Improving Portfolio Selection Using Option-Implied Volatility and Skewness*

Victor DeMiguel‡ Yuliya Plyakha§ Raman Uppal¶ Grigory Vilkov§

This version: February 27, 2011

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 four metrics: volatility, Sharpe ratio, certainty-equivalent return, and turnover. Our empirical evidence shows that using option-implied volatility helps to reduce portfolio volatility, but does not improve the Sharpe ratio or certainty-equivalent return; option-implied correlation does not improve any of the metrics; however, expected returns estimated using information in the volatility risk premium and option-implied skewness increase substantially both the Sharpe ratio and certainty-equivalent return, 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

*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, Lorenzo Garlappi, Nicolae Gârleanu, Rene Garcia, 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), 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 and Western Finance Association.

‡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-Pf H 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.

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 returns data 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 in the literature 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.[3] 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, as in Bali and Hovakimian (2009), Goyal and Saretto (2009), and Cremers and Weinbaum (2010), but also using the model-free approach developed in Bakshi, Kapadia, and Madan (2003) and Carr and Wu (2009).

First, we consider the use of option implied volatilities and correlations to improve out-of-sample performance of mean-variance portfolios. 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.[4] 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 of the covariance matrix (as in Ledoit and Wolf (2004a,b)),

[1]For evidence of this poor performance, see DeMiguel, Garlappi, and Uppal (2009) and the references therein.

[2]These approaches include: imposing a factor structure on returns (Chan, Karceski, and Lakonishok, 1999; MacKinlay and Pástor, 2000), using data for daily rather than monthly returns (Jagannathan and Ma, 2003), using Bayesian methods (Jobson, Korkie, and Ratti, 1979; Jorion, 1986; Pástor, 2000; Pástor and Stambaugh, 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).

[3]For other examples of the use of option-implied volatility and skewness, see Christoffersen and Chang (2009), who use implied volatility and skewness to forecast future realized betas.

[4]Jagannathan 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.

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 the minimum-variance portfolios based on the risk-premium-corrected implied volatilities attain an out-of-sample portfolio volatility that is about 5% lower than the traditional portfolios based on the historical stock-return data, while the improvement in Sharpe ratio and certainty-equivalent return is insignificant.[5] The main reason why option-implied volatility fails to improve out-of-sample portfolio performance substantially is that, being based exclusively on current option prices, it is an estimator with high variance.

Next, we examine the use of option-implied correlations and find that these do not lead to any improvement in performance. To understand the reason for this, note that a covariance matrix that improves portfolio performance will be one that contains enough information about future covariances *and* is stable (with a small condition number and, therefore, leading to less extreme portfolio weights). Our empirical results indicate that, while option-implied correlations are better than their historical counterparts at forecasting the future realizations of these moments, the gains are not substantial enough to offset the loss from the increased instability of the covariance matrix, the effect of which is reflected in the much higher turnover of portfolios based on these estimates.

Finally, we consider the use of the risk premium for stochastic volatility and option-implied skewness to improve the out-of-sample performance of mean-variance portfolios. Because these quantities have been shown in the literature to help forecast expected returns,[6] it makes sense to explore their effect in the framework of mean-variance portfolios. But it is well known that the weights of the classical mean-variance portfolio are very sensitive to errors in estimates of expected returns, and this is confirmed in our experiments.[7] Recently, Brandt, Santa-Clara, and Valkanov (2009) have developed "parametric portfolios" to address this problem and so we apply this methodology to mean-variance portfolios in order to evaluate the gains from using

[5]Because option-implied volatilities estimate the risk-neutral volatilities, they are biased estimators of the real-world (objective) volatilities that are needed for portfolio optimization, with the gap between the two being the volatility risk premium, as explained in Chernov (2007).

[6]See, for instance, Bali and Hovakimian (2009), Goyal and Saretto (2009), Xing, Zhang, and Zhao (2009), and Cremers and Weinbaum (2010).

[7]It is well known that it is much more difficult to estimate expected returns than second moments of stock returns (Merton, 1980), and also that mean-variance portfolio weights are more sensitive to errors in estimates of expected returns than second moments.

option-implied information to forecast stock returns in the presence of transactions costs and shortsale constraints.

We find that the volatility risk premium and implied skewness improve the performance in terms of both Sharpe ratio and certainty equivalent return (net of transactions costs) of mean-variance portfolios constructed using the parametric-portfolio methodology, over and above the gains obtained from using the "size" and "value" characteristics identified in Fama and French (1992), and "momentum" identified in Jegadeesh and Titman (1993). The use of option-implied characteristics leads to higher turnover, but even after adjusting for transactions costs and prohibiting shortsales, mean-variance portfolios based on option-implied characteristics outperform portfolios that ignore these characteristics.

Our analysis is carried out in a comprehensive fashion. We consider two data sets: with 100 assets and 561 assets; two data frequencies: daily and intraday; three portfolio rebalancing (holding) periods: daily, weekly, and monthly; four performance metrics: portfolio volatility, Sharpe ratio, certainty-equivalent return, and turnover; and eight benchmark portfolios: the "$1/N$" equally-weighted portfolio; sample-based mean-variance portfolio; minimum-variance portfolio based on the sample covariance matrix; shortsale-constrained minimum-variance portfolio; minimum-variance portfolio with shrinkage of the covariance matrix; minimum-variance portfolio with correlations set equal to zero; minimum-variance portfolio with correlations set equal to the mean correlation across all asset pairs; and mean-variance portfolios implemented using the parametric-portfolio methodology developed in Brandt, Santa-Clara, and Valkanov (2009).

We conclude this introduction by discussing the relation of our work to the existing literature. The idea that option prices contain information about future asset returns has been understood ever since the work of Black and Scholes (1972) and Merton (1973).[8] The focus of our work is to investigate how the information implied by option prices can be used to improve

[8]For example, Latane and Rendleman (1976), Lamoureux and Lastrapes (1993), and Christensen and Prabhala (1998) find that implied volatility outperforms historical volatility in forecasting future volatility, and Poon and Granger (2005) provide a comprehensive survey of this literature. Bakshi, Kapadia, and Madan (2003) explain how one can use option prices to infer also higher moments of the return distribution, such as skewness. Driessen, Maenhout, and Vilkov (2009) show, in the working paper version of their article, how one can also obtain implied correlations from the prices of options on individual stocks and on the index, while Bali and Hovakimian (2009), Bollerslev, Tauchen, and Zhou (2008), Cremers and Weinbaum (2010), Goyal and Saretto (2009), Rehman and Vilkov (2009), and Xing, Zhang, and Zhao (2009) show that options can also be used to forecast future returns of the underlying asset. Of course, one can extract not just particular moments of returns, but also the probability distribution function, as shown by Jackwerth and Rubinstein (1996), Aït-Sahalia and Lo (1998), Jackwerth (2000), Bliss and Panigirtzoglou (2004), Panigirtzoglou and Skiadopoulos (2004), and Benzoni (1998), while Chernov and Ghysels (2000) show how to estimate jointly both the objective measure and the risk-neutral measure.

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. This paper finds that optimal consumption and portfolio rules based on option-implied information are different from those obtained using standard return dynamics; however, its focus is not on finding the optimal portfolio with superior out-of-sample performance. The second, which is by Kostakis, Panigirtzoglou, and Skiadopoulos (2009), studies the *asset-allocation* problem of allocating wealth between the S&P500 index and a riskless asset. The paper uses options on the index to first back out the implied risk-neutral distribution of returns and then transforms this to the objective distribution. This paper finds that the out-of-sample performance of the portfolio based on this distribution 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; in particular, we consider portfolios with 100 stocks and 561 stocks. It is not clear how one would extend the methodology of Kostakis, Panigirtzoglou, and Skiadopoulos (2009) 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. The robustness checks we undertake are described in Section 6, and we conclude in Section 7. Appendix A explains the construction of model-free option-implied moments; Appendix B explains how to compute variances and covariances for high-frequency intraday data; and Appendix C describes the method used for shrinkage and regularization of the covariance matrix.

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) and NYSE's Trades-And-Quotes (TAQ) database. To implement the parametric-portfolio methodology, we also use data from Compustat. Our data for options are from IvyDB (OptionMetrics).

2.1 Data on Stock Returns

Our sample period is January 3, 1995, to June 29, 2007. 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 is from the daily file of the CRSP, and we have in our sample a total of 3,146 trading days. We also use high-frequency *intraday* stock-price data consisting of transaction prices of the S&P500 constituents; these data are from the NYSE's Trades-And-Quotes database. We use the intraday data because several studies have highlighted the advantage of using high-frequency data to measure volatility of financial returns, and also as a robustness check for the results obtained from daily data.[9]

To improve the quality of the raw data used in our analysis, we apply the following filters and data-cleaning rules. For the daily stock returns of the S&P500 constituents from the CRSP daily file, we remove the observations with standard missing codes (SAS missing codes A,B,C,D, and E) as described in the Wharton Research Data Services documentation on CRSP. For the intraday stock-price raw data, we filter data for each day from the official opening at 9:30 EST until 16:00 EST, and delete entries with: a bid, ask, or transaction price equal to zero; corrected trades (trades with a correction indicator, "corr" $\neq 0$); an abnormal sale condition (trades where the variable "cond" has a letter code, except for "E" and "F")[10]; and prices that are above the ask plus the bid-ask spread or prices that are below the bid minus the bid-ask spread. See Barndorff-Nielsen, Hansen, Lunde, and Shephard (2009) for the details and discussion of these rules.[11] After cleaning the data, we construct a regularly spaced one-minute price grid for every trading day using the volume-weighted average of all transactions within a given minute. If there is no price for a given minute, we fill it in with the previous available price.

[9]For a survey of the literature on using high-frequency data to estimate moments of asset returns, see Andersen, Bollerslev, and Diebold (2009).

[10]See the TAQ 3 Users Guide for additional details about sale conditions.

[11]See rules P1, P2, T1, T2 and an adjusted version of T4.

Counting by IvyDB (OptionMetrics) identifiers, we have data for a maximum of 810 stocks, from which we choose those stocks for which at least 2,000 records of intraday volatilities and model-free implied volatilities are available, which gives us 561 stocks. Of these 561 stocks, there are 219 stocks for which the intraday volatilities and model-free implied volatilities are available for the *entire* time series. For robustness, we consider two datasets in our analysis. The first consists of the entire 561 stocks,[12] and the second consists of 100 stocks out of the 219 for which data are available for all dates; to select these 100 stocks, we first order the 219 stocks with respect to the security identifier code of IvyDB, and then select the first 100.[13]

2.2 Data on Stock Options

For stock options we use IvyDB that contains data on all U.S.-listed index and equity options. We use data from January 4, 1996, to June 29, 2007.[14]

While we do not use option *prices* directly in our analysis, we wish to use option-based information to obtain the moments of the option-implied distributions, and for this reason it is important for us to have the maximum number of options for a given maturity. Therefore, we choose for our analysis not the raw data on prices of options, but the volatility surface file, which contains a smoothed implied-volatility surface for a range of standard maturities and a set of option delta points.[15]

From the surface file we select for our sample 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 standard maturities of 30 and 60 days, which we consider to be the most suitable.[16] For each date, underlying stock, and time to maturity, we have from the surface data 13 implied volatilities, which are then used to calculate the moments of the risk-neutral distribution.[17] Some of the option-based characteristics also

[12]At each point in time, we consider only those stocks that have no missing data, which means that this sample has a variable number of stocks; on average, there are about 400 stocks at each point in time.

[13]In addition to the reported results, we have also checked our results on different subsamples of 50 and 100 stocks out of the 219 for which data are available for all dates, and these subsamples deliver similar results; details of this are provided in Section 6.1.

[14]Note that our data for stocks start in 1995, but we need 750 data points to compute the covariance matrix, so our portfolio optimization starts only at the beginning of 1998.

[15]We calculated implied moments also from the raw data on option prices, and the results are similar.

[16]The 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 is that selecting options that are out of the money reduces the effect of the premium for early exercise for these American options.

[17]There are 13 implied volatilities given for standard delta points for each call and put. For puts, these 13 deltas are $\{-0.80, -0.75, -0.70, -0.65, -0.60, -0.55, -0.50, -0.45, -0.40, -0.35, -0.30. - 0.25, -0.20\}$, and for calls the delta points are the same, but positive. We select calls with a delta less than or equal to 0.5 and for puts greater than -0.5, which gives a total of 13 implied volatilities for out-of-the-money options—a mix of calls and puts.

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 closest to 0.5.

2.3 Data on Stock Characteristics

For our analysis of mean-variance portfolios using the parametric-portfolio methodology of Brandt, Santa-Clara, and Valkanov (2009), 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, we first use the Compustat Quarterly Fundamentals file (from 1994 to 2008) to calculate the book value of equity, which is total assets (ACTQ) minus liabilities (LCTQ) minus preferred/preference stock — redeemable (PSTKRQ) plus deferred taxes and investment tax credit (TXDITCQ), and then divide the book equity by the market value of equity computed earlier. The three-month momentum characteristic is measured using daily returns data from CRSP. To get better distributional properties of the constructed characteristics, we take the logarithm of size and value characteristics.[18]

3 Option-Implied Information

In this section, we explain how we compute the option-implied moments that we use for portfolio selection, and compare their ability to forecast the actual realized moments to that of the moments based on historical return data. We consider the following measures: (i) model-free option-implied volatility and model-free volatility risk premium; (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.

[18] 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 the 3rd percentile to all values below the 3rd percentile and do the same for values higher than the 97th percentile. And we normalize all characteristics to have zero mean and unit standard deviation.

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.[19] In contrast to the model-specific Black-Scholes 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.

Model-free implied volatility is given by a single number and it subsumes information in the whole Black-Scholes implied volatility smile. Theoretical and empirical research (see Jiang and Tian (2005) and Vanden (2008)) finds that model-free implied volatility is better at predicting the future realized volatility than the Black-Scholes implied volatility, and it is used by the CBOE to compute VIX, which is the ticker symbol for the CBOE Volatility Index that gives the implied volatility of S&P500 index options. To compute the model-free implied volatility, we first calculate the option prices from the interpolated volatility surface data. We then use these prices to find the value of the "variance contract," following the approach in Bakshi, Kapadia, and Madan (2003); the formula for the variance contract and the procedure used to compute it is provided in Appendix A.[20] The square root of the variance contract then gives the model-free implied volatility.

To confirm the intuition that the model-free implied volatility is a better predictor of realized volatility in the future, relative to using the volatility estimate based on historical returns, we regress realized volatility on the model-free implied volatility and compare the root mean squared error (RMSE) and R^2 to that when volatility based on historical data is used as a predictor. The prediction regression we estimate is $\mathrm{RV} = \alpha + \beta\,\widehat{\mathrm{RV}}$, where we regress the 30-days realized volatility (RV) on various volatility predictors ($\widehat{\mathrm{RV}}$). The volatility predictors we consider are: historical daily volatility (based on the past 250 days for the 100-stock sample and the past 750 days for the 561-stock sample), historical intraday volatility (based on the past 30 days), model-free implied volatility, and model-free implied volatility adjusted for the volatility risk premium. We compute the RMSE under the restrictions $\alpha = 0$ and $\beta = 1$. We see from

[19]Note 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, Engle (1982), Bollerslev (1986), Engle (2002), and the survey articles by Bollerslev, Chou, and Kroner (1992), Engle (1993), Poon and Granger (2005), Andersen, Bollerslev, Christoffersen, and Diebold (2006), and Andersen, Bollerslev, and Diebold (2009).

[20]For a discussion of how to compute the model-free implied volatility, see also Dumas (1995), Carr and Madan (1998, 2001), and Britten-Jones and Neuberger (2000).

Panel A of Table 1 that when regressing the 30-day realized volatility in the future on (i) 750-day historical daily volatility, (ii) 30-day intraday historical volatility, and (iii) model-free implied volatility, the R^2 for the model-free implied volatility is higher than that for intraday historical volatility, which is higher than for daily historical volatility. This is true for both the dataset with 100 stocks and that with 561 stocks. For example, in the case of the data with 561 stocks, the R^2 for historical daily volatility is 17.91%, for intraday historical volatility is 31.16%, and for model-free implied volatility is 40.52%. Also, the RMSE for implied volatility is smaller than that for historical daily volatility and historical intraday volatility.

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 make a correction to the model-free implied volatility in order to get the volatility under the objective measure.

The difference between the model-free implied volatility and the expected volatility is the *volatility risk premium*. Bollerslev, Gibson, and Zhou (2004), Carr and Wu (2009), and others have shown that one can use the realized volatility (RV), instead of the expected volatility, to estimate the volatility risk premium. Assuming that the magnitude of the variance risk premium is proportional to the level of the variance under the actual probability measure (as it is in the Heston (1993) model), we estimate the historical volatility risk premium (HVRP) for a particular stock as the square root of the average variance risk premium for that stock for the past T trading days:[21]

$$\mathrm{HVRP}_t^2 = \frac{1}{T - \Delta t} \sum_{i=t-T+1}^{t-\Delta t} \frac{\mathrm{MFIV}_{i,i+\Delta t}^2}{\mathrm{RV}_{i,i+\Delta t}^2}. \tag{1}$$

In our analysis, we estimate the historical volatility risk premium on each day over the past year (-252 days to -21 days) using the model-free implied volatility and realized volatility, 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 (1), one can obtain the *prediction* of the future realized volatility, $\widehat{\mathrm{RV}}_t$:

$$\widehat{\mathrm{RV}}_{t,t+\Delta t} = \frac{\mathrm{MFIV}_{t,t+\Delta t}}{\mathrm{HVRP}_t}. \tag{2}$$

[21]Note that because HVRP_t^2 is calculated as the average of the ratio of $\mathrm{MFIV}_{i,i+\Delta t}^2$ and $\mathrm{RV}_{i,i+\Delta t}^2$, both of which are calculated over Δt days, as a result we will have only $T - \Delta t$ observations when computing the average.

Panel A of Table 1 shows that for the data with 561 stocks the R^2 for the regression of the risk-premium-corrected implied volatility is equal to 40.22%, which is about the same as the R^2 for the model-free implied volatility (40.52%), suggesting that there is no additional improvement in predictive ability from the risk premium correction; however, the risk-premium-corrected implied volatility is expected to have smaller bias with respect to the realized volatility, which can be seen from its lower RMSE and also by comparing the time series for the different volatility measures in Figure 1, where we plot the historical volatility based on the last 250 days (solid blue line), historical volatility based on the last 750 days (dot-dashed blue line), model-free implied volatility (dashed red line), risk-premium-corrected model-free implied volatility (solid pink line), and the 30-day realized volatility (thick black line). The figure is based on the cross-sectional equally-weighted average volatilities across the 561 stocks at each point in time. The figure shows that the risk-premium-corrected model-free implied volatility tracks realized volatility quite closely. The model-free implied volatility (without any risk-premium correction) tracks the realized volatility, but there is a distinct gap between the two. And the historical 750-day realized volatility does not track realized volatility very closely. Observe also that the variability of each of these volatility series is quite different.

3.2 Predicting Correlations Using Options

The second piece of option-implied information that we consider is implied correlation. Note that if a portfolio is composed of N individual stocks with weights w_i, $i = \{1, \ldots, N\}$, we can write the variance of the portfolio, σ_p^2, as follows:

$$\sigma_p^2 = \sum_{1=1}^{N} w_i^2 \sigma_i^2 + \sum_{i=1}^{N} \sum_{j \neq i}^{N} w_i w_j \sigma_i \sigma_j \rho_{ij}. \tag{3}$$

This equality holds under the objective probability measure P, and also under the risk-neutral probability measure Q. Hence, we can rewrite

$$\left(\sigma_p^Q\right)^2 = \sum_{i=1}^{N} w_i^2 \left(\sigma_i^Q\right)^2 + \sum_{i=1}^{N} \sum_{j \neq i}^{N} w_i w_j \sigma_i^Q \sigma_j^Q \rho_{ij}^Q. \tag{4}$$

The volatilities under the risk-neutral measure Q can be computed from the observed option prices as the model-free implied volatilities (MFIV), or as Black-Scholes implied volatilities (IV). Once we substitute into the equation above the implied volatilities, we have one equation and $N \times (N-1)/2$ unknown Q-correlations, ρ_{ij}^{Q}. Thus, to compute all pairwise correlations under the Q measure we need to make some identifying assumptions. We explain below two approaches that can be used to compute implied correlations.

One approach, which is adopted in Driessen, Maenhout, and Vilkov (2009), is to assume that all pairwise correlations are the same: $\rho_{ij}^{Q} = \rho^{Q}$. This assumption gives us a "homogeneous" implied-correlation matrix (HOMIC), with the correlation identified from Equation (4):

$$\rho^{Q} = \frac{\left(\sigma_{p}^{Q}\right)^{2} - \sum_{i=1}^{N} w_i^2 \left(\sigma_i^{Q}\right)^{2}}{\sum_{i=1}^{N} \sum_{j \neq i} w_i w_j \sigma_i^{Q} \sigma_j^{Q}}. \tag{5}$$

The resulting homogeneous implied-covariance matrix is positive definite under mild restrictions on ρ^{Q}.

An alternative is to use the approach proposed in Buss and Vilkov (2008), who compute a "heterogeneous" implied-correlation matrix (HETIC). It is well known that empirically volatilities and correlations are stochastic, and that both second moments are typically higher under the Q measure than under the objective P measure. The difference stems from the volatility and correlation risk premiums, respectively, and we can write the correlation risk premium CRP_{ij} for the pair of stocks i and j:

$$\text{CRP}_{ij} = \rho_{ij}^{Q} - \rho_{ij}^{P}. \tag{6}$$

Substituting (6) into (4) gives:

$$\left(\sigma_{p}^{Q}\right)^{2} = \sum_{i=1}^{N} w_i^2 \left(\sigma_i^{Q}\right)^{2} + \sum_{i=1}^{N} \sum_{j \neq i}^{N} w_i w_j \sigma_i^{Q} \sigma_j^{Q} \left(\rho_{ij}^{P} + CRP_{ij}\right). \tag{7}$$

Buss and Vilkov (2008) assume that all pairwise correlation risk premiums are driven by a single factor, ψ^{Q}, so that

$$\text{CRP}_{ij} = \left(1 - \rho_{ij}^{P}\right) \times \psi^{Q}. \tag{8}$$

Using this assumption, we can identify the heterogeneous implied-correlation matrix (HETIC) by first determining ψ^{Q} from (7):

$$\psi^{Q} = \frac{\left(\sigma_{p}^{Q}\right)^{2} - \sum_{i=1}^{N} w_i^2 \left(\sigma_i^{Q}\right)^{2} - \sum_{i=1}^{N} \sum_{j \neq i}^{N} w_i w_j \sigma_i^{Q} \sigma_j^{Q} \rho_{ij}^{P}}{\sum_{i=1}^{N} \sum_{j \neq i}^{N} w_i w_j \sigma_i^{Q} \sigma_j^{Q} \left(1 - \rho_{ij}^{P}\right)}, \tag{9}$$

and then obtaining each pairwise correlation by substituting (8) into Equation (6),

$$\rho_{ij}^{Q} = \rho_{ij}^{P} + \left(1 - \rho_{ij}^{P}\right) \times \psi^{Q}, \tag{10}$$

where ψ^Q is obtained from (9). The resulting heterogeneous implied-covariance matrix is positive definite under mild restrictions on ψ^Q.

After computing implied correlation using the above approach, we examine whether it is superior at predicting realized correlation. To do this, we estimate the correlation prediction regression: $\text{corr} = \alpha + \beta\,\widehat{\text{corr}}$, where we regress the 30-days realized correlation on the following correlation predictors ($\widehat{\text{corr}}$): (i) 250-day (750-day) historical daily correlations for the dataset with 100 (561) stocks; (ii) 30-day intraday correlations; (iii) daily heterogeneous implied correlations; and (iv) intraday heterogeneous implied correlations. The RMSE is computed imposing the restriction $\alpha = 0$ and $\beta = 1$. We see from Panel B of Table 1 that for the data with 561 stocks, the R^2 for historical daily correlations is 4.60%, for intraday correlations is 6.55%, for daily heterogeneous implied correlations is 9.90%, and for intraday heterogeneous implied correlations is 9.33%; the results are similar for the data with 100 stocks. These results suggest that implied correlations are better than historical correlations (both daily and intraday) at predicting realized correlations, though the magnitude of the improvement in R^2 is smaller than it was for predicting realized volatilities.[22] Note, however, that RMSE is smallest for historical daily correlation.

In Figure 2, we plot the historical correlation based on the last 750 days (dashed blue line), implied correlation (solid red line), and 30-day realized correlations (thick black line). Just as in the figure for volatilities, the plot is based on the cross-sectional equally-weighted average of average correlations across 561 stocks. There are two observations about these series: first, implied correlation follows the level of realized correlation much more closely than historical correlation; second, implied correlation is much more volatile (that is, contains more noise) than realized correlation, while historical correlation is smoother (but contains less current information).

3.3 Predicting Returns Using Options

There are four option-based quantities that we use to predict returns; the first two measure the risk premium for stochastic volatility and the second two measure option-implied skewness. We

[22]We do not consider time-series models of dynamic conditional correlations, such as the ones proposed in Engle (2002), because it is difficult to estimate these models for the case where the number of assets is large. See also Footnote 19.

describe each of them and then undertake a simple empirical test to see if these characteristics indeed have significant power to predict returns for our sample.

The first option-based characteristic we use is the volatility risk premium, which is the difference between implied volatility and realized volatility. Bali and Hovakimian (2009) and Goyal and Saretto (2009) have documented that stocks with a large spread between implied volatility and realized volatility tend to outperform those with low spreads. The first approach we use to measure the volatility risk premium is model free, and has already been described above in Equation (1) of Section 3.1. We refer to this as the *historical* volatility risk premium (HVRP) because to measure it we use an average of the ratio of historical model-free implied volatility to realized (historical) volatility, rather than the difference between the *current* Black-Scholes implied volatility and realized volatility. We also compute the volatility risk premium 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).[23] We use the abbreviation IRVS to refer to this implied-realized volatility spread. Note that because IRVS is based on the *current* value of implied volatility rather than the historical average, IRVS is likely to be a more noisy measure of the volatility risk premium than HVRP.

The second characteristic we consider is option-implied skewness. The first measure of this is *model-free implied skewness*, which represents a nonparametric estimate of the risk-neutral stock-return skewness, and it is this skewness that gives rise to the Black-Scholes implied volatility smirk.[24] Rehman and Vilkov (2009) find that stocks with high option-implied skewness outperform stocks with low option-implied skewness.[25] Our calculation of the model-free implied skewness (MFIS) parallels that of the model-free implied volatility. We first calculate the option prices from the interpolated volatility surface data. We then use these prices to determine the model-free implied skewness, as in Bakshi, Kapadia, and Madan (2003); the formula for this and the procedure used to compute it are provided in Appendix A.

[23] This closely resembles the measure of Goyal and Saretto (2009), though in the their paper the horizon for computing the realized volatility is longer (one year, or 251 trading days).

[24] For 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).

[25] Some researchers (for example, Xing, Zhang, and Zhao (2009)) 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.

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)). From the put-call parity relation, we can see that this measure is related to skewness. Cremers and Weinbaum (2010) construct this measure as the difference in implied volatility between pairs of call and put options using options of all available strikes for a given maturity, and as a result their spread contains information about the volatility smile. In our analysis, we are already using the model-free implied skewness (MFIS) that represents the risk-neutral skewness, and hence, reflects the volatility smile. Because of this, we do not use options of multiple strikes to compute the call-put volatility spread, and instead follow the methodology of Bali and Hovakimian (2009) to compute the call-put volatility spread (CPVS) as the difference between the Black-Scholes implied volatilities of the one-month at-the-money call and put options.

In order to see if these four option-implied measures are useful as univariate predictors of cross-sectional returns for our sample, we look at the decile long-short portfolio returns 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 one month). In Table 2 we show the annualized returns for each portfolio, along with the p-value, based on the Newey and West (1987) autocorrelated-adjusted standard error 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) and book-to-market (BTM) studied in Fama and French (1992), and momentum (MOM) studied in Jegadeesh and Titman (1993).

Table 2 confirms that the conventional Fama-French characteristics (SIZE and BTM) predict returns in the expected direction. MOM is significant only for the sample of 100 stocks, and the sign indicates a negative relation between the characteristic and the cross-sectional return. More interestingly, most option-based characteristics also lead to significant returns on the long-short decile portfolios, consistent with evidence in the existing literature. The strongest results — in terms of the magnitude of returns and significance — are for the portfolios based on measures of implied skewness, model free implied skewness (MFIS) and the call-put implied volatility spread (CPVS), and as expected, high decile stocks outperform the low decile ones for these measures. The implied-realized-volatility spread (IRVS) is also positively and significantly related to returns, while the historical measure of volatility risk premium (HVRP) is not significant for either sample.[26] Looking across the three rebalancing periods, we see that the magnitude of the return decreases as the rebalancing period goes from daily to weekly, and from weekly to monthly.

4 Construction of Portfolios and Performance Metrics

In this section, we explain the construction of the various portfolios we consider and also the performance 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. The first is the equally-weighted ($1/N$) portfolio. The next four are variations of the minimum-variance portfolio: the unconstrained minimum-variance portfolio based on the sample covariance matrix; the shortsale-constrained minimum-variance portfolio; the unconstrained minimum-variance portfolio with shrinkage of the covariance matrix; and the unconstrained minimum-variance portfolio with all correlations set equal to zero. And, the last one is the mean-variance portfolio implemented using the parametric-portfolio methodology. In principle, one could consider also the classical sample-based mean-variance portfolio as a benchmark, but, as we discuss below (see Footnote 37), this performs much worse than the benchmark portfolios listed above, and so we do not report in the tables the results for the sample-based mean-variance portfolios. We conclude this section by explaining how the performance metrics are constructed.

4.1 Equally-Weighted Portfolio

For the "*equally-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) 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.[27]

[26] The main reason for the poor performance of historical volatility risk premium (HVRP) is that the sign of the relation changes in the middle of the sample.

[27] For a discussion of why capitalization-weighted portfolios may be inefficient, see Haugen and Baker (1991).

4.2 Minimum-Variance Portfolios

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 all expected returns to be equal 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.[28]

$$\min_{w} \quad w^\top \hat{\Sigma} w, \tag{11}$$

$$\text{s.t.} \quad w^\top e = 1, \tag{12}$$

where $w \in \mathbb{R}^N$ is the vector of portfolio weights invested in stocks, $\Sigma \in \mathbb{R}^{N\times N}$ is the covariance matrix, and $e \in \mathbb{R}^N$ is the vector of ones. The solution to the above problem is:

$$w_{min} = \frac{\hat{\Sigma}^{-1} e}{e^\top \hat{\Sigma}^{-1} e}. \tag{13}$$

Note that the covariance matrix $\hat{\Sigma}$ in (13) can be decomposed into volatility and correlation matrices,

$$\hat{\Sigma} = \text{diag}(\hat{\sigma})\, \hat{\Omega}\ \text{diag}(\hat{\sigma}), \tag{14}$$

where diag $(\hat{\sigma})$ denotes 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 (13) there are two quantities that need to be estimated: volatilities ($\hat{\sigma}$) and correlations ($\hat{\Omega}$). We will use information implied by prices of options to estimate both quantities.

Observe that the minimum-variance portfolio w_{min} in (13) is based on an estimate of the "*sample covariance*" matrix, that is, the historical volatilities and correlations. For *daily data* we compute the conventional sample estimators of (co)variances using data over the past 750 days for the data with 561 stocks and 250 days for the data with 100 stocks. For *intraday data* we use the filtered and calendar-time aligned transaction prices over the last 30 trading days to estimate the (co)variances.[29]

In the existing literature, several methods have been proposed to improve the out-of-sample performance of the minimum-variance portfolio based on the sample (co)variances. We consider

[28] The constraint that the portfolio weights sum to one is required because we consider the case without the risk-free asset, given our objective is to explore how to use option-implied information to select the portfolio of *risky* stocks.

[29] There are several issues that have to be addressed when estimating moments from intraday data; the approach we use is consistent with the "second-best" approach of Zhang, Mykland, and Aït-Sahalia (2005), and the details of our procedure are provided in Appendix B.

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. To compute these portfolio weights, we solve the problem in (11) subject to the constraint in (12), after imposing the additional constraint that all weights have to be non-negative.

The second approach we consider is the "*shrinkage*" portfolio, where we compute the minimum-variance portfolio weights *after* shrinking the covariance matrix.[30] First, the sample covariance matrix for daily data and intraday 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), where they show how one can compute the *optimal* shrinkage of the covariance matrix under certain assumptions about the distribution of returns. For intraday data, instead of shrinkage, we use the regularization approach of Zumbach (2009) because the distribution of intraday returns is different from that of daily returns and does not satisfy the assumptions of Ledoit and Wolf (2004a,b).[31]

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 four performance metrics when compared to portfolios obtained from the first method.

4.3 Mean-Variance Parametric Portfolios

The recent literature on empirical asset pricing (see, for example, Bali and Hovakimian (2009), Goyal and Saretto (2009), Xing, Zhang, and Zhao (2009), and Cremers and Weinbaum (2010)) 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 these variables in improving the performance of traditional

[30] We do not consider the norm-constrained approach of DeMiguel, Garlappi, Nogales, and Uppal (2009) because we already consider the shortsale-constrained and shrinkage portfolios, which are special cases of the norm-constrained portfolios.

[31] Details of the shrinkage and regularization methods we use are provided in Appendix C, and the results for intraday data are summarized in Section 6.2.

mean-variance portfolios. But, it is well known in the literature 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.[32] Therefore, we evaluate the performance of portfolios based on these option-implied quantities using the "parametric-portfolio methodology," which has been developed by Brandt, Santa-Clara, and Valkanov (2009) to deal explicitly with this problem. We do this by using the historical variance risk premium (HVRP), the 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 characteristics (size, value and momentum), to construct parametric portfolios based on mean-variance utility assuming a risk aversion of unity. This is in contrast to Brandt, Santa-Clara, and Valkanov (2009), who use power utility.

In the parametric portfolios, the weight of an asset is a function of its weight in the benchmark $1/N$ portfolio,[33] denoted by $\omega_{i,1/N}$ and, in addition, the sum of the weight of the stock within each factor multiplied by the factor loadings θ:

$$\omega_i = \omega_{i,1/N} + (\theta_1 \times x_{i,1} + \theta_2 \times x_{i,2} + \ldots)\,, \tag{15}$$

where θ_1 is the factor loading of the first characteristic, and $x_{i,1}$ is the first characteristic of stock i. To construct investable factors we follow the approach in Brandt, Santa-Clara, and Valkanov (2009) where all returns on each day are multiplied by the normalized characteristics on the previous day (to avoid a look-ahead bias), within the estimation window of 250 days. We then sum these for each date to get a time series of investable factors.[34]

Note that θ_1 is *not* asset-specific, but is the same for all assets in the portfolio. We choose the vector $\theta = (\theta_1, \theta_2, \ldots)$ optimally by maximizing the average daily mean-variance utility during the estimation period; we use 250 days of data to estimate the vector θ. Because it is difficult to short stocks, we constrain shortsales and we choose the factor loadings θ such that the projected weight of each stock for the next day $t + 1$, after optimizing at day t, is non-negative: $\omega_i = \omega_{i,1/N} + \theta_1\, x_{i,1} + \theta_2, x_{i,2} + ... \geqslant 0$.

[32] See, for example, Frost and Savarino (1986, 1988), Michaud (1989), Best and Grauer (1991), Chopra and Ziemba (1993), Broadie (1993), Litterman (2003), and DeMiguel, Garlappi, and Uppal (2009).

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

[34] We also use a second approach based on constructing characteristic deciles. In this approach, we sort all stocks with respect to the given characteristic, assign a value of $+1$ to the characteristic for all the highest decile stocks and -1 to all the lowest decile stocks (and 0 to the middle 80% of the stocks), then multiply all returns by these new "binary" characteristics on each date within the estimation window (250 days for the smaller sample of 100 stocks and 750 days for the larger sample of 561 stocks), and sum these for each date to get a time series of investable factors. The results for this approach are similar to the results that are reported.

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, $1/N$ + FF + MOM, where "FF" denotes the size and value characteristics identified in Fama and French (1992), and "MOM" denotes 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 FF and MOM characteristics with the option-implied characteristics: the historical volatility risk premium (HVRP), model-free implied skewness (MFIS), call-put implied volatility spread (CPVS), and implied-realized volatility spread (IRVS). Second, in order to study the *incremental* value of option-implied information over and above the FF and MOM characteristics, we also consider the effect of including these option implied characteristics *in addition* to the FF and MOM factors.

The out-of-sample performance of these benchmark portfolios is reported in Table 3 and discussed in Section 5.1, using a variety of metrics that are described next.

4.4 Portfolio-Performance Metrics

We evaluate performance of the various portfolios using four criteria. These are the (i) out-of-sample portfolio volatility (standard deviation); (ii) out-of-sample portfolio Sharpe ratio; (iii) out-of-sample portfolio certainty-equivalent return; and (iv) portfolio turnover (trading volume). The reason for using the certainty-equivalent return, in addition to the Sharpe ratio, is that the Sharpe ratio considers only the mean and volatility of returns, while the certainty-equivalent return considers also the higher moments of returns.

We use the following "rolling-horizon" procedure for computing the portfolio weights and evaluating their performance. First, we choose a window over which to perform the estimation. We denote the length of the estimation window by $\tau < T$, where T is the total number of returns in the dataset. For our experiments, we use an estimation window of $\tau = 750$ data points for the sample with 561 stocks, and $\tau = 250$ data points for the sample with 100 stocks, which for daily data corresponds to three years and one year, respectively.[35] Second, using the return data over the estimation window τ, we compute the various portfolios we wish to compare. Third, we repeat this "rolling-window" procedure for the next day, by including the data for the next day and dropping the data for the earliest day. We continue doing this until the end

[35] Because our samples consists of 561 stocks and 100 stocks, estimation window lengths shorter than $\tau = 750$ for the data with 561 stocks and shorter than $\tau = 250$ for the data with 100 stocks often give singularities in the covariance matrix.

of the dataset is reached. At the end of this process, we have generated $T-\tau$ portfolio-weight vectors for each strategy; that is, $w_t^{strategy}$ for $t=\tau,\ldots,T-1$ and for each strategy.

We consider three rebalancing intervals: daily, weekly, and monthly. For the weekly and monthly rebalancing interval, we find the new set of weights daily, but hold that portfolio for 7 or 30 calendar days (5 or 21 trading days), which corresponds to the average of 5 or 21 daily returns; the advantage of this approach for monthly rebalancing is that it is not sensitive to the particular day on which the portfolio is formed.

Following this "rolling horizon" methodology, holding the portfolio $w_t^{strategy}$ for one day (or for one week or one month) gives the *out-of-sample* return at time $t+1$: that is, $r_{t+1}^{strategy} = w_t^{strategy} r_{t+1}$, where r_{t+1} denotes the returns from t to $t+1$. After collecting the time series of $T-\tau$ returns, $r_t^{strategy}$, the out-of-sample mean, volatility ($\widehat{\sigma}$), Sharpe ratio of returns (SR), and certainty-equivalent return (ce) are:

$$\widehat{\mu}^{strategy} = \frac{1}{T-\tau}\sum_{t=\tau}^{T-1} r_{t+1}^{strategy}, \tag{16}$$

$$\widehat{\sigma}^{strategy} = \left(\frac{1}{T-\tau-1}\sum_{t=\tau}^{T-1}\left(r_{t+1}^{strategy}-\widehat{\mu}^{strategy}\right)^2\right)^{1/2}, \tag{17}$$

$$\widehat{\mathrm{SR}}^{strategy} = \frac{\widehat{\mu}^{strategy}}{\widehat{\sigma}^{strategy}}, \tag{18}$$

$$\widehat{\mathrm{ce}}^{strategy} = u^{-1}\left(\frac{1}{T-\tau}\sum_{t=\tau}^{T-1} u\left(r_{t+1}^{strategy}\right)\right), \tag{19}$$

where u denotes the power utility function with a relative risk aversion of unity, and the certainty-equivalent return (ce) is the riskless return that an investor is willing to accept instead of investing in the risky strategy.

To measure the statistical significance of the difference in the volatility, Sharpe ratio, and certainty-equivalent return 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 monthly rebalancing we make an additional adjustment, as in Politis and Romano (1994), to account for the autocorrelation arising from overlapping returns.[36]

[36]Specifically, consider two portfolios i and n, with μ_i, μ_n, σ_i, σ_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)

Finally, we wish to obtain a measure of portfolio turnover. 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 $t+1$, and $w_{j,t+1}^{strategy}$ the desired portfolio weight at time $t+1$ (after rebalancing). Then, turnover, which is the average percentage of wealth traded per rebalancing interval (daily, weekly, or monthly), is defined as the sum of the absolute value of the rebalancing trades across the N available stocks and over the $T-\tau-1$ trading dates, normalized by the total number of trading dates:

$$\text{Turnover} = \frac{1}{T-\tau-1} \sum_{t=\tau}^{T-1} \sum_{j=1}^{N} \left(\left| w_{j,t+1}^{strategy} - w_{j,t^+}^{strategy} \right| \right). \tag{20}$$

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, Sharpe ratio, or certainty equivalent return) of a given strategy with that for the benchmark strategy. To find this equivalent transaction cost, we proceed as follows: First, for each level of transaction cost, we compute the time series of net returns $\tilde{r}_t^{strategy}$ for a given strategy and the benchmark:

$$\tilde{r}_t^{strategy} = r_t^{strategy} - \sum_{j=1}^{N} \left| w_{j,t}^{strategy} - w_{j,t^+}^{strategy} \right| \times TC^{strategy,benchmark}. \tag{21}$$

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 benchmark strategy. We consider two benchmark strategies when computing this equivalent transaction cost: the $1/N$ portfolio and the parametric portfolio strategy that uses only the FF and MOM characteristics.

than that of the benchmark portfolio n, that is, $\text{H}_0 : \mu_i/\sigma_i - \mu_n/\sigma_n \leqslant 0$. 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 monthly returns. We choose $B = 10{,}000$ for both cases and the block size equal to the number of overlaps in a series, that is, 20 for monthly data. If $\hat{F}$ denotes the empirical distribution function of the B bootstrap pairs corresponding to $\widehat{\mu}_i/\widehat{\sigma}_i - \widehat{\mu}_n/\widehat{\sigma}_n$, then a one-sided P-value for the previous null hypothesis is given by $\hat{p} = \hat{F}(0)$, and we will reject it for a small $\hat{p}$. In a similar way, to test the hypothesis that the variance of the portfolio i is greater (worse) than the variance of the benchmark portfolio n, $\text{H}_0 : \sigma_i^2/\sigma_n^2 \geqslant 1$, 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).

5 Out-of-Sample Performance of Portfolios

In this section, we discuss the major 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 volatility 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 Table 3 we report the performance of the benchmark strategies, all of which do *not* use data on option prices: the $1/N$ portfolio, four variants of the minimum-variance benchmark portfolio ("Sample cov," "Constrained," "Shrinkage," and "Zero correlation"), and the mean-variance portfolio implemented using the parametric-portfolio methodology with the FF and MOM characteristics.[37] In Panel A, 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 a month. The p-value for the comparison with the $1/N$ benchmark is reported in parenthesis under each performance metric. Each p-value is for the *one-sided* null hypothesis that the portfolio being evaluated is *worse* 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 worse than the benchmark).

[37]We examined also three sample-based *mean*-variance policies, corresponding to the first three minimum-variance policies described above. The first mean-variance policy is based on the sample covariance matrix, the second has shortsale constraints, and the third is computed with shrinkage applied to the covariance matrix, as in Ledoit and Wolf (2004a,b). All three mean-variance portfolios perform very poorly along all metrics, and this is especially true for the portfolios that do not have shortsale constraints. For example, while the volatility of the minimum-variance portfolios is less than 0.1433 (see Panel A of Table 3), the volatility of the corresponding three mean-variance portfolios is always higher, and, for the case without shortsale constraints, is several times higher. Similarly, the Sharpe ratio of the shortsale-constrained mean-variance portfolios is less than that of the shortsale-constrained minimum-variance portfolios, and it is negative for the other mean-variance portfolios. Finally, the turnover of the mean-variance portfolios is substantially higher than that of the minimum-variance portfolios.

From Table 3, we see that, compared to the $1/N$ portfolio, most of the strategies based on the minimum-variance portfolio achieve significantly lower volatility ($\widehat{\sigma}$) out of sample. For example, in Panel A with results for "Daily rebalancing," we see that for the data with 561 stocks, the volatility of the $1/N$ portfolio is 0.1745 and that of the minimum-variance portfolio with daily data is 0.1333, for the minimum-variance portfolio with constraints it is 0.1161, and for the minimum-variance portfolio with shrinkage it is 0.1164. The portfolio obtained from setting all correlations equal to zero has a volatility of 0.1493.[38] The p-values indicate that the volatilities of the minimum-variance portfolios are significantly lower than that of $1/N$. The results for the data with 100 assets and in Panels B and C for "Weekly rebalancing and " "Monthly rebalancing" are similar.

However, the Sharpe ratio (sr), certainty-equivalent return (ce), 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 minimum-variance portfolios based on shrinkage for the case of the data with 100 stocks, and the minimum-variance portfolio obtained by setting all correlations equal to zero for the case of 561 stocks; but, for both cases the differences are not statistically significant.[39] The mean-variance parametric portfolio typically has a higher Sharpe ratio than the minimum-variance strategies, though for the data with 561 stocks the zero-correlation minimum-variance portfolio outperforms it for the case of "Monthly rebalancing" and the shrinkage portfolio outperforms it for the cases of "Weekly rebalancing" and "Monthly rebalancing."

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 turnover that is slightly higher than that of the $1/N$ portfolio but is lower than the turnover of the mean-variance parametric 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.

[38]The volatility of the mean-variance portfolio implemented using the parametric-portfolio methodology is higher than that of the minimum-variance portfolios, which is not surprising, given that the objective of this portfolio is not to minimize volatility.

[39]It 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.

5.2 Performance of Portfolios Using Option-Implied Volatilities

Motivated by the findings in Section 3.1 about the predictive power of model-free implied volatilities (MFIV), we use them in diag($\hat{\sigma}$) to obtain the covariance matrix given in (14); that is, we use as the covariance matrix: $\hat{\Sigma} = \text{diag}(\widehat{\text{MFIV}})\,\hat{\Omega}\,\text{diag}(\widehat{\text{MFIV}})$. Using this covariance matrix, and setting the expected returns to be equal across all stocks, we determine the portfolio in (13), 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 performance of these portfolios is reported in Table 4. This table, and also all the tables that follow, have 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 3. There are three striking conclusions from Table 4, all of which are *negative*. First, the reduction in portfolio volatility compared to the benchmark portfolios is relatively small and is significant in only 3 out of the 24 cases ("Constrained" portfolio in Panel A for 100 stocks, and the "Zero correlation" portfolios for 561 stocks in Panels B and C). Second, turnovers are substantially higher for the portfolios in Table 4 that are based on implied volatility, compared to those in Table 3, which rely only on historical estimates of volatility; because option-implied volatility is an estimator with higher variance than historical volatility, it generates substantial changes in portfolio positions. Finally, the Sharpe ratios and the certainty-equivalent returns are substantially lower for the portfolios in Table 4 that are based on implied volatility, compared to those in Table 3, which rely only on historical estimates of volatility.

One reason for these negative results is that the model-free implied volatility from options is a measure of the *risk-neutral* volatility, which is the sum of the volatility risk premium and expected volatility under the objective measure; thus, implied volatility is a biased estimator of expected volatilities under the objective distribution. Moreover, assuming the same level of expected volatility under the objective measure, the implied volatility is relatively higher for stocks with high volatility risk premium, than for stocks with low volatility risk premium. Hence, when we use risk-neutral implied volatilities, we underweight the stocks with high volatility risk premium (because they have a higher implied volatility) in comparison to the stocks with low volatility risk premium. Given the findings of Bali and Hovakimian (2009) and Goyal and

Saretto (2009) that stocks with high volatility risk premium have higher returns, this explains the reduction in the portfolio's realized return and, hence, its Sharpe ratio.

To find the optimal portfolios using volatility under the objective measure, we repeat the above analysis but with volatility of stock returns estimated from the model-free implied volatility *after* it is corrected for the risk premium. The results for this case are given in Table 5. Comparing the results in this table to those in Table 4, where implied volatility is not corrected for the risk premium, we see that using the risk-premium-corrected implied volatilities leads to a reduction in portfolio volatility, especially for the case of 561 assets.

Next, comparing the results in Table 5 to those in Table 3 for the benchmark portfolios that do not rely on option-implied information, we observe that the portfolios with the risk-premium-corrected implied volatilities attain a lower out-of-sample volatility, with the difference being around 0–7% for daily and weekly rebalancing, and around 1%–11% for monthly rebalancing, with the difference being greater for the data with 561 stocks.

More importantly, comparing Table 5 with Table 4, we see that there is a significant improvement in the Sharpe ratio, certainty-equivalent return, and turnover when using the risk-premium-corrected implied volatilities, which confirms our motivation for making this correction. However, for almost all the cases considered in Table 5, the $1/N$ portfolio has a better Sharpe ratio, certainty-equivalent return and turnover than the portfolios based on option-implied volatility (the exceptions are the Sharpe ratio for the data with 561 stocks for the "Zero correlation" portfolio in all three panels and the "Shrinkage" portfolio in Panel C, though the p-value is not significant for any of these cases). The reason for the relatively poor Sharpe ratio, certainty-equivalent return, and turnover of the portfolios based on implied volatility is that, with or without the risk-premium correction, implied volatility is highly variable over time (see Figure 1), which increases the variability of portfolio weights and reduces the gains from having a better predictor of realized volatility.

5.3 Performance of Portfolios Using Option-Implied Correlations

We now investigate the benefit of using option-implied correlations in portfolio selection. In order to isolate the effect of using implied correlations on portfolios, when computing the portfolio weights we use volatilities from historical data. Recall that in Section 3.2 we described two approaches for estimating option-implied correlations. In the first approach, we assumed that all pairwise correlations are the same and equal to ρ^Q, which is given by Equation (5); the performance of portfolios obtained from this approach is reported in Table 6. In the second approach, we allow for different correlations between assets, which are given by Equation (10); the performance of portfolios obtained from this approach is reported in Table 7. In both these tables we consider three portfolios: 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" in these two tables, replaces the "shrinkage" portfolio considered in the earlier tables. Because we are using option-implied correlations, we do not know the distribution of returns for the resulting covariance matrix, and therefore, cannot use the shrinkage results of Ledoit and Wolf (2004a,b). Thus, instead of shrinkage, we use the regularization approach of Zumbach (2009), details of which are given in Appendix C.

Table 6, where we restrict correlations to be the same across all asset pairs, shows that of the three portfolios based on implied correlations that we consider, the shortsale-constrained portfolio has the best portfolio volatility, Sharpe ratio, and certainty-equivalent return for both datasets and also for all three trading frequencies, and performs substantially better than the portfolios based on the sample covariance matrix and regularization. However, even this portfolio performs worse in terms of all the performance metrics compared to the benchmark shortsale-constrained portfolio based on historical data in Table 3.

Table 7, where we allow correlations to be heterogeneous across asset pairs, shows that of the three portfolios based on implied correlations that we consider, the portfolio based on "regularization" has the lowest portfolio volatility for both datasets and for all three trading frequencies (though the difference with the volatility of the shortsale-constrained portfolio is small). The portfolio based on regularization also has the highest Sharpe ratio and certainty-equivalent return for the dataset with 100 stocks. However, for the dataset with 561 stocks, the constrained portfolio has the highest Sharpe ratio and certainty-equivalent return; the constrained portfolio also has the lowest turnover for both datasets and for both trading frequencies. The portfolio based on the "sample covariance" matrix without constraints and without regularization performs the worst across all four metrics.

Comparing the results in Table 7, where we allow correlations to be heterogeneous across asset pairs, to those in Table 6, where we restrict correlations to be the same across all asset pairs, we see that allowing correlations to be different across asset pairs leads to portfolios with a superior volatility, Sharpe ratio, and certainty-equivalent return, but turnover that is about twice as large for the data with 100 stocks and about three times as large for the data with 561 stocks.

However, in both Tables 6 and 7, the three portfolios we consider typically fail to outperform the $1/N$ portfolio (whose performance is reported in Table 3) in terms of Sharpe ratio, certainty-equivalent return, and turnover; the only exception is the constrained strategy, which achieves a higher Sharpe ratio and certainty-equivalent return for the case of 561 stocks in Panels A and B of Table 7, though the improvement in Sharpe ratio relative to the $1/N$ portfolio is not statistically significant.

We conclude that there are no significant benefits to using option-implied correlations for choosing portfolios. One explanation for the poor performance of portfolios based on implied correlations is that by replacing historical correlations with the implied correlations, we are essentially increasing the magnitude of the off-diagonal elements of the covariance matrix making it less stable; moreover, as can be seen in Figure 2, the implied correlations are also much more variable than the other correlation series. 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 historical volatility risk premium (HVRP); 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). Given the sensitivity of portfolio weights to the errors in estimating expected returns, the traditional mean-variance portfolios fail to outperform the benchmarks; hence, to reduce the impact of estimation error on the portfolio weights, we use the parametric-portfolio methodology of Brandt, Santa-Clara, and Valkanov (2009).

The out-of-sample performance of the mean-variance portfolios that use option-implied characteristics based on the parametric-portfolio methodology is reported in Table 8. We consider two benchmark portfolios: the $1/N$ portfolio and the mean-variance parametric portfolio "1/N + FF + MOM" 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 volatility and turnover but the parametric portfolio has a better Sharpe ratio and certainty equivalent return.

From the table we see that when we use the historical volatility risk premium (HVRP) as a characteristic instead of FF + MOM, the portfolio has a lower volatility than both benchmark portfolios, and more importantly, has a higher Sharpe ratio than both benchmarks. This is true for all three panels for daily, weekly, and monthly rebalancing. The turnover of this portfolio is also less than that of the benchmark parametric mean-variance portfolio. However, when we use the spread between implied and realized volatility (IRVS) to measure the volatility risk premium, the gains are smaller than those from using HVRP. When we use the model-free implied skewness (MFIS) as the characteristic, both the Sharpe ratio and certainty equivalent return are better than those for the two benchmark portfolios. For example, in Panel A for daily rebalancing, when using the characteristic MFIS the Sharpe ratio is 1.1903 compared to 1.0351 for the data with 100 stocks and is 1.0013 compared to 0.7928 for the data with 561 stocks. In both cases, the p-value for the difference in Sharpe ratios is 0.00. However, the portfolio based on MFIS has higher turnover and volatility than the two benchmark portfolios. The Sharpe ratio and certainty equivalent return are even higher when using the call-put implied volatility spread (CPVS). For example, in Panel A for daily rebalancing, the Sharpe ratio is 1.2336 compared to 0.7928 for the benchmark parametric portfolio for the data with 561 stocks. However, these gains are accompanied by higher turnover. Note that for all the option-implied quantities discussed above, as the rebalancing period increases, the improvement decreases, but so does the turnover.

When all four option-implied characteristics are used together, the improvement in performance is substantial. For the case of daily rebalancing and 100 stocks, the Sharpe ratio is 1.4093 compared to the Sharpe ratio of 0.9286 for the $1/N$ portfolio and 1.0351 for the portfolio that uses FF and MOM as the characteristics. While this improvement diminishes with the rebalancing period, the difference with the Sharpe ratio of the $1/N$ portfolio continues to have a p-value of 0.00 and with the FF+MOM portfolio the p-value is 0.00 for weekly rebalancing and 0.08 for monthly rebalancing. The portfolio using all four option-implied characteristics has also a higher certainty equivalent return and lower volatility compared to both benchmark portfolios; its main disadvantage is that it also has higher turnover, though the difference in turnover decreases as the rebalancing interval goes from daily to weekly to monthly.

We can also ask the question whether the option-implied characteristics improve performance if one is already using the 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 FF and MOM characteristics considered in the benchmark portfolio. From the lower part of each panel, we see that MFIS and CPVS improve the Sharpe ratio and certainty-equivalent return beyond the gains obtained from using only the FF and MOM characteristics. This improvement is statistically significant both for daily and weekly rebalancing, though not for monthly rebalancing. Moreover, this improvement is accompanied by an increase in portfolio turnover. The evidence for the improvement in performance from including HVRP or IRVS, when FF and MOM are already being used, is mixed and the differences are not statistically significant. Finally, looking at the last row of each panel where we add all four option-implied characteristics to the portfolio that already considers the FF and MOM characteristics, we observe that the portfolio volatility improves slightly, while the Sharpe ratio, certainty equivalent return, and turnover increase substantially.

To understand if the 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, Sharpe ratio, or certainty equivalent return) of a given strategy with that for the benchmark strategy. We consider two benchmark strategies when computing this cost: the $1/N$ portfolio and the parametric portfolio strategy that uses only the FF and MOM characteristics. The equivalent transactions costs for these two strategies are given in Table 9. For comparison, note that the typical cost for trading the stocks that are in our datasets is less than 20 basis points, with the actual cost depending on the size of the trade and the execution capability of the trader.

In Table 9, there are three possibilities when comparing portfolios that use option-implied information with the benchmark strategies. The first possibility is that the portfolio using option-implied information does not outperform the benchmark portfolio and yet the option-based portfolio has higher turnover than the benchmark; in this case, the transaction cost reported in the table is a *negative* number, implying that the only way the performance of the benchmark portfolio can be matched by the portfolio using option-implied information is if it is *given* money (rather than being charged) for making transactions. The second 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. Finally, it is possible that the portfolio using option-implied information has a better performance than the benchmark portfolio *and* also lower turnover; in this case, the performance of the option-based portfolio is better than that of the benchmark for any level of transactions cost, and we indicate this in the table with the symbol "$\star$".

In Table 9, for each dataset and rebalancing frequency, we report two sets of numbers, each set having three 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 FF and MOM characteristics. The first number in each set indicates the transaction cost that equates the mean return on the portfolio using option-implied characteristics to the mean return of the benchmark portfolio. The second number is the transaction cost that equates the Sharpe ratio of the two portfolios, and the third number is the transaction cost that equates the certainty equivalent return.

We observe from Panel A of Table 9 that for the data with 561 stocks, the strategy that uses the historical volatility risk premium (HVRP) as a characteristic will match the *mean* of the $1/N$ benchmark portfolio for a transaction cost of 30 basis points, the *Sharpe ratio* for a transaction cost of 49 basis points, and the *certainty equivalent return* for a transaction cost of 35 basis points. Similarly, when the benchmark is the parametric portfolio using FF and MOM as characteristics, the strategy that uses the historical volatility risk premium (HVRP) as a characteristic will outperform the benchmark portfolio for any level of transactions costs. The portfolio performance with daily rebalancing when using the other option-implied characteristics is not as good. For example, the portfolios using MFIS or CPVS as characteristics can outperform the benchmark portfolios only for transactions costs of up to 12 to 15 basis points. And, for the case of daily rebalancing, even if we use all four option characteristics, the portfolios based on these characteristics can outperform the benchmark portfolios for transaction costs of up to 12 basis points.

Studying the results in Panel B of Table 9 we see that the results are better because the turnover decreases faster than the benefits from using the option-implied characteristics. Looking at the results for the data with 561 stocks, we see that now the portfolio using HVRP does even better, achieving a higher Sharpe ratio than that for the $1/N$ portfolio for a transactions cost of up to 246 basis points and that of the "FF + MOM" portfolio for any transactions cost. The performance of the portfolio using the MFIS characteristic has also improved; for example, in the data with 561 stocks, it can match the Sharpe ratio of the $1/N$ portfolio for a transaction cost of 25 basis points and that of the FF and MOM portfolio for 19 basis points. The portfolios based on CPVS perform slightly worse than those based on MFIS, while those based on IRVS fail to beat the portfolio based on FF and MOM. The portfolio based on all four option characteristics outperforms two benchmark portfolios in terms of Sharpe ratio for up to 20 basis points.

Panel C of Table 9 illustrates that the results net of transactions costs are even stronger when the rebalancing frequency is monthly. In this case, the portfolio using HVRP as a characteristic attains a higher Sharpe ratio than that of the $1/N$ portfolio for very large transactions costs, and that of the portfolio using the FF and MOM characteristics for any transactions cost. Similarly, the portfolio based on MFIS can match the Sharpe ratio of the $1/N$ portfolio for transactions costs of up to 58 basis points and that of the portfolio using FF and MOM characteristics for up to 42 basis points. The portfolio based on all four option characteristics achieves a higher Sharpe ratio than the two benchmark portfolios for up to about 50 basis points.

The results for the *incremental* effect of option-implied characteristics are similar. When we consider all four option-implied characteristics *in addition to* the FF and MOM characteristics, we see from the last two rows of each panel in Table 9 that these portfolios outperform the two benchmark portfolios for up to 11 basis points for the case of daily trading, up to about 15 basis points for the case of weekly trading, and up to about 25 basis points for the case of monthly trading.

In summary, Table 9 suggests that in the presence of transactions costs, HVRP and MFIS are useful characteristics for choosing portfolios, CPVS is less useful, and IRVS usually fails to match the performance of the benchmark portfolio based on the FF and MOM characteristics.

Overall, the empirical evidence suggests that using information in the model-free historical volatility risk premium and model-free implied skewness can lead to an improvement in the out-of-sample portfolio Sharpe ratio and certainty-equivalent return, even after adjusting for the higher transactions costs incurred by these strategies.

6 Robustness Tests

In this section, we describe the various tests that we have undertaken to verify the robustness of the results of our empirical analysis in Section 5.

6.1 Different Datasets

Ideally, one would like to study more than a single dataset. We are limited in our desire to consider additional datasets because we have option prices only for U.S. stocks. To overcome this limitation, we have reported results for two datasets, where one is a subset of the other. The first dataset consists of 100 stocks out of the 219 for which data are available for all dates, where these 100 stocks are selected by first sorting the 219 stocks with respect to the security identifier code of IvyDB, and then selecting the first 100. The results for this dataset are reported in Columns 2–4 of each table. The second consists of the entire 561 stocks in our dataset, and the results for this dataset are reported in the last four columns of each table. In addition to the results reported for these two samples, we have also evaluated the performance of the different portfolios for 20 additional samples, where the first 10 samples are constructed by randomly choosing 100 stocks from the 219-stock dataset, and the second 10 samples are constructed by randomly choosing 250 stocks from the 561 stock dataset. The results from these additional samples confirm the findings reported in the tables regarding the relative improvement in Sharpe ratio and certainty-equivalent return of the option-implied portfolios relative to the benchmark portfolios.

6.2 Different Data Frequencies

We consider both *daily* stock returns from the CRSP daily file and *intraday* stock-price data from the NYSE's Trades-And-Quotes (TAQ) database. Results for the daily data are reported in the tables. Results for the high-frequency transaction data are not reported in order to conserve space, but can be obtained from the authors. The main insight from intraday data is that using high-frequency data to compute the covariance estimators rarely improves the out-of-sample performance of resulting portfolios, and in our analysis the intraday data give significantly better results relative to daily data only for the sample covariance matrix for 561 assets when neither shrinkage is applied to the covariance matrix nor shortsale constraints are imposed on the portfolio weights.

6.3 Different Rebalancing Frequencies

We report results for three rebalancing frequencies: daily, weekly, and monthly. We also computed the results for a holding period of two weeks, but have not included these results in order to conserve space. The results for the fortnightly rebalancing frequency are in between the results for weekly and monthly balancing. Overall, we find that the results are similar for the four holding periods, although Sharpe ratios and certainty-equivalent returns are higher for daily rebalancing, while turnover is lower for monthly rebalancing (note that in the tables where we report the results about the performance of various portfolios, in Panel A the turnover numbers are per day, in Panel B they are per week, and in Panel C they are per month).

6.4 Different Benchmark Portfolios

We consider eight benchmark portfolios, of which six are listed in Table 3. The first is the equally-weighted ($1/N$) portfolio in which one invests an equal amount of wealth across all N available stocks each period. As DeMiguel, Garlappi, and Uppal (2009) have shown, this portfolio performs extremely well. For example, its Sharpe ratio over our sample period exceeds 0.90, while that for the S&P500 index over the corresponding period is only 0.35. Then there are four strategies that are based on minimizing the portfolio variance. The second benchmark portfolio is the minimum-variance portfolio using daily data to compute the sample covariance matrix. The third benchmark portfolio is the minimum-variance portfolio with shortsale constraints, which is motivated by the finding in Jagannathan and Ma (2003) that imposing shortselling constraints can lead to substantial gains in performance. The fourth benchmark portfolio is the minimum-variance portfolio with "shrinkage," using the approach in Ledoit and Wolf (2004a,b) for daily data, and with regularization using the approach of Zumbach (2009) for intraday data.[40] The fifth benchmark portfolio is the minimum-variance portfolio with all correlations set equal to zero. Finally, we consider the parametric portfolios developed in Brandt and Santa-Clara (2005).

There are two other benchmark portfolios that we consider but whose results we do not report. The first is the mean-variance portfolio. We do not report the performance of the mean-variance portfolio in each table, because the performance of this portfolio is quite poor, as already documented extensively in the literature; see, for instance, DeMiguel, Garlappi, and

[40]The results with intraday data are available on request.

Uppal (2009). The second is the minimum-variance portfolio with all correlations set equal to the mean correlation across all asset pairs; again, the performance of this portfolio is not reported because it performs quite poorly relative to the other benchmarks we consider.

7 Conclusion

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, certainty-equivalent return, and turnover. The benchmark portfolios are the $1/N$ portfolio; four types of minimum-variance portfolios based on historical data: the first on the sample covariance matrix, the second with shortsale constraints, the third with shrinkage applied to the covariance matrix, and the fourth with correlations set equal to zero; and, the parametric portfolios of Brandt, Santa-Clara, and Valkanov (2009) based on the "size" and "value" characteristics identified in Fama and French (1992), and "momentum" identified in Jegadeesh and Titman (1993).

We find that forming portfolios using expected returns that exploit information in option-implied model-free skewness and volatility risk premium outperform, in terms of Sharpe ratio and certainty-equivalent return, portfolios that ignore option-implied information. This improvement in performance is present even after adjusting for transactions costs. Using option-implied volatilities one can achieve also a small improvement in portfolio volatility, but not in the Sharpe ratio. However, option-implied correlations are highly variable and lead to poorly behaved and unstable covariance matrices, resulting in portfolios with fluctuating weights that fail to outperform the benchmarks in terms of either portfolio volatility or Sharpe ratio.

Based on our empirical analysis, we conclude that prices of stock options contain information that can be used to improve the out-of-sample performance of portfolios. In this paper, we have explored only very simple ways of incorporating information implied by option prices into static portfolios; more sophisticated ways of incorporating this information should lead to even larger gains in out-of-sample performance.

A The Construction of the Risk-Neutral Implied Moments

The formulas in this appendix follow closely Bakshi, Kapadia, and Madan (2003) and are reproduced here only for completeness; for more details, please refer to the original paper.

Let $S(t)$ be the stock price at time t and $R(t,\tau)$ the τ-period return (seen at time $t+\tau$) given by the log-price relative:

$$R(t,\tau) \equiv \ln S(t+\tau) - \ln S(t). \tag{A1}$$

Let r be the interest rate, $C(t,\tau;K)$ and $P(t,\tau;K)$ the prices of call and put options written on the stock with current price $S(t)$, τ the time to maturity, and K the strike price.

Let $V(t,\tau) \equiv \mathcal{E}_t^*\{e^{-r\tau}R(t,\tau)^2\}$, $W(t,\tau) \equiv \mathcal{E}_t^*\{e^{-r\tau}R(t,\tau)^3\}$, and $X(t,\tau) \equiv \mathcal{E}_t^*\{e^{-r\tau}R(t,\tau)^4\}$ represent the fair value of the variance, cubic, and quartic contracts, respectively. Then, the price of the variance contract is given by

$$\begin{aligned} V(t,\tau) = & \int_{S(t)}^{\infty} \frac{2\left(1-\log\left(\frac{K}{S(t)}\right)\right)}{K^2} \cdot C(t,\tau;K)dK \qquad \text{(A2)} \\ & + \int_0^{S(t)} \frac{2\left(1-\log\left(\frac{K}{S(t)}\right)\right)}{K^2} \cdot P(t,\tau;K)dK, \end{aligned}$$

the price of the cubic contract is

$$\begin{aligned} W(t,\tau) = & \int_{S(t)}^{\infty} \frac{6\log\left(\frac{K}{S(t)}\right) - 3\left(\log\left(\frac{K}{S(t)}\right)\right)^2}{K^2} \cdot C(t,\tau;K)dK \\ & - \int_0^{S(t)} \frac{6\log\left(\frac{K}{S(t)}\right) + 3\left(\log\left(\frac{K}{S(t)}\right)\right)^2}{K^2} \cdot P(t,\tau;K)dK, \qquad \text{(A3)} \end{aligned}$$

and the price of the quartic contract is

$$\begin{aligned} X(t,\tau) = & \int_{S(t)}^{\infty} \frac{12\left(ln[\frac{K}{S(t)}]\right)^2 - 4\left(ln[\frac{K}{S(t)}]\right)^3}{K^2} \cdot C(t,\tau;K)dK \\ & + \int_0^{S(t)} \frac{12\left(ln[\frac{S(t)}{K}]\right)^2 + 4\left(ln[\frac{S(t)}{K}]\right)^3}{K^2} \cdot P(t,\tau;K)dK. \qquad \text{(A4)} \end{aligned}$$

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

$$\mathrm{MFIV}(t,\tau) = \big(V(t,\tau)\big)^{1/2}, \tag{A5}$$

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

$$\mathrm{MFIS}(t,\tau) = \frac{e^{r\tau}W(t,\tau) - 3\mu(t,\tau)e^{r\tau}V(t,\tau) + 2(\mu(t,\tau))^3}{(e^{r\tau}V(t,\tau) - (\mu(t,\tau))^2)^{\frac{3}{2}}}. \tag{A6}$$

To calculate the integrals in (A2), (A3), and (A4) precisely, we need a continuum of option prices. We discretize the respective integrals and approximate them using the available options. As mentioned earlier, we normally have 13 out-of-the-money call and put implied volatilities for each maturity. Using cubic splines, we interpolate them inside the available moneyness range, and extrapolate using the last known (boundary for each side) value to fill in a total of 1001 grid points in the moneyness range from 1/3 to 3.[41] 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 (A5) and (A6), respectively.

B Computing (Co)variances from Intraday Data

Consistent with the literature on estimating moments from intraday data (see, for example, Brown (1990), Zhou (1996), and Corsi, Zumbach, Müller, and Dacorogna (2001)), we assume that instead of the true price, $X_i(t)$, we observe $Y_i(t)$, which is contaminated with noise; that is, $Y_i(t) = X_i(t) + \epsilon_i(t)$, where the noise process ϵ_i is assumed to be i.i.d and independent also of X_i. A common estimator for the integrated (co)variance of the efficient price process $\langle X_i, X_j \rangle$ is given by the *Realized (Co)Variance* (RV/RC):

$$\widehat{\langle Y_i, Y_j \rangle}^{(\Delta)} = \sum_{k=1}^{n} r_i(k\Delta) \cdot r_j\,(k\Delta)\,, \tag{B1}$$

where the sampling frequency n is defined as $n = T/\Delta$, and $r_i(t)$ denotes the observed stock return for a time interval of length Δ; that is, $r_i(t) = Y_i(t) - Y_i(t-\Delta)$.

In the absence of microstructure noise, this estimator is consistent as the sampling frequency n increases (Jacod (1994), and Jacod and Protter (1998)). However, it is inconsistent under

[41] The reason for choosing such a wide grid is that our simulation studies have shown that with a narrower grid we may not be estimating the skew and kurtosis of the risk-neutral distribution well enough. Decreasing the number of points in the grid also leads to a deterioration in accuracy.

real market conditions where there is noise and asynchronous trading (Barndorff-Nielsen and Shephard (2002), and Zhang, Mykland, and Aït-Sahalia (2005)).[42] To mitigate this problem, we use the "second-best" estimator from Zhang, Mykland, and Aït-Sahalia (originally derived to estimate realized variances) and apply it to realized (co)variances. The idea underlying this estimator is to compute the realized (co)variance estimator in (B1) at a low frequency in order to mitigate the problems induced by microstructure noise and non-synchronicity. When we sample at lower frequencies, we discard some observations; to overcome this problem, Zhang, Mykland, and Aït-Sahalia suggest computing the realized (co)variance estimator in (B1) over different subsamples and then averaging the estimators obtained for these subsamples. The "second-best" estimator is given by:

$$\widehat{\langle Y_i, Y_j \rangle}^{(avg,K)} = \frac{1}{K} \sum_{k=1}^{K} \widehat{\langle Y_i, Y_j \rangle}^{(\Delta,k)}. \tag{B2}$$

We introduce one more averaging step to eliminate the chance of choosing the wrong sampling frequency by calculating the estimator (B2) over several frequencies and taking the mean. Because our sample also includes less frequently traded stocks, especially early in the sample period, we choose relatively low sampling frequencies from 240 to 390 minutes (which corresponds to the number of minutes in a typical trading day) with a step size of 10 minutes for the estimator (B2) to get the final realized (co)variance estimator:

$$\widehat{\langle Y_i, Y_j \rangle}^{(avg,\bar{K})} = \frac{1}{dim(\hat{K})} \sum_{s=1}^{dim(\bar{K})} \widehat{\langle Y_i, Y_j \rangle}^{(avg,\bar{K}(s))}. \tag{B3}$$

C Shrinkage and Regularization of Covariance Matrix

We consider a sample with a large number of stocks (100 and 561); therefore, sample covariance matrices estimated from a limited history of daily stock returns are likely to be poorly behaved. To improve the sample covariance estimate for daily returns, we apply the shrinkage methodology of Ledoit and Wolf (2004a,b):

$$\widehat{\Sigma}_{Shrunk} = (1 - \phi)\widehat{\Sigma} + \phi S, \tag{C1}$$

[42] The non-synchronicity of the data induces an additional bias, known as the Epps effect (Epps (1979)), which drives covariances to zero as the sampling frequency increases.

where we shrink the sample estimate of the covariance matrix $\widehat{\Sigma}$ toward a diagonal matrix with the cross-sectional average variance on the diagonal, defined as the target S.[43] We minimize the Frobenius norm between the shrinkage estimator and the true covariance matrix in order to find the optimal shrinkage intensity parameter ϕ, using the time series of 750 points; details can be found in Ledoit and Wolf (2004a,b).

The asymptotic properties of intraday data are different from those of daily data, as shown by Zhang, Mykland, and Aït-Sahalia (2005) among others, and the intraday data do not satisfy the distributional assumptions of Ledoit and Wolf (2004a,b). Therefore, to improve the properties of the covariance matrix estimated for intraday returns, we apply regularization of the inverse covariance proposed by Zumbach (2009). He uses the spectral decomposition of the covariance matrix estimator $\widehat{\Sigma}$, which is:

$$\widehat{\Sigma} = \sum_{n=1}^{N} \lambda_n V_n V_n^{'}, \tag{C2}$$

where $\{\lambda_1, ..., \lambda_N\}$ are eigenvalues and $\{V_1, ..., V_N\}$ are eigenvectors (pairwise orthogonal) for the set of N stocks. We order the eigenvalues by decreasing values, such that λ_1 is the largest eigenvalue. The inverse square root covariance can then be written as:

$$\widehat{\Sigma}^{-1/2} = \sum_{n=1}^{N} \frac{1}{\sqrt{\lambda_n}} V_n V_n^{'}, \tag{C3}$$

where one can see that for $\lambda_n \approx 0$ the singularity problem arises. To overcome this problem, we define

$$\widetilde{\Sigma}^{-1/2} = \sum_{n=1}^{k} \frac{1}{\sqrt{\lambda_n}} V_n V_n^{'} + \frac{1}{\sqrt{\lambda_{k+1}}} \sum_{n=k+1}^{N} V_n V_n^{'}, \tag{C4}$$

and use $\widetilde{\Sigma}$ as the estimator of the covariance matrix. This approach allows us to maintain the covariance structure by keeping all eigenvectors $\{V_1, ..., V_N\}$ and to eliminate the singularity problem by substituting all eigenvalues that are smaller than λ_{k+1} by $\lambda_{k+1} \neq 0$. We choose the parameter k so that the first k eigenvalues explain 75% of the overall variance.

[43] We also used the cross-sectional average covariances matrix as the target, but the first target performs better out of sample.

Table 1: Prediction of Volatility, Correlation

In this table, we report the results of predicting volatilities and correlations. The results of the predictive regressions are reported in terms of both RMSE and R^2. In Panel A of this table, we report the results of the volatilities prediction regressions $RV = \alpha + \beta \widehat{RV}$, where we regress the 30-days realized volatilities (RV) on the volatility predictors ($\widehat{RV}$). To calculate the RMSE we assume $\alpha = 0$ and $\beta = 1$. The volatilities predictors are: historical daily volatility (based on past 250 days for the 100-stock sample and past 750 days for the 561-stock sample), historical intraday volatility (based on past 30 days), implied volatility, and implied volatility adjusted for the variance risk premium. In Panel B, we show the results of the correlation prediction regressions $corr = \alpha + \beta \widehat{corr}$, where we regress the 30-days realized correlation on the correlation predictors ($\widehat{corr}$). To calculate the RMSE we assume $\alpha = 0$ and $\beta = 1$. The correlation predictors are: historical daily correlation (based on past 250 days for 100 stock sample and past 750 days for 561 stock-sample), historical intraday correlation (based on past 30 days), implied correlation based on daily data, and implied correlation based on intraday data.

	100 stocks		561 stocks	
Predictor	RMSE	R^2	RMSE	R^2
Panel A: Prediction of volatility				
Historical daily volatility	0.1293	0.1629	0.1604	0.1791
Historical intraday volatility	0.1165	0.3434	0.1298	0.3116
Implied volatility	0.1067	0.4468	0.1215	0.4052
Implied volatility (HVRP corrected)	0.0990	0.4438	0.1118	0.4022
Panel B: Prediction of correlation				
Historical daily correlation	0.1873	0.0653	0.1924	0.0460
Historical intraday correlation	0.2205	0.0652	0.2188	0.0655
Implied daily correlation	0.2254	0.0791	0.2040	0.0990
Implied intraday correlation	0.2147	0.0831	0.2127	0.0933

Table 2: Return Predictability

In this table, we report the results of cross-sectional return predictability of various stock characteristics: size (SIZE), book-to-market (BTM), momentum (MOM), historical volatility risk premium (HVRP), model-free implied skewness (MFIS), call-put-volatility spread (CPVS), and implied-realized-volatility spread (IRVS). For each stock subsample (100 and 561 stocks), 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.

	100 Stocks			561 Stocks		
Characteristic	1 day	1 week	1 month	1 day	1 week	1 month
SIZE	-0.2441	-0.2311	-0.2331	-0.1960	-0.2023	-0.2072
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
BTM	0.3232	0.2708	0.2117	0.2192	0.2309	0.2174
	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)
MOM	-0.1841	-0.1342	-0.0927	-0.1147	-0.0522	0.0030
	(0.01)	(0.02)	(0.06)	(0.09)	(0.26)	(0.49)
HVRP	0.0586	0.0543	0.0421	-0.0350	-0.0333	-0.0227
	(0.16)	(0.13)	(0.19)	(0.32)	(0.30)	(0.35)
IRVS	0.0588	0.0891	0.0504	0.1468	0.0824	0.0573
	(0.17)	(0.02)	(0.08)	(0.00)	(0.03)	(0.04)
MFIS	0.4051	0.2329	0.1581	0.4133	0.2033	0.1244
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
CPVS	0.6234	0.2124	0.0525	0.9251	0.2687	0.0731
	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)

Table 3: Benchmark Portfolios That Do Not Use Option-Implied Information

In this table, we evaluate the performance of various benchmark portfolios that are based on historical returns and do *not* rely on prices of options. The $1/N$ portfolio is the equally-weighted strategy where one invests an equal amount of wealth across all N available stocks each period. 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 two p-values in parenthesis, the first with respect to the $1/N$ portfolio, and the second with respect to the "Sample cov" portfolio, with the null hypothesis being that the portfolio being evaluated is worse than the benchmark (so a small p-value suggests *rejecting* the null hypothesis that the portfolio being evaluated is worse than the benchmark).

	100 stocks				561 stocks			
Strategy	std	sr	ce	trn	std	sr	ce	trn
				Panel A: Daily rebalancing				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7490	0.1155	0.0144
Sample cov	0.1331	0.9066	0.1118	0.3064	0.1333	0.2998	0.0311	0.5342
	(0.00)	(0.53)	(0.71)		(0.00)	(0.90)	(0.94)	
Constrained	0.1199	0.8251	0.0917	0.0580	0.1161	0.7435	0.0796	0.0364
	(0.00)	(0.70)	(0.93)		(0.00)	(0.51)	(0.85)	
Shrinkage	0.1197	1.0323	0.1164	0.1716	0.1164	0.4440	0.0449	0.3050
	(0.00)	(0.35)	(0.68)		(0.00)	(0.83)	(0.92)	
Zero Correlation	0.1433	0.8820	0.1161	0.0131	0.1493	0.7808	0.1055	0.0124
	(0.00)	(0.78)	(0.97)		(0.00)	(0.33)	(0.75)	
1/N + FF + MOM	0.1636	1.0351	0.1559	0.0805	0.1796	0.7928	0.1263	0.0710
	(1.00)	(0.02)	(0.01)		(1.00)	(0.15)	(0.08)	
				Panel B: Weekly rebalancing				
1/N	0.1616	0.9214	0.1358	0.0290	0.1766	0.7384	0.1148	0.0325
Sample cov	0.1310	0.9353	0.1139	0.3334	0.1344	0.2760	0.0280	0.6083
	(0.00)	(0.48)	(0.72)		(0.00)	(0.95)	(0.98)	
Constrained	0.1197	0.8189	0.0909	0.0644	0.1213	0.7129	0.0791	0.0456
	(0.00)	(0.74)	(0.97)		(0.00)	(0.55)	(0.90)	
Shrinkage	0.1190	1.0512	0.1180	0.1938	0.1195	0.4081	0.0416	0.3653
	(0.00)	(0.29)	(0.70)		(0.00)	(0.91)	(0.97)	
Zero Correlation	0.1434	0.8775	0.1156	0.0266	0.1511	0.7716	0.1052	0.0273
	(0.00)	(0.80)	(0.98)		(0.00)	(0.30)	(0.77)	
1/N + FF + MOM	0.1661	0.9817	0.1492	0.0885	0.1820	0.7880	0.1268	0.0815
	(1.00)	(0.06)	(0.03)		(1.00)	(0.11)	(0.06)	
				Panel C: Monthly rebalancing				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7337	0.1107	0.0673
Sample cov	0.1313	0.9323	0.1138	0.4045	0.1329	0.2452	0.0237	0.7894
	(0.02)	(0.52)	(0.70)		(0.00)	(0.97)	(0.98)	
Constrained	0.1168	0.8409	0.0914	0.0802	0.1239	0.6852	0.0771	0.0663
	(0.00)	(0.72)	(0.95)		(0.00)	(0.59)	(0.87)	
Shrinkage	0.1172	1.0625	0.1177	0.2492	0.1236	0.3706	0.0381	0.5072
	(0.00)	(0.31)	(0.67)		(0.00)	(0.93)	(0.98)	
Zero Correlation	0.1342	0.9076	0.1127	0.0524	0.1444	0.7794	0.1020	0.0562
	(0.00)	(0.70)	(0.97)		(0.00)	(0.26)	(0.74)	
1/N + FF + MOM	0.1614	0.9938	0.1472	0.1098	0.1788	0.7576	0.1193	0.1072
	(1.00)	(0.09)	(0.03)		(0.99)	(0.26)	(0.14)	

Table 4: Portfolios Using Option-Implied Volatility

In this table, we report the performance of various minimum-variance portfolios that use the model-free implied volatility calculated from option prices, but with correlations estimated from historical data. The $1/N$ portfolio is the equally-weighted strategy where an equal amount of wealth is invested in the N available stocks each period. 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; "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 covariance 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. We report two p-values in parenthesis, the first with respect to the $1/N$ portfolio, and the second with respect to the corresponding benchmark minimum-variance portfolio in Table 3, with the null hypothesis being that the portfolio being evaluated is worse than the benchmark (so a small p-value suggests *rejecting* the null hypothesis that the portfolio being evaluated is worse than the benchmark).

	100 stocks				561 stocks			
Strategy	std	sr	ce	trn	std	sr	ce	trn
Panel A: Daily rebalancing								
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7490	0.1155	0.0144
Sample cov	0.1351	0.4122	0.0466	0.7945	0.1312	0.4297	0.0478	1.5501
	(0.00)	(0.95)	(0.97)		(0.00)	(0.85)	(0.91)	
	(0.84)	(1.00)	(1.00)		(0.31)	(0.31)	(0.32)	
Constrained	0.1175	0.6872	0.0738	0.2267	0.1237	0.6142	0.0683	0.2763
	(0.00)	(0.90)	(0.98)		(0.00)	(0.69)	(0.87)	
	(0.05)	(0.86)	(0.88)		(0.89)	(0.72)	(0.67)	
Shrinkage	0.1197	0.5388	0.0573	0.5210	0.1161	0.3487	0.0338	1.1091
	(0.00)	(0.93)	(0.98)		(0.00)	(0.93)	(0.97)	
	(0.50)	(1.00)	(1.00)		(0.45)	(0.66)	(0.66)	
Zero Correlation	0.1433	0.8085	0.1056	0.0521	0.1468	0.7230	0.0954	0.0574
	(0.00)	(0.98)	(1.00)		(0.00)	(0.64)	(0.91)	
	(0.57)	(1.00)	(1.00)		(0.00)	(0.99)	(1.00)	
Panel B: Weekly rebalancing								
1/N	0.1616	0.9214	0.1358	0.0290	0.1766	0.7384	0.1148	0.0325
Sample cov	0.1336	0.4728	0.0542	0.8074	0.1356	0.4431	0.0511	1.5711
	(0.00)	(0.97)	(0.99)		(0.00)	(0.88)	(0.94)	
	(0.84)	(1.00)	(1.00)		(0.50)	(0.21)	(0.20)	
Constrained	0.1196	0.6942	0.0758	0.2289	0.1214	0.6495	0.0715	0.2786
	(0.00)	(0.93)	(0.99)		(0.00)	(0.65)	(0.91)	
	(0.45)	(0.90)	(0.90)		(0.48)	(0.65)	(0.65)	
Shrinkage	0.1202	0.6165	0.0669	0.5313	0.1224	0.3883	0.0401	1.1272
	(0.00)	(0.92)	(0.98)		(0.00)	(0.94)	(0.98)	
	(0.75)	(1.00)	(1.00)		(0.61)	(0.55)	(0.53)	
Zero Correlation	0.1436	0.8099	0.1060	0.0589	0.1479	0.7302	0.0970	0.0647
	(0.00)	(0.99)	(1.00)		(0.00)	(0.54)	(0.91)	
	(0.69)	(1.00)	(1.00)		(0.00)	(0.98)	(1.00)	
Panel C: Monthly rebalancing								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7337	0.1107	0.0673
Sample cov	0.1320	0.5923	0.0695	0.8495	0.1283	0.4244	0.0463	1.6371
	(0.01)	(0.92)	(0.96)		(0.00)	(0.92)	(0.96)	
	(0.57)	(0.98)	(0.99)		(0.24)	(0.19)	(0.20)	
Constrained	0.1142	0.7240	0.0761	0.2368	0.1178	0.7090	0.0766	0.2849
	(0.00)	(0.91)	(0.99)		(0.00)	(0.54)	(0.88)	
	(0.19)	(0.85)	(0.90)		(0.14)	(0.47)	(0.55)	
Shrinkage	0.1188	0.7130	0.0777	0.5641	0.1190	0.3829	0.0385	1.1832
	(0.00)	(0.85)	(0.96)		(0.00)	(0.96)	(0.99)	
	(0.69)	(1.00)	(1.00)		(0.18)	(0.47)	(0.50)	
Zero Correlation	0.1351	0.8481	0.1054	0.0768	0.1405	0.7617	0.0971	0.0840
	(0.00)	(0.96)	(1.00)		(0.00)	(0.34)	(0.85)	
	(0.90)	(1.00)	(0.99)		(0.00)	(0.80)	(0.93)	

Table 5: Portfolios Using Risk-Premium-Corrected Option-Implied Volatility

In this table, we evaluate the performance of various portfolios that use the risk-premium-corrected model-free implied volatility calculated from option prices, while correlations are estimated from historical data. The $1/N$ portfolio is the equally-weighted strategy where each period one invests an equal amount across all N available stocks. 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 3, with the null hypothesis being that the portfolio being evaluated is worse than the benchmark (so a small p-value suggests *rejecting* the null hypothesis that the portfolio being evaluated is worse than the benchmark).

	100 stocks				561 stocks			
Strategy	std	sr	ce	trn	std	sr	ce	trn
				Panel A: Daily rebalancing				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7490	0.1155	0.0144
Sample cov	0.1323	0.8085	0.0982	0.6635	0.1250	0.6619	0.0749	1.1773
	(0.00)	(0.66)	(0.80)		(0.00)	(0.60)	(0.80)	
	(0.35)	(0.72)	(0.73)		(0.06)	(0.11)	(0.12)	
Constrained	0.1190	0.8016	0.0883	0.2188	0.1175	0.6726	0.0722	0.2437
	(0.00)	(0.72)	(0.93)		(0.00)	(0.59)	(0.83)	
	(0.28)	(0.57)	(0.58)		(0.56)	(0.61)	(0.60)	
Shrinkage	0.1176	0.8630	0.0946	0.4249	0.1110	0.5843	0.0587	0.8476
	(0.00)	(0.61)	(0.86)		(0.00)	(0.74)	(0.92)	
	(0.06)	(0.91)	(0.93)		(0.02)	(0.27)	(0.29)	
Zero Correlation	0.1384	0.8528	0.1085	0.0541	0.1399	0.7496	0.0951	0.0618
	(0.00)	(0.85)	(0.99)		(0.00)	(0.50)	(0.87)	
	(0.00)	(0.86)	(0.97)		(0.00)	(0.81)	(0.96)	
				Panel B: Weekly rebalancing				
1/N	0.1616	0.9214	0.1358	0.0290	0.1766	0.7384	0.1148	0.0325
Sample cov	0.1314	0.8382	0.1015	0.6747	0.1278	0.6488	0.0749	1.1940
	(0.00)	(0.63)	(0.83)		(0.00)	(0.65)	(0.85)	
	(0.55)	(0.77)	(0.77)		(0.19)	(0.04)	(0.05)	
Constrained	0.1226	0.8067	0.0914	0.2210	0.1161	0.7346	0.0786	0.2452
	(0.00)	(0.75)	(0.95)		(0.00)	(0.50)	(0.84)	
	(0.95)	(0.54)	(0.48)		(0.18)	(0.46)	(0.52)	
Shrinkage	0.1190	0.9146	0.1017	0.4347	0.1161	0.5997	0.0630	0.8626
	(0.00)	(0.51)	(0.87)		(0.00)	(0.76)	(0.94)	
	(0.47)	(0.93)	(0.93)		(0.23)	(0.11)	(0.12)	
Zero Correlation	0.1389	0.8503	0.1085	0.0605	0.1412	0.7552	0.0967	0.0684
	(0.00)	(0.88)	(0.99)		(0.00)	(0.42)	(0.88)	
	(0.00)	(0.91)	(0.99)		(0.00)	(0.72)	(0.96)	
				Panel C: Monthly rebalancing				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7337	0.1107	0.0673
Sample cov	0.1294	0.9363	0.1128	0.7127	0.1214	0.6795	0.0752	1.2476
	(0.01)	(0.51)	(0.71)		(0.00)	(0.60)	(0.84)	
	(0.30)	(0.49)	(0.54)		(0.05)	(0.01)	(0.02)	
Constrained	0.1150	0.8527	0.0914	0.2282	0.1097	0.8483	0.0871	0.2499
	(0.00)	(0.69)	(0.94)		(0.00)	(0.27)	(0.77)	
	(0.26)	(0.46)	(0.51)		(0.02)	(0.21)	(0.38)	
Shrinkage	0.1162	0.9868	0.1080	0.4660	0.1152	0.6010	0.0627	0.9102
	(0.00)	(0.43)	(0.79)		(0.00)	(0.76)	(0.94)	
	(0.35)	(0.79)	(0.82)		(0.05)	(0.05)	(0.08)	
Zero Correlation	0.1304	0.8929	0.1079	0.0779	0.1350	0.7924	0.0978	0.0865
	(0.00)	(0.76)	(0.99)		(0.00)	(0.23)	(0.80)	
	(0.00)	(0.79)	(0.97)		(0.00)	(0.29)	(0.85)	

Table 6: Portfolios Using Homogeneous Option-Implied Correlation

In this table, we evaluate the performance of various portfolios that use option-implied correlation, under the restriction that the correlation is the same across all asset pairs, as computed in Driessen, Maenhout, and Vilkov (2009), while volatilities are estimated from historical data. The $1/N$ portfolio is the equally-weighted strategy where one invests an equal amount of wealth across all N available stocks each period. The "Sample cov" portfolio is the minimum-variance portfolio based on the sample covariance matrix, but with option-implied correlations that are 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 3, with the null hypothesis being that the portfolio being evaluated is worse than the benchmark (so a small p-value suggests *rejecting* the null hypothesis that the portfolio being evaluated is worse than the benchmark).

	100 stocks				561 stocks			
Strategy	std	sr	ce	trn	std	sr	ce	trn
				Panel A: Daily rebalancing				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7490	0.1155	0.0144
Sample cov	0.1594	0.1115	0.0051	0.0810	0.1915	0.2589	0.0312	0.0511
	(0.37)	(0.98)	(0.98)		(1.00)	(0.87)	(0.85)	
	(1.00)	(0.99)	(0.99)		(1.00)	(0.55)	(0.50)	
Constrained	0.1319	0.6850	0.0816	0.0710	0.1439	0.6582	0.0844	0.0561
	(0.00)	(0.81)	(0.91)		(0.00)	(0.61)	(0.72)	
	(1.00)	(0.79)	(0.68)		(1.00)	(0.65)	(0.43)	
Regularization	0.1579	0.1193	0.0064	0.1062	0.2027	0.2308	0.0262	0.0819
	(0.23)	(0.99)	(0.99)		(1.00)	(0.89)	(0.87)	
	(1.00)	(1.00)	(1.00)		(1.00)	(0.73)	(0.62)	
				Panel B: Weekly rebalancing				
1/N	0.1616	0.9214	0.1358	0.0290	0.1766	0.7384	0.1148	0.0325
Sample cov	0.1662	0.0837	0.0001	0.1105	0.2052	0.2329	0.0268	0.0994
	(0.75)	(0.99)	(0.99)		(1.00)	(0.91)	(0.89)	
	(1.00)	(1.00)	(1.00)		(1.00)	(0.56)	(0.51)	
Constrained	0.1296	0.6831	0.0802	0.0742	0.1510	0.6138	0.0813	0.0590
	(0.00)	(0.85)	(0.94)		(0.00)	(0.67)	(0.77)	
	(1.00)	(0.82)	(0.72)		(1.00)	(0.71)	(0.46)	
Regularization	0.1636	0.1346	0.0086	0.1363	0.2145	0.2397	0.0286	0.1343
	(0.61)	(0.99)	(1.00)		(1.00)	(0.92)	(0.89)	
	(1.00)	(1.00)	(1.00)		(1.00)	(0.74)	(0.61)	
				Panel C: Monthly rebalancing				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7337	0.1107	0.0673
Sample cov	0.1654	0.0956	0.0020	0.1807	0.2085	0.2304	0.0260	0.1990
	(0.85)	(0.99)	(0.99)		(0.98)	(0.90)	(0.87)	
	(1.00)	(1.00)	(0.99)		(1.00)	(0.52)	(0.48)	
Constrained	0.1249	0.7132	0.0813	0.0827	0.1449	0.5968	0.0760	0.0677
	(0.00)	(0.80)	(0.91)		(0.01)	(0.68)	(0.78)	
	(0.93)	(0.77)	(0.69)		(1.00)	(0.70)	(0.52)	
Regularization	0.1618	0.1089	0.0043	0.2092	0.2195	0.2277	0.0253	0.2467
	(0.74)	(0.99)	(0.99)		(0.99)	(0.91)	(0.88)	
	(1.00)	(1.00)	(1.00)		(1.00)	(0.71)	(0.60)	

Table 7: Portfolios Using Heterogeneous Option-Implied Correlation

In this table, we evaluate the performance of various portfolios that use option-implied correlation without restricting correlations to be the same across asset pairs, as computed in Buss and Vilkov (2008), while volatilities are estimated from historical data. The $1/N$ portfolio is the equally-weighted strategy where one invests an equal amount of wealth across all N available stocks each period. 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 3, with the null hypothesis being that the portfolio being evaluated is worse than the benchmark (so a small p-value suggests *rejecting* the null hypothesis that the portfolio being evaluated is worse than the benchmark).

	100 stocks				561 stocks			
Strategy	std	sr	ce	trn	std	sr	ce	trn
Panel A: Daily rebalancing								
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7490	0.1155	0.0144
Sample cov	0.5200	0.5197	0.1949	2.0310	0.7564	-0.5270	-251.0000	14.2573
	(0.69)	(0.82)	(0.34)		(1.00)	(1.00)	(1.00)	
	(1.00)	(0.89)	(0.26)		(1.00)	(0.98)	(1.00)	
Constrained	0.1220	0.7839	0.0882	0.1457	0.1209	0.9117	0.1030	0.1792
	(0.00)	(0.72)	(0.91)		(0.00)	(0.26)	(0.62)	
	(0.95)	(0.64)	(0.60)		(1.00)	(0.13)	(0.10)	
Regularization	0.1219	0.8402	0.0950	0.1924	0.1176	0.5876	0.0622	0.2539
	(0.00)	(0.62)	(0.84)		(0.00)	(0.69)	(0.86)	
	(0.91)	(0.87)	(0.85)		(0.67)	(0.27)	(0.27)	
Panel B: Weekly rebalancing								
1/N	0.1616	0.9214	0.1358	0.0290	0.1766	0.7384	0.1148	0.0325
Sample cov	1.2015	0.2371	0.1473	1.8090	1.3837	-0.1986	-50.2000	15.1728
	(0.68)	(0.89)	(0.45)		(1.00)	(1.00)	(1.00)	
	(1.00)	(0.97)	(0.30)		(1.00)	(0.98)	(1.00)	
Constrained	0.1206	0.7299	0.0808	0.1483	0.1250	0.7796	0.0896	0.1821
	(0.00)	(0.83)	(0.96)		(0.00)	(0.42)	(0.77)	
	(0.70)	(0.85)	(0.84)		(0.92)	(0.29)	(0.24)	
Regularization	0.1225	0.7748	0.0874	0.2060	0.1189	0.4262	0.0436	0.2816
	(0.00)	(0.74)	(0.93)		(0.00)	(0.88)	(0.96)	
	(0.91)	(0.98)	(0.97)		(0.41)	(0.47)	(0.47)	
Panel C: Monthly rebalancing								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7337	0.1107	0.0673
Sample cov	1.4607	0.1453	0.1144	1.8248	0.5288	-0.2353	-11.9524	18.6061
	(0.73)	(0.88)	(0.64)		(1.00)	(1.00)	(1.00)	
	(1.00)	(0.95)	(0.51)		(1.00)	(0.99)	(1.00)	
Constrained	0.1191	0.7216	0.0789	0.1556	0.1224	0.7263	0.0814	0.1913
	(0.00)	(0.84)	(0.96)		(0.00)	(0.51)	(0.80)	
	(0.75)	(0.90)	(0.88)		(0.40)	(0.39)	(0.39)	
Regularization	0.1159	0.7732	0.0829	0.2438	0.1179	0.4055	0.0409	0.3535
	(0.00)	(0.75)	(0.92)		(0.00)	(0.87)	(0.94)	
	(0.37)	(0.98)	(0.99)		(0.21)	(0.42)	(0.44)	

Table 8: Portfolios Using Option-Implied Information about Expected Returns

In this table, we evaluate the performance of mean-variance portfolios that rely on option-implied characteristics about expected returns using the parametric-portfolio 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 a particular option-implied characteristic rather than the traditional FF and MOM characteristics. The four quantities we consider are the historical volatility risk premium (HVRP), the spread between implied and realized volatilities (IRVS), model-free implied skewness (MFIS), and the spread between call and put implied volatilities (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 FF and MOM 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+FF+MOM" portfolio, with the null hypothesis being that the portfolio being evaluated is worse than the benchmark.

	100 stocks				561 stocks			
Strategy	std	sr	ce	trn	std	sr	ce	trn
			Panel A: Daily rebalancing					
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7490	0.1155	0.0144
1/N + FF + MOM	0.1636	1.0351	0.1559	0.0805	0.1796	0.7928	0.1263	0.0710
	(1.00)	(0.02)	(0.01)		(1.00)	(0.15)	(0.08)	
	(0.50)	(0.50)	(0.50)		(0.50)	(0.50)	(0.50)	
1/N + HVRP	0.1532	1.0615	0.1509	0.0296	0.1608	0.8950	0.1310	0.0320
	(0.00)	(0.00)	(0.03)		(0.00)	(0.01)	(0.07)	
	(0.00)	(0.35)	(0.66)		(0.00)	(0.09)	(0.37)	
1/N + IRVS	0.1621	0.9539	0.1415	0.0850	0.1737	0.8123	0.1260	0.0915
	(1.00)	(0.22)	(0.17)		(0.05)	(0.01)	(0.02)	
	(0.05)	(0.92)	(0.93)		(0.00)	(0.35)	(0.52)	
1/N + MFIS	0.1647	1.1903	0.1825	0.1733	0.1809	1.0013	0.1648	0.1774
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.93)	(0.00)	(0.00)		(0.90)	(0.00)	(0.00)	
1/N + CPVS	0.1610	1.2749	0.1923	0.2572	0.1758	1.2336	0.2014	0.2824
	(0.61)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.00)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
1/N + HVRP + IRVS + MFIS + CPVS	0.1598	1.4093	0.2125	0.3661	0.1704	1.2437	0.1974	0.3003
	(0.06)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
	(0.00)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
1/N + FF + MOM + HVRP	0.1605	1.0006	0.1478	0.0971	0.1720	0.7698	0.1176	0.0853
	(0.35)	(0.08)	(0.09)		(0.00)	(0.33)	(0.40)	
	(0.00)	(0.85)	(0.93)		(0.00)	(0.72)	(0.87)	
1/N + FF + MOM + IRVS	0.1614	1.0325	0.1537	0.1349	0.1721	0.8163	0.1257	0.1041
	(0.76)	(0.02)	(0.02)		(0.00)	(0.05)	(0.09)	
	(0.00)	(0.53)	(0.64)		(0.00)	(0.28)	(0.53)	
1/N + FF + MOM + MFIS	0.1647	1.1758	0.1801	0.2117	0.1800	0.9502	0.1549	0.1790
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.98)	(0.00)	(0.00)		(0.76)	(0.00)	(0.00)	
1/N + FF + MOM + CPVS	0.1632	1.3756	0.2112	0.3788	0.1766	1.2061	0.1974	0.3102
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.32)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
1/N + FF + MOM + HVRP + IRVS + MFIS + CPVS	0.1628	1.3884	0.2129	0.4221	0.1758	1.2432	0.2031	0.3318
	(1.00)	(0.00)	(0.00)		(0.98)	(0.00)	(0.00)	
	(0.19)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	

... *Table 8 continued*

Strategy	100 stocks std	sr	ce	trn	561 stocks std	sr	ce	trn
Panel B: Weekly rebalancing								
1/N	0.1616	0.9214	0.1358	0.0290	0.1766	0.7384	0.1148	0.0325
1/N + FF + MOM	0.1661	0.9817	0.1492	0.0885	0.1820	0.7880	0.1268	0.0815
	(1.00)	(0.06)	(0.03)		(1.00)	(0.11)	(0.06)	
	(0.50)	(0.50)	(0.50)		(0.50)	(0.50)	(0.50)	
1/N + HVRP	0.1543	1.0437	0.1492	0.0441	0.1641	0.8727	0.1297	0.0477
	(0.00)	(0.00)	(0.02)		(0.00)	(0.00)	(0.04)	
	(0.00)	(0.14)	(0.52)		(0.00)	(0.10)	(0.41)	
1/N + IRVS	0.1623	0.9360	0.1388	0.0915	0.1779	0.7634	0.1200	0.0993
	(0.88)	(0.26)	(0.22)		(0.93)	(0.13)	(0.10)	
	(0.02)	(0.84)	(0.91)		(0.03)	(0.71)	(0.78)	
1/N + MFIS	0.1653	1.0614	0.1618	0.1767	0.1845	0.8392	0.1378	0.1820
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.31)	(0.03)	(0.04)		(0.91)	(0.13)	(0.10)	
1/N + CPVS	0.1616	1.0365	0.1545	0.2589	0.1776	0.8796	0.1404	0.2844
	(0.50)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.01)	(0.11)	(0.25)		(0.02)	(0.02)	(0.06)	
1/N + HVRP + IRVS	0.1602	1.1507	0.1715	0.3674	0.1739	0.9223	0.1453	0.3022
+ MFIS + CPVS	(0.08)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
	(0.00)	(0.00)	(0.00)		(0.00)	(0.00)	(0.02)	
1/N + FF + MOM + HVRP	0.1641	0.9572	0.1435	0.1041	0.1746	0.7568	0.1168	0.0944
	(0.95)	(0.20)	(0.14)		(0.11)	(0.33)	(0.39)	
	(0.02)	(0.86)	(0.93)		(0.00)	(0.85)	(0.95)	
1/N + FF + MOM + IRVS	0.1652	0.9607	0.1450	0.1397	0.1754	0.7683	0.1193	0.1109
	(0.99)	(0.17)	(0.10)		(0.20)	(0.21)	(0.25)	
	(0.21)	(0.75)	(0.79)		(0.00)	(0.73)	(0.89)	
1/N + FF + MOM + MFIS	0.1665	1.0415	0.1595	0.2147	0.1831	0.8050	0.1306	0.1833
	(1.00)	(0.00)	(0.00)		(1.00)	(0.01)	(0.00)	
	(0.68)	(0.01)	(0.01)		(0.91)	(0.26)	(0.22)	
1/N + FF + MOM + CPVS	0.1659	1.0645	0.1628	0.3800	0.1806	0.8794	0.1425	0.3121
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.45)	(0.01)	(0.01)		(0.22)	(0.01)	(0.02)	
1/N + FF + MOM + HVRP	0.1653	1.0795	0.1647	0.4231	0.1798	0.8973	0.1452	0.3335
+ IRVS+ MFIS + CPVS	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.28)	(0.00)	(0.00)		(0.08)	(0.00)	(0.00)	

... Table 8 continued

Strategy	100 stocks std	sr	ce	trn	561 stocks std	sr	ce	trn
Panel C: Monthly rebalancing								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7337	0.1107	0.0673
1/N + FF + MOM	0.1614	0.9938	0.1472	0.1098	0.1788	0.7576	0.1193	0.1072
	(1.00)	(0.09)	(0.03)		(0.99)	(0.26)	(0.14)	
	(0.50)	(0.50)	(0.50)		(0.50)	(0.50)	(0.50)	
1/N + HVRP	0.1453	1.0675	0.1446	0.0724	0.1559	0.8990	0.1279	0.0797
	(0.00)	(0.00)	(0.03)		(0.00)	(0.00)	(0.02)	
	(0.00)	(0.10)	(0.60)		(0.00)	(0.01)	(0.24)	
1/N + IRVS	0.1536	0.9572	0.1352	0.1096	0.1718	0.7570	0.1153	0.1212
	(0.76)	(0.20)	(0.15)		(0.82)	(0.12)	(0.08)	
	(0.01)	(0.79)	(0.92)		(0.05)	(0.54)	(0.70)	
1/N + MFIS	0.1572	1.0392	0.1510	0.1888	0.1808	0.7819	0.1250	0.1975
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.09)	(0.14)	(0.30)		(0.71)	(0.30)	(0.24)	
1/N + CPVS	0.1541	0.9603	0.1361	0.2661	0.1720	0.7722	0.1180	0.2933
	(0.95)	(0.08)	(0.02)		(0.98)	(0.00)	(0.00)	
	(0.01)	(0.79)	(0.90)		(0.04)	(0.38)	(0.56)	
1/N + HVRP + IRVS	0.1528	1.0702	0.1519	0.3730	0.1679	0.8343	0.1260	0.3104
+ MFIS + CPVS	(0.47)	(0.00)	(0.00)		(0.04)	(0.00)	(0.00)	
	(0.02)	(0.08)	(0.30)		(0.01)	(0.05)	(0.24)	
1/N + FF + MOM + HVRP	0.1593	1.0068	0.1476	0.1233	0.1689	0.7517	0.1126	0.1176
	(0.99)	(0.05)	(0.02)		(0.26)	(0.33)	(0.40)	
	(0.08)	(0.28)	(0.46)		(0.00)	(0.60)	(0.89)	
1/N + FF + MOM + MFIS	0.1603	1.0163	0.1500	0.2253	0.1795	0.7452	0.1176	0.1979
	(1.00)	(0.01)	(0.00)		(1.00)	(0.30)	(0.10)	
	(0.33)	(0.19)	(0.27)		(0.69)	(0.69)	(0.64)	
1/N + FF + MOM + CPVS	0.1611	1.0055	0.1490	0.3858	0.1760	0.7917	0.1238	0.3202
	(1.00)	(0.02)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.44)	(0.35)	(0.38)		(0.20)	(0.19)	(0.27)	
1/N + FF + MOM + IRVS	0.1594	0.9786	0.1432	0.1547	0.1701	0.7521	0.1133	0.1310
	(1.00)	(0.15)	(0.05)		(0.41)	(0.32)	(0.35)	
	(0.09)	(0.75)	(0.84)		(0.00)	(0.59)	(0.85)	
1/N + FF + MOM + HVRP	0.1604	1.0111	0.1493	0.4283	0.1758	0.7986	0.1249	0.3411
+ IRVS + MFIS + CPVS	(1.00)	(0.02)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.34)	(0.30)	(0.37)		(0.16)	(0.13)	(0.20)	

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

In this table, we evaluate the effect of transactions costs on the mean, Sharpe ratio, and certainty equivalent return of mean-variance portfolios that rely on option-implied characteristics using the parametric-portfolio methodology proposed in Brandt, Santa-Clara, and Valkanov (2009). The four quantities we consider are the historical volatility risk premium (HVRP), the spread between implied and realized volatilities (IRVS), model-free implied skewness (MFIS), and the spread between call and put implied volatilities (CPVS). We also consider the case where the investor uses all four option-implied characteristics together. In the second part of each panel, we consider the effect of considering these option-implied characteristics (individually and together) in addition to the FF and MOM 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 + FF + MOM" parametric portfolio. Negative transactions costs indicate that the portfolio using option-implied information performs worse than the benchmark and has higher turnover, and so needs to receive (rather than pay) the reported number in order to match the performance of the benchmark. A $\star$ indicates that the portfolio using option-implied information performs better than the benchmark and also has lower turnover than the benchmark, in which case the performance of the portfolio using option-implied information will be better than that of the benchmark for any level of transaction cost.

	100 stocks			561 stocks		
Strategy	mean	sr	ce	mean	sr	ce
Panel A: Daily rebalancing						
1/N + HVRP	0.32	0.47	0.34	0.30	0.49	0.35
	0.05	$\star$	0.04	$\star$	$\star$	$\star$
1/N + IRVS	0.03	0.02	0.03	0.05	0.06	0.05
	-1.27	-0.91	-1.24	-0.03	0.06	-0.00
1/N + MFIS	0.12	0.11	0.11	0.12	0.11	0.12
	0.11	0.11	0.11	0.15	0.14	0.14
1/N + CPVS	0.09	0.09	0.09	0.13	0.13	0.13
	0.08	0.09	0.08	0.14	0.15	0.14
1/N + HVRP+ IRVS	0.09	0.09	0.09	0.11	0.12	0.11
+ MFIS + CPVS	0.08	0.08	0.08	0.12	0.13	0.12
1/N + FF + MOM + HVRP	0.05	0.05	0.05	0.01	0.02	0.01
	-0.21	-0.12	-0.20	-0.28	-0.09	-0.24
1/N + FF + MOM + IRVS	0.06	0.05	0.06	0.04	0.05	0.05
	-0.02	-0.00	-0.02	-0.02	0.04	-0.01
1/N + FF + MOM + MFIS	0.09	0.08	0.09	0.10	0.09	0.10
	0.07	0.07	0.07	0.11	0.10	0.11
1/N + FF + MOM + CPVS	0.08	0.08	0.08	0.11	0.11	0.11
	0.07	0.07	0.07	0.12	0.12	0.12
1/N + FF + MOM + HVRP	0.07	0.07	0.07	0.11	0.11	0.11
+ IRVS + MFIS + CPVS	0.07	0.07	0.07	0.12	0.12	0.12

... Table 9 continued

Strategy	100 stocks mean	sr	ce	561 stocks mean	sr	ce
Panel B: Weekly rebalancing						
1/N + HVRP	1.61	2.26	1.76	1.68	2.46	1.95
	0.09	⋆	0.00	0.01	⋆	⋆
1/N + IRVS	0.10	0.08	0.09	0.16	0.13	0.16
	-7.31	-2.63	-6.56	-0.85	-0.45	-0.76
1/N + MFIS	0.36	0.31	0.35	0.33	0.25	0.31
	0.28	0.30	0.28	0.23	0.19	0.22
1/N + CPVS	0.16	0.16	0.16	0.20	0.20	0.20
	0.05	0.10	0.06	0.13	0.16	0.13
1/N + HVRP+ IRVS	0.21	0.22	0.21	0.22	0.24	0.23
+ MFIS + CPVS	0.15	0.19	0.16	0.15	0.21	0.17
1/N + FF + MOM + HVRP	0.22	0.16	0.20	0.06	0.10	0.07
	-0.77	-0.47	-0.72	-1.73	-0.66	-1.53
1/N + FF + MOM + IRVS	0.18	0.12	0.16	0.11	0.13	0.12
	-0.17	-0.13	-0.16	-0.58	-0.21	-0.51
1/N + FF + MOM + MFIS	0.26	0.22	0.25	0.23	0.16	0.21
	0.16	0.16	0.16	0.08	0.06	0.08
1/N + FF + MOM + CPVS	0.16	0.14	0.15	0.20	0.18	0.20
	0.09	0.09	0.09	0.13	0.14	0.14
1/N + FF + MOM + HVRP	0.15	0.13	0.15	0.20	0.19	0.20
+ IRVS + MFIS + CPVS	0.09	0.10	0.09	0.14	0.15	0.15

... Table 9 continued

Strategy	100 stocks mean	100 stocks sr	100 stocks ce	561 stocks mean	561 stocks sr	561 stocks ce
Panel C: Monthly rebalancing						
1/N + HVRP	7.25	9.34	7.94	10.05	11.62	11.49
	1.17	⋆	0.60	⋆	⋆	⋆
1/N + IRVS	0.52	0.45	0.51	0.74	0.64	0.72
	731.31	-6.95	68.95	-3.21	-0.04	-2.47
1/N + MFIS	1.26	1.03	1.22	1.03	0.58	0.93
	0.32	0.73	0.40	0.55	0.42	0.53
1/N + CPVS	0.16	0.13	0.16	0.28	0.25	0.27
	-0.66	-0.27	-0.60	-0.12	0.11	-0.06
1/N + HVRP + IRVS	0.52	0.54	0.53	0.51	0.58	0.53
+ MFIS + CPVS	0.10	0.37	0.15	0.19	0.52	0.27
1/N + FF + MOM + HVRP	2.16	1.45	2.02	0.27	0.50	0.32
	0.01	1.08	0.22	-6.85	-0.50	-5.46
1/N + FF + MOM + IRVS	1.06	0.55	0.97	0.34	0.41	0.35
	-0.82	-0.42	-0.75	-2.65	-0.27	-2.11
1/N + FF + MOM + MFIS	0.96	0.63	0.90	0.54	0.14	0.44
	0.19	0.26	0.20	-0.16	-0.21	-0.17
1/N + FF + MOM + CPVS	0.46	0.27	0.43	0.46	0.34	0.43
	0.05	0.06	0.05	0.15	0.24	0.18
1/N + FF + MOM + HVRP	0.41	0.26	0.39	0.46	0.35	0.44
+ IRVS + MFIS + CPVS	0.05	0.07	0.05	0.18	0.26	0.20

Figure 1: Volatilities: Realized, historical, and implied

In this figure, we plot the historical volatility based on the past 250 days (solid blue line), past 750 days (dot-dashed blue line), model-free implied volatility (dashed red line), risk-premium-corrected model-free implied volatility (solid pink line), and the 30-day realized volatility (thick black line). The figure is based on the cross-sectional equally-weighted average volatilities across 561 stocks. The figure shows that risk-premium-correct implied volatility tracks realized volatility quite closely. The model-free implied volatility (without any risk-premium correction) tracks realized volatility, but there is a distinct gap between the two. The gap between the historical 250-day volatility is larger, and this gap is even larger for the historical 750-day volatility. Note also that all these volatility series have different levels of variability: the implied and risk-premium-corrected implied volatilities are slightly more variable than the realized volatility, while the 750-day historical volatility is the smoothest.

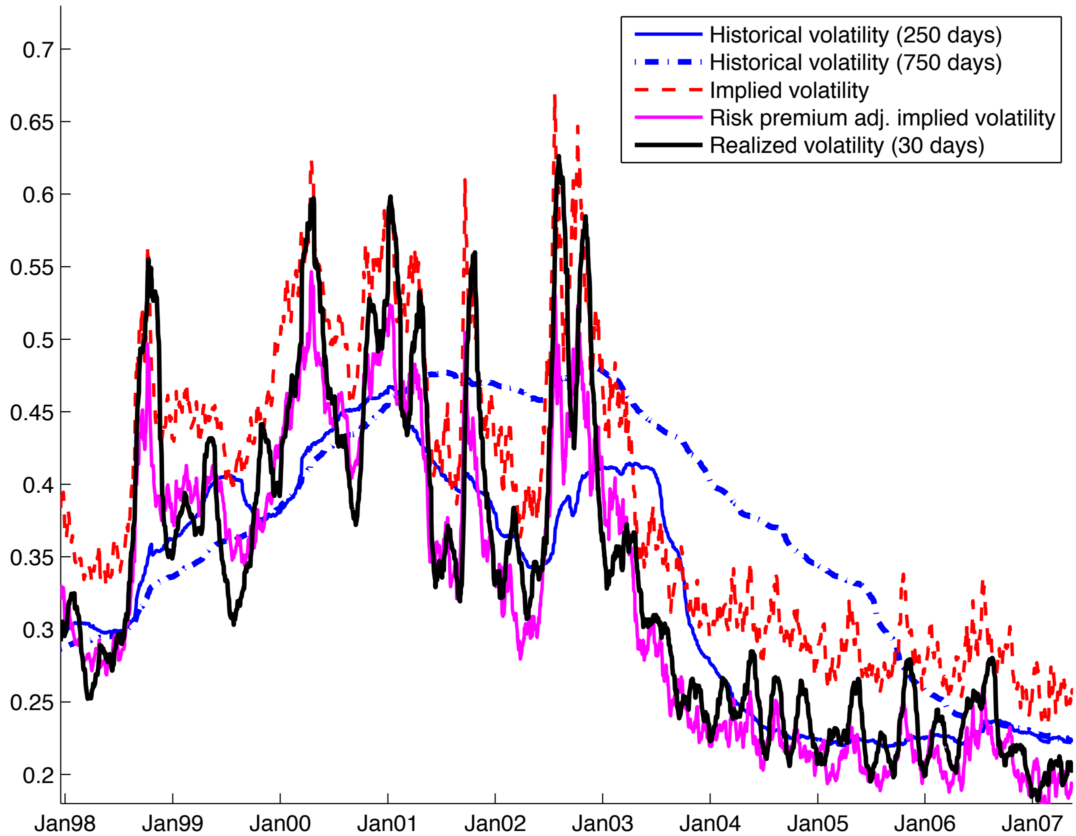

Figure 2: Correlations: Realized, historical, and implied

In this figure, we plot the historical correlation based on the past 250 days (solid blue line), past 750 days (dashed blue line), implied correlation (solid red line), and 30-day realized correlation (thick black line). The plot is based on the cross-sectional equally-weighted average of average correlations across 561 stocks. There are two observations about these series: first, implied correlation follows the level of realized correlation much more closely than historical correlation; two, implied correlation is much more volatile than realized correlation, while historical correlation is even smoother.

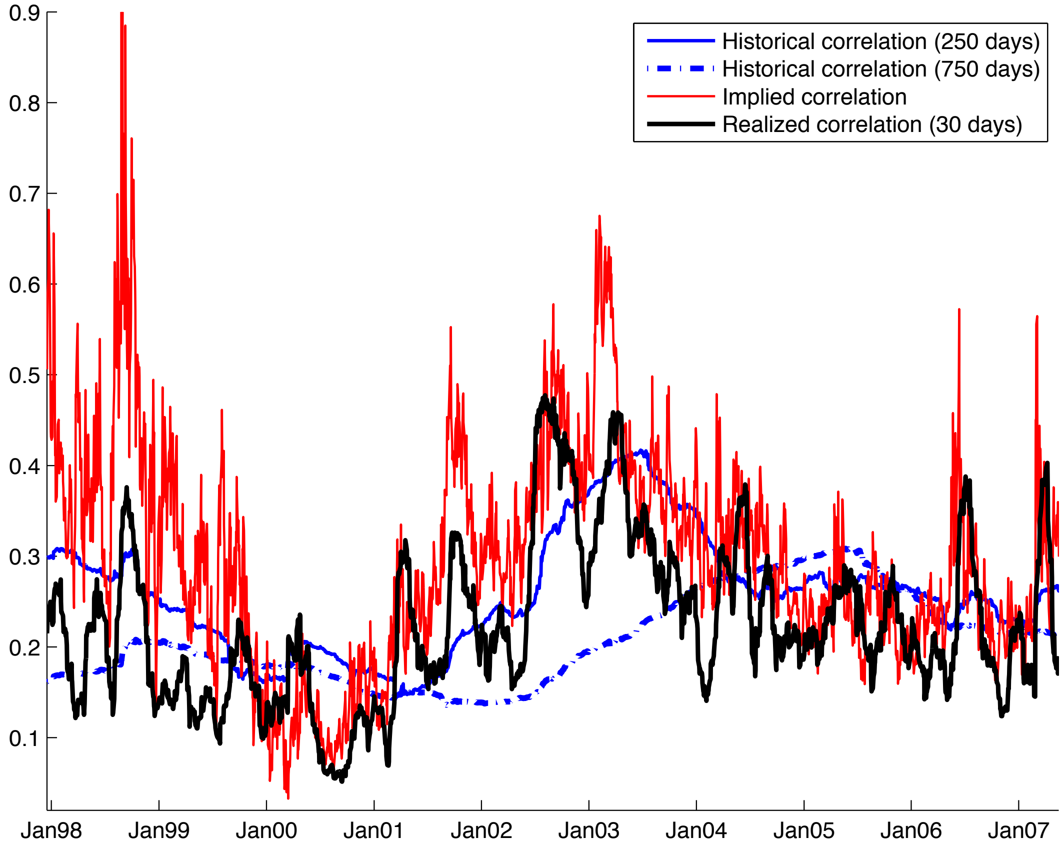

References

Aït-Sahalia, Y., and M. W. Brandt, 2008, "Consumption and Portfolio Choice with Option-Implied State Prices," NBER Working Paper.

Aït-Sahalia, Y., and A. W. Lo, 1998, "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices," *Journal of Finance*, 53, 499–547.

Andersen, T. G., T. Bollerslev, P. Christoffersen, and F. X. Diebold, 2006, "Volatility and Correlation Forecasting," in Graham Elliott, Clive W. J. Granger, and Allan G. Timmermann (ed.), *Handbook of Economic Forecasting*, North Holland, Amsterdam.

Andersen, T. G., T. Bollerslev, and F. X. Diebold, 2009, "Parametric and Nonparametric Measurement of Volatility," in Yacine Aït-Sahalia, and Lars Peter Hansen (ed.), *Handbook of Financial Econometrics*, North Holland, Amsterdam.

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.

Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard, 2009, "Realized Kernels in Practice: Trades and Quotes," *Econometrics Journal*, 12, C1–C32.

Barndorff-Nielsen, O. E., and N. Shephard, 2002, "Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models," *Journal of Royal Statistical Society Series B*, 64, 253–280.

Benzoni, L., 1998, "Pricing Options Under Stochastic Volatility: An Econometric Analysis," Working Paper, Northwestern University.

Best, M. J., and R. R. Grauer, 1991, "On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results," *Review of Financial Studies*, 4, 315–342.

Black, F., and M. S. Scholes, 1972, "The Valuation of Option Contracts and a Test of Market Efficiency," *Journal of Finance*, 27, 399–417.

Bliss, R. R., and N. Panigirtzoglou, 2004, "Option-Implied Risk Aversion Estimates," *Journal of Finance*, 59, 407–446.

Bollerslev, T., 1986, "Generalized Autoregressive Conditional Heteroscadicity," *Journal of Econometrics*, 31, 34–105.

Bollerslev, T., R. Y. Chou, and K. Kroner, 1992, "ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 307–327.

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. E. Tauchen, and H. Zhou, 2008, "Expected Stock Returns and Variance Risk Premia," forthcoming in *Review of Financial Studies*.

Boyer, B. H., T. Mitton, and K. Vorkink, 2009, "Expected Idiosyncratic Skewness," forthcoming in *Review of Financial Studies*.

Brandt, M. W., and P. Santa-Clara, 2005, "Dynamic Portfolio Selection by Augmenting the Asset Space," Working paper, UCLA.

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.

Britten-Jones, M., and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance*, 55, 839–866.

Broadie, M., 1993, "Computing Efficient Frontiers using Estimated Parameters," *Annals of Operations Research*, 45, 21–58.

Brown, S. J., 1990, "Estimating Volatility," in Stephen Figlewski, William L. Silber, and Martin G. Subrahmanyam (ed.), *Financial Options: From Theory to Practice*, Business One-Irwin, Homewood, IL.

Buss, A., and G. Vilkov, 2008, "Option-Implied Correlation and Factor Betas Revisited," Working Paper, Goethe University, Frankfurt.

Carr, P., and L. Wu, 2009, "Variance Risk Premiums," *Review of Financial Studies*, 22, 1311–1341.

Carr, P. P., and D. B. Madan, 1998, "Towards a Theory of Volatility Trading," in Robert A. Jarrow (ed.), *Volatility: New Estimation Techniques for Pricing Derivatives*, RISK Publications, London.

Carr, P. P., and D. B. Madan, 2001, "Optimal Positioning in Derivative Securities," *Quantitative Finance*, 1, 19–37.

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.

Chernov, M., 2007, "On the Role of Risk Premia in Volatility Forecasting," *Journal of Business & Economic Statistics*, 25, 411–426.

Chernov, M., and E. Ghysels, 2000, "A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation," *Journal of Financial Economics*, 56, 407–458.

Chopra, V. K., and W. T. Ziemba, 1993, "The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice," *Journal of Portfolio Management*, 19, 6–11.

Christensen, B. J., and N. R. Prabhala, 1998, "The Relation Between Implied and Realized Volatility," *Journal of Financial Economics*, 50, 125–150.

Christoffersen, P., and B.-Y. Chang, 2009, "Option-Implied Measures of Equity Risk," Working Paper, McGill University.

Corsi, F., G. Zumbach, U. Müller, and M. Dacorogna, 2001, "Consistent High-Precision Volatility From High-Frequency Data," *Economic Notes*, 30, 183–204.

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.

Dumas, B., 1995, "The Meaning of the Implicit Volatility Function in Case of Stochastic Volatility," Working Paper, HEC, Paris.

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.

Engle, R. F., 1982, "Autoregressive Conditional Heteroscadicity with Estimates of the Variance of U.K. Inflation," *Econometrica*, 50, 987–1008.

Engle, R. F., 1993, "Statistical Models for Financial Volatility," *Financial Analysts Journal*, 49, 72–78.

Engle, R. F., 2002, "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business and Economic Statistics*, 20, 339–350.

Epps, T. W., 1979, "Comovements in Stock Prices in the Very Short Run," *Journal of the American Statistical Association*, 74, 291–298.

Fama, E. F., and K. R. French, 1992, "The Cross-Section of Expected Stock Returns," *Journal of Finance*, 47, 427–465.

Frost, P. A., and J. E. Savarino, 1986, "An Empirical Bayes Approach to Efficient Portfolio Selection," *Journal of Financial and Quantitative Analysis*, 21, 293–305.

Frost, P. A., and J. E. Savarino, 1988, "For Better Performance Constrain Portfolio Weights," *Journal of Portfolio Management*, 15, 29–34.

Goyal, A., and A. Saretto, 2009, "Option Returns and Volatility Mispricing," forthcoming in *Journal of Financial Economics*.

Harvey, C. R., and A. Siddique, 2000, "Conditional Skewness in Asset Pricing Tests," *Journal of Finance*, 55, 1263–1295.

Haugen, R. A., and N. L. Baker, 1991, "The Efficient Market Inefficiency of Capitalization-Weighted Stock Portfolios," *Journal of Portfolio Management*, Spring, 35–40.

Heston, S., 1993, "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327–343.

Jackwerth, J. C., 2000, "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, 13, 433–451.

Jackwerth, J. C., and M. Rubinstein, 1996, "Recovering Probability Distributions from Option Prices," *Journal of Finance*, 51, 1611–1632.

Jacod, J., 1994, "Limit of Random Measures Associated with the Increments of a Brownian Semimartingale," *Preprint University of Paris VI.*

Jacod, J., and P. Protter, 1998, "Asymptotic Error Distributions for the Euler Method for Stochastic Differential Equations," *Annals of Probability*, 26, 267–307.

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.

Jiang, G. J., and Y. S. Tian, 2005, "The Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, 18, 1305–1342.

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, 2009, "Asset Allocation with Option-Implied Distributions: A Forward-Looking Approach," SSRN-id1288103.

Kraus, A., and R. H. Litzenberger, 1976, "Skewness Preference and the Valuation of Risk Assets," *Journal of Finance*, 31, 1085–1100.

Lamoureux, C. G., and W. D. Lastrapes, 1993, "Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities," *Review of Financial Studies*, 6, 293–326.

Latane, H. A., and R. J. Rendleman, Jr., 1976, "Standard Deviations of Stock Price Ratios Implied in Option Prices," *Journal of Finance*, 31, 369–381.

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.

Litterman, R., 2003, *Modern Investment Management: An Equilibrium Approach*, Wiley, New York.

MacKinlay, A. C., and Ľ. Pástor, 2000, "Asset Pricing Models: Implications for Expected Returns and Portfolio Selection," *Review of Financial Studies*, 13, 883–916.

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.

Merton, R. C., 1980, "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323–361.

Michaud, R. O., 1989, "The Markowitz Optimization Enigma: Is Optimized Optimal?," *Financial Analysts Journal*, 45, 31–42.

Newey, W. K., and K. D. West, 1987, "A simple, positive-semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica*, 55, 703–708.

Panigirtzoglou, N., and G. Skiadopoulos, 2004, "A New Approach to Modeling the Dynamics of Implied Distributions: Theory and Evidence from the S&P 500 Options," *Journal of Banking and Finance*, 28, 1499–1520.

Pástor, Ľ., 2000, "Portfolio Selection and Asset Pricing Models," *Journal of Finance*, 55, 179–223.

Pástor, Ľ., and R. F. Stambaugh, 2000, "Comparing Asset Pricing Models: An Investment Perspective," *Journal of Financial Economics*, 56, 335–381.

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, "Option-Implied Skewness as a Stock-Specific Sentiment Measure," Working Paper, Goethe University, Frankfurt, SSRN id1301648.

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, 2009, "What Does Individual Option Volatility Smirk Tell Us About Future Equity Returns?," forthcoming in *Journal of Financial and Quantitative Analysis*.

Zhang, L., P. A. Mykland, and Y. Aït-Sahalia, 2005, "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data," *Journal of the American Statistical Association*, 100, 1394–1411.

Zhou, B., 1996, "High-Frequency Data and Volatility in Foreign-Exchange Rates," *Journal of Business and Economic Statistics*, 14, 45–52.

Zumbach, G., 2009, "Inference on Multivariate ARCH Processes with Large Sizes," *SSRN eLibrary*.