

Improving Portfolio Selection Using Option-Implied Volatility and Skewness

Victor DeMiguel, Yuliya Plyakha, Raman Uppal,
and Grigory Vilkov*

Abstract

Our objective in this paper is to examine whether one can use option-implied information to improve the selection of mean-variance portfolios with a large number of stocks, and to document which aspects of option-implied information are most useful to improve their out-of-sample performance. Portfolio performance is measured in terms of volatility, Sharpe ratio, and turnover. Our empirical evidence shows that using option-implied volatility helps to reduce portfolio volatility. Using option-implied correlation does not improve any of the metrics. Using option-implied volatility, risk premium, and skewness to adjust expected returns leads to a substantial improvement in the Sharpe ratio, even after prohibiting short sales and accounting for transaction costs.

I. Introduction

To determine the optimal mean-variance portfolio of an investor, one needs to estimate the moments of asset returns, such as means, volatilities, and correlations.

*DeMiguel, avmiguel@london.edu, London Business School, 6 Sussex Place, Regent's Park, London NW1 4SA, United Kingdom; Plyakha, plyakha@gmx.de, Luxembourg School of Finance, University of Luxembourg, 4, rue Albert Borschette, L-1246 Luxembourg; Uppal, raman.uppal@edhec.edu, Edhec Business School, 10 Fleet Place, Ludgate, London EC4M 7RB, United Kingdom and Center for Economic and Policy Research (CEPR); and Vilkov, vilkov@vilkov.net, Goethe University Frankfurt, Finance Department, Grüneburgplatz 1 / Uni-Pf H 25, Frankfurt am Main, D-60323, Germany. We gratefully acknowledge financial support from INQUIRE-Europe; however, this article represents the views of the authors and not of INQUIRE. We acknowledge detailed feedback from Luca Benzoni, Massimo Guidolin (the referee), Pascal Maenhout, and Christian Schlag. We received helpful comments and suggestions from Alexander Alekseev, Hendrik Bessembinder (the editor), Michael Brandt, Mike Chernov, Engelbert Dockner, Bernard Dumas, Wayne Ferson, Rene Garcia, Lorenzo Garlappi, Nicolae Gârleanu, Amit Goyal, Jakub Jurek, Nikunj Kapadia, Ralph Koijen, Lionel Martellini, Vasant Naik, Stavros Panageas, Andrew Patton, Boryana Racheva-Iotova, Marcel Rindisbacher, Paulo Rodrigues, Pedro Santa-Clara, Bernd Scherer, Peter Schotman, George Skiadopoulos, Luis Viceira, Josef Zechner, and seminar participants at AHL (Man Investments), BlackRock, Goethe University Frankfurt, London School of Economics, University of Mainz, University of Piraeus, University of St. Gallen, Vienna University of Economics and Business Administration, CEPR European Summer Symposium on Financial Markets, Duke-University of North Carolina at Chapel Hill Asset Pricing Conference, EDHEC-Risk Seminar on Advanced Portfolio Construction, Financial Econometrics Conference at Toulouse School of Economics, Stockholm Institute for Financial Research Conference on Asset Allocation and Pricing in Light of the Recent Financial Crisis, and meetings of the European Finance Association, Western Finance Association, and the Joint Conference of INQUIRE-UK and INQUIRE-Europe 2011.

Traditionally, historical data on returns are used to estimate these moments, but researchers have found that portfolios based on sample estimates perform poorly out of sample.¹ Several approaches have been proposed for improving the performance of portfolios based on *historical data*.²

In this paper, instead of trying to improve the quality of the moments estimated from historical data, we use *forward-looking* moments of stock-return distributions that are implied by option prices. The main contribution of our work is to evaluate empirically *which aspects* of option-implied information are particularly useful for improving the out-of-sample performance of portfolios with a large number of stocks. Specifically, we consider option-implied volatility, correlation, skewness, and the risk premium for stochastic volatility, and we obtain these not just from the Black-Scholes (1973) model, but also using the model-free approach, which has the benefit that the measurement error resulting from model misspecification is reduced.

In selecting portfolios, we use a variety of moments implied by prices of options. First, we consider the use of option-implied volatilities and correlations to improve out-of-sample performance of mean-variance portfolios invested in only *risky* stocks. When evaluating the benefits of using option-implied volatilities and correlations, we set expected returns to be the same across all assets so that the results are not confounded by the large errors in estimating expected returns.³ Consequently, the mean-variance portfolio reduces to the minimum-variance portfolio. In addition to considering the minimum-variance portfolio based on the sample covariance matrix, we consider also the short-sale-constrained minimum-variance portfolio, the minimum-variance portfolio with shrinkage of the covariance matrix (as in Ledoit and Wolf (2004a), (2004b)), and the minimum-variance portfolio obtained by assuming all correlations are equal to 0 or with correlations set equal to the mean correlation across all asset pairs (as suggested by Elton, Gruber, and Spitzer (2006)). We find that using risk-premium-corrected option-implied volatilities in minimum-variance portfolios improves the out-of-sample volatility by more than 10% compared to the traditional portfolios based on only historical stock-return data, while the changes in the Sharpe ratio are insignificant. Thus, using option-implied volatility allows one to reduce the out-of-sample portfolio volatility significantly.

Next, we examine the use of risk-premium-corrected option-implied correlations to improve the performance of minimum-variance portfolios. We find that in most cases option-implied correlations do not lead to any improvement

¹For evidence of this poor performance, see DeMiguel, Garlappi, and Uppal (2009), Jacobs, Muller, and Weber (2010), and the references therein.

²These approaches include: imposing a factor structure on returns (Chan, Karceski, and Lakonishok (1999)), using data for daily rather than monthly returns (Jagannathan and Ma (2003)), using Bayesian methods (Jobson, Korkie, and Ratti (1979), Jorion (1986), Pástor (2000), and Ledoit and Wolf (2004b)), constraining short sales (Jagannathan and Ma), constraining the norm of the vector of portfolio weights (DeMiguel, Garlappi, Nogales, and Uppal (2009)), and using stock-return characteristics such as size, book-to-market ratio, and momentum to choose parametric portfolios (Brandt, Santa-Clara, and Valkanov (2009)).

³Jagannathan and Ma ((2003), pp. 1652–1653) write, “The estimation error in the sample mean is so large nothing much is lost in ignoring the mean altogether when no further information about the population mean is available.”

in performance. Our empirical results indicate that the gains from using implied correlations are not substantial enough to offset the higher turnover resulting from the increased instability over time of the covariance matrix when it is estimated using option-implied correlations.

Finally, to improve the out-of-sample performance of mean-variance portfolios, we consider the use of option-implied volatility, risk premium for stochastic volatility, and option-implied skewness. These characteristics have been shown in the literature to help explain the cross section of expected returns.⁴ Therefore, it makes sense to explore their effect in the framework of mean-variance portfolios. Using these characteristics to rank stocks and adjusting by a scaling factor the expected returns of the stocks, or using these characteristics with the parametric-portfolio methodology of Brandt et al. (2009), leads to a substantial improvement in the Sharpe ratio, even after prohibiting short sales and accounting for transaction costs.

We conclude this Introduction by discussing the relation of our work to the existing literature. The idea that option prices contain information about future moments of asset returns has been understood ever since the work of Black and Scholes (1973) and Merton (1973); Poon and Granger (2005) provide a comprehensive survey of this literature. The focus of our work is to investigate how the information implied by option prices can be used to improve portfolio selection. Two other papers study this question. The first, by Ait-Sahalia and Brandt (2008), uses option-implied state prices to solve for the intertemporal consumption and portfolio choice problem using the Cox and Huang (1989) martingale representation formulation rather than the Merton (1971) dynamic-programming formulation; however, the focus of the paper is not on finding the optimal portfolio with superior out-of-sample performance. The second, by Kostakis, Panigirtzoglou, and Skiadopoulos (2011), studies the *asset-allocation* problem of allocating wealth between the Standard & Poor's (S&P) 500 index and a riskless asset. That paper finds that the out-of-sample performance of the portfolio based on the return distribution inferred from option prices is better than that of a portfolio based on the historical distribution. However, there is an important difference between that work and ours: Rather than considering the problem of how to allocate wealth between the S&P 500 index and the risk-free asset, we consider the *portfolio-selection* problem of allocating wealth across a large number of individual stocks. It is not clear how one would extend the methodology of Kostakis et al. to accommodate a large number of risky assets. They also need to make other restrictive assumptions, such as the existence of a representative investor and the completeness of financial markets, which are not required in our analysis.

The rest of the paper is organized as follows: In Section II, we describe the data on stocks and options that we use. In Section III, we explain how we use data on options to predict volatilities, correlations, and expected returns. In Section IV,

⁴For instance, Bollerslev, Tauchen, and Zhou (2009) document a positive relation between the variance risk premium and future returns. Bali and Hovakimian (2009) and Goyal and Saretto (2009) show that stocks with a large spread between Black-Scholes (1973) implied volatility and realized volatility tend to outperform those with low spreads. Bali and Hovakimian (2009), Xing, Zhang, and Zhao (2010), and Cremers and Weinbaum (2010) find a positive relation between various measures of option-implied skewness and future stock returns.

we describe the construction of the various portfolios we evaluate along with the benchmark portfolios, and the metrics used to compare the performance of these portfolios. Our main findings about the performance of various portfolios that use option-implied information are given in Section V. We conclude in Section VI.

II. Data

In this section, we describe the data on stocks and stock options that we use in our study. Our data on stocks are from the Center for Research in Security Prices (CRSP). To implement the parametric-portfolio methodology, we also use data from Compustat. Our data for options are from IvyDB (OptionMetrics). Our sample period is from Jan. 1996 to Oct. 2010.⁵

A. Data on Stock Returns and Stock Characteristics

We study stocks that are in the S&P 500 index at any time during our sample period. The *daily* stock returns of the S&P 500 constituents are from the daily file of CRSP, and we have in our sample a total of 3,986 trading days.⁶ Counting by CRSP identifiers (PERMNO), we have data for a maximum of 961 stocks. Out of these 961 stocks, there are 143 stocks for which implied volatilities are available for the entire time series. For robustness, we consider two samples in our analysis. The first, which we label “Sample 1,” consists of the 143 stocks for which there are no missing data. The second, “Sample 2,” consists of all the stocks that are part of the S&P 500 index on a particular day and that have no missing data on that day (such as prices of options on the same underlying and the same maturity but across different strikes, which are needed to compute model-free implied volatility (MFIV) and model-free implied skewness (MFIS)). Thus, the second sample has a variable number of stocks; on average, there are about 400 stocks at each date in this sample.⁷

We measure size (market value of equity) as the price of the stock per share multiplied by shares outstanding; both variables are obtained from the CRSP database. For measuring value or book-to-market (BTM) characteristic, we use the Compustat Quarterly Fundamentals file. The 12-month momentum (MOM) characteristic is measured for each day t using daily returns data from CRSP as the cumulative return from day $t - 251 - 21$ to day $t - 21$.⁸

⁵We carry out all the tests included in the manuscript also for the pre-crisis period from Jan. 1996 to Dec. 2007, and the crisis period from Jan. 2008 to June 2009 (identified as a recession by the National Bureau of Economic Research (NBER)). The main insights of our analysis do not change with the choice of sample period. These results are available in the working paper version of the manuscript.

⁶We also use high-frequency *intraday* stock-price data consisting of transaction prices for the S&P 500 constituents from the NYSE's Trade and Quote database; the results for these data are very similar to those using daily data.

⁷The main difference between Sample 1 and Sample 2 is that for estimating the parameters of the covariance matrix for Sample 2, which has more stocks, one needs a longer estimation window. Thus, while the estimation window for Sample 1 is 250 trading days, for Sample 2 it is 750 days. As a result of the longer estimation window, the weights are relatively more stable over time for Sample 2. On the other hand, because the covariance matrix for Sample 2 is of a larger dimension, its condition number is different from that of the covariance matrix for Sample 1.

⁸To get better distributional properties of the size and BTM characteristics we construct, we take the logarithm of size and value characteristics. In order to prepare these characteristics so that they

B. Data on Stock Options

For stock options, we use IvyDB, which contains data on all U.S.-listed index and equity options, most of which are American. We use the volatility surface file, which contains a smoothed implied-volatility surface for a range of standard maturities and a set of option delta points. From the surface file we select the out-of-the-money implied volatilities for calls and puts (we take implied volatilities for calls with deltas smaller than or equal to 0.5, and implied volatilities for puts with deltas bigger than -0.5) for a maturity of 30 days.⁹ For each date, underlying stock, and time to maturity, we have 13 implied volatilities from the surface data, which are used to calculate the moments of the risk-neutral distribution. Some of the option-based characteristics also use the parametric Black-Scholes (1973) implied volatilities for at-the-money options. We compute the at-the-money volatility as the average volatility for a put and a call with absolute delta level equal to 0.5.

III. Option-Implied Information

In this section, we explain how we compute the option-implied moments that we use for portfolio selection; we compare the ability of option-implied moments and the historical moments to forecast the actual realized moments. We consider the following measures: i) model-free option-implied volatility; ii) the volatility risk premium, measured as the spread between realized and Black-Scholes (1973) option-implied volatility; iii) option-implied correlation; iv) model-free option-implied skewness; and v) a proxy for skewness, measured as the spread between the Black-Scholes implied volatility obtained from calls and that from puts.

A. Predicting Volatilities Using Options

When option prices are available, an intuitive first step is to use this information to back out implied volatilities and use them to predict volatility.¹⁰ In contrast to the model-specific Black and Scholes (1973) implied volatility, we use for this purpose MFIV, which represents a nonparametric estimate of the risk-neutral expected stock-return volatility until the option's expiration. It subsumes information in the whole Black-Scholes implied volatility smile (Vanden (2008)) and is expected to predict the realized volatility (RV) better than the Black-Scholes volatility. We compute MFIV as the square root of the variance contract of Bakshi et al. (2003), as explained later.

can be used to compute the parametric-portfolio weights, we also winsorize the characteristics by assigning the value of the 3rd percentile to all values below the 3rd percentile and do the same for values higher than the 97th percentile. Finally, we normalize all characteristics to have zero mean and unit standard deviation.

⁹The use of out-of-the-money options is standard in this literature; see, for instance, Bakshi, Kapadia, and Madan (2003) and Carr and Wu (2009). The reason for selecting options that are out of the money is that it reduces the effect of the premium for early exercise of American options.

¹⁰Note that our objective is *only* to show that the option-implied moments provide better forecasts than the estimators based on historical sample data, rather than to demonstrate that option-implied moments provide the *best* forecasts of future volatility and correlations. There is a very large literature on forecasting stock-return volatility and correlations; see, for instance, the survey article by Andersen, Bollerslev, and Diebold (2009).

Let $S(t)$ be the stock price at time t , $R(t, \tau) \equiv \ln S(t + \tau) - \ln S(t)$ the τ -period log return, and r the risk-free interest rate. Let $V(t, \tau) \equiv \mathbb{E}_t^* \{e^{-r\tau} R(t, \tau)^2\}$, $W(t, \tau) \equiv \mathbb{E}_t^* \{e^{-r\tau} R(t, \tau)^3\}$, and $X(t, \tau) \equiv \mathbb{E}_t^* \{e^{-r\tau} R(t, \tau)^4\}$ represent the fair value of the variance, cubic, and quartic contracts, respectively, as defined in Bakshi et al. (2003). Then, the τ -period MFIV can be calculated as

$$(1) \quad \text{MFIV}(t, \tau) = (V(t, \tau))^{1/2},$$

and the τ -period MFIS as

$$(2) \quad \text{MFIS}(t, \tau) = \frac{e^{r\tau} W(t, \tau) - 3\mu(t, \tau)e^{r\tau} V(t, \tau) + 2(\mu(t, \tau))^3}{(e^{r\tau} V(t, \tau) - (\mu(t, \tau))^2)^{3/2}},$$

where $\mu(t, \tau)$ denotes the risk-neutral expectation of the τ -period log return:

$$\mu(t, \tau) \equiv \mathbb{E}_t^* \left[\ln \frac{S(t + \tau)}{S(t)} \right] = e^{r\tau} - 1 - \frac{e^{r\tau}}{2} V(t, \tau) - \frac{e^{r\tau}}{6} W(t, \tau) - \frac{e^{r\tau}}{24} X(t, \tau).$$

To compute the integrals that give the values of the variance, cubic, and quartic contracts precisely, we need a continuum of option prices. We discretize the respective integrals and approximate them using the available options. As mentioned earlier, we normally have 13 out-of-the-money call and put implied volatilities for each maturity. Using cubic splines, we interpolate them inside the available moneyness range and extrapolate using the last known (boundary for each side) value to fill in a total of 1,001 grid points in the moneyness range from $1/3$ to 3. Then we calculate the option prices from the interpolated volatilities using the known interest rate for a given maturity and use these prices to compute the MFIV and MFIS as in equations (1) and (2), respectively.

However, what we need for portfolio selection is not the risk-neutral implied volatility of stock returns but the expected volatility under the *objective* distribution. We now explain how to use information in the MFIV in order to get the volatility under the objective measure.

The implied volatility differs from the expected volatility under the true measure by the volatility risk premium. Bollerslev, Gibson, and Zhou (2004) and Carr and Wu (2009) show that one can use RV, instead of the expected volatility, to estimate the volatility risk premium. Assuming that the magnitude of the volatility risk premium is proportional to the level of volatility under the true probability measure, we estimate the monthly historical volatility risk premium adjustment (HVRP) for a particular stock as the ratio of average monthly implied and realized volatilities for that stock for the past $T + \Delta t$ trading days:¹¹

$$(3) \quad \text{HVRP}_t = \frac{\sum_{i=t-T-\Delta t+1}^{t-\Delta t} \text{MFIV}_{i,i+\Delta t}}{\sum_{i=t-T-\Delta t+1}^{t-\Delta t} \text{RV}_{i,i+\Delta t}}.$$

¹¹Note that because HVRP_t is calculated as the ratio of the average $\text{MFIV}_{i,i+\Delta t}$ and $\text{RV}_{i,i+\Delta t}$, both of which are calculated over Δt days, we have only T observations when computing the sum.

In our analysis, we estimate the HVRP on each day over the past year (-272 days to -21 days) using the MFIV and RV from daily returns, each measured over 21 trading days and each annualized appropriately. Then, assuming that in the next period, from t to $t + \Delta t$, the prevailing volatility risk premium will be well approximated by the historical volatility risk premium in equation (3), one can obtain the *prediction* of the future realized volatility, \widehat{RV}_t , which we call the *risk-premium-corrected implied volatility*:¹²

$$\widehat{RV}_{t,t+\Delta t} = \frac{\text{MFIV}_{t,t+\Delta t}}{\text{HVRP}_t}.$$

We now wish to confirm the intuition that risk-premium-corrected implied volatility is better than historical volatility at predicting RV. To do this, we consider, as a predictor first historical volatility and then risk-premium-corrected implied volatility, for the monthly RV from daily returns for each stock. We compare the performance of each predictor in terms of root mean squared error (RMSE) and mean prediction error (ME). We find that the average RMSE in Sample 1 for the risk-premium-corrected implied volatility is 0.1274, which is smaller than the RMSE of 0.1671 for historical daily volatility. The ME for both predictors is negative, indicating that on average both measures are biased upward with respect to the RV; however, the ME of -0.0047 for the risk-premium-corrected implied volatility is one order of magnitude smaller than the ME of -0.0185 for historical volatility.

B. Predicting Correlations Using Options

The second piece of option-implied information that we consider is implied correlation; because we need the correlation under the *objective* measure for the portfolio optimization, we discuss directly how to obtain option-implied correlation corrected for the risk premium.

If a portfolio is composed of N individual stocks with weights w_i , $i = \{1, \dots, N\}$, we can write the variance of the portfolio p under the objective (physical) probability measure P as

$$(4) \quad (\sigma_{p,t}^P)^2 = \sum_{i=1}^N w_i^2 (\sigma_{i,t}^P)^2 + \sum_{i=1}^N \sum_{j \neq i}^N w_i w_j \sigma_{i,t}^P \sigma_{j,t}^P \rho_{ij,t}^P.$$

Assume that we have estimated the expectation of the future volatilities of the portfolio $\widehat{\sigma}_{p,t}^P$ and of its components $\widehat{\sigma}_{i,t}^P$, and we want to estimate the set of expected correlations $\widehat{\rho}_{ij,t}^P$ that turn equation (4) into identity. Once we substitute the expected volatilities into equation (4), we have one equation with $N \times (N - 1) / 2$ unknown correlations, $\widehat{\rho}_{ij,t}^P$. Thus, to compute all pairwise correlations we need to make some identifying assumptions.

¹²Another method for obtaining the predictor of future realized volatility is to use a modified version of the heterogeneous autoregressive model of RV proposed by Corsi (2009). We find that the root mean squared error (RMSE) and mean prediction error (ME) with this measure are larger than those using the approach we adopt.

We use the approach proposed in Buss and Vilkov (2012), who compute a heterogeneous implied-correlation matrix, where all pairwise correlations are allowed to be different.¹³ Under their approach, the expected correlation is assumed to differ from the historical correlation by a fixed proportion ψ of the distance between the historical correlation and the maximum correlation of 1:

$$(4a) \quad \rho_{ij,t}^P - \widehat{\rho}_{ij,t}^P = \psi_t(1 - \rho_{ij,t}^P),$$

which implies that

$$(4b) \quad \widehat{\rho}_{ij,t}^P = \rho_{ij,t}^P - \psi_t(1 - \rho_{ij,t}^P).$$

When we substitute this into equation (4) above, we get

$$(5) \quad (\widehat{\sigma}_{p,t}^P)^2 = \sum_i \sum_j w_i w_j \widehat{\sigma}_{i,t}^P \widehat{\sigma}_{j,t}^P (\rho_{ij,t}^P - \psi_t(1 - \rho_{ij,t}^P)),$$

from which one can derive an explicit expression for the parameter ψ_t ,

$$(5a) \quad \psi_t = - \frac{(\widehat{\sigma}_{p,t}^P)^2 - \sum_i \sum_j w_i w_j \widehat{\sigma}_{i,t}^P \widehat{\sigma}_{j,t}^P \rho_{ij,t}^P}{\sum_i \sum_j w_i w_j \widehat{\sigma}_{i,t}^P \widehat{\sigma}_{j,t}^P (1 - \rho_{ij,t}^P)},$$

and then the expected correlations $\widehat{\rho}_{ij,t}^P$ from equation (4b). Thus, we construct the “heterogeneous” implied-correlation matrix corrected for the risk premium, inferred from expected index and individual volatilities, which contains the up-to-date market perception of future correlation under the true measure.

To determine whether option-implied correlation is superior to historical correlation at predicting realized correlation, we compute the RMSE and ME for these two predictors of the 21-day realized correlation. In both Samples 1 and 2, we find that the RMSE for historical correlation is about 0.25, slightly smaller than the RMSE of 0.26 for the option-implied correlation; note, however, that the RMSE for both predictors is only slightly smaller than the average realized correlation of 0.29 for our sample, implying that there is very little predictability. For Sample 1, the ME of 0.0039 for historical correlation is of the same order of magnitude as the ME of -0.0068 for implied correlation, while for Sample 2 the ME of 0.0342 for historical correlation is one order of magnitude greater than the ME of 0.0071 for implied correlation.

C. Explaining Returns Using Options

We use four option-based quantities to explain returns in the cross section; the first one is option-implied volatility, the next is based on the risk premium for stochastic volatility, and the last two are based on option-implied skewness. We first describe each of these quantities and then test empirically if these characteristics have significant power to explain the cross section of returns in our samples.

¹³An alternative approach is to compute homogeneous implied correlations, where all pairwise correlations are assumed to be the same: $\widehat{\rho}_{ij,t}^P = \widehat{\rho}_t^P, \forall i \neq j$. This is the approach used in Driessen, Maenhout, and Vilkov (2009). We also consider this approach, but the portfolios constructed using this approach perform worse than those when correlations are allowed to vary across assets.

The first option-based characteristic we use is option-implied volatility. Ang, Bali, and Cakici (2010) show that stocks with high current levels of option-implied volatility earn, in the next periods, higher returns than stocks with low levels of implied volatility. To maximize the information content of the option-implied volatility proxy, we use the MFIV described in Section III.A.

The second option-based characteristic we use is the variance risk premium, which is defined as the difference between risk-neutral (implied) and objective (expected or realized) variances. Previous research (see the papers cited in footnote 4) documents a positive relation between the variance risk premium and future stock returns. We use the implied-realized volatility spread (IRVS) as a measure of the volatility risk premium. We compute IRVS using the approach in Bali and Hovakimian (2009) as the spread between the Black-Scholes (1973) implied volatility averaged across call and put options and the realized stock-return volatility for the past month (21 trading days).

The third characteristic we consider is option-implied skewness, for which we use two measures. The first, MFIS, as defined in equation (2), represents a nonparametric estimate of the risk-neutral stock-return skewness.¹⁴ Rehman and Vilkov (2009) find that stocks with high option-implied skewness outperform stocks with low option-implied skewness.¹⁵ The second measure of skewness we consider is the spread between the Black-Scholes (1973) implied volatility for pairs of calls and puts, which is studied in Bali and Hovakimian (2009) and Cremers and Weinbaum (2010). We follow the methodology of Bali and Hovakimian to compute the call-put volatility spread (CPVS) as the difference between the current Black-Scholes implied volatilities of the 1-month at-the-money call and put options.

In order to evaluate if each of these four option-implied measures is useful for explaining the cross section of returns for our samples, we examine the returns of long-short decile portfolios for each characteristic separately. The long-short strategies are rebalanced daily based on the characteristic value at the end of a day, and each portfolio is held for the particular holding period we are considering (1 day, 1 week, or 1 fortnight). In Table 1 we show the annualized returns for each portfolio, along with the p -values, based on the Newey and West (1987) standard errors with a lag equal to the number of overlapping portfolio returns for each holding period. For completeness, we also include standard characteristics such as size (SIZE), book-to-market (BTM), and 12-month momentum (MOM).

¹⁴For the relation between expected stock returns and skewness measured directly, as opposed to option-implied skewness, see Rubinstein (1973), Kraus and Litzenberger (1976), Harvey and Siddique (2000), and Boyer, Mitton, and Vorkink (2010). For a study of asset allocation that takes into account time variations in risk premia, volatility, correlations, skewness, kurtosis, co-skewness, and co-kurtosis measured directly from stock returns, see Guidolin and Timmermann (2008).

¹⁵Some researchers (e.g., Xing et al. (2010)) use as a simple measure of skewness the difference between the implied volatilities for out-of-the-money put and at-the-money call options. However, that measure does not take into account the whole distribution, but rather just the left tail. Moreover, it is based on only two options and, hence, may be less informative than implied skewness, measured using the entire range of out-of-the money options. Rehman and Vilkov (2009) find that risk-neutral skewness contains information about future stock returns above and beyond that contained in the simple measure of skewness.

Table 1 confirms that the Fama-French characteristics (SIZE and BTM) explain returns in the expected direction, while momentum (MOM) is not significant. More interestingly, most option-based characteristics lead to significant returns on the long-short decile portfolios. The strongest results, in terms of the magnitude of returns, persistence across holding periods, and significance of returns, are for the portfolios based on the two measures of implied skewness, model-free implied skewness (MFIS) and the call-put implied volatility spread (CPVS), and for model-free implied volatility (MFIV); as expected, high-decile stocks outperform the low-decile ones for these measures. The IRVS is also positively and significantly related to returns, but at the 10% level.

TABLE 1
Return Predictability

In Table 1, we report the results of various stock characteristics: size (SIZE), book-to-market (BTM), momentum (MOM), model-free implied volatility (MFIV), implied-realized-volatility spread (IRVS), model-free implied skewness (MFIS), and call-put volatility spread (CPVS) to explain the cross section of returns. For each sample, on a daily basis, we sort the stocks by a particular characteristic, form the long-short decile portfolio, and hold this portfolio for a particular holding period (1 day, 1 week, or 1 fortnight). Below, we show the annualized mean holding return for each decile-based portfolio and in the parentheses the p -value for the hypothesis that the mean return is not different from 0. The p -values are based on the Newey and West (1987) autocorrelation-adjusted standard errors with the lag equal to the number of overlapping periods in portfolio holding.

Characteristic	Sample 1			Sample 2		
	1 Day	1 Week	1 Fortnight	1 Day	1 Week	1 Fortnight
SIZE	-0.1705 (0.00)	-0.1729 (0.00)	-0.1623 (0.00)	-0.1543 (0.00)	-0.1543 (0.00)	-0.1511 (0.00)
BTM	0.2050 (0.00)	0.1983 (0.00)	0.1853 (0.00)	0.1463 (0.00)	0.1365 (0.00)	0.1337 (0.00)
MOM	0.0026 (0.49)	-0.0148 (0.42)	-0.0325 (0.34)	-0.0034 (0.48)	-0.0060 (0.46)	-0.0160 (0.41)
MFIV	0.3055 (0.00)	0.2674 (0.00)	0.2243 (0.00)	0.1694 (0.07)	0.1497 (0.06)	0.1174 (0.09)
IRVS	0.1477 (0.01)	0.0959 (0.01)	0.0610 (0.06)	0.1535 (0.01)	0.0911 (0.01)	0.0479 (0.09)
MFIS	0.3721 (0.00)	0.1651 (0.00)	0.1465 (0.00)	0.4640 (0.00)	0.2337 (0.00)	0.1798 (0.00)
CPVS	0.7581 (0.00)	0.1899 (0.00)	0.1061 (0.00)	0.8699 (0.00)	0.2477 (0.00)	0.1421 (0.00)

Looking across the three rebalancing periods, we see that the magnitude of the returns for the option-implied characteristic portfolios decreases as the holding period increases. However, the rate of change of return with the holding period is not the same across characteristics: For example, in Sample 2, for daily rebalancing, it is the CPVS portfolio that earns the highest return of 86.99% per annum (p.a.), while for the fortnightly holding period, it is the MFIS portfolio that delivers the highest return of 17.98% p.a. Thus, the various characteristics will lead to different out-of-sample portfolio performance for the daily, weekly, and fortnightly rebalancing periods that we consider.

IV. Portfolio Construction and Performance Metrics

In this section, we explain the construction of the various portfolios we consider and also the metrics used to compare the performance of the benchmark

portfolios with that of portfolios based on option-implied information. For robustness, we consider several benchmark portfolios that do *not* rely on option-implied information.

A. Equal-Weighted Portfolio

For the *equal-weighted* ($1/N$) portfolio, in each period one allocates an equal amount of wealth across all N available stocks. The reason for considering this portfolio is that DeMiguel, Garlappi, and Uppal (2009) and Jacobs et al. (2010) show that it performs quite well even though it does not rely on any optimization; for example, the Sharpe ratio of the $1/N$ portfolio is more than double that of the S&P 500 over our sample period.

B. Minimum-Variance Portfolios

In this section, we study minimum-variance portfolios. We start by describing the mean-variance problem and then explain that this reduces to the minimum-variance problem if mean returns on all the assets are assumed to be equal. Then, in Section IV.C we consider the case where the mean returns on all assets are not assumed to be the same.

The *mean-variance* optimization problem can be written as

$$(6) \quad \min_w \quad w^\top \hat{\Sigma} w - w^\top \hat{\mu},$$

$$(7) \quad \text{s.t.} \quad w^\top e = 1,$$

where $w \in \mathbb{R}^N$ is the vector of portfolio weights invested in stocks, $\hat{\Sigma} \in \mathbb{R}^{N \times N}$ is the estimated covariance matrix, $\hat{\mu} \in \mathbb{R}^N$ is the estimated vector of expected returns, and $e \in \mathbb{R}^N$ is the vector of ones. The objective in equation (6) is to minimize the difference between the variance of the portfolio return, $w^\top \hat{\Sigma} w$, and its mean, $w^\top \hat{\mu}$. The constraint $w^\top e = 1$ in equation (7) ensures that the portfolio weights for the risky assets sum to 1; we consider the case without the risk-free asset because our objective is to explore how to use option-implied information to select the portfolio of only risky stocks.

In light of our discussion in the Introduction about the difficulty in forecasting expected returns, when we are studying the benefits of using option-implied second moments we assume that the expected return for each asset is equal to the grand mean return across all assets. In this case, the mean-variance portfolio problem in equation (6) reduces to finding the portfolio that minimizes the variance of the portfolio return, subject to the constraints that the portfolio weights sum to 1. The solution to the resulting *minimum-variance* portfolio problem is

$$(8) \quad w_{\min} = \frac{\hat{\Sigma}^{-1} e}{e^\top \hat{\Sigma}^{-1} e}.$$

The covariance matrix $\hat{\Sigma}$ in equation (8) can be decomposed into volatility and correlation matrices,

$$(9) \quad \hat{\Sigma} = \text{diag}(\hat{\sigma}) \hat{\Omega} \text{diag}(\hat{\sigma}),$$

where $\text{diag}(\hat{\sigma})$ is the diagonal matrix with volatilities of the stocks on the diagonal, and $\hat{\Omega}$ is the correlation matrix. Thus, to obtain the optimal portfolio weights

in equation (8) based on the *sample covariance* matrix, two quantities need to be estimated: volatilities ($\hat{\sigma}$) and correlations ($\hat{\Omega}$).

In the existing literature, several methods have been proposed to improve the out-of-sample performance of the minimum-variance portfolio based on sample (co)variances. We consider four approaches. The first is to impose constraints on the portfolio weights, which Jagannathan and Ma (2003) show can lead to substantial gains in performance. Thus, our next benchmark is the *constrained* portfolio, where we compute the short-sale-constrained minimum-variance portfolio weights.

The second approach we consider is the *shrinkage* portfolio, where we compute the minimum-variance portfolio weights *after* shrinking the covariance matrix. First, the sample covariance matrix for daily data is computed using the same approach that is described above. Then, to shrink the covariance matrix for daily returns, we use the approach in Ledoit and Wolf (2004a), (2004b), which uses as the covariance matrix a weighted average of the sample covariance matrix and a low-variance estimator of the covariance matrix (we use the identity matrix), with the weights assigned to these two covariance matrices determined optimally.

We also consider two other methods proposed in the literature for improving the behavior of the covariance matrix (see Elton et al. (2006) and the references therein). The first relies on setting all correlations equal to 0 so that the covariance matrix contains only estimates of variances. The second relies on setting the correlations equal to the mean of the estimated correlations; we do not report the performance of portfolios based on the second method, because they perform worse in terms of all three performance metrics when compared to portfolios obtained from the first method.

C. Mean-Variance Portfolios

In the previous section, we assumed that all assets had the same expected return. However, the recent literature on empirical asset pricing (see, e.g., the papers cited in footnote 4) finds that quantities that can be inferred from option prices such as the volatility risk premium and option-implied skewness are useful for predicting returns on stocks. But, it is well known in the literature (see, e.g., DeMiguel, Garlappi, and Uppal (2009), Jacobs et al. (2010)) that the weights of the traditional mean-variance portfolio are very sensitive to errors in estimates of expected returns, and that these portfolios perform poorly out of sample. Therefore, when evaluating the benefits of using option-based characteristics to form mean-variance portfolios, we use two alternative approaches. In the first approach, we use option-based characteristics to rank stocks and then adjust the mean returns of the top and bottom decile portfolios by a constant factor. In the second approach, we use the parametric-portfolio methodology of Brandt et al. (2009), which can be interpreted as a method where the adjustment of stock returns is done in an optimal fashion. We now describe these two approaches.

1. Mean-Variance Portfolios with Characteristic-Adjusted Returns

In the first approach, we assume that the conditional expected return $\mathbb{E}_t[r_{i,t+1}] = \mu_{i,t}$ of stock i at time t can be written as a function of the stock

characteristics $k = \{1, \dots, K\}$. More precisely, we specify that

$$\mu_{i,t} = \mu_{\text{BENCH},t} \left(1 + \sum_{k=1}^K \delta_{k,t} x_{ik,t} \right),$$

where $\mu_{\text{BENCH},t}$ is the expected benchmark return at t (we choose the benchmark return to be the grand mean return across all stocks), the value of $x_{ik,t}$ depends on the sorting index of stock i with respect to characteristic k at t , and the parameter $\delta_{k,t}$ denotes the intensity at t of the effect of the characteristic k on the conditional mean. In our analysis, we adjust the mean returns for only the stocks in the top and bottom deciles. That is, in our empirical exercise, with the characteristic defined so that it is positively related to returns, we set $x_{ik,t}$ equal to -1 if the stock is in the bottom decile in the cross section of all companies at date t , to $+1$ if it is located in the top decile, and 0 otherwise. Moreover, to isolate the effect of each option-implied characteristic, we consider each characteristic individually; that is, we set the mean return for each asset to be

$$\mu_{i,t} = \mu_{\text{BENCH},t} (1 + \delta_{k,t} x_{ik,t}).$$

In our empirical analysis, we report the results for the intensity $\delta_{k,t} = \delta_k = 0.10$.

2. Mean-Variance Parametric Portfolios

In the second approach, we apply the parametric-portfolio methodology of Brandt et al. (2009) by using the MFIV, MFIS, CPVS, and IRVS, in addition to the traditional stock characteristics (size, value, and momentum), to construct parametric portfolios based on mean-variance utility.

The parametric-portfolio methodology has been developed to deal with the problem of poor out-of-sample performance of portfolios because of estimation error. In the parametric portfolios, the weight of an asset is a linear function of its weight in the benchmark portfolio and the value of characteristics:

$$\omega_{i,t} = \omega_{i,t}^{1/N} + \sum_{k=1}^K \theta_{k,t} x_{ik,t},$$

where $\omega_{i,t}^{1/N}$ is the weight of the asset i in the equal-weighted benchmark portfolio at t , $\theta_{k,t}$ is the loading on characteristic k at t , and $x_{ik,t}$ is the value of characteristic k for stock i at t .¹⁶ Following Brandt et al. (2009), we normalize the characteristics to have zero mean and unit variance. Note that $\theta_{k,t}$ is *not* asset-specific, but it is the same for all assets in the portfolio. We choose the vector $\theta_t = (\theta_{1,t}, \theta_{2,t}, \dots)$ optimally by maximizing the average daily mean-variance utility using a rolling window procedure with an estimation window of 250 days. Because it is difficult to short stocks, we constrain short sales; that is, we choose the loadings θ_t such that $\omega_{i,t} \geq 0$.

To determine the parametric portfolios, we start with the same characteristics as the ones in Brandt et al. (2009), but using the $1/N$ portfolio as the benchmark:

¹⁶In addition to using the $1/N$ portfolio as the benchmark, we also consider the value-weighted portfolio, and the findings are similar with this benchmark portfolio.

That is, $1/N + \text{FFM}$, where FFM denotes the size and value characteristics identified in Fama and French (1992) and the momentum characteristic identified in Jegadeesh and Titman (1993). Then, to study the effect of option-implied information, we first consider the effect of *replacing* the FFM characteristics with the following option-implied characteristics: MFIV, IRVS, MFIS, and CPVS. Second, in order to study the *incremental* value of option-implied information over and above the FFM characteristics, we also consider the effect of including these option-implied characteristics *in addition* to the FFM factors.

Using a variety of metrics that are described next, the out-of-sample performance of the benchmark portfolios described in Sections IV.A–IV.C (reported in Tables 2 and 3) and discussed in Section V.A.

D. Portfolio-Performance Metrics

We evaluate performance of the various portfolios using three criteria. These are the i) out-of-sample portfolio volatility (standard deviation); ii) out-of-sample portfolio Sharpe ratio;¹⁷ and, iii) portfolio turnover (trading volume).

We consider three rebalancing intervals: daily, weekly, and fortnightly. Typically, for weekly and fortnightly holding periods, one would form the optimal portfolio on a particular day and then compute the return from holding that portfolio for a week or a fortnight by multiplying the optimal weights on that particular day by the cumulative returns of each asset over the following week or fortnight. One concern when doing this is that the performance of the portfolios for the weekly and fortnightly holding periods would depend on the particular date chosen for forming the portfolio. In order to address this concern for the portfolios with weekly and fortnightly rebalancing, we find a new set of weights *daily* and then hold that portfolio for a week or fortnight. Thus, we have a series of overlapping portfolio returns. To compute the annualized performance metrics, we multiply the overlapping portfolio returns by the number of rebalancing periods in a year; that is, we multiply by the ratio of 251 to Δt , where for weekly rebalancing $\Delta t = 5$, and for fortnightly rebalancing $\Delta t = 10$.

We use the “rolling-horizon” procedure for computing the portfolio weights and evaluating their performance, with the estimation-window length for daily data being $\tau = 250$ days for Sample 1 and $\tau = 750$ days for Sample 2. Holding the portfolio w_t^{STRATEGY} for the period Δt gives the *out-of-sample* return at time $t + \Delta t$: That is, $r_{t+\Delta t}^{\text{STRATEGY}} = (w_t^{\text{STRATEGY}})^{\top} r_{t+\Delta t}$, where $r_{t+\Delta t}$ denotes the returns from t to $t + \Delta t$, and Δt is 1 day, 1 week, or 1 fortnight. After collecting the time series of $T - \tau - \Delta t$ returns, r_t^{STRATEGY} , the annualized out-of-sample mean, volatility ($\hat{\sigma}$), and Sharpe ratio of returns (SR) are, respectively,

$$\hat{\mu}^{\text{STRATEGY}} = \left(\frac{1}{T - \tau - \Delta t - 1} \right) \left(\frac{251}{\Delta t} \right) \sum_{t=\tau}^{T-\Delta t} r_{t+\Delta t}^{\text{STRATEGY}},$$

¹⁷We also compute the certainty equivalent return of an investor with power utility in order to evaluate the effect of higher moments; we find that the insights using this measure are the same as those from using the Sharpe ratio.

$$\widehat{\sigma}^{\text{STRATEGY}} = \left[\left(\frac{1}{T - \tau - \Delta t - 1} \right) \left(\frac{251}{\Delta t} \right) \times \sum_{t=\tau}^{T-\Delta t} (r_{t+\Delta t}^{\text{STRATEGY}} - \widehat{\mu}^{\text{STRATEGY}})^2 \right]^{1/2},$$

$$\widehat{\text{SR}}^{\text{STRATEGY}} = \frac{\widehat{\mu}^{\text{STRATEGY}}}{\widehat{\sigma}^{\text{STRATEGY}}}.$$

To measure the statistical significance of the difference in the volatility and Sharpe ratio of a particular portfolio from that of another portfolio that serves as a benchmark, we also report the p -values for these differences. For calculating the p -values for the case of daily rebalancing, we use the bootstrapping methodology described in Efron and Tibshirani (1993), and for weekly or fortnightly rebalancing we make an additional adjustment, as in Politis and Romano (1994), to account for the autocorrelation arising from overlapping returns.¹⁸

Finally, we wish to obtain a measure of portfolio turnover per holding period. Let $w_{j,t}^{\text{STRATEGY}}$ denote the portfolio weight in stock j chosen at time t for a particular strategy, $w_{j,t+\Delta t}^{\text{STRATEGY}}$ the portfolio weight *before* rebalancing but at $t + \Delta t$, and $w_{j,t+\Delta t}^{\text{STRATEGY}}$ the desired portfolio weight at time $t + \Delta t$ (after rebalancing). Then, turnover, which is the average percentage of wealth traded per rebalancing interval (daily, weekly, or fortnightly), is defined as the sum of the absolute value of the rebalancing trades across the N available stocks and over the $T - \tau - \Delta t$ trading dates, normalized by the total number of trading dates, and, because our portfolios for weekly and fortnightly rebalancing periods are created each day and hence are overlapping, normalized further by the number of overlapping periods:

$$\text{TURNOVER} = \left(\frac{1}{T - \tau - \Delta t} \right) \left(\frac{1}{\Delta t} \right) \sum_{t=\tau}^{T-\Delta t} \sum_{j=1}^N \left(|w_{j,t+\Delta t}^{\text{STRATEGY}} - w_{j,t}^{\text{STRATEGY}}| \right).$$

The strategies that rely on forecasts of expected returns based on option-implied characteristics have much higher turnover compared to the benchmark strategies. In order to understand whether or not the option-based strategies would outperform the benchmarks even after adjusting for transaction costs, we also

¹⁸Specifically, consider two portfolios i and n , with $\mu_i, \mu_n, \sigma_i, \sigma_n$ as their true means and volatilities. We wish to test the hypothesis that the Sharpe ratio (or certainty-equivalent return) of portfolio i is worse (smaller) than that of the benchmark portfolio n , that is, $H_0: \mu_i/\sigma_i - \mu_n/\sigma_n \leq 0$. To do this, we obtain B pairs of size $T - \tau$ of the portfolio returns i and n by simple resampling with replacement for daily returns, and by blockwise resampling with replacement for overlapping weekly and fortnightly returns. We choose $B = 10,000$ for both cases and the block size equal to the number of overlaps in a series, that is, 4 for weekly and 9 for fortnightly data. If \hat{F} denotes the empirical distribution function of the B bootstrap pairs corresponding to $\widehat{\mu}_i/\widehat{\sigma}_i - \widehat{\mu}_n/\widehat{\sigma}_n$, then a one-sided p -value for the previous null hypothesis is given by $\hat{p} = \hat{F}(0)$, and we will reject it for a small \hat{p} . In a similar way, to test the hypothesis that the variance of portfolio i is greater (worse) than the variance of the benchmark portfolio n , $H_0: \sigma_i^2/\sigma_n^2 \geq 1$, if \hat{F} denotes the empirical distribution function of the B bootstrap pairs corresponding to $\widehat{\sigma}_i^2/\widehat{\sigma}_n^2$, then a one-sided p -value for this null hypothesis is given by $\hat{p} = 1 - \hat{F}(1)$, and we reject the null for a small \hat{p} . For a nice discussion of the application of other bootstrapping methods to tests of differences in portfolio performance, see Ledoit and Wolf (2008).

compute the *equivalent transaction cost*: that is, the transaction cost level in basis points that equates the particular performance metric (mean return or Sharpe ratio) of a given strategy with that for the benchmark strategy. To find this equivalent transaction cost, we adopt the following approach, which is similar to that in Grundy and Martin (2001). First, for each level of transaction cost, we compute the time series of net returns $\tilde{r}_{t+\Delta t}^{\text{STRATEGY}}$ for a given strategy and the benchmark:

$$\tilde{r}_{t+\Delta t}^{\text{STRATEGY}} = r_{t+\Delta t}^{\text{STRATEGY}} - \sum_{j=1}^N \left| w_{j,t+\Delta t}^{\text{STRATEGY}} - w_{j,t}^{\text{STRATEGY}} \right| \times \text{TC}^{\text{STRATEGY,BENCHMARK}}.$$

Then, we compute the performance metrics using these returns. Finally, we search for the level of transaction costs that makes the performance metric the same for the strategy being evaluated and the appropriate benchmark strategy.

V. Out-of-Sample Performance

In this section, we discuss the major empirical findings of our paper about the ability of forward-looking information implied in option prices to improve the out-of-sample performance of stock portfolios. We start, in Section V.A, by discussing the performance of the benchmark portfolios that do not use information from option prices. In Section V.B, we report the performance of portfolios obtained using option-implied volatilities. In Section V.C, we report the performance of portfolios that use option-implied correlations. Finally, in Section V.D, we report the improvement in out-of-sample portfolio performance from using option-implied quantities that explain the cross section of returns, such as the variance risk premium and option-implied skewness. In each of these sections, we use option-implied information about only one moment at a time (volatility, correlation, or expected return) in order to isolate the magnitude of the gains from using option-implied information to estimate that particular moment.

A. Performance of Benchmark Portfolios

In Tables 2 and 3 we report the performance of several benchmark strategies, all of which do *not* use data on option prices. Table 2 gives the performance of *minimum-variance* portfolios, and Table 3 gives the performance of *mean-variance* portfolios; both tables also report the performance of the $1/N$ portfolio. In Panel A of each table, we report the results for daily rebalancing, in Panel B for weekly rebalancing, and in Panel C for the case in which the portfolio is held for 1 fortnight. We report three performance metrics in the table: the volatility (STD) of portfolio returns, the Sharpe ratio (SR), and the turnover (TRN) of the portfolio. The p -value for the comparison with the $1/N$ benchmark is reported in parentheses under each performance metric. The p -value is for the *one-sided* null hypothesis that the portfolio being evaluated is *no better* than the $1/N$ benchmark for a given performance metric (so a small p -value suggests *rejecting* the null hypothesis that the portfolio being evaluated is no better than the benchmark).

Table 2 reports the performance of four variants of the minimum-variance benchmark portfolio: “Sample cov,” “Constrained,” “Shrinkage,” and “Zero correlation.” We see from this table that, compared to the $1/N$ portfolio, all of the

TABLE 2
Minimum-Variance Portfolios without Option-Implied Information

In Table 2, we report the performance of the $1/N$ portfolio and various minimum-variance benchmark portfolios that are based on historical returns and do *not* rely on prices of options. The “Sample cov” portfolio is the minimum-variance portfolio based on the sample covariance matrix; “Constrained” is the minimum-variance portfolio based on the sample covariance matrix but with short sales constrained; “Shrinkage” is the minimum-variance portfolio where shrinkage has been applied to the sample covariance matrix using the Ledoit and Wolf (2004a), (2004b) methodology; and “Zero correlation” is the minimum-variance portfolio where all correlations are set equal to 0. We report p -values in parentheses with respect to the $1/N$ portfolio, with the null hypothesis being that the portfolio being evaluated is no better than the benchmark (so a small p -value suggests *rejecting* the null hypothesis that the portfolio being evaluated is no better than the benchmark).

Strategy	Sample 1			Sample 2		
	STD	SR	TRN	STD	SR	TRN
<i>Panel A. Daily Rebalancing</i>						
$1/N$	0.2083	0.5903	0.0134	0.2254	0.4963	0.0145
Sample cov	0.1636 (0.00)	0.2632 (0.84)	0.6131	0.1429 (0.00)	0.2094 (0.82)	0.5130
Constrained	0.1367 (0.00)	0.4489 (0.80)	0.0552	0.1383 (0.00)	0.6558 (0.15)	0.0303
Shrinkage	0.1332 (0.00)	0.4943 (0.63)	0.2723	0.1263 (0.00)	0.2939 (0.74)	0.3046
Zero correlation	0.1771 (0.00)	0.5512 (0.77)	0.0133	0.1899 (0.00)	0.5161 (0.35)	0.0125
<i>Panel B. Weekly Rebalancing</i>						
$1/N$	0.1973	0.6045	0.0302	0.2127	0.5030	0.0325
Sample cov	0.1613 (0.00)	0.2893 (0.90)	0.6556	0.1475 (0.00)	0.2535 (0.84)	0.5961
Constrained	0.1305 (0.00)	0.4591 (0.83)	0.0630	0.1322 (0.00)	0.6768 (0.08)	0.0408
Shrinkage	0.1343 (0.00)	0.5257 (0.64)	0.3047	0.1326 (0.00)	0.3196 (0.77)	0.3724
Zero correlation	0.1663 (0.00)	0.5697 (0.76)	0.0268	0.1782 (0.00)	0.5331 (0.23)	0.0276
<i>Panel C. Fortnightly Rebalancing</i>						
$1/N$	0.1905	0.6081	0.0427	0.2055	0.5036	0.0462
Sample cov	0.1587 (0.00)	0.2439 (0.93)	0.6961	0.1506 (0.00)	0.2690 (0.83)	0.6678
Constrained	0.1262 (0.00)	0.4542 (0.85)	0.0696	0.1297 (0.00)	0.6726 (0.11)	0.0492
Shrinkage	0.1333 (0.00)	0.4910 (0.70)	0.3340	0.1372 (0.00)	0.3262 (0.77)	0.4303
Zero correlation	0.1599 (0.00)	0.5778 (0.71)	0.0372	0.1719 (0.00)	0.5388 (0.21)	0.0390

strategies based on the minimum-variance portfolio achieve significantly lower volatility ($\hat{\sigma}$) out of sample. For example, in Panel A with results for “Daily rebalancing,” we see that for Sample 2 with daily rebalancing, the volatility of the $1/N$ portfolio is 0.2254, for the minimum-variance portfolio based on the sample-covariance matrix it is 0.1429, for the minimum-variance portfolio with short-sale constraints it is 0.1383, for the minimum-variance portfolio with shrinkage it is 0.1263, and for the portfolio obtained from setting all correlations equal to 0 it is 0.1899. The p -values indicate that the volatilities of the minimum-variance portfolios are significantly lower than that of $1/N$. The results are similar for Sample 1 and in Panels B and C for “Weekly Rebalancing” and “Fortnightly Rebalancing.”

However, the Sharpe ratio (SR) and turnover (TRN) are typically better for the $1/N$ portfolio compared to the four minimum-variance portfolios, with the

TABLE 3
Mean-Variance Portfolios without Option-Implied Information

In Table 3, we report the performance of the $1/N$ portfolio and various mean-variance portfolios that are based on historical returns and do *not* rely on prices of options. The "Sample cov" portfolio is the mean-variance portfolio based on the sample covariance matrix; "Constrained" is the mean-variance portfolio based on the sample covariance matrix but with short sales constrained; "Shrinkage" is the mean-variance portfolio where shrinkage has been applied to the sample covariance matrix using the Ledoit and Wolf (2004a), (2004b) methodology; "Zero correlation" is the mean-variance portfolio where all correlations are set equal to 0; and " $1/N + \text{FFM}$ " denotes the parametric benchmark portfolio, where we start with the " $1/N$ " initial portfolio and adjust it optimally using the Fama-French and momentum characteristics. We report p -values in parentheses with respect to the $1/N$ portfolio, with the null hypothesis being that the portfolio being evaluated is no better than the benchmark (so a small p -value suggests *rejecting* the null hypothesis that the portfolio being evaluated is no better than the benchmark).

Strategy	Sample 1			Sample 2		
	STD	SR	TRN	STD	SR	TRN
<i>Panel A. Daily Rebalancing</i>						
$1/N$	0.2083	0.5903	0.0134	0.2254	0.4963	0.0145
Sample cov	0.4204 (1.00)	-0.2277 (0.99)	3.7940	0.6481 (1.00)	-0.9411 (1.00)	4.6965
Constrained	0.1370 (0.00)	0.3553 (0.92)	0.0624	0.1377 (0.00)	0.6083 (0.23)	0.0321
Shrinkage	0.6747 (1.00)	-0.1630 (0.99)	3.9762	0.9947 (1.00)	-0.6261 (1.00)	6.9215
Zero correlation	0.5700 (1.00)	-0.0683 (0.98)	1.6410	0.3427 (1.00)	0.1929 (0.91)	0.3008
$1/N + \text{FFM}$	0.2079 (0.38)	0.6589 (0.06)	0.0515	0.2285 (1.00)	0.5453 (0.14)	0.0372
<i>Panel B. Weekly Rebalancing</i>						
$1/N$	0.1973	0.6045	0.0302	0.2127	0.5030	0.0325
Sample cov	0.4404 (1.00)	-0.0664 (0.99)	4.0024	0.7402 (1.00)	-0.6171 (1.00)	7.0693
Constrained	0.1301 (0.00)	0.4018 (0.92)	0.0703	0.1301 (0.00)	0.6513 (0.13)	0.0426
Shrinkage	0.8527 (1.00)	-0.0660 (0.99)	5.3350	1.1251 (1.00)	-0.4198 (1.00)	26.4889
Zero correlation	0.6571 (1.00)	0.0195 (0.99)	2.5679	0.3980 (1.00)	0.2036 (0.95)	0.4230
$1/N + \text{FFM}$	0.1979 (0.66)	0.6725 (0.05)	0.0624	0.2156 (0.94)	0.5617 (0.08)	0.0508
<i>Panel C. Fortnightly Rebalancing</i>						
$1/N$	0.1905	0.6081	0.0427	0.2055	0.5036	0.0462
Sample cov	0.4330 (1.00)	-0.0352 (0.99)	4.2037	0.8165 (1.00)	-0.5013 (1.00)	19.2794
Constrained	0.1250 (0.00)	0.4124 (0.90)	0.0769	0.1267 (0.00)	0.6582 (0.13)	0.0508
Shrinkage	0.8992 (1.00)	-0.0386 (0.98)	7.0368	1.1933 (1.00)	-0.3280 (0.99)	24.4343
Zero correlation	0.6616 (1.00)	0.1669 (0.96)	9.9765	0.4239 (1.00)	0.1704 (0.96)	0.3918
$1/N + \text{FFM}$	0.1903 (0.48)	0.6749 (0.06)	0.0719	0.2075 (0.77)	0.5657 (0.08)	0.0623

only exceptions being the Sharpe ratio for the constrained minimum-variance portfolio, and for the minimum-variance portfolio obtained by setting all correlations equal to 0, in the case of Sample 2; but, for both cases the differences are not statistically significant.¹⁹

¹⁹It might seem strange to evaluate the Sharpe ratio of minimum-variance portfolios, whose objective is to only minimize the volatility of the portfolio. This comparison is motivated by the statement in Jagannathan and Ma ((2003), p. 1653) that "the global minimum variance portfolio has

Of the four minimum-variance portfolios that we consider, the short-sale-constrained portfolio and the portfolio obtained by setting all correlations equal to 0 have turnover that is comparable to that of the $1/N$ portfolio and substantially lower than the turnover of the unconstrained “Sample cov” portfolio. This is true also in the tables that follow, where we use option-implied information.

Table 3 reports the performance of four mean-variance portfolios, “Sample cov,” “Constrained,” “Shrinkage,” and “Zero correlation,” and the mean-variance portfolio implemented using the parametric-portfolio methodology with the Fama and French characteristics along with momentum, which in the table is labeled “ $1/N + \text{FFM}$.” The three mean-variance portfolios that do *not* have constraints on short selling (“Sample cov,” “Shrinkage,” and “Zero correlation”) perform very poorly along all metrics. The mean-variance strategy with short-sale constraints achieves a lower volatility than the $1/N$ portfolio, but it has a lower Sharpe ratio and higher turnover than the $1/N$ portfolio and also the short-sale-constrained minimum-variance portfolio considered in Table 2. The parametric portfolio usually has the best Sharpe ratio compared to the $1/N$ portfolio and the other mean-variance portfolios (though the difference is not always statistically significant), with a turnover that is comparable to that of $1/N$. The volatility of the parametric portfolio is higher than that of the short-sale-constrained mean-variance portfolio and the minimum-variance portfolios considered in Table 2, which is not surprising, given that this portfolio is not designed with the objective of minimizing volatility.

B. Performance of Portfolios Using Option-Implied Volatility

Motivated by the findings in Section V.A about the predictive power of model-free implied volatilities after correction for the risk premium, $\widehat{R\bar{V}}$, we use them in $\text{diag}(\hat{\sigma})$ to obtain the covariance matrix given in equation (9); that is, $\hat{\Sigma} = \text{diag}(\widehat{R\bar{V}}) \hat{\Omega} \text{diag}(\widehat{R\bar{V}})$. Using this covariance matrix, and setting the expected return on each asset equal to the grand mean across all stocks, we determine the minimum-variance portfolio in equation (8), along with the portfolios where short sales are constrained, where shrinkage is applied to this covariance matrix, and where we impose the restriction that all correlations are equal to 0. In computing these portfolios, we continue to use historical correlations (except for the last portfolio, where correlations are set equal to 0).

The results for the minimum-variance portfolios based on risk-premium-corrected option-implied volatility are given in Table 4. In this table we report two sets of p -values: the first with respect to the $1/N$ portfolio, and the second with respect to the corresponding benchmark portfolio in Table 2. For Sample 2, comparing the volatility (STD) of portfolio returns across the different portfolio strategies, we see that the “Shrinkage” portfolio always achieves the lowest

as large an out-of-sample Sharpe ratio as other efficient portfolios when past historical average returns are used as proxies for expected returns.” DeMiguel, Garlappi, and Uppal (2009) also find that the minimum-variance portfolio performs surprisingly well in terms of Sharpe ratio when compared to other portfolios that rely on estimates of expected returns.

TABLE 4
Minimum-Variance Portfolios Using Option-Implied Volatility

In Table 4, we report the performance of the $1/N$ portfolio and various minimum-variance portfolios that use the risk-premium-corrected model-free implied volatility calculated from option prices, while correlations are estimated from historical data. The "Sample cov" portfolio is the minimum-variance portfolio based on the sample covariance matrix, but where historical volatility is replaced by option-implied volatility corrected for the volatility risk premium; "Constrained" is the minimum-variance portfolio based on the same covariance matrix as for "Sample-cov" but with short sales constrained; "Shrinkage" is the minimum-variance portfolio based on the same covariance matrix as for "Sample-cov" but with shrinkage applied to the "Sample-cov" matrix; and "Zero correlation" is the minimum-variance portfolio based on a covariance matrix where all correlations are set equal to 0 and historical volatility is replaced by option-implied volatility corrected for the volatility risk premium. We report two p -values in parentheses, the first with respect to the $1/N$ portfolio, and the second with respect to the corresponding minimum-variance benchmark portfolio in Table 2, with the null hypothesis being that the portfolio being evaluated is no better than the benchmark (so a small p -value suggests *rejecting* the null hypothesis that the portfolio being evaluated is no better than the benchmark).

Strategy	Sample 1			Sample 2		
	STD	SR	TRN	STD	SR	TRN
<i>Panel A. Daily Rebalancing</i>						
$1/N$	0.2083	0.5903	0.0134	0.2254	0.4963	0.0145
Sample cov	0.1554 (0.00) (0.00)	0.2037 (0.91) (0.67)	1.1646	0.1212 (0.00) (0.00)	0.4963 (0.51) (0.10)	1.3598
Constrained	0.1351 (0.00) (0.11)	0.3427 (0.94) (0.94)	0.1960	0.1342 (0.00) (0.00)	0.4302 (0.68) (1.00)	0.1911
Shrinkage	0.1320 (0.00) (0.29)	0.2849 (0.88) (0.96)	0.6466	0.1197 (0.00) (0.02)	0.2828 (0.84) (0.53)	0.9467
Zero correlation	0.1728 (0.00) (0.00)	0.5018 (0.93) (1.00)	0.0551	0.1759 (0.00) (0.00)	0.4858 (0.57) (0.91)	0.0615
<i>Panel B. Weekly Rebalancing</i>						
$1/N$	0.1973	0.6045	0.0302	0.2127	0.5030	0.0325
Sample cov	0.1545 (0.00) (0.01)	0.2366 (0.95) (0.70)	1.1823	0.1189 (0.00) (0.00)	0.5714 (0.36) (0.03)	1.3774
Constrained	0.1295 (0.00) (0.21)	0.3783 (0.95) (0.93)	0.2014	0.1268 (0.00) (0.00)	0.5046 (0.49) (1.00)	0.1987
Shrinkage	0.1327 (0.00) (0.19)	0.3461 (0.90) (0.98)	0.6603	0.1177 (0.00) (0.00)	0.3678 (0.78) (0.37)	0.9629
Zero correlation	0.1624 (0.00) (0.00)	0.5240 (0.93) (1.00)	0.0615	0.1650 (0.00) (0.00)	0.5216 (0.36) (0.73)	0.0680
<i>Panel C. Fortnightly Rebalancing</i>						
$1/N$	0.1905	0.6081	0.0427	0.2055	0.5036	0.0462
Sample cov	0.1504 (0.00) (0.00)	0.1984 (0.96) (0.69)	1.2019	0.1207 (0.00) (0.00)	0.5890 (0.34) (0.02)	1.3957
Constrained	0.1263 (0.00) (0.54)	0.3687 (0.94) (0.96)	0.2062	0.1257 (0.00) (0.01)	0.5123 (0.47) (1.00)	0.2048
Shrinkage	0.1305 (0.00) (0.13)	0.3309 (0.92) (0.98)	0.6745	0.1192 (0.00) (0.00)	0.4044 (0.71) (0.29)	0.9799
Zero correlation	0.1561 (0.00) (0.00)	0.5394 (0.89) (1.00)	0.0678	0.1595 (0.00) (0.00)	0.5395 (0.24) (0.48)	0.0747

volatility, and this is significantly lower than that of the $1/N$ portfolio and also the “Shrinkage” benchmark strategy in Table 2, which uses historical volatility; however, the “Shrinkage” strategy has the lowest Sharpe ratio of all the strategies in Table 4, and also its turnover is quite high. For Sample 1, again the lowest volatility for daily rebalancing is achieved by the “Shrinkage” strategy, but for weekly and fortnightly rebalancing, it is the constrained strategy that has the lowest volatility. For both samples and all three rebalancing frequencies, of the four minimum-variance portfolios, it is the “Zero correlation” portfolio that achieves the lowest turnover and the highest Sharpe ratio; for Sample 2, this Sharpe ratio is higher than even that of the $1/N$ portfolio.²⁰

We conclude that volatility of stock returns estimated from risk-premium-corrected implied volatility is successful in achieving a significant reduction in portfolio volatility.²¹

C. Performance of Portfolios with Option-Implied Correlations

In this section, we investigate the gains from using option-implied correlations in portfolio selection. The performance of portfolios obtained from using the risk-premium-corrected option-implied correlations in equation (9), instead of historical correlations, is reported in Table 5. In order to isolate the effect of

TABLE 5
Minimum-Variance Portfolios Using Option-Implied Correlation

In Table 5, we report the performance of the $1/N$ portfolio and various minimum-variance portfolios that use risk-premium-corrected option-implied correlation without restricting correlations to be the same across asset pairs, as computed in Buss and Vilkov (2012), while volatilities are estimated from historical data. The “Sample cov” portfolio is the minimum-variance portfolio based on the sample covariance matrix, but with option-implied correlations that are not assumed to be the same across all asset pairs; “Constrained” is the minimum-variance portfolio based on the same covariance matrix as for “Sample-cov” but with short sales constrained; and “Regularization” is the minimum-variance portfolio based on the same covariance matrix as for “Sample-cov” but with regularization applied to the “Sample-cov” matrix using the Zumbach (2009) methodology. We report two p -values in parentheses, the first with respect to the $1/N$ portfolio, and the second with respect to the corresponding minimum-variance benchmark portfolio in Table 2, with the null hypothesis being that the portfolio being evaluated is no better than the benchmark (so a small p -value suggests *rejecting* the null hypothesis that the portfolio being evaluated is no better than the benchmark).

Strategy	Sample 1			Sample 2		
	STD	SR	TRN	STD	SR	TRN
<i>Panel A. Daily Rebalancing</i>						
$1/N$	0.2083	0.5903	0.0134	0.2254	0.4963	0.0145
Sample cov	0.2374 (1.00) (1.00)	0.2771 (0.84) (0.47)	3.6129	0.2204 (0.27) (1.00)	0.2838 (0.76) (0.37)	4.4839
Constrained	0.1549 (0.00) (1.00)	0.8235 (0.17) (0.03)	0.0843	0.1803 (0.00) (1.00)	0.5867 (0.36) (0.62)	0.0895
Regularization	0.1367 (0.00) (0.82)	0.6334 (0.43) (0.25)	0.2659	0.1506 (0.00) (1.00)	0.6321 (0.26) (0.07)	0.4240

(continued on next page)

²⁰The reason for the relatively poor Sharpe ratio and turnover of the other portfolios based on implied volatility is that implied volatility is a highly unstable estimator of future volatility; this instability increases the error in portfolio weights and reduces the gains from having a better predictor of RV.

²¹For the portfolio constructed using implied volatility *without* correcting for the risk premium, the volatility, Sharpe ratio, and turnover are worse than for the case where one uses the risk-premium-corrected implied volatility.

TABLE 5 (continued)
Minimum-Variance Portfolios Using Option-Implied Correlation

Strategy	Sample 1			Sample 2		
	STD	SR	TRN	STD	SR	TRN
<i>Panel B. Weekly Rebalancing</i>						
1/N	0.1973	0.6045	0.0302	0.2127	0.5030	0.0325
Sample cov	0.2307 (1.00) (1.00)	0.4974 (0.67) (0.16)	3.6233	0.2034 (0.16) (1.00)	0.3755 (0.71) (0.24)	4.5216
Constrained	0.1416 (0.00) (0.99)	0.7545 (0.22) (0.03)	0.0879	0.1651 (0.00) (1.00)	0.5673 (0.36) (0.71)	0.0945
Regularization	0.1321 (0.00) (0.28)	0.6429 (0.43) (0.22)	0.2756	0.1454 (0.00) (0.99)	0.6261 (0.23) (0.05)	0.4381
<i>Panel C. Fortnightly Rebalancing</i>						
1/N	0.1905	0.6081	0.0427	0.2055	0.5036	0.0462
Sample cov	0.2160 (0.98) (1.00)	0.4381 (0.77) (0.17)	3.6494	0.2117 (0.71) (1.00)	0.2819 (0.83) (0.47)	4.5829
Constrained	0.1347 (0.00) (0.94)	0.7294 (0.27) (0.05)	0.0909	0.1608 (0.00) (1.00)	0.5446 (0.40) (0.73)	0.0983
Regularization	0.1277 (0.00) (0.12)	0.6214 (0.47) (0.18)	0.2857	0.1431 (0.00) (0.83)	0.6051 (0.28) (0.05)	0.4522

using implied correlations, we use volatilities calculated from historical data when computing the portfolio weights.

We consider three minimum-variance portfolios in Table 5: The first is based on the sample-covariance matrix with option-implied correlations; the second is the same as the first, but with short sales constrained; and the third, which is labeled “regularization,” replaces the “shrinkage” portfolio considered in the earlier tables. We use the regularization approach of Zumbach (2009) because we are using option-implied correlations and, hence, do not know the distribution of returns for the resulting covariance matrix, which means that we cannot use the shrinkage results of Ledoit and Wolf (2004a), (2004b) that rely on particular distributional assumptions.

We observe from Table 5 that using the risk-premium-corrected implied correlations does not lead to much of an improvement in the out-of-sample performance of the minimum-variance portfolios.²² While the volatility of the portfolios with short-sale constraints and regularization is less than that of the 1/N for both Sample 1 and Sample 2, it exceeds that of the corresponding benchmark portfolios studied in Table 2. The portfolios based on option-implied correlations also have higher turnover. The only positive result is that the Sharpe ratio of the constrained portfolio is greater than that of the 1/N portfolio and also the corresponding benchmark portfolio in Table 2, though the improvement is not always statistically significant.

²²The performance of portfolios based on implied correlations *without* the correction for the risk premium is slightly worse.

Thus, we conclude that using the option-implied correlations does not lead to a significant improvement in portfolio performance. The reason for the poor performance of portfolios based on implied correlations is that the covariance matrix based on these correlations is highly unstable over time. Consequently, the resulting portfolio weights are highly variable and perform poorly out of sample.

D. Performance of Portfolios with Returns Predicted Using Options

In this section, we examine the effect on portfolio performance of using four option-implied quantities that explain the cross section of stock returns: MFIV; volatility risk premium, measured as the spread between the currently observed Black-Scholes (1973) option-implied volatility and realized (historical) volatility (IRVS); MFIS; and skewness, measured as the spread between the Black-Scholes implied volatilities for calls and for puts (CPVS). There are two ways in which we use these quantities to improve the performance of portfolios. In the first, described in Section V.D.1, we rank all stocks based on each characteristic and adjust the returns of the top and bottom decile portfolios. In the second, described in Section V.D.2, we use the parametric-portfolio methodology of Brandt et al. (2009).

1. Performance of a Mean-Variance Portfolio with Option Characteristics

The out-of-sample performance of the mean-variance portfolios that use option-implied characteristics to adjust returns is reported in Table 6. There are four portfolios, each with short-sale constraints, considered in this table corresponding to the following four option characteristics: MFIV, IRVS, MFIS, and CPVS. We compare performance of these portfolios to two sets of benchmarks: the $1/N$ portfolio and the constrained minimum-variance portfolio reported in Table 2, which does not rely on option prices.²³

From Table 6, we see that other than MFIV, the portfolios whose returns are adjusted based on any of the other three characteristics have a significantly higher Sharpe ratio than the $1/N$ portfolio and the benchmark portfolios in Table 3. The difference in Sharpe ratios is largest for daily rebalancing, and it declines as the rebalancing frequency decreases. For example, for Sample 1, the Sharpe ratio for the $1/N$ portfolio is 0.5903, while for the portfolio using IRVS it is 0.9232, for the portfolio using MFIS it is 1.0092, and for the portfolio using CPVS it is 1.4291. However, the improvement in the Sharpe ratio is accompanied by an increase in turnover.

To understand whether the mean-variance portfolios using option-implied characteristics to forecast expected returns outperform the benchmark portfolios even in the presence of transaction costs, we compute the *equivalent transaction cost* for each portfolio. Recall from Section IV.D that this is the transaction cost level that equates the performance metric (mean return or Sharpe ratio) of a given strategy with that for the benchmark strategy.

²³We use as a benchmark the short-sale-constrained minimum-variance portfolio rather than the constrained mean-variance portfolios because the constrained minimum-variance portfolio has better performance in terms of all three metrics.

TABLE 6
Mean-Variance Portfolios with Option-Characteristic-Adjusted Returns

In Table 6, we report the performance of the $1/N$ portfolio and various short-sale-constrained mean-variance portfolios based on mean returns that are adjusted using the option-implied characteristics: model-free implied volatility (MFIV), implied-realized volatility spread (IRVS), model-free implied skewness (MFIS), and call-put volatility spread (CPVS). The methodology used to adjust returns is described in Section V.D.1. We report two p -values in parentheses, the first with respect to the $1/N$ portfolio, and the second with respect to the corresponding constrained minimum-variance benchmark portfolio in Table 2, with the null hypothesis being that the portfolio being evaluated is no better than the benchmark (so a small p -value suggests *rejecting* the null hypothesis that the portfolio being evaluated is no better than the benchmark).

Strategy	Sample 1			Sample 2		
	STD	SR	TRN	STD	SR	TRN
<i>Panel A. Daily Rebalancing</i>						
$1/N$	0.2083	0.5903	0.0134	0.2254	0.4963	0.0145
Constrained + MFIV	0.3959 (1.00) (1.00)	0.6785 (0.30) (0.19)	0.0983	0.4523 (1.00) (1.00)	0.4121 (0.73) (0.85)	0.0983
Constrained + IRVS	0.2499 (1.00) (1.00)	0.9232 (0.01) (0.02)	0.2265	0.2634 (1.00) (1.00)	0.7650 (0.01) (0.29)	0.2097
Constrained + MFIS	0.2520 (1.00) (1.00)	1.0092 (0.00) (0.00)	0.2859	0.2681 (1.00) (1.00)	1.0130 (0.00) (0.02)	0.2701
Constrained + CPVS	0.2518 (1.00) (1.00)	1.4291 (0.00) (0.00)	0.6209	0.2726 (1.00) (1.00)	1.3596 (0.00) (0.00)	0.6088
<i>Panel B. Weekly Rebalancing</i>						
$1/N$	0.1973	0.6045	0.0302	0.2127	0.5030	0.0325
Constrained + MFIV	0.3898 (1.00) (1.00)	0.5974 (0.53) (0.28)	0.1245	0.4387 (1.00) (1.00)	0.3657 (0.89) (0.94)	0.1286
Constrained + IRVS	0.2567 (1.00) (1.00)	0.7982 (0.03) (0.03)	0.2454	0.2689 (1.00) (1.00)	0.6216 (0.08) (0.65)	0.2316
Constrained + MFIS	0.2549 (1.00) (1.00)	0.8397 (0.01) (0.01)	0.3023	0.2690 (1.00) (1.00)	0.8063 (0.00) (0.20)	0.2876
Constrained + CPVS	0.2472 (1.00) (1.00)	0.9270 (0.00) (0.00)	0.6338	0.2578 (1.00) (1.00)	0.9058 (0.00) (0.06)	0.6229
<i>Panel C. Fortnightly Rebalancing</i>						
$1/N$	0.1905	0.6081	0.0427	0.2055	0.5036	0.0462
Constrained + MFIV	0.3745 (1.00) (1.00)	0.5850 (0.58) (0.31)	0.1445	0.4158 (1.00) (1.00)	0.3300 (0.94) (0.95)	0.1513
Constrained + IRVS	0.2590 (1.00) (1.00)	0.7070 (0.15) (0.08)	0.2594	0.2714 (1.00) (1.00)	0.5271 (0.39) (0.81)	0.2479
Constrained + MFIS	0.2520 (1.00) (1.00)	0.7867 (0.04) (0.02)	0.3151	0.2701 (1.00) (1.00)	0.7291 (0.00) (0.37)	0.3010
Constrained + CPVS	0.2392 (1.00) (1.00)	0.8132 (0.01) (0.02)	0.6436	0.2515 (1.00) (1.00)	0.7453 (0.00) (0.33)	0.6337

In Table 7, for each data set and rebalancing frequency, we report two sets of numbers, each set consisting of two numbers. The first set of numbers is for the case where the benchmark portfolio is the $1/N$ portfolio; the second set of numbers is for the case where the benchmark is the constrained minimum-variance portfolio reported in Table 2. The first number in each set indicates the transaction cost that equates the *mean return* on the portfolio using option-implied

characteristics to that of the benchmark portfolio. The second number is the transaction cost that equates the *Sharpe ratio* of the two portfolios. For comparison, note that the typical cost for trading the stocks that are in our data sets is about 10 basis points (bp) (see French (2008)), with the actual cost depending on the size of the trade and the execution capability of the trader.

TABLE 7
Equivalent Transaction Costs for Mean-Variance Portfolios
with Option-Characteristic Adjusted Returns

In Table 7, we evaluate the effect of transaction costs on the mean return and Sharpe ratio of mean-variance portfolios that rely on returns adjusted using option-implied characteristics, as described in Section 1. The four characteristics we consider are: model-free implied volatility (MFIV), implied-realized volatility spread (IRVS), model-free implied skewness (MFIS), and call-put volatility spread (CPVS). We report the transaction cost (in basis points) that makes the performance of the parametric portfolio equal to that of the benchmark portfolio, and we consider two benchmark portfolios: the 1/N portfolio and the short-sale-constrained minimum-variance portfolio. A "*" indicates the portfolio using option-implied information performs worse than the benchmark *and* also has higher turnover than the benchmark. A "—" indicates the performance of both portfolios, net of transaction costs, is negative.

Strategy	Sample 1		Sample 2	
	TC Mean	TC SR	TC Mean	TC SR
<i>Panel A. Daily Rebalancing</i>				
Constrained + MFIV	0.68	0.19	0.35	*
	1.92	—	0.56	*
Constrained + IRVS	0.20	0.16	0.18	0.15
	0.39	0.38	0.25	0.08
Constrained + MFIS	0.19	0.16	0.25	0.22
	0.33	0.31	0.30	0.18
Constrained + CPVS	0.16	0.14	0.17	0.16
	0.21	0.19	0.19	0.14
<i>Panel B. Weekly Rebalancing</i>				
Constrained + MFIV	2.40	*	1.11	*
	5.60	—	1.61	*
Constrained + IRVS	0.79	0.48	0.60	0.33
	1.58	1.43	0.81	*
Constrained + MFIS	0.69	0.45	0.86	0.66
	1.28	1.08	1.03	0.34
Constrained + CPVS	0.36	0.27	0.43	0.35
	0.59	0.45	0.49	0.22
<i>Panel C. Fortnightly Rebalancing</i>				
Constrained + MFIV	4.04	*	1.28	*
	8.60	—	1.95	*
Constrained + IRVS	1.24	0.51	0.78	0.14
	2.64	2.24	1.12	*
Constrained + MFIS	1.20	0.69	1.46	1.01
	2.29	1.89	1.74	0.31
Constrained + CPVS	0.52	0.33	0.57	0.42
	0.95	0.67	0.68	0.14

In Table 7, there are two possibilities when comparing the performance of portfolios that use option-implied information to the performance of benchmark strategies. The first possibility is that the portfolio using option-implied information performs better than the benchmark portfolio but has a higher turnover; in this case, the positive number reported in the table indicates the transaction cost that the portfolio using option-implied information can incur before its performance drops to the level of the benchmark, which is also considered net of transaction costs. The second possibility is that the benchmark portfolio performs better *and*

has better turnover; in this case, there is no positive level of transaction costs that equates the performance of the two strategies, and we indicate this in the table with the symbol “*”. The third possibility is that the performance of both portfolios turns negative after adjustment for transaction costs; in the table, we use “—” to represent this case.

We observe from Panel A of Table 7 that for the case of daily rebalancing, the equivalent transaction cost for the Sharpe ratio ranges from 14 bp to 38 bp; for example, in the case of Sample 2, the portfolio with returns adjusted using MFIS has a higher Sharpe ratio than the $1/N$ portfolio for transaction costs of up to 22 bp. In Panels B and C, we see that as the rebalancing frequency decreases, and the equivalent transaction cost increases; for example, in the case of Sample 2, the portfolio with returns adjusted using MFIS has a higher Sharpe ratio than the $1/N$ portfolio for transaction costs of up to 66 bp for weekly rebalancing, and up to 101 bp for fortnightly rebalancing.

We conclude from this analysis that information in MFIV, IRVS, MFIS, and CPVS can be used to improve the Sharpe ratio of mean-variance portfolios even after adjusting for the higher transaction cost as a consequence of higher turnover in implementing the option-based strategy.

2. Performance of a Parametric Portfolio Based on Option Characteristics

Next, we examine the out-of-sample performance of the *parametric portfolios* that use option-implied characteristics. We consider two benchmark portfolios: the $1/N$ portfolio and the mean-variance parametric portfolio “ $1/N + \text{FFM}$ ” that starts with the $1/N$ portfolio and uses the Fama-French and momentum characteristics to adjust the portfolio weights; we do not allow for short sales. Comparing the two benchmarks, we see that the $1/N$ portfolio has better turnover and volatility, but the parametric portfolio has a better Sharpe ratio.

From Table 8, we see that when we use the option-implied characteristics alone instead of the FFM, then the implied skewness measures (MFIS and CPVS) improve the risk-return tradeoff significantly compared to the two benchmarks. Moreover, the improvement over $1/N$ is significant for daily, weekly, and fortnightly holding periods, while the improvement over the parametric portfolio is significant only for daily rebalancing. For instance, in Panel A for Sample 1, the Sharpe ratio for the $1/N$ portfolio is 0.5903 and for the parametric portfolio with FFM factors it is 0.6589, while for the parametric portfolio with IRVS it is 0.6433, for the parametric portfolio with MFIS it is 0.7908, and for the parametric portfolio with CPVS it is 0.9390; the results are similar for Sample 2. However, these gains are accompanied by higher turnover. Note that as the rebalancing period increases, the improvement decreases, while turnover *per rebalancing period* stays at about the same level. This implies that the total transaction cost paid over the entire time period is decreasing.

We can also ask whether the option-implied characteristics improve performance if one is already using the traditional size, value, and momentum characteristics when selecting the parametric portfolio. This question is answered in the second part of each panel of Table 8, where we consider each option-implied characteristic *in addition to* the FFM characteristics considered in the benchmark portfolio. From the lower part of each panel, we see that MFIS and CPVS improve the

TABLE 8
Parametric Portfolios Using Option-Implied Information

In Table 8, we evaluate the performance of parametric portfolios that rely on option-implied characteristics about expected returns, using the methodology proposed in Brandt et al. (2009). In each panel, we first consider the effect of choosing a portfolio that is based on option-implied characteristic (individually and together), rather than the traditional FFM characteristics. The characteristics we consider are: model-free implied volatility (MFIV), implied-realized volatility spread (IRVS), model-free implied skewness (MFIS), and call-put volatility spread (CPVS). In the second part of each panel, we consider the effect of considering these option-implied characteristics (individually and together) *in addition to* the FFM characteristics. We report two p -values in parentheses, the first with respect to the $1/N$ portfolio and the second with respect to the "1/N + FFM" portfolio, with the null hypothesis being that the portfolio being evaluated is no better than the benchmark (so a small p -value suggests *rejecting* the null hypothesis that the portfolio being evaluated is no better than the benchmark).

Strategy	Sample 1			Sample 2		
	STD	SR	TRN	STD	SR	TRN
<i>Panel A. Daily Rebalancing</i>						
1/N	0.2083	0.5903	0.0134	0.2254	0.4963	0.0145
1/N + FFM	0.2079 (0.38) (0.50)	0.6589 (0.06) (0.50)	0.0515	0.2285 (1.00) (0.50)	0.5453 (0.14) (0.50)	0.0372
1/N + MFIV	0.2142 (1.00) (1.00)	0.6180 (0.33) (0.75)	0.0780	0.2269 (0.84) (0.16)	0.5063 (0.44) (0.74)	0.0605
1/N + IRVS	0.2062 (0.00) (0.09)	0.6433 (0.01) (0.63)	0.0904	0.2246 (0.19) (0.00)	0.5413 (0.04) (0.54)	0.0855
1/N + MFIS	0.2143 (1.00) (1.00)	0.7908 (0.00) (0.01)	0.2033	0.2378 (1.00) (1.00)	0.7075 (0.00) (0.00)	0.1950
1/N + CPVS	0.2093 (0.98) (0.87)	0.9390 (0.00) (0.00)	0.2944	0.2300 (1.00) (0.92)	0.8516 (0.00) (0.00)	0.2935
1/N + ALL IMPLIED	0.2178 (1.00) (1.00)	0.9269 (0.00) (0.00)	0.3352	0.2346 (1.00) (1.00)	0.8493 (0.00) (0.00)	0.2844
1/N + FFM + MFIV	0.2140 (1.00) (1.00)	0.6635 (0.12) (0.47)	0.0891	0.2270 (0.85) (0.07)	0.5714 (0.09) (0.23)	0.0613
1/N + FFM + IRVS	0.2070 (0.08) (0.02)	0.6645 (0.04) (0.39)	0.0932	0.2257 (0.60) (0.00)	0.5556 (0.09) (0.29)	0.0607
1/N + FFM + MFIS	0.2091 (0.79) (0.98)	0.7996 (0.00) (0.00)	0.2011	0.2340 (1.00) (1.00)	0.6649 (0.00) (0.00)	0.1683
1/N + FFM + CPVS	0.2127 (1.00) (1.00)	1.0174 (0.00) (0.00)	0.3497	0.2333 (1.00) (1.00)	0.8655 (0.00) (0.00)	0.3047
1/N + FFM + ALL IMPLIED	0.2167 (1.00) (1.00)	1.0184 (0.00) (0.00)	0.3999	0.2351 (1.00) (1.00)	0.8619 (0.00) (0.00)	0.3164
<i>Panel B. Weekly Rebalancing</i>						
1/N	0.1973	0.6045	0.0302	0.2127	0.5030	0.0325
1/N + FFM	0.1979 (0.66) (0.50)	0.6725 (0.05) (0.50)	0.0624	0.2156 (0.94) (0.50)	0.5617 (0.08) (0.50)	0.0508
1/N + MFIV	0.2046 (1.00) (1.00)	0.6349 (0.30) (0.76)	0.0880	0.2157 (0.89) (0.53)	0.4965 (0.55) (0.89)	0.0723
1/N + IRVS	0.1966 (0.23) (0.25)	0.6270 (0.12) (0.84)	0.0970	0.2144 (0.94) (0.30)	0.5250 (0.13) (0.81)	0.0938
1/N + MFIS	0.2045 (1.00) (1.00)	0.6757 (0.00) (0.47)	0.2064	0.2269 (1.00) (1.00)	0.5884 (0.00) (0.26)	0.1992
1/N + CPVS	0.1978 (0.91) (0.46)	0.7034 (0.00) (0.23)	0.2960	0.2167 (1.00) (0.75)	0.6031 (0.00) (0.13)	0.2956

(continued on next page)

TABLE 8 (continued)
 Parametric Portfolios Using Option-Implied Information

Strategy	Sample 1			Sample 2		
	STD	SR	TRN	STD	SR	TRN
<i>Panel B. Weekly Rebalancing (continued)</i>						
1/N + ALL IMPLIED	0.2074 (1.00) (1.00)	0.6915 (0.01) (0.33)	0.3372	0.2215 (1.00) (1.00)	0.6097 (0.00) (0.11)	0.2869
1/N + FFM + MFIV	0.2042 (1.00) (1.00)	0.6697 (0.11) (0.52)	0.0973	0.2149 (0.80) (0.34)	0.5732 (0.07) (0.35)	0.0724
1/N + FFM + IRVS	0.1963 (0.31) (0.01)	0.6700 (0.04) (0.56)	0.1006	0.2141 (0.79) (0.02)	0.5505 (0.11) (0.78)	0.0712
1/N + FFM + MFIS	0.1990 (0.88) (0.90)	0.7075 (0.00) (0.06)	0.2044	0.2237 (1.00) (1.00)	0.5760 (0.01) (0.31)	0.1729
1/N + FFM + CPVS	0.2021 (1.00) (1.00)	0.7538 (0.00) (0.01)	0.3511	0.2209 (1.00) (1.00)	0.6168 (0.00) (0.04)	0.3068
1/N + FFM + ALL IMPLIED	0.2075 (1.00) (1.00)	0.7600 (0.00) (0.01)	0.4014	0.2235 (1.00) (1.00)	0.6284 (0.00) (0.01)	0.3185
<i>Panel C. Fortnightly Rebalancing</i>						
1/N	0.1905	0.6081	0.0427	0.2055	0.5036	0.0462
1/N + FFM	0.1903 (0.48) (0.50)	0.6749 (0.06) (0.50)	0.0719	0.2075 (0.77) (0.50)	0.5657 (0.08) (0.50)	0.0623
1/N + MFIV	0.1979 (0.99) (1.00)	0.6441 (0.29) (0.71)	0.0972	0.2084 (0.82) (0.61)	0.5023 (0.53) (0.88)	0.0830
1/N + IRVS	0.1913 (0.74) (0.64)	0.6226 (0.20) (0.87)	0.1036	0.2085 (0.98) (0.66)	0.5114 (0.33) (0.90)	0.1020
1/N + MFIS	0.1986 (1.00) (1.00)	0.6631 (0.00) (0.59)	0.2103	0.2205 (1.00) (1.00)	0.5585 (0.00) (0.58)	0.2041
1/N + CPVS	0.1911 (0.92) (0.61)	0.6627 (0.00) (0.60)	0.2984	0.2094 (1.00) (0.79)	0.5559 (0.00) (0.59)	0.2985
1/N + ALL IMPLIED	0.2008 (1.00) (1.00)	0.6662 (0.05) (0.58)	0.3398	0.2137 (1.00) (0.99)	0.5678 (0.00) (0.49)	0.2901
1/N + FFM + MFIV	0.1973 (0.98) (1.00)	0.6712 (0.15) (0.55)	0.1051	0.2078 (0.76) (0.57)	0.5665 (0.11) (0.49)	0.0825
1/N + FFM + IRVS	0.1897 (0.36) (0.22)	0.6681 (0.07) (0.67)	0.1076	0.2068 (0.70) (0.23)	0.5480 (0.13) (0.91)	0.0808
1/N + FFM + MFIS	0.1926 (0.86) (0.99)	0.7030 (0.00) (0.12)	0.2084	0.2163 (1.00) (1.00)	0.5484 (0.07) (0.72)	0.1782
1/N + FFM + CPVS	0.1957 (1.00) (1.00)	0.7165 (0.00) (0.11)	0.3533	0.2146 (1.00) (1.00)	0.5731 (0.00) (0.41)	0.3097
1/N + FFM + ALL IMPLIED	0.2009 (1.00) (1.00)	0.7286 (0.00) (0.06)	0.4034	0.2166 (1.00) (1.00)	0.5865 (0.00) (0.24)	0.3212

Sharpe ratio beyond the gains obtained from using only the FFM characteristics. This improvement in Sharpe ratio for daily rebalancing is statistically significant for both Samples 1 and 2 when using MFIS and CPVS; for weekly rebalancing, it is statistically significant for CPVS for both samples, and for MFIS only for Sample 1; and it is not statistically significant for fortnightly rebalancing. Finally, looking at the last row of each panel, where we add all four option-implied characteristics to the portfolio that already considers the FFM characteristics, we observe that the Sharpe ratio increases substantially, and this improvement is statistically significant for daily and weekly rebalancing for both samples; for fortnightly rebalancing it is significant for Sample 1.

To understand if the parametric portfolios using option-implied characteristics outperform the benchmark portfolios even in the presence of transaction costs, just as before, we compute the *equivalent transaction cost* for each portfolio. In Table 9, for each data set and rebalancing frequency, we report two sets of numbers, each set having two numbers. The first set of numbers is for the case where the benchmark portfolio is the $1/N$ portfolio; the second set of numbers is for the case where the benchmark is the parametric portfolio with the FFM characteristics.

TABLE 9
Equivalent Transaction Costs for Parametric Portfolios Using Option-Implied Information

In Table 9, we evaluate the effect of transaction costs on the mean return and Sharpe ratio of parametric portfolios that rely on option-implied characteristics, using the methodology proposed in Brandt et al. (2009). The four characteristics we consider (individually and together) are: model-free implied volatility (MFIV), implied-realized volatility spread (IRVS), model-free implied skewness (MFIS), and call-put volatility spread (CPVS). In the second part of each panel, we consider the effect of considering these option-implied characteristics (individually and together) in addition to the FFM characteristics. We report the transaction cost (in basis points) that makes the performance of the parametric portfolio equal to that of the benchmark portfolio, and we consider two benchmark portfolios: the $1/N$ portfolio and the " $1/N + \text{FFM}$ " parametric portfolio. A "*" indicates the portfolio using option-implied information performs worse than the benchmark and also has higher turnover than the benchmark.

Strategy	Sample 1		Sample 2	
	TC Mean	TC SR	TC Mean	TC SR
<i>Panel A. Daily Rebalancing</i>				
1/N + MFIV	0.06 *	0.04 *	0.03 *	0.02 *
1/N + IRVS	0.05 *	0.06 *	0.05 *	0.06 *
1/N + MFIS	0.10 0.09	0.09 0.07	0.12 0.11	0.11 0.10
1/N + CPVS	0.10 0.10	0.10 0.10	0.12 0.11	0.12 0.11
1/N + ALL IMPLIED	0.10 0.09	0.09 0.08	0.13 0.12	0.12 0.12
1/N + FFM + MFIV	0.10 0.05	0.08 0.01	0.15 0.08	0.15 0.10
1/N + FFM + IRVS	0.07 0.00	0.08 0.01	0.12 0.01	0.12 0.04
1/N + FFM + MFIS	0.09 0.08	0.09 0.08	0.11 0.09	0.10 0.09
1/N + FFM + CPVS	0.11 0.11	0.11 0.10	0.12 0.12	0.12 0.11
1/N + FFM + ALL IMPLIED	0.10 0.10	0.10 0.09	0.12 0.11	0.11 0.11

(continued on next page)

TABLE 9 (continued)
 Equivalent Transaction Costs for Parametric Portfolios Using Option-Implied Information

Strategy	Sample 1		Sample 2	
	TC Mean	TC SR	TC Mean	TC SR
<i>Panel B. Weekly Rebalancing</i>				
1/N + MFIV	0.37 *	0.22 *	0.01 *	* *
1/N + IRVS	0.12 *	0.13 *	0.18 *	0.15 *
1/N + MFIS	0.21 0.07	0.17 0.01	0.32 0.17	0.23 0.08
1/N + CPVS	0.15 0.05	0.15 0.05	0.18 0.08	0.16 0.07
1/N + ALL IMPLIED	0.16 0.07	0.12 0.03	0.22 0.12	0.19 0.09
1/N + FFM + MFIV	0.52 0.21	0.40 *	0.81 0.19	0.76 0.23
1/N + FFM + IRVS	0.35 *	0.36 *	0.56 *	0.53 *
1/N + FFM + MFIS	0.25 0.11	0.23 0.10	0.31 0.13	0.23 0.05
1/N + FFM + CPVS	0.21 0.13	0.19 0.11	0.21 0.12	0.18 0.10
1/N + FFM + ALL IMPLIED	0.21 0.14	0.17 0.11	0.23 0.14	0.20 0.11
<i>Panel C. Fortnightly Rebalancing</i>				
1/N + MFIV	0.85 *	0.53 *	0.13 *	* *
1/N + IRVS	0.21 *	0.18 *	0.22 *	0.12 *
1/N + MFIS	0.38 0.09	0.26 *	0.50 0.16	0.31 *
1/N + CPVS	0.17 *	0.16 *	0.20 *	0.17 *
1/N + ALL IMPLIED	0.24 0.08	0.16 *	0.29 0.07	0.23 0.01
1/N + FFM + MFIV	1.06 0.48	0.82 *	1.56 0.07	1.46 0.03
1/N + FFM + IRVS	0.67 *	0.70 *	1.14 *	1.07 *
1/N + FFM + MFIS	0.47 0.20	0.44 0.16	0.46 0.04	0.30 *
1/N + FFM + CPVS	0.31 0.17	0.27 0.12	0.29 0.09	0.23 0.03
1/N + FFM + ALL IMPLIED	0.34 0.22	0.27 0.13	0.34 0.15	0.26 0.07

We observe from Panel A of Table 9 that for the case of daily rebalancing, the benefits of using option-implied information would be eliminated if one had to pay a transaction cost of more than 11 bp. This is because of the relatively higher turnover of the strategies using option-implied information. However, when we rebalance less frequently, the benefits of using option-implied information survive higher transaction costs. For example, the results in Panel C for fortnightly rebalancing indicate that MFIS improves performance relative to the 1/N strategy even when we pay transaction costs of about 30 bp. However, option-implied characteristics improve performance relative to the benchmark parametric portfolio for

transaction costs of up to 10 bp. Using all of the implied characteristics is beneficial in the presence of transaction costs of about 20 bp with respect to the $1/N$ benchmark, and about 10 bp for the parametric-portfolio benchmark.

In summary, Table 9 suggests that in the presence of transaction costs, MFIS and CPVS are useful characteristics for choosing portfolios, while the value of using other characteristics is smaller, and their contribution is not robust across the two samples. Overall, the empirical evidence suggests that using option-implied skewness can lead to an improvement in the out-of-sample portfolio Sharpe ratio, even after adjusting for the higher transaction costs incurred by these strategies.

VI. Conclusion

Mean-variance portfolio weights depend on estimates of volatilities, correlations, and expected returns of stocks. In this paper, we have studied how information implied in prices of stock options can be used to improve estimates of these three moments in order to improve the out-of-sample performance of portfolios with a large number of stocks. Performance is measured in terms of portfolio volatility, Sharpe ratio, and turnover. The benchmark portfolios are the $1/N$ portfolio; four types of minimum-variance portfolios and four types of mean-variance portfolios based on historical returns; and the parametric portfolios of Brandt et al. (2009), based on historical returns and size, value, and momentum characteristics.

We find that using option-implied volatilities can lead to a significant improvement in portfolio volatility; however, option-implied correlations are less useful in reducing portfolio volatility. We also find that forming portfolios using expected returns that exploit information in option-implied model-free skewness and implied volatility achieve a higher Sharpe ratio than portfolios that ignore option-implied information; this improvement in performance is present even after adjusting for transaction costs. Based on our empirical analysis, we conclude that prices of stock options contain information that can be useful for improving the out-of-sample performance of portfolios.

References

- Ait-Sahalia, Y., and M. W. Brandt. "Consumption and Portfolio Choice with Option-Implied State Prices." NBER Working Paper (2008).
- Andersen, T. G.; T. Bollerslev; and F. X. Diebold. "Parametric and Nonparametric Measurement of Volatility." In *Handbook of Financial Econometrics*, Y. Ait-Sahalia, and L. P. Hansen, eds. Amsterdam: North Holland (2009).
- Ang, A.; T. G. Bali; and N. Cakici. "The Joint Cross Section of Stocks and Options." Working Paper, Fordham University (2010).
- Bakshi, G. S.; N. Kapadia; and D. B. Madan. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies*, 16 (2003), 101–143.
- Bali, T. G., and A. Hovakimian. "Volatility Spreads and Expected Stock Returns." *Management Science*, 55 (2009), 1797–1812.
- Black, F., and M. S. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), 637–659.
- Bollerslev, T.; M. Gibson; and H. Zhou. "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities." *Finance and Economics Discussion Series*, 2004-56, Board of Governors of the Federal Reserve System, Washington, DC (2004).
- Bollerslev, T.; G. Tauchen; and H. Zhou. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies*, 22 (2009), 4463–4492.

- Boyer, B. H.; T. Mitton; and K. Vorkink. "Expected Idiosyncratic Skewness." *Review of Financial Studies*, 23 (2010), 169–202.
- Brandt, M. W.; P. Santa-Clara; and R. Valkanov. "Parametric Portfolio Policies: Exploiting Characteristics in the Cross Section of Equity Returns." *Review of Financial Studies*, 22 (2009), 3411–3447.
- Buss, A., and G. Vilkov. "Measuring Equity Risk with Option-Implied Correlations." *Review of Financial Studies*, 25 (2012), 3113–3140.
- Carr, P., and L. Wu. "Variance Risk Premiums." *Review of Financial Studies*, 22 (2009), 1311–1341.
- Chan, L. K. C.; J. Karceski; and J. Lakonishok. "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model." *Review of Financial Studies*, 12 (1999), 937–974.
- Corsi, F. "A Simple Approximate Long-Memory Model of Realized Volatility." *Journal of Financial Econometrics*, 7 (2009), 174–196.
- Cox, J. C., and C.-F. Huang. "Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process." *Journal of Economic Theory*, 49 (1989), 33–83.
- Cremers, M., and D. Weinbaum. "Deviations from Put-Call Parity and Stock Return Predictability." *Journal of Financial and Quantitative Analysis*, 45 (2010), 335–367.
- DeMiguel, V.; L. Garlappi; F. Nogales; and R. Uppal. "A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms." *Management Science*, 55 (2009), 798–812.
- DeMiguel, V.; L. Garlappi; and R. Uppal. "Optimal versus Naive Diversification: How Inefficient Is the $1/N$ Portfolio Strategy?" *Review of Financial Studies*, 22 (2009), 1915–1953.
- Driessen, J.; P. Maenhout; and G. Vilkov. "The Price of Correlation Risk: Evidence from Equity Options." *Journal of Finance*, 64 (2009), 1375–1404.
- Efron, B., and R. Tibshirani. *An Introduction to the Bootstrap*. Boca Raton, FL: Chapman and Hall (1993).
- Elton, E. J.; M. J. Gruber; and J. Spitzer. "Improved Estimates of Correlation Coefficients and Their Impact on Optimum Portfolios." *European Financial Management*, 12 (2006), 303–318.
- Fama, E. F., and K. R. French. "The Cross-Section of Expected Stock Returns." *Journal of Finance*, 47 (1992), 427–465.
- French, K. R. "Presidential Address: The Cost of Active Investing." *Journal of Finance*, 63 (2008), 1537–1573.
- Goyal, A., and A. Saretto. "Cross-Section of Option Returns and Volatility." *Journal of Financial Economics*, 94 (2009), 310–326.
- Grundy, B. D., and S. Martin. "Understanding the Nature of the Risks and the Source of the Rewards to Momentum Investing." *Review of Financial Studies*, 14 (2001), 29–78.
- Guidolin, M., and A. Timmermann. "International Asset Allocation under Regime Switching, Skew and Kurtosis Preferences." *Review of Financial Studies*, 21 (2008), 889–935.
- Harvey, C. R., and A. Siddique. "Conditional Skewness in Asset Pricing Tests." *Journal of Finance*, 55 (2000), 1263–1295.
- Jacobs, H.; S. Muller; and M. Weber. "How Should Private Investors Diversify? An Empirical Evaluation of Alternative Asset Allocation Policies to Construct a 'World Market Portfolio'." Working Paper, University of Mannheim (2010).
- Jagannathan, R., and T. Ma. "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps." *Journal of Finance*, 58 (2003), 1651–1684.
- Jegadeesh, N., and S. Titman. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance*, 48 (1993), 65–91.
- Jobson, J. D.; R. Korkie; and V. Ratti. "Improved Estimation for Markowitz Portfolios Using James-Stein Type Estimators." *Proceedings of the American Statistical Association*, 41 (1979), 279–284.
- Jorion, P. "Bayes-Stein Estimation for Portfolio Analysis." *Journal of Financial and Quantitative Analysis*, 21 (1986), 279–292.
- Kostakis, A.; N. Panigirtzoglou; and G. Skiadopoulos. "Market Timing with Option-Implied Distributions: A Forward-Looking Approach." *Management Science*, 57 (2011), 1231–1249.
- Kraus, A., and R. H. Litzenberger. "Skewness Preference and the Valuation of Risk Assets." *Journal of Finance*, 31 (1976), 1085–1100.
- Ledoit, O., and M. Wolf. "Honey, I Shrunk the Sample Covariance Matrix." *Journal of Portfolio Management*, 30 (2004a), 110–119.
- Ledoit, O., and M. Wolf. "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices." *Journal of Multivariate Analysis*, 88 (2004b), 365–411.
- Ledoit, O., and M. Wolf. "Robust Performance Hypothesis Testing with the Sharpe Ratio." *Journal of Empirical Finance*, 15 (2008), 850–859.
- Merton, R. C. "Optimum Consumption and Portfolio Rules in a Continuous-Time Model." *Journal of Economic Theory*, 3 (1971), 373–413.
- Merton, R. C. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science*, 4 (1973), 141–183.

- Newey, W. K., and K. D. West. "A Simple, Positive-Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (1987), 703–708.
- Pástor, Ľ. "Portfolio Selection and Asset Pricing Models." *Journal of Finance*, 55 (2000), 179–223.
- Politis, D., and J. Romano. "The Stationary Bootstrap." *Journal of the American Statistical Association*, 89 (1994), 1303–1313.
- Poon, S.-H., and C. W. J. Granger. "Practical Issues in Forecasting Volatility." *Financial Analysts Journal*, 61 (2005), 45–56.
- Rehman, Z., and G. Vilkov. "Risk-Neutral Skewness: Return Predictability and Its Sources." Working Paper, Goethe University Frankfurt (2009).
- Rubinstein, M. E. "The Fundamental Theorem of Parameter-Preference Security Valuation." *Journal of Financial and Quantitative Analysis*, 8 (1973), 61–69.
- Vanden, J. M. "Information Quality and Options." *Review of Financial Studies*, 21 (2008), 2635–2676.
- Xing, Y.; X. Zhang; and R. Zhao. "What Does Individual Option Volatility Smirk Tell Us About Future Equity Returns?" *Journal of Financial and Quantitative Analysis*, 45 (2010), 641–662.
- Zumbach, G. "Inference on Multivariate ARCH Processes with Large Sizes." Available at <http://dx.doi.org/10.2139/ssrn.1346416> (2009).