Improving Portfolio Selection Using Option-Implied Volatility and Skewness*

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

This version: September 30, 2010

Abstract

Our objective in this paper is to examine whether one can use option-implied information to improve the selection of 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, while using option-implied volatility and correlation does *not* improve significantly the portfolio volatility, Sharpe ratio, and certainty-equivalent return, exploiting information contained in the volatility risk premium and option-implied skewness increases substantially both the Sharpe ratio and certainty-equivalent return, although this is accompanied by higher turnover. Moreover, the volatility risk premium and option-implied skewness help improve not just the performance of mean-variance portfolios, but also the performance of parametric portfolios developed in Brandt, Santa-Clara, and Valkanov (2009).

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 also received helpful comments and suggestions from Alexander Alekseev, Luca Benzoni, Michael Brandt, Mike Chernov, Engelbert Dockner, Bernard Dumas, Wayne Ferson, Lorenzo Garlappi, Nicolae Gârleanu, Rene Garcia, Amit Goyal, Jakub Jurek, Nikunj Kapadia, Ralph Koijen, Vasant Naik, Stavros Panageas, Andrew Patton, Marcel Rindisbacher, Paulo Rodrigues, Pedro Santa-Clara, Christian Schlag, Bernd Scherer, George Skiadopoulos, Josef Zechner, and seminar participants at Goethe University Frankfurt, AHL (Man Investments), London School of Economics, University of Mainz, University of Piraeus, Vienna University of Economics and Business Administration, Duke-UNC Asset Pricing Conference, Financial Econometrics Conference at Toulouse School of Economics, and meetings of the Western Finance Association.

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

Improving Portfolio Selection Using Option-Implied Volatility and Skewness

This version: September 30, 2010

Abstract

Our objective in this paper is to examine whether one can use option-implied information to improve the selection of 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, while using option-implied volatility and correlation does *not* improve significantly the portfolio volatility, Sharpe ratio, and certainty-equivalent return, exploiting information contained in the volatility risk premium and option-implied skewness increases substantially both the Sharpe ratio and certainty-equivalent return, although this is accompanied by higher turnover. Moreover, the volatility risk premium and option-implied skewness help improve not just the performance of mean-variance portfolios, but also the performance of parametric portfolios developed in Brandt, Santa-Clara, and Valkanov (2009).

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

JEL: G11, G12, G13, G17

1 Introduction

To determine the optimal 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 for this estimation, 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* option-implied moments of stock-return distributions.[3] The main contribution of our work is to demonstrate empirically *how* one can use option-implied information to improve portfolio selection with a large number of stocks, and to document *which aspects* of option-implied information are particularly useful. Specifically, we study how one can use option-implied volatility, correlation, skewness, and the volatility risk premium to adjust the volatility and correlation of stock returns in order to improve the out-of-sample performance of static portfolios. We find that the improvement in portfolio performance from using option-implied volatilities and correlations is small, contrary to what one may have expected. However, the use of option-implied skewness and the volatility risk premium can lead to substantial improvements in Sharpe ratios and certainty-equivalent returns (even when shortsales are constrained), but this is accompanied by an increase in portfolio turnover.

It is well known that it is much more difficult to estimate expected returns than second moments of stock returns (Merton, 1980), and, as a result, much recent research has focused on minimum-variance portfolios, which rely solely on estimates of covariances. In fact, 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. For example, 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. In view of this, we focus our attention on global minimum variance portfolios in this study." Just like Jagannathan and Ma (2003), we too focus on minimum-variance portfolios. However, the methodology we develop applies

[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 daily data rather than monthly data (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, and Uppal, 2009), and using stock-return characteristics such as size, momentum, and the book-to-market ratio (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.

also to mean-variance portfolios, to portfolios obtained from the maximization of more general utility functions, and to the parametric portfolios of Brandt, Santa-Clara, and Valkanov (2009).

To determine the minimum-variance portfolio, one needs to estimate for each stock its volatility and correlations with all the other stocks. We undertake our analysis in three steps. In step one, we determine the optimal portfolio using volatilities implied by option prices. In step two, we find the optimal portfolio using correlations implied by option prices. In step three, we find the optimal portfolio when volatilities are scaled based on option-implied skewness and the volatility risk premium. We summarize below the findings from these three steps.

In the first step, we find that using option-implied volatilities to compute the optimal portfolio does *not* lead to a substantial reduction in the out-of-sample portfolio volatility or to an increase in the Sharpe ratio and certainty-equivalent return. This is surprising because there is a large literature that documents that implied volatility can predict stock-return volatility better than sample volatility (see, for example, Blair, Poon, and Taylor (2001) and Jiang and Tian (2005)). We explain that there are two reasons why option-implied volatility fails to improve portfolio performance. First, the implied volatilities are estimators with large variances because they are based exclusively on current option prices. Second, because the implied volatilities estimate the risk-neutral volatilities, they are biased estimators of the real-world (objective) volatilities, with the gap between the two being the volatility risk premium, as explained in Chernov (2007). However, we find that even the portfolios based on the risk-premium-corrected implied volatilities attain an out-of-sample portfolio volatility that is only about 5% lower than the traditional portfolios based on the historical stock-return data, while the improvement in Sharpe ratio is still insignificant.

In the second step, we find that the benefits from using option-implied correlations are even smaller than the gains from using option-implied volatilities. To understand the reason for this, note that the covariance matrix that improves portfolio performance will be the one that contains enough information about future covariances *and* is stable (with a small condition number and, correspondingly, less volatile portfolio weights). Our empirical results indicate that, while option-implied volatilities and 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 portfolio turnover.

Finally, in the third step, we study how two other sources of option-implied information can be used to improve portfolio selection. The first is the historical volatility risk premium, and its choice is motivated by the empirical regularity documented by Bali and Hovakimian (2009) and Goyal and Saretto (2009) that assets with high volatility risk premium tend to outperform those with low volatility risk premium. Our empirical evidence shows that portfolios based on volatilities scaled by the volatility risk premia outperform traditional portfolios. The second source of information is *option-implied* skewness, whose choice is motivated by the finding in Rehman and Vilkov (2009) that stocks with high option-implied skewness outperform stocks with low option-implied skewness.[4] We find that portfolios that use volatilities scaled by implied skewness achieve significantly higher Sharpe ratios than those of traditional portfolios (even in the presence of shortsale constraints), but these gains are accompanied by higher portfolio turnover. The volatility risk premium and implied skewness improve the performance also of the Brandt, Santa-Clara, and Valkanov (2009) parametric portfolios, 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).

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; two portfolio rebalancing periods: daily and monthly; four performance metrics: portfolio volatility, Sharpe ratio, certainty-equivalent return, and turnover; and nine 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 in the covariance matrix set equal to zero; minimum-variance portfolio with correlations in the covariance matrix set equal to the mean correlation across all asset pairs; the parametric portfolios developed in Brandt, Santa-Clara, and Valkanov (2009); and the portfolio based on the maximization of expected utility.

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).[5] The focus of

[4] 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).

[5] 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 (2008), 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.

our work is to investigate how the information implied by option prices can be used to improve portfolio selection. There are two other papers that study this. The first, by Aït-Sahalia and Brandt (2008), uses option-implied state prices to solve for the intertemporal consumption and portfolio choice problem, using the Cox and Huang (1989) martingale representation formulation, rather than the Merton (1971) dynamic-programming formulation. 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 divided into a number of short distinct sections. In Section 2, we provide a brief background to the portfolio selection problem. In Section 3, we describe the data on stock returns and options that we use. In Section 4, we explain the performance metrics we use to evaluate portfolios. The construction and performance of our benchmark portfolios that do *not* use option-implied information are described in Section 5. How we compute the quantities implied by option prices that we use for portfolio selection is explained in Section 6. Our main findings about the performance of various portfolios that use option-implied information are given in Section 7. The robustness checks we undertake are described in Section 8, and we conclude in Section 9. Appendix A explains how to compute variances and covariances for high-frequency intraday data; Appendix B describes the method used for shrinkage and regularization of the covariance matrix; and Appendix C explains the construction of model-free option-implied moments.

2 Portfolio Selection Problem

The classic mean-variance optimization problem can be written as

$$\min_{w} \quad w^\top \hat{\Sigma} w - \frac{1}{\gamma} w^\top \hat{\mu}, \tag{1}$$

$$\text{s.t.} \quad w^\top e = 1, \tag{2}$$

where $w \in \mathbb{R}^N$ is the vector of portfolio weights invested in stocks, $\hat{\Sigma} \in \mathbb{R}^{N \times N}$ is the estimated covariance matrix, γ is the investor's risk aversion, $\hat{\mu} \in \mathbb{R}^N$ is the estimated vector of expected returns, and $e \in \mathbb{R}^N$ is the vector of ones. The objective in (1) is to minimize the difference between the variance of the portfolio return, $w^\top \hat{\Sigma} w$, and its mean, $w^\top \hat{\mu}$, after taking into account the risk aversion of the agent. The constraint $w^\top e = 1$ in (2) ensures that the portfolio weights sum to one; we consider the case without the risk-free asset because our objective is to explore how to use option-implied information to select the portfolio of risky stocks.

In light of our discussion in the introduction about the difficulty in forecasting expected returns, we assume all expected returns to be equal to the same constant, $\hat{\mu}_i = \bar{\mu}$. Then the mean-variance objective in (1) reduces to minimizing the variance of the portfolio return:

$$\min_{w} \quad w^\top \hat{\Sigma} w, \tag{3}$$

subject to the constraint in (2). The solution to the above problem is:

$$w_{min} = \frac{\hat{\Sigma}^{-1} e}{e^\top \hat{\Sigma}^{-1} e}. \tag{4}$$

Note that the covariance matrix $\hat{\Sigma}$ in (4) can be decomposed into volatility and correlation matrices,

$$\hat{\Sigma} = \text{diag}(\hat{\sigma})\, \hat{\Omega}\, \text{diag}(\hat{\sigma}), \tag{5}$$

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 (4) 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.

3 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 policies, we also use data from Compustat. Our data for options are from IvyDB (OptionMetrics).

3.1 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.[6]

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")[7]; 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.[8] 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.

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,[9] and the second consists of 100 stocks out of the

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

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

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

[9] 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.

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.[10]

3.2 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.[11]

While we do not use option *prices* directly in our analysis, we wish to use option-based information only 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.[12]

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.[13] 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.[14]

When working with data on option prices and the volatility surface, for several calculations we need a proxy for the riskfree rate for the maturity of a particular option. For this, we use the certificate-of-deposit yields for maturities between one day and one year from the IvyDB and interpolate them linearly to get the appropriate yield.

3.3 Data on Stock Characteristics

For our analysis of the Brandt, Santa-Clara, and Valkanov (2009) parametric portfolios, we measure size (market value of equity) as the price of the stock per share multiplied by shares

[10] 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 8.1.

[11] 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.

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

[13] 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.

[14] 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.

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.[15]

4 Description of 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.[16] 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 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 two rebalancing intervals: daily and monthly. For the monthly rebalancing interval, we find the new set of weights daily, but hold that portfolio for 30 calendar days (21 trading days), which corresponds to the average of 21 daily returns; the advantage of this

[15] 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.

[16] 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.

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 month, when we consider a monthly holding period) 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:

$$\widehat{\mu}^{strategy} = \frac{1}{T-\tau}\sum_{t=\tau}^{T-1} r_{t+1}^{strategy}, \tag{6}$$

$$\widehat{\sigma}^{strategy} = \left(\frac{1}{T-\tau-1}\sum_{t=\tau}^{T-1}\left(r_{t+1}^{strategy} - \widehat{\mu}^{strategy}\right)^2\right)^{1/2}, \tag{7}$$

$$\widehat{\mathrm{SR}}^{strategy} = \frac{\widehat{\mu}^{strategy}}{\widehat{\sigma}^{strategy}}, \tag{8}$$

$$\widehat{\mathrm{ce}}^{strategy} = u^{-1}\left(\frac{1}{T-\tau}\sum_{t=\tau}^{T-1} u\left(r_{t+1}^{strategy}\right)\right), \tag{9}$$

where u denotes the power utility function with a relative risk aversion of $\gamma = 1$, 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 monthly rebalancing we make an additional adjustment, as in Politis and Romano (1994), to account for the autocorrelation arising from overlapping returns.[17]

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 under strategy *strategy*, $w_{j,t^+}^{strategy}$ the portfolio weight *before*

[17]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) than that of the benchmark portfolio n, that is, $\mathrm{H}_0 : \mu_i/\sigma_i - \mu_n/\sigma_n \leqslant 0$. 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. 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, $\mathrm{H}_0 : \sigma_i^2/\sigma_n^2 \geqslant 1$, 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).

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 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:

$$\text{Turnover} = \frac{1}{T-\tau-1} \sum_{t=\tau}^{T-1} \sum_{j=1}^{N} \left(\left| w_{j,t+1}^{strategy} - w_{j,t^+}^{strategy} \right| \right). \tag{10}$$

5 Benchmark Portfolios That Do Not Use Option Prices

For robustness, we consider five benchmark portfolios that are *not* based on option-implied information. These are: (i) the equally-weighted ($1/N$) portfolio; (ii) the unconstrained minimum-variance portfolio; (iii) the shortsale-constrained minimum-variance portfolio; (iv) the unconstrained minimum-variance portfolio with shrinkage of the covariance matrix; and (v) the unconstrained minimum-variance portfolio with all correlations set equal to zero. The construction of these benchmark portfolios is described below. In principle, one could also consider the mean-variance portfolio as a benchmark, but, as we discuss in Section 5.3, this performs much worse than the five benchmark portfolios listed above, and so we do not report the results for it in the tables. In addition to these five benchmark policies, we study also the parametric portfolios developed in Brandt, Santa-Clara, and Valkanov (2009), and the portfolios obtained from maximizing expected utility; these are discussed in Sections 8.5 and 8.6, respectively.

5.1 The Equally-Weighted Portfolio

For the "*equally-weighted*" ($1/N$) portfolio, one invests each period 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.[18]

5.2 Four Minimum-Variance Portfolios Based on Historical Returns

In addition to the $1/N$ portfolio, we also consider several minimum-variance portfolios. The first minimum-variance portfolio that we consider is w_{min}, which is given in (4) and is based on an estimate of the "*sample covariance*" matrix, that is, the realized volatilities and correlations.

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

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.[19]

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 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 (3) subject to the constraint in (2), after imposing the additional constraint that all weights have to be non-negative.

The next benchmark is the "*shrinkage*" portfolio, where we compute the minimum-variance portfolio weights *after* shrinking the covariance matrix.[20] 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). The reason for using regularization is that the distribution of intraday returns is different from that of daily returns and does not satisfy the assumptions of Ledoit and Wolf (2004a,b).[21]

We also considered two other methods proposed in the literature (see Elton, Gruber, and Spitzer (2006) and the references therein) for improving the behavior of the covariance matrix. 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.

[19]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 A.

[20]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 particular cases of the norm-constrained portfolios.

[21]Details of the shrinkage and regularization methods we use are provided in Appendix B, and the results for intraday data are summarized in Section 8.2.

In Table 1 we report the performance of the $1/N$ portfolio and four variants of the minimum-variance benchmark portfolio, all of which do *not* use data on option prices. In Panel A, we report the results for daily rebalancing, and in Panel B, we report the results when the portfolio is held for a month. Two p-values are reported in parenthesis under each performance metric. The first p-value is relative to the $1/N$ benchmark, and the second p-value in this table is relative to the "Sample-cov" benchmark. Each p-value is for the *one-sided* null hypothesis that the portfolio being evaluated is *worse* than the 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).

From Table 1, 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. The first set of p-values indicate that the volatilities of the minimum-variance portfolios are significantly lower than that of $1/N$; the second set of p-values indicate that the constrained and shrinkage portfolios have a lower portfolio volatility than the portfolio based on the sample covariance. The results for the data with 100 assets and in Panel B for "Monthly rebalancing" are similar.

However, the Sharpe Ratio (sr), certainty-equivalent return (ce), and turnover (trn) are typically better for the $1/N$ portfolio, 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; however, for both cases the differences are not statistically significant. 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 substantially lower turnover (which is true also in the tables that follow, where we use option-implied information).

5.3 The Mean-Variance Portfolio Based on Historical Returns

For completeness, we also discuss briefly the results for the *mean-variance* portfolios. Of the three variants of the mean-variance portfolio we consider, the first 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 three minimum-variance portfolios is less than 0.1350 (see Table 1), the volatility of the corresponding three mean-variance portfolios is always higher, and, in the absence of 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 two mean-variance portfolios. Finally, the turnover of the three mean-variance portfolios is higher than that of the minimum-variance portfolios. Consistent with the findings documented in the existing literature, we conclude that relative to the $1/N$ portfolio and also the minimum-variance benchmarks, the mean-variance portfolios perform much worse across all four metrics. Therefore, we do not report the performance of these portfolios.

6 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 option-implied information: (i) option-implied volatility and the volatility risk premium; (ii) option-implied correlation; and (iii) option-implied skewness.[22]

6.1 Option-Implied Volatility and Volatility Risk Premium

When option prices are available, an intuitive first step is to use this information to back out implied volatilities. In contrast to the model-specific Black-Scholes implied volatility, we use 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,

[22]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).

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 C.[23] 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 RMSE and R^2 to that when volatility based on historical data is used as a predictor. The prediction regression we estimate is $RV = \alpha + \beta \widehat{RV}$, where we regress the 30-days realized volatility (RV) on various volatility predictors ($\widehat{RV}$). The volatility predictors we consider 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 volatility risk premium. We compute the RMSE under the restrictions $\alpha = 0$ and $\beta = 1$. We see from Panel A of Table 2 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)

[23]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).

for a particular stock as the square root of the average variance risk premium for that stock for the past T trading days:[24]

$$\mathrm{HVRP}_t^2 = \frac{1}{T-\Delta t} \sum_{i=t-T+1}^{t-\Delta t} \frac{\mathrm{MFIV}_{i,i+\Delta t}^2}{\mathrm{RV}_{i,i+\Delta t}^2}. \tag{11}$$

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 (11), one can obtain the *prediction* of the future realized volatility, $\widehat{\mathrm{RV}}_t$:

$$\widehat{\mathrm{RV}}_{t,t+\Delta t} = \frac{\mathrm{MFIV}_{t,t+\Delta t}}{\mathrm{HVRP}_t}. \tag{12}$$

Panel A of Table 2 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.

6.2 Option-Implied Correlation

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

[24] 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.

write the variance of the portfolio, σ_p^2, as follows:

$$\sigma_p^2 = \sum_{1=1}^{N} w_i^2 \sigma_i^2 + \sum_{i=1}^{N} \sum_{j \neq i}^{N} w_i w_j \sigma_i \sigma_j \rho_{ij}. \tag{13}$$

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

$$\left(\sigma_p^Q\right)^2 = \sum_{i=1}^{N} w_i^2 \left(\sigma_i^Q\right)^2 + \sum_{i=1}^{N} \sum_{j \neq i}^{N} w_i w_j \sigma_i^Q \sigma_j^Q \rho_{ij}^Q. \tag{14}$$

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 way is to use the approach in Driessen, Maenhout, and Vilkov (2009), where it is assumed 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 (14):

$$\rho^Q = \frac{\left(\sigma_p^Q\right)^2 - \sum_{i=1}^{N} w_i^2 \left(\sigma_i^Q\right)^2}{\sum_{i=1}^{N} \sum_{j \neq i} w_i w_j \sigma_i^Q \sigma_j^Q}. \tag{15}$$

The resulting HOMIC 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:

$$CRP_{ij} = \rho_{ij}^Q - \rho_{ij}^P. \tag{16}$$

Substituting (16) into (14) gives:

$$\left(\sigma_p^Q\right)^2 = \sum_{i=1}^{N} w_i^2 \left(\sigma_i^Q\right)^2 + \sum_{i=1}^{N} \sum_{j \neq i}^{N} w_i w_j \sigma_i^Q \sigma_j^Q \left(\rho_{ij}^P + CRP_{ij}\right). \tag{17}$$

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

$$CRP_{ij} = \left(1 - \rho_{ij}^P\right) \times \psi^Q. \tag{18}$$

Using this assumption, we can identify the heterogeneous implied correlation matrix (HETIC) by first determining ψ^Q from (17):

$$\psi^Q = \frac{\left(\sigma_p^Q\right)^2 - \sum_{i=1}^N w_i^2 \left(\sigma_i^Q\right)^2 - \sum_{i=1}^N \sum_{j \neq i}^N w_i w_j \sigma_i^Q \sigma_j^Q \rho_{ij}^P}{\sum_{i=1}^N \sum_{j \neq i}^N w_i w_j \sigma_i^Q \sigma_j^Q \left(1 - \rho_{ij}^P\right)}, \tag{19}$$

and then obtaining each pairwise correlation by substituting (18) into Equation (16),

$$\rho_{ij}^Q = \rho_{ij}^P + \left(1 - \rho_{ij}^P\right) \times \psi^Q, \tag{20}$$

where ψ^Q is obtained from (19). The resulting HETIC 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 $corr = \alpha + \beta\, \widehat{corr}$, where we regress the 30-days realized correlation on the following correlation predictors ($\widehat{corr}$): (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 report in Panel B of Table 2 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.[25] 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

[25] 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 22.

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).

6.3 Option-Implied Skewness

The *model-free implied skewness* 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. Some researchers 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 (Xing, Zhang, and Zhao, 2009). 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 the implied skewness measured using the whole range of out-of-the money options. It has been shown in Rehman and Vilkov (2009) that risk-neutral skewness contains information about future stock returns above and beyond that in the simple measure of skew.

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 C.

7 Portfolios That Use Option-Implied Information

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. As explained above, to determine the minimum-variance portfolio weights, one needs to estimate for each stock its volatility and correlations with all other stocks. In Sections 7.1 and 7.2, we examine how option-implied volatilities and option-implied correlations can be used for portfolio selection. In Sections 7.3 and 7.4, we study how the volatility risk premium and option-implied skewness can be used for adjusting estimates of historical volatility in order to choose portfolios with superior out-of-sample performance.

In each of these sections, we use only one source of option-implied information at a time, in order to identify the magnitude of the gains from that particular source of information.

7.1 Portfolios Using Option-Implied Volatilities

Motivated by the findings in Section 6.1 about the predictive power of model-free implied volatilities, we use them in $\text{diag}(\hat{\sigma})$ to obtain the covariance matrix given in (5); 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, we then determine the minimum-variance portfolio in (4), 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). Effectively, we are computing the minimum-variance portfolio but with implied volatilities instead of historical volatilities in the covariance matrix.

The performance of these portfolios is reported in Table 3. 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 minimum-variance benchmark portfolio in Table 1.

Table 3 shows that the portfolios based on the model-free implied volatility have a lower volatility than the $1/N$ portfolio, but the Sharpe ratio, certainty-equivalent return, and turnover is better for the $1/N$ portfolio. Compared to the traditional minimum-variance portfolios in Table 1 that are based on historical returns data, we find that replacing the estimates of stock-return volatilities by their model-free implied counterparts often helps to reduce the out-of-sample volatility of the portfolio (especially for the data with 561 stocks), although the reduction is relatively small and often statistically insignificant. However, the Sharpe ratio, certainty-equivalent return, and turnover are superior for portfolios based only on historical returns.

There are two striking conclusions from Table 3, both of which are *negative*. First, the reduction in portfolio volatility is very small. Second, and more disconcerting, is the observation that the Sharpe ratios and also the certainty-equivalent returns are substantially *smaller* for the portfolios in Table 3 that are based on implied volatility, compared to those in Table 1, which rely only on historical estimates of volatility. One reason for these negative results is that when we derive the model-free implied volatility from options, we are getting 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.

Table 4 gives the results where volatility of stock returns is estimated as the model-free implied volatility *after* it is corrected for the risk premium. Comparing the results in this table to those in Table 3, 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. Comparing the results in Table 4 to the traditional minimum-variance portfolios in Table 1, we observe that the portfolios with the risk-premium-corrected implied volatilities attain a lower out-of-sample volatility, with the difference being around 6%–10% for monthly rebalancing and around 0–7% for daily rebalancing.

More importantly, comparing Table 4 with Table 3, 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 4, the $1/N$ portfolio has a higher Sharpe ratio and certainty-equivalent return, and substantially lower turnover (the exceptions are the Sharpe ratio for the data with 561 stocks for the zero-correlation portfolio in Panel A and the shortsale-constrained and zero-correlation portfolios in Panel B, though the p-value is not significant). The reason for the relatively poor Sharpe ratio and certainty-equivalent return 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.

7.2 Portfolios Using Option-Implied Correlations

Next, we investigate the benefit of using option-implied correlations in portfolio selection. In order to isolate the effect of using implied correlations, when computing portfolios we use volatilities from historical data. Recall that in Section 6.2 we described two approaches to estimate 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 (15); the performance of portfolios obtained from this approach is reported in Table 5. In the second approach, we allow for different correlations between assets, which are given by Equation (20); the performance of portfolios obtained from this approach is reported in Table 6.

Table 5, where we restrict correlations to be the same across all asset pairs, shows that of the three minimum-variance 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 both trading frequencies, and performs substantially better than the portfolios based on the sample covariance matrix and regularization.[26] However, even the best performing minimum-variance portfolio fails to achieve a Sharpe ratio that is significantly better than that of the $1/N$ portfolio. Even more surprising is the finding that the portfolios based on option-implied correlations fail to achieve a higher Sharpe ratio than the corresponding benchmark portfolios considered in Table 1, which do not use option-implied information.

Table 6, where we allow correlations to be heterogeneous across asset pairs, shows that of the three minimum-variance portfolios based on implied correlations that we consider, the portfolio based on "regularization" has the lowest portfolio volatility for both datasets and also for both 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 6, where we allow correlations to be heterogeneous across asset pairs, to those in Table 5, 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 higher turnover.

However, in both Tables 5 and 6, the three minimum-variance portfolios we consider typically fail to outperform the $1/N$ portfolio 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 Panel A, though the improvement in Sharpe ratio relative to the $1/N$ portfolio is not statistically significant. One

[26] Because we are using option-implied correlations, we do not know the distribution of returns for the resulting covariance matrix. Not knowing the distribution of returns, we cannot use the shrinkage results of Ledoit and Wolf (2004a,b), and so we use instead the regularization approach of Zumbach (2009). Details of the regularization method we use are given in Appendix B.

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.

7.3 Portfolios Using Volatility Risk Premium

In this section, we examine the effect of adjusting volatilities based on the volatility risk premium. Bali and Hovakimian (2009) and Goyal and Saretto (2009) have documented that assets with high volatility risk premium tend to outperform those with low volatility risk premium.[27] We now study whether this empirical regularity can be exploited to improve portfolio performance.

To incorporate the above finding for portfolio selection, we sort all the stocks into deciles using the characteristic "historical volatility risk premium," and then scale the volatility, $\hat{\sigma}$, of the top decile to $\hat{\sigma}(1-\delta)$ and of the bottom decile to $\hat{\sigma}(1+\delta)$.[28] This scaling, based on option-implied information, increases the weights of assets that are expected to outperform other stocks and decreases the weights of stocks that are expected to underperform. In our empirical implementation, we consider $\delta = \{0.1, 0.2, \ldots, 0.9\}$ but report the results only for two values, $\delta_1 = 0.40$ and $\delta_2 = 0.60$; the tables for the other values of δ are available from the authors. To compute the minimum-variance portfolio, we then use a covariance matrix with the scaled volatilities, along with correlations estimated from historical data.

Tables 7 and 8 give the results when the estimates of the historical volatilities have been scaled based on the volatility risk premium for each stock, using $\delta = 0.40$ and $\delta = 0.60$, respectively. The tables show that scaling the stock return volatilities improves performance: the Sharpe ratio and certainty-equivalent return are typically greater than for the benchmark minimum-variance portfolios in Table 1. Moreover, in most cases the Sharpe ratios exceed that of the $1/N$ portfolio (with several p-values being below 0.10). Comparing Tables 7 and 8, we see that increasing δ *reduces* turnover, increases portfolio volatility, and has a mixed effect

[27]See also Cremers and Weinbaum (2008) for another characteristic of options that is useful for predicting stock returns. They find that deviations from put-call parity contain information about future stock returns. They measure these deviations from the difference in implied volatility between pairs of call and put options and find that stocks with relatively expensive calls outperform stocks with relatively expensive puts.

[28]Observe that, in principle, one could have adjusted expected returns instead of volatilities using the characteristic "historical volatility risk premium." We analyzed this case, too, but we do not report the results because they are similar; in fact, because portfolio weights are very sensitive to changes in expected returns, the weights are more stable when we adjust volatilities, and this is reflected in slightly better performance when we adjust volatilities.

on Sharpe ratios and certainty-equivalent returns.[29] We conclude from the comparison of the results in Tables 7 and 8 with the results for the benchmark portfolios in Table 1 that the volatility risk premium is useful for improving portfolio performance.

7.4 Portfolios Using Option-Implied Skewness

The final approach we consider for identifying portfolios with high out-of-sample performance is motivated by the empirical results of Rehman and Vilkov (2009), who find that stocks with high option-implied skewness outperform stocks with low option-implied skewness.

To exploit this feature of the data, we proceed in the same manner as in Section 7.3, but now adjust volatilities based on the characteristic "option-implied skewness."[30] We first sort all the stocks in the data set by the characteristic "model-free implied skewness" into deciles, and then scale the volatility, $\hat{\sigma}$, of the top decile to $\hat{\sigma}(1-\delta)$ and the bottom decile to $\hat{\sigma}(1+\delta)$. As before, we consider $\delta = \{0.1, 0.2, \ldots, 0.9\}$ but report the results only for $\delta_1 = 0.40$ and $\delta_2 = 0.60$; the tables for the other values of δ are available from the authors.

Comparing the performance of portfolios in Tables 9 and Tables 10, where $\delta = 0.40$ and $\delta = 0.60$, respectively, to the benchmark portfolios in Table 1, we see that using option-implied skewness leads to a significant increase in the out-of-sample Sharpe ratio: almost all the p-values in Tables 9 and 10 for the comparison of Sharpe ratios with the $1/N$ portfolio and the benchmark portfolios in Table 1, which do not use option-implied information, are significant.

Comparing Tables 9 and 10 to Tables 7 and 8, we see that the Sharpe ratios obtained from using implied skewness are typically higher than those obtained from using the volatility risk premium (especially for the data with 561 stocks), although there is a considerable increase in turnover for portfolios based on implied skewness. Observe that while the adjustment based on implied skewness increases Sharpe ratio *and* turnover, the adjustment based on the volatility risk premium achieves improvement in Sharpe ratios while often *reducing* portfolio turnover. The reason for this is that the historical volatility risk premium estimator has much lower estimation variance than the model-free implied skewness. This is because the model-free implied skewness is estimated purely from current option prices,[31] and therefore is based on the market's expectations about the future, while the historical volatility risk premium is computed

[29] Note that the effect of increasing δ on the portfolio performance metrics is not always monotonic.

[30] Just as in the case of the "historical volatility risk premium," we could have adjusted expected returns based on implied skewness, instead of adjusting volatilities. We analyzed this case, too, but we do not report the results because the results are similar.

[31] In order to reduce variability of implied skewness, we use its average value over the last five days but it is still quite variable.

from historical return data as well as historical model-free implied volatilities, and therefore is more stable.

Overall, the empirical evidence demonstrates that using information in the volatility risk premium, and especially model-free implied skewness, can lead to a substantial improvement in the out-of-sample portfolio Sharpe ratio and certainty-equivalent return, with the main challenge being how to control turnover.

8 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 7.

8.1 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.

8.2 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.

8.3 Different Rebalancing Frequencies

We consider two rebalancing frequencies in our analysis: *daily* rebalancing, the results for which are given in Panel A of each table, and *monthly* rebalancing, the results for which are given in Panel B of each table. We find that the results are in the same direction for the two holding periods, although Sharpe ratios and certainty-equivalent returns are higher for daily rebalancing, while turnover is lower for monthly rebalancing (note that in Panel A the turnover numbers are per day, while in Panel B they are per month).

8.4 Different Benchmark Portfolios

We consider nine benchmark portfolios, of which five are listed in Table 1. 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. 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.[32] The fifth benchmark portfolio is the minimum-variance portfolio with all correlations set equal to zero.

There are four other benchmark portfolios that we consider. 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 Uppal (2009). For completeness, the performance of the mean-variance portfolio for daily returns is summarized in Section 5.3. The

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

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. The third is the parametric portfolio policy of Brandt, Santa-Clara, and Valkanov (2009) that is studied in Section 8.5, and the last is based on the maximization of expected utility, the results for which are reported in Section 8.6.

8.5 Different Portfolio Policies: Parametric Portfolios

We complement our analysis of minimum-variance portfolios by studying the effect of using option-implied information in the parametric portfolio policies developed by Brandt, Santa-Clara, and Valkanov (2009). We do this by using two option-implied stock characteristics, namely the historical variance risk premium (HVRP) and the model-free implied skewness (MFIS), in addition to the traditional stock characteristics (size, value and momentum) used to construct the parametric portfolio weights in Brandt, Santa-Clara, and Valkanov.

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 window (250 days for the smaller sample of 100 stocks and 750 days for the sample of 561 stocks). We then sum these for each date to get a time series of investable factors.[33]

To construct parametric portfolios, we start with the same factors as the ones in Brandt, Santa-Clara, and Valkanov (2009): (i) Market + 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 consider the following combinations of factors: (ii) Market + HVRP; (iii) Market + MFIS; and (iv) Market + HVRP + MFIS. Finally, in order to study the *incremental* value of option-implied information over and above the factors considered by Brandt, Santa-Clara, and Valkanov, we also consider the following combinations of factors: (v) Market + FF + MOM + HVRP; (vi) Market + FF + MOM + MFIS; and (vii) Market + FF + MOM + HVRP + MFIS.

In the parametric portfolios, the weight of an asset is a function of its market weight $\omega_{i,market}$ and, in addition, the sum of the weight of the stock within each factor multiplied by the factor

[33]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.

loadings θ:

$$\omega_i = \omega_{i,market} + (\theta_1 \times x_{i,1} + \theta_2 \times x_{i,2} + \ldots), \tag{21}$$

where θ_1 is the factor loading of the first characteristic, and $x_{i,1}$ is the first characteristic of stock i. Note that the factor loading θ_1 is *not* asset-specific, but is the same for all assets in the portfolio. For the "market portfolio" we use the equally-weighted portfolio (we also considered the value-weighted portfolio, and the findings are similar).

We choose the vector $\theta = (\theta_1, \theta_2, \ldots)$ optimally by maximizing the mean daily log-utility during the estimation period: for the 100 (561) stock sample we use 250 (750) days of data to estimate the vector θ. We consider optimization both with shortsales allowed and shortsales constrained. For the constrained optimization, 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: $\omega_i = \omega_{i,market} + \theta_1\, x_{i,1} + \theta_2, x_{i,2} + ... \geqslant 0$.

The out-of-sample performance of the parametric portfolios is reported in Table 11. We first compare the performance of portfolios based on "FF + MOM," that is, the two Fama and French (1992) characteristics and the Jegadeesh and Titman (1993) momentum characteristic, with portfolios based on the two option-implied characteristics: volatility risk premium and implied skewness. From the table, we see that the parametric portfolios based on "MFIS" attain a Sharpe ratio and certainty-equivalent return that is significantly higher than that of the $1/N$ portfolio and also the parametric portfolio based on "FF + MOM." The parametric portfolios based on "HVRP" also attain a Sharpe ratio and certainty-equivalent return that is always higher than that of the $1/N$ portfolio and typically higher than that for the parametric portfolio based on "FF + MOM," but this difference is usually not significant. However, portfolios based on "HVRP" (in addition to the market) achieve a lower portfolio volatility and turnover than the parametric portfolios based on "FF+MOM," while the volatility of the returns on the portfolio based on "MFIS" exceeds that of the "FF+MOM" portfolio, though the difference is statistically insignificant.

We can also ask the question whether the option-implied characteristics improve performance if one is already using the size, value, and momentum characteristics in parametric policies. This question is answered in the last three rows of Panels A and B of Table 11. Comparing the row "FF + MOM + MFIS" to the row "FF + MOM," we see that implied skewness improves the Sharpe ratio and certainty-equivalent return beyond the gains obtained from using only "FF + MOM," though this is accompanied by an increase in portfolio turnover. On the other hand, using HVRP sometimes leads to an improvement in Sharpe ratio and certainty-equivalent return (for the dataset with 100 assets), but other times the gains are not significant (for the dataset with 561 assets). Using both HVRP and MFIS in addition to FF + MOM leads to significant improvement in portfolio volatility, Sharpe ratio, and certainty-equivalent return, but with slightly higher turnover compared to the case where only MFIS is used, and about 2–4 times higher turnover compared to when only "FF + MOM" are used.

Overall, the analysis of option characteristics in parametric portfolios confirms the findings based on minimum-variance portfolios in Section 7.

8.6 Different Objective Functions

In our analysis in Section 7, our objective was minimum-variance optimization. However, one could have used a utility function instead. When we use the log utility function (that is, the power utility function with risk aversion γ equal to 1), we find for the data with 100 assets that the optimal portfolio has extreme weights and it performs extremely poorly, just as the mean-variance portfolio. In order to compare the weights with those from the optimization of the minimum-variance objective function, we "demean" the returns and set expected returns on all assets to be equal. We then repeat the maximization of the log utility function and find that the weights are virtually identical to the minimum-variance weights. When maximizing log utility, we also consider scaling the volatilities of stock returns by the same scaling factor, δ, in a manner similar to that used for the analysis of minimum-variance portfolios in Tables 7–10. We find a distinct improvement in portfolio performance, and the scaling works in the same direction as reported in these tables.

8.7 Risk Exposures of the Option-Implied Portfolios

To investigate the exposures of the option-implied portfolios to different risk factors, we estimate the betas of the portfolio returns with respect to the Fama and French (1993) and Carhart (1997) momentum factors. Compared to the corresponding betas of the benchmark portfolios, we find only a small increase in the size beta for the returns on portfolios where standard deviations are scaled by historical volatility premium and for the returns on portfolios with volatilities scaled by implied skewness. The value beta decreases slightly and the momentum beta increases slightly for the option-implied portfolios. Overall, the difference in the exposure to systematic factors between the benchmark portfolios and the portfolios where volatilities have been "scaled" based on option-implied information lies mainly in the exposure to the market factor, which is very small on average.[34]

[34] The tables with details of these risk exposures are available from the authors.

9 Conclusion

In this paper, we have studied how information implied in prices of stock options can be used to improve estimates of stock-return volatilities and correlations 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, with the benchmarks being the $1/N$ portfolio and 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.

We find that using *option-implied correlations* and *option-implied volatilities*, even after correcting for the volatility-risk premium, does *not* improve portfolio performance significantly. The reason why the improvement in performance is small is that the estimates of implied volatilities and implied correlations are highly variable and give poorly behaved and unstable covariance matrices that lead to highly variable weights that fail to outperform the benchmarks.

We then investigate the effect of adjusting the estimates of historical volatilities of stock returns using two sources of option-implied information. One, we use the *volatility risk premium* of each stock motivated by the empirical finding that stocks with high volatility risk premium tend to outperform those with low volatility risk premium. Our empirical evidence shows that the portfolios where volatilities have been scaled using the volatility risk premium outperform the traditional portfolios in terms of Sharpe ratio and certainty-equivalent return, but with an increase in turnover. Two, we use the *model-free option-implied skewness* to scale volatilities in the same manner as for the volatility risk premium. We find that portfolios based on implied skewness outperform even more strongly the traditional portfolios in terms of Sharpe ratio and certainty-equivalent return, but this increase is accompanied by an increase in turnover and portfolio volatility.

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 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):

$$\widehat{\langle Y_i, Y_j \rangle}^{(\Delta)} = \sum_{k=1}^{n} r_i(k\Delta) \cdot r_j(k\Delta), \tag{A1}$$

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 real market conditions where there is noise and asynchronous trading (Barndorff-Nielsen and Shephard (2002), and Zhang, Mykland, and Aït-Sahalia (2005)).[35] 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 (A1) 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 (A1) over different subsamples and then averaging the estimators obtained for these subsamples. The "second-best" estimator is given by:

$$\widehat{\langle Y_i, Y_j \rangle}^{(avg,K)} = \frac{1}{K} \sum_{k=1}^{K} \widehat{\langle Y_i, Y_j \rangle}^{(\Delta,k)}. \tag{A2}$$

We introduce one more averaging step to eliminate the chance of choosing the wrong sampling frequency by calculating the estimator (A2) 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

[35] 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.

the estimator (A2) to get the final realized (co)variance estimator:

$$\widehat{\langle Y_i, Y_j \rangle}^{(avg,\bar{K})} = \frac{1}{dim(\hat{K})} \sum_{s=1}^{dim(\bar{K})} \widehat{\langle Y_i, Y_j \rangle}^{(avg,\bar{K}(s))}. \tag{A3}$$

B 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):

$$\widehat{\Sigma}_{Shrunk} = (1-\phi)\widehat{\Sigma} + \phi S, \tag{B1}$$

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.[36] 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:

$$\widehat{\Sigma} = \sum_{n=1}^{N} \lambda_n V_n V_n^{'}, \tag{B2}$$

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:

$$\widehat{\Sigma}^{-1/2} = \sum_{n=1}^{N} \frac{1}{\sqrt{\lambda_n}} V_n V_n^{'}, \tag{B3}$$

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

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

$$\widetilde{\Sigma}^{-1/2} = \sum_{n=1}^{k} \frac{1}{\sqrt{\lambda_n}} V_n V_n^{'} + \frac{1}{\sqrt{\lambda_{k+1}}} \sum_{n=k+1}^{N} V_n V_n^{'}, \tag{B4}$$

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.

C 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:

$$R(t,\tau) \equiv \ln S(t+\tau) - \ln S(t). \tag{C1}$$

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

$$V(t,\tau) = \int_{S(t)}^{\infty} \frac{2\left(1 - \log\left(\frac{K}{S(t)}\right)\right)}{K^2} \cdot C(t,\tau;K)dK \tag{C2}$$
$$+ \int_{0}^{S(t)} \frac{2\left(1 - \log\left(\frac{K}{S(t)}\right)\right)}{K^2} \cdot P(t,\tau;K)dK,$$

the price of the cubic contract is

$$W(t,\tau) = \int_{S(t)}^{\infty} \frac{6\log\left(\frac{K}{S(t)}\right) - 3\left(\log\left(\frac{K}{S(t)}\right)\right)^2}{K^2} \cdot C(t,\tau;K)dK$$
$$- \int_{0}^{S(t)} \frac{6\log\left(\frac{K}{S(t)}\right) + 3\left(\log\left(\frac{K}{S(t)}\right)\right)^2}{K^2} \cdot P(t,\tau;K)dK, \tag{C3}$$

and the price of the quartic contract is

$$X(t,\tau) = \int_{S(t)}^{\infty} \frac{12\left(ln[\frac{K}{S(t)}]\right)^2 - 4\left(ln[\frac{K}{S(t)}]\right)^3}{K^2} \cdot C(t,\tau;K)dK$$
$$+ \int_{0}^{S(t)} \frac{12\left(ln[\frac{S(t)}{K}]\right)^2 + 4\left(ln[\frac{S(t)}{K}]\right)^3}{K^2} \cdot P(t,\tau;K)dK. \tag{C4}$$

Define

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

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

$$\text{MFIV}(t,\tau) = \left(V(t,\tau)\right)^{1/2}, \tag{C5}$$

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

$$\text{MFIS}(t,\tau) = \frac{e^{r\tau}W(t,\tau) - 3\mu(t,\tau)e^{r\tau}V(t,\tau) + 2(\mu(t,\tau))^3}{(e^{r\tau}V(t,\tau) - (\mu(t,\tau))^2)^{\frac{3}{2}}}. \tag{C6}$$

To calculate the integrals in (C2), (C3), and (C4) 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.[37] 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 (C5) and (C6), respectively.

[37]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.

Table 1: 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 trading				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.1331	0.9066	0.1118	0.3076	0.1333	0.2998	0.0311	0.5365
	(0.00)	(0.53)	(0.71)		(0.00)	(0.89)	(0.94)	
	(0.50)	(0.50)	(0.50)		(0.50)	(0.50)	(0.50)	
Constrained	0.1199	0.8251	0.0917	0.0588	0.1161	0.7435	0.0796	0.0373
	(0.00)	(0.70)	(0.93)		(0.00)	(0.51)	(0.85)	
	(0.00)	(0.63)	(0.75)		(0.00)	(0.06)	(0.09)	
Shrinkage	0.1197	1.0323	0.1164	0.1730	0.1164	0.4440	0.0449	0.3074
	(0.00)	(0.35)	(0.68)		(0.00)	(0.81)	(0.92)	
	(0.00)	(0.10)	(0.36)		(0.00)	(0.06)	(0.15)	
Zero Correlation	0.1433	0.8820	0.1161	0.0137	0.1493	0.7808	0.1055	0.0125
	(0.00)	(0.78)	(0.97)		(0.00)	(0.32)	(0.74)	
	(1.00)	(0.54)	(0.47)		(1.00)	(0.08)	(0.06)	
				Panel B: Monthly trading				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	0.1313	0.9323	0.1138	0.4051	0.1329	0.2452	0.0237	0.7897
	(0.00)	(0.53)	(0.71)		(0.00)	(0.92)	(0.95)	
	(0.50)	(0.50)	(0.50)		(0.50)	(0.50)	(0.50)	
Constrained	0.1168	0.8409	0.0914	0.0805	0.1239	0.6852	0.0771	0.0666
	(0.00)	(0.71)	(0.93)		(0.00)	(0.50)	(0.84)	
	(0.00)	(0.65)	(0.77)		(0.00)	(0.04)	(0.06)	
Shrinkage	0.1172	1.0625	0.1177	0.2496	0.1236	0.3706	0.0381	0.5078
	(0.00)	(0.37)	(0.71)		(0.00)	(0.83)	(0.93)	
	(0.00)	(0.10)	(0.36)		(0.00)	(0.07)	(0.14)	
Zero Correlation	0.1342	0.9076	0.1127	0.0524	0.1444	0.7794	0.1020	0.0562
	(0.00)	(0.76)	(0.96)		(0.00)	(0.34)	(0.73)	
	(0.97)	(0.53)	(0.45)		(1.00)	(0.05)	(0.03)	

Table 2: Prediction of volatility and 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 3: Portfolios using option-implied volatility

In this table, we evaluate the performance of various 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 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 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 "Sample-cov" matrix using the Ledoit and Wolf (2004a,b) methodology; 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 minimum-variance benchmark portfolio in Table 1, 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 trading								
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.1351	0.4122	0.0466	0.7980	0.1312	0.4297	0.0478	1.5553
	(0.00)	(0.95)	(0.97)		(0.00)	(0.85)	(0.92)	
	(0.84)	(1.00)	(1.00)		(0.32)	(0.32)	(0.32)	
Constrained	0.1175	0.6872	0.0738	0.2285	0.1237	0.6142	0.0683	0.2778
	(0.00)	(0.90)	(0.98)		(0.00)	(0.70)	(0.88)	
	(0.05)	(0.86)	(0.88)		(0.88)	(0.71)	(0.66)	
Shrinkage	0.1197	0.5388	0.0573	0.5236	0.1161	0.3487	0.0338	1.1131
	(0.00)	(0.93)	(0.98)		(0.00)	(0.93)	(0.97)	
	(0.49)	(1.00)	(1.00)		(0.46)	(0.66)	(0.66)	
Zero Correlation	0.1433	0.8085	0.1056	0.0558	0.1468	0.7230	0.0954	0.0610
	(0.00)	(0.98)	(1.00)		(0.00)	(0.62)	(0.90)	
	(0.56)	(1.00)	(1.00)		(0.00)	(0.99)	(1.00)	
Panel B: Monthly trading								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	0.1320	0.5923	0.0695	0.8527	0.1283	0.4244	0.0463	1.6414
	(0.00)	(0.96)	(0.98)		(0.00)	(0.86)	(0.92)	
	(0.81)	(1.00)	(1.00)		(0.35)	(0.29)	(0.31)	
Constrained	0.1142	0.7240	0.0761	0.2385	0.1178	0.7090	0.0766	0.2862
	(0.00)	(0.90)	(0.99)		(0.00)	(0.70)	(0.90)	
	(0.07)	(0.86)	(0.88)		(0.78)	(0.72)	(0.68)	
Shrinkage	0.1188	0.7130	0.0777	0.5665	0.1190	0.3829	0.0385	1.1863
	(0.00)	(0.92)	(0.98)		(0.00)	(0.93)	(0.97)	
	(0.50)	(1.00)	(1.00)		(0.45)	(0.67)	(0.67)	
Zero Correlation	0.1351	0.8481	0.1054	0.0792	0.1405	0.7617	0.0971	0.0863
	(0.00)	(0.98)	(1.00)		(0.00)	(0.62)	(0.89)	
	(0.55)	(1.00)	(1.00)		(0.00)	(0.99)	(1.00)	

Table 4: Portfolios using risk-premium-corrected 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. Correlations 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 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 using the Ledoit and Wolf (2004a,b) methodology; 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 1, 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 trading				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.1323	0.8085	0.0982	0.6665	0.1250	0.6619	0.0749	1.1819
	(0.00)	(0.66)	(0.80)		(0.00)	(0.60)	(0.79)	
	(0.34)	(0.72)	(0.73)		(0.07)	(0.11)	(0.12)	
Constrained	0.1190	0.8016	0.0883	0.2204	0.1175	0.6726	0.0722	0.2447
	(0.00)	(0.72)	(0.93)		(0.00)	(0.60)	(0.83)	
	(0.29)	(0.56)	(0.58)		(0.56)	(0.60)	(0.59)	
Shrinkage	0.1176	0.8630	0.0946	0.4272	0.1110	0.5843	0.0587	0.8511
	(0.00)	(0.61)	(0.86)		(0.00)	(0.74)	(0.92)	
	(0.06)	(0.91)	(0.93)		(0.02)	(0.27)	(0.30)	
Zero Correlation	0.1384	0.8528	0.1085	0.0575	0.1399	0.7496	0.0951	0.0649
	(0.00)	(0.85)	(0.99)		(0.00)	(0.49)	(0.87)	
	(0.00)	(0.86)	(0.97)		(0.00)	(0.81)	(0.96)	
				Panel B: Monthly trading				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	0.1294	0.9363	0.1128	0.7155	0.1214	0.6795	0.0752	1.2513
	(0.00)	(0.66)	(0.82)		(0.00)	(0.61)	(0.82)	
	(0.35)	(0.73)	(0.74)		(0.13)	(0.08)	(0.10)	
Constrained	0.1150	0.8527	0.0914	0.2296	0.1097	0.8483	0.0871	0.2508
	(0.00)	(0.73)	(0.94)		(0.00)	(0.60)	(0.85)	
	(0.30)	(0.58)	(0.60)		(0.52)	(0.62)	(0.63)	
Shrinkage	0.1162	0.9868	0.1080	0.4681	0.1152	0.6010	0.0627	0.9129
	(0.00)	(0.60)	(0.87)		(0.00)	(0.74)	(0.93)	
	(0.07)	(0.92)	(0.94)		(0.06)	(0.23)	(0.27)	
Zero Correlation	0.1304	0.8929	0.1079	0.0801	0.1350	0.7924	0.0978	0.0885
	(0.00)	(0.84)	(0.98)		(0.00)	(0.48)	(0.86)	
	(0.00)	(0.88)	(0.99)		(0.00)	(0.84)	(0.98)	

Table 5: Portfolios using homogeneous 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 1, 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 trading				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.1594	0.1115	0.0051	0.0793	0.1915	0.2589	0.0312	0.0502
	(0.37)	(0.98)	(0.98)		(1.00)	(0.86)	(0.84)	
	(1.00)	(0.99)	(0.99)		(1.00)	(0.55)	(0.50)	
Constrained	0.1319	0.6850	0.0816	0.0714	0.1439	0.6582	0.0844	0.0564
	(0.00)	(0.81)	(0.91)		(0.00)	(0.61)	(0.73)	
	(1.00)	(0.78)	(0.67)		(1.00)	(0.65)	(0.43)	
Regularization	0.1579	0.1193	0.0064	0.1046	0.2027	0.2308	0.0262	0.0808
	(0.22)	(0.99)	(0.99)		(1.00)	(0.89)	(0.87)	
	(1.00)	(1.00)	(1.00)		(1.00)	(0.73)	(0.62)	
				Panel B: Monthly trading				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	0.1654	0.0956	0.0020	0.1790	0.2085	0.2304	0.0260	0.1984
	(0.41)	(0.97)	(0.97)		(0.93)	(0.85)	(0.83)	
	(1.00)	(0.99)	(0.98)		(1.00)	(0.55)	(0.49)	
Constrained	0.1249	0.7132	0.0813	0.0830	0.1449	0.5968	0.0760	0.0679
	(0.00)	(0.80)	(0.90)		(0.00)	(0.61)	(0.72)	
	(1.00)	(0.76)	(0.65)		(1.00)	(0.66)	(0.43)	
Regularization	0.1618	0.1089	0.0043	0.2078	0.2195	0.2277	0.0253	0.2460
	(0.30)	(0.98)	(0.98)		(1.00)	(0.88)	(0.85)	
	(1.00)	(1.00)	(0.99)		(1.00)	(0.74)	(0.62)	

Table 6: Portfolios using heterogeneous 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 1, 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 trading				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.5200	0.5197	0.1949	1.7923	0.7564	−0.5270	−0.3878	12.0527
	(0.69)	(0.82)	(0.33)		(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	0.1463	0.1209	0.9117	0.1030	0.1796
	(0.00)	(0.72)	(0.90)		(0.00)	(0.26)	(0.61)	
	(0.94)	(0.66)	(0.61)		(0.99)	(0.13)	(0.10)	
Regularization	0.1219	0.8402	0.0950	0.1922	0.1176	0.5876	0.0622	0.2534
	(0.00)	(0.62)	(0.85)		(0.00)	(0.70)	(0.87)	
	(0.91)	(0.88)	(0.86)		(0.68)	(0.28)	(0.27)	
				Panel B: Monthly trading				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	1.4607	0.1453	0.1144	1.8168	0.5288	−0.2353	−0.1770	14.5847
	(0.72)	(0.82)	(0.33)		(1.00)	(1.00)	(1.00)	
	(1.00)	(0.89)	(0.25)		(1.00)	(0.97)	(1.00)	
Constrained	0.1191	0.7216	0.0789	0.1560	0.1224	0.7263	0.0814	0.1916
	(0.00)	(0.73)	(0.91)		(0.00)	(0.27)	(0.63)	
	(0.92)	(0.65)	(0.61)		(0.99)	(0.16)	(0.11)	
Regularization	0.1159	0.7732	0.0829	0.2431	0.1179	0.4055	0.0409	0.3530
	(0.00)	(0.63)	(0.84)		(0.00)	(0.68)	(0.84)	
	(0.85)	(0.89)	(0.87)		(0.64)	(0.28)	(0.27)	

Table 7: Portfolios using historical volatility scaled by volatility risk premium: $\delta = 0.40$

In this table, we evaluate the performance of portfolios computed with historical volatilities that have been scaled by the volatility risk premium using the following procedure. We sort all stocks by the characteristic "volatility risk premium" into deciles, and then change the volatility, $\hat{\sigma}$, of the top decile to $\hat{\sigma}(1-\delta)$ and of the bottom decile to $\hat{\sigma}(1+\delta)$, with $\delta = 0.40$. Correlations 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 where historical volatility is scaled by the volatility risk premium using the procedure described above; "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 using the Ledoit and Wolf (2004a,b) methodology; 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 scaled by the volatility risk premium using the procedure described above. 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 1, 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 trading				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.1390	1.3144	0.1731	0.2860	0.1323	0.6496	0.0772	0.5177
	(0.00)	(0.07)	(0.18)		(0.00)	(0.65)	(0.83)	
	(1.00)	(0.02)	(0.01)		(0.36)	(0.09)	(0.09)	
Constrained	0.1321	1.0239	0.1265	0.0652	0.1216	0.9418	0.1072	0.0510
	(0.00)	(0.34)	(0.62)		(0.00)	(0.20)	(0.58)	
	(1.00)	(0.10)	(0.04)		(1.00)	(0.08)	(0.05)	
Shrinkage	0.1295	1.3781	0.1700	0.1725	0.1254	0.7026	0.0803	0.3298
	(0.00)	(0.04)	(0.19)		(0.00)	(0.58)	(0.82)	
	(1.00)	(0.02)	(0.01)		(1.00)	(0.13)	(0.11)	
Zero Correlation	0.1375	0.9648	0.1232	0.0187	0.1427	0.8009	0.1041	0.0161
	(0.00)	(0.34)	(0.81)		(0.00)	(0.28)	(0.73)	
	(0.00)	(0.02)	(0.13)		(0.00)	(0.24)	(0.62)	
				Panel B: Monthly trading				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	0.1384	1.2062	0.1573	0.3521	0.1382	0.6169	0.0758	0.6316
	(0.00)	(0.09)	(0.19)		(0.00)	(0.65)	(0.83)	
	(0.99)	(0.02)	(0.01)		(0.38)	(0.07)	(0.07)	
Constrained	0.1276	0.9932	0.1186	0.0862	0.1285	0.8787	0.1047	0.0756
	(0.00)	(0.35)	(0.60)		(0.00)	(0.22)	(0.58)	
	(1.00)	(0.10)	(0.04)		(1.00)	(0.06)	(0.04)	
Shrinkage	0.1263	1.2907	0.1550	0.2306	0.1367	0.6483	0.0793	0.4308
	(0.00)	(0.04)	(0.18)		(0.00)	(0.57)	(0.82)	
	(1.00)	(0.03)	(0.01)		(1.00)	(0.11)	(0.09)	
Zero Correlation	0.1280	0.9834	0.1177	0.0555	0.1395	0.7927	0.1008	0.0580
	(0.00)	(0.35)	(0.79)		(0.00)	(0.31)	(0.72)	
	(0.00)	(0.03)	(0.14)		(0.00)	(0.27)	(0.62)	

Table 8: Portfolios using historical volatility scaled by volatility risk premium: $\delta = 0.60$

In this table, we evaluate the performance of portfolios computed with historical volatilities that have been scaled by the volatility risk premium using the following procedure. We sort all stocks by the characteristic "volatility risk premium" into deciles, and then change the volatility, $\hat{\sigma}$, of the top decile to $\hat{\sigma}(1-\delta)$ and of the bottom decile to $\hat{\sigma}(1+\delta)$, with $\delta = 0.60$. Correlations 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 where historical volatility is scaled by the volatility risk premium using the procedure described above; "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 using the Ledoit and Wolf (2004a,b) methodology; 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 scaled by the volatility risk premium using the procedure described above. 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 1, 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 trading				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.1438	1.3044	0.1773	0.2096	0.1348	0.7390	0.0905	0.3276
	(0.00)	(0.06)	(0.14)		(0.00)	(0.52)	(0.77)	
	(1.00)	(0.06)	(0.03)		(0.68)	(0.07)	(0.07)	
Constrained	0.1368	1.0442	0.1336	0.0598	0.1230	0.9223	0.1059	0.0508
	(0.00)	(0.31)	(0.53)		(0.00)	(0.23)	(0.58)	
	(1.00)	(0.11)	(0.04)		(1.00)	(0.11)	(0.07)	
Shrinkage	0.1385	1.3503	0.1775	0.1336	0.1324	0.7644	0.0925	0.2179
	(0.00)	(0.04)	(0.12)		(0.00)	(0.47)	(0.75)	
	(1.00)	(0.07)	(0.02)		(1.00)	(0.12)	(0.09)	
Zero Correlation	0.1348	1.0439	0.1316	0.0237	0.1379	0.8186	0.1034	0.0198
	(0.00)	(0.17)	(0.59)		(0.00)	(0.26)	(0.71)	
	(0.00)	(0.02)	(0.10)		(0.00)	(0.26)	(0.59)	
				Panel B: Monthly trading				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	0.1408	1.2137	0.1611	0.2604	0.1379	0.7293	0.0911	0.4010
	(0.00)	(0.08)	(0.16)		(0.00)	(0.52)	(0.76)	
	(1.00)	(0.05)	(0.03)		(0.62)	(0.05)	(0.05)	
Constrained	0.1320	1.0199	0.1260	0.0812	0.1297	0.8715	0.1046	0.0751
	(0.00)	(0.34)	(0.54)		(0.00)	(0.25)	(0.60)	
	(1.00)	(0.12)	(0.05)		(1.00)	(0.10)	(0.05)	
Shrinkage	0.1349	1.2557	0.1604	0.1828	0.1398	0.7328	0.0927	0.2879
	(0.00)	(0.05)	(0.14)		(0.00)	(0.47)	(0.74)	
	(1.00)	(0.09)	(0.02)		(1.00)	(0.09)	(0.05)	
Zero Correlation	0.1253	1.0510	0.1239	0.0583	0.1364	0.8027	0.1001	0.0598
	(0.00)	(0.19)	(0.59)		(0.00)	(0.29)	(0.71)	
	(0.00)	(0.04)	(0.12)		(0.00)	(0.28)	(0.58)	

Table 9: Portfolios using historical volatility scaled by implied skewness: $\delta = 0.40$

In this table, we evaluate the performance of portfolios computed with historical volatilities that have been scaled by the option-implied skewness using the following procedure. We sort all stocks by the characteristic "model free implied skewness" into deciles, and then change the volatility, $\hat{\sigma}$, of the top decile to $\hat{\sigma}(1-\delta)$ and of the bottom decile to $\hat{\sigma}(1+\delta)$, with $\delta = 0.40$. Correlations 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 where historical volatility is scaled by option-implied skewness using the procedure described above; "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 using the Ledoit and Wolf (2004a,b) methodology; 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 scaled by the option-implied skewness using the procedure described above. 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 1, 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 trading				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.1439	1.5154	0.2078	0.9520	0.1417	1.3132	0.1761	1.9354
	(0.00)	(0.02)	(0.05)		(0.00)	(0.02)	(0.07)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
Constrained	0.1322	1.2760	0.1600	0.2704	0.1321	1.1654	0.1453	0.2888
	(0.00)	(0.05)	(0.23)		(0.00)	(0.03)	(0.20)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
Shrinkage	0.1331	1.5949	0.2035	0.5996	0.1323	1.4435	0.1823	1.2994
	(0.00)	(0.00)	(0.03)		(0.00)	(0.00)	(0.04)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
Zero Correlation	0.1445	0.9762	0.1307	0.0484	0.1507	0.8550	0.1175	0.0416
	(0.00)	(0.23)	(0.70)		(0.00)	(0.07)	(0.44)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
				Panel B: Monthly trading				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	0.1441	1.1856	0.1606	0.9915	0.1483	0.8868	0.1206	2.0016
	(0.00)	(0.01)	(0.03)		(0.00)	(0.01)	(0.06)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
Constrained	0.1262	0.9551	0.1126	0.2836	0.1420	0.8189	0.1062	0.3039
	(0.00)	(0.04)	(0.21)		(0.00)	(0.04)	(0.20)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
Shrinkage	0.1323	1.2823	0.1610	0.6365	0.1476	0.9491	0.1293	1.3606
	(0.00)	(0.00)	(0.02)		(0.00)	(0.00)	(0.03)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
Zero Correlation	0.1359	0.9507	0.1199	0.0837	0.1476	0.8117	0.1089	0.0834
	(0.00)	(0.24)	(0.69)		(0.00)	(0.10)	(0.45)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	

Table 10: Portfolios using historical volatility scaled by implied skewness: $\delta = 0.60$

In this table, we evaluate the performance of portfolios computed with historical volatilities that have been scaled by the option-implied skewness using the following procedure. We sort all stocks by the characteristic "model free implied skewness" into deciles, and then change the volatility, $\hat{\sigma}$ of the top decile to $\hat{\sigma}(1-\delta)$ and of the bottom decile to $\hat{\sigma}(1+\delta)$, with $\delta = 0.60$. Correlations 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 where historical volatility is scaled by option-implied skewness using the procedure described above; "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 using the Ledoit and Wolf (2004a,b) methodology; 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 scaled by the option-implied skewness using the procedure described above. 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 1, 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 trading				
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
Sample cov	0.1656	1.6324	0.2567	1.0502	0.1554	1.3939	0.2045	1.5894
	(0.92)	(0.00)	(0.00)		(0.00)	(0.00)	(0.01)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
Constrained	0.1541	1.4439	0.2107	0.3635	0.1440	1.1674	0.1578	0.3302
	(0.02)	(0.01)	(0.02)		(0.00)	(0.04)	(0.13)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.01)	(0.00)	
Shrinkage	0.1583	1.7080	0.2579	0.7298	0.1536	1.4606	0.2126	1.1502
	(0.23)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
Zero Correlation	0.1467	1.0927	0.1495	0.0910	0.1521	0.9352	0.1307	0.0750
	(0.00)	(0.02)	(0.15)		(0.00)	(0.01)	(0.14)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
				Panel B: Monthly trading				
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
Sample cov	0.1637	1.1771	0.1795	1.0766	0.1670	0.9769	0.1494	1.6275
	(0.89)	(0.00)	(0.00)		(0.00)	(0.00)	(0.01)	
	(1.00)	(0.01)	(0.00)		(1.00)	(0.00)	(0.00)	
Constrained	0.1461	0.9776	0.1322	0.3755	0.1549	0.8210	0.1152	0.3447
	(0.04)	(0.01)	(0.01)		(0.00)	(0.04)	(0.12)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.01)	(0.00)	
Shrinkage	0.1558	1.2523	0.1832	0.7566	0.1731	0.9965	0.1577	1.1886
	(0.25)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
	(1.00)	(0.01)	(0.00)		(1.00)	(0.00)	(0.00)	
Zero Correlation	0.1388	0.9961	0.1286	0.1222	0.1520	0.8412	0.1163	0.1150
	(0.00)	(0.02)	(0.15)		(0.00)	(0.02)	(0.16)	
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	

Table 11: Parametric portfolios using option-implied information

In this table, we evaluate the performance of parametric portfolios proposed in Brandt, Santa-Clara, and Valkanov (2009) using option-implied information. We report two p-values in parenthesis, the first with respect to the $1/N$ portfolio, and the second with respect to the "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 trading					
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7470	0.1151	0.0144
FF + MOM	0.1644	1.0439	0.1581	0.0757	0.1799	0.8487	0.1365	0.0534
	(1.00)	(0.01)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.50)	(0.50)	(0.50)		(0.50)	(0.50)	(0.50)	
HVRP	0.1536	1.0669	0.1521	0.0254	0.1675	0.7979	0.1197	0.0261
	(0.00)	(0.00)	(0.02)		(0.00)	(0.16)	(0.30)	
	(0.00)	(0.37)	(0.70)		(0.00)	(0.80)	(0.94)	
MFIS	0.1647	1.1903	0.1825	0.1713	0.1809	0.9988	0.1643	0.1749
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.62)	(0.00)	(0.00)		(0.96)	(0.00)	(0.00)	
HVRP + MFIS	0.1600	1.2608	0.1889	0.2034	0.1779	1.0079	0.1635	0.1945
	(0.09)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.00)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
FF + MOM + HVRP	0.1605	1.0770	0.1600	0.0899	0.1781	0.8108	0.1286	0.0622
	(0.33)	(0.00)	(0.00)		(1.00)	(0.03)	(0.01)	
	(0.00)	(0.20)	(0.39)		(0.00)	(0.94)	(0.96)	
FF + MOM + MFIS	0.1660	1.1914	0.1840	0.2069	0.1796	0.9979	0.1631	0.1780
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(1.00)	(0.00)	(0.00)		(0.30)	(0.00)	(0.00)	
FF + MOM + HVRP + MFIS	0.1617	1.2411	0.1877	0.2327	0.1761	1.0048	0.1614	0.1930
	(0.89)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.00)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	
			Panel B: Monthly trading					
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7312	0.1102	0.0672
FF + MOM	0.1634	0.9552	0.1427	0.1049	0.1801	0.7942	0.1267	0.0925
	(1.00)	(0.01)	(0.01)		(1.00)	(0.00)	(0.00)	
	(0.50)	(0.50)	(0.50)		(0.50)	(0.50)	(0.50)	
HVRP	0.1455	1.0694	0.1450	0.0690	0.1623	0.8016	0.1169	0.0745
	(0.00)	(0.00)	(0.02)		(0.01)	(0.15)	(0.29)	
	(0.00)	(0.39)	(0.69)		(0.00)	(0.78)	(0.91)	
MFIS	0.1572	1.0392	0.1510	0.1872	0.1809	0.7804	0.1248	0.1957
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.58)	(0.00)	(0.00)		(0.83)	(0.00)	(0.00)	
HVRP + MFIS	0.1522	1.0981	0.1555	0.2174	0.1782	0.8009	0.1268	0.2137
	(0.18)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.01)	(0.00)	(0.00)		(0.01)	(0.00)	(0.00)	
FF + MOM + HVRP	0.1598	1.0142	0.1492	0.1166	0.1759	0.7796	0.1216	0.0987
	(0.42)	(0.01)	(0.01)		(0.99)	(0.03)	(0.02)	
	(0.01)	(0.24)	(0.41)		(0.16)	(0.93)	(0.95)	
FF + MOM + MFIS	0.1627	0.9958	0.1487	0.2208	0.1803	0.7863	0.1255	0.1973
	(1.00)	(0.00)	(0.00)		(1.00)	(0.00)	(0.00)	
	(0.97)	(0.00)	(0.00)		(0.37)	(0.00)	(0.00)	
FF + MOM + HVRP + MFIS	0.1575	1.0628	0.1550	0.2448	0.1770	0.7957	0.1252	0.2108
	(0.85)	(0.00)	(0.00)		(0.99)	(0.00)	(0.00)	
	(0.03)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)	

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.

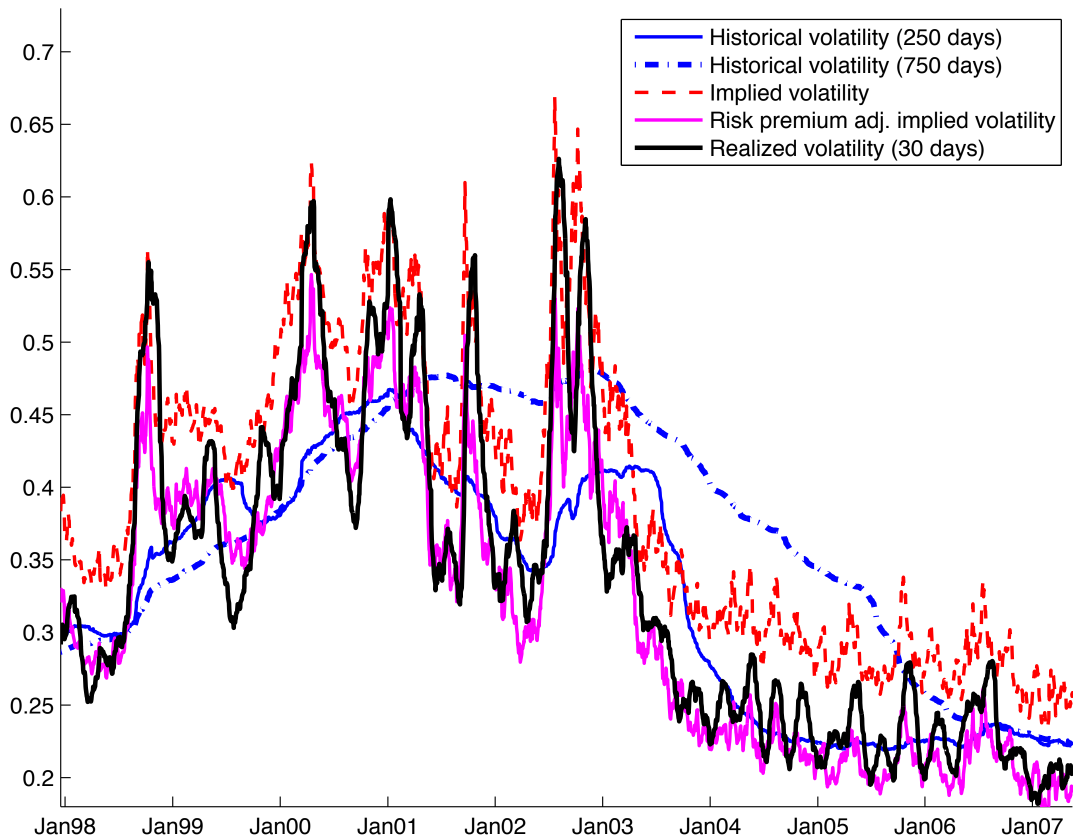

Figure 2: Correlations: Realized, historical, and implied

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.

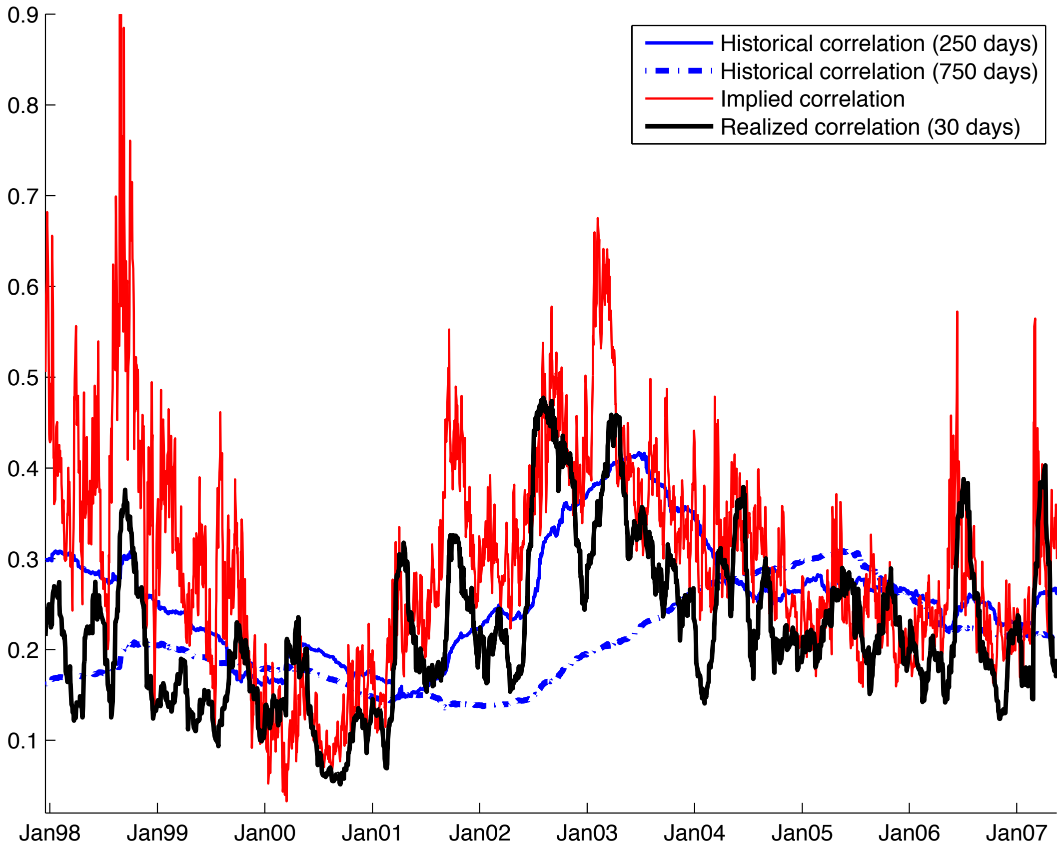

References

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(2), 499–547.

Andersen, T. G., T. Bollerslev, P. Christoffersen, and F. X. Diebold, 2006, "Volatility and Correlation Forecasting," in *Handbook of Economic Forecasting*, ed. by G. Elliott, C. W. J. Granger, and A. G. Timmermann. North Holland, Amsterdam.

Andersen, T. G., T. Bollerslev, and F. X. Diebold, 2009, "Parametric and Nonparametric Measurement of Volatility," in *Handbook of Financial Econometrics*, ed. by Y. Aït-Sahalia, and L. P. Hansen. 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(11), 1797–1812.

Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard, 2009, "Realised Kernels in Practice: Trades and Quotes," *Econometrics Journal*, 4, 1–32.

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(2), 253–280.

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

Black, F., and M. S. Scholes, 1972, "The Valuation of Option Contracts and a Test of Market Efficiency," *Journal of Finance*, 27(2), 399–417.

Blair, B. J., S.-H. Poon, and S. J. Taylor, 2001, "Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High-Frequency Index Returns," *Journal of Econometrics*, 105, 5–26.

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(1), 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., P. Santa-Clara, and R. Valkanov, 2009, "Parametric Portfolio Policies: Exploiting Characteristics in the Cross Section of Equity Returns," *Review of Financial Studies*, 22(9), 3411–3447.

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

Brown, S. J., 1990, "Estimating Volatility," in *Financial Options: From Theory to Practice*, ed. by S. Figlewski, W. L. Silber, and M. G. Subrahmanyam. Business One-Irwin, Homewood, IL.

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

Carhart, M. M., 1997, "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52(1), 57–82.

Carr, P., and L. Wu, 2009, "Variance Risk Premiums," *Review of Financial Studies*, 22(3), 1311–1341.

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

——— , 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(5), 937–974.

Chernov, M., 2007, "On the Role of Risk Premia in Volatility Forecasting," *Journal of Business & Economic Statistics*, 25(4), 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(3), 407–458.

Christensen, B. J., and N. R. Prabhala, 1998, "The Relation Between Implied and Realized Volatility," *Journal of Financial Economics*, 50(2), 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(1), 33–83.

Cremers, M., and D. Weinbaum, 2008, "Deviations from Put-Call Parity and Stock Return Predictability," forthcoming in *Journal of Financial and Quantitative Analysis*.

DeMiguel, V., L. Garlappi, F. J. Nogales, and R. Uppal, 2009, "A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms," *Management Science*, 55(5), 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(5), 1915–1953.

Driessen, J., P. Maenhout, and G. Vilkov, 2009, "The Price of Correlation Risk: Evidence from Equity Options," *Journal of Finance*, 64(3), 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(3), 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(1), 72–78.

——— , 2002, "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business and Economic Statistics*, 20(3), 339–350.

Epps, T. W., 1979, "Comovements in Stock Prices in the Very Short Run," *Journal of the American Statistical Association*, 74(366), 291–298.

Fama, E. F., and K. R. French, 1992, "The Cross-Section of Expected Stock Returns," *Journal of Finance*, 47(2), 427–465.

——— , 1993, "Common Risk Factors in the Returns on Stock and Bonds," *Journal of Financial Economics*, 33(1), 3–56.

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(3), 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(2), 433–451.

Jackwerth, J. C., and M. Rubinstein, 1996, "Recovering Probability Distributions from Option Prices," *Journal of Finance*, 51(5), 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(1), 267–307.

Jagannathan, R., and T. Ma, 2003, "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps," *Journal of Finance*, 58(4), 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(3), 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(4), 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(2), 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(2), 369–381.

Ledoit, O., and M. Wolf, 2004a, "Honey, I Shrunk the Sample Covariance Matrix," *Journal of Portfolio Management*, 30(4), 110–119.

——— , 2004b, "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices," *Journal of Multivariate Analysis*, 88, 365–411.

——— , 2008, "Robust Performance Hypothesis Testing with the Sharpe Ratio," *Journal of Empirical Finance*, 15, 850–859.

MacKinlay, A. C., and Ľ. Pástor, 2000, "Asset Pricing Models: Implications for Expected Returns and Portfolio Selection," *Review of Financial Studies*, 13(4), 883–916.

Merton, R. C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.

——— , 1973, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4(1), 141–183.

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

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(1), 179–223.

Pástor, Ľ., and R. F. Stambaugh, 2000, "Comparing Asset Pricing Models: An Investment Perspective," *Journal of Financial Economics*, 56(3), 335–381.

Politis, D., and J. Romano, 1994, "The Stationary Bootstrap," *Journal of the American Statistical Association*, 89(428), 1303–1313.

Poon, S.-H., and C. W. J. Granger, 2005, "Practical Issues in Forecasting Volatility," *Financial Analysts Journal*, 61(1), 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(1), 61–69.

Vanden, J. M., 2008, "Information Quality and Options," *Review of Financial Studies*, 21(6), 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(472), 1394–1411.

Zhou, B., 1996, "High-Frequency Data and Volatility in Foreign-Exchange Rates," *Journal of Business and Economic Statistics*, 14(1), 45–52.

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