

Stock Market Return Predictability Dormant in Option Panels*

Yoosoon Chang[†] Youngmin Choi[‡] Soohun Kim[§] Joon Park[¶]

July 2021

Abstract

This paper offers a novel approach to identify the relationship between extensive option panels and market returns using functional predictive regression. Employing our approach on the options and realized returns of the S&P 500, we achieve a remarkable performance in predicting S&P 500 monthly returns, yielding a 4.720% (6.198%) in-sample (out-of-sample) R^2 . The performance of our approach is superior to that of other well-known predictors and equilibrium models. The out-of-sample performance delivers substantial utility gains over historical averages. We find that both the use of option panels and the adoption of functional regression are indispensable for the outperformance.

JEL classification codes: G12, G17

Keywords: functional predictive regression, return predictability, risk-neutral measure, option market, market risk premium

*We are grateful to Liuren Wu, Linda Allen, Lin Peng, Xi Dong, Dexin Zhou, Jaeram Lee, Hogyu Jhang, and the conference and seminar participants at Baruch College, the Korea Advanced Institute of Science and Technology (KAIST), the 2020 Asia-Pacific Association of Derivatives (APAD) conference, the 2020 Financial Management Association (FMA) conference, the 2021 Korea-American Economic Association (KAEA) Macro-finance Virtual seminar, and the 2021 joint conference with the Allied Korea Finance Association (AKFA). Any remaining errors are solely our own. The usual disclaimer applies. Youngmin Choi gratefully acknowledges the support of a grant from the City University of New York PSC-CUNY Research Award Program.

[†]Department of Economics, Indiana University, Bloomington, IN. E-mail: yoosoon@iu.edu.

[‡]Department of Economics and Finance, Zicklin School of Business at Baruch College, the City University of New York, New York, NY. E-mail: youngmin.choi@baruch.cuny.edu.

[§]College of Business, Korea Advanced Institute of Science and Technology, Seoul, South Korea. E-mail: soohun@kaist.ac.kr.

[¶]Department of Economics, Indiana University, Bloomington, IN. E-mail: joon@iu.edu.

1 Introduction

Humans have been fascinated by the prospect of predicting future events since time immemorial. According to a tale in Aristotle’s ‘Politics,’ Thales monetized his forecast of the olive price through forward contracts on the exclusive use of olive presses. Since the adoption of modern financial markets, where firm ownership is publicly traded, the topic of forecasting stock returns has enticed many traders as well as academics. From the viewpoint of practitioners, it is necessary to exploit real-time forecasts of stock returns to ensure successful investment performance. Hence, it is natural for finance practitioners to eagerly employ various variables and adopt novel methodologies for the purpose of forecasting stock returns. From the perspective of academics, by analyzing the nature of stock return forecastability, we can deepen our understanding of market participants’ assessment of risks and their aversion toward those risks.

Tackling the holy grail of predicting stock market returns, this paper exploits one of the most universal phenomena over rational asset pricing models; the return distribution of an asset in a real world is supposed to be related to that in a risk-neutral world. Such a relation is commonly referred to a stochastic discount factor or pricing kernel, and various models in the literature have explained empirically observed behavior of prices by taking specific stances in the relation. In contrast, we let the data speak. We observe the risk-neutral density from the cross-section of options prices as of the end of each month and examine the empirically observed relation between the two densities over many periods.

Fundamentally, our approach is based on two theoretical claims in finance: (i) the risk-neutral density and the physical density are equivalent (Harrison and Kreps (1979)), and an asset pricing model can be interpreted as the change in measure between the two measures (Hansen and Richard (1987)); and (ii) cross-sectional option prices contain information on the risk-neutral density of an underlying asset (Banz and Miller (1978) and Breeden and Litzenberger (1978)). We study the predictability of the physical density using the risk-

neutral density resorting to the well-established two claims above in the literature, although the approach proposed in this paper can be applied to other economics in general.

Based on (i), we aim to identify the relation between the physical and risk-neutral densities. To this end, we do not take a stance on a particular asset pricing model but utilize the functional regression method.¹ In particular, we construct the physical density of S&P 500 index monthly returns by bootstrapping daily returns of the index in a given month. Using the functional regression method, we regress the physical density on the risk-neutral density observed in the previous month. From this, we identify the relation between the two densities and predict the physical density over the following month through the observed risk-neutral density. To exploit (ii)², we adopt the method of Ait-Sahalia and Duarte (2003), which imposes no-arbitrage restrictions nonparametrically and estimates the risk-neutral density of the S&P 500 index with a cross-section of options on the index that are expiring in a month.

Our approach using the functional predictive regression with the risk-neutral density shows a statistically and economically significant performance in predicting the market risk premium. In particular, we focus on examining the predictability of the first moment of the market excess return given the broad interest in prediction of the market risk premium, although we can use this methodology to forecast the entire physical density, not just a specific moment.³ Using the estimated functional relation and the risk-neutral density in month $t - 1$, we predict the mean of a physical density in month t with in-sample R^2 statistics of 4.720%,⁴ eclipsing the performance of well-known predictors, including the dividend yield, earnings-price ratio, and other variables examined in Welch and Goyal (2008).

¹For a detailed discussion of the method, see Bosq (2000) and Park and Qian (2012).

²See Bliss and Panigirtzoglou (2002) and Jackwerth (2004) for comprehensive reviews on this topic.

³We extend the usage of the functional predictive regression by further forecasting the second moment of a physical density of the S&P 500 when discussing the economic significance of the main findings in Section 4.3.

⁴The number of functional principal components is a choice variable in the functional regression framework, while the predictive ability of our approach remains stable with respect to a different choice of the number of factors. In particular, when we use three, four, and five factors, the in-sample R^2 statistics are 4.375%, 4.487%, and 4.720%, respectively.

More strikingly, our approach delivers an even stronger performance for out-of-sample prediction than in-sample prediction. The robustness of predictability for both in- and out-of-sample is particularly important given that many predictive regressions have often performed poorly in out-of-sample forecasts (see Welch and Goyal (2008)). Using the out-of-sample forecast assessment of Campbell and Thompson (2008), our prediction model achieves a 6.198% out-of-sample R^2 .⁵ Campbell and Thompson (2008) show that imposing weak restrictions on predictive regression can improve the forecasting performance of economic variables. However, regardless of the restrictions on the regression, the methodology proposed in this paper exhibits a superior out-of-sample performance relative to that of any of the economic forecasting variables discussed in Welch and Goyal (2008).

We further show that the statistical significance of our forecasting performance can translate into substantial economic gains for an investor. Although the R^2 is a useful statistical metric, it is not a proper tool to measure economic gain because it does not explicitly account for the risk borne by an investor over the out-of-sample period. Hence, as a metric of economic gain, we propose the realized utility gain of a mean-variance optimizing investor on an out-of-sample basis.⁶ When an investor fully exploits the advantages of our functional predictive regression approach that uses option panel data to forecast the mean as well as variance in the stock market return distribution, she obtains a 107.050% (68.769%) cumulative return over the five-year out-of-sample forecasting period assuming a risk aversion coefficient of 3 (10).⁷ This performance of our approach can be interpreted as a 5.577% (5.272%) annual risk-free return relative to the performance under the historical average model for a risk aversion coefficient of 3 (10).

Furthermore, we examine whether existing well-known equilibrium models explain the

⁵The out-of-sample forecasting performance of our approach remains intact when we change the number of functional factors used in the estimation. In particular, when we use three, four, and five factors, the out-of-sample R^2 statistics are 6.012%, 5.749%, and 6.198%, respectively.

⁶A similar approach is used in Marquering and Verbeek (2004), Welch and Goyal (2008), Campbell and Thompson (2008), Wachter and Warusawitharana (2009), and Rapach et al. (2010).

⁷See Figure 5 for more details. Over the same period, a mean-variance optimizing investor using the historical average of market returns and the rolling-window estimate for the variance can obtain 53.97% (12.51%) when assuming a risk aversion coefficient of 3 (10).

observed predictability using option panels. The approach proposed in this paper is fully nonparametric. Thus, imposing an asset pricing model’s restrictions on our approach should improve the forecasting performance if the model of interest were to explain the true data generating process. However, we find that the forecasting performance worsens when an equilibrium model’s restrictions are imposed on the functional regression framework. In particular, we investigate return predictability with well-known equilibrium models: (i) a model based on the recovery theorem of Ross (2015), (ii) a model with constant relative risk aversion (CRRA), (iii) the long-run risk model of Bansal and Yaron (2004) and (iv) the external habit model of Campbell and Cochrane (1999). Overall, these findings show that existing predictors or models cannot match the achievements of our novel approach.

Additional analysis reveals that our superior performance is not attributable solely to either the methodology or the use of option data. Applying functional predictive regression to the physical density itself to forecast the next month’s return does not deliver meaningful forecasting power. The finite moments (the mean, variance, skewness, and kurtosis) extracted from the risk-neutral density also fail to provide any predictive ability in forecasting the next month’s return. Furthermore, various option-related measures, such as the Chicago Board Options Exchange (CBOE) Volatility Index (VIX) and the variance risk premium, do not deliver a forecasting ability comparable to that of our proposed approach. The combination of the advanced econometric methodology and the potent data source is the key driver of the main findings of the present paper.

This paper lies at the intersection of two strands of literature: return predictability and options. Academic work predicting stock returns goes back to Cowles (1933) and Cowles and Jones (1937). In the early literature, return predictability was interpreted as being contradictory to market efficiency (Fama (1965), Fama (1970) and Samuelson (1965)). However, Fama (1991) harmonizes the empirical findings on return predictability with market efficiency. Over the past decades, researchers have proposed various models featuring return predictability: external habits (Campbell and Cochrane (1999)), dynamic risk-sharing op-

portunities among heterogeneous agents (Lustig and Nieuwerburgh (2005)), long-run risks (Bansal and Yaron (2004), Bansal et al. (2010)), and time-varying disaster risks (Gabaix (2012)). Furthermore, advanced methodologies have been widely used: structural vector autoregression models (VARs) (Cochrane (2008), van Binsbergen and Koijen (2010)), model combinations (Rapach et al. (2010) and Dangel and Halling (2012)), and structural breaks (Guidolin and Timmermann (2007), Henkel et al. (2011)). We contribute to the return predictability literature by revealing the inherent kinship between risk-neutral density and option prices.⁸ Compared to the existing accomplishments in the return predictability literature, we achieve phenomenal and robust predictive ability through the vast amount of information embedded in option price data and provide a novel statistical tool.

This paper is not the first attempt to relate the information in the option market to the physical dynamics of an underlying asset. Christoffersen and Mazzotta (2005) investigate the relationship between currency options and currency dynamics. Christoffersen et al. (2013) propose a pricing kernel that unifies stock returns and anomalies in option prices. Roll et al. (2010), Johnson and So (2012), and Hu (2014) suggest that option trading activities contain information on future stock returns. Furthermore, given the documented return predictability based on the information of higher moments in option panels,⁹ this paper is related to the literature on the pricing of higher moments studied in Rehman and Vilkov (2012), Chang et al. (2013), and Amaya et al. (2015). These studies focus on examining a cross-section of returns or individual stock returns through risk-neutral density. In contrast, we highlight the delivery of meaningful out-of-sample predictive ability through the rich information embedded in full risk-neutral density.

⁸See, for example, Banz and Miller (1978) and Breeden and Litzenberger (1978). Researchers have proposed various methods of capturing this relation (Jackwerth and Rubinstein (1996), Ait-Sahalia and Lo (1998), Ait-Sahalia and Duarte (2003)). Furthermore, Rosenberg and Engle (2002) and Jackwerth (2000, 2004) show how we can learn about investor risk aversion by jointly observing option markets and the return dynamics of the underlying market. As recent endeavors in this line of research, Ross (2015) and Carr and Yu (2012) propose how to recover both the risk-neutral and physical densities from only option panels under certain restrictions.

⁹See Bollerslev et al. (2009), Bollerslev and Todorov (2011), Bollerslev and Todorov (2014), and Bollerslev et al. (2015). Recently, Feunou et al. (2018) and Kilic and Shaliastovich (2019) investigate the relationship between equity premium and different dimensions of variance risk premium.

In summary, the key takeaways of this paper are twofold. First, we highlight the information dormant in option panels that is valuable for predicting market returns. Recall that the transition between risk-neutral density and physical density is pinned down by the risk preference of pricing agents (Hansen and Jagannathan (1991)). Then, a rather mild assumption of persistent preference naturally implies a stable relation between the risk-neutral and physical densities. To the best of our knowledge, this is the first paper aiming to predict the market risk premium by exploiting this relation. Second, we introduce a novel prediction method that handles predictors in a high-dimensional space, such as the risk-neutral density function. Given that the size of relevant datasets such as social network site (SNS) posts, household-level consumption data, and demographic indicators continues to grow, the ability to extract relevant information from big data is crucial. The methodology that we use to connect the two densities can be easily applied to connect any two high-dimensional objects.

This paper is organized as follows. In Section 2, we explain the data that we use for our empirical analysis. Section 3 describes our prediction methodology. The empirical results are reported in Section 4. Section 5 concludes.

2 Data

To extract the risk-neutral density of the aggregate stock market, we use data on S&P 500 index options from January 1996 to December 2015 from the Option Metrics Database. In particular, we collect data on implied volatility, strike prices, expiration dates, dividend yields, the price of the underlying asset (S&P 500 index), and the risk-free rate. For each option contract, we filter out zero trading volume, zero open interest, or zero or missing implied volatility data. We also eliminate options with average bid and ask quotes of less than \$3/8. We use only put option data. The object of prediction is the mean of the physical density of the aggregate stock market. However, one of the empirical challenges from this perspective is that in reality, we observe only one data realization of monthly returns. To overcome this issue, we first collect daily returns of the S&P 500 index from the Center

for Research in Security Prices (CRSP). Then, we construct the density of monthly returns using a bootstrap method, which we explain in detail in Section 3.2.

We describe how we select the observation date and the corresponding time to maturity because we aim to predict the next month's stock market return using the information embedded in option data without hindsight. From the option data filtered as described above, we use the options whose days until expiration are 30 days so that the horizon of return prediction analysis is monthly. Figure 2 shows the dates for which we collect option data with a certain time to maturity and the way in which we aggregate market return data for our analysis. For example, when the expiration date of the index options is January 17, 2016, we collect the option data observed on December 18, 2015, whose time to expiration date is 30 days.¹⁰ In this example, the *observation date* of the option data is December 18, 2015, and the *expiration date* of the options is January 17, 2016. Similar adjustments are made to the dividend yield and risk-free rate data. That is, we interpolate the dividend yields and risk-free rate data provided by Option Metrics to obtain a 30-day dividend yield and risk-free rate on the *observation date*.

When the *observation* and *expiration dates* of options do not correspond to the first or last date of a month, we make an analogous adjustment in constructing the market return data. Recall that our main objective is to investigate the predictive ability of risk-neutral density obtained from option data on the physical density of stock market returns. Hence, we examine whether the risk-neutral density extracted from option data on the *observation date* has predictive ability in explaining the stock market return realized over the *observation date* to the *expiration date* of the options used for the prediction. To this end, we collect daily returns of the S&P 500 index from the *observation* to *expiration dates* and use a bootstrap method to construct the physical density of monthly returns. Sections 3.1 and 3.2 provide detailed descriptions of the construction of the risk-neutral and physical densities, respectively.

¹⁰If there are no options whose time to expiration is exactly 30 days, we collect options whose time to expiration is closest to 30 days.

Table 1 provides descriptive statistics for the option data used in our analysis. The second column of Table 1 provides annual and overall averages of the S&P 500 index, and the next column shows the number of put options used in our analysis, exhibiting a dramatic increase in recent years. The next four and last four columns display information on strike prices and implied volatility of put options, respectively. Similarly, the range of strike prices and implied volatility exhibits wider dispersion in the second half of the sample period than in the first half.

3 Methodology

3.1 Risk-Neutral Density Construction

This subsection describes our approach to extracting the risk-neutral density (Q -density) from the panel data of option prices. The value of an option contract is the expected payoff on the expiration date discounted to the present. Under risk neutrality, the value of a call option at time t can be written as

$$C_t = \int_K^\infty e^{-r_f(T-t)}(S_T - K)q(S_T)dS_T, \quad (1)$$

where K is the strike price, r_f is the risk-free rate, T is the date of expiration, S_T is the price of the underlying asset, and $q(\cdot)$ is the risk-neutral density. Breeden and Litzenberger (1978) and Banz and Miller (1978) show that from Equation (1), the risk-neutral density can be obtained by taking a second-order derivative with respect to the strike price. That is,

$$q(S_T) = e^{r_f(T-t)} \frac{\partial^2 C_t}{\partial K^2}. \quad (2)$$

Practical application of the above approach to extracting the risk-neutral density has several empirical challenges. First, we observe only a limited number of option contracts with discrete prices. Second, an option contract is traded based on bid and ask prices with

microstructure noise. Third, there is a limited range of available strike prices. These issues make data on option prices coarse and noisy. Furthermore, the problems worsen when we take the second-order derivative, which is our main object of interest.

In this paper, we obtain a risk-neutral density by imposing no-arbitrage constraints: monotonicity and convexity of call option prices following Ait-Sahalia and Duarte (2003). In particular, from the positivity of the density and its integrability to one, two constraints can be written as follows:

$$-e^{-r(T-t)} \leq \frac{\partial C_t}{\partial K} \leq 0 \quad \text{and} \quad \frac{\partial^2 C_t}{\partial K^2} \geq 0.$$

In particular, following Ait-Sahalia and Duarte (2003), a risk-neutral density is obtained using constrained least squares regression with a local polynomial kernel smoothing approach. Using the option panel data described in Section 2, we estimate the risk-neutral density of the S&P 500 index from a cross-section of option contracts on the S&P 500 that expire in a month. The top two plots in Figure 1 display our estimated Q -density (left) and demeaned Q -density (right). Since we apply the no-arbitrage constraints in extracting the risk-neutral density from the option panel data, the estimated Q -density function is well behaved.

3.2 Physical Density Construction

In this subsection, we describe the bootstrap method we use to construct the physical density of S&P 500 index monthly returns. Recall that we aim to predict the physical density of the S&P 500 index return in a given month using the risk-neutral density of option prices observed at the end of the previous month. In particular, we collect daily returns of the S&P 500 index over the period between the *observation date* and the *expiration date* described in Figure 2, which corresponds to the lifetime of the option contracts used to construct the risk-neutral density.

Suppose that, as of observation date t , there are N days until expiration date $t + 1$.

More precisely, a cross-section of prices of options on the S&P 500 index expiring in month $t + 1$ is observed in month t , and there are N number of days between the observation and expiration dates. In this setup, the realized monthly return of a given month t is written as $r_t = \prod_{i=1}^N (1 + r_{i,t}) - 1$, where $r_{i,t}$ is the daily return on day i of month t . To obtain the monthly physical density function of the S&P 500 index required for our study, we collect the daily returns in a given month, $\{r_{1,t}, r_{2,t}, \dots, r_{i,t}, \dots, r_{N,t}\}$, and construct a series of bootstrapped samples, $\{r_{1,t}^j, r_{2,t}^j, \dots, r_{i,t}^j, \dots, r_{N,t}^j\}$, where $j = 1, \dots, B$ with B denoting the number of bootstrapped samples. Then, we simulate the monthly returns as $r_t^j = \prod_{i=1}^N (1 + r_{i,t}^j) - 1$ for $j = 1, \dots, B$, using the bootstrapped samples. In our empirical analysis, we set $B = 10,000$ and 10,000 simulated monthly returns are generated for a given month, which are used to construct the physical density function of the S&P 500 index. The bottom two plots in Figure 1 show our physical density of the S&P 500 index (left) estimated using the bootstrapping method and the demeaned density (right).

3.3 Projection of the Physical Density onto the Risk-Neutral Density

Finally, we explain how we can relate the two densities described in the previous two subsections. Recall that \mathbf{X}_R is the set of 1024 potential one-month-horizon S&P 500 returns used for the P -density estimation. Let $\mathbf{q}_t : \mathbf{X}_R \rightarrow \mathbb{R}$ denote the Q -density function mapping the one-month-horizon S&P 500 returns from the end of month t to the end of month $t + 1$ to a real number, which is constructed using one-month-horizon option prices observed at time t , as described in Section 3.1. Let $\mathbf{p}_{t+1} : \mathbf{X}_R \rightarrow \mathbb{R}$ denote the P -density function mapping the one-month-horizon S&P 500 returns from month t to $t + 1$ to a real number, which is constructed using realized daily returns from the end of month t to the end of month $t + 1$, as described in Section 3.2. Instead of including a constant functional term in our functional regression model, we directly work with the demeaned versions of the density

functions, $\mathbf{d}_{\mathbf{p},t+1}$ and $\mathbf{d}_{\mathbf{q},t}$ corresponding to \mathbf{p}_{t+1} and \mathbf{q}_t , respectively.¹¹ For the in-sample (out-of-sample) estimation, we use the densities over the whole (past) time-series to compute the mean densities.

We consider the functional regression as follows:

$$\mathbf{d}_{\mathbf{p},t+1} = \mathbf{A}\mathbf{d}_{\mathbf{q},t} + \boldsymbol{\varepsilon}_{t+1}, \quad (3)$$

where \mathbf{A} is a mapping from a space of real-valued functions to itself.¹² We assume all technical conditions required in Bosq (2000) and Park and Qian (2012). Because our main interest lies in out-of-sample prediction, we describe the estimation procedures for performing out-of-sample prediction.

Step 1. We wavelet-transform the demeaned density functions of $\mathbf{d}_{\mathbf{p},s+1}$, $\mathbf{d}_{\mathbf{q},s}$ into the sequences of coefficients of $\mathbf{w}_{\mathbf{p},s+1}$, $\mathbf{w}_{\mathbf{q},s}$ for $s \leq t$.

Step 2. We find the first K functional principal components from the matrix $\mathbf{W}_{\mathbf{q}}\mathbf{W}_{\mathbf{q}}'$, where $\mathbf{W}_{\mathbf{q}} = [\mathbf{w}_{\mathbf{q},1} \ \mathbf{w}_{\mathbf{q},2} \ \cdots \ \mathbf{w}_{\mathbf{q},t-1}]$. Let λ_k and \mathbf{e}_k denote the k -th eigenvalue and eigenvector, respectively, for $k \leq K$.

Step 3. Using the regularized regressors from Step 2, we find

$$\widehat{\mathbf{A}}_w = \left(\sum_{s \leq t} \mathbf{w}_{\mathbf{p},s+1} \mathbf{w}_{\mathbf{q},s}' \right) \left(\sum_{k \leq K} \lambda_k^{-1} \mathbf{e}_k \mathbf{e}_k' \right).$$

Step 4. From the estimated mapping in Step 3, we make a prediction $\widehat{\mathbf{w}}_{\mathbf{p},t+1} = \widehat{\mathbf{A}}_w \mathbf{w}_{\mathbf{q},t}$ in the wavelet space

Step 5. We transform the estimated coefficient sequence $\widehat{\mathbf{w}}_{\mathbf{p},t+1}$ in the wavelet space to $\widehat{\mathbf{d}}_{\mathbf{p},t+1}$ in the function space by reversing the procedures in Step 1.

¹¹Note that the demeaned density functions $\mathbf{d}_{\mathbf{p},t}$ and $\mathbf{d}_{\mathbf{q},t}$ are *not* densities anymore. In particular, they are not nonnegative, and they are integrated not up to one but up to zero.

¹²The functional regression (3) is equivalent to $\mathbf{p}_{t+1} = \mathbf{b} + \mathbf{A}\mathbf{q}_t + \boldsymbol{\varepsilon}_{t+1}$, where \mathbf{b} is the functional intercept term for the purpose of inference on \mathbf{A} by the least squares method.

Step 6. We add back the historical mean to $\hat{\mathbf{d}}_{\mathbf{p},t+1}$ to obtain $\hat{\mathbf{p}}_{t+1} = \hat{\mathbf{d}}_{\mathbf{p},t+1} + \frac{1}{t} \sum_{s=1}^t \mathbf{p}_s$.

The intuition behind the suggested procedures follows. Steps 1-3 stabilize the estimation outcomes. Recall that our target of \mathbf{A} is a high-dimensional object. Hence, a brute-force regression approach is infeasible or highly unstable. Instead, we propose an alternative route. We transform a real-value vector to a vector of elements corresponding to a wavelet basis in Step 1 and summarize the information in the regressor by the most important K components in Step 2. As a result of this regularization, the estimator in Step 3 does not suffer from the ill-posed inverse problem.¹³ Steps 4 and 5 simply reverse the pretreatments.¹⁴

A discussion on the tuning parameters follows. In Step 2, we choose the number of functional principal components, K . Figure 3 shows the scree plot of eigenvalues of the variance estimator. We find that the first few functional principal components are able to explain a significant portion of the variation in matrix $\mathbf{W}_{\mathbf{q}}\mathbf{W}_{\mathbf{q}}'$. In particular, the first three (five) principal components explain 87.1% (93.3%) of the total variation in the matrix. In the empirical application, we further show that the main results of this paper are robust to the choice of K within a reasonable range.

Furthermore, the result in Step 2, combined with a functional principal component analysis (FPCA), illuminates which features of the Q -density carry useful and timely information for predicting a subsequent period's P -density. Figure 4 provides the result of the FPCA on Q -density extracted from Section 3.1. The top, middle, and bottom left figures plot the first ($\nu_{T,1}$), second ($\nu_{T,2}$), and third ($\nu_{T,3}$) functional principal components extracted from the risk-neutral density over our sample period, respectively. The top, middle, and bottom right figures in each panel show how each of the three factors move the average risk-neutral density function. In each of the right figures, the blue line (*mean*) represents the time-series

¹³For the validity of this regularization, see Bosq (2000). It is required to set $K \rightarrow \infty$ at an appropriate rate as the sample size increases.

¹⁴Park and Qian (2012) demonstrate that the regression statistics, such as the R^2 , of the functional regression can be obtained and interpreted analogously to those of an ordinary regression. All the technical assumptions in their paper are applied in the present paper.

average of the risk-neutral density, $\bar{\mathbf{q}} = (1/T)\Sigma_{t=1}^T \mathbf{q}_t$. The green line ($\text{mean} + \min \times \nu_{T,k}$) and the red line ($\text{mean} + \max \times \nu_{T,k}$) represent $\bar{\mathbf{q}} + \langle \nu_{T,k}, \mathbf{q}_t \rangle_{\min} \nu_{T,k}$ and $\bar{\mathbf{q}} + \langle \nu_{T,k}, \mathbf{q}_t \rangle_{\max} \nu_{T,k}$ for $k = 1, 2$, and 3 , respectively.

The functional principal component analysis results depicted in Figure 4 suggest that each principal component places different weights on different dimensions in the return space. In particular, the first principal component slightly scales down medium-sized returns, while the third component put some emphasis on the tails. The second principal component highlights asymmetric weights on positive vs. negative returns. To show the variations in Q -density caused by its principal components more explicitly, we also present the time averaged Q -density added by each principal component with its historically maximum and minimum loadings. The first principal component with positive loadings concentrates the Q -density around 0. However, its impact with negative loadings is not so simple: it introduces bimodality and asymmetry as well as less concentration around 0. The second and third principal components yield impacts that appear to be more complicated and harder to summarize: they not only shift the means but also change the modes and modalities. Nevertheless, it is clear that the leading principal components introduce very complicated variations in shapes of the Q -density. Overall, our analysis here suggests that functional regression using Q -density allows us to capture profound and unique features in option data, unlike existing approaches that use a few finite moments of the Q -density in a simple predictive regression framework. In a later section, we provide evidence that using the finite moments of the Q -density does not provide a meaningful forecasting performance for predicting next-month stock returns.

4 Empirical Results

4.1 In-Sample Prediction

For the sample period from January 1996 to December 2015, we use the functional regres-

sion framework of Bosq (2000) and Park and Qian (2012) to predict the next-month returns of the S&P 500 index. While the functional regression framework in Equation (3) relates the two density functions with a mapping of \mathbf{A} , we focus on the first moment of the predicted P -density function due to the wide range of interest in return predictability.¹⁵ A predicted return \hat{r}_{t+1} is computed as the first moment of the predicted physical density function of the S&P 500 index $\hat{\mathbf{p}}_{t+1}$. For the in-sample analysis, we use the whole sample when estimating the functional regression following the procedure described in Section 3.3. Also, when adding back the mean of the density for the in-sample estimation, we add the mean computed from the whole sample to obtain $\hat{\mathbf{p}}_{t+1}$.¹⁶ Using the predicted returns estimated by the functional regression, we provide R^2 statistics in Table 2. The R^2 statistics of the functional predictive regression for different numbers of functional principal components used in the estimation are provided in Panel A. The in-sample estimation exhibits significant predictive ability of the risk-neutral distribution extracted from option prices on the S&P 500 index return, with the R^2 statistics ranging from 4.375% to 4.720%, with very strong statistical significance measured by the F -statistics (and p -values) of the predictive regressions.

The existing literature has documented that even in the in-sample prediction, most of the well-known predictors have poor predictive ability for the market risk premium (see, among many others, Welch and Goyal (2008), Campbell and Thompson (2008), and Rapach et al. (2010)). To compare the performance of our approach to that of existing predictors in the literature, we also compute the in-sample R^2 statistics of the well-known predictors used in Welch and Goyal (2008).¹⁷ In particular, we obtain the R^2 statistics from the following

¹⁵It follows from Equation (3) that

$$\langle \mathbf{m}, \mathbf{d}_{\mathbf{p},t+1} \rangle = \langle \mathbf{m}, \mathbf{A} \mathbf{d}_{\mathbf{q},t} \rangle + \langle \mathbf{m}, \boldsymbol{\varepsilon}_{t+1} \rangle = \langle \mathbf{A}^* \mathbf{m}, \mathbf{d}_{\mathbf{q},t} \rangle + \langle \mathbf{m}, \boldsymbol{\varepsilon}_{t+1} \rangle \quad (4)$$

where \mathbf{m} is the identity function, $\mathbf{m}(r) = r$, $\langle \cdot, \cdot \rangle$ is the functional inner product, $\langle \mathbf{a}, \mathbf{b} \rangle = \int \mathbf{a}(r) \mathbf{b}(r) dr$, and \mathbf{A}^* is the adjoint of \mathbf{A} . Note that $\langle \mathbf{m}, \mathbf{d}_{\mathbf{p},t+1} \rangle$ is the demeaned first moment of \mathbf{p}_{t+1} . Therefore, Equation (3) implies in particular that the first moment of the P -density function is predicted by the Q -density function for the S&P 500 index.

¹⁶This process is analogous to Step 6 of the procedure described in Section 3.3.

¹⁷The data are available from Amit Goyal's homepage: <http://www.hec.unil.ch/agoyal/>

predictive regression:

$$r_t = \alpha + \beta X_{t-1} + \varepsilon_t, \quad (5)$$

where r_t is the excess return on the S&P 500 index of period t , X_{t-1} is a set of predictors observed in period $t - 1$, and ε_t is an error term. Panel B of Table 2 provides the in-sample R^2 of regressions with the 14 forecasting variables used in Welch and Goyal (2008). The estimated R^2 statistics range from a low of 0.001% for the Treasury bill rate to a high of 2.084% for the stock variance. Only three forecasting variables (dividend yield, stock variance, and net equity expansion) exhibit marginal statistical significance measured by p -values for the corresponding F -statistics of predictive regressions, and the other 11 variables do not show any statistically significant forecasting ability for stock market returns.

When we include all 14 variables in the regression (a “Kitchen sink” model), the in-sample R^2 obviously increases to as much as 13.578%. However, the out-of-sample forecasting ability of the kitchen sink model deserves mention and is investigated in the next subsection. Overall, the results provided in Table 2 suggest that our approach using information embedded in option prices to predict market returns shows an impressive prediction performance in comparison with that of existing and well-known predictors in the financial market.

4.2 Out-of-Sample Prediction

While numerous economic variables have been proposed as predictors of stock returns, including valuation ratios and other variables as documented in Welch and Goyal (2008), the existence of out-of-sample predictability has still been controversial. In this subsection, we evaluate the out-of-sample performance of our approach to predicting stock market returns using functional predictive regression with risk-neutral density extracted from option data.

In out-of-sample prediction, we run the functional predictive regression (Equation(3)) using an expanding window with the forecasting period starting in January 2011, which covers the last five years of our main sample. For each estimation, we generate an out-of-sample forecast of the stock market return and compare the forecast with the realized stock

market return. That is, when estimating an out-of-sample prediction of \hat{r}_t , we use available options data up to time t to estimate $\hat{\mathbf{d}}_{\mathbf{p},t}$ in Step 6 of the procedure described in Section 3.3. Following Campbell and Thompson (2008) and Welch and Goyal (2008), we use the historical average of S&P 500 returns (\bar{r}_t) as a natural benchmark model.

To assess the out-of-sample forecasting power of our approach, we use the out-of-sample R^2 as suggested by Campbell and Thompson (2008) to compare the forecast from the functional regression (\hat{r}_t) and the forecast using the historical average of stock market returns (\bar{r}_t). The out-of-sample R^2 of Campbell and Thompson (2008) is computed as follows:

$$R_{OOS}^2 = 1 - \frac{\sum_{t=1}^T (r_t - \hat{r}_t)^2}{\sum_{t=1}^T (r_t - \bar{r}_t)^2}, \quad (6)$$

where \hat{r}_t is a fitted value from the functional regression estimated through period $t - 1$ in an out-of-sample manner and \bar{r}_t is the historical average of stock market returns through period $t - 1$. That is, the above R_{OOS}^2 measures the reduction in the mean squared forecast error (MSFE) for the functional predictive regression relative to that of the historical average forecast. To estimate the historical average of stock returns, we use the long historical data on the S&P 500 index returns starting in 1927, giving the historical average model an advantage in terms of data availability.

To provide a statistical test of whether the functional predictive regression provides a significantly better forecast over the historical average forecast, we compute the *MSFE-adjusted* statistics proposed by Clark and West (2007). Clark and West (2007)'s test in our context examines the null hypothesis that the MSFE of the historical average model is less than or equal to the MSFE of the functional predictive regression against the alternative hypothesis that the historical average MSFE is greater than the functional predictive regression MSFE.

Table 3 provides the out-of-sample R^2 statistics, R_{OOS}^2 , of our proposed approach (Panel A) and predictors used in Welch and Goyal (2008) (Panel B). In particular, Panel A reports out-of-sample R^2 statistics of the functional predictive regression obtained with different

numbers of functional principal components used in the estimation. Panel B provides the unrestricted predictive regression of each forecasting variable (columns under Unrestricted) and the restricted regression with theoretical restrictions proposed in Campbell and Thompson (2008) (columns under Campbell and Thompson (2008) restrictions).¹⁸ Along with the R^2_{OOS} , we also provide p -values (in brackets) for Clark and West (2007) *MSFE-adjusted* statistics for the functional predictive regression in Panel A and for the unrestricted and restricted predictive regressions of conventional variables in Panel B.

Our functional predictive regression model achieves a 6.198% (6.102%) out-of-sample R^2 when we use five (three) functional principal components in the functional predictive regression with strong statistical significance, which indicates that the functional predictive regression delivers a significantly lower MSFE than that of the historical average forecast. Consistent with the well documented findings in the existing literature, the 14 predictors exhibit very poor performance in predicting stock market returns out of sample, with even some negative out-of-sample R^2 statistics. Furthermore, due to an overfitting problem in the out-of-sample prediction, the kitchen sink regression using all 14 variables delivers a significantly negative out-of-sample R^2 of -9.190% without a regression restriction and -9.212% with the restrictions of Campbell and Thompson (2008). Overall, our approach using the functional regression in predicting stock market returns with option data provides an unprecedentedly high predictive ability, even in the out-of-sample analysis.

4.3 Economic Significance of Findings

In the previous section, we showed that the functional predictive regression relating the risk-neutral density from option prices to stock market return predictability displays significantly positive in-sample as well as out-of-sample predictive power. Compared to forecasting variables that have been discussed in the literature, our functional regression approach using option panel data exhibits superior performance. However, a limitation of

¹⁸Campbell and Thompson (2008) suggest to impose sign restrictions on $\hat{\beta}$ and \hat{r}_t in Equation 5. In particular, if $\hat{\beta}$ has an unexpected sign, then it is set $\hat{\beta} = 0$. Also, if $\hat{r}_t < 0$, the forecast is set to zero.

the R^2 statistic is that it does not explicitly take into account the risk borne by an investor over the out-of-sample period. Thus, following Marquering and Verbeek (2004), Welch and Goyal (2008), Campbell and Thompson (2008), Wachter and Warusawitharana (2009), and Rapach et al. (2010), we investigate the realized utility gain for a mean-variance optimizing investor on a real out-of-sample basis.

Consider a mean-variance investor with a relative risk aversion coefficient of γ who allocates her portfolio between equity and a risk-free asset based on a forecast. On the one hand, if the investor were to use the historical average to forecast the equity market return, then at the end of period t , she would allocate the following share of a portfolio to the stock market for the period of $t + 1$:

$$\alpha_{0,t} = \left(\frac{1}{\gamma} \right) \left(\frac{\bar{r}_{t+1}}{\hat{\sigma}_{t+1}^2} \right), \quad (7)$$

where \bar{r}_{t+1} is the historical average of stock market returns using data upto time t and $\hat{\sigma}_{t+1}^2$ is the forecast of the variance, which is provided in more detail below. Over the out-of-sample period, the investor realizes the utility level of

$$\hat{\nu}_0 = \hat{\mu}_0 - \left(\frac{1}{2} \right) \gamma \hat{\sigma}_0 \quad (8)$$

where $\hat{\mu}_0$ and $\hat{\sigma}_0$ are, respectively, the sample mean and variance over the out-of-sample period for the return in the benchmark portfolio formed using forecasts based on the historical average.

On the other hand, we can compute the average utility for the same investor when she utilizes our functional predictive regression approach with option panel data, instead of the historical average, to forecast the equity market return. Then, the share of her portfolio allocated to equity would be

$$\alpha_t = \left(\frac{1}{\gamma} \right) \left(\frac{\hat{r}_{t+1}}{\hat{\sigma}_{t+1}^2} \right), \quad (9)$$

where \hat{r}_{t+1} is the predicted mean of stock returns and $\hat{\sigma}_{t+1}^2$ is the predicted variance for the

period $t + 1$. The utility level of the investor would be

$$\hat{\nu} = \hat{\mu} - \left(\frac{1}{2}\right) \gamma \hat{\sigma}^2, \quad (10)$$

where $\hat{\mu}$ and $\hat{\sigma}^2$ are the sample mean and variance over the out-of-sample period for the return on the portfolio constructed using the functional predictive regression, respectively.

This exercise requires the investor to forecast the variance of stock returns, $\hat{\sigma}^2$. When the investor uses the historical average to forecast stock market returns, we assume that she estimates the variance using a five-year rolling window of monthly returns, similar to the procedure in Campbell and Thompson (2008) and Rapach et al. (2010). The functional predictive regression allows us to relate not only a function to a single variable but also a function to another function. That is, in our primary analysis, our approach uses a risk-neutral density to forecast the first moment of the physical density of the S&P 500. However, this approach can be extended to map a risk-neutral density to predict a full distribution of S&P 500 returns. Therefore, the investor has an additional advantage—she can further estimate higher moments from a predicted full distribution of S&P 500 returns.¹⁹ Therefore, the investor can either use the variance estimate based on a five-year rolling window of monthly returns or further exploit the advantage of the functional predictive regression to forecast the variance of stock returns by predicting the full distribution.

Following Rapach et al. (2010), we measure the utility gain as the difference between Equations (10) and (8) and multiply this difference by 1200 to express it as an annualized percentage return. This utility gain can be interpreted as the portfolio management fee that an investor would be willing to pay to choose to use our functional predictive regression approach exploiting the relation between the risk-neutral density and the physical density of the S&P 500, instead of using the historical average approach. We provide results for the two risk aversion parameter values, $\gamma = 3$ and $\gamma = 10$, in the next paragraphs, tables, and

¹⁹For the prediction of the k -th moment of S&P 500 returns, we may use Equation (4) in Footnote 15 with \mathbf{m} redefined as $\mathbf{m}(r) = r^k$.

figures. The results are robust and qualitatively similar to other reasonable choices of γ .

The predictive ability of the functional regression approach delivers economically meaningful outcomes to investors. Table 4 reports the utility gain (in terms of the certainty equivalent return, expressed in a percentage per annum) of an investor who utilizes the approach proposed in this paper instead of the historical average approach. In particular, Panel A of Table 4 provides the results when an investor uses the functional predictive regression approach to forecast the mean, i.e., the first moment, of the next month's S&P 500 index return distribution (under the column Mean forecast) and to forecast the mean and variance in the distribution (under the column Mean-Var forecast). Regardless of the choice of the number of functional principal components used in the estimation, the functional predictive regression of the risk-neutral density provides significant utility gains to the investors. When an investor fully utilizes our approach to forecast the mean and variance of the stock market return, her utility gains are well above 7% per annum, which is a very substantial profit. Even if an investor used the approach to forecast the mean of the stock market returns, she could obtain a sizable utility gain well above 2.5% per annum.

Figure 5 further highlights the economic gain of an investor. The two plots in Figure 5 show the portfolio performance of a mean-variance maximizing investor who uses the functional regression approach compared to using the historical average model. The top and bottom plots correspond to the cumulative returns of the trading strategies of a mean-variance optimizing investor with a risk aversion coefficient of 3 and 10, respectively. The red line represents the cumulative gain of investing one dollar using functional regression with option data to forecast \hat{r}_{t+1} and $\hat{\sigma}_{t+1}$, and the blue line shows the cumulative gain using the historical average model with the rolling-window estimate for $\hat{\sigma}_{t+1}$. In Panel A and Panel B, the investor obtains the cumulative returns of 107.050% with $\gamma = 3$ and 68.769% with $\gamma = 10$ over the out-of-sample forecasting period from 2011 to 2016, respectively.²⁰ However, the historical average model delivers 53.97% and 12.51% cumulative return over the same

²⁰In this analysis, we fix $K = 5$, but the results are qualitatively stable with $K = 3$ and 4. These results are available upon request.

period with $\gamma = 3$ and $\gamma = 10$, respectively.

We also provide evidence that the economic gains from our functional predictive regression using risk-neutral density are consistent across different choices of regularization and superior to those gains obtained from existing approaches using forecasting variables with economic restrictions. Panel B of Table 4 provides the utility gains of an investor when she forecasts using the economic variables discussed in Welch and Goyal (2008) with the unrestricted regression (under the column Unrestricted) and with the economically restricted regression (under the column Campbell and Thompson (2008) restriction). Consistent with the results on the in-sample and out-of-sample R^2 statistics discussed in the previous section, the utility gains from forecasting using the economic variables are mostly not sizable and sometimes even negative. The dividend-price ratio and dividend yield are the only two variables that deliver marginally meaningful utility gains, while many of the individual predictors with negative out-of-sample R^2 values (for example, inflation, the Treasury bill rate, and net equity expansion) fail by a substantial margin to outperform the historical average benchmark.

4.4 Where Does This Superior Performance Come From?

In this subsection, we investigate the source of the outstanding performance of the forecasting approach proposed in this paper. The critical elements of our approach are twofold: the econometric method using the functional regression and the rich information embedded in the extensive option panel data. Thus, this section examines how much the functional regression framework and the use of the full distribution of the risk-neutral density contribute to the forecasting performance.

4.4.1 Functional Predictive Regression Using Physical Density

One of the natural candidates for stock market return prediction is the stock market return itself. Thus, we examine the predictability of the S&P 500 index monthly return

using the index return from the previous month. In particular, we consider the functional autoregression as follows:

$$\mathbf{d}_{\mathbf{p},t+1} = \mathbf{A}\mathbf{d}_{\mathbf{p},t} + \varepsilon_{t+1}, \quad (11)$$

where $\mathbf{d}_{\mathbf{p},t+1}$ and $\mathbf{d}_{\mathbf{p},t}$ are the demeaned versions of P -density functions constructed using daily returns over the periods from t to $t + 1$ and from $t - 1$ to t , respectively. As we focus on the first moment of the left-hand side of the above functional regression in the empirical analysis, we obtain a fitted P -density $\hat{\mathbf{d}}_{\mathbf{p},t+1}$ using Equation (11) and the first moment of the density. Then, we examine its forecasting power using in-sample and out-of-sample R^2 statistics, as described in Section 4.1 and Section 4.2.

Table 5 reports the result of the functional autoregression analysis using Equation (11). In particular, Panel A and Panel B of Table 5 provide R^2 statistics and corresponding p -values for the in-sample and out-of-sample predictions, respectively. Each row in Panel A and Panel B corresponds to the result with different numbers of functional principal components used to estimate the functional autoregression. The R^2 statistics of the in-sample estimation range from 2.290% to 2.881%, depending on the number of principal components in the functional autoregression. For the out-of-sample prediction (Panel B), the R^2 statistics become negative. That is, the predictive functional autoregression fails to deliver meaningful forecasting power out of sample. This finding is consistent with the existing evidence that stock market return predictability using historical data often deteriorates in an out-of-sample context.

The results using the previous month's physical density in Table 5 suggest that the functional regression method alone is not able to provide enough forecasting power. In comparison with the superior out-of-sample performance documented in Section 4.2, the analysis in this section shows that the forecasting power of the approach proposed in this paper is attributable to the linkage between the physical density and the risk-neutral density of stock market returns.

4.4.2 Predictive Regression Using Finite Moments of the Risk-Neutral Density

There have been many attempts to utilize option price data to build a link to asset prices. Roll et al. (2010), Johnson and So (2012), and Hu (2014) suggest that option trading activities contain information on future stock returns, while Cremers and Weinbaum (2010), Rehman and Vilkov (2012), and Chang et al. (2013) show that the price and implied moments of options predict the cross-section of stock returns. As we argue that using a risk-neutral density embedded in option price data offers significant predictive ability on aggregate market returns, the implied moments in option data, such as implied volatility, skewness and kurtosis, seem to be other natural candidates as predictors.

Table 6 shows the simple predictive in-sample regression results using moments of the risk-neutral density to predict the next month’s return of the S&P 500 index. In particular, the dependent variable in the regression is the mean of the physical density of the S&P 500 index return in month $t + 1$ constructed as in Section 3.2. The independent variables are the moments of the risk-neutral density extracted from the option panel data, as described in Section 3.1. That is, $Mean_t^Q$, $Variance_t^Q$, $Skewness_t^Q$, and $Kurtosis_t^Q$ are the mean, variance, skewness, and kurtosis of the risk-neutral density of the option prices in month t . The regression result shows that the finite moments embedded in the option data do not exhibit any significant return predictive ability. While the regression coefficients on the mean and variance are much larger than those on the third and fourth moments, none of these moments of the risk-neutral density produce statistically significant regression coefficients or R^2 values.

As our approach exploits the rich information from option data to predict stock market returns, one might naturally think of other alternative variables from option data. The CBOE VIX is often referred to as the “investor fear gauge” and is shown to have components that are useful in forecasting stock market returns (for example, among many others, see Ang et al. (2006), Bekaert and Hoerova (2014), and Bardgett et al. (2019)). Additionally, the variance risk premium (VRP), which is defined as the difference between the actual and risk-

neutral expectation of stock market variation, is known to help forecast future stock market returns (see Carr and Wu (2009), Bakshi and Kapadia (2003), and Bollerslev et al. (2015), among many others). Thus, we investigate the performances of three variables, namely, the VIX, a change in the VIX, and the VRP, in forecasting the next month’s stock market return. To construct the variance risk premium, the risk-neutral expectation of variance is estimated as the de-annualized VIX-squared ($VIX^2/12$), and the realized variance is constructed as the sum of squared daily log returns of the S&P 500 index over the month.

Table 7 provides the results of the predictive regressions with the three option-related measures on the next month’s S&P 500 index return. For the in-sample estimation, the VRP exhibits statistically significant performance, while the VIX and the change in the VIX show marginal and statistically non-significant predictive power, respectively. However, when we move to the out-of-sample forecasting performance, the VRP fails to deliver statistically significant prediction power. Furthermore, only the VIX is able to exhibit marginally significant out-of-sample predictive ability at the 5% level. The evidence documented in Table 7 implies that the option variables widely known to exhibit some degree of return predictability still fail to deliver satisfactory forecasting performance, especially in comparison with the superior performance of the approach in this paper.

We emphasize that using a finite number of moments extracted from option price data does not provide any meaningful predictive power. This finding suggests that the rich information embedded in the full risk-neutral density is the main driver of the superior return predictive power. Along with the evidence provided in the previous subsection, the result in this section shows that neither the econometric method using functional regression nor the exploitation of option data alone can account for the superior performance documented in this paper; that is, the combination of exploiting extensive option data with an advanced econometric method is a crucial contributor to the superior predictive power.

4.5 In Relation to Other Equilibrium Models

In this subsection, we examine the validity of frequently used equilibrium models using our prediction algorithm. A specific equilibrium model suggests a kernel that relates the P - and Q - densities. Therefore, if we impose such a relation in our estimation procedure, the functional regression results may be improved through the regularization of the equilibrium model. We consider four different equilibrium models: a model based on the recovery theorem of Ross (2015), a CRRA model, the long-run risk model of Bansal and Yaron (2004), and the external habit model of Campbell and Cochrane (1999). For the long-run risk model and the external habit model,²¹ we use the parameter values proposed in the original papers. For the risk aversion parameter in the CRRA model and the state variables in the long-run risk model and the external habit model, we calibrate them to minimize the distance between the observed P -density and the density implied by an equilibrium model.

Table 8 reports the prediction results of the equilibrium models. For all four models, the in-sample estimations deliver marginally and statistically significant R^2 statistics that range from 2.224% to 2.925%. However, in regard to the out-of-sample prediction, all four models exhibit very poor performance. When comparing the results in Table 8 to the performance of the approach proposed in this paper (Table 2 and Table 3), we observe stark improvements in the in-sample as well as the out-of-sample forecasting performance. This finding suggests that a linkage between the risk-neutral density from the option panel data and the physical density provides unique benefits in forecasting stock market returns. Thus, we conclude that the predictability of market returns through option panel data is a phenomenon not captured by these existing equilibrium models.

²¹In particular, the state variable in the external habit model of Campbell and Cochrane (1999) is the log surplus consumption ratio (s_t). The state variables in the long-run risk model of Bansal and Yaron (2004) are (1) a small persistent predictable component (x_t) in consumption and dividend growth rates and (2) consumption volatility (σ_t).

4.6 Risk-Neutral Density and Existing Forecasting Variables: LASSO Analysis

Finally, we examine the relationship between the information embedded in the option panel data that we utilized to predict stock market returns and conventional predictors that have been frequently employed in the return predictability literature. In doing so, we use the least absolute shrinkage and selection operator (LASSO) method (Tibshirani (1996)). Aiming to identify relationships between risk-neutral density and widely used predictors in the literature, we start from the 14 variables used in Welch and Goyal (2008), stacked in X_t . Using the risk-neutral density extracted from the option panel, we construct the k -th factor loading, f_t^k , which is the time-series of coefficients multiplied to the corresponding eigenvector (the k -th factor) in each time period. With X_t and f_t^k , we estimate the following LASSO problem to identify the variables that have significant effects on f_t^k :

$$\min_{\beta_0, \beta} \left(\sum_{t=1}^T (f_t^k - \beta_0 - \beta' X_t)^2 + \lambda \|\beta\| \right), \quad (12)$$

where λ is a nonnegative regularization parameter and $\|\cdot\|$ is the standard ℓ^1 -norm.

Table 9 provides the result for the LASSO analysis of the first three factors on existing forecasting variables. The first, second, and third columns contain the top five forecasting variables that are strongly associated with the first, second, and third factors extracted from the time-series of the risk-neutral densities. The results indicate that among the well-known predictors used in Welch and Goyal (2008), the dividend yield spread and stock variance are most strongly associated with all three factors of the risk-neutral density dynamics. In addition, inflation, net equity expansion, and book-to-market ratio are the next three most important variables in explaining the first factor, while the term spread, long-term yield, and net equity expansion have a significant association with the second and third factors of the risk-neutral density. This evidence implies that the main factors embedded in risk-neutral density contain unique features that are not captured by any single existing forecasting

variable and the nonlinear relation with those variables.

5 Conclusion

We propose a novel methodology that fully exploits the rich information embedded in option price data to predict aggregate stock market returns. Our methodology is easy to apply and statistically robust. In particular, our approach combines the risk-neutral density extraction method of Ait-Sahalia and Duarte (2003) and the functional regression method of Bosq (2000) and Park and Qian (2012). Applying the proposed method to a large panel of option data, we obtain statistically significant predictive power in forecasting the next month's stock market return.

Not only the statistical significance of our predictions but also the economic gains from using our proposed approach are very substantial. Regardless of whether we impose the restrictions of Campbell and Thompson (2008), our approach exhibits considerable forecasting performance. In particular, our approach delivers an in-sample forecasting performance with an R^2 of 4.720% , and the out-of-sample forecast achieves an R^2 of 6.198%. Furthermore, this superior performance can easily be translated into substantial utility gains if a mean-variance investor exploits our functional predictive regression approach. Overall, our findings imply that our approach exploits the rich information embedded in the full distribution of option price data, which has not yet been fully discovered in the literature.

Our analysis reveals that applying an advanced econometric methodology to a potent data source is the key driver of the main findings of this paper. Thus, we see several avenues for future research. A natural next step is to examine the predictability of other macro variables, such as interest rates or exchange rates, using the data from options markets whose underlying assets reflect dynamics of those macro variables. As machine learning techniques are recently widely used in the various strands of finance literature, a horse race between our approach and different machine learning methods would be worth investigating. Furthermore, although we focus on the risk premium of market returns, the proposed approach

can be further extended to analyze the risk premium of higher moments of a risk-neutral density such as variance, skewness, and kurtosis.

References

- Ait-Sahalia, Yacine, and Jefferson Duarte, 2003, Nonparametric option pricing under shape restrictions, *Journal of Econometrics* 116, 9–47.
- Ait-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.
- Amaya, Diego, Peter Christoffersen, Kris Jacobs, and Aurelio Vasquez, 2015, Does realized skewness predict the cross-section of equity returns?, *Journal of Financial Economics* 118, 135–167.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.
- Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-hedged gains and the negative market volatility risk premium, *Review of Financial Studies* 16, 527–566.
- Bansal, Ravi, Dana Kiku, and Amir Yaron, 2010, Long run risks, the macroeconomy, and asset prices, *American Economic Review* 100, 542–546.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Banz, Rolf W., and Merton H. Miller, 1978, Prices for State-contingent Claims: Some Estimates and Applications, *Journal of Business* 51, 653–672.
- Bardgett, Chris, Elise Gourier, and Markus Leippold, 2019, Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets, *Journal of Financial Economics* 131, 593–618.
- Bekaert, Geert, and Marie Hoerova, 2014, The VIX, the variance premium and stock market volatility, *Journal of Econometrics* 183, 181–192.
- Bliss, Robert R., and Nikolaos Panigirtzoglou, 2002, Testing the stability of implied probability density functions, *Journal of Banking & Finance* 26, 381–422.
- Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.

- Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears, and risk premia, *Journal of Finance* 66, 2165–2211.
- Bollerslev, Tim, and Viktor Todorov, 2014, Time-varying jump tails, *Journal of Econometrics* 183, 168–180.
- Bollerslev, Tim, Viktor Todorov, and Lai Xu, 2015, Tail risk premia and return predictability, *Journal of Financial Economics* 118, 113–134.
- Bosq, Denis, 2000, *Linear Processes in Function Spaces: Theory and Applications*, Lecture Notes in Statistics 149 (Springer-Verlag, New York).
- Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.
- Campbell, John Y., and John Cochrane, 1999, Force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Samuel B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509–1531.
- Carr, Peter, and Liuren Wu, 2009, Variance risk premiums, *Review of Financial Studies* 22, 1311–1341.
- Carr, Peter, and Jiming Yu, 2012, Risk, return, and Ross recovery, *Journal of Derivatives* 20, 38.
- Chang, Bo Young, Peter Christoffersen, and Kris Jacobs, 2013, Market skewness risk and the cross section of stock returns, *Journal of Financial Economics* 107, 46–68.
- Christoffersen, Peter, Steven Heston, and Kris Jacobs, 2013, Capturing option anomalies with a variance-dependent pricing kernel, *Review of Financial Studies* 26, 1963–2006.
- Christoffersen, Peter, and Stefano Mazzotta, 2005, The accuracy of density forecasts from foreign exchange options, *Journal of Financial Econometrics* 3, 578–605.
- Clark, Todd E., and Kenneth D. West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.


- Cochrane, John H., 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533–1575.
- Cowles, Alfred, 1933, Can stock market forecasters forecast?, *Econometrica* 1, 309–324.
- Cowles, Alfred, and Herbert E. Jones, 1937, Some a posteriori probabilities in stock market action, *Econometrica* 5, 280–294.
- Cremers, Martijn, and David Weinbaum, 2010, Deviations from put-call parity and stock return predictability, *Journal of Financial and Quantitative Analysis* 45, 335–367.
- Dangl, Thomas, and Michael Halling, 2012, Predictive regressions with time-varying coefficients, *Journal of Financial Economics* 106, 157–181.
- Fama, Eugene F., 1965, The behavior of stock-market prices, *Journal of Business* 38, 34–105.
- Fama, Eugene F., 1970, Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance* 25, 383–417.
- Fama, Eugene F., 1991, Efficient capital markets: II, *Journal of Finance* 46, 1575–1617.
- Feunou, Bruno, Mohammad R Jahan-Parvar, and CÃ©dric Okou, 2018, Downside Variance Risk Premium, *Journal of Financial Econometrics* 16, 341–383.
- Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.
- Guidolin, Massimo, and Allan Timmermann, 2007, Asset allocation under multivariate regime switching, *Journal of Economic Dynamics and Control* 31, 3503–3544.
- Hansen, Lars Peter, and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.
- Hansen, Lars Peter, and Scott F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.
- Harrison, J. Michael, and David M. Kreps, 1979, Martingales and arbitrage in multiperiod securities markets, *Journal of Economic Theory* 20, 381–408.
- Henkel, Sam James, J. Spencer Martin, and Federico Nardari, 2011, Time-varying short-horizon predictability, *Journal of Financial Economics* 99, 560–580.

- Hu, Jianfeng, 2014, Does option trading convey stock price information?, *Journal of Financial Economics* 111, 625–645.
- Jackwerth, Jens Carsten, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433–451.
- Jackwerth, Jens Carsten, 2004, *Option-Implied Risk-Neutral Distributions and Risk Aversion* (Charlottesville : Research Foundation of AIMR).
- Jackwerth, Jens Carsten, and Mark Rubinstein, 1996, Recovering probability distributions from option prices, *Journal of Finance* 51, 1611–1632.
- Johnson, Travis L., and Eric C. So, 2012, The option to stock volume ratio and future returns, *Journal of Financial Economics* 106, 262–286.
- Kilic, Mete, and Ivan Shaliastovich, 2019, Good and Bad Variance Premia and Expected Returns, *Management Science* 67, 2522–2544.
- Lustig, Hanno N., and Stijn G. Van Nieuwerburgh, 2005, Housing collateral, consumption insurance, and risk premia: An empirical perspective, *Journal of Finance* 60, 1167–1219.
- Marquering, Wessel, and Marno Verbeek, 2004, The economic value of predicting stock index returns and volatility, *Journal of Financial and Quantitative Analysis* 39, 407–429.
- Park, Joon Y., and Junhui Qian, 2012, Functional regression of continuous state distributions, *Journal of Econometrics* 167, 397–412.
- Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2010, Out-of-sample equity premium prediction: Combination forecasts and links to the real economy, *Review of Financial Studies* 23, 821–862.
- Rehman, Zahid, and Grigory Vilkov, 2012, Risk-neutral skewness: Return predictability and its sources, *Available at SSRN 1301648* .
- Roll, Richard, Eduardo Schwartz, and Avanidhar Subrahmanyam, 2010, O/S: The relative trading activity in options and stock, *Journal of Financial Economics* 96, 1–17.
- Rosenberg, Joshua V., and Robert F. Engle, 2002, Empirical pricing kernels, *Journal of Financial Economics* 64, 341–372.
- Ross, Steve, 2015, The recovery theorem, *Journal of Finance* 70, 615–648.


- Samuelson, Paul A., 1965, Proof that properly anticipated prices fluctuate randomly, *IMR; Industrial Management Review (pre-1986)* 6, 41.
- Tibshirani, Robert, 1996, Regression shrinkage and selection via the Lasso, *Journal of the Royal Statistical Society* 58, 267–288.
- van Binsbergen, Jules H., and Ralph S.J. Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance* 65, 1439–1471.
- Wachter, Jessica A., and Missaka Warusawitharana, 2009, Predictable returns and asset allocation: Should a skeptical investor time the market?, *Journal of Econometrics* 148, 162–178.
- Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.

Figure 1: Estimated Risk-neutral (Q) and Physical (P) Densities from Sample Data

Panel A: Q -density




Panel B: P -density




The plots above display the estimated Q -density (Panel A) and the P -density (Panel B) from our sample data. The Q -density is estimated by following Ait-Sahalia and Duarte (2003) as described in Section 3.1. The P -density is obtained using daily returns of the S&P 500 index as described in Section 3.2. In each top and bottom section, we provide the estimated Q - and P -densities along with their demeaned densities, which are used in our main predictive analysis in Section 4.

Figure 2: Time Matching and Aggregation of Option and Stock Market Return Data



The above timeline displays how the observation and expiration dates of the option data are coordinated and how the stock market return data are aggregated accordingly. The option data are collected on *observation dates*, which are 30 days before the option *expiration dates*. That is, the collected S&P 500 index options have a 30-day time to maturity. Once these *observation* and *expiration* dates are specified, daily returns on the S&P 500 from the *observation* to *expiration* dates are collected.


Figure 3: Scree Plot: Risk-neutral Density Extracted from S&P 500 Index Options




The figure presents the scree plot of eigenvalues of the sample variance plotted against the number of factors as described in Step 2 of Section 3.3. The x -axis of the plot represents the number of factors from the risk-neutral density, and the y -axis represents the factors obtained from the functional principal component analysis.

Figure 4: Functional Principal Component Analysis


Panel A: First Factor



Panel B: Second Factor




Panel C: Third Factor




This plot represents the functional principal component analysis. The top, middle, and bottom left figures plot the first ($\nu_{T,1}$), second ($\nu_{T,2}$), and third ($\nu_{T,3}$) functional principal components extracted from the risk-neutral density over our sample period, respectively. The top, middle, and bottom right figures show how each of the three factors move the average risk-neutral density function. In each of the right figures, the blue line (*mean*) represents the time-series average of the risk-neutral density, $\bar{\mathbf{q}} = (1/T)\sum_{t=1}^T \mathbf{q}_t$. The green line ($\text{mean} + \min \times \nu_{T,k}$) and the red line ($\text{mean} + \max \times \nu_{T,k}$) represent $\bar{\mathbf{q}} + \langle \nu_{T,k}, \mathbf{q}_t \rangle_{\min} \nu_{T,k}$ and $\bar{\mathbf{q}} + \langle \nu_{T,k}, \mathbf{q}_t \rangle_{\max} \nu_{T,k}$ for $k = 1, 2, 3$, respectively.

Figure 5: Portfolio Performance using Functional Regression Approach

Panel A: Cumulative gains with $\gamma=3$



Panel B: Cumulative gains with $\gamma=10$



The two figures show the portfolio performances of the functional regression approach proposed in this paper. Each line represents the time-series of the cumulative gains of investing one dollar at the beginning of the out-of-sample forecasting period. The red line (*Functional Regression Forecast*) represents the cumulative gain of a mean-variance optimizing investor (with a risk aversion coefficient of 3 in Panel A and 10 in Panel B) who utilizes our proposed approach to forecast \hat{r}_{t+1} and $\hat{\sigma}_{t+1}^2$ using the functional regression, and the blue line (*Historical Average Forecast*) shows the portfolio performance using the historical average of market returns in the mean-variance optimization with the rolling-window estimate of the variance.

Table 1: Descriptive Statistics

Year	S&P 500 Index Avg.	No. of Put Options	Strike Price Range				Implied Volatility			
			Min	Median	Mean	Max	Min	Median	Mean	Max
1996	674.85	287	375	645.0	641.59	775	0.0865	0.1778	0.1974	0.7993
1997	875.86	370	400	830.0	831.68	1075	0.1500	0.2354	0.2620	0.9156
1998	1087.86	386	400	1025.0	1005.47	1275	0.1096	0.2600	0.3019	1.3293
1999	1330.58	353	600	1275.0	1235.42	1550	0.1009	0.2598	0.2916	0.9613
2000	1419.73	297	850	1395.0	1365.74	1800	0.0813	0.2470	0.2692	0.7016
2001	1185.75	305	700	1130.0	1137.72	1800	0.1350	0.2900	0.3209	1.0501
2002	988.59	331	600	975.0	977.05	1800	0.1487	0.3064	0.3289	0.9350
2003	967.93	328	500	930.0	920.70	1300	0.1139	0.2280	0.2535	0.6290
2004	1133.97	365	700	1090.0	1074.12	1300	0.0848	0.1611	0.1860	0.5351
2005	1207.77	410	800	1170.0	1152.28	1400	0.0741	0.1443	0.1632	0.5458
2006	1318.31	505	800	1255.0	1244.95	1500	0.0575	0.1505	0.1662	0.5459
2007	1478.10	660	900	1395.0	1383.64	1700	0.0782	0.2098	0.2179	0.5905
2008	1215.22	870	200	1162.5	1107.11	1900	0.1393	0.3282	0.4147	1.6948
2009	948.52	902	300	835.0	829.77	1500	0.1392	0.3819	0.4103	1.1949
2010	1130.68	1016	400	1005.0	989.55	1500	0.1104	0.3082	0.3393	1.0780
2011	1280.76	1014	400	1130.0	1099.66	1700	0.0997	0.3364	0.3699	1.2981
2012	1386.51	902	500	1250.0	1223.90	1550	0.0907	0.2444	0.2709	1.1321
2013	1652.29	1114	500	1460.0	1453.11	2000	0.0606	0.2329	0.2453	1.1972
2014	1944.41	1256	1000	1725.0	1700.35	2175	0.0450	0.2305	0.2540	0.7937
2015	2051.93	2016	300	1720.0	1707.86	2250	0.0556	0.3003	0.3236	1.9824
All	1263.98	13687	200	1225	1259.15	2250	0.0450	0.2579	0.2956	1.9824

The table shows descriptive statistics of the put options used in the main analyses. Each row represents annual averages, and the last row provides statistics of the data for the full sample period from 1996-2015. The second and third columns show the average of the S&P 500 index and the total number of put options on the index used to extract the risk-neutral probability distribution, respectively. The next four columns (the last four columns) provide minimum, median, mean, and maximum strike prices (implied volatility) of the put options.

Table 2: In-Sample Prediction Results

Panel A: Functional Regression		
Number of Factors	In-Sample R^2 (%)	p-value
K = 3	4.375***	[0.001]
K = 4	4.487***	[0.000]
K = 5	4.720***	[0.000]
Panel B: Goyal and Welch (2008) Variables		
Variable	In-Sample R^2 (%)	p-value
Dividend-Price Ratio	1.113	[0.104]
Dividend Yield	1.437*	[0.064]
Earnings-Price Ratio	0.246	[0.445]
Dividend Payout Ratio	0.004	[0.922]
Stock Variance	2.084**	[0.026]
Book-to-Market Ratio	0.174	[0.521]
Net Equity Expansion	1.840**	[0.036]
Treasury Bill Rate	0.001	[0.961]
Long-Term Yield	0.068	[0.687]
Long-Term Return	0.134	[0.573]
Term Spread	0.098	[0.630]
Default Yield Spread	0.606	[0.231]
Default Return Spread	0.477	[0.288]
Inflation	0.805	[0.167]
Kitchen Sink (All)	13.578***	[0.001]

The table reports the R^2 statistics of the functional predictive regression provided in Section 3 and R^2 statistics of variables used in Welch and Goyal (2008). In Panel A, the in-sample R^2 statistics from the functional predictive regression are computed using Equation (3) in Section 3.3. The value of K in the first column represents the number of functional principal components used in the estimation of the functional regression. In Panel B, the in-sample R^2 statistics for the variables from Welch and Goyal (2008) are computed from the predictive regression of Equation (5). The sample period of estimation spans from January 1996 to December 2015. The numbers in brackets are the p -values for the F -statistics of the regressions. Asterisks denote the significance of the in-sample regression as measured by its corresponding p -value. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 3: Out-of-Sample Prediction Results

Panel A: Functional Regression				
Number of Factors	Out-of-Sample R^2 (%)	p-value		
K = 3	6.012***	[0.000]		
K = 4	5.749***	[0.000]		
K = 5	6.198***	[0.000]		

Panel B: Goyal and Welch (2008) Variables				
	Unrestricted		Campbell and Thompson (2008) Restrictions	
	Out-of-Sample R^2 (%)	p-value	Out-of-Sample R^2 (%)	p-value
Dividend-Price Ratio	1.175**	[0.028]	1.175**	[0.028]
Dividend Yield	1.061*	[0.059]	1.061*	[0.059]
Earnings-Price Ratio	-0.868	[0.718]	-0.868	[0.718]
Dividend Payout Ratio	-1.798	[0.729]	-1.798	[0.729]
Stock Variance	0.547	[0.362]	0.000	[-]
Book-to-Market Ratio	-2.212	[0.778]	0.277	[0.292]
Net Equity Expansion	-5.326	[0.870]	-4.846	[0.855]
Treasury Bill Rate	-3.745	[0.738]	-3.862	[0.753]
Long-Term Yield	-4.460	[0.804]	-0.459	[0.795]
Long-Term Return	-1.879	[0.811]	-1.176	[0.685]
Term Spread	-1.874	[0.652]	-1.874	[0.652]
Default Yield Spread	-2.745	[0.811]	-2.745	[0.811]
Default Return Spread	-1.955	[0.573]	-2.560	[0.694]
Inflation	-7.456	[0.936]	-6.809	[0.926]
Kitchen Sink	-9.190	[0.278]	-9.212	[0.279]

The table reports the out-of-sample R^2 statistics of the functional predictive regression approach provided in Section 3. The out-of-sample R^2 statistics are computed following Campbell and Thompson (2008) as in Equation (6). Panel A reports the out-of-sample R^2 statistics (and p -values in brackets) of the functional predictive regression using different numbers of functional principal components. Panel B provides the out-of-sample R^2 statistics (and p -values in brackets) of the unrestricted and restricted (with Campbell and Thompson (2008) restrictions) predictive regressions of the forecasting variables used in Welch and Goyal (2008). The period of the out-of-sample prediction is the last 5 years of our sample period, starting in January 2011. The value of K in the first column represents the number of functional principal components used in the estimation of the functional regression. The sample period spans from January 1996 to December 2015. The numbers in brackets are the p -values for the Clark and West (2007) *MSFE (mean squared forecast error)-adjusted* statistic for testing the null hypothesis that the historical average MSFE is less than or equal to the predictive regression MSFE against the alternative that the historical average MSFE is greater than the predictive regression MSFE. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 4: Economic Significance of the Functional Regression Forecast

Panel A: Functional Regression		
Number of Factors	Utility Gain (% per annum)	
	Mean Forecast	Mean-Var Forecast
K = 3	2.743	7.265
K = 4	3.414	7.668
K = 5	2.577	6.838

Panel B: Goyal and Welch (2008) Variables		
	Utility Gain (% per annum)	
	Unrestricted	Campbell and Thompson (2008) Restrictions
Dividend-Price Ratio	2.359	2.359
Dividend Yield	2.791	2.791
Earnings-Price Ratio	0.877	0.877
Dividend Payout Ratio	-0.307	-0.307
Stock Variance	1.814	-
Book-to-Market Ratio	0.626	0.845
Net Equity Expansion	-1.010	-1.010
Treasury Bill Rate	-1.105	-1.051
Long-Term Yield	0.049	0.737
Long-Term Return	-0.585	-0.532
Term Spread	-0.418	-0.418
Default Yield Spread	-0.051	-0.051
Default Return Spread	0.355	0.355
Inflation	-2.721	-2.721
Kitchen Sink (All)	4.478	4.478

This table reports the utility gains (in terms of the certainty equivalent return, expressed in a percentage per annum) of a mean-variance investor who allocates her portfolio between stocks and risk-free bills based on predictions using functional regression (Panel A) and the well-known economic variables in Welch and Goyal (2008) (Panel B). The utility gain in each panel can be considered the portfolio management fee (in annualized percent return) that an investor with a mean-variance preference and a risk aversion coefficient of five would be willing to pay to have access to the predictive forecast based on our functional regression approach (Panel A) or the economic variables (Panel B) in place of the historical average benchmark forecast. In Panel A, the second column (Mean Forecast) reports the average utility gains of a mean-variance investor who utilizes the functional predictive regression to obtain the forecast \hat{r}_{t+1} and the rolling-window estimate of the variance for $\hat{\sigma}_{t+1}^2$. Under the third column (Mean-Var Forecast), we report the average utility gains of a mean-variance investor who further utilizes the functional predictive regression to obtain the forecast $\hat{\sigma}_{t+1}^2$. Panel B provides the average utility gain of the investor when she uses the economic variables in Welch and Goyal (2008) with the unrestricted and restricted (Campbell and Thompson (2008) restrictions) predictive regressions.

Table 5: Functional Predictive Regression with the Physical Density

Panel A: In-Sample Estimation		
Number of Factors	In-Sample $R^2(\%)$	p-value
K = 3	2.290**	[0.019]
K = 4	2.476**	[0.015]
K = 5	2.881***	[0.008]
Panel B: Out-of-Sample Prediction		
Number of Factors	Out-of-Sample $R^2(\%)$	p-value
K = 3	-6.586	[0.706]
K = 4	-5.292	[0.637]
K = 5	-3.665	[0.536]

This table provides the R^2 statistics of the functional predictive regression of next month's return of the S&P 500 index on the physical density constructed using the current month's S&P daily returns. The dependent variable, the next month's return r_{t+1} , is computed as the mean of the P -density in month $t + 1$ constructed as described in Section 3.2. The independent variable is the P -density in month t . The value of K in the first column represents the number of functional principal components used in the functional regression estimation. In Panel A, we provide the in-sample estimation results over the full sample period from January 1997 to December 2015. The in-sample R^2 statistics in Panel A are computed using Equation (3) in Section 3.3. The numbers in brackets in Panel A report the p -values for the F -statistics of the regressions. In Panel B, we report the out-of-sample R^2 statistics of the predictive functional regression over the last 5 years of our sample period, January 2011 to December 2015. The out-of-sample R^2 statistics are computed as in Equation (6), following Campbell and Thompson (2008). The numbers in brackets in Panel B are the p -values for the Clark and West (2007) *MSFE (mean squared forecast error)-adjusted* statistic for testing the null hypothesis that the historical average MSFE is less than or equal to the predictive regression MSFE against the alternative that the historical average MSFE is greater than the predictive regression MSFE. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Predictive Regression with Finite Moments of Risk-Neutral Density

Independent Variable	Dependent Variable: r_{t+1}			
	(1)	(2)	(3)	(4)
$Mean_t^Q$	0.356 (1.316)	-0.066 (-0.160)	0.100 (0.165)	0.085 (0.139)
$Variance_t^Q$		1.187 (1.356)	1.041 (1.085)	1.002 (1.029)
$Skewness_t^Q$			0.000 (-0.377)	0.000 (-0.254)
$Kurtosis_t^Q$				0.000 (-0.257)
R-squares (%)	0.725	1.492	1.552	1.579
Adj R-squares (%)	0.306	0.657	0.295	-0.103

This table provides the regression results for a simple predictive regression of the next month's return of the S&P 500 index on moments of the risk-neutral density extracted from the current month's options data. The dependent variable, the next month's stock market return r_{t+1} , is computed as the mean of the P -density in month $t+1$ constructed following Section 3.2. The dependent variables are the first four moments of the risk-neutral density extracted from the option price data, as described in Section 3.1. In particular, $Mean_t^Q$, $Variance_t^Q$, $Skewness_t^Q$, and $Kurtosis_t^Q$ represent the mean, variance, skewness, and kurtosis of the risk-neutral density estimated in month t . The numbers in parentheses are the t -statistics of the regression coefficients. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 7: Predictive Regression with Option-related Variables

	In-Sample		Out-of-Sample	
	$R^2(\%)$	p -value	$R_{OOS}^2(\%)$	p -value
VIX	1.710**	[0.043]	4.374**	[0.047]
Δ VIX	0.440	[0.306]	1.439	[0.139]
VRP	6.448***	[0.000]	1.406	[0.225]

This table provides the results for a predictive regression of the next month's return of the S&P 500 index on other option-related variables that are known to exhibit some predictive power. The option-related variables include the CBOE VIX, the change in the VIX (Δ VIX), and the variance risk premium (VRP). The VRP is defined as the difference between the risk-neutral and objective expectations of the realized variance, where the risk-neutral expectation of the variance is estimated as the de-annualized VIX-squared ($\text{VIX}^2/12$) and the realized variance is constructed as the sum of daily log returns of the S&P index over the month. The two columns under In-sample and the other two columns under Out-of-sample provide the R^2 statistics and corresponding p -values for the in-sample estimation and the out-of-sample prediction, respectively. The numbers in parentheses are the t -statistics of the regression coefficients. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 8: Equilibrium Models

Functional Predictive Regression Result				
Models	In-Sample		Out-of-Sample	
	R^2 (%)	p -value	R^2_{OOS} (%)	p -value
Ross (2015)	2.877***	0.008	-0.463	0.159
CRRA model	2.224**	0.021	-0.736	0.401
Long-run risk model	2.925***	0.008	-0.530	0.159
Habit model	2.913***	0.008	-4.010	0.383

This table reports in-sample and out-of-sample R^2 statistics for the functional predictive regression based on densities estimated from equilibrium models. In particular, four different equilibrium models are considered: (i) a model based on the recovery theorem of Ross (2015), (ii) a CRRA model, (iii) the long-run risk model of Bansal and Yaron (2004) and (iv) the external habit model of Campbell and Cochrane (1999). For the in-sample R^2 statistics, the numbers in brackets are the p -values for the F -statistics of the regressions. For the out-of-sample R^2 statistics, the numbers in brackets are the p -values for the Clark and West (2007) *MSFE (mean squared forecast error)-adjusted* statistic for testing the null hypothesis that the historical average MSFE is less than or equal to the predictive regression MSFE against the alternative that the historical average MSFE is greater than the predictive regression MSFE. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 9: Selected Variables for the First Three Factor of the Risk-Neutral Density

First Factor	Second Factor	Third Factor
Default Yield Spread	Stock Variance	Stock Variance
Stock Variance	Default Yield Spread	Default Yield Spread
Inflation	Term Spread	Term Spread
Net Equity Expansion	Long Term Yield	Long Term Yield
Book to Market Ratio	Net Equity Expansion	Net Equity Expansion

This table presents the variables selected to explain the first three principal components extracted from the dynamics of the risk-neutral density. A complete set of predictors used in the LASSO analysis includes the 14 variables documented in Welch and Goyal (2008). Among all 14 predictors, the table reports the five most significant variables for the three principal components in each column.