

Improving Portfolio Selection Using Option-Implied Volatility and Skewness*

Victor DeMiguel[‡]

Yuliya Plyakha[§]

Raman Uppal[¶]

Grigory Vilkov[§]

This version: June 11, 2012

*We gratefully acknowledge financial support from Inquire-Europe; however, this article represents the views of the authors and not of Inquire. We would like to acknowledge detailed feedback from Luca Benzoni, Pascal Maenhout, and Christian Schlag. We received helpful comments and suggestions from Alexander Alekseev, 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-UNC Asset Pricing Conference, EDHEC-Risk Seminar on Advanced Portfolio Construction, Financial Econometrics Conference at Toulouse School of Economics, SIFR 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.

[‡]London Business School, 6 Sussex Place, Regent's Park, London, United Kingdom NW1 4SA; E-mail: avmiguel@london.edu.

[§]Goethe University Frankfurt, Finance Department, Grüneburgplatz 1 / Uni-PfH 25, D-60323 Frankfurt am Main, Germany; Email: plyakha@finance.uni-frankfurt.de and vilkov@vilkov.net.

[¶]CEPR and Edhec Business School, 10 Fleet Place, Ludgate, London, United Kingdom EC4M 7RB; Email: raman.uppal@edhec.edu.

Improving Portfolio Selection Using Option-Implied Volatility and Skewness

This version: June 11, 2012

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 for improving 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 shortsales and accounting for transactions costs.

Keywords: mean variance, option-implied volatility, variance risk premium, option-implied skewness, portfolio optimization

JEL: G11, G12, G13, G17

1 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. Traditionally, historical data on returns have been 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 model, but also using the model-free approach developed in Bakshi, Kapadia, and Madan (2003).

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 shortsale-constrained minimum-variance portfolio, the minimum-variance portfolio with shrinkage

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; Ledoit and Wolf, 2004b), constraining shortsales (Jagannathan and Ma, 2003), 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 that: “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.”

of the covariance matrix (as in Ledoit and Wolf (2004a,b)), and the minimum-variance portfolio obtained by assuming all correlations are equal to zero 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 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, Santa-Clara, and Valkanov (2009), leads to a substantial improvement in the Sharpe ratio, even after prohibiting shortsales and accounting for transactions 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

⁴For instance, Bollerslev, Tauchen, and Zhou (2009) have documented 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 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.

literature. The focus of our work is to investigate how the information implied by option prices can be used to improve portfolio selection. There are two other papers that study this. The first, by Aït-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, its focus 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 S&P500 index and a riskless asset. This 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 this work and ours: rather than considering the problem of how to allocate wealth between the S&P500 index and the riskfree 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, Panigirtzoglou, and Skiadopoulos (2011) 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 2, we describe the data on stocks and options that we use. In Section 3, we explain how we use data on options to predict volatilities, correlations, and expected returns. In Section 4, 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 5. We conclude in Section 6.

2 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 January 1996 to October 2010.⁵

2.1 Data on Stock Returns and Stock Characteristics

We study stocks that are in the S&P500 index at any time during our sample period. The *daily* stock returns of the S&P500 constituents are from the daily file of the 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 the data are available for all dates. The second, “Sample 2,” consists of all the stocks that are part of the S&P500 index on that day and which have no missing data on that day (such as prices of options on the same underlying for a given maturity and across different strikes that are needed to compute model-free implied volatilities), which means that the second sample has a variable number of stocks; on average, this sample has about 400 stocks at each point in time.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 twelve-month momentum (MOM) characteristic is measured for each day t using daily returns data from CRSP as the cumulative return from day $t - 21 - 251$ to day $t - 21$.⁸

⁵We carry out all the tests included in the manuscript also for the pre-crisis period from January 1996 to December 2007, and the crisis period from January 2008 to June 2009 (identified as a recession by the NBER). The main insights of our analysis do not change with the choice of sample period.

⁶We also use high-frequency *intraday* stock-price data consisting of transaction prices for the S&P500 constituents from the NYSE’s Trades-And-Quotes 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 constructed characteristics, we take the logarithm of size and value characteristics. In order to prepare these characteristics so that they can be used to compute the parametric portfolio weights, we also winsorize the characteristics by assigning the value of

2.2 Data on Stock Options

For stock options we use IvyDB that 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 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 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.

3 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 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.

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.

3.1 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 the *model-free implied volatility* (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 better than the Black-Scholes volatility. We compute MFIV as the square root of the variance contract of Bakshi, Kapadia, and Madan (2003), as explained below.

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, Kapadia, and Madan (2003). Define

$$\mu(t, \tau) = e^{r\tau} - 1 - \frac{e^{r\tau}}{2} V(t, \tau) - \frac{e^{r\tau}}{6} W(t, \tau) - \frac{e^{r\tau}}{24} X(t, \tau).$$

Then, the τ -period model-free implied volatility (MFIV) can be calculated as

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

and the τ -period model-free implied skewness (MFIS) as

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

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

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).

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 1001 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 model-free implied volatility and model-free implied skewness as in (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 model-free implied volatility 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), Carr and Wu (2009), and others have shown that one can use realized volatility (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

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

In our analysis, we estimate the historical volatility risk premium adjustment on each day over the past year (-272 days to -21 days) using the model-free implied volatility and realized volatility 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 (3), one can obtain the *prediction* of the future realized volatility, $\widehat{\text{RV}}_t$,

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 will have only T observations when computing the sum.

which we call the *risk-premium-corrected implied volatility*:¹²

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

We now wish to confirm the intuition that risk-premium-corrected implied volatility is better than historical volatility at predicting realized volatility. To do this, we consider as a predictor, first historical volatility and then risk-premium-corrected implied volatility, for the monthly realized volatility 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 upwards with respect to the realized volatility; 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.

3.2 Predicting Correlations Using Options

The second piece of option-implied information that we consider is implied correlation, and, because for the portfolio optimization we need the correlation under the *objective* measure, 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:

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

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

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

$\widehat{\sigma}_{i,j,t}^P$ that turn Equation (4) into identity. Once we substitute into (4) the expected volatilities, we have one equation with $N \times (N - 1) / 2$ unknown correlations, $\widehat{\rho}_{i,j,t}^P$. Thus, to compute all pairwise correlations we need to make some identifying assumptions.

We use the approach proposed in Buss and Vilkov (2011), 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 one:

$$\rho_{i,j,t}^P - \widehat{\rho}_{i,j,t}^P = \psi_t(1 - \rho_{i,j,t}^P),$$

which implies that

$$\widehat{\rho}_{i,j,t}^P = \rho_{i,j,t}^P - \psi_t(1 - \rho_{i,j,t}^P). \quad (5)$$

When we substitute this into Equation (4) above, we get:

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

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

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

and then the expected correlations $\widehat{\rho}_{i,j,t}^P$ from equation (5). 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-days realized correlation. In both Sample 1 and 2, we find that the RMSE for historical correlation is about 0.25 and is slightly smaller than the RMSE of 0.26 for the option-implied correlation; note, however, that the RMSE for both predictors is only

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

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.

3.3 Explaining Returns Using Options

There are four option-based quantities that we use 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.

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 model-free implied volatility (MFIV) described in Section 3.1.

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 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. We use two measures of this. The first, model-free implied skewness (MFIS), as defined in Equation (2), represents a nonparametric estimate of the risk-neutral stock-return skewness.¹⁴ Rehman

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 (2009). For a study of asset allocation that takes into account

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 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 (2009) to compute the call-put volatility spread (CPVS) as the difference between the current Black-Scholes implied volatilities of the one-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-sectional 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 (one day, one week, or two weeks). In Table 1 we show the annualized returns for each portfolio, along with the p-value, 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 in the table standard characteristics such as size (SIZE), book-to-market (BTM), and 12-month momentum (MOM).

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

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 (for example, Xing, Zhang, and Zhao (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 skew.

implied-realized-volatility spread (IRVS) is also positively and significantly related to returns, but at the 10% level.

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 the daily holding, it is the CPVS portfolio that earns the highest return of 86.99% 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 investment horizons that we consider.

4 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 six benchmark portfolios that do *not* rely on option-implied information.

4.1 Equal-Weighted Portfolio

For the “*equal-weighted*” ($1/N$) portfolio, 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, Muller, and Weber (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&P500 over our sample period.

4.2 Minimum-Variance Portfolios

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

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

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

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 (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 (7) ensures that the portfolio weights sum to one; 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 reduces to finding the portfolio that minimizes the variance of the portfolio return, subject to the constraints that the portfolio weights sum to one. The solution to the resulting *minimum-variance* portfolio problem is:

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

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

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

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 (8) based on the *sample covariance* matrix, there are two quantities that 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 shortsale-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,b).

We also consider two other methods proposed in the literature for improving the behavior of the covariance matrix (see Elton, Gruber, and Spitzer (2006) and the references therein). The first relies on setting all correlations equal to zero 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.

4.3 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, for example, the papers cited in Footnote 4) has found 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. At first sight, one may want to study the effect of using these variables directly in improving the performance of traditional mean-variance portfolios obtained from solving the problem in (6) subject to the constraint in (7). But, it is well known in the literature (see, for example, DeMiguel, Garlappi, and Uppal (2009) and Jacobs, Muller, and Weber (2010)) that the weights of the traditional mean-variance portfolio are very sensitive to errors in estimates of expected returns, and 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 form portfolio deciles, and then adjust the returns of only the top and bottom decile portfolios by a constant factor. In the second approach, we use the ‘parametric-portfolio methodology,’ of Brandt, Santa-Clara, and Valkanov (2009), which has been developed to deal with the problem of poor out-of-sample performance of

mean-variance portfolios because of estimation error, and which can be interpreted as a method where the adjustment of stock returns is done in an optimal fashion. We provide details of these two approaches below.

4.3.1 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

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

where $\mu_{bench,t}$ 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 to 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

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

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

4.3.2 Mean-Variance Parametric Portfolios

In the second approach, we apply the “parametric-portfolio methodology,” of Brandt, Santa-Clara, and Valkanov (2009) by using model-free implied volatility (MFIV), model-free implied skewness (MFIS), the call-put-implied volatility spread (CPVS), and the implied-realized-volatility spread (IRVS), in addition to the traditional stock character-

istics (size, value and momentum), to construct parametric portfolios based on mean-variance utility.

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, Santa-Clara, and Valkanov (2009), we normalize the characteristics to have zero mean and unit variance. Note that $\theta_{k,t}$ is *not* asset-specific, but 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 shortsales; 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, Santa-Clara, and Valkanov (2009), but using the $1/N$ portfolio as the benchmark: 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: model-free implied volatility (MFIV), implied-realized volatility spread (IRVS), model-free implied skewness (MFIS), and call-put implied volatility spread (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 4.1–4.3 is reported in Tables 2 and 3 and discussed in Section 5.1.

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

4.4 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. For the weekly and fortnightly rebalancing interval, we find the new set of weights daily, but hold that portfolio for 7 or 14 calendar days, which corresponds to the average of 5 or 10 daily returns; the advantage of computing new weights daily for weekly and fortnightly rebalancing is that the results are then not sensitive to the particular day chosen to form the portfolio.

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, $r_{t+\Delta t}^{strategy} = w_t^{strategy} r_{t+\Delta t}$, where $r_{t+\Delta t}$ denotes the returns from t to $t + \Delta t$, and Δt is one day, one week, or one fortnight. After collecting the time series of $T - \tau - \Delta t$ returns, $r_t^{strategy}$, the out-of-sample mean, volatility ($\hat{\sigma}$), and Sharpe ratio of returns (SR) are:

$$\begin{aligned}\hat{\mu}^{strategy} &= \frac{1}{T - \tau - \Delta t} \sum_{t=\tau}^{T-\Delta t} r_{t+\Delta t}^{strategy}, \\ \hat{\sigma}^{strategy} &= \left(\frac{1}{T - \tau - \Delta t - 1} \sum_{t=\tau}^{T-\Delta t} (r_{t+\Delta t}^{strategy} - \hat{\mu}^{strategy})^2 \right)^{1/2}, \\ \widehat{\text{SR}}^{strategy} &= \frac{\hat{\mu}^{strategy}}{\hat{\sigma}^{strategy}}.\end{aligned}$$

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 report also 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 ad-

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.

justment, 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 $w_{j,t}^{strategy}$ denote the portfolio weight in stock j chosen at time t for a particular strategy, $w_{j,t+}^{strategy}$ the portfolio weight *before* rebalancing but at $t + \Delta t$, and $w_{j,t+\Delta t}^{strategy}$ 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:

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

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 beat the benchmarks even after adjusting for transactions costs, we also 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,

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 $\hat{\mu}_i/\hat{\sigma}_i - \hat{\mu}_n/\hat{\sigma}_n$, 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 the 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: $\hat{\sigma}_i^2/\hat{\sigma}_n^2$, then a one-sided P-value for this null hypothesis is given by $\hat{p} = 1 - \hat{F}(1)$, and we will 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).

we compute the time series of net returns $\tilde{r}_{t+\Delta t}^{strategy}$ for a given strategy and the benchmark:

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

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.

5 Out-of-Sample Performance of Portfolios

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 5.1, by discussing the performance of the benchmark portfolios that do not use information from option prices. In Sections 5.2, we report the performance of portfolios obtained using option-implied volatilities. In 5.3, we report the performance of portfolios that use option-implied correlations. Finally, in Section 5.4 we report the improvement in out-of-sample portfolio performance from using option-implied quantities that predict 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.

5.1 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 two weeks. 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 parenthesis under each performance metric. Two p-values are reported in parenthesis under each performance metric: the first p-value is relative to the $1/N$ benchmark, and the second p-value in this table is relative to the “Sample-cov” benchmark. 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 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, the volatility of the $1/N$ portfolio is 0.2254 and that of the minimum-variance portfolio based on the sample-covariance matrix with daily rebalancing is 0.1429, for the minimum-variance portfolio with shortsale constraints it is 0.1383, and for the minimum-variance portfolio with shrinkage it is 0.1263, and for the portfolio obtained from setting all correlations equal to zero 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 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 zero, in the case of Sample 2; but, for both cases the differences are not statistically significant.¹⁹

Of the four minimum-variance portfolios that we consider, the shortsale-constrained portfolio and the portfolio obtained by setting all correlations equal to zero have a

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 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.

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 shortselling (“Sample cov,” “Shrinkage,” and “Zero correlation”) perform very poorly along all metrics. The mean-variance strategy with shortsale constraints achieves a lower volatility than the $1/N$ portfolio, but has a lower Sharpe ratio and higher turnover than the $1/N$ portfolio and also the shortsale-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 shortsale-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.

5.2 Performance of Portfolios Using Option-Implied Volatility

Motivated by the findings in Section 3.1 about the predictive power of model-free implied volatilities after correction for the risk premium, \widehat{RV} , we use them in $\text{diag}(\hat{\sigma})$ to obtain the covariance matrix given in (9); that is, $\hat{\Sigma} = \text{diag}(\widehat{RV}) \hat{\Omega} \text{diag}(\widehat{RV})$. 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 shortsales are constrained, where shrinkage is applied to this covariance matrix, and where we impose the restriction that all correlations are equal to zero. In computing these portfolios, we continue to use historical correlations (except for the last portfolio, where correlations are set equal to zero).

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 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 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

5.3 Performance of Portfolios With Option-Implied Correlation

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 (9), instead of historical correlations, is reported in Table 5. In order to isolate the effect of 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 shortsales constrained; and, the third, which is labeled “regularization,” replaces the “shrinkage” portfolio considered in the earlier tables. We use the regulariza-

²⁰The reason for the relatively poor Sharpe ratio and turnover of the other portfolios based on implied volatility is that implied volatility is highly variable over time, which increases the variability of portfolio weights and reduces the gains from having a better predictor of realized volatility.

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.

tion 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,b) that rely on particular distributional assumptions.

We observe from Table 5 that using the 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 shortsale 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.

Thus, we conclude that using the option-implied correlations does not lead to a significant improvement in portfolio performance. that the implied correlations almost fully eliminate the mean prediction error with respect to realized correlations. 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.

5.4 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 help forecast returns: the model-free implied volatility (MFIV); the volatility risk premium, measured as the spread between the currently observed Black-Scholes option-implied volatility and realized (historical) volatility (IRVS); model-free option-implied skewness (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,

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

described in Section 4.3.1, we adjust the returns of the top and bottom decile portfolios based on each of these characteristics. In the second, described in Section 4.3.2, we use the parametric-portfolio methodology of Brandt, Santa-Clara, and Valkanov (2009).

5.4.1 Performance of 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 shortsale constraints, considered in this table corresponding to the following four option characteristics: MFIV, IRVS, MFIS, 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 do 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 declines as the rebalancing frequency decreases. For example, in the case of 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 if the mean-variance portfolios using option-implied characteristics to forecast expected returns outperform the benchmark portfolios even in the presence of transactions costs, we compute the *equivalent transaction cost* for each portfolio. Recall from Section 4.4 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.

In Table 7, for each dataset and rebalancing frequency, we report two sets of numbers, each set having two numbers. The first set of numbers are for the case where the benchmark portfolio is the $1/N$ portfolio; the second set of numbers are for the case where the

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

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 datasets is about 10 basis points (see French (2008)), with the actual cost depending on the size of the trade and the execution capability of the trader.

In Table 7, there are two possibilities when comparing the performance of portfolios that use option-implied information with the 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 the portfolio using option-implied information can incur before its performance drops to the level of the benchmark, which is also considered net of transactions costs. The second possibility is that the benchmark portfolio performs better, *and* has a better turnover. In this case there is no positive level of transaction costs that equates the two strategies, and we indicate this in the table with the symbol “★”. The third possibility is where the performance of both portfolios, net of transaction costs, turns negative; 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 to 56 basis points; 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 transactions costs of up to 22 basis points. In Panels B and C, we see that as the rebalancing frequency decreases, 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 transactions costs of up to 66 basis points for weekly rebalancing, and up to 101 basis points 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.

5.4.2 Performance of 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 FF and MOM characteristics to adjust the portfolio weights; we do not allow for shortsales. Comparing the two benchmarks, we see that the $1/N$ portfolio has a 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, 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 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 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 Sample 1 and Sample 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, 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, and for fortnightly rebalancing it is significant for Sample 1.

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

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 basis points. This is because of the relatively higher turnover of the strategies using option-implied information. However, when we rebalance less frequently, then the benefits of using option-implied information survive higher transaction costs. For example, the results in Panel C of Table 9 for fortnightly rebalancing indicate that MFIS improves performance relative to the $1/N$ strategy even when we pay transaction costs of about 30 basis points. However, option-implied characteristics improve performance relative to the benchmark parametric portfolio for transaction costs of up to 10 basis points. Using all of the implied characteristics is beneficial in the presence of transaction costs of about 20 basis points with respect to the $1/N$ benchmark, and about 10 basis points for the parametric portfolio benchmark.

In summary, Table 9 suggests that in the presence of transactions 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 transactions costs incurred by these strategies.

6 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, Santa-Clara, and Valkanov (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 transactions 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.

Table 1: Return Predictability

In this table, we report the results of various stock characteristics: size (SIZE), book-to-market (BTM), momentum (MOM), model-free implied volatility (MFIV), model-free implied skewness (MFIS), call-put-volatility spread (CPVS), and implied-realized-volatility spread (IRVS) 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 period for a particular holding period (one week, two weeks, or one month). Below, we show the annualized mean holding return for each decile-based portfolio and in the parenthesis the p-value for the hypothesis that the mean return is not different from zero. 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	2 weeks	1 day	1 week	2 weeks
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)

Table 2: Minimum-Variance Portfolios Without Option-Implied Information

In this table, 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 shortsales constrained; “Shrinkage” is the minimum-variance portfolio where shrinkage has been applied to the sample covariance matrix using the Ledoit and Wolf (2004a,b) methodology; and “Zero Correlation” is the minimum-variance portfolio where all correlations are set equal to zero. We report a p-value in parenthesis 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	0.2632	0.6131	0.1429	0.2094	0.5130
	(0.00)	(0.84)		(0.00)	(0.82)	
Constrained	0.1367	0.4489	0.0552	0.1383	0.6558	0.0303
	(0.00)	(0.80)		(0.00)	(0.15)	
Shrinkage	0.1332	0.4943	0.2723	0.1263	0.2939	0.3046
	(0.00)	(0.63)		(0.00)	(0.74)	
Zero Correlation	0.1771	0.5512	0.0133	0.1899	0.5161	0.0125
	(0.00)	(0.77)		(0.00)	(0.35)	
<i>Panel B: Weekly rebalancing</i>						
1/N	0.1973	0.6045	0.0302	0.2127	0.5030	0.0325
Sample cov	0.1613	0.2893	0.6556	0.1475	0.2535	0.5961
	(0.00)	(0.90)		(0.00)	(0.84)	
Constrained	0.1305	0.4591	0.0630	0.1322	0.6768	0.0408
	(0.00)	(0.83)		(0.00)	(0.08)	
Shrinkage	0.1343	0.5257	0.3047	0.1326	0.3196	0.3724
	(0.00)	(0.64)		(0.00)	(0.77)	
Zero Correlation	0.1663	0.5697	0.0268	0.1782	0.5331	0.0276
	(0.00)	(0.76)		(0.00)	(0.23)	
<i>Panel C: Fortnightly rebalancing</i>						
1/N	0.1905	0.6081	0.0427	0.2055	0.5036	0.0462
Sample cov	0.1587	0.2439	0.6961	0.1506	0.2690	0.6678
	(0.00)	(0.93)		(0.00)	(0.83)	
Constrained	0.1262	0.4542	0.0696	0.1297	0.6726	0.0492
	(0.00)	(0.85)		(0.00)	(0.11)	
Shrinkage	0.1333	0.4910	0.3340	0.1372	0.3262	0.4303
	(0.00)	(0.70)		(0.00)	(0.77)	
Zero Correlation	0.1599	0.5778	0.0372	0.1719	0.5388	0.0390
	(0.00)	(0.71)		(0.00)	(0.21)	

Table 3: Mean-Variance Portfolios Without Option-Implied Information

In this table, 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 shortsales constrained; “Shrinkage” is the mean-variance portfolio where shrinkage has been applied to the sample covariance matrix using the Ledoit and Wolf (2004a,b) methodology; and “Zero Correlation” is the mean-variance portfolio where all correlations are set equal to zero. “ $1/N + \text{FFM}$ ” 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 a p-value in parenthesis 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 + 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 + 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 + FFM	0.1903 (0.48)	0.6749 (0.06)	0.0719	0.2075 (0.77)	0.5657 (0.08)	0.0623

Table 4: Minimum-Variance Portfolios Using Option-Implied Volatility

In this table, 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 shortsales 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 zero and historical volatility is replaced by option-implied volatility corrected for the volatility risk premium. We report two p-values in parenthesis, 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

Table 5: Minimum-Variance Portfolios Using Option-Implied Correlation

In this table, 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 (2011), 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 shortsales 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 parenthesis, 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
<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

Table 6: Mean-Variance Portfolios With Option-Characteristic-Adjusted Returns

In this table, we report the performance of the $1/N$ portfolio and various shortsale-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 4.3.1. We report two p-values in parenthesis, 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)	0.6785 (0.30)	0.0983	0.4523 (1.00)	0.4121 (0.73)	0.0983 (0.85)
Constrained + IRVS	0.2499 (1.00)	0.9232 (0.01)	0.2265	0.2634 (1.00)	0.7650 (0.01)	0.2097 (0.29)
Constrained + MFIS	0.2520 (1.00)	1.0092 (0.00)	0.2859	0.2681 (1.00)	1.0130 (0.00)	0.2701 (0.02)
Constrained + CPVS	0.2518 (1.00)	1.4291 (0.00)	0.6209	0.2726 (1.00)	1.3596 (0.00)	0.6088 (0.00)
<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)	0.5974 (0.53)	0.1245	0.4387 (1.00)	0.3657 (0.89)	0.1286 (0.94)
Constrained + IRVS	0.2567 (1.00)	0.7982 (0.03)	0.2454	0.2689 (1.00)	0.6216 (0.08)	0.2316 (0.65)
Constrained + MFIS	0.2549 (1.00)	0.8397 (0.01)	0.3023	0.2690 (1.00)	0.8063 (0.00)	0.2876 (0.20)
Constrained + CPVS	0.2472 (1.00)	0.9270 (0.00)	0.6338	0.2578 (1.00)	0.9058 (0.00)	0.6229 (0.06)
<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)	0.5850 (0.58)	0.1445	0.4158 (1.00)	0.3300 (0.94)	0.1513 (0.95)
Constrained + IRVS	0.2590 (1.00)	0.7070 (0.15)	0.2594	0.2714 (1.00)	0.5271 (0.39)	0.2479 (0.81)
Constrained + MFIS	0.2520 (1.00)	0.7867 (0.04)	0.3151	0.2701 (1.00)	0.7291 (0.00)	0.3010 (0.37)
Constrained + CPVS	0.2392 (1.00)	0.8132 (0.01)	0.6436	0.2515 (1.00)	0.7453 (0.00)	0.6337 (0.33)

Table 7: Equivalent Transaction Cost for Mean-Variance Portfolios With Option-Characteristic Adjusted Returns

In this table, we evaluate the effect of transactions 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 4.3.1. 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 transactions 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 shortsale-constrained minimum-variance portfolio. A “★” indicates that the portfolio using option-implied information performs worse than the benchmark *and* also has higher turnover than the benchmark. A “–” indicates that 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

Table 8: Parametric Portfolios Using Option-Implied Information

In this table, we evaluate the performance of parametric portfolios that rely on option-implied characteristics about expected returns, using the methodology proposed in Brandt, Santa-Clara, and Valkanov (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 parenthesis, 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

... Table 8 continued

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
1/N + FFM	0.1979	0.6725	0.0624	0.2156	0.5617	0.0508
	(0.66)	(0.05)		(0.94)	(0.08)	
	(0.50)	(0.50)		(0.50)	(0.50)	
1/N + MFIV	0.2046	0.6349	0.0880	0.2157	0.4965	0.0723
	(1.00)	(0.30)		(0.89)	(0.55)	
	(1.00)	(0.76)		(0.53)	(0.89)	
1/N + IRVS	0.1966	0.6270	0.0970	0.2144	0.5250	0.0938
	(0.23)	(0.12)		(0.94)	(0.13)	
	(0.25)	(0.84)		(0.30)	(0.81)	
1/N + MFIS	0.2045	0.6757	0.2064	0.2269	0.5884	0.1992
	(1.00)	(0.00)		(1.00)	(0.00)	
	(1.00)	(0.47)		(1.00)	(0.26)	
1/N + CPVS	0.1978	0.7034	0.2960	0.2167	0.6031	0.2956
	(0.91)	(0.00)		(1.00)	(0.00)	
	(0.46)	(0.23)		(0.75)	(0.13)	
1/N + ALL IMPLIED	0.2074	0.6915	0.3372	0.2215	0.6097	0.2869
	(1.00)	(0.01)		(1.00)	(0.00)	
	(1.00)	(0.33)		(1.00)	(0.11)	
1/N + FFM + MFIV	0.2042	0.6697	0.0973	0.2149	0.5732	0.0724
	(1.00)	(0.11)		(0.80)	(0.07)	
	(1.00)	(0.52)		(0.34)	(0.35)	
1/N + FFM + IRVS	0.1963	0.6700	0.1006	0.2141	0.5505	0.0712
	(0.31)	(0.04)		(0.79)	(0.11)	
	(0.01)	(0.56)		(0.02)	(0.78)	
1/N + FFM + MFIS	0.1990	0.7075	0.2044	0.2237	0.5760	0.1729
	(0.88)	(0.00)		(1.00)	(0.01)	
	(0.90)	(0.06)		(1.00)	(0.31)	
1/N + FFM + CPVS	0.2021	0.7538	0.3511	0.2209	0.6168	0.3068
	(1.00)	(0.00)		(1.00)	(0.00)	
	(1.00)	(0.01)		(1.00)	(0.04)	
1/N + FFM + ALL IMPLIED	0.2075	0.7600	0.4014	0.2235	0.6284	0.3185
	(1.00)	(0.00)		(1.00)	(0.00)	
	(1.00)	(0.01)		(1.00)	(0.01)	

... Table 8 continued

Strategy	Sample 1			Sample 2		
	std	sr	trn	std	sr	trn
<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	0.6749	0.0719	0.2075	0.5657	0.0623
	(0.48)	(0.06)		(0.77)	(0.08)	
	(0.50)	(0.50)		(0.50)	(0.50)	
1/N + MFIV	0.1979	0.6441	0.0972	0.2084	0.5023	0.0830
	(0.99)	(0.29)		(0.82)	(0.53)	
	(1.00)	(0.71)		(0.61)	(0.88)	
1/N + IRVS	0.1913	0.6226	0.1036	0.2085	0.5114	0.1020
	(0.74)	(0.20)		(0.98)	(0.33)	
	(0.64)	(0.87)		(0.66)	(0.90)	
1/N + MFIS	0.1986	0.6631	0.2103	0.2205	0.5585	0.2041
	(1.00)	(0.00)		(1.00)	(0.00)	
	(1.00)	(0.59)		(1.00)	(0.58)	
1/N + CPVS	0.1911	0.6627	0.2984	0.2094	0.5559	0.2985
	(0.92)	(0.00)		(1.00)	(0.00)	
	(0.61)	(0.60)		(0.79)	(0.59)	
1/N + ALL IMPLIED	0.2008	0.6662	0.3398	0.2137	0.5678	0.2901
	(1.00)	(0.05)		(1.00)	(0.00)	
	(1.00)	(0.58)		(0.99)	(0.49)	
1/N + FFM + MFIV	0.1973	0.6712	0.1051	0.2078	0.5665	0.0825
	(0.98)	(0.15)		(0.76)	(0.11)	
	(1.00)	(0.55)		(0.57)	(0.49)	
1/N + FFM + IRVS	0.1897	0.6681	0.1076	0.2068	0.5480	0.0808
	(0.36)	(0.07)		(0.70)	(0.13)	
	(0.22)	(0.67)		(0.23)	(0.91)	
1/N + FFM + MFIS	0.1926	0.7030	0.2084	0.2163	0.5484	0.1782
	(0.86)	(0.00)		(1.00)	(0.07)	
	(0.99)	(0.12)		(1.00)	(0.72)	
1/N + FFM + CPVS	0.1957	0.7165	0.3533	0.2146	0.5731	0.3097
	(1.00)	(0.00)		(1.00)	(0.00)	
	(1.00)	(0.11)		(1.00)	(0.41)	
1/N + FFM + ALL IMPLIED	0.2009	0.7286	0.4034	0.2166	0.5865	0.3212
	(1.00)	(0.00)		(1.00)	(0.00)	
	(1.00)	(0.06)		(1.00)	(0.24)	

Table 9: Equivalent Transactions Cost for Parametric Portfolios Using Option-Implied Information

In this table, we evaluate the effect of transactions costs on the mean return and Sharpe ratio of parametric portfolios that rely on option-implied characteristics, using the methodology proposed in Brandt, Santa-Clara, and Valkanov (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 transactions 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 + FFM” parametric portfolio. A \star indicates that 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 \star	0.04 \star	0.03 \star	0.02 \star
1/N + IRVS	0.05 \star	0.06 \star	0.05 \star	0.06 \star
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

... Table 9 continued

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

... Table 9 continued

Strategy	Sample 1		Sample 2	
	TC mean	TC sr	TC mean	TC sr
<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

References

- Ait-Sahalia, Y., and M. W. Brandt, 2008, "Consumption and Portfolio Choice with Option-Implied State Prices," NBER Working Paper.
- Andersen, T. G., T. Bollerslev, and F. X. Diebold, 2009, "Parametric and Nonparametric Measurement of Volatility," in Yacine Ait-Sahalia, and Lars Peter Hansen (ed.), *Handbook of Financial Econometrics*, North Holland, Amsterdam.
- Ang, A., T. G. Bali, and N. Cakici, 2010, "The Joint Cross Section of Stocks and Options," Working Paper.
- Bakshi, G. S., N. Kapadia, and D. B. Madan, 2003, "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options," *Review of Financial Studies*, 16, 101–143.
- Bali, T. G., and A. Hovakimian, 2009, "Volatility Spreads and Expected Stock Returns," *Management Science*, 55, 1797–1812.
- Black, F., and M. S. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–659.
- Bollerslev, T., M. Gibson, and H. Zhou, 2004, "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, D.C.
- Bollerslev, T., G. Tauchen, and H. Zhou, 2009, "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22, 4463–4492.
- Boyer, B. H., T. Mitton, and K. Vorkink, 2009, "Expected Idiosyncratic Skewness," forthcoming in *Review of Financial Studies*.
- Brandt, M. W., P. Santa-Clara, and R. Valkanov, 2009, "Parametric Portfolio Policies: Exploiting Characteristics in the Cross Section of Equity Returns," *Review of Financial Studies*, 22, 3411–3447.
- Buss, A., and G. Vilkov, 2011, "Measuring Equity Risk with Option-Implied Correlations," Working Paper, Goethe University Frankfurt.
- Carr, P., and L. Wu, 2009, "Variance Risk Premiums," *Review of Financial Studies*, 22, 1311–1341.
- Chan, L. K. C., J. Karceski, and J. Lakonishok, 1999, "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model," *Review of Financial Studies*, 12, 937–974.
- Corsi, F., 2009, "A Simple Approximate Long-Memory Model of Realized Volatility," *Journal of Financial Econometrics*, pp. 174–196.
- Cox, J. C., and C.-f. Huang, 1989, "Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process," *Journal of Economic Theory*, 49, 33–83.
- Cremers, M., and D. Weinbaum, 2010, "Deviations from Put-Call Parity and Stock Return Predictability," *Journal of Financial and Quantitative Analysis*, 45, 335–367.

- DeMiguel, V., L. Garlappi, F. Nogales, and R. Uppal, 2009, “A Generalized Approach to Portfolio Optimization: Improving Performance By Constraining Portfolio Norms,” *Management Science*, 55, 798–812.
- DeMiguel, V., L. Garlappi, and R. Uppal, 2009, “Optimal Versus Naive Diversification: How Inefficient Is the $1/N$ Portfolio Strategy?,” *Review of Financial Studies*, 22, 1915–1953.
- Driessen, J., P. Maenhout, and G. Vilkov, 2009, “The Price of Correlation Risk: Evidence from Equity Options,” *Journal of Finance*, 64, 1375–1404.
- Efron, B., and R. Tibshirani, 1993, *An Introduction to the Bootstrap*, Chapman and Hall, Boca Raton, FL.
- Elton, E. J., M. J. Gruber, and J. Spitzer, 2006, “Improved Estimates of Correlation Coefficients and their Impact on Optimum Portfolios,” *European Financial Management*, 12, 303–318.
- Fama, E. F., and K. R. French, 1992, “The Cross-Section of Expected Stock Returns,” *Journal of Finance*, 47, 427–465.
- French, K. R., 2008, “Presidential Address: The Cost of Active Investing,” *Journal of Finance*, 63, 1537–1573.
- Goyal, A., and A. Saretto, 2009, “Cross-Section of Option Returns and Volatility,” *Journal of Financial Economics*, 94, 310–326.
- Grundy, B. D., and S. Martin, 2001, “Understanding the nature of the risks and the source of the rewards to momentum investing,” *Review of Financial Studies*, 14, 29–78.
- Guidolin, M., and A. Timmermann, 2008, “International Asset Allocation under Regime Switching, Skew and Kurtosis Preferences,” *Review of Financial Studies*, 21, 889–935.
- Harvey, C. R., and A. Siddique, 2000, “Conditional Skewness in Asset Pricing Tests,” *Journal of Finance*, 55, 1263–1295.
- Jacobs, H., S. Muller, and M. Weber, 2010, “How Should Private Investors Diversify? An Empirical Evaluation of Alternative Asset Allocation Policies to Construct a “World Market Portfolio”,” Working Paper, University of Mannheim.
- Jagannathan, R., and T. Ma, 2003, “Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps,” *Journal of Finance*, 58, 1651–1684.
- Jegadeesh, N., and S. Titman, 1993, “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *Journal of Finance*, 48, 65–91.
- Jobson, J. D., R. Korkie, and V. Ratti, 1979, “Improved Estimation for Markowitz Portfolios Using James-Stein Type Estimators,” *Proceedings of the American Statistical Association*, 41, 279–292.
- Jorion, P., 1986, “Bayes-Stein Estimation for Portfolio Analysis,” *Journal of Financial and Quantitative Analysis*, 21, 279–292.

- Kostakis, A., N. Panigirtzoglou, and G. Skiadopoulos, 2011, "Market Timing with Option-Implied Distributions: A Forward-Looking Approach," *Management Science*, pp. 1231–1249.
- Kraus, A., and R. H. Litzenberger, 1976, "Skewness Preference and the Valuation of Risk Assets," *Journal of Finance*, 31, 1085–1100.
- Ledoit, O., and M. Wolf, 2004a, "Honey, I Shrunk the Sample Covariance Matrix," *Journal of Portfolio Management*, 30, 110–119.
- Ledoit, O., and M. Wolf, 2004b, "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices," *Journal of Multivariate Analysis*, 88, 365–411.
- Ledoit, O., and M. Wolf, 2008, "Robust Performance Hypothesis Testing with the Sharpe Ratio," *Journal of Empirical Finance*, 15, 850–859.
- Merton, R. C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.
- Merton, R. C., 1973, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.
- Newey, W. K., and K. D. West, 1987, "A simple, positive-semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica*, 55, 703–708.
- Pástor, Ľ., 2000, "Portfolio Selection and Asset Pricing Models," *Journal of Finance*, 55, 179–223.
- Politis, D., and J. Romano, 1994, "The Stationary Bootstrap," *Journal of the American Statistical Association*, 89, 1303–1313.
- Poon, S.-H., and C. W. J. Granger, 2005, "Practical Issues in Forecasting Volatility," *Financial Analysts Journal*, 61, 45–56.
- Rehman, Z., and G. Vilkov, 2009, "Risk-Neutral Skewness: Return Predictability and Its Sources," Working Paper, Goethe University Frankfurt.
- Rubinstein, M. E., 1973, "The Fundamental Theorem of Parameter-Preference Security Valuation," *Journal of Financial and Quantitative Analysis*, 8, 61–69.
- Vanden, J. M., 2008, "Information Quality and Options," *Review of Financial Studies*, 21, 2635–2676.
- Xing, Y., X. Zhang, and R. Zhao, 2010, "What Does Individual Option Volatility Smirk Tell Us About Future Equity Returns?," *Journal of Financial and Quantitative Analysis*, 45, 641–662.
- Zumbach, G., 2009, "Inference on Multivariate ARCH Processes with Large Sizes," *SSRN eLibrary*.