirical distributions for several of the forecasting variables are highly non-normal, as indicated by the skewness and kurtosis statistics. In the standard univariate predictive regression setting, non-normal forecasting variables directly imply that conditional expected stock returns are also non-normal and tend to take on extreme values in a few sample months. Deviations from normality are particularly acute for several of the new predictor variables. Both PLS and VRP, for example, are highly negatively skewed. These variables are positive return predictors, such that negative skewness translates into extreme negative expected return forecasts for months in which PLS and VRP take on their lowest values. The PLS and VRP predictors also exhibit fat tails, with kurtosis measures of 32.48 and 56.17, respectively. Second, there is substantial variation in measures of persistence across the predictors. Fifteen of the 25 forecasting variables in Table I have monthly autocorrelation coefficients that exceed 0.95. Four predictors, in contrast, have autocorrelation coefficients below 0.50 in magnitude. Autocorrelation is an important summary measure because more highly persistent predictors tend to be more influential in settings that require multi-period return forecasts.

Table II presents results from standard univariate predictive regressions. We specifically regress the log excess stock market return on the lagged value of each predictor. For each regression, the table reports the OLS estimate of the slope coefficient, its corresponding two-tailed p-value with heteroscedasticity-consistent standard errors, and the regression R^2. For the Goyal and Welch (2008) predictors in Panel A, the statistical evidence in favor of predictability tends to be modest. Only four of the 14 variables are statistically significant return predictors at the 10% level, and just one variable (DY) is significant at the 5% level. The monthly R^2 values tend to be small, ranging from 0.000 to 0.006. As emphasized by Kandel and Stambaugh (1996), however, even weak statistical evidence of predictable returns can have an economically large impact on asset allocation decisions.

Panel B of Table II shows that the new predictors generate considerably stronger statistical support for return predictability. Eight of the 11 variables are statistically significant at the 10% level, and six remain significant at the 5% level. The regression R^2s are also more impressive than those in Panel A, with R^2 values as high as 0.031 for PLS and 0.049 for VRP. Based on the analysis in Kandel and Stambaugh (1996), these results suggests that many of the new predictors should be of considerable economic value to investors making asset allocation decisions. At the same time, however, the non-normal distributions for these variables and some of their other time-series properties may limit their value in portfolio applications.

4 Results

In this section, we examine the market return predictors from the perspective of Bayesian investors who each believe in one of the models introduced in Section 2. We begin our analysis by focusing on the economic value of the dividend-price ratio (DP) predictor in a setting that closely matches that of Kandel and Stambaugh (1996) to illustrate their findings that even weak statistical evidence of return predictability may be economically important. As shown in Table II, DP qualifies as a statistically weak predictor with a p-value of 0.295 for the OLS predictability coefficient and a monthly predictive regression R^2 of 0.002. To maintain consistency with Kandel and Stambaugh (1996), we examine DP from the perspective of an investor who believes that stock market volatility is constant.

Figure 1 shows properties of the predictive return distribution and optimal portfolio weights in the constant-volatility framework. The top panel shows quantiles of the predictive return distribution from the P-CV model. The model produces a predictive return distribution for each month in the sample period conditional on the value of DP, and the figure plots the median (solid line) and 25th and 75th percentiles (dashed lines) of the distribution. The bottom panel shows the difference between the optimal portfolio weight under the P-CV model and the optimal weight for the NP-CV model. As such, this weight difference isolates the effect of incorporating information from the DP predictor variable on the optimal portfolio weight.

The results in Figure 1 illustrate that even statistically weak evidence of return predictability can have a large effect on the optimal portfolio weight of a Bayesian investor who considers predictable returns. The optimal portfolio weight in stocks for the NP-CV model is 46% and constant across periods (because returns have constant moments). An investor who believes in the P-CV model with the DP predictor varies her weight in stocks between 4% (September 2000) and 98% (July 1932) during the sample period. Despite the weak evidence of predictability, the location of the predictive return distribution shifts over time from the perspective of the P-CV investor as she optimally considers information from the predictor rather than rejecting it based on a statistical test.

To measure the economic value of the DP predictor for the P-CV investor, we calculate the difference between the investor's CERs under the optimal weights for the P-CV and NP-CV models. This annualized CER difference is 0.24%, which represents the utility gain from considering information in the DP variable. This finding is consistent with Kandel and Stambaugh's (1996) conclusion that even a predictor with statistically weak evidence can still provide economic value.

Table III reports the full set of CER differences across predictors, models, and horizons. Panel A shows results for the Goyal-Welch predictors and Panel B contains corresponding results for the new predictors. For each predictor, we report four CER differences that measure the economic value of the predictor variable under a specific model-horizon combination. The constant volatility cases report CER differences for the P-CV investor who compares weights from the P-CV and NP-CV models for a given horizon, and the stochastic volatility cases are from the perspective of the P-SV investor who considers the P-SV and NP-SV model weights. We report annualized CER differences for one-month and three-month investment horizons.

Beginning with the one-month, constant-volatility CER differences, the results indicate substantial variation in economic value across predictors. The CER differences for the Goyal-Welch predictors in Panel A of Table III range from 0.03% (DE) to 0.69% ($NTIS$) per year. Many of the new predictors in Panel B have larger CER differences, consistent with the tendency of stronger statistical evidence for these predictors from the OLS regressions in Table II. The smallest CER difference is 0.55% (DOW) and the largest is 4.37% (VRP) per year. Among the 11 new predictors, CER differences that exceed 1.00% per year are achieved by seven predictors: PLS, GAP, NOS, $TAIL$, GP, VRP, and LJV. These findings indicate that the new predictors have considerable economic value to investors who believe in the constant-volatility framework studied by Kandel and Stambaugh (1996).

Comparisons of these initial CER differences with those from the remaining cases in Table III demonstrate that the economic value of a given predictor can differ substantially depending on the investor's model and horizon. Focusing on extreme examples with a one-month horizon, the CER difference for the inflation predictor variable ($INFL$) is 0.16% for the P-CV investor and 1.34% for the P-SV investor, whereas the CER difference for the PLS variable is 3.38% for P-CV versus only 0.53% for P-SV. These cases illustrate that considering stochastic volatility can either increase or decrease the economic value of a predictor variable. Overall, the CER difference is larger under stochastic volatility compared with constant volatility for seven of the 14 Goyal-Welch predictors and four of the 11 new predictors. Considering different horizons within the constant-volatility framework shows that the value of many of the predictors is similar for one-month and three-month horizons (e.g., 0.24% per year for DP for each horizon) but substantially lower for a few predictors (e.g., 4.37% for VRP at a one-month horizon and 1.00% at a three-month horizon).

Figure 2 provides a visual illustration of the CER differences across the various predictors and cases. The figure contains four panels with scatterplots of CER differences compared to the OLS R^2 from Table II for each predictor. The top-left panel shows CER differences for the P-CV investor with a one-month horizon. The plot demonstrates that OLS R^2 is a very good indicator of the economic value of a given predictor in the one-month, constant-volatility setup considered by Kandel and Stambaugh (1996). This finding reinforces that our test design is congruent with our goal of examining the in-sample economic value of predictors, as there is a close correspondence between CER differences and in-sample OLS R^2. The two predictors that visually deviate the most from a linear relation between OLS R^2 and CER difference across the predictors are VRP and LJV, which are the two predictors with the shortest sample periods (VRP and LJV data begin in January 1990 and June 1996, respectively). Given the relatively short sample periods, the predictive return distributions for these predictors have somewhat fatter tails compared with the distributions for other predictors, which reduces their economic value. Nonetheless, there is a clear connection between the strength of statistical evidence and the economic value of a predictor in the Kandel and Stambaugh (1996) setup.

The remaining plots in Figure 2 demonstrate that OLS R^2 provides a murkier indication of economic value for investors who consider stochastic volatility or who have longer investment horizons. In each of these panels, CER differences are somewhat positively related to OLS R^2, but there are clear exceptions of predictors that are statistically strong and economically weak or economically strong and statistically weak (according to OLS). These results indicate that the economic value of a predictor to investors may not be fully captured by the strength of statistical evidence from OLS.

In the remainder of this section, we further investigate the economic mechanisms that produce differences in the perceived economic value of a predictor across investor types. Section 4.1 discusses the effects of incorporating stochastic volatility into the model when evaluating a predictor variable. Section 4.2 demonstrates the impact of investor horizon on the value of return predictability evidence.

4.1 Return Predictability with Stochastic Volatility

The results in Table III show that information from the same predictor variable can have substantially different degrees of economic value for investors in the constant-volatility and stochastic-volatility frameworks. In this section, we discuss two mechanisms that contribute to these differences. First, a purely statistical effect captures how Bayesian investors learn about market return predictability. The P-SV investor considers conditional volatility while weighing information from the sample and may, therefore, arrive at different conclusions about the statistical evidence of market return predictability for a given predictor variable compared with the P-CV investor. Second, given information from a predictor variable about the conditional expected market return, the P-SV investor will also consider information about the conditional market variance while making her portfolio decision. To the extent that the timing of extreme expected return predictions corresponds to periods of high market volatility, the P-SV investor may optimally choose to moderate her bets on stocks relative to the P-CV investor who believes in constant market variance.

Accounting for time variation in stock market variance can affect inferences about market return predictability. Bayesian investors who believe in constant market volatility effectively weight squared residuals equally across periods, similar to OLS regression in the frequentist framework. In contrast, Bayesian investors who believe in stochastic volatility will weight information from each month according to its precision. That is, if a very high market return occurs with a high value of the predictor variable during a high-volatility period, OLS treats the observation as strong evidence for a positive predictive relation whereas a Bayesian stochastic-volatility investor may attribute much of the high return to a positive realization of the high-variance residual rather than to a high expected return. In this sense, the P-SV investor estimates the predictive regression in a similar fashion to a frequentist WLS approach that weights observations by the inverse of estimated conditional market return variance (see Johnson (2018) for an application of WLS to market return predictability).

The effect of time-varying volatility on learning about a predictor variable depends on whether the strength of its predictability evidence is concentrated in high- or low-volatility periods. Depending on the relation, the posterior mean of the predictive regression coefficient β could increase or decrease in magnitude in the stochastic-volatility model relative to the constant-volatility model. Assuming that predictability evidence is unrelated to market return variance, weighting observations by their precision also improves the efficiency of estimating β. As such, accounting for time variation in return variance while estimating the return predictability relation can affect both the location and the scale of the posterior distribution of β.

Figure 3 shows posterior distributions of predictive regression coefficients that relate log excess stock market return to lagged predictor variables following equation (6). Panel A shows posteriors for the Goyal-Welch predictors and Panel B reports results for the new predictors. For each predictor, we show a box-and-whiskers plot for the posterior distribution of β for both the P-CV and P-SV models. In each box-and-whiskers plot, the red line shows the posterior median, the box represents a 50% credible interval, and the whiskers span a 95% credible interval. We provide a dotted line at zero in each plot for convenience of determining whether zero falls within a credible interval.

The posteriors in Figure 3 indicate that investor beliefs about the strength of return predictability for a given predictor often vary across the P-CV and P-SV investors. In Panel A, the posterior medians noticeably shift toward zero for the DY, EP, BM, $NTIS$, TMS, and DFR predictors. As a specific example, the posterior median for DY under the P-SV model is about 58% of the magnitude of the median for the P-CV model, such that the magnitude of variation in conditional expected return forecasts is about two-fifths lower when learning about predictability while considering stochastic volatility. In contrast, variables related to interest rate levels including TBL, LTY, LTR, and $INFL$ appear to better forecast returns with stochastic volatility. The $INFL$ variable, for which the CER difference is much larger for the P-SV investor compared with the P-CV investor in Table III, has posterior medians of -0.36 and -0.98 for the P-CV and P-SV models, respectively. The new predictors in Panel B also show effects, as the posteriors under the P-SV model are centered noticeably closer to zero for the PLS, GAP, $TAIL$, SII, GP, and LJV predictors. Evidence of market predictability with the OIL variable is somewhat stronger after considering stochastic volatility.

Changes in inferences about market return predictability have clear implications for the economic value of a given predictor. A shift in the posterior distribution of the predictive regression coefficient directly affects the conditional expected market return in each period. Further, to the extent that the posterior is less (more) diffuse under the P-SV model compared with the P-CV model, the predictive return distribution is also less (more) diffuse. In particular, equation (5) shows that reducing uncertainty about β has the effect of lowering predictive return variance by decreasing the estimation risk component of variance from uncertainty about expected return, $Var[E(R_{T+1}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T]$. Thus, the degree of certainty about the market return predictability relation has a direct impact on the perceived risk of investing in stocks from the perspective of a Bayesian investor.

The second effect of stochastic volatility on the Bayesian investor's problem is that the investor must consider the conditional risk of investing in stocks when determining whether (and how much to) bet on a view about expected market return. The impact of this issue depends on the relation between the conditional expected market return and the conditional volatility of market return. To the extent that a predictor variable produces a few extreme values of conditional expected market return that correspond to periods with high market volatility, the investor is likely to adopt moderate investment positions and associate a relatively low economic value with the predictor.

To illustrate these effects, we more closely examine the PLS predictor variable. Table III shows that the P-CV investor has a CER difference of 3.38% for this variable, whereas the P-SV investor's CER difference is only 0.53%. A portion of this decline in economic value of PLS is attributable to a shift in the investor's views about the statistical evidence of return predictability that is apparent in Figure 3. The posterior median of the predictive regression coefficient is 0.028 for the P-CV model compared with 0.016 for the P-SV model. Moreover, uncertainty about β is higher in the P-SV model with a posterior standard deviation of 0.007 versus 0.005 for the P-CV model, such that the perceived riskiness of stocks is somewhat higher under the P-SV model. Nonetheless, statistical evidence that PLS positively forecasts market returns remains strong in the P-SV model and over 98% of posterior draws of β are positive.

The more important effect of stochastic volatility on the economic value of PLS is in the portfolio decision step. Figure 4 illustrates the relation between the predictor variable and the conditional variance of market returns. The top panel plots the time series of PLS and the bottom panel shows the time series of the annualized standard deviation of market returns implied by the posterior mean of the stochastic volatility process. Consistent with past empirical work, market volatility is highly time varying with large spikes that generally correspond to times of economic uncertainty. Volatility peaked during the Great Depression, with the annualized standard deviation reaching as high as 55%. This period of extreme market volatility corresponds closely with the most extreme values of PLS. In particular, 90% of PLS observations fall within the range of -1.30 to -0.45, but the variable drops as low as -3.44 during the market volatility spike in the early 1930s. The strong relation between PLS and market volatility is likely to affect its economic value in making portfolio decisions.

Figure 5 shows predictive return distributions and weight differences for the P-CV and P-SV investors who consider PLS. The top panels show the median (solid line) and 25th and 75th percentiles (dashed lines) of the predictive return distribution from a given model. The bottom panel on the left (right) plots the difference between the optimal portfolio weight under the P-CV (P-SV) model and the optimal weight for the NP-CV (NP-SV) model. The weight differences represent the effect of information from PLS on the optimal portfolio weight.

The results in Figure 5 demonstrate stark differences in the predictive return distribution and optimal portfolio weights across the constant-volatility and stochastic-volatility cases. Beginning with the P-CV model, the median of the predictive return distribution varies substantially over the sample period with the largest shift occurring in the early 1930s. In the month with the lowest conditional return forecast (January 1932), the median of the predictive return distribution is -6.72% and 90% of return draws from the predictive distribution are negative. The P-CV investor thus concludes that the stock market return is highly likely to be large and negative, and she uses this predictive distribution to calculate an optimal weight of -392% in stocks (relative to a weight of 46% for the NP-CV investor).[5] After the portfolio weight stabilizes in the early 1940s, the investor continues to shift her investment position with weight differences as low as -45% (March 2000) and as high as 84% (September 1960), but the lion's share of the investor's response to information in PLS occurs early in the sample.

The investor who believes in the P-SV model responds to PLS quite differently compared with the P-CV investor. In January 1932, the median of this investor's predictive return distribution reaches a low of -3.57%. However, given that this period corresponds to a time of high market volatility, the standard deviation of monthly return from the predictive distribution is 13.69% and the investor believes there is a 39% chance that the market return will be positive. Thus, despite the large, negative expected return forecast, the P-SV investor adopts a relatively modest portfolio

[5]It is notable that the investor's weights in stocks were -75% and -78% in July and August 1932 when the monthly stock market excess returns were 34% and 37%, suggesting the weights of around -400% in stocks during the first half of 1932 produced substantial risk of financial ruin.

weight of −25% in stocks compared with a weight of 18% for the NP-SV investor during the same month. Under less volatile market conditions, the P-SV investor often responds similarly to information in PLS compared with the P-CV investor. For example, weight differences of −30% in March 2000 and 66% in September 1960 are similar in magnitude to those reported above for the constant-volatility investor. Overall, however, the P-SV investor makes fewer large bets based on PLS during the sample period. The investor's muted response to the predictor variable accounts for the relatively low CER difference of 0.53% under stochastic volatility compared with 3.38% for constant volatility.

Discussion

The case of the PLS predictor shows that the presence of stochastic volatility can have a substantial impact on the perceived economic value of predictability evidence. The economic value of PLS is lower with stochastic volatility because (i) the estimated variation in expected return is lower because the predictability coefficient is shifted toward zero, (ii) uncertainty about expected return is higher because the posterior of the predictability coefficient is more diffuse, and (iii) periods with more extreme expected return forecasts tend to correspond to times of high market volatility. This example is somewhat extreme because the predictor is affected in the same direction by each of the three effects. We see several other cases in Table III in which considering stochastic volatility either increases or decreases perceived economic value. In particular, the Goyal-Welch variables related to interest rate levels along with the OIL variable have substantially higher value with stochastic volatility, whereas there are notable declines in value for a few of the newer predictors that bear a relation to market volatility including PLS, $TAIL$, VRP, and LJV. To diagnose whether a new predictor is likely to be susceptible to these issues, we recommend considering both OLS and WLS for estimating predictive relations to assess statistical evidence as well as visually inspecting whether the expected return forecasts are closely related to estimated market return volatility to gauge the potential impact on portfolio choice.

4.2 Return Predictability and Investor Horizon

We now consider the economic value of market return predictability from the perspective of Bayesian investors with multi-period horizons. Specifically, we study buy-and-hold investors who form beliefs about market returns based on one of the four models studied above (i.e., NP-CV, P-CV, NP-SV, or P-SV) but who have three-month investment horizons. Studying predictability from the perspective of multi-period investors has practical appeal since many investors do not rebalance their portfolios on a monthly basis. Further, it is interesting in this setting because the joint dynamics of the market return and the predictor affect the predictive return distribution at horizons longer than one month, so additional characteristics of predictors are considered by longer-horizon investors.

Our investigation of multi-period Bayesian investors relates to several previous studies. Stambaugh (1999), Barberis (2000), and Avramov (2002) consider multi-period, buy-and-hold Bayesian investors similar to our setup, and Johannes, Korteweg, and Polson (2014) and Hoevenaars, Molenaar, Schotman, and Steenkamp (2014) study investors who periodically rebalance their portfolios.[6] A related literature, including studies by Pástor and Stambaugh (2012), Avramov, Cederburg, and Lučivjanská (2018), and Carvalho, Lopes, and McCulloch (2018), investigates the long-horizon predictive variance of stock market returns in the presence of market return predictability.

We largely focus on different issues about multi-period investment horizons compared with the previous literature. Prior studies primarily focus on the dynamics of highly persistent predictor variables, and they document important effects like mean reversion in returns using persistent predictors and long investment horizons. Our broad sample of predictors contains several variables with relatively low persistence, and we find that these low-persistence predictors produce interesting dynamics and portfolio choice implications for multi-period investors. Even over a relatively short horizon of three months, these predictors generate interesting patterns for expected return and risk, and we focus much of our attention on these effects among low-persistence predictors.

We first describe our method for calculating optimal portfolio weights and CER differences for three-month, buy-and-hold Bayesian investors. These investors maximize expected utility by choosing a weight ω_t to invest in stocks at the beginning of the three-month period. Expected utility is calculated with respect to the predictive distribution of three-month, buy-and-hold excess stock market returns. To generate three-month cumulative returns, we first note that the VAR structure specified in equations (6) and (7) allows us to draw (r, x) pairs for periods $t+1$, $t+2$, and $t+3$ based on the evolution of the state variable and the distribution of the log excess stock return conditional on a value of the lagged state variable. Once we have draws of log excess returns r_{t+1}, r_{t+2}, and r_{t+3}, we convert them into excess returns and calculate the buy-and-hold excess return as $R_{t,t+3} = (R_{f,t+1} + R_{t+1})(R_{f,t+2} + R_{t+2})(R_{f,t+3} + R_{t+3}) - R_{f,t+1}R_{f,t+2}R_{f,t+3}$. The optimal portfolio weight ω_t maximizes average utility across 100,000 draws from the predictive distribution of three-month buy-and-hold excess market returns.[7] Finally, we calculate CER differences analogously to

[6] A related literature, discussed and contributed to by Campbell and Viceira (2002) and Wachter (2010), considers optimal allocations by multi-period investors in the absence of estimation risk.

[7] We choose to study a relatively short three-month horizon because we are able to illustrate our main points without encountering problems with the calculation of expected utility. Given that there is no closed-form solution for the predictive return distribution, we estimate expected utility by calculating average utility across 100,000 draws.

the one-month case and annualize by multiplying by four.

As previously mentioned, our analysis in this section largely concentrates on those predictors that are relatively less persistent. To identify these predictors, Figure 6 shows posteriors of the autoregression coefficient β_x from equation (7) for each predictor variable. Panel A shows posteriors for the Goyal-Welch predictors and Panel B reports results for the new predictors. For each predictor, we produce a box-and-whiskers plot for the posterior distribution of β for both the P-CV and P-SV models. In each box-and-whiskers plot, the red line shows the posterior median, the box represents a 50% credible interval, and the whiskers span a 95% credible interval.

The posteriors in Figure 6 show that most of the predictor variables are highly persistent with autoregression coefficient posteriors that are concentrated near one. Perhaps unsurprisingly, the persistent state variables tend to maintain relatively similar levels of economic value with a three-month horizon as in the one-month case. Other variables, however, have lower levels of persistence. The variables with a posterior median of β_x below 0.90 in either the P-CV or P-SV models are $SVAR$, LTR, DFR, and $INFL$ among the Goyal-Welch predictors in Panel A and NOS, $TAIL$, OIL, and VRP for the new predictors in Panel B. The CER differences in Table III for each of these predictors are noticeably smaller for three-month investors compared with the corresponding one-month investors. Specializing to the stochastic-volatility case, the three-month CER differences are less than half of those at a one-month horizon for $SVAR$ (0.11% for three-month versus 0.23% for one-month), LTR (0.12% versus 0.82%), $INFL$ (0.40% versus 1.34%), NOS (0.64% versus 1.50%), OIL (0.30% versus 1.13%), and VRP (1.31% versus 3.41%).

The effects of longer horizons on the economic value of low-persistence predictors work through two primary channels. The first channel is a relatively straightforward effect on expected market return of the three months in the holding period. Given a predictor with high persistence, the investor's expectation of market returns will also be highly persistent. For example, an investor who observes a high value of DP in month t believes that the expected return in month $t+1$ is higher than usual. When forming expectations about market returns in months $t+2$ and $t+3$, the investor anticipates that DP will remain high and therefore continues to expect relatively high market returns throughout the holding period. In contrast, market return forecasts based on low-persistence predictors quickly converge toward the long-run mean as the horizon increases. Thus, when presented with evidence that the expected market return is high in month $t+1$, the investor believes that the predictor does not contain much information about market returns in months

This approach is stable when there is a low probability of very large, negative returns that correspond to extremely large negative utility. At long horizons, the probability of these large, negative events increases and the expected utility estimation deteriorates. We find that the estimates with a three-month horizon are stable.

$t + 2$ and $t + 3$. The investor is therefore likely to perceive less of an economic gain from the low-persistence predictor when she has a longer horizon.

The second channel of persistence on the economic value of a predictor is the effect on risk. As shown by Pástor and Stambaugh (2012), uncertainty about future expected returns has a positive impact on predictive return variance from the perspective of a Bayesian investor. This risk component arises because there is uncertainty about the evolution of expected return over the investment horizon, and it is relatively small at short horizons such as three months when predictors are highly persistent. For low-persistence predictors, however, we show that this effect can be quite large even over short time frames. In particular, we formally demonstrate below that low-persistence predictors tend to be highly volatile compared with high-persistence predictors. As such, expected return can change a great deal from month to month as the predictor varies over time. An investor who is investing for three months when the conditional expected return for the next month is high must also acknowledge that the expected return could end up being quite low in the second and third months of the holding period. Thus, longer-horizon bets on stocks using low-persistence predictors are perceived to be risky from the perspective of the Bayesian investor.

We now demonstrate that low-persistence predictors tend to be volatile relative to high-persistence predictors for a given level of market return predictability. The OLS R^2 is a measure of the variance of the fitted expected return as a proportion of return variance. Given the structure of the BVAR in equations (6) and (7) with constant volatility, the unconditional variance of the fitted expected excess log return is given by

$$Var(E(r_{t+1}|\mathcal{M}_i, \theta, x_t)|\mathcal{M}_i, \theta) = \frac{\beta^2}{1 - \beta_x}\sigma_x^2. \tag{16}$$

When β_x is near one for a persistent predictor, the quantity $1/(1 - \beta_x)$ contributes substantially to this quantity (e.g., $1/(1 - \beta_x) = 100$ when $\beta_x = 0.99$, which is a common range for β_x). For low-persistence predictors, however, the quantity $\beta^2\sigma_x^2$ must be relatively large to produce a non-trivial OLS R^2. This quantity also drives the contribution of uncertainty about future expected returns to predictive return variance. Formally, uncertainty about future expected returns over a three-month horizon is driven by uncertainty about the expected return in months two and three of the holding period, and it is given by $E(\beta^2((1+\beta_x)^2+1)\sigma_x^2|\mathcal{M}_i, D_T)$ (see the Appendix for derivation). Relative to the unconditional variance of expected return given by equation (16), this quantity is very small for high β_x but large for low β_x. Uncertainty about future expected returns can thus contribute substantially to predictive return variance even over short horizons for predictors with low persistence. Moreover, the effects can be further amplified in the stochastic-volatility model when $\sigma^2_{x,t+1}$ is large.

To illustrate the potential effects of investment horizon on the economic value of a predictor variable, we focus on the VRP predictor. The annualized CER difference for P-CV investors is 4.37% at a one-month horizon but only 1.00% at a three-month horizon. Similarly, the P-SV investor with a one-month horizon has a CER difference of 3.41% compared with only 1.31% for the three-month investor. This variable exhibits low persistence, substantial stochastic volatility in the predictor process, and a tendency for extreme values within the sample period. Each of these features works against the value of the variable for longer-term investors.

Figure 7 shows the VRP predictor variable along with the stochastic volatility processes for returns and the state variable. The VRP variable displays a noticeable downward spike in October 2008, but we note that this observation post-dates the sample periods of the Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011) studies such that the initial evidence for VRP in the literature is not driven by the outlier. Nonetheless, this observation contributes to a pattern observed throughout the sample that the conditional volatility of VRP is highly variable. There is also a tendency for spikes in the volatility of VRP to correspond with spikes in market volatility.

Figure 8 plots predictive return distributions and weight differences for one-month and three-month horizons. Panel A (Panel B) shows results for the constant-volatility (stochastic-volatility) models. The predictive return distribution plots report the median and 25th and 75th percentiles of the distribution of one-month or three-month cumulative returns. The weight differences are the difference between optimal weights for the P-CV and NP-CV models in Panel A and the P-SV and NP-SV models in Panel B.

The one-month, constant-volatility results show that this investor aggressively shifts her portfolio weights in response to information from VRP. The weight in stocks for the P-CV investor ranges from -667% (November 2008) to 535% (September 1998) compared with the NP-CV weight of 76%. The strong in-sample predictability evidence for VRP thus produces substantial shifts in optimal weights and an economically large CER difference of 4.37% for VRP.

The optimal weights for the three-month P-CV investor are muted relative to the one-month results. The weight in stocks varies between -272% and 288% during the sample period (the NP-CV optimal weight is 75%). Further, the absolute magnitude of the weight difference is less than 20% for over half of months, reflecting the tendency of the highly kurtotic VRP to have many moderate periods with a few large spikes. The most important impact of the longer horizon for this investor is a lack of perceived persistence in VRP, reflected by the relatively low value of β_x shown in Figure 6. The median posterior draw of β_x is only 0.28, such that any variation in expected return is seen to be quite short-lived and primarily affects only the first month of the holding period. The more moderate positions taken by the three-month P-CV investor are reflected in the relatively low CER difference of 1.00% (compared with 4.37% for the one-month investor).

The results in Panel B of Figure 8 for the stochastic-volatility cases also show large horizon effects. The weight difference for the three-month P-SV investor is noticeably devoid of large spikes, although there is some variation in optimal weights that produces a CER difference of 1.31%. The mechanism at work in the stochastic-volatility case is somewhat different from that in the constant-volatility results. After considering stochastic volatility in returns and the state variable, the investor perceives VRP to be more persistent with a posterior median β_x of 0.63. An unusually high expected market return at the one-month horizon is thus accompanied by relatively high expected returns for the last two months of the holding period. However, the optimal portfolio weights are still relatively moderate for the three-month P-SV investor because of effects of horizon on beliefs about risk. In particular, investing for multi-period horizons becomes very risky from the perspective of the Bayesian investor when the predictor variable takes extreme values and is highly volatile. The investor thus moderates her positions and avoids taking large bets.

Figure 9 illustrates the effects of horizon on the riskiness of stocks given the NP-CV, P-CV, NP-SV, and P-SV models with the VRP predictor. For each model, we plot the three-month predictive variance ratio, which is calculated as the variance of the three-month predictive return distribution divided by three times the variance of the one-month predictive return distribution. As shown by Pástor and Stambaugh (2012) and Avramov, Cederburg, and Lučivjanská (2018), predictive variance ratios can be pushed above one by uncertainty about future expected return and estimation risk, whereas mean reversion can have a negative effect on predictive variance ratios. The predictive variance ratios for VRP are nearly uniformly greater than one across models, suggesting that uncertainty about future expected return and estimation risk are outweighing any effect of mean reversion. The predictive variance ratios for the NP-CV and P-CV cases are relatively stable and only slightly higher than one, such that these investors perceive roughly the same per-period risk over a three-month horizon compared with a one-month horizon.

The more interesting effects in Figure 9 are for the stochastic-volatility models. The predictive variance ratio for the NP-SV model exceeds that of the constant-volatility cases because of the well-known effect that stochastic volatility in returns produces fatter tails and higher variance in multi-period returns. Relative to the NP-SV model, the P-SV model produces several large spikes in the predictive variance ratio. Each of these spikes corresponds to a relatively high level of stochastic volatility in VRP observable in Figure 7. In these cases, the expected market returns in months two and three of the holding period are particularly uncertain given the high conditional volatility of the predictor variable. Uncertainty about future expected return is quite high in these times, which can have a substantial impact on the riskiness of stocks from the investor's perspective. The largest predictive variance ratio spike reaches over three in November 2008, which means that the predictive return variance at a three-month horizon exceeded nine times the predictive variance of one-month returns.[8] The three-month investor is understandably cautious with her portfolio decisions in this period despite the relatively large negative expected three-month return.

Discussion

The VRP predictor results show that multi-period horizons can impact the economic value of a predictor. Among the new forecasting variables, we also see large declines in economic value for the NOS and OIL variables. The horizon effects for these predictors are caused both by investor beliefs about the persistence of expected return and by additional risk from the perspective of the investor. The impact on expected return in future periods is relatively obvious and easy to diagnose with an autoregression coefficient. Less persistent variables will be less valuable to multi-period investors because of the low persistence of expected return. The risk channel is somewhat more subtle, since it arises from the investor's beliefs that a predictor and the implied conditional mean return could swing wildly from month to month. Predictor variables with low persistence that also display substantial time variation in stochastic volatility will tend to produce the largest effects on perceived risk over multi-period horizons. This type of variable will tend to have high kurtosis because of these time-series dynamics (e.g., the kurtosis of VRP is 56.17 as reported in Table I), which can be used to diagnose whether the effects of horizon on risk are likely to be important.

5 Conclusion

We evaluate the economic value of stock market return predictors from the perspective of Bayesian investors while accounting for realistic features of the data and investors. In the classic one-month, constant-volatility setting, OLS R^2 from a predictive return regression is a strong indicator of economic value, and even weak statistical evidence can produce non-trivial economic gains. With a one-month horizon, we show that stochastic volatility in returns has a large impact on the economic value of many predictors through both statistical and portfolio choice channels.

[8]Note that higher-order moments are also affected by the same channels. For example, the kurtosis of the three-month predictive return distribution for the P-SV model is 28.78 compared with kurtosis of one-month returns of 3.54 in November 2008.

In particular, many of the predictors correlate to stock market volatility in systematic ways that affect the amount and usefulness of information about expected market return. In multi-period settings, the persistence level and stochastic volatility of predictor variables have strong impacts on expected return and risk from the investor's perspective. Many of these effects are strongest among the newer predictors that occasionally take on extreme values contemporaneous to high return volatility. Overall, we complement Kandel and Stambaugh's (1996) result that even weak statistical evidence of predictability is economically important by showing that there are several predictors for which strong statistical findings correspond with relatively little economic value.

References

Avramov, Doron, 2002, Stock return predictability and model uncertainty, *Journal of Financial Economics* 64, 423–458.

——— , 2004, Stock return predictability and asset pricing models, *Review of Financial Studies* 17, 699–738.

Avramov, Doron, Scott Cederburg, and Katarína Lučivjanská, 2018, Are stocks riskier over the long run? Taking cues from economic theory, *Review of Financial Studies* 31, 556–594.

Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225–264.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.

Bollerslev, Tim, Viktor Todorov, and Lai Xu, 2015, Tail risk premia and return predictability, *Journal of Financial Economics* 118, 113–164.

Campbell, John Y., and Luis M. Viceira, 2002, *Strategic asset allocation: Portfolio choice for long-term investors*. Oxford University Press, Oxford, UK.

Campbell, John Y., and Tuomo Vuolteenaho, 2004, Bad beta, good beta, *American Economic Review* 94, 1249–1275.

Carvalho, Carlos M., Hedibert F. Lopes, and Robert E. McCulloch, 2018, On the long-run volatility of stocks, *Journal of the American Statistical Association* 113, 1050–1069.

Cooper, Ilan, and Richard Priestly, 2009, Time-varying risk premiums and the output gap, *Review of Financial Studies* 22, 2801–2833.

Cremers, K. J. Martijn, 2003, Stock return predictability: A Bayesian model selection perspective, *Review of Financial Studies* 15, 1223–1249.

Del Negro, Marco, and Giorgio E. Primiceri, 2015, Time varying structural vector autoregressions and monetary policy: A corrigendum, *Review of Economic Studies* 82, 1342–1345.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it?, *Review of Financial Studies* 24, 1–45.

Driesprong, Gerben, Ben Jacobsen, and Benjamin Maat, 2008, Striking oil: Another puzzle?, *Journal of Financial Economics* 89, 307–327.

Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 2001, The economic value of volatility timing, *Journal of Finance* 56, 329–352.

——— , 2003, The economic value of volatility timing using "realized" volatility, *Journal of Financial Economics* 67, 473–509.

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

Hill, Bruce M., 1975, A simple general approach to inference about the tail of a distribution, *Annals of Statistics* 3, 1163–1174.

Hoevenaars, Roy P. P. M., Roderick D. J. Molenaar, Peter C. Schotman, and Tom B. M. Steenkamp, 2014, Strategic asset allocation for long-term investors: Parameter uncertainty and prior information, *Journal of Applied Econometrics* 29, 353–376.

Huang, Darien, and Mete Kilic, 2018, Gold, platinum, and expected stock returns, Forthcoming in *Journal of Financial Economics*.

Johannes, Michael, Arthur Korteweg, and Nicholas Polson, 2014, Sequential learning, predictability, and optimal portfolio returns, *Journal of Finance* 69, 611–644.

Johnson, Travis L., 2018, A fresh look at return predictability using a more efficient estimator, Forthcoming in *Review of Asset Pricing Studies*.

Jones, Christopher S., and Selale Tuzel, 2013, New orders and asset prices, *Review of Financial Studies* 26, 115–157.

Kandel, Shmuel, and Robert F. Stambaugh, 1996, On the predictability of stock returns: An asset allocation perspective, *Journal of Finance* 51, 385–424.

Kelly, Bryan, and Hao Jiang, 2014, Tail risk and asset prices, *Review of Financial Studies* 27, 2841–2871.

Kelly, Bryan, and Seth Pruitt, 2013, Market expectations in the cross-section of present values, *Journal of Finance* 68, 1721–1756.

Kim, Sangjoon, Neil Shephard, and Siddhartha Chib, 1998, Stochastic volatility: Likelihood inference and comparison with ARCH models, *Review of Economic Studies* 65, 361–393.

Klein, Roger W., and Vijay S. Bawa, 1976, The effect of estimation risk on optimal portfolio choice, *Journal of Financial Economics* 3, 215–231.

Koop, Gary, and Dimitris Korobilis, 2010, Bayesian multivariate time series methods for empirical macroeconomics, *Foundations and Trends in Econometrics* 3, 267–358.

Li, Jun, and Jianfeng Yu, 2012, Investor attention, psychological anchors, and stock return predictability, *Journal of Financial Economics* 104, 401–419.

Pástor, Ľuboš, and Robert F. Stambaugh, 2000, Comparing asset pricing models: An investment perspective, *Journal of Financial Economics* 56, 335–381.

———, 2012, Are stocks really less volatile in the long run?, *Journal of Finance* 67, 431–477.

Pettenuzzo, Davide, Allan Timmermann, and Rossen Valkanov, 2014, Forecasting stock returns under economic constraints, *Journal of Financial Economics* 114, 517–553.

Pollett, Joshua M., and Mungo Wilson, 2010, Average correlation and stock market returns, *Journal of Financial Economics* 96, 364–380.

Primiceri, Giorgio E., 2005, Time varying structural vector autoregressions and monetary policy, *Review of Economic Studies* 72, 821–852.

Rapach, David E., Matthew C. Ringgenberg, and Guofu Zhou, 2016, Short interest and aggregate stock returns, *Journal of Financial Economics* 121, 46–65.

Shanken, Jay, and Ane Tamayo, 2012, Payout yield, risk, and mispricing: A Bayesian analysis, *Journal of Financial Economics* 105, 131–152.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Tu, Jun, and Guofu Zhou, 2010, Incorporating economic objectives into Bayesian priors: Portfolio choice under parameter uncertainty, *Review of Financial Studies* 45, 959–986.

Wachter, Jessica A., 2010, Asset allocation, *Annual Review of Financial Economics* 2, 175–206.

Wachter, Jessica A., and Missaka Warusawitharana, 2009, Predictable returns and asset allocation: Should a skeptical investor time the market?, *Journal of Econometrics* 148, 162–178.

——— , 2015, What is the chance that the equity premium varies over time? Evidence from regressions on the dividend-price ratio, *Journal of Econometrics* 186, 74–93.

Table I: Summary statistics for predictor variables.
The table reports summary statistics for stock market return predictor variables. Panel A shows summary statistics for the Goyal-Welch predictors and Panel B displays summary statistics for the new predictors. All predictor variables are monthly such that the summary statistics reflect monthly values.

Predictor	Sample start	Mean	Standard deviation	Skewness	Kurtosis	Autocorrelation
Panel A: Goyal and Welch (2008) predictors						
DP	1927:01	−3.37	0.46	−0.22	2.65	0.99
DY	1927:01	−3.32	0.45	−0.46	2.72	0.99
EP	1927:01	−2.74	0.42	−0.60	5.60	0.99
DE	1927:01	−0.64	0.33	1.51	9.00	0.99
$SVAR$	1927:01	0.00	0.01	5.79	46.45	0.63
BM	1927:01	0.57	0.27	0.78	4.45	0.99
$NTIS$	1927:01	0.02	0.03	1.65	11.20	0.98
TBL	1927:01	0.03	0.03	1.08	4.27	0.99
LTY	1927:01	0.05	0.03	1.08	3.59	1.00
LTR	1927:01	0.00	0.02	0.59	7.66	0.04
TMS	1927:01	0.02	0.01	−0.29	3.16	0.96
DFY	1927:01	0.01	0.01	2.48	11.82	0.98
DFR	1927:01	0.00	0.01	−0.39	10.75	−0.12
$INFL$	1927:01	0.00	0.01	1.08	16.75	0.48
Panel B: New predictors						
PLS	1927:01	−0.71	0.34	−4.68	32.48	0.96
GAP	1947:12	0.00	0.07	−0.03	1.98	0.99
NOS	1958:02	0.01	0.04	−0.01	4.70	0.66
DOW	1960:01	0.90	0.10	−1.11	3.91	0.94
$TAIL$	1963:01	0.00	1.00	−0.49	2.78	0.82
COR	1963:03	0.26	0.11	0.86	4.34	0.90
SII	1973:01	0.00	0.25	0.38	2.96	0.97
GP	1975:01	−0.20	0.28	−0.57	2.50	0.99
OIL	1983:04	0.00	0.09	−0.21	5.29	0.17
VRP	1990:01	16.20	20.40	−3.70	56.17	0.28
LJV	1996:06	0.00	0.00	2.94	14.19	0.96

Table II: Predictive regression coefficients from OLS.

The table reports predictive regression coefficients, associated p-values against the null of no predictability, and regression R^2s from OLS regressions. Panel A shows predictive regression coefficients for the Goyal-Welch predictors and Panel B displays coefficients for the new predictors. The predictive regressions are monthly log stock market excess returns on lagged values of the predictor.

Predictor	Coefficient	p-value	R^2
Panel A: Goyal and Welch (2008) predictors			
DP	0.005	0.295	0.002
DY	0.008	0.026	0.004
EP	0.008	0.063	0.003
DE	−0.002	0.756	0.000
$SVAR$	−0.385	0.536	0.002
BM	0.013	0.203	0.004
$NTIS$	−0.158	0.077	0.006
TBL	−0.093	0.085	0.003
LTY	−0.073	0.190	0.001
LTR	0.113	0.101	0.003
TMS	0.189	0.127	0.002
DFY	0.152	0.786	0.000
DFR	0.188	0.325	0.002
$INFL$	−0.365	0.363	0.001
Panel B: New predictors			
PLS	0.028	0.003	0.031
GAP	−0.091	0.000	0.019
NOS	−0.119	0.009	0.011
DOW	−0.032	0.155	0.005
$TAIL$	0.004	0.018	0.009
COR	0.030	0.089	0.006
SII	−0.015	0.054	0.007
GP	0.019	0.003	0.014
OIL	−0.035	0.200	0.006
VRP	0.046	0.000	0.049
LJV	4.278	0.120	0.016

Table III: CER differences for return predictability.
The table reports CER differences that capture the economic value of return predictability in models with constant volatility or stochastic volatility.

	CER differences			
	One-month horizon		Three-month horizon	
Predictor	Constant volatility	Stochastic volatility	Constant volatility	Stochastic volatility
Panel A: Goyal and Welch (2008) predictors				
DP	0.24	0.39	0.24	0.36
DY	0.53	0.35	0.52	0.30
EP	0.42	0.26	0.41	0.23
DE	0.03	0.02	0.03	0.03
$SVAR$	0.21	0.23	0.10	0.11
BM	0.51	0.23	0.49	0.22
$NTIS$	0.69	0.52	0.65	0.40
TBL	0.35	1.18	0.34	1.01
LTY	0.17	0.83	0.17	0.73
LTR	0.31	0.82	0.04	0.12
TMS	0.25	0.19	0.23	0.16
DFY	0.05	0.06	0.04	0.06
DFR	0.27	0.01	0.03	0.02
$INFL$	0.16	1.34	0.05	0.40
Panel B: New predictors				
PLS	3.38	0.53	2.73	0.55
GAP	2.34	2.10	2.27	1.68
NOS	1.29	1.50	0.61	0.64
DOW	0.55	0.58	0.48	0.43
$TAIL$	1.03	0.52	0.68	0.33
COR	0.72	0.87	0.57	0.63
SII	0.81	0.69	0.76	0.50
GP	1.70	1.31	1.68	0.99
OIL	0.65	1.13	0.10	0.30
VRP	4.37	3.41	1.00	1.31
LJV	1.67	0.88	1.26	0.98

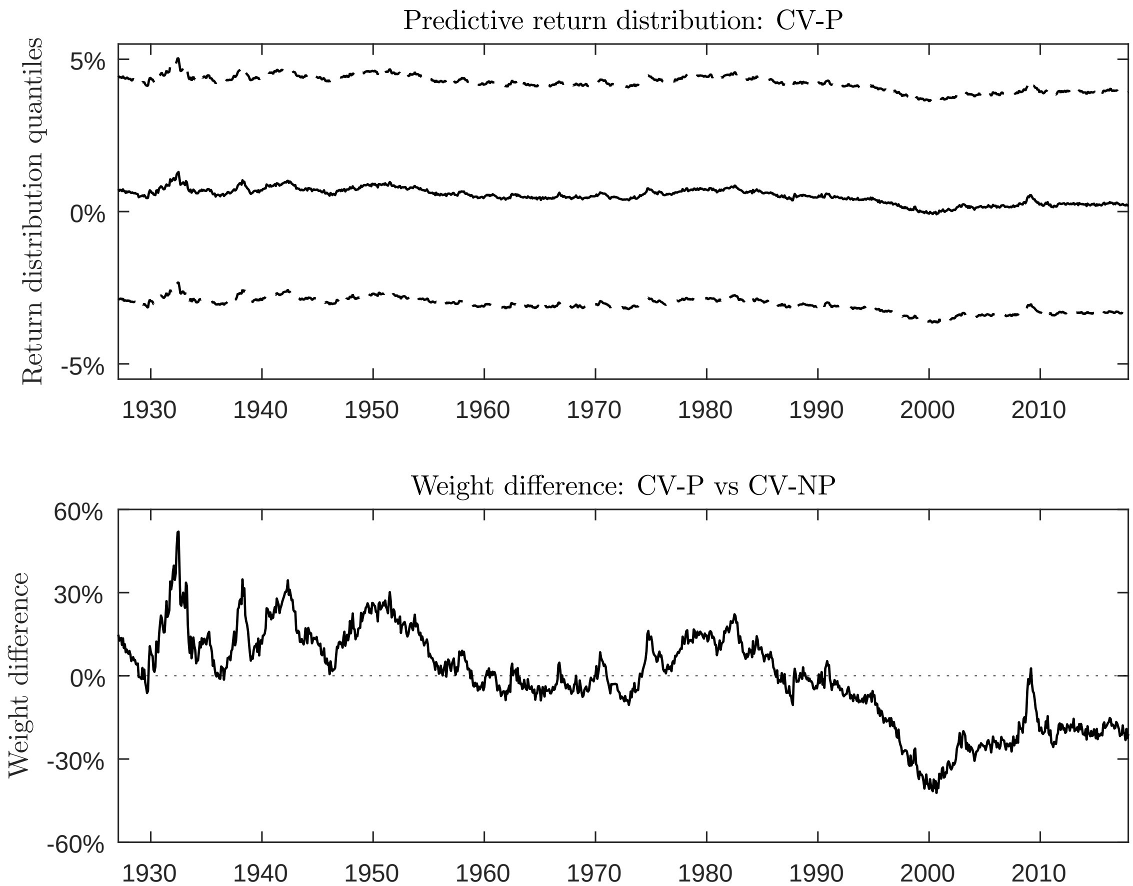

Figure 1: Predictive return distribution quantiles and effects of predictability on portfolio weight for the DP predictor.
The figure shows quantiles of the predictive return distributions and weight differences for the constant-volatility models when DP is the market return predictor variable. The top panel shows quantiles of the predictive return distribution for the P-CV model. The bottom panel shows the differences in optimal portfolio weights between the P-CV and NP-CV models. The solid line in the return distribution represents the median and the dashed lines are the 25th and 75th percentiles of the predictive return distribution in each month. The weight differences represent the effect of including return predictability in the model.

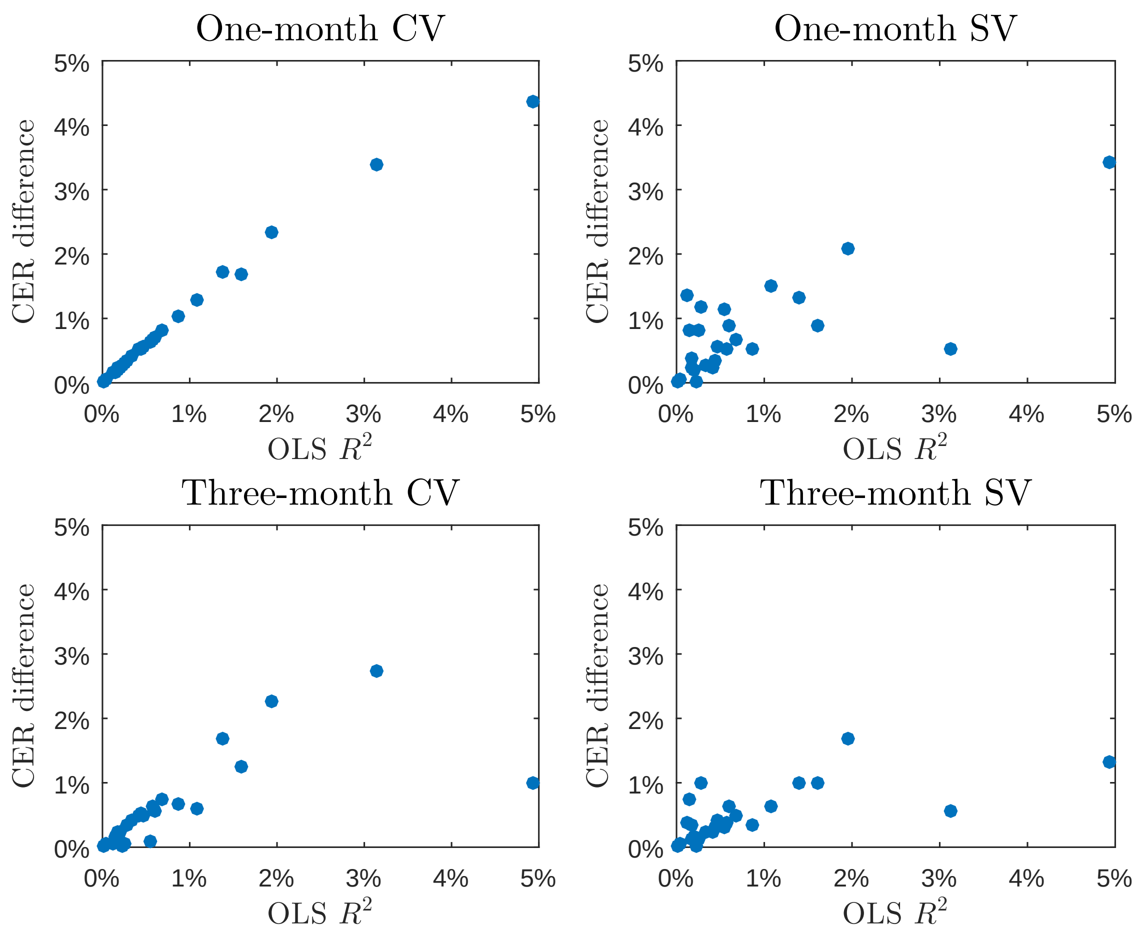

Figure 2: OLS R^2s and CER differences.
The figure show scatter plots of the OLS R^2 from a predictive regression of log market excess returns on a predictor variable and the CER differences that captures the economic value of return predictability in each of four frameworks. The one-month CV CER difference for a predictor is calculated as the difference between the CERs of the P-CV investor for the optimal strategies under the P-CV and NP-CV models. The one-month SV CER differences are calculated analogously with the P-SV and NP-SV models, and the three-month CER differences use three-month predictive return distributions and optimal weights in the constant-volatility and stochastic-volatility cases. The predictive R^2s are from monthly regressions. The CER differences are annualized by multiplying the monthly CER differences by 12.

Panel A: Posteriors of predictive regression coefficients for Goyal-Welch predictors

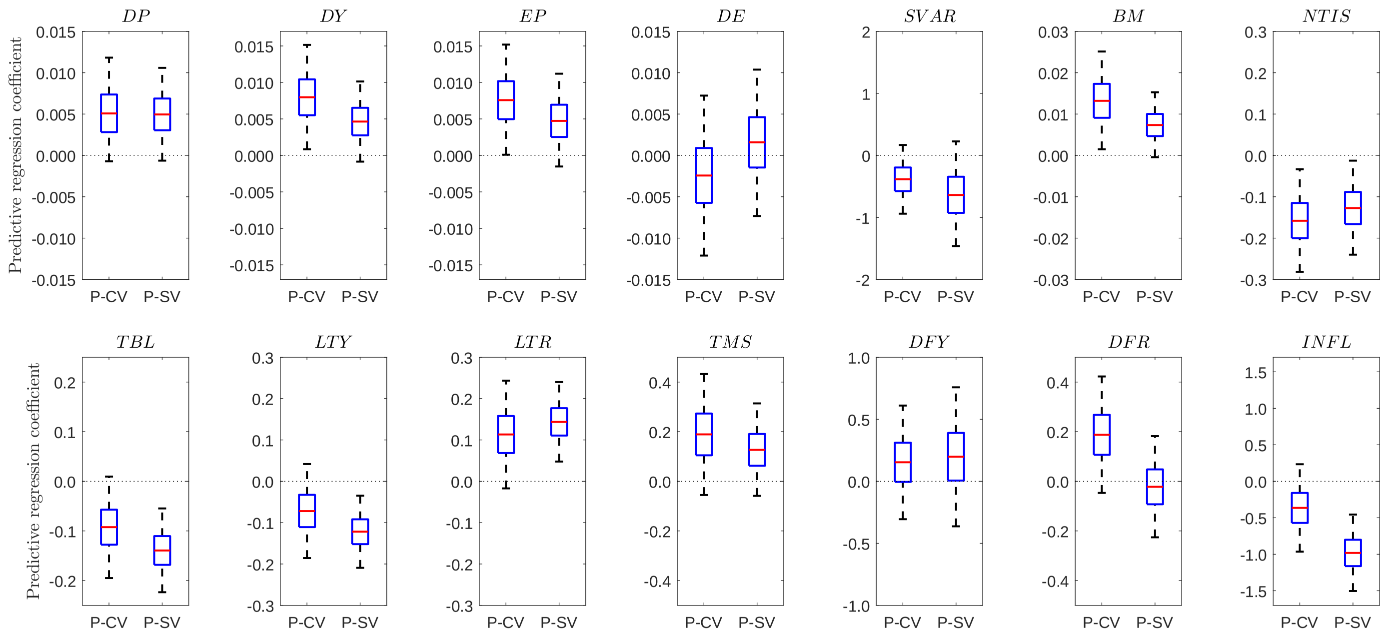

Figure 3: Posteriors of predictive regression coefficients in the constant-volatility and stochastic-volatility models. The figure shows a box-and-whiskers plot for the posterior distribution of the predictive regression coefficient for each model and predictor combination. Panel A shows posteriors for the Goyal-Welch predictors and Panel B shows posteriors for the new predictors. Each predictor has posteriors for the P-CV and the P-SV models. In each box-and-whiskers plot, the red line shows the posterior median, the box represents a 50% credible interval, and the whiskers span a 95% credible interval.

Panel B: Posteriors of predictive regression coefficients for new predictors

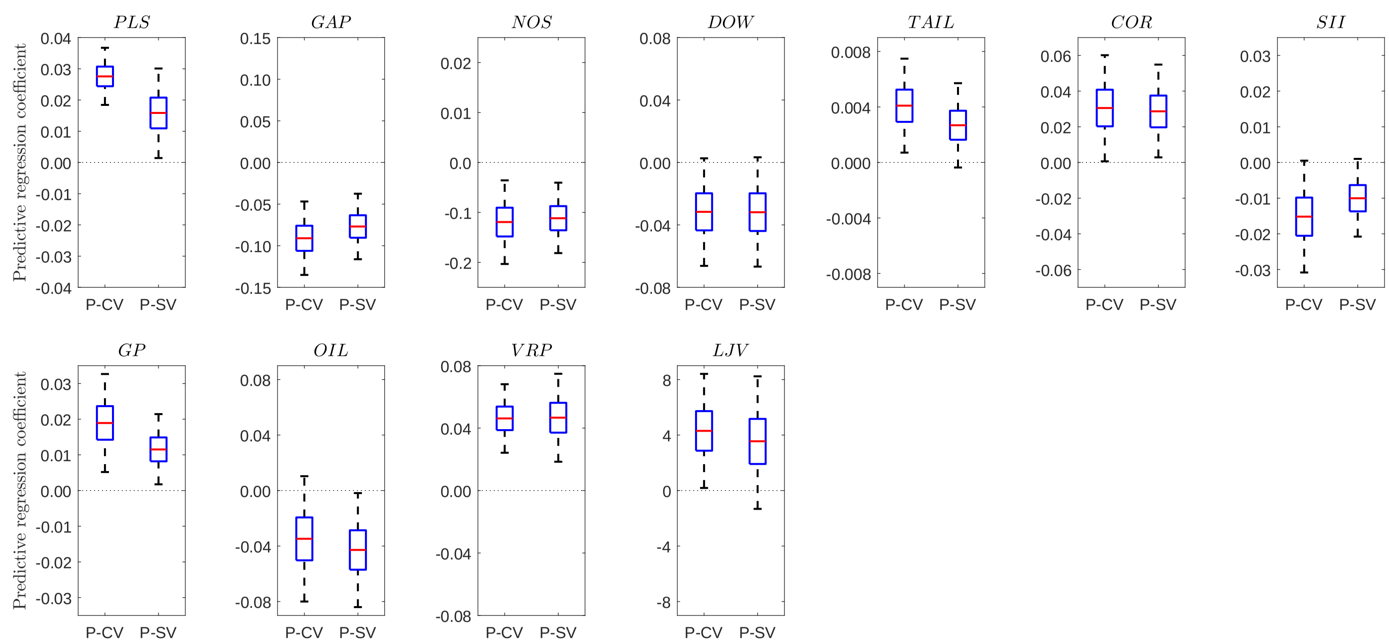

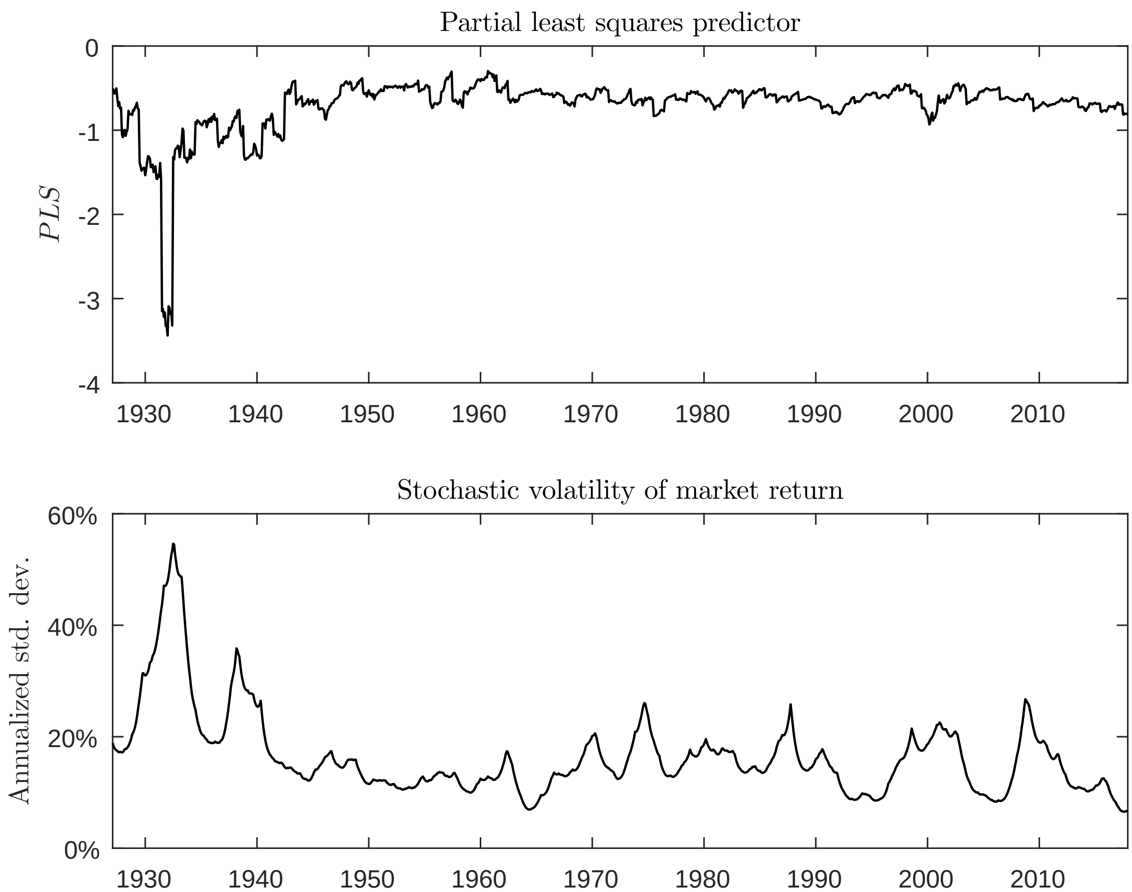

Figure 4: ***PLS* predictor variable and stochastic volatility process.**
The figure shows the *PLS* predictor variable in the top panel and the posterior mean of the annualized standard deviation of the stock market return from the stochastic volatility process in the bottom panel. The monthly standard deviation is annualized by multiplying by $\sqrt{12}$.

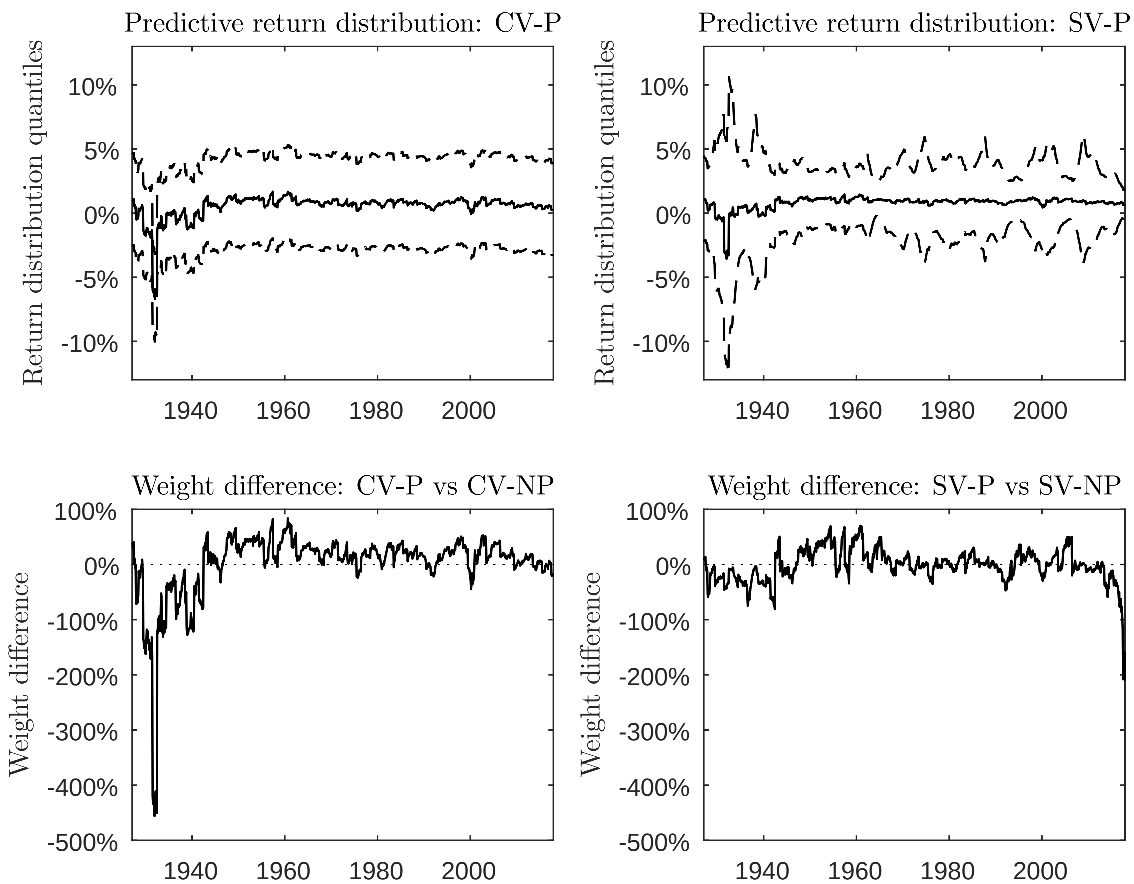

Figure 5: Predictive return distribution quantiles and effects of predictability on portfolio weight for the PLS predictor.
The figure shows quantiles of the predictive return distributions for the P-CV and P-SV models in the top panels and the differences in optimal portfolio weights between the P-CV and NP-CV models and the P-SV and NP-SV models in the bottom panels when PLS is the market return predictor variable. The solid line in the return distribution represents the median and the dashed lines are the 25th and 75th percentiles of the predictive return distribution in each month. The weight differences represent the effect of including return predictability in the model for the constant volatility and stochastic volatility cases.

Panel A: Posteriors of autoregression coefficients for Goyal-Welch predictors

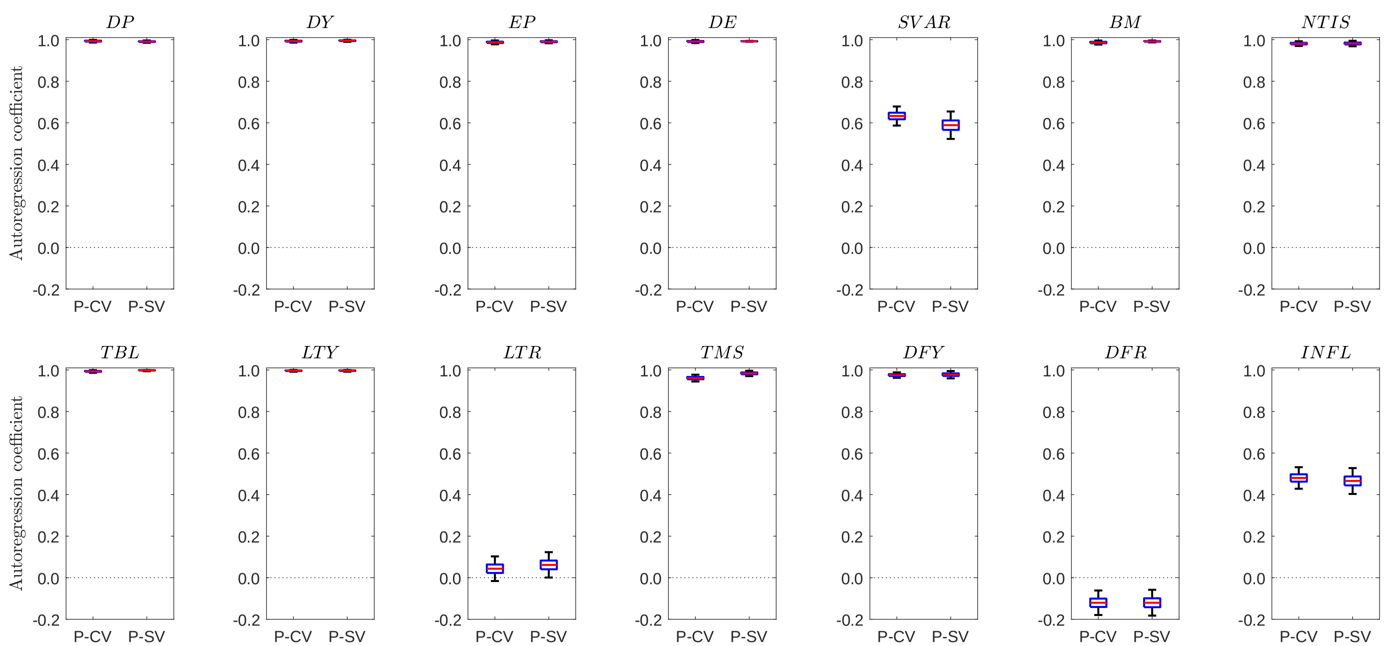

Figure 6: Posteriors of autoregression coefficients in the constant-volatility and stochastic-volatility models.
The figure shows a box-and-whiskers plot for the posterior distribution of the autoregression coefficient for each model and predictor combination. Panel A shows posteriors for the Goyal-Welch predictors and Panel B shows posteriors for the new predictors. Each predictor has posteriors for the P-CV and the P-SV models. In each box-and-whiskers plot, the red line shows the posterior median, the box represents a 50% credible interval, and the whiskers span a 95% credible interval.

Panel B: Posteriors of autoregression coefficients for new predictors

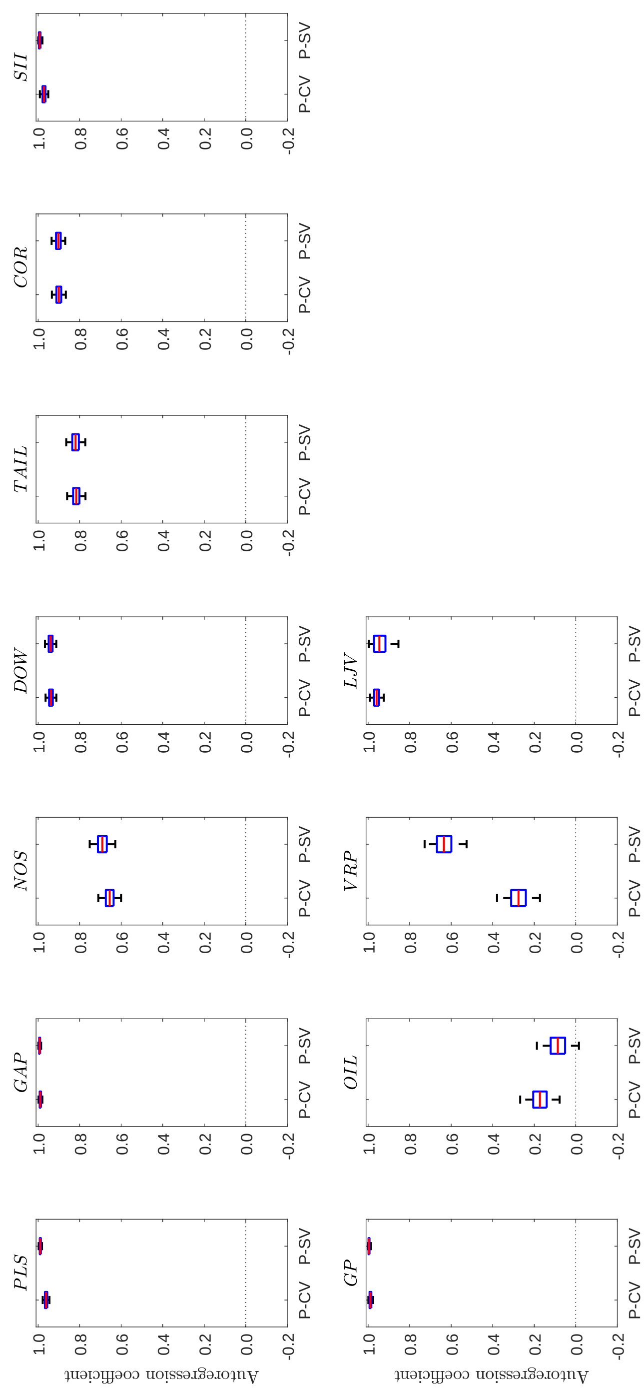

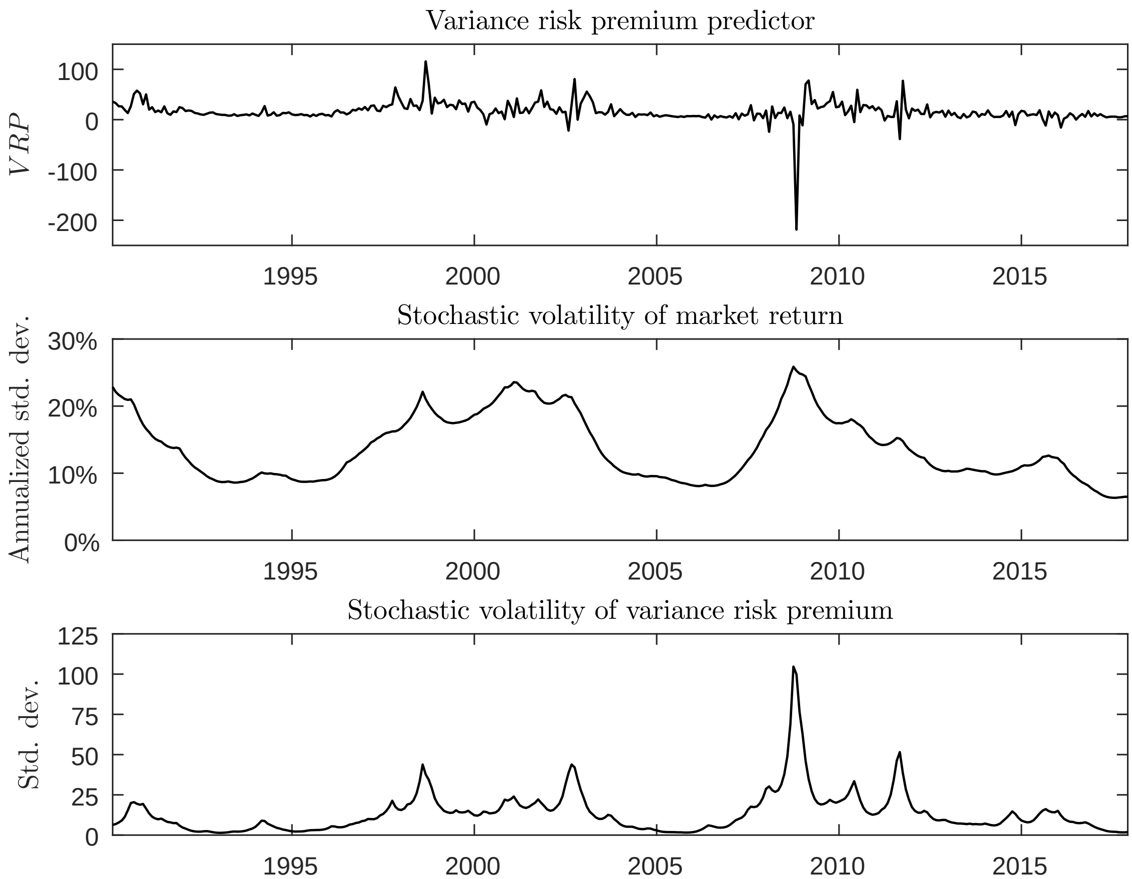

Figure 7: VRP predictor variable and stochastic volatility processes.
The figure shows the VRP predictor variable in the top panel, the posterior mean of the annualized standard deviation of the stock market return from the stochastic volatility process in the middle panel, and the posterior mean of the standard deviation of the VRP predictor variable from the stochastic volatility process in the bottom panel. The monthly standard deviation of market returns is annualized by multiplying by $\sqrt{12}$.

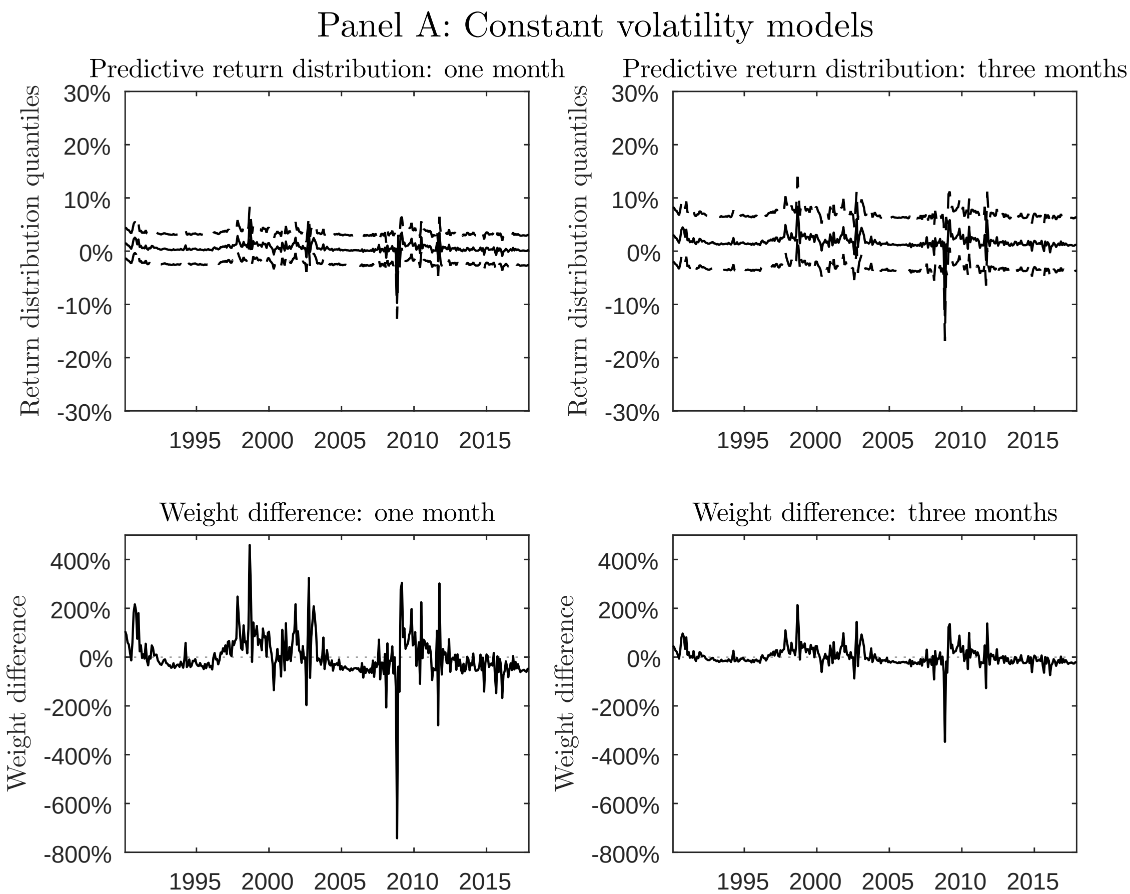

Figure 8: Predictive return distribution quantiles and effects of predictability on portfolio weight for the VRP predictor.
The figure shows quantiles of the predictive return distributions and weight differences for the constant-volatility models in Panel A and the stochastic-volatility models in Panel B when VRP is the market return predictor variable. In Panel A (Panel B), the top panels show quantiles of the predictive return distribution for the P-CV (P-SV) model at the one-month and three-month horizons. The bottom panels of Panel A (Panel B) show the differences in optimal portfolio weights between the P-CV and NP-CV (P-SV and NP-SV) models at the one-month and three-month horizons. The solid line in the return distribution represents the median and the dashed lines are the 25th and 75th percentiles of the predictive return distribution in each month. The weight differences represent the effect of including return predictability in the model for each case.

Panel B: Stochastic volatility models

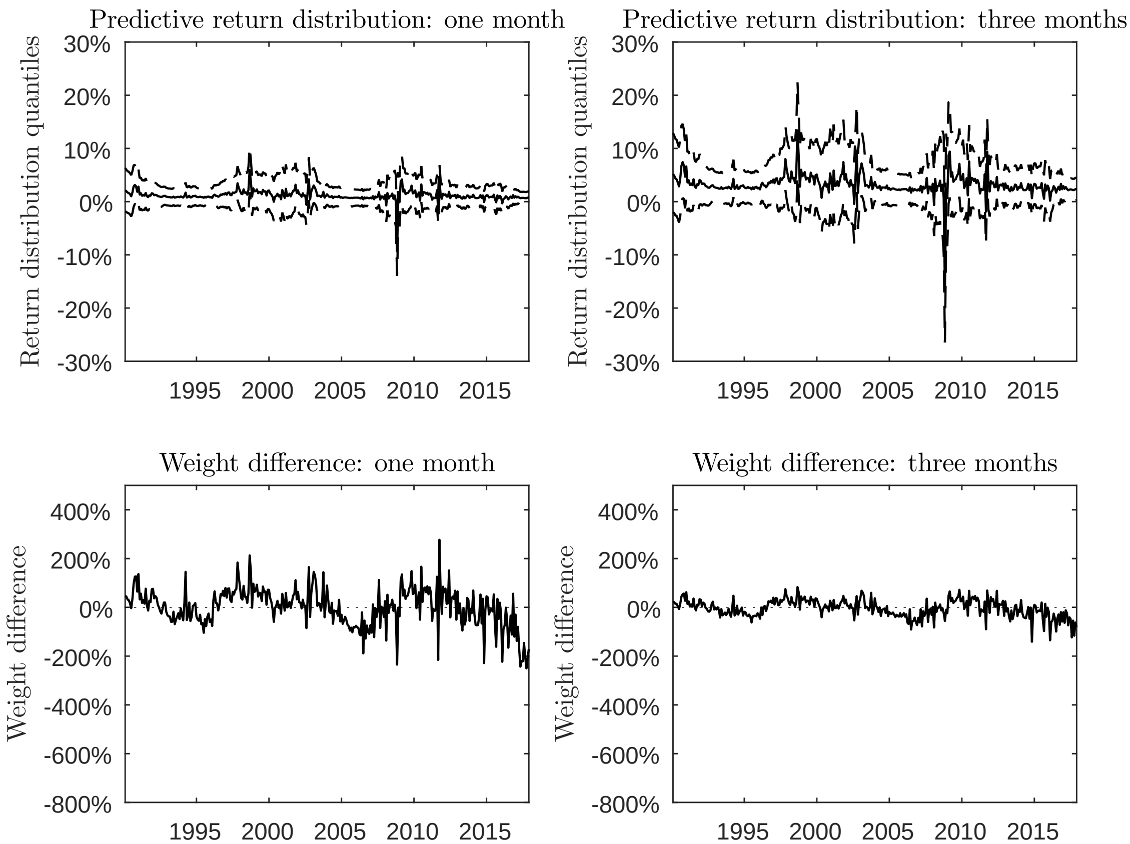

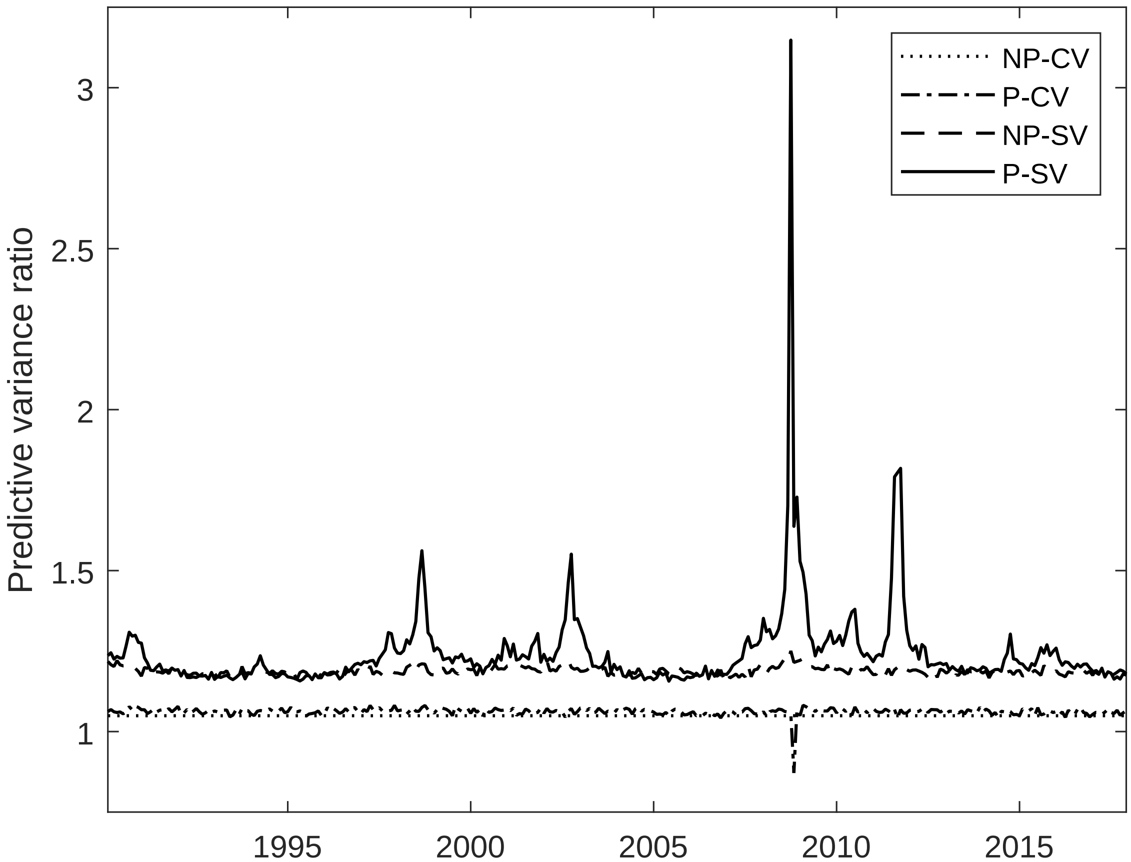

Figure 9: Predictive variance ratio for the VRP predictor.
The figure shows the predictive variance ratio for the NP-CV, P-CV, NP-SV, and P-SV models when VRP is the market return predictor variable. The predictive variance ratio is calculated as the variance of the three-month predictive return distribution divided by three times the variance of the one-month predictive return distribution.

A Estimation Appendix

We estimate the NP-CV, P-CV, NP-SV, and P-SV models using Markov chain Monte Carlo (MCMC) techniques. For each model, we draw 110,000 sets of parameters from the posterior distribution, and we discard the first 10,000 draws as a burn-in period to produce the 100,000 posterior draws that we use in the analysis. This appendix describes the MCMC chain and provides information about prior parameters.

The $(T \times 1)$ vector of log excess returns is denoted by r. The predictor variable x_t has $T+1$ observations, and we refer to the $(T \times 1)$ vector of the first T observations as x_l to denote its use as a lagged predictor variable and the $(T \times 1)$ vector of the last T observations as x. Let the matrix of regressors for the P-CV and P-SV models be $Z_P = [\iota_T \otimes I_2 \quad x_l \otimes I_2]$, where ι_T is a $(T \times 1)$ vector of ones and I_2 is a (2×2) identity matrix, such that Z_P is a $(2T \times 4)$ matrix. Further let the corresponding matrix for the NP-CV and NP-SV models be $Z_{NP} = \left[\iota_T \otimes I_2 \quad x_l \otimes \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right]$, such that Z_{NP} is a $(2T \times 3)$ matrix. Define a $(2T \times 1)$ vector y by stacking the (2×1) vectors $\begin{bmatrix} r_t \\ x_t \end{bmatrix}$ for $t = 1, 2, \ldots, T$. For the models with constant volatility (NP-CV and P-CV), we define

$$\tilde{\Sigma}_{NP} = \begin{bmatrix} \Sigma & 0_2 & \cdots & 0_2 \\ 0_2 & \Sigma & & 0_2 \\ \vdots & & \ddots & \vdots \\ 0_2 & \cdots & \cdots & \Sigma \end{bmatrix}, \tag{A1}$$

where 0_2 denotes a (2×2) matrix of zeros. Similarly, for the stochastic volatility models (NP-SV and P-SV), we define

$$\tilde{\Sigma}_{P} = \begin{bmatrix} \Sigma_1 & 0_2 & \cdots & 0_2 \\ 0_2 & \Sigma_2 & & 0_2 \\ \vdots & & \ddots & \vdots \\ 0_2 & \cdots & \cdots & \Sigma_T \end{bmatrix}. \tag{A2}$$

Given these definitions, the VAR in equations (6) and (7) can be restated as

$$y = Zb + \epsilon, \qquad \epsilon \sim N(0, \tilde{\Sigma}), \tag{A3}$$

inserting the appropriate Z and $\tilde{\Sigma}$ that correspond to the given model. In this setup, b is a (3×1) vector for the no-predictability models with Z_{NP} and a (4×1) vector for the predictability models with Z_P.

We draw the regression coefficients b from a normal distribution conditional on the draw of $\tilde{\Sigma}$. Specifically,

$$b \sim N\left(\overline{H}^{-1}(Z'\tilde{\Sigma}^{-1}y + \underline{H}\underline{b}), \overline{H}^{-1}\right), \tag{A4}$$

where $\overline{H} = (Z'\tilde{\Sigma}^{-1}Z + \underline{H})$ and the constraint $-1 < \beta_x < 1$ ensures that the x_t process is stationary. We specify a diffuse prior for b with $\underline{b} = 0_k$ and $\underline{H} = 10^{-6} I_k$, where k equals three for the no-predictability models and four for the predictability models. We use an accept-reject algorithm to produce draws of b such that $-1 < \beta_x < 1$.

We draw Σ for the constant volatility models and Σ_t for the stochastic volatility models conditional on the draw of b. For the NP-CV and P-CV models, we specify a conjugate normal-inverse-Wishart prior. Conditional on b, we calculate the $(T \times 1)$ vectors of residuals ϵ^r and ϵ^x according to equations (6) and (7), and we define $\epsilon = \begin{bmatrix} \epsilon^r & \epsilon^x \end{bmatrix}$ to be the $(T \times 2)$ matrix of residuals. The draw of Σ is then given by

$$\Sigma \sim IW\left(\underline{V} + \epsilon'\epsilon, \underline{v} + T\right). \tag{A5}$$

To minimize the influence of the prior, we specify an empirical Bayes prior with $\underline{v} = k + 3$ and $\underline{V} = \underline{v}\hat{\Sigma}$, where k is equal to three for NP-CV and four for P-CV and $\hat{\Sigma}$ is the covariance matrix of the OLS residuals for equations (6) and (7).

We draw from the stochastic volatility processes for the NP-SV and P-SV models following Primiceri (2005), who generally uses the approach of Kim, Shephard, and Chib (1998). See Primiceri (2005), Koop and Korobilis (2010), and Del Negro and Primiceri (2015) for details.[9] Our prior parameters are specified as follows. The initial states a_0, $\log(\sigma_{r,0})$, and $\log(\tilde{\sigma}_{x,0})$ are independently normally distributed with mean zero and variance of four. For the W and S parameters (see Primiceri (2005) for definitions), we specify prior means of $0.01^2\iota_2$ for W and 0 for S with prior variances of two for the diagonal elements of W and two for S. We check that these scales for the priors are consistent with the scales of our variables. Following Del Negro and Primiceri (2015), we iterate through the posterior draws by drawing the b, a, and σ parameters in order conditional on the other parameters.

B Data Appendix

This section describes the predictor variables used in our empirical tests and notes the relevant data sources. Some of the results in the paper are based on a monthly forecast horizon. As such, we focus on predictor variables that are available at a monthly frequency. We also require that a given forecasting variable has data availability through December 2017. The first group of 14 predictors consists of updated versions of the variables considered in Goyal and Welch (2008).[10] The sample period for each of these variables is from January 1927 to December 2017.

1. **Dividend-price ratio** (DP). The difference between the log of dividends paid on the S&P 500 index over the prior 12 months and the log of the index level.

2. **Dividend yield** (DY). The difference between the log of dividends paid on the S&P 500 index over the prior 12 months and the log of the lagged index level.

3. **Earnings-price ratio** (EP). The difference between the log of earnings on the S&P 500 index over the prior 12 months and the log of the index level.

4. **Dividend-earnings ratio** (DE). The difference between the log of dividends paid on the S&P 500 index over the prior 12 months and the log of earnings on the S&P 500 index over the prior 12 months.

5. **Stock variance** ($SVAR$). The sum of squared daily returns on the S&P 500 index over the prior month.

6. **Book-to-market ratio** (BM). The ratio of book value to market value for the Dow Jones Industrial Average.

7. **Net equity expansion** ($NTIS$). The ratio of the sum of net issues by NYSE stocks over the prior 12 months to the total market capitalization of NYSE stocks.

8. **Treasury bill yield** (TBL). The yield on a three-month Treasury security.

[9]We use functions from Dimitris Korobilis to draw a_t and σ_t, and we thank him for sharing this code.

[10]Data on these predictors are from Amit Goyal's website at http://www.hec.unil.ch/agoyal/. We thank Amit Goyal for making these data available.

9. **Long-term Treasury bond yield** (LTY). The yield on a long-term Treasury security.

10. **Long-term Treasury bond return** (LTR). The return on a long-term Treasury security over the prior month.

11. **Term spread** (TMS). The difference between the yield on a long-term Treasury security and the yield on a three-month Treasury security.

12. **Default yield spread** (DFY). The difference between the yield on BAA-rated corporate bonds and the yield on AAA-rated corporate bonds.

13. **Default return spread** (DFR). The difference between the return on long-term corporate bonds and the return on long-term Treasury bonds over the prior month.

14. **Inflation** ($INFL$). The monthly growth rate in the Consumer Price Index (CPI) for all urban consumers. The inflation variable is lagged by one month to account for the reporting lag in the CPI data.

We supplement the Goyal and Welch (2008) predictors with a second group of forecasting variables from more recent literature. We search articles published in the *Journal of Finance*, *Journal of Financial Economics*, and *Review of Financial Studies* subsequent to Goyal and Welch (2008) and include 11 additional variables that satisfy our screen of being available at a monthly frequency through December 2017. These predictors are summarized below.

15. **Partial least squares aggregated book-to-market ratio** (PLS). Following Kelly and Pruitt (2013), PLS is extracted from the cross section of portfolio-level book-to-market ratios using a two-step procedure. We collect data on returns and book-to-market ratios for 100 portfolios formed on size and book-to-market from Kenneth French's website. We also collect return data for the CRSP value-weighted index from the same source. We convert all portfolio returns to log returns and compute the log book-to-market ratio for each portfolio at the end of June of each year (i.e., at the portfolio formation date). Following Campbell and Vuolteenaho (2004), each portfolio's log book-to-market ratio from July to May is constructed by subtracting the portfolio's cumulative log return (from the previous June) from its end-of-June log book-to-market ratio. In the first step of the PLS procedure, we run a time-series regression for each portfolio i of log book-to-market ratio on future log market return:

$$\log(BM_{i,t}) = \phi_{i,0} + \phi_{i,1} r_{t+1} + e_{i,t}. \tag{B1}$$

In the second step, we estimate a cross-sectional regression for each month t of book-to-market ratio on the corresponding first-stage loading:

$$\log(BM_{i,t}) = c_t + F_t \hat{\phi}_{i,1} + w_{i,t}, \tag{B2}$$

and PLS_t is the estimate of the latent factor F_t. The PLS series starts in January 1927.

16. **Output gap** (GAP). Following Cooper and Priestly (2009), GAP is the trend deviation of the log of industrial production. We obtain monthly data on the Industrial Production Index from the Federal Reserve Bank of St. Louis website.[11] The industrial production data are lagged by one month to account for the publishing lag from the Federal Reserve. We estimate GAP from the following regression:

$$y_t = a + bt + ct^2 + \nu_t, \tag{B3}$$

[11] See https://fred.stlouisfed.org/.

where t indexes sample period months, y_t is the log of industrial production, and $GAP_t = \nu_t$ is the regression residual. The GAP series starts in December 1947.

17. **New orders-to-shipments of durable goods** (NOS). Following Jones and Tuzel (2013), NOS is the difference between the log of new orders of durable goods and the log of shipments of durable goods. The data on new orders and shipments are from the Census Bureau's Survey of Manufacturers' Shipments, Inventories, and Orders.[12] We use data based on Standard Industrial Classification (SIC) codes through February 1992 and data based on North American Industry Classification (NAICS) codes starting in March 1992. We adjust the NAICS shipments data by subtracting shipments of semiconductors. The data are lagged by one month to account for reporting delay. The NOS series starts in February 1958.

18. **Nearness to Dow historical high** (DOW). Following Li and Yu (2012), DOW is the ratio of the current level of the Dow Jones Industrial Average (DJIA) index to its historical high. The calculation is based on daily, end-of-day prices, and DOW for a given month is the value at the last trading day of that month. The DJIA data are from Thomson Reuters Eikon. The DOW series starts in January 1960.

19. **Tail risk** ($TAIL$). Following Kelly and Jiang (2014), we estimate $TAIL$ using Hill's (1975) power law estimator. The underlying data are daily stock returns from the Center for Research in Security Prices (CRSP) database. The sample includes all NYSE, Amex, and NASDAQ stocks with share codes 10 and 11. For a given month t, we use the pooled cross section of daily returns to compute the common time-varying component of return tails as

$$TAIL_t = \frac{1}{K_t} \sum_{k=1}^{K_t} \log\left(\frac{R_{k,t}}{u_t}\right), \tag{B4}$$

where $R_{k,t}$ is the kth daily return that falls below the extreme value threshold for month t (u_t), and K_t is the number of such events in month t. The threshold parameter, u_t, is the fifth percentile of the pooled cross-sectional distribution of daily returns in month t. The $TAIL$ measure is standardized to have a mean of zero and a variance of one. The $TAIL$ series starts in January 1963.

20. **Average correlation** (COR). Following Pollett and Wilson (2010), COR is the value-weighted average correlation for the largest 500 stocks by market capitalization in the CRSP universe at the end of a given month. The sample includes all NYSE, Amex, and NASDAQ stocks with share codes 10 and 11. We compute the correlation between each pair of stocks i and j at the end of month t, $\hat{\rho}_{ij,t}$, using the prior three months of daily returns. The average correlation measure is

$$COR_t = \sum_{i=1}^{500} \sum_{j \neq i} w_{i,t} w_{j,t} \hat{\rho}_{ij,t}, \tag{B5}$$

where $w_{i,t}$ and $w_{j,t}$ are the market capitalizations of stocks i and j divided by the sum of the market capitalizations of the largest 500 stocks at the end of month t. The COR series starts in March 1963.

21. **Short interest index** (SII). Following Rapach, Ringgenberg, and Zhou (2016), SII is the trend deviation of the log of equal-weighted short interest. We collect mid-month, firm-level short interest data from Compustat and compute the percentage of shares held short in a given firm by dividing by the firm's number of shares outstanding from CRSP. The sample

[12] See https://www.census.gov/manufacturing/m3/historical_data/index.html.

excludes assets with a share price below \$5 and assets with a market capitalization below the fifth percentile of NYSE market capitalization. Aggregate short interest for month t, $EWSI_t$, is the equal-weighted average percentage short interest across all assets. We estimate SII from the following regression:

$$\log(EWSI_t) = a + bt + \nu_t, \tag{B6}$$

where t indexes sample period months, and $SII_t = \nu_t$ is the regression residual. The SII series starts in January 1973.

22. **Gold-to-platinum price ratio** (GP). Following Huang and Kilic (2018), GP is the difference between the log gold price and the log platinum price. Gold prices are the monthly averages of daily 10:30 A.M. fixing prices in the London Bullion Market. These data are obtained from the Federal Reserve Bank of St. Louis website. Platinum prices from April 1990 to December 2017 are the monthly averages of daily 9:45 A.M. fixing prices in the London Platinum and Palladium Market. These data are obtained from Thomson Reuters Eikon. Monthly platinum prices prior to April 1990 are from Kitco.[13] The GP series starts in January 1975.

23. **Oil price change** (OIL). Following Driesprong, Jacobsen, and Maat (2008), OIL is the monthly change in oil price, measured as a log return. We use end-of-month prices for NYMEX Light Sweet Crude Oil from Thomson Reuters Eikon. The OIL series starts in April 1983.

24. **Variance risk premium** (VRP). Following Bollerslev, Tauchen, and Zhou (2009), VRP is the difference between the squared VIX and the monthly realized variance of the S&P 500 index computed from intraday data. The VRP data are from Hao Zhou's website.[14] The VRP series starts in January 1990.

25. **Left jump tail variation** (LJV). Following Bollerslev, Todorov, and Xu (2015), LJV is an estimate of the left jump tail variation over the coming month. If we let $O_{t,\tau}(k)$ denote the time t price of an out-of-the-money put option on the S&P 500 index with time to expiration τ and log-moneyness k, LJV is computed using

$$LJV_{t,\tau} = \tau \phi_t^- e^{-\alpha_t^- |k_t|} \left(\alpha_t^- k_t \left(\alpha_t^- k_t + 2 \right) + 2 \right) \left(\alpha_t^- \right)^{-3}, \tag{B7}$$

$$\alpha_t^- = \arg\min_{\alpha} \frac{1}{N_t^-} \sum_{i=1}^{N_t} \left| \log \left(\frac{O_{t,\tau}(k_{t,i})}{O_{t,\tau}(k_{t,i-1})} \right) (k_{t,i} - k_{t,i-1})^{-1} - (1+\alpha) \right|, \tag{B8}$$

$$\begin{aligned} \phi_t^- = \arg\min_{\phi} \frac{1}{N_t^-} \sum_{i=1}^{N_t} \Bigg| \log \left(\frac{e^{r_{t,\tau}} O_{t,\tau}(k_{t,i})}{\tau F_{t-,\tau}} \right) \\ - (1+\alpha_t^-) k_{t,i} + \log(\alpha_t^- + 1) + \log(\alpha_t^-) - \log(\phi) \Bigg|, \end{aligned} \tag{B9}$$

where N_t^- denotes the number of puts used in the estimation, with $0 \leq -k_{t,1} < \cdots < -k_{t,N_t^-}$. We follow Bollerslev, Todorov, and Xu (2015) and implement this estimator weekly for all S&P 500 options with times to maturity between 8 and 45 calendar days and meeting liquidity screens. We then average across weeks ending in a calendar month. Bollerslev, Todorov, and

[13] See http://www.kitco.com/charts/historicalplatinum.html.

[14] See https://sites.google.com/site/haozhouspersonalhomepage/. We thank Hao Zhou for making these data available.

Xu's (2015) main predictive regression results use a forecast horizon of six months. We therefore construct LJV as the six-month moving average of the left jump tail variation measure. The LJV series starts in June 1996.

C Derivation Appendix

The three-month predictive return variance is given by an analogue of equation (5),

$$Var(R_{T,T+3}|\mathcal{M}_i, D_T) = E[Var(R_{T,T+3}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T] + Var[E(R_{T,T+3}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T]. \tag{C1}$$

To better understand the components of $E[Var(R_{T,T+3}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T]$, we work with the version of this expression in log excess returns, $E[Var(r_{T,T+3}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T]$. This term can be written as

$$E[Var(r_{T,T+3}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T] = E[Var(r_{T+1} + r_{T+2} + r_{T+3}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T], \tag{C2}$$

and

$$r_{T+1} + r_{T+2} + r_{T+3} = \alpha + \beta x_T + \epsilon^r_{T+1} + \alpha + \beta(\alpha_x + \beta_x x_T + \epsilon^x_{T+1}) + \epsilon^r_{T+2} \tag{C3}$$

$$+ \alpha + \beta(\alpha_x + \beta_x(\alpha_x + \beta_x x_T + \epsilon^x_{T+1}) + \epsilon^x_{T+2}) + \epsilon^r_{T+3}. \tag{C4}$$

The expression for the portion of $r_{T+1} + r_{T+2} + r_{T+3}$ that relates to uncertainty about future expected return is

$$\beta(1 + \beta_x)\epsilon^x_{T+1} + \beta\epsilon^x_{T+2}, \tag{C5}$$

such that the contribution of uncertainty about future expected return in the constant-volatility case is

$$E[Var(\beta(1 + \beta_x)\epsilon^x_{T+1} + \beta\epsilon^x_{T+2}|\mathcal{M}_i, \theta, D_T)|\mathcal{M}_i, D_T] = E[\beta^2((1 + \beta_x)^2 + 1)\sigma^2_x|\mathcal{M}_i, D_T]. \tag{C6}$$

D Robustness Appendix

The results in Table III correspond to an investor with a risk-aversion parameter of $\gamma = 5$. Table D.I presents results for $\gamma = 2$, and Table D.II presents results for $\gamma = 8$.

Table D.I: CER differences for return predictability with $\gamma = 2$.
The table reports CER differences that capture the economic value of return predictability in models with constant volatility or stochastic volatility for an investor with $\gamma = 2$.

	CER differences			
	One-month horizon		Three-month horizon	
Predictor	Constant volatility	Stochastic volatility	Constant volatility	Stochastic volatility
Panel A: Goyal and Welch (2008) predictors				
DP	0.60	0.93	0.59	0.77
DY	1.32	0.83	1.30	0.69
EP	1.04	0.61	1.03	0.50
DE	0.07	0.04	0.07	0.04
$SVAR$	0.51	0.58	0.24	0.24
BM	1.27	0.55	1.23	0.46
$NTIS$	1.72	1.26	1.63	1.06
TBL	0.87	2.88	0.86	2.52
LTY	0.42	2.01	0.42	1.76
LTR	0.78	1.98	0.10	0.23
TMS	0.63	0.44	0.59	0.36
DFY	0.11	0.15	0.10	0.12
DFR	0.68	0.02	0.06	0.01
$INFL$	0.39	3.25	0.13	0.91
Panel B: New predictors				
PLS	8.51	1.32	8.16	1.17
GAP	5.85	5.00	5.67	4.17
NOS	3.22	3.60	1.54	1.57
DOW	1.37	1.40	1.20	0.99
$TAIL$	2.59	1.27	1.72	0.76
COR	1.79	2.09	1.44	1.41
SII	2.02	1.73	1.90	1.49
GP	4.26	3.22	4.20	2.91
OIL	1.62	2.69	0.26	0.34
VRP	12.09	8.34	2.62	3.28
LJV	4.13	2.06	3.17	1.42

Table D.II: CER differences for return predictability with $\gamma = 8$.
The table reports CER differences that capture the economic value of return predictability in models with constant volatility or stochastic volatility for an investor with $\gamma = 8$.

	CER differences			
	One-month horizon		Three-month horizon	
Predictor	Constant volatility	Stochastic volatility	Constant volatility	Stochastic volatility
Panel A: Goyal and Welch (2008) predictors				
DP	0.15	0.24	0.15	0.22
DY	0.33	0.22	0.32	0.18
EP	0.26	0.16	0.25	0.13
DE	0.02	0.01	0.02	0.01
$SVAR$	0.13	0.15	0.06	0.06
BM	0.32	0.14	0.30	0.13
$NTIS$	0.43	0.33	0.40	0.25
TBL	0.22	0.74	0.21	0.64
LTY	0.11	0.52	0.10	0.45
LTR	0.20	0.52	0.02	0.06
TMS	0.16	0.12	0.14	0.09
DFY	0.03	0.04	0.03	0.03
DFR	0.17	0.00	0.02	0.00
$INFL$	0.10	0.84	0.03	0.24
Panel B: New predictors				
PLS	2.11	0.33	1.71	0.31
GAP	1.46	1.32	1.41	1.07
NOS	0.80	0.95	0.38	0.41
DOW	0.34	0.36	0.30	0.27
$TAIL$	0.64	0.33	0.42	0.20
COR	0.45	0.55	0.36	0.39
SII	0.50	0.44	0.47	0.33
GP	1.06	0.83	1.04	0.64
OIL	0.40	0.71	0.06	0.11
VRP	2.71	2.16	0.62	0.74
LJV	1.04	0.55	0.77	0.58