



Equity premium prediction: The role of information from the options market

Antonios K. Alexandridis^a, Iraklis Apergis^b, Ekaterini Panopoulou^c, Nikolaos Voukelatos^{b,*}

^a Department of Accounting and Finance, University of Macedonia, 156 Egnatia St, Thessaloniki 54636, Greece

^b Kent Business School, University of Kent, Sibson, Canterbury CT2 7FS, UK

^c Essex Business School, University of Essex, Wivenhoe Park, Colchester CO43SQ, UK

ARTICLE INFO

JEL classification:

C53
C58
G14
G17

Keywords:

Equity premium
Forecasting
Options
Quantile regression

ABSTRACT

We examine the role of information from the options market in forecasting the equity premium. We provide evidence that the equity premium is predictable out-of-sample using a set of CBOE strategy benchmark indices as predictors. We use a range of econometric approaches to generate point, quantile, and density forecasts of the equity premium. We find that models based on option variables consistently outperform the historical average benchmark. In addition to statistical gains, using option predictors results in substantial economic benefits for a mean-variance investor, delivering up to a fivefold increase in certainty equivalent returns over the benchmark during the 1996–2021 sample period.

1. Introduction

The topic of equity premium predictability has long been of significant interest to academics and practitioners alike. However, the literature has yet to reach a consensus about the optimal set of predictors, and even about the extent to which the equity premium can actually be predicted. In this paper, we contribute to this ongoing debate by examining the role of forward-looking information from the options market in forecasting excess market returns.

A number of studies in the earlier literature had argued that the equity premium is, to a large extent, predictable using a set of financial and economic variables, such as dividend yields, earning-price ratios, book-to-market ratios, term spreads, and default spreads (Campbell and Shiller, 1988; Fama and French, 1988, 1989; Kothari and Shanken, 1997; Pontiff and Schall, 1998; Lettau and Ludvigson, 2001; Cochrane, 2008). However, Goyal and Welch (2008) challenged that commonly-held view and argued that these variables fail to consistently provide accurate predictions over time, with the associated models having an unstable and overall poor forecasting performance in-sample and out-of-sample. The finding that standard economic variables produce forecasts with unstable and short-lived accuracy, particularly when compared to the historical average benchmark, is further supported by Lettau and Van Nieuwerburgh (2008), Timmermann (2008), and Baetje and Menkhoff (2016), among others.

Subsequent studies explored whether alternative predictors can provide consistently more accurate forecasts of the equity premium. For instance, Neely et al. (2014) and Baetje and Menkhoff (2016) examine a set of technical indicators and find that they result in more efficient and stable forecasts compared to standard economic indicators. Other studies document the forecasting power

* Corresponding author.

E-mail addresses: alexandridis@uom.edu.gr (A.K. Alexandridis), I.Apergis@kent.ac.uk (I. Apergis), A.Panopoulou@essex.ac.uk (E. Panopoulou), N.Voukelatos@kent.ac.uk (N. Voukelatos).

<https://doi.org/10.1016/j.finmar.2022.100801>

Received 14 June 2022; Received in revised form 21 December 2022; Accepted 22 December 2022

Available online 29 December 2022

1386-4181/© 2022 The Author(s).

Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Published by Elsevier B.V. This is an open access article under the CC BY license

of investor sentiment (Huang et al., 2015), cross-sectional return dispersion (Maio, 2016), manager sentiment (Jiang et al., 2019), oil price increases (Wang et al., 2019), news extracted from newspaper articles (Adämmer and Schüssler, 2020), and cross-sectional higher moments (Stöckl and Kaiser, 2020).

Other studies focus on whether equity premium predictability can be improved by adopting alternative econometric approaches to generate forecasts from a given set of variables. In particular, Rapach et al. (2010) examine a standard set of macroeconomic variables and show that forecast combinations result in statistically and economically significant gains in forecastability, in contrast to the poor performance of individual forecasts; see also Meligkotsidou et al. (2014) and Adämmer and Schüssler (2020). Moreover, Meligkotsidou et al. (2014, 2021) and Pedersen (2015) adopt a quantile regression approach and find that point forecasts that have been generated by aggregating across a set of quantiles significantly improve forecast accuracy. Following a different approach, Pettenuzzo et al. (2014), Li and Tsiakas (2017), and Tsiakas et al. (2020) show that forecast accuracy can be significantly improved by imposing economic and statistical constraints on equity premium forecasts.

Our paper contributes to the literature on equity premium predictability by exploring the predictive ability of a set of variables extracted from the options market. Given that option contracts are forward-looking by construction, it would be reasonable to expect that they contain important information about the future returns distribution of the underlying asset, in addition to information that is already contained in the historical record of returns or other contemporaneously observed variables. To this end, we focus on twelve strategy benchmark indices quoted by the Chicago Board Options Exchange (CBOE) as potential predictors of the equity premium. These indices reflect the performance of trading strategies that have been constructed using options written on the S&P 500, and they are designed to serve as benchmarks for investors trading index options. More importantly, the CBOE indices reflect investors' aggregate expectations about the distribution of future market returns. Considering that each CBOE strategy is based on a different mix of option contracts, the full set of strategy benchmarks is likely to contain rich information about several aspects of the distribution of market returns, such as tail risk, volatility risk, expected skewness, etc. Overall, we argue that the forward-looking nature and the increasing liquidity of index options make the CBOE benchmark indices natural candidates for the set of equity premium predictors.

Our paper also contributes to the literature by evaluating equity premium predictability at a daily frequency. Previous studies have tended to examine financial and economic predictors that are observed weekly, monthly or even at a lower frequency. However, the daily dynamics of predictors are likely to contain incremental information about the future evolution of the equity premium. Our dataset consists of the daily time-series of the CBOE benchmark indices and we focus on forecasting the daily equity premium.

Despite the extensive search for optimal predictors of the equity premium in the literature, there has been surprisingly little interest in exploiting the forward-looking information embedded in index options. In this sense, our paper is directly related to a small number of recent studies that have explored option-related information in the context of forecasting market returns. More specifically, Buss et al. (2017, 2019) show that implied correlation is a robust predictor of aggregate market returns at long horizons. Furthermore, Andersen et al. (2020) find that the tail risk premium extracted from index options can predict future market returns, while Cao et al. (2020) report that the implied volatility spread outperforms several well-established predictors at horizons of up to six months. Our research question is in a similar vein but, in contrast to focusing on a specific option-implied variable, we seek to exploit a richer information set about several features of the market returns distribution by using a large set of variables based on trading index options.

We examine the daily time-series of the equity premium from January 1996 to April 2021. We focus on the twelve CBOE strategy benchmark indices as our main predictors of interest, and we also consider some of the standard economic variables of Goyal and Welch (2008) as well as the VIX and the variance risk premium as commonly used predictors. We employ a methodology that is consistent with recent developments in the literature, allowing us to generate point forecasts, quantile forecasts and density forecasts of the equity premium. We evaluate the performance of alternative forecast models out-of-sample, in order to reflect investors' real-time decisions, with the emphasis on identifying models that provide statistical as well as economic gains in forecastability.

The findings provide strong support for the hypothesis that information from the options market can significantly improve the predictability of the equity premium. First, we find that almost all option variables outperform the historical average benchmark when used in univariate linear models, as evidenced by positive and highly significant out-of-sample R^2 . This finding is particularly important considering the relatively poor performance of the Goyal and Welch (2008) economic fundamentals and the fact that the historical average is notoriously hard to beat (Campbell and Thompson, 2008). The highest improvement in forecasting power is offered by the risk reversal strategy benchmark, highlighting the importance of skewness premia for predicting future market returns. Interestingly, we find that using the mean forecast across all univariate models outperforms the historical average benchmark, but to a lesser extent compared to most univariate models, while a kitchen sink model including all potential predictors actually underperforms.

Exploiting the entire distributional information of each predictor from quantile regressions and employing variable selection/shrinkage techniques leads to additional improvements in forecasting power. For instance, the results are even stronger when we examine point forecasts obtained by aggregating quantile forecasts, with universally positive out-of-sample R-squares and even lower p -values for all option variables. Moreover, we find that the least absolute shrinkage and selection operator (LASSO) scheme offers one of the highest improvements over the historical average benchmark, with an out-of-sample R^2 of 1.37%. Applying the Lima and Meng (2017) post LASSO quantile combinations (PLQC) scheme further improves forecasting performance, with the three-quantile PLQC3 scheme offering an R^2_{OS} of 1.48%, which is the highest across all the models that we examine. In contrast to the findings in Meligkotsidou et al. (2021), we find that fixed-weight PLQC schemes deliver more significant gains compared to schemes with time-varying weights.

The quantile regression results provide further support for the role of option variables in forecasting the equity premium, as all option predictors consistently outperform the historical average benchmark across all the quantiles that we consider. However, we find that the alternative quantile constant benchmark is relatively harder to beat compared to that of the historical average, with univariate models using option variables failing to offer statistically significant improvements in the left part of the distribution. Nevertheless, we find that the forecasting power of the CBOE indices gradually increases from the low to high quantiles, and the majority of univariate option-based models significantly outperform both benchmarks in the right part of the distribution. Importantly, when used in a multivariate setting, option predictors result in highly significant forecasting gains in both parts of the distribution.

The results are qualitatively similar when forecasting the entire density, although the performance varies depending on the weighting scheme used to aggregate across quantiles. In brief, we find that every option predictor outperforms the historical average benchmark, with the highest improvements offered by the mean forecast model. The quantile constant benchmark is again found to be more difficult to beat, but we still find that 9 of the 12 option variables offer significant gains under weighting schemes that place greater emphasis on the right part or the middle of the distribution, while skewness premia and multivariate models are the only cases of outperformance when we place greater emphasis on the left part of the distribution.

Finally, we find that forecasts generated by the option variables are not only statistically significant, but also economically valuable. For instance, a standard mean–variance investor generating forecasts using the BXMD buy-write index in a univariate setting would earn a certainty equivalent return (CER) of 5.6% per annum, representing an almost fourfold improvement over and above the CER of 1.27% offered by the historical mean benchmark. Applying variable selection via PLQC would offer a more than fivefold increase in the investor's Sharpe ratio, reaching 0.60 when using time-varying weights, compared to a Sharpe ratio of 0.12 offered by the benchmark. Generally, using option predictors delivers important economic gains in terms of returns per unit of risk, as evidenced by Sharpe ratios that are between three and five times higher than that of the historical average benchmark across all our models.

The remainder of the paper is organized as follows. In Section 2, we discuss the econometric methodology that is applied to generate and evaluate forecasts of the equity premium under competing models. In Section 3, we describe the data used in the empirical analysis, while in Section 4 we present the empirical results. Finally, we conclude in Section 5.

2. Methodology

In this section, we describe the methodological approaches we use to produce and evaluate the forecasts of the equity premium. We begin by discussing a number of alternative forecasting approaches, ranging from the simple linear model to several types of quantile forecast combinations. We then present the criteria we use to evaluate the accuracy of these forecasts.

2.1. Forecasting approaches

2.1.1. Univariate linear model

We begin with the traditional predictive linear regression model, where we regress the equity premium against a lagged predictor as follows:

$$r_{t+1} = \alpha + \beta x_t + \varepsilon_{t+1}, \quad (1)$$

where r_{t+1} denotes the equity premium at time $t + 1$, x_t is the value of the predictive variable at t , and ε_{t+1} is a random error term.

We generate out-of-sample forecasts of the equity premium by estimating Eq. (1) recursively; see also Rapach et al. (2010) and Goyal and Welch (2008). More specifically, we start by estimating the linear model using an initial window consisting of the first m observations, regressing $\{r_t\}_{t=2}^m$ against a constant and $\{x_t\}_{t=1}^{m-1}$. Using the estimated parameters and the predictive variable's realized value x_m , we produce the first out-of-sample forecast of the equity premium at time $m + 1$, given by:

$$\hat{r}_{m+1} = \hat{\alpha} + \hat{\beta}x_m. \quad (2)$$

We obtain the time-series of out-of-sample equity premium forecasts by repeating these steps using a sequence of expanding windows. Finally, we generate one time-series of equity premium forecasts separately for each variable x_t in our set of predictors.

2.1.2. Kitchen sink

In addition to the univariate linear model in Eq. (1), we combine the information contained in all the predictive variables by estimating kitchen sink forecasts. We start by estimating multivariate regressions of the equity premium against all k lagged predictors as follows:

$$r_{t+1} = \alpha + \sum_{i=1}^k \beta_i x_{i,t} + \varepsilon_{t+1}, \quad (3)$$

where $x_{i,t}$ denotes the value of the i th predictor at time t . The kitchen sink forecast of the equity premium is given by:

$$\hat{r}_{t+1} = \hat{\alpha} + \sum_{i=1}^k \hat{\beta}_i x_{i,t}. \quad (4)$$

Similarly to the univariate case in Section 2.1.1, we estimate kitchen sink forecasts in a recursive fashion.

2.1.3. LASSO regression

The multivariate predictive regressions in Section 2.1.2 are likely to suffer from overfitting, especially when the number of predictive variables is relatively large (see [Rapach and Zhou, 2021](#)).¹ To address this concern, we adopt the [Tibshirani \(1996\)](#) LASSO approach. The LASSO is a shrinkage technique that can be used for variable selection. More specifically, this approach applies a penalty to the estimated slope coefficients by forcing the sum of the absolute values of these coefficients to be below a predetermined threshold. In the process, one or more coefficients could be shrunk to zero, effectively excluding the associated variables from the set of predictors. The LASSO regression coefficients can be obtained by solving the following optimization problem:

$$\min_{\alpha, \beta_1, \dots, \beta_k} \left[\sum_{i=1}^{T-1} (r_{i+1} - \alpha - \sum_{i=1}^k \beta_i x_{i,t})^2 + \lambda \sum_{i=1}^k |\beta_i| \right], \quad (5)$$

where λ denotes a penalty parameter that determines the degree of shrinkage. If $\lambda = 0$, then the LASSO coefficients are identical to those obtained by the standard kitchen sink estimation in Eq. (3). As λ increases, $\hat{\beta}$ shrinks to zero. While selecting an appropriate value for the shrinkage parameter is not straightforward, the main objective is to optimize the trade-off between reducing overfitting and discarding potentially useful information. When estimating LASSO regressions, we set $\lambda = \frac{1}{n}$, where n is the number of in-sample observations in a particular step.²

2.1.4. Elastic net

We use the [Zou and Hastie \(2005\)](#) elastic net (ENet) approach as an alternative shrinkage technique. Similarly to the LASSO forecast, the ENet forecast is based on a penalized regression that addresses potential overfitting. However, the ENet penalty term consists of two components: a LASSO component λ_1 and a ridge component λ_2 ([Hoerl and Kennard, 1970](#)). We follow [Dong et al. \(2022\)](#) to select the value of the parameter that determines the degree of shrinkage, with the ENet coefficients being obtained by solving the following system:

$$\min_{\alpha, \beta_1, \dots, \beta_k} \frac{1}{2} \sum_{i=1}^{T-1} (r_{i+1} - \alpha - \sum_{i=1}^k \beta_i x_{i,t})^2 \text{ subject to } \sum_{i=1}^k |\beta_i| \leq \lambda_1 \text{ and } \sum_{i=1}^k (\beta_i)^2 \leq \lambda_2. \quad (6)$$

2.1.5. Principal component analysis

We consider forecasts based on principal component analysis (PCA), which has often been adopted in equity premium prediction as a way to extract a common source of variation among a large set of predictors (e.g., [Neely et al., 2014](#); [Li and Tsiakas, 2017](#)). We also apply the more recent scaled PCA (sPCA) approach developed by [Huang et al. \(2021\)](#). In contrast to the equal weighting of predictors in the PCA, the sPCA scales each predictor according to its predictive power over a particular target that is being forecasted. In other words, instead of maximizing the extent to which a principal component can explain the variation among the predictors, the sPCA attempts to maximize a principal component's forecasting accuracy on a specific target. As a result of using the target variable information to guide dimension reduction, the sPCA has been found to generate more efficient forecasts of market returns ([Huang et al., 2021](#); [Chen et al., 2022](#)). We follow [Dong et al. \(2022\)](#) and extract the first principal component from the full set of predictors in our sample, for both the PCA and the sPCA forecasts.

2.1.6. Partial least squares

We use the partial least squares (PLS) approach to construct a single forecasting variable from our large set of predictors. Following [Kelly and Pruitt \(2013, 2015\)](#), we apply a three-pass regression filter to construct a factor as a linear combination of individual predictors. In a similar spirit to the sPCA, the PLS technique uses information from the forecasting target to create an optimal predictor. More specifically, the weight of each individual predictor in the PLS combination is determined by its covariance with the forecast target (i.e., the equity premium), resulting in a univariate predictor that is expected to maximize the correlation with the target variable (see also [Dong et al., 2022](#)).

2.1.7. Quantile regression

The linear models discussed in the previous sections can generate forecasts about the mean of the returns distribution. However, a number of studies document significant non-linear predictability patterns in stock returns (e.g., [Guidolin et al., 2009](#); [Henkel et al., 2011](#)). To capture the potentially non-linear relationship between the equity premium and the set of predictors, we adopt a quantile regression approach that allows us to explore the equity premium's predictability across different parts of its distribution, including the center as well as the tails (see [Meligkotsidou et al., 2014, 2019](#); [Pedersen, 2015](#)).

The quantile regression model is given by:

$$r_{t+1} = \alpha^{(\tau)} + \beta^{(\tau)} x_t + \varepsilon_{t+1}, \quad (7)$$

¹ As a first, simple way of addressing overfitting, we consider the mean of all univariate forecasts generated in sub Section 2.1.1 (Mean), as well as the cross-sectional mean of individual predictors (PredAvg).

² We also considered selecting the optimal value for λ by running a tenfold cross-validation and then choosing the value for λ that minimizes the respective mean squared error. The results are relatively similar to, albeit weaker than, those obtained when simply setting λ equal to $\frac{1}{n}$. Therefore, we omit these latter results for brevity, but they are available upon request.

where $\tau \in (0, 1)$, while $\alpha^{(\tau)}$ and $\beta^{(\tau)}$ are quantile-varying parameters. The errors ε_{t+1} are assumed to be independent and drawn from an error distribution $g_\tau(\varepsilon)$ with the τ th quantile equal to 0 (i.e., $\int_{-\infty}^0 g_\tau(\varepsilon)d\varepsilon = \tau$).

These quantile regressions are estimated separately for each predictive variable in our set. The estimated parameters $\hat{\alpha}^{(\tau)}$ and $\hat{\beta}^{(\tau)}$ are obtained by minimizing the sum $\sum_{t=0}^{T-1} \rho_\tau(r_{t+1} - \alpha^{(\tau)} - \beta^{(\tau)}x_t)$, where $\rho_\tau(u)$ is an asymmetric linear loss function given by:

$$\rho_\tau(u) = u(\tau - I(u < 0)) = \frac{1}{2} [|u| + (2\tau - 1)u]. \tag{8}$$

Once the parameters in Eq. (7) have been estimated, the forecast of the τ th quantile of the equity premium distribution at $t + 1$ can be obtained as $\hat{r}_{t+1}^{(\tau)} = \hat{\alpha}^{(\tau)} + \hat{\beta}^{(\tau)}x_t$. For example, in the symmetric case of $\tau = \frac{1}{2}$, we obtain the median of the forecasted distribution. Moreover, we follow Meligkotsidou et al. (2014) and Lima and Meng (2017) and forecast the mean of the equity premium distribution as the weighted average of a set of quantiles. The weights used can be interpreted as the probabilities associated with different quantile forecasts, indicating how likely a particular regression quantile is to predict the equity premium over the next period. More specifically, we compute the point forecast of the equity premium as:

$$\hat{r}_{t+1} = \sum_{\tau \in S} p_\tau \hat{r}_{t+1}^{(\tau)}, \quad \sum_{\tau \in S} p_\tau = 1, \tag{9}$$

where p_τ denotes the weight associated with quantile τ , and S denotes the full set of quantiles that are being aggregated. Regarding the specific choice of weights p_τ , we follow Gastwirth (1966) and use the three-quantile (Q_3) and the five-quantile (Q_5) estimators given by:

$$Q_3 : \hat{r}_{t+1} = \frac{1}{3} \hat{r}_{t+1}^{(0.3)} + \frac{1}{3} \hat{r}_{t+1}^{(0.5)} + \frac{1}{3} \hat{r}_{t+1}^{(0.7)}. \tag{10}$$

$$Q_5 : \hat{r}_{t+1} = \frac{1}{5} \hat{r}_{t+1}^{(0.3)} + \frac{1}{5} \hat{r}_{t+1}^{(0.4)} + \frac{1}{5} \hat{r}_{t+1}^{(0.5)} + \frac{1}{5} \hat{r}_{t+1}^{(0.6)} + \frac{1}{5} \hat{r}_{t+1}^{(0.7)}. \tag{11}$$

2.1.8. Post LASSO quantile combinations

Including predictors with very small effects on the equity premium in the forecasting equation is likely to have a significant negative impact on forecast accuracy. We address the issue of potentially weak predictors by applying the post LASSO quantile combination (PLQC) approach that was first proposed by Lima and Meng (2017). The PLQC is a methodology that attempts to minimize the negative impact of weak predictors and estimation errors by applying an averaging scheme to quantiles that are based on LASSO-selected predictors.

We begin by applying the ℓ_1 -penalized LASSO following Chernozhukov et al. (2010). We then select only the predictors whose coefficients have not been shrunk to zero (i.e., excluding any weak predictors).³ We then use the selected predictors in a quantile regression, with the post-LASSO forecast for quantile τ being given by:

$$\hat{r}_{t+1}^{(\tau)} = \hat{\alpha}^{(\tau)} + \sum_{i=1}^{k^*} \hat{\beta}_i^{(\tau)} x_{i,t}^*, \tag{12}$$

where $\hat{\alpha}^{(\tau)}$ denotes the estimated intercept and $\hat{\beta}_i^{(\tau)}$ the estimated slope for the i th predictor in the quantile regression. Furthermore, $x_{i,t}^*$ is the value of the i th predictor selected at time t , while k^* is the total number of selected predictors. We use this approach to obtain one forecast for each quantile at time t . We combine all these quantile forecasts to construct the PLQC forecast of the mean of the equity premium distribution as:

$$\hat{r}_{t+1} = \sum_{j=1}^J \omega_{t,\tau_j} \hat{r}_{t+1,t}^{(\tau_j)}, \tag{13}$$

where ω denotes the averaging scheme we use to combine quantile forecasts and J is the total number of quantiles we use to aggregate the forecasts. In our analysis, we use both fixed and time-varying schemes. With respect to fixed weights, we consider a discrete set of quantiles $\tau \in (\tau_1, \tau_2, \dots, \tau_J)$ and compute the PLQC forecast as the simple arithmetic average with equal weighing $\omega_\tau = \omega$. Consistent with the quantile models presented in Section 2.1.7, we compute these equal-weighted PLQC forecasts as:

$$PLQC_3 : \hat{r}_{t+1} = \frac{1}{3} \hat{r}_{t+1,t}^{(0.3)} + \frac{1}{3} \hat{r}_{t+1,t}^{(0.5)} + \frac{1}{3} \hat{r}_{t+1,t}^{(0.7)}. \tag{14}$$

$$PLQC_5 : \hat{r}_{t+1} = \frac{1}{5} \hat{r}_{t+1,t}^{(0.3)} + \frac{1}{5} \hat{r}_{t+1,t}^{(0.4)} + \frac{1}{5} \hat{r}_{t+1,t}^{(0.5)} + \frac{1}{5} \hat{r}_{t+1,t}^{(0.6)} + \frac{1}{5} \hat{r}_{t+1,t}^{(0.7)}. \tag{15}$$

In addition, we apply time-varying weights to account for the possibility that the contribution of specific quantiles in the optimal forecast varies across time. Similarly to Lima and Meng (2017), we determine the weighting scheme $\omega_{t,j}$ by estimating a constrained OLS regression of r_{t+1} on $\hat{r}_{t+1,t}^{(\tau)}$, $\tau \in (\tau_1, \tau_2, \dots, \tau_J)$. The resulting PLQC forecasts with time-varying (TW) weights are given by:

$$TW_3 : \hat{r}_{t+1} = \omega_{t,\tau_1} \hat{r}_{t+1,t}^{(0.3)} + \omega_{t,\tau_2} \hat{r}_{t+1,t}^{(0.5)} + \omega_{t,\tau_3} \hat{r}_{t+1,t}^{(0.7)}. \tag{16}$$

³ The ℓ_1 -penalized LASSO method can be used to classify predictors into three groups. *Strong* predictors are those that are selected in all quantiles, *weak* predictors are those that are selected in a subset of all quantiles, while *fully weak* predictors are those that are not selected in any quantile. When obtaining post-LASSO forecasts for a given quantile, we exclude all weak predictors, without trying to distinguish between weak and fully weak ones across the entire set of quantiles.

$$\begin{aligned}
& \text{s.t. } \omega_{i,\tau_1} + \omega_{i,\tau_2} + \omega_{i,\tau_3} = 1. \\
TW_5 : \quad & \hat{r}_{t+1} = \omega_{i,\tau_1} \hat{r}_{t+1,t}^{(0.3)} + \omega_{i,\tau_2} \hat{r}_{t+1,t}^{(0.4)} + \omega_{i,\tau_3} \hat{r}_{t+1,t}^{(0.5)} + \omega_{i,\tau_4} \hat{r}_{t+1,t}^{(0.6)} + \omega_{i,\tau_5} \hat{r}_{t+1,t}^{(0.7)} \\
& \text{s.t. } \omega_{i,\tau_1} + \omega_{i,\tau_2} + \omega_{i,\tau_3} + \omega_{i,\tau_4} + \omega_{i,\tau_5} = 1.
\end{aligned} \tag{17}$$

2.2. Forecast evaluation criteria

The methodologies discussed in Section 2.1 allow us to generate a set of point, quantile, and density forecasts of the equity premium. We obtain these time-series of equity premium forecasts using a recursive (expanding) window. In particular, we begin by using the model parameters we obtain in the initial estimation period that consists of the first five years, as well as the predictors' values on the last day of the estimation period, in order to forecast the equity premium one day ahead. We continue to generate out-of-sample forecasts by continuously updating the estimation period, adding one observation at a time. This approach allows us to generate a time-series of one-day-ahead out-of-sample forecasts of the equity premium, under a set of competing models. We evaluate the forecast accuracy of each model based on the following criteria.

2.2.1. Point forecast accuracy

We evaluate the accuracy of point forecasts using the standard out-of-sample R^2 , computed as:

$$R_{OS}^2 = 1 - \frac{MSFE_i}{MSFE_0}, \tag{18}$$

where $MSFE_i$ and $MSFE_0$ denote the mean squared forecast errors of the i th model and the benchmark model, respectively. Positive values of the R_{OS}^2 are indicative of the proposed model outperforming the benchmark, while negative values indicate the opposite. We evaluate the statistical significance of a model's R_{OS}^2 by performing the Clark and West (2007) test.

2.2.2. Quantile forecast accuracy

Gneiting and Raftery (2007) and Gneiting and Ranjan (2011) suggest that forecast evaluation should be based on the same loss function as the one used in model estimation (i.e., Eq. (8) used in the quantile regressions in this study). In order to evaluate the accuracy of quantile forecasts, we follow Manzan (2015) and compute the forecasts' quantile score (QS). More specifically, we compute the QS of the τ th quantile forecast generated by the i th model as:

$$QS_{t+1|t}^i(\tau) = \left[r_{t+1} - \hat{r}_{t+1}^{(\tau)} \right] \left[\tau - I(r_{t+1} - \hat{r}_{t+1}^{(\tau)} < 0) \right], \tag{19}$$

where $\hat{r}_{t+1}^{(\tau)}$ denotes the i th model's forecast of quantile τ , and I is an indicator function that takes the value of 1 if the argument is true and the value of 0 otherwise. A lower QS indicates superior forecast accuracy.

We evaluate the statistical significance of quantile forecast accuracy by following Giacomini and White (2006) and Amisano and Giacomini (2007). To this end, we compute the quantile score's test statistic as

$$t_i^{QS}(\tau) = \frac{\overline{QS}_i(\tau) - \overline{QS}_0(\tau)}{\hat{\sigma}}, \tag{20}$$

where $\overline{QS}_i(\tau)$ and $\overline{QS}_0(\tau)$ denote the mean QS for quantile τ generated by the i th model and the benchmark model, respectively, while $\hat{\sigma}$ is the standard error of the quantile score difference. The null hypothesis is that both models' quantile scores are equal, with model i outperforming the benchmark if the null is rejected with a negative $t_i^{QS}(\tau)$ and the benchmark outperforming model i if the null is rejected with a positive $t_i^{QS}(\tau)$.

2.2.3. Density forecast accuracy

We interpolate across the set of quantile forecasts to approximate the entire density of the equity premium, without the constraint of assuming a particular distribution. The forecast accuracy of these density forecasts is evaluated via the weighted quantile score (WQS) by integrating the QS across the set of quantiles τ , with a function ω assigning different weights to different parts of the distribution. We compute the WQS of model i as:

$$WQS_{t+1|t}^i = \int_0^1 QS_{t+1|t}^i(\tau) \omega(\tau) d\tau. \tag{21}$$

Given that our analysis is based on a discrete set of quantiles, we replace the continuous version of WQS in Eq. (21) with a discrete version that aggregates across the obtained quantiles. More specifically, we employ four different weighting functions ω , namely:

1. WQS_1 : $\omega(\tau) = 1$, assigning uniform weights across the entire distribution.
2. WQS_2 : $\omega(\tau) = \tau(1 - \tau)$, assigning greater weights to the middle of the distribution.
3. WQS_3 : $\omega(\tau) = (1 - \tau)^2$, assigning greater weights to the left tail of the distribution.
4. WQS_4 : $\omega(\tau) = \tau^2$, assigning greater weights to the right tail of the distribution.

Finally, we evaluate the statistical significance of the weighted quantile score by replacing QS with WQS in Eq. (20).

3. Data and main variables

We focus on the period January 4, 1996 to April 15, 2021, for a total of 6312 daily observations. Our main variable of interest is the equity premium r_t , defined as the log return of the market index in excess of the risk-free rate. In particular, the equity premium at time t is given by $r_t = \ln(1 + r_{mkt,t}) / \ln(1 + r_{f,t})$, where $r_{mkt,t}$ and $r_{f,t}$ denote the return of the S&P 500 index and the 1-month Treasury bill rate, respectively, at t . We use the first five years (1248 daily observations) as our initial estimation period in the forecasting exercise, with the out-of-sample period starting on January 3, 2001.

The set of predictors consists of the daily returns of 12 strategy benchmark indices trading S&P 500 options, the VIX, the variance risk premium (VRP), and 3 economic variables. We obtain data on the strategy benchmark indices directly from the CBOE, which reports the daily returns of a number of passive strategies involving options written on the S&P 500. These indices are designed to serve as benchmarks for investors trading options on the exchange, and they can be split into seven categories according to the type of information that they are expected to reflect. Each category contains several strategies that are designed to broadly pursue the same objective by adopting variations of the same construction methodology. Given the significant degree of correlation among same-category strategies, we select from each group only a subset of strategy indices that can capture the main properties of that category. More specifically, our set of option predictors consists of the Buy-Write (BXM, BXMC, BXMD, and BXY), Put-Write (PUT and PUTY), Combo (CMBO), Butterfly (BFLY), Condor (CNDR), Collar (CLL), Put Protection (PPUT), and Risk Reversal (RXM) strategies. Table A1 in the Online Appendix provides more information about the construction of each strategy benchmark index.

While a theoretical framework that links these option variables to subsequent equity returns has yet to be developed, our choice of options-based strategies as potential predictors is motivated primarily by the substantial literature that has emerged on the informational content of the options market. For instance, [Vanden \(2008\)](#) shows that option prices subsume market expectations about future investment opportunities, while several studies find that information extracted from options has significant forecasting power over stock price dynamics (e.g., [Shackleton et al., 2010](#); [Christoffersen et al., 2012](#); [Buss et al., 2017, 2019](#)).

The forward-looking nature of option contracts suggests that they are, in theory, expected to contain information about investors' expectations of the future state of the underlying market index. Moreover, each benchmark strategy reflects aggregate expectations about different parts and/or moments of the index's distribution, determined by the specific types of option contracts that it trades. Buy-Write strategies, for instance, provide long equity and short volatility exposure ([Israelov and Nielsen, 2014](#)). By going long in the underlying market index and short in a covered call, these strategies perform well when the market fluctuates little in the short-term and rises in the longer-term. As such, these option variables reflect investors' beliefs about the future short-term and longer-term movements of the underlying index, regarding the market's level and volatility. Importantly, each of the four Buy-Write strategies (BXM, BXMC, BXMD, and BXY) has different exposure to short volatility, determined by the moneyness of the options used in its construction, so that using all four variables captures a large part of the forecasting power embedded in the options' implied volatility smile; see also [Whaley \(2002\)](#) and [Israelov and Nielsen \(2014, 2015\)](#).

Put-Write strategies, on the other hand, reflect the performance of providing crash insurance on the market index, a practice that has traditionally offered very high returns ([Bondarenko, 2014](#); [Kelly et al., 2016](#)). In this sense, PUT and PUTY capture investors' expectations about the likelihood of a market crash, as well as their associated aversion to crash risk. This variable is expected to correlate significantly with the subsequent equity premium, with higher returns of put-writing strategies being indicative of investors attaching a higher probability to a subsequent market crash.

Butterfly (BFLY) and Condor (CNDR) are standard volatility trading strategies that reflect investors' appetite for insurance against volatility risk. Previous studies show that volatility expectations embedded in the VIX and the VRP have significant forecasting power over the equity premium, documenting a positive relation between volatility expectations and subsequent market returns ([Bollerslev et al., 2009, 2014](#); [Buss et al., 2017](#)). The volatility strategies that we explore can be similarly seen as complimentary measures of aggregate uncertainty and risk aversion extracted from option prices. In fact, BFLY and CNDR represent direct ways of trading on the variance risk premium using options, as opposed to going short in VIX futures. Therefore, we expect these volatility strategies to contain incremental information about future market returns, in line with the findings of [Bollerslev et al. \(2009\)](#) of high volatility premia being associated with high future returns. The Combo strategy (CMBO) is essentially a 50/50 combination of the PUT and BFLY strategies, jointly reflecting investors' beliefs about crash risk and volatility risk.

Collar (CLL) and put protection (PUTY) strategies are typically considered as a cost-efficient way of hedging the underlying market index. By going long in an index put and short in an index call, CLL is designed to protect a long position in the underlying index at a relatively low cost, at the expense of capping upside potential. Given its construction, CLL is expected to perform well when investors have bearish short-term forecasts but bullish long-term ones about the market index. PUTY offers an alternative, and relatively more expensive, way of hedging the market index. Put protection strategies tend to somewhat underperform during good states of the market but substantially outperform during crashes, thereby reflecting investors' expectations of large drops of the market index.

Finally, risk reversals (RXM) are designed to reverse the risk stemming from the well-documented negative skewness of index returns. As such, RXM returns directly reflect the implied skewness premium and, by extension, capture investors' expectations about future market skewness. Although skewness should, in theory, be negatively related to future stock returns, the empirical results have been somewhat mixed. For instance, [Bali and Murray \(2013\)](#), [Chang et al. \(2013\)](#), [Conrad et al. \(2013\)](#), and [Kim and Park \(2018\)](#) confirm a negative relation between implied skewness and subsequent stock returns, while a positive relation is reported




Fig. 1. Time-evolution of main variables.

Notes: We plot the time evolution of the equity premium (EP) and the options predictors. Each subplot shows the cumulative daily return of the respective variable. The sample period is from January 3, 1996 to April 15, 2021.

by Bali and Hovakimian (2009), Xing et al. (2010), and Chordia and Lin (2021). We expect RXM to reflect investors' beliefs about the likelihood of a market crash (Doran et al., 2007) and to act as a proxy for investor sentiment (Han, 2008; Cao et al., 2020).⁴

Following the literature on the predictive power of the VIX and the VRP over market returns, we include both of these variables in our set of predictors. We obtain daily data on the VIX directly from the CBOE, while data on the VRP are from the personal webpage of Grigory Vilkov.⁵ In addition to the above option predictors, we use a set of variables that have often been found to contain information about future market returns. In particular, Goyal and Welch (2008) show that a set of economic variables can have a significant predictive power over the equity premium. Given that our forecasting exercise is based on daily data, we include in our set of predictors only the Goyal and Welch (2008) variables that are available at a daily frequency. More specifically, we include the term spread (TMS), defined as the difference between the yield of long-term government bonds and the yield of T-bills, the TED spread (TED), defined as the difference between the 3-month T-bill rate and 3-month LIBOR, and the default yield spread (DFY), given as the difference between the yields of BAA and AAA corporate bonds. We obtain the data on these economic variables from Bloomberg.

Fig. 1 shows the time evolution of the equity premium and the option predictors during our sample period while Table 1 reports the descriptive statistics. The equity premium has a mean of 0.02% on a daily basis, with a standard deviation of 1.22%. Consistent with earlier findings in the literature, we find that the equity premium is characterized by negative skewness and excess kurtosis. All the option-related predictors are also negatively skewed (with PPUT and VIX being the only exceptions) and substantially leptokurtic, while the economic predictors have positive skewness. As can be seen in Table 2, the strategy benchmark indices generally tend to be strongly correlated with one another. Interestingly, the option predictors are negatively and weakly correlated with the VIX, and almost completely uncorrelated with the three macroeconomic predictors. These relatively weak correlations with variables that have traditionally been used to forecast the equity premium suggests that the strategy benchmark indices are likely to contain incremental information about future market returns.

⁴ Cao et al. (2020) find that an alternative measure of implied skewness (the implied volatility spread) is significantly positively related to several measures of market expectations and investor sentiment, namely the Gallup investor survey, American Association of Individual Investors survey, Crash Confidence Index from the Yale School of Management, and the Baker and Wurgler (2006) sentiment index. Implied skewness was also found to be significantly negatively related to eight measures of market uncertainty, including macroeconomic, political, and financial uncertainty.

⁵ The VRP data can be found at <https://www.vilkov.net/codedata.html>. This dataset does not cover our entire sample period, as it ends in December 2017. Therefore, we include the analysis of the VRP in the paper's robustness section.

Table 1
Descriptive statistics.

	Mean	Median	St.Dev	Min	Max	Skew	Kurt	LBQ	JB
<i>EP</i>	0.0002	0.0006	0.0122	-0.1277	0.1024	-0.53	12.72	0.00	0.00
<i>BXM</i>	0.0003	0.0006	0.0087	-0.1296	0.0898	-1.53	28.44	0.00	0.00
<i>BXMC</i>	0.0003	0.0007	0.0092	-0.1283	0.1097	-1.11	26.29	0.00	0.00
<i>BXMD</i>	0.0003	0.0008	0.0105	-0.1272	0.0930	-0.88	17.60	0.00	0.00
<i>BXY</i>	0.0004	0.0008	0.0112	-0.2954	0.2969	-0.75	169.24	0.00	0.00
<i>PUT</i>	0.0003	0.0005	0.0083	-0.1218	0.0903	-1.62	32.70	0.00	0.00
<i>PUTY</i>	0.0002	0.0003	0.0071	-0.1228	0.0908	-2.35	52.96	0.00	0.00
<i>CMBO</i>	0.0003	0.0007	0.0090	-0.1240	0.0892	-1.41	24.65	0.00	0.00
<i>BFLY</i>	0.0001	0.0006	0.0070	-0.0501	0.0628	-0.48	9.90	0.09	0.00
<i>CNDR</i>	0.0001	0.0005	0.0048	-0.0462	0.0431	-1.40	23.02	0.00	0.00
<i>CLL</i>	0.0002	0.0003	0.0081	-0.2208	0.1165	-3.10	100.39	0.02	0.00
<i>PPUT</i>	0.0003	0.0003	0.0089	-0.0799	0.0679	0.02	7.11	0.00	0.00
<i>RXM</i>	0.0002	0.0003	0.0074	-0.1189	0.0892	-2.01	42.03	0.00	0.00
<i>VIX</i>	0.2033	0.0845	0.8269	0.1868	0.0914	2.09	10.54	0.00	0.00
<i>TMS</i>	0.0158	0.0155	0.0112	-0.0095	0.0385	0.05	2.04	0.00	0.00
<i>TED</i>	0.0046	0.0035	0.0039	0.0009	0.0458	3.35	21.71	0.00	0.00
<i>DFY</i>	0.0099	0.0090	0.0041	0.0050	0.0350	3.07	15.81	0.00	0.00

Notes: This table reports the descriptive statistics for the equity premium (*EP*) and the predictive variables. The set of predictors includes the returns of twelve strategy benchmark indices based on index options, the VIX, and three macroeconomic variables. The option strategy predictors consist of the Buy-Write Index (*BXM*), Conditional Buy-Write Index (*BXMC*), 30-Delta Buy-Write Index (*BXMD*), 2% OTM Buy-Write Index (*BXY*), Put-Write Index (*PUT*), 2% OTM Put-Write Index (*PUTY*), Covered Combo Index (*CMBO*), Iron Butterfly Index (*BFLY*), Iron Condor Index (*CNDR*), 95–110 Collar Index (*CLL*), 5% Put Protection Index (*PPUT*), and the Risk Reversal Index (*RXM*). The macroeconomic variables consist of the term spread (*TMS*), TED spread (*TED*), and the default yield spread (*DFY*). The descriptive statistics reported include the mean, median, standard deviation, minimum, maximum, skewness, and kurtosis of the daily time-series. The last two columns report the *p*-values of the [Ljung and Box \(1978\)](#) test for serial correlation and the [Jarque and Bera \(1987\)](#) test for normality. The sample period is from January 3, 1996 to April 15, 2021.

Table 2
Correlation matrix.

	<i>EP</i>	<i>BXM</i>	<i>BXMC</i>	<i>BXMD</i>	<i>BXY</i>	<i>PUT</i>	<i>PUTY</i>	<i>CMBO</i>	<i>BFLY</i>	<i>CNDR</i>	<i>CLL</i>	<i>PPUT</i>	<i>RXM</i>	<i>VIX</i>	<i>TMS</i>	<i>TED</i>	<i>DFY</i>
<i>EP</i>	1.00	0.91	0.22	0.97	0.83	0.88	0.84	0.93	0.17	0.37	0.73	0.90	0.90	-0.15	0.00	-0.05	-0.01
<i>BXM</i>	0.91	1.00	0.22	0.96	0.86	0.97	0.94	0.98	0.42	0.60	0.62	0.70	0.86	-0.14	-0.01	-0.05	-0.02
<i>BXMC</i>	0.22	0.22	1.00	0.23	0.20	0.23	0.24	0.24	0.06	0.09	0.19	0.16	0.23	-0.13	0.00	-0.03	-0.01
<i>BXMD</i>	0.97	0.96	0.23	1.00	0.86	0.94	0.90	0.98	0.31	0.52	0.69	0.82	0.86	-0.14	0.00	-0.04	-0.01
<i>BXY</i>	0.83	0.86	0.20	0.86	1.00	0.84	0.80	0.87	0.30	0.49	0.61	0.68	0.75	-0.13	-0.01	-0.04	-0.01
<i>PUT</i>	0.88	0.97	0.23	0.94	0.84	1.00	0.97	0.97	0.43	0.63	0.58	0.64	0.89	-0.14	-0.01	-0.05	-0.02
<i>PUTY</i>	0.84	0.94	0.24	0.90	0.80	0.97	1.00	0.94	0.39	0.65	0.49	0.55	0.90	-0.15	0.00	-0.06	-0.02
<i>CMBO</i>	0.93	0.98	0.24	0.98	0.87	0.97	0.94	1.00	0.38	0.59	0.65	0.74	0.88	-0.15	0.00	-0.05	-0.01
<i>BFLY</i>	0.17	0.42	0.06	0.31	0.30	0.43	0.39	0.38	1.00	0.64	0.05	0.08	0.17	-0.01	0.00	0.01	-0.01
<i>CNDR</i>	0.37	0.60	0.09	0.52	0.49	0.63	0.65	0.59	0.64	1.00	0.22	0.15	0.44	-0.07	0.00	-0.01	-0.01
<i>CLL</i>	0.73	0.62	0.19	0.69	0.61	0.58	0.49	0.65	0.05	0.22	1.00	0.78	0.57	-0.10	-0.01	-0.02	-0.01
<i>PPUT</i>	0.90	0.70	0.16	0.82	0.68	0.64	0.55	0.74	0.08	0.15	0.78	1.00	0.68	-0.10	0.00	-0.01	0.00
<i>RXM</i>	0.90	0.86	0.23	0.86	0.75	0.89	0.90	0.88	0.17	0.44	0.57	0.68	1.00	-0.15	-0.01	-0.05	-0.02
<i>VIX</i>	-0.15	-0.14	-0.13	-0.14	-0.13	-0.14	-0.15	-0.15	-0.01	-0.07	-0.10	-0.10	-0.15	1.00	0.13	0.46	0.59
<i>TMS</i>	0.00	-0.01	0.00	0.00	-0.01	-0.01	0.00	0.00	0.00	0.00	-0.01	0.00	-0.01	0.13	1.00	-0.19	0.31
<i>TED</i>	-0.05	-0.05	-0.03	-0.04	-0.04	-0.05	-0.06	-0.05	0.01	-0.01	-0.02	-0.01	-0.05	0.46	-0.19	1.00	0.34
<i>DFY</i>	-0.01	-0.02	-0.01	-0.01	-0.01	-0.02	-0.02	-0.01	-0.01	-0.01	-0.01	0.00	-0.02	0.59	0.31	0.34	1.00

Notes: This table reports the correlation matrix among the EP and the predictive variables. The set of predictors includes the returns of twelve strategy benchmark indices based on index options, the VIX, and three macroeconomic variables. The option strategy predictors consist of the Buy-Write Index (*BXM*), Conditional Buy-Write Index (*BXMC*), 30-Delta Buy-Write Index (*BXMD*), 2% OTM Buy-Write Index (*BXY*), Put-Write Index (*PUT*), 2% OTM Put-Write Index (*PUTY*), Covered Combo Index (*CMBO*), Iron Butterfly Index (*BFLY*), Iron Condor Index (*CNDR*), 95–110 Collar Index (*CLL*), 5% Put Protection Index (*PPUT*), and the Risk Reversal Index (*RXM*). The macroeconomic variables consist of the term spread (*TMS*), TED spread (*TED*), and the default yield spread (*DFY*). The descriptive statistics reported include the mean, median, standard deviation, minimum, maximum, skewness, and kurtosis of the daily time-series. The sample period is from January 3, 1996 to April 15, 2021.

4. Results

4.1. In-sample predictability

We begin by exploring the in-sample predictability of the equity premium. To this end, [Table 3](#) reports the estimated coefficients from univariate predictive regressions of the equity premium against each predictor in turn. The table reports the results from simple OLS estimations as well as those from quantile regressions.

The univariate regression results provide initial support for the hypothesis that the options market contains significant information about the future evolution of the underlying equity index. Estimating a simple linear predictive model produces slope

Table 3
In-sample parameters.

	OLS	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90
<i>BXM</i>	-0.148***	-0.012	-0.039	-0.056***	-0.077***	-0.096***	-0.131***	-0.134***	-0.155***	-0.188***
<i>BXMC</i>	-0.097***	0.081***	0.056***	0.016	0.005	-0.038***	-0.076***	-0.119***	-0.141***	-0.168***
<i>BXMD</i>	-0.121***	0.012	-0.016	-0.033**	-0.062***	-0.079***	-0.105***	-0.122***	-0.146***	-0.159***
<i>BXY</i>	-0.091***	0.018	-0.007	-0.008	-0.045***	-0.065***	-0.088***	-0.106***	-0.106***	-0.129***
<i>PUT</i>	-0.163***	-0.022	-0.045*	-0.064***	-0.085***	-0.103***	-0.133***	-0.155***	-0.173***	-0.208***
<i>PUTY</i>	-0.219***	-0.083*	-0.114***	-0.125***	-0.118***	-0.156***	-0.181***	-0.202***	-0.241***	-0.283***
<i>CMBO</i>	-0.147***	-0.013	-0.025	-0.052***	-0.076***	-0.096***	-0.128***	-0.145***	-0.169***	-0.190***
<i>BFLY</i>	-0.036	-0.128***	-0.128***	-0.076***	-0.041***	-0.013*	0.021	0.043**	0.039*	-0.013
<i>CNDR</i>	-0.097***	-0.158**	-0.131***	-0.095***	-0.052***	-0.021	0.000	0.022	-0.009	-0.068
<i>CLL</i>	-0.055***	0.159***	0.081***	0.021	-0.026*	-0.067***	-0.101***	-0.136***	-0.157***	-0.149***
<i>PPUT</i>	-0.062***	0.083**	0.045**	0.006	-0.040**	-0.077***	-0.107***	-0.119***	-0.138***	-0.150***
<i>RXM</i>	-0.213***	-0.015	-0.030	-0.066***	-0.101***	-0.162***	-0.195***	-0.225***	-0.272***	-0.301***
<i>VIX</i>	0.002	-0.087	-0.055	-0.029	-0.011	0.005*	0.021	0.038	0.057	0.080
<i>TMS</i>	-0.009	-0.067**	-0.040**	-0.017	-0.007	-0.005	0.007	0.002	0.001	0.060**
<i>TED</i>	-0.099**	-0.972***	-0.606***	-0.362***	-0.209***	-0.051	0.116**	0.212***	0.452***	0.760***
<i>DFY</i>	-0.022	-1.079***	-0.665***	-0.356***	-0.129**	0.029	0.159***	0.308***	0.583***	0.987***

Notes: This table reports the estimated slope coefficients from in-sample univariate predictive regressions of the equity premium against a set of predictors. The first column reports the estimated coefficients from univariate OLS regressions, and the remaining columns report the results from quantile regressions for quantiles ranging from $\tau = 0.10$ to $\tau = 0.90$. The sample period is from January 3, 1996 to April 15, 2021.

*Statistical significance at the 10% level.

**Statistical significance at the 5% level.

***Statistical significance at the 1% level.

coefficients that are statistically significant at the 1% level across all option-related variables (with the exception of BFLY). For context, the coefficient of the VIX is insignificant while TED is the only macroeconomic variable with a statistically significant coefficient.

The quantile regression results confirm the non-linear impact of the predictors on the equity premium distribution. For instance, looking at the results for $\tau = 0.10$, it appears that, even though some option predictors have statistically insignificant coefficients, the left tail of the equity premium distribution is affected by at least one option strategy index from each category group, as well as by the Goyal and Welch (2008) economic variables. As we move further to the right of the equity premium distribution, the importance of option predictors increases substantially, as evidenced by the fact that almost all variables have statistically significant coefficients across all quantiles. In contrast to previous studies, we find that the VIX is largely insignificant as a univariate predictor, both in an OLS and in a quantile setting.

Interestingly, we find that BFLY and CNDR are significant in-sample predictors mainly for quantiles ranging from $\tau = 0.10$ to $\tau = 0.50$, potentially suggesting that volatility premia embedded in option prices are driven by investors' concerns about downside risk. The opposite pattern can be seen in the coefficients of Buy-Write strategies (BXM, BXMC, BXMD, and BXY), indicating that these option variables can more accurately predict the right part of the equity premium distribution compared to the left part. Overall, the results indicate that option variables have a significant non-linear impact on in-sample equity premium predictability.

4.2. Point forecasts

We next examine the out-of-sample predictability of the equity premium. While in-sample forecasts are generally associated with higher statistical power, out-of-sample performance is typically considered a more appropriate evaluation measure as it avoids issues of overfitting and, importantly, it is based on information that would have been available to investors at the time the forecast was generated. Table 4 reports the out-of-sample R-square (R_{OS}^2) of point forecasts under a set of competing forecasting models. In particular, the table reports the performance of forecasts based on univariate and multivariate linear models, as well as those generated by quantile forecast combination schemes. The statistical significance of the R_{OS}^2 is evaluated against a benchmark model using the Clark and West (2007) test. We adopt the natural benchmark of the historical average (HA), which Campbell and Thompson (2008) and Goyal and Welch (2008) argue is hard to beat, as the HA generally outperforms an extensive set of commonly-used predictors out-of-sample.

Table 4 shows that all option strategy indices outperform the benchmark of the historical average, as evidenced by positive out-of-sample R-squares (BFLY is the only exception with an R_{OS}^2 of -0.03). Importantly, this outperformance relative to the HA benchmark is highly significant in most cases, with the majority of the Clark and West (2007) p -values being lower than 1%. The greatest improvement in forecasting power is offered by the risk reversal strategy RXM, which has the highest R_{OS}^2 at 1.91 (p -value = 0.1%). Given that risk reversals reflect the level of implied skewness embedded in index options, it seems that investors' demand for downside protection contains substantial information about subsequent market returns, in excess of the information contained in the historical returns record.

To put these results into context, it is worth highlighting that the three Goyal and Welch (2008) economic variables have a negative R^2 , suggesting that they are inferior predictors of the equity premium relative to the HA benchmark. Therefore, our out-of-sample results provide strong support for the hypothesis that information from the options market can markedly improve the predictability of the equity premium, relative to the historical average benchmark and commonly-used economic predictors.

Table 4
Point forecasts: Out-of-sample R^2 .

Panel A: Individual predictors						
	Linear		Q_3		Q_5	
	R^2_{OS}	p -value	R^2_{OS}	p -value	R^2_{OS}	p -value
<i>BXM</i>	1.26*	0.0032*	1.20*	0.0012*	1.24*	0.0011*
<i>BXMC</i>	0.39*	0.0557*	0.29*	0.0335*	0.25*	0.0420*
<i>BXMD</i>	1.24*	0.0012*	1.13*	0.0006*	1.16*	0.0005*
<i>BXY</i>	0.72*	0.0029*	0.73*	0.0015*	0.75*	0.0012*
<i>PUT</i>	1.37*	0.0035*	1.26*	0.0016*	1.28*	0.0015*
<i>PUTY</i>	1.80*	0.0033*	1.70*	0.0018*	1.70*	0.0017*
<i>CMBO</i>	1.32*	0.0022*	1.21*	0.0011*	1.24*	0.0010*
<i>BFLY</i>	-0.03	0.2510	0.03*	0.0817*	0.04*	0.0831*
<i>CNDR</i>	0.04*	0.1683	0.04*	0.1176	0.06*	0.1011
<i>CLL</i>	0.06*	0.1229	0.22*	0.0065*	0.24*	0.0049*
<i>PPUT</i>	0.18*	0.0142*	0.29*	0.0033*	0.29*	0.0031*
<i>RXM</i>	1.91*	0.0014*	1.75*	0.0008*	1.78*	0.0007*
<i>VIX</i>	-0.31	0.9488	-0.12	0.4388	-0.10	0.3630
<i>TMS</i>	-0.06	0.8602	-0.03	0.2543	-0.03	0.2539
<i>TED</i>	-0.14	0.4216	-0.13	0.4285	-0.19	0.5784
<i>DFY</i>	-0.33	0.5558	-0.19	0.4177	-0.14	0.4157
Mean	0.72*	0.0027*	0.68*	0.0009*	0.69*	0.0009*
Kitchen sink	-0.59	0.0181*	0.23*	0.0073*	0.14*	0.0079*

Panel B: Variable selection and combination schemes			
	R^2_{OS}	p -value	
PredAvg	0.36*	0.0671*	
LASSO	1.37*	0.0151*	
ENet	0.13*	0.0188*	
PCA	1.17*	0.0018*	
sPCA	1.41*	0.0031*	
PLS	1.11*	0.0044*	
PLQC3	1.48*	0.0166*	
PLQC5	1.38*	0.0137*	
TW3	0.98*	0.0191*	
TW5	0.89*	0.0245*	

Notes: This table reports the out-of-sample performance for a number of forecasting models that use a set of option-related and economic predictors. Panel A reports the results for forecasts produced by linear regression models (including univariate models, taking the mean of univariate models' forecasts, and a kitchen sink estimation), as well as quantile combination models (Q_3 and Q_5). Panel B reports the results produced by variable selection and combination schemes, namely predictor average (PredAvg), LASSO, elastic net (ENet), principal component analysis (PCA), scaled PCA (sPCA), partial least squares (PLS), and post LASSO quantile combination (PLQC) models with equally-weighted schemes (PLQC3 and PLQC5) and schemes with time-varying weights (TW3 and TW5). The table reports the models' out-of-sample R-square (R^2_{OS}) and the respective p -value based on the Clark and West (2007) test. The Clark and West (2007) test evaluates each model's forecast accuracy against the benchmark of the historical average.

*Model outperformance against the benchmark (i.e., $R^2_{OS} > 0$) is denoted by *. Statistical significance at the 10% level (i.e., p -value < 0.10) is denoted by *.

Furthermore, we find that taking the mean of the set of forecasts produced by the univariate linear models also improves the forecastability of the equity premium, relative to the HA benchmark (see also Rapach et al., 2010). As can be seen in Table 4, the R^2_{OS} of this forecast mean is positive (0.72) and highly significant (p -value = 0.3%). With an R^2_{OS} of 0.36, using the mean of all predictors as a single forecasting variable (PredAvg) also outperforms the historical average, but to a lesser extent compared to averaging across forecasts. By contrast, including all variables in a single kitchen sink model results in a markedly lower forecasting power, with a negative R^2_{OS} of -0.59.

In Fig. 2, we plot the time-series of the differences between the cumulative mean squared forecast error of the HA benchmark minus that of each univariate linear model. To conserve space, we present the plots of the six best-performing predictors (RXM, PUTY, PUT, CMBO, BXM, and BXMD). In general, the subplots show that forecasts based on the option strategy indices consistently outperform the historical average benchmark. The forecasting performance of option predictors is somewhat modest, and some even underperform slightly, in the beginning of the evaluation period. However, all predictors exhibit substantially increasing predictive ability post-2008, and they experience another noticeable improvement towards the end of the sample period.

The role of option variables in forecasting the equity premium is further supported by the results of aggregating forecasts across quantiles. More specifically, the three-quantile (Q_3) and the five-quantile (Q_5) estimators produce forecasts that significantly outperform the HA benchmark across all option variables (including the BFLY, which underperformed in the linear setting). The option variables' positive R^2 s are generally highly significant, with p -values that are universally lower than those produced by

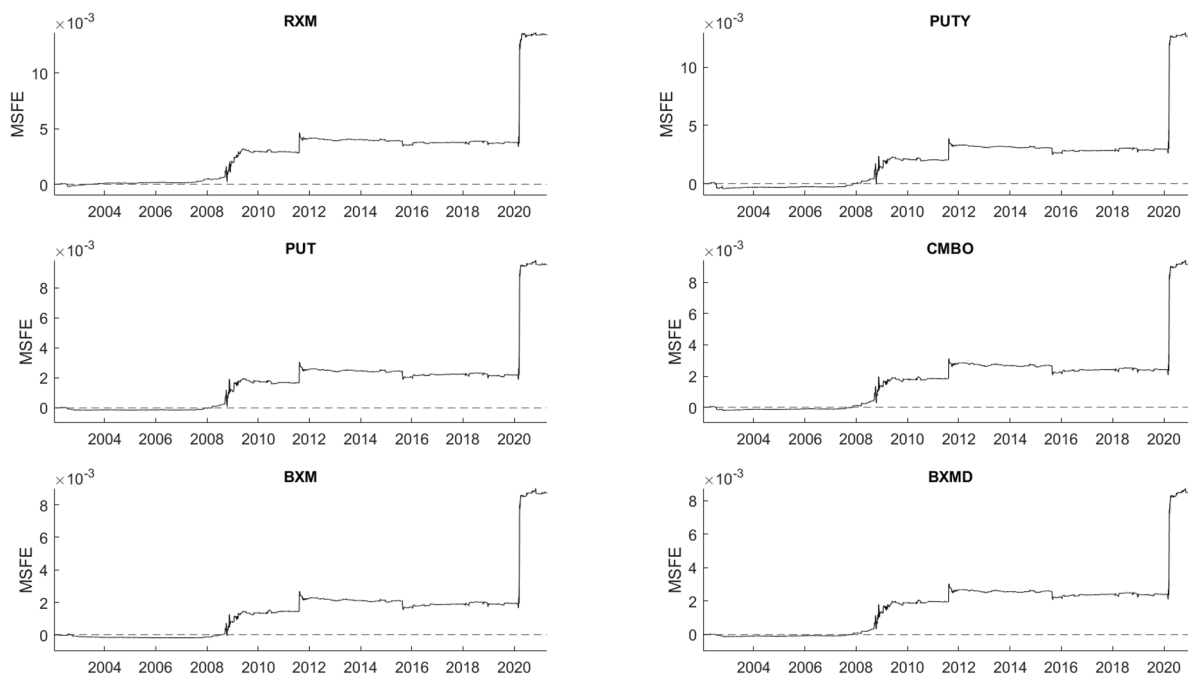


Fig. 2. Cumulative MSFE — Benchmark minus univariate linear models.

Notes: We plot the cumulative mean squared forecast error of the historical average benchmark minus that of a selection of univariate linear models. The sample period is from January 3, 1996 to April 15, 2021, with the out-of-sample period starting in January 3, 2001.

the respective linear models. Interestingly, the choice between aggregating across three or five quantiles does not appear to be particularly important, as both estimators produce comparable R^2 s, in terms of magnitude as well as statistical significance.

Finally, the results from variable selection schemes are equally strong (Panel B of Table 4). All the schemes that we examine produce forecasts with an R^2 that is positive and significant at the 5% level (with the only exception being PredAvg, which is significantly positive at the 10% level). The greatest improvement relative to the HA benchmark is offered by the equally-weighted PLQC3 scheme ($R_{OS}^2 = 1.48$, p -value = 1.7%), followed by the sPCA and PLQC5 schemes (R_{OS}^2 of 1.41 and 1.38, respectively). LASSO forecasts also substantially outperform the benchmark ($R_{OS}^2 = 1.37$), while ENet results in a relatively low, albeit still significantly positive, out-of-sample R-square. Consistent with Huang et al. (2021) and Chen et al. (2022), we find that using information from the equity premium to guide dimension reduction leads to more accurate forecasts, as evidenced by the sPCA outperforming the standard PCA. Finally, the results suggest that the equally-weighted PLQC forecasts (PLQC₃ and PLQC₅) are associated with a greater improvement in forecast accuracy compared to the PLQC schemes with time-varying weights (TW₃ and TW₅).

4.3. Quantile forecasts

In this subsection, we evaluate the accuracy of quantile forecasts. Table 5 reports the t -statistics of the quantile scores (Q_S) produced by the set of competing models, using Eqs. (19) and (20). We examine the forecast accuracy of the competing models across the equally-spaced quantiles $\tau = \{0.05, 0.10, \dots, 0.95\}$, with the table reporting the results for the respective deciles, for brevity. We evaluate the predictive ability of each model against two different benchmarks. The first benchmark (Panel A of Table 5) refers to quantile forecasts generated by the historical average (HA) model. The second benchmark (Panel B of Table 5) refers to a simple quantile constant (QC) model, obtained by setting $\beta^{(\tau)}x_t = 0$ in Eq. (7).

Panel A of Table 5 shows that all the predictors outperform the HA benchmark across all the quantiles of the equity premium distribution, as evidenced by universally negative and highly significant QS t -statistics. The mean forecast model offers the most significant improvement, while the kitchen sink model also generally outperforms the historical average in contrast to the earlier results from point forecasts (Table 4). Generating PLQ forecasts via Eq. (12) significantly improves forecasting performance as well, with the respective quantile scores being more significant than those associated with the majority of univariate models.

The results in Panel B of Table 5 show that the QC benchmark is harder to beat. When we consider the left tail of the equity premium distribution ($\tau \leq 0.3$), almost all option strategy predictors underperform relative to the benchmark, with quantile scores that are either positive or negative but statistically insignificant. By contrast, the VIX performs very well across almost all quantiles. Nevertheless, combining information across the entire set of predictors still improves forecast accuracy, with significantly negative Q_S t -statistics associated with the mean forecast, the kitchen sink model, and the PLQ forecasts in the lower quantiles. Interestingly, as we move towards the right part of the distribution, the majority of the options predictors begin to outperform the QC benchmark.

Table 5
Quantile forecasts: QS t -statistics.

Panel A: HA quantile benchmark									
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
<i>BXM</i>	-4.58*	-10.87*	-13.11*	-11.11*	-4.03*	-4.09*	-8.22*	-10.11*	-8.27*
<i>BXMC</i>	-4.50*	-11.05*	-13.15*	-11.18*	-2.36*	-2.70*	-7.47*	-9.24*	-7.05*
<i>BXMD</i>	-4.86*	-10.95*	-13.09*	-11.14*	-4.26*	-4.69*	-9.00*	-10.46*	-8.86*
<i>BXY</i>	-4.57*	-10.65*	-12.74*	-10.69*	-3.15*	-3.93*	-8.62*	-10.30*	-8.89*
<i>PUT</i>	-5.18*	-11.17*	-13.31*	-11.08*	-3.51*	-3.43*	-7.99*	-9.44*	-7.79*
<i>PUTY</i>	-5.85*	-11.12*	-13.03*	-10.65*	-3.05*	-3.55*	-7.73*	-8.89*	-7.44*
<i>CMBO</i>	-4.91*	-11.07*	-12.94*	-10.86*	-3.99*	-4.33*	-8.62*	-9.90*	-8.14*
<i>BFLY</i>	-5.34*	-10.82*	-13.06*	-11.16*	-2.87*	-2.92*	-8.01*	-9.37*	-7.77*
<i>CNDR</i>	-5.51*	-10.82*	-12.92*	-11.01*	-3.05*	-2.76*	-7.80*	-9.02*	-8.10*
<i>CLL</i>	-4.81*	-10.96*	-12.98*	-11.12*	-4.85*	-5.94*	-10.10*	-10.81*	-9.08*
<i>PPUT</i>	-4.43*	-10.75*	-12.90*	-11.09*	-4.39*	-5.47*	-9.98*	-11.13*	-9.35*
<i>RXM</i>	-5.11*	-11.26*	-13.26*	-10.97*	-3.92*	-4.61*	-8.70*	-9.24*	-7.35*
<i>VIX</i>	-11.31*	-13.26*	-13.44*	-11.06*	-2.09*	-5.38*	-10.28*	-13.20*	-14.08*
<i>TMS</i>	-4.70*	-11.60*	-13.04*	-11.28*	-2.79*	-2.98*	-7.91*	-9.51*	-8.33*
<i>TED</i>	-8.27*	-12.42*	-13.31*	-10.96*	-2.02*	-1.77*	-6.89*	-8.74*	-8.90*
<i>DFY</i>	-6.29*	-11.08*	-12.01*	-10.37*	-1.49*	-3.98*	-10.34*	-12.23*	-10.64*
Mean	-14.68*	-15.60*	-15.11*	-12.14*	-4.75*	-6.62*	-11.55*	-14.55*	-16.97*
Kitchen sink	-9.59*	-11.08*	-9.62*	-5.22*	0.07*	-3.12*	-8.41*	-12.19*	-13.39*
PLQ	-9.50*	-10.81*	-10.91*	-10.37*	-2.67*	-5.44*	-9.95*	-12.60*	-13.51*
Panel B: Constant quantile benchmark									
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
<i>BXM</i>	1.89	1.21	-1.21	-2.33*	-2.66*	-2.57*	-3.07*	-2.58*	-2.39*
<i>BXMC</i>	-0.01	-0.17	0.02	0.25	0.15	-0.56	-1.22	-1.61*	-1.22
<i>BXMD</i>	1.53	1.24	-0.91	-2.53*	-2.90*	-3.23*	-4.10*	-3.13*	-2.81*
<i>BXY</i>	1.59	1.45	1.55	-0.23	-1.51	-2.25*	-3.34*	-2.74*	-2.66*
<i>PUT</i>	1.53	0.53	-2.32*	-2.25*	-2.06*	-1.87*	-2.75*	-1.94*	-2.06*
<i>PUTY</i>	0.16	-0.78	-1.91*	-2.05*	-1.85*	-2.23*	-2.78*	-1.91*	-2.16*
<i>CMBO</i>	1.94	0.92	-0.11	-2.07*	-2.65*	-2.90*	-3.68*	-2.56*	-2.41*
<i>BFLY</i>	-0.81	-1.44	-0.65	-0.32	0.47	0.38	-0.38	-1.11	1.31
<i>CNDR</i>	-0.02	-0.16	-0.46	0.00	0.46	0.84	1.21	1.75	-0.17
<i>CLL</i>	0.00	0.02	0.86	-0.53	-3.61*	-4.64*	-5.76*	-4.82*	-2.96*
<i>PPUT</i>	1.16	0.91	1.15	-0.98	-2.95*	-3.96*	-5.41*	-4.66*	-3.11*
<i>RXM</i>	1.26	-0.05	-1.69*	-2.39*	-2.89*	-3.43*	-4.21*	-3.10*	-2.54*
<i>VIX</i>	-9.37*	-7.49*	-5.17*	-1.99*	0.47	-4.20*	-6.77*	-8.56*	-10.24*
<i>TMS</i>	0.85	0.89	1.65	1.32	1.37	0.36	1.28	1.12	-0.28
<i>TED</i>	-6.17*	-6.26*	-5.19*	-3.06*	1.00	0.39	-1.32	-3.06*	-4.47*
<i>DFY</i>	-4.27*	-3.27*	-1.41	-0.16	1.20	-1.32	-2.97*	-3.72*	-5.04*
Mean	-9.99*	-9.78*	-6.90*	-4.16*	-3.34*	-5.55*	-8.46*	-9.20*	-9.53*
Kitchen sink	-8.04*	-5.86*	-2.83*	0.26	0.72	-2.39*	-5.52*	-7.96*	-9.71*
PLQ	-7.96*	-5.94*	-3.69*	-1.35	-1.41	-4.48*	-6.85*	-8.38*	-9.91*

Notes: This table reports the quantile score t -statistics for a number of forecasting models that use a set of option-related and economic predictors. The QS t -statistics are computed following [Giacomini and White \(2006\)](#) and [Amisano and Giacomini \(2007\)](#). The null hypothesis is that the QS of the candidate model is equal to that of the benchmark (equal forecast accuracy). A negative t -statistic indicates that the candidate model outperforms the benchmark, while a positive t -statistic indicates the opposite. Panel A reports the results when the performance of candidate models is evaluated against the benchmark of the Historical Average (HA). Panel B reports the results when the performance of candidate models is evaluated against the Quantile Constant (QC) model.

*Statistical significance at the 5% level.

More specifically, all strategy benchmark indices have significantly negative QS t -statistics for quantiles $\tau \geq 0.5$, with the exception of *BXMC*, *BFLY*, and *CNDR*.

Overall, these results confirm the role of options predictors in forecasting the equity premium. Single-predictor models seem to significantly improve forecast accuracy in the right tail of the distribution, while aggregation schemes result in significantly improved forecasts across the entire distribution of the equity premium.

4.4. Density forecasts

In this subsection, we examine the forecastability of the entire distribution. [Table 6](#) reports the t -statistics of the weighted quantile scores (WQS) associated with each candidate model, under four different weighting schemes that are intended to place different emphasis on specific regions of the distribution. Similarly to the previous subsection, we focus on the quantiles $\tau = \{0.05, 0.10, \dots, 0.95\}$, and we evaluate forecast accuracy relative to the HA and the QC benchmarks.

When we evaluate performance against the HA benchmark, we find that all option predictors significantly improve the accuracy of density forecasts under all four weighting schemes. While the relative importance of specific variables somewhat varies across

Table 6
Density forecasts: *WQS* *t*-statistics.

	Benchmark: HA				Benchmark: QC			
	<i>WQS</i> ₁	<i>WQS</i> ₂	<i>WQS</i> ₃	<i>WQS</i> ₄	<i>WQS</i> ₁	<i>WQS</i> ₂	<i>WQS</i> ₃	<i>WQS</i> ₄
<i>BXM</i>	-15.79*	-16.07*	-12.45*	-10.68*	-2.65*	-2.75*	-0.86	-2.81*
<i>BXMC</i>	-16.15*	-17.05*	-11.87*	-9.57*	-1.25	-1.19	-0.08	-1.27
<i>BXMD</i>	-16.33*	-16.66*	-12.60*	-11.12*	-3.28*	-3.40*	-1.34	-3.34*
<i>BXY</i>	-16.42*	-16.58*	-12.05*	-11.32*	-2.20*	-2.11*	0.48	-2.76*
<i>PUT</i>	-14.97*	-15.25*	-12.60*	-9.94*	-2.35*	-2.37*	-1.42	-2.41*
<i>PUTY</i>	-13.06*	-13.23*	-12.07*	-8.94*	-2.32*	-2.32*	-1.63	-2.36*
<i>CMBO</i>	-15.39*	-15.65*	-12.58*	-10.37*	-2.81*	-2.90*	-1.3	-2.91*
<i>BFLY</i>	-17.05*	-17.74*	-11.82*	-11.60*	-0.70	-0.76	-0.90	0.33
<i>CNDR</i>	-17.07*	-17.66*	-11.93*	-11.53*	0.21	0.24	-0.09	0.57
<i>CLL</i>	-18.87*	-19.58*	-12.24*	-13.48*	-4.44*	-4.94*	-0.64	-4.74*
<i>PPUT</i>	-18.47*	-18.84*	-12.14*	-13.28*	-4.19*	-4.39*	-0.31	-4.50*
<i>RXM</i>	-13.92*	-14.27*	-12.90*	-9.27*	-3.39*	-3.53*	-2.52*	-3.27*
<i>VIX</i>	-23.60*	-23.91*	-14.61*	-14.76*	-16.02*	-15.40*	-9.51*	-10.13*
<i>TMS</i>	-18.12*	-18.67*	-12.43*	-12.08*	1.71	1.96	1.23	0.86
<i>TED</i>	-18.19*	-18.16*	-13.12*	-11.27*	-8.24*	-6.89*	-6.53*	-4.55*
<i>DFY</i>	-17.44*	-18.39*	-9.88*	-14.26*	-6.25*	-5.04*	-3.73*	-5.31*
Mean	-25.66*	-23.95*	-19.42*	-19.16*	-13.23*	-11.67*	-11.29*	-9.35*
Kitchen sink	-14.70*	-12.18*	-11.03*	-12.66*	-9.61*	-7.04*	-6.98*	-8.71*
PLQ	-19.77*	-19.65*	-12.04*	-13.89*	-13.40*	-12.32*	-7.77*	-9.76*

Notes: This table reports the weighted quantile score *t*-statistics for a number of forecasting models that use a set of option-related and economic predictors. Each model's *WQS* is computed under four different weighting schemes, namely (1) *WQS*₁ : $\omega(\tau) = 1$, (2) *WQS*₂ : $\omega(\tau) = \tau(1 - \tau)$, (3) *WQS*₃ : $\omega(\tau) = (1 - \tau)^2$, and *WQS*₄ : $\omega(\tau) = \tau^2$. The *WQS* *t*-statistics are computed following [Giacomini and White \(2006\)](#) and [Amisano and Giacomini \(2007\)](#). The null hypothesis is that the *WQS* of the candidate model is equal to that of the benchmark (equal forecast accuracy). A negative *t*-statistic indicates that the candidate model outperforms the benchmark, while a positive *t*-statistic indicates the opposite. The first four columns report the results when the performance of candidate models is evaluated against the benchmark of the Historical Average (HA), while the last four columns report the results when the performance of candidate models is evaluated against the Quantile Constant (QC) model.

*Statistical significance at the 5% level.

weighting schemes, *VIX*, *CNDR*, *CLL*, and *PPUT* are generally associated with the most significant *WQS* values. Interestingly, the mean forecast results in consistently greater improvements compared to that offered by univariate models, while the kitchen sink model offers a less significant improvement.

Consistent with the quantile forecast results in [Table 5](#), improving forecast accuracy is more difficult when evaluated against the QC benchmark. Nevertheless, the last four columns of [Table 6](#) show that the majority of the options predictors still result in more accurate density forecasts, while the *VIX* still results in the highest improvement. For instance, when we place uniform emphasis across the distribution under the equally-weighted scheme *WQS*₁, we find that 9 out of 12 strategy benchmark indices significantly outperform the QC benchmark, with *CCL* and *PPUT* being the most significant. The results are very similar when we place greater emphasis on the middle or the right tail of the distribution, under *WQS*₂ and *WQS*₄, respectively.

The results under *WQS*₃ are the only notable exception, with *RXM* the only strategy benchmark predictor with a significantly negative weighted quantile score. This finding suggests that using option variables in a univariate setting is less likely to improve the accuracy of forecasts of the left tail of the equity premium distribution. However, forecasts obtained by using information from the full set of option predictors still outperform the QC benchmark under all four weighting schemes, including *WQS*₃. In fact, the mean forecast model offers its most significant improvement in forecast accuracy under *WQS*₃, suggesting that combinations of individual forecasts are especially useful when forecasting the left tail of the distribution. Finally, the performance of *PLQ* forecasts is consistently somewhere between that of the mean forecast and the kitchen sink model.

4.5. Economic evaluation

In this subsection, we evaluate the economic performance of competing forecast models. [Campbell and Thompson \(2008\)](#) argue that even small predictability gains, in a statistical sense, could translate to an economically meaningful increase in predictability, resulting in an increase in portfolio returns for a mean–variance investor who maximizes expected utility. We follow this utility-based approach within a standard asset allocation framework to evaluate the economic performance of competing models in a way that captures an investor's risk–return trade-off.

More specifically, we consider a mean–variance investor who optimally allocates their wealth between equities and the risk-free asset based on forecasts of the equity premium. In the standard mean–variance framework, the solution to this maximization problem yields the following weight w_t for the investor's wealth to be invested in the risky asset at time t :

$$w_t = \frac{\hat{r}_{t+1}}{\gamma \hat{\sigma}_{t+1}^2}, \quad (22)$$

Table 7
Economic evaluation.

Panel A: Individual predictors									
	Linear			Q_3			Q_5		
	\bar{r}_p	CER	SR	\bar{r}_p	CER	SR	\bar{r}_p	CER	SR
<i>BXM</i>	9.55	4.08	0.50	12.55	5.59	0.58	12.66	5.59	0.58
<i>BXMC</i>	6.14	0.55	0.29	11.19	3.56	0.49	11.48	3.71	0.50
<i>BXMD</i>	11.27	5.59	0.58	14.12	6.92	0.64	14.44	7.16	0.66
<i>BXY</i>	9.43	4.03	0.49	11.32	4.42	0.53	11.66	4.55	0.53
<i>PUT</i>	7.89	2.24	0.39	11.65	4.43	0.53	11.76	4.58	0.53
<i>PUTY</i>	9.22	3.38	0.46	12.78	5.26	0.57	12.85	5.25	0.57
<i>CMBO</i>	10.81	4.96	0.55	12.94	5.50	0.58	12.99	5.68	0.59
<i>BFLY</i>	3.60	-0.71	0.17	6.56	-0.35	0.30	6.65	-0.33	0.30
<i>CNDR</i>	4.57	-0.11	0.22	7.79	0.64	0.35	7.88	0.62	0.34
<i>CLL</i>	2.32	-1.21	0.09	6.20	0.43	0.32	6.58	0.35	0.33
<i>PPUT</i>	8.03	2.83	0.42	10.36	3.33	0.48	10.59	3.33	0.48
<i>RXM</i>	11.26	5.16	0.56	14.17	5.89	0.60	14.35	6.00	0.60
<i>VIX</i>	-0.34	-7.82	-0.08	5.67	-0.91	0.23	5.62	-1.26	0.22
<i>TMS</i>	2.46	-0.69	0.11	5.54	0.18	0.29	5.12	-0.21	0.27
<i>TED</i>	5.42	0.40	0.26	4.70	0.32	0.33	4.59	-0.36	0.30
<i>DFY</i>	2.51	0.46	0.17	7.36	2.03	0.36	7.74	2.25	0.38
Mean	8.20	3.43	0.45	13.36	6.58	0.63	13.23	6.13	0.61
Kitchen sink	8.60	2.05	0.39	9.40	2.39	0.42	9.37	2.14	0.41

Panel B: Variable selection and combination schemes			
	\bar{r}_p	CER	SR
PredAvg	4.67	0.93	0.26
LASSO	8.44	1.86	0.36
Enet	9.32	2.33	0.41
PCA	11.00	4.54	0.52
sPCA	10.65	4.03	0.49
PLS	9.30	3.16	0.44
PLQC3	10.77	4.80	0.54
PLQC5	10.36	4.24	0.50
TW3	11.48	5.54	0.58
TW5	11.56	5.84	0.60

Notes: This table reports the economic performance of a mean–variance investor who allocates their wealth between equities and the risk-free asset based on forecasts of the equity premium that have been generated by a set of competing models. Panel A reports the results for forecasts produced by linear regression models, including univariate models, taking the mean of univariate models’ forecasts, and a Kitchen Sink estimation. Panel B reports the results produced by variable selection and combination schemes. The table reports the mean daily return \bar{r}_p , earned by the investor’s portfolio (annualized, in percentages), the associated certainty equivalent return (CER, annualized in percentages), and the portfolio’s Sharpe ratio (SR, annualized). The CER has been computed using a risk aversion parameter $\gamma = 3$. For comparison, the CER and the SR of the historical average benchmark are equal to 1.27 and 0.12, respectively.

where \hat{r}_{t+1} is the equity premium forecast generated at time t for $t+1$, γ is the relative risk aversion (RRA) coefficient that determines the investor’s appetite for risk, and $\hat{\sigma}_{t+1}^2$ is the forecasted volatility of the equity premium at $t+1$. The volatility forecast at each point in time is generated as the conditional standard deviation between symmetric quantiles, given by $\frac{\hat{r}_{t+1}^{(0.95)} - \hat{r}_{t+1}^{(0.05)}}{3.25}$ (see also Meligkotsidou et al., 2021). Following the extant literature, we set the risk aversion coefficient γ equal to 3. At each point in time, the investor allocates w_t of their wealth in the equity index and the remaining $(1 - w_t)$ in risk-free T-bills. The portfolio’s return at $t + 1$ is given by:

$$r_{p,t+1} = w_t r_{t+1} + (1 - w_t) r_{f,t+1}. \tag{23}$$

Table 7 presents the economic performance of the mean–variance investor who trades based on equity premium forecasts generated by a set of competing models. The table reports each portfolio’s annualized mean daily return, certainty equivalent return (CER), and Sharpe ratio (SR). The portfolio’s CER and SR are computed as:

$$CER = \bar{r}_p - \frac{1}{2} \gamma \sigma_p^2, \tag{24}$$

$$SR = \frac{\bar{r}_p - r_f}{\sigma_p}, \tag{25}$$

where \bar{r}_p and σ_p denote the portfolio’s mean and standard deviation of returns, respectively, during the forecast evaluation period.

The results in Panel A show that all option strategy indices result in portfolios with positive mean returns in the univariate linear setting. Importantly, these portfolios consistently outperform the one based on the HA benchmark, which offers a mean return of 2.16% per annum. The highest performing portfolios are based on forecasts computed using the buy-write strategy BXMD (11.27%)

and the risk reversal strategy RXM (11.26%). Interestingly, exactly half the option predictors outperform the mean forecast model (8.63%), while the portfolio based on the kitchen sink model forecasts offers a somewhat lower mean return (7.01%).

The ranking among competing models is fairly similar when we account for the portfolios' risk and investors' risk aversion. More specifically, forecasts based on BXMD result in the highest CER (5.59% per annum), followed by RXM with 5.16%. These two predictors also result in portfolios with the highest Sharpe ratios (0.58 and 0.56, respectively). In other words, a mean–variance investor would maximize their returns (non-adjusted and adjusted for risk) by exploiting the informational content of strategies based on hedging the index and selling insurance against index skewness. For comparison, a portfolio based on forecasts produced by the HA benchmark offers a CER of 1.27% and a Sharpe ratio of 0.12. As such, the benchmark's risk-adjusted performance seems to be inferior to that of the majority of univariate linear models, given that 8 out of 12 predictors offer higher CERs and 11 out of 12 offer higher SRs.

The results are even stronger when point forecasts are generated by aggregating across quantiles. More specifically, both the Q3 and Q5 estimators result in portfolios that outperform the ones based on linear models in terms of CER and SR for every option predictor. Similarly, portfolios constructed based on variable selection, dimensionality reduction, and combination schemes (Panel B of Table 7) often outperform those based on linear univariate models. For instance, the time-varying-weighted PLQC schemes TW3 and TW5 offer the highest mean returns (11.48% and 11.56%, respectively), Sharpe ratios (0.58 and 0.60, respectively), and CERs (5.54% and 5.84%, respectively).

To explore the stability, or otherwise, of these portfolios, in Figs. A3 and A4 in the Online Appendix we plot the time evolution of the mean–variance investor's cumulative returns under a set of competing models. Fig. A3 refers to portfolios using univariate linear models, while Fig. A4 refers to portfolios based on variable selection/shrinkage approaches. For comparison purposes, each subplot also shows the cumulative returns of the portfolio based on forecasts generated by the HA benchmark.

Figs. A3 and A4 show that the economic performance of option predictors is relatively stable and, importantly, it does not reflect an empirical relationship that prevailed in the distant past. In fact, forecasts based on the CBOE indices performed rather modestly in the first few years of the evaluation period, followed by a period of noticeable outperformance that starts around 2010. Consistent with the performance measures reported in Table 7, the highest performance is offered by variable selection models with time-varying weights (TW3 and TW5), followed by the BXMD and RXM univariate models.

4.6. Robustness

We perform a set of additional empirical estimations to determine the robustness of our main findings regarding the forecasting ability of option variables over the equity premium. We briefly summarize the results of these robustness checks, with the full results reported in Tables A1–A8 in the Online Appendix.

1. *Risk aversion:* We re-examine the economic significance of options-based forecasts using alternative risk aversion coefficients. The main results in Table 7 are based on the risk aversion coefficient γ taking the value of 3, in line with the most common choice in the literature. We re-evaluate the economic performance of competing forecast models using the alternative value of $\gamma = 5$ (Table A2 in the Online Appendix). We also use the Bekaert et al. (2021) risk aversion index as an alternative proxy for investors' risk aversion (Table A3 in the Online Appendix). The results confirm that option variables can deliver significant economic gains, with most option variables outperforming the HA benchmark in quantile aggregation schemes, and PLQC quantile combination schemes offering the highest economic performance overall.
2. *Volatility forecasts:* The economic evaluation results in Table 7 are potentially sensitive to the way in which volatility forecasts $\hat{\sigma}_{t+1}^2$ are generated. We further explore the robustness of our economic evaluation results by generating alternative volatility forecasts via the simplest and arguably most commonly adopted approach in the literature. More specifically, at each point in time we construct volatility forecasts as the standard deviation of portfolio returns over the previous five years. The results are reported in Table A4 in the Online Appendix, and they are qualitatively similar to the ones reported in the main analysis.
3. *Business cycle:* Table 8 and Tables A5–A9 in the Online Appendix present the results of an evaluation of the statistical gains of competing forecast models for recession and expansion periods as defined by the NBER. Consistent with Rapach et al. (2010), Neely et al. (2014), and Li and Tsiakas (2017), among others, we find that equity premium predictability is generally stronger during recessions. For instance, point forecasts tend to have a higher R_{OS}^2 , which is more likely to be statistically significant in recessions compared to expansions (Table 8 and Tables A5–A6 in the Online Appendix). Furthermore, during recessions quantile forecasts tend to have higher and more significant quantile scores QS when compared against the tougher-to-beat QC benchmark (Tables A7–A8 in the Online Appendix). Somewhat surprisingly, we find that during expansions, options variables have a more significant QS relative to the HA benchmark, which is nonetheless easier to beat compared to the QC benchmark. A similar picture emerges from the WQS under the density forecast evaluation (Table A9 in the Online Appendix), with option variables being more likely to outperform the QC benchmark during recessions, but more likely to outperform the HA benchmark during expansions. Interestingly, PPUT and RXM perform significantly better in recessionary periods compared to expansions, as evidenced by the substantially higher out-of-sample R-squares. These results are consistent with the intuition that predictors which capture investors' beliefs about crash risk play a more important role during market downturns.

Table 8
Business cycle.

Panel A: Individual predictors										
	Recessions					Expansions				
	R_{OS}^2	WQS_1	WQS_2	WQS_3	WQS_4	R_{OS}^2	WQS_1	WQS_2	WQS_3	WQS_4
<i>BXM</i>	2.73*	1.68	0.99	4.87	-0.02	0.07*	-20.97*	-20.38*	-16.31*	-14.30*
<i>BXMC</i>	1.30*	5.06	4.86	5.41	1.83	-0.33	-22.20*	-22.01*	-17.28*	-13.73*
<i>BXMD</i>	2.61*	1.79	1.15	5.15	-0.06	0.12*	-20.93*	-20.62*	-16.44*	-13.64*
<i>BXY</i>	1.75*	3.30	2.51	5.90	0.65	-0.11	-20.91*	-20.17*	-16.00*	-14.63*
<i>PUT</i>	3.08*	1.26	0.67	4.29	-0.17	0.00	-19.95*	-19.55*	-15.93*	-13.09*
<i>PUTY</i>	3.95*	0.35	-0.05	2.67	-0.57	0.10*	-19.10*	-18.42*	-15.98*	-12.95*
<i>CMBO</i>	2.81*	1.52	0.91	4.79	-0.09	0.12*	-20.73*	-20.07*	-16.57*	-14.08*
<i>BFLY</i>	0.32*	7.93	6.86	4.53	5.76	-0.29	-21.88*	-21.72*	-15.40*	-15.79*
<i>CNDR</i>	0.15*	8.47	7.61	5.30	5.59	-0.05	-22.59*	-22.32*	-16.21*	-15.88*
<i>CLL</i>	0.23*	7.60	6.20	6.01	3.69	-0.08	-23.54*	-23.46*	-16.22*	-17.02*
<i>PPUT</i>	0.53*	6.06	4.56	6.11	2.58	-0.10	-23.60*	-22.89*	-16.70*	-17.27*
<i>RXM</i>	4.12*	-0.11	-0.59	3.54	-1.07	0.16*	-19.58*	-18.77*	-16.40*	-13.37*
<i>VIX</i>	-0.64	-6.64*	-6.01*	-4.25*	-3.82*	-0.07	-27.25*	-25.82*	-18.23*	-19.85*
<i>TMS</i>	0.01*	9.79	9.39	5.54	5.42	-0.11	-21.83*	-21.92*	-14.98*	-15.26*
<i>TED</i>	-0.26	-3.86*	-1.67*	-3.13*	-2.41*	0.00	-17.91*	-18.38*	-13.15*	-11.19*
<i>DFY</i>	-0.58	-3.90*	-2.85*	-2.52*	-3.28*	-0.11	-24.84*	-24.09*	-16.67*	-20.84*
Mean	1.58*	-2.10*	-1.11	-1.34	-2.39*	0.10	-27.50*	-25.78*	-19.81*	-21.45*
Kitchen sink	4.43*	-5.09*	-4.08*	-3.86*	-3.92*	-4.27	-16.52*	-12.88*	-14.09*	-16.15*

Panel B: Variable selection and combination schemes					
	Recessions		Expansions		
	R_{OS}^2	<i>p</i> -value	R_{OS}^2	<i>p</i> -value	
PredAvg	1.12*	0.0503*	-0.21	0.6776	
LASSO	3.26*	0.0228*	-0.01	0.1363	
ENet	4.26*	0.0226*	-2.88	0.2525	
PCA	2.68*	0.0046*	0.07*	0.0753*	
sPCA	3.36*	0.0580*	-0.02	0.1183	
PLS	3.33*	0.0041*	-0.52	0.331	
PLQC3	3.52*	0.0211*	-0.02	0.0696	
PLQC5	3.24*	0.0255*	0.02*	0.0677*	
TW3	2.85*	0.0318*	-0.39	0.1453	
TW5	2.67*	0.0368*	-0.41	0.1798	

Notes: This table reports the out-of-sample performance for a number of forecasting models that use a set of option-related and economic predictors. The results are reported separately for recession and expansion periods, as defined by NBER. Panel A reports the results for forecasts produced by linear regression models (including univariate models, taking the mean of univariate models' forecasts, and a kitchen sink estimation). Panel A reports the models' out-of-sample R-square (R_{OS}^2) and the weighted quantile score *t*-statistics under four different weighting schemes. Panel B reports the R_{OS}^2 and the respective *p*-values produced by variable selection and combination schemes.

*Panel A: Statistical outperformance of a model's R_{OS}^2 relative to that of the historical average benchmark at the 10% level (based on the Clark and West, 2007 test) is denoted by *. Statistical outperformance of a model's *WSQ* relative to that of the historical average benchmark at the 10% level (based on the Giacomini and White, 2006; Amisano and Giacomini, 2007 test) is denoted by *.

Panel B: Statistical outperformance relative to the historical average benchmark at the 10% level (i.e., $R_{OS}^2 > 0$ and *p*-value < 0.10) is denoted by *.

4. *Variance risk premium:* Previous studies show that the VRP is an efficient predictor of the equity premium (e.g., Bollerslev et al., 2014; Buss et al., 2017, 2019). To explore whether the options strategy indices contain information that is incremental to that of the VRP, we add the latter to our set of option predictors (for the period January 3, 1996 to December 29, 2017). Tables A10–A12 in the Online Appendix show that several option variables have a higher forecasting performance compared to the VRP. For example, 6 out of 12 strategy indices produce out-of-sample R-squares that are higher than that of the VRP in the univariate linear setting, with a similar pattern in the Q3 and Q5 combination models (Table A10). Moreover, the VRP tends to have a higher *QS* *t*-statistic in the left part of the equity premium distribution, but the CBOE indices generally outperform in the right part of the distribution (Table A11). Finally, virtually all option strategies outperform the VRP in terms of density forecasting, as evidenced by larger *WQS* *t*-statistics in Table A12.

5. Conclusion

The predictability of the equity premium has received a substantial amount of attention in the literature. Despite an original consensus that the equity premium can be forecasted using a set of standard economic variables, Goyal and Welch (2008) demonstrate that these predictors performed rather poorly after the 1970s. In this study, we contribute to the ongoing debate on equity premium predictability by examining the forecasting performance of variables from the options market.

Given that option contracts are forward looking by construction, option variables are natural candidates when considering potential predictors of the equity premium. Our results provide strong support for this intuitive hypothesis. In contrast to the limited, or at least time-varying, predictive ability of standard economic variables (Goyal and Welch, 2008; Baetje and Menkhoff, 2016), we find that using a set of CBOE strategy benchmark indices based on S&P 500 options consistently results in significant improvements in forecasting performance.

More specifically, we apply a range of approaches to generate point, quantile, and density forecasts. These forecasts are based on univariate and multivariate linear models, variable selection schemes, and quantile forecast combination schemes. The results show that forecasts generated by option variables significantly outperform the historical average benchmark, across all the different frameworks we examine. In addition to a highly significant improvement in statistical accuracy, we find that option-based forecasts result in substantial economic gains for a standard mean–variance investor, markedly higher than those associated with the historical average benchmark. Considering how notoriously hard it is to consistently beat the historical average (Campbell and Thompson, 2008), our results strongly support the use of information from the options market to significantly improve equity premium predictability.

Data availability

Main options data is publicly available on the CBOE webpage

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.finmar.2022.100801>.

References

- Adämmer, P., Schüssler, R., 2020. Forecasting the equity premium: Mind the news!. *Rev. Finance* 24, 1313–1355.
- Amisano, G., Giacomini, R., 2007. Comparing density forecasts via weighted likelihood ratio tests. *J. Bus. Econom. Statist.* 25, 177–190.
- Andersen, T., Fusari, N., Todorov, V., 2020. The pricing of tail risk and the equity premium: Evidence from international option markets. *J. Bus. Econom. Statist.* 38, 662–678.
- Baetje, F., Menkhoff, L., 2016. Equity premium prediction: Are economic and technical indicators unstable? *Int. J. Forecast.* 32 (4), 1193–1207.
- Baker, M., Wurgler, J., 2006. Investor sentiment and the cross-section of stock returns. *J. Finance* 61, 1645–1680.
- Bali, T., Hovakimian, A., 2009. Volatility spreads and expected stock returns. *Manage. Sci.* 55, 1797–1812.
- Bali, T., Murray, S., 2013. Does risk-neutral skewness predict the cross-section of equity option portfolio returns? *J. Financ. Quant. Anal.* 48, 1145–1171.
- Bekaert, G., Engstrom, E., Xu, N., 2021. The time variation in risk appetite and uncertainty. *Manage. Sci.* 68, 3975–4753.
- Bollerslev, T., Marrone, L., Xu, L., Zhou, H., 2014. Stock return predictability and variance risk premia: Statistical inference and international evidence. *J. Financ. Quant. Anal.* 49, 633–661.
- Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. *Rev. Financ. Stud.* 22, 4463–4492.
- Bondarenko, O., 2014. Why are put options so expensive? *Q. J. Financ.* 4, 1–50.
- Buss, A., Schoenleber, L., Vilkov, G., 2017. Option-Implied Correlations, Factor Models, and Market Risk, INSEAD Working Paper.
- Buss, A., Schoenleber, L., Vilkov, G., 2019. Expected Correlation and Future Market Returns, CEPR Working Paper.
- Campbell, Y., Shiller, R., 1988. The dividend-price ratio and expectations of future dividends and discount factors. *Rev. Financ. Stud.* 1, 195–228.
- Campbell, J., Thompson, S., 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? *Rev. Financ. Stud.* 21, 1509–1540.
- Cao, C., Simin, T., Xiao, H., 2020. Predicting the equity premium with the implied volatility spread. *J. Financial Mark.* 51.
- Chang, Y., Christoffersen, P., Jacobs, K., 2013. Market skewness risk and the cross section of stock returns. *J. Financ. Econ.* 107, 46–68.
- Chen, J., Tang, G., Yao, J., Zhou, G., 2022. Investor attention and stock returns. *J. Financ. Quant. Anal.* 57, 455–484.
- Chernozhukov, C., Fernandez-Val, I., Galichon, A., 2010. Quantile and probability curves without crossing. *Econometrica* 78, 1093–1125.
- Chordia, T., Lin, T., 2021. Risk-neutral skewness, informed trading, and the cross section of stock returns. *J. Financ. Quant. Anal.* 56, 1713–1737.
- Christoffersen, P., Jacobs, K., Chang, B.Y., 2012. Forecasting with option-implied information. In: Elliott, G., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*, Volume 2.
- Clark, T.E., West, K.D., 2007. Approximately normal tests for equal predictive accuracy in nested models. *J. Econometrics* 138, 291–311.
- Cochrane, J., 2008. The dog that did not bark: A defence of return predictability. *Rev. Financ. Stud.* 21, 1533–1575.
- Conrad, J., Dittmar, R., Ghysels, E., 2013. Ex ante skewness and expected stock returns. *J. Finance* 68, 85–124.
- Dong, X., Li, Y., Rapach, D., Zhou, G., 2022. Anomalies and the expected market return. *J. Finance* 77, 639–681.
- Doran, J., Peterson, D., Tarrant, B., 2007. Is there information in the volatility skew? *J. Futures Mark.* 27, 921–960.
- Fama, E., French, K., 1988. Dividend yields and expected stock returns. *J. Financ. Econ.* 22, 3–25.
- Fama, E., French, K., 1989. Business conditions and expected returns on stocks and bonds. *J. Financ. Econ.* 25, 23–49.
- Gastwirth, J., 1966. On robust procedures. *J. Amer. Statist. Assoc.* 61, 929–948.
- Giacomini, R., White, H., 2006. Tests of conditional predictive ability. *Econometrica* 74, 1545–1578.
- Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* 102, 359–378.
- Gneiting, T., Ranjan, R., 2011. Comparing density forecasts using threshold and quantile weighted scoring rules. *J. Bus. Econom. Statist.* 29, 411–422.
- Goyal, A., Welch, I., 2008. A Comprehensive Look at The Empirical Performance of Equity Premium Prediction. *Rev. Financ. Stud.* 21 (4), 1455–1508.
- Guidolin, M., Hyde, S., Mcmillan, D., Ono, S., 2009. Non-linear predictability in stock and bond returns: When and where is it exploitable? *Int. J. Forecast.* 25, 373–399.
- Han, B., 2008. Investor sentiment and option prices. *Rev. Financ. Stud.* 21, 387–414.
- Henkel, S., Martin, J., Nadari, F., 2011. Time-varying short-horizon predictability. *J. Financ. Econ.* 99, 560–580.
- Hoerl, A., Kennard, W., 1970. Ridge regression: Applications to nonorthogonal problems. *Technometrics* 12, 69–82.
- Huang, D., Jiang, F., Li, K., Tong, G., Zhou, G., 2021. Scaled PCA: A new approach to dimension reduction. *Manage. Sci.* 68, 1678–1695.
- Huang, D., Jiang, F., Tu, J., Zhou, G., 2015. Investor sentiment aligned: A powerful predictor of stock returns. *Rev. Financ. Stud.* 28, 791–837.
- Israelov, R., Nielsen, L., 2014. Covered call strategies: One fact and eight myths. *Financ. Anal. J.* 70, 23–31.
- Israelov, R., Nielsen, L., 2015. Covered calls explained. *Financ. Anal. J.* 71, 44–57.
- Jarque, C., Bera, A., 1987. A test for normality of observations and regression residuals. *Internat. Statist. Rev.* 55, 163–172.

- Jiang, F., Lee, J., Xiumin, M., Zhou, G., 2019. Manager sentiment and stock returns. *J. Financ. Econ.* 132, 126–149.
- Kelly, B., Lustig, H., Van Nieuwerburgh, S., 2016. Too-systemic-to-fail: What option markets imply about sector-wide government guarantees. *Amer. Econ. Rev.* 106, 1278–1319.
- Kelly, B., Pruitt, S., 2013. Market expectations in the cross-section of present values. *J. Finance* 68, 1721–1756.
- Kelly, B., Pruitt, S., 2015. The three-pass regression filter: A new approach to forecasting using many predictors. *J. Econometrics* 186, 294–316.
- Kim, T., Park, H., 2018. Is stock return predictability of option-implied skewness affected by the market state? *J. Futures Mark.* 38, 1024–1042.
- Kothari, S., Shanken, J., 1997. Book-to-market, dividend yield, and expected market returns: A time-series analysis. *J. Financ. Econ.* 44, 169–203.
- Lettau, M., Ludvigson, S., 2001. Consumption, aggregate wealth, and expected stock returns: What do the data tell us? *J. Finance* 56, 815–849.
- Lettau, M., Van Nieuwerburgh, S., 2008. Reconciling the return predictability evidence. *Rev. Financ. Stud.* 21, 1607–1652.
- Li, J., Tsiakas, I., 2017. Equity premium prediction: The role of economic and statistical constraints. *J. Financial Mark.* 36, 56–75.
- Lima, L.R., Meng, F., 2017. Out-of-sample return predictability: A quantile combination approach. *J. Appl. Econometrics* 32 (4), 877–895.
- Ljung, G., Box, G., 1978. On a measure of a lack of fit in time series models. *Biometrika* 65, 297–303.
- Maio, P., 2016. Cross-sectional return dispersion and the equity premium. *J. Financial Mark.* 29, 87–109.
- Manzan, S., 2015. Forecasting the distribution of economic variables in a data-rich environment. *J. Bus. Econom. Statist.* 33, 144–164.
- Meligkotsidou, L., Panopoulou, E., Vrontos, I.D., Vrontos, S.D., 2014. A quantile regression approach to equity premium prediction. *J. Forecast.* 33 (7), 558–576.
- Meligkotsidou, L., Panopoulou, E., Vrontos, I.D., Vrontos, S.D., 2019. Quantile forecast combinations in realised volatility prediction. *J. Oper. Res. Soc.* 70 (10), 1720–1733.
- Meligkotsidou, L., Panopoulou, E., Vrontos, I.D., Vrontos, S.D., 2021. Out-of-sample equity premium prediction: A complete subset quantile regression approach. *Eur. J. Finance* 27 (1–2), 110–135.
- Neely, C.J., Rapach, D.E., Tu, J., Zhou, G., 2014. Forecasting the equity risk premium: The role of technical indicators. *Manage. Sci.* 60 (7), 1772–1791.
- Pedersen, T.Q., 2015. Predictable return distributions. *J. Forecast.* 34, 114–132.
- Pettenuzzo, D., Timmermann, A., Valkanov, R., 2014. Forecasting stock returns under economic constraints. *J. Financ. Econ.* 114 (3), 517–553.
- Pontiff, J., Schall, L., 1998. Book-to-market ratios as predictors of market returns. *J. Financ. Econ.* 49, 141–160.
- Rapach, D., Strauss, J., Zhou, G., 2010. Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Rev. Financ. Stud.* 23 (2), 821–862.
- Rapach, D., Zhou, G., 2021. **Asset pricing: Time-series predictability**, Working Paper.
- Shackleton, M., Taylor, S., Yu, P., 2010. A multi-horizon comparison of density forecasts for the S&P 500 using index returns and option prices. *J. Bank. Financ.* 34, 2678–2693.
- Stöckl, S., Kaiser, L., 2020. Higher moments matter! cross-sectional (higher) moments and the predictability of stock returns. *Rev. Financ. Econ.* 39, 455–481.
- Tibshirani, R., 1996. Regression shrinkage via the LASSO. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 23, 267–288.
- Timmermann, A., 2008. Elusive return predictability. *Int. J. Forecast.* 24, 1–18.
- Tsiakas, I., Li, J., Zhang, H., 2020. Equity premium prediction and the state of the economy. *J. Empir. Financ.* 58, 75–95.
- Vanden, J., 2008. Information quality and options. *Rev. Financ. Stud.* 21, 2635–2676.
- Wang, Y., Pan, Z., Liu, L., Wu, C., 2019. Oil price increases and the predictability of equity premium. *J. Bank. Financ.* 102, 43–58.
- Whaley, R., 2002. Return and risk of CBOE buy write monthly index. *J. Derivatives* 10, 35–42.
- Xing, Y., Zhang, X., Zhao, R., 2010. What does the individual option volatility smirk tell us about future equity returns? *J. Financ. Quant. Anal.* 45, 641–662.
- Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 67, 301–320.