

Improving Portfolio Selection Using Option-Implied Volatility and Skewness*

Victor DeMiguel[‡]

Yuliya Plyakha[§]

Raman Uppal[‡]

Grigory Vilkov[§]

This version: March 3, 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. Our empirical evidence shows that, while using the option-implied volatilities and correlations does *not* improve significantly the portfolio variance, Sharpe ratio, and certainty-equivalent return, exploiting information contained in the volatility risk premium and option-implied skewness increases substantially Sharpe ratios and certainty-equivalent returns, but this is accompanied by higher portfolio turnover.

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

JEL: G11, G12, G13, G17

*We gratefully acknowledge comments from Alexander Alekseev, Mike Chernov, Engelbert Dockner, Bernard Dumas, Lorenzo Garlappi, Nicolae Gârleanu, Vasant Naik, Stavros Panageas, Paulo Rodrigues, Christian Schlag, George Skiadopoulos, Josef Zechner, and seminar participants at Goethe University Frankfurt, Man Investments, London School of Economics, University of Mainz, University of Piraeus, and Vienna University of Economics and Business Administration.

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

Improving Portfolio Selection Using Option-Implied Volatility and Skewness

This version: March 3, 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. Our empirical evidence shows that, while using the option-implied volatilities and correlations does *not* improve significantly the portfolio variance, Sharpe ratio, and certainty-equivalent return, exploiting information contained in the volatility risk premium and option-implied skewness increases substantially Sharpe ratios and certainty-equivalent returns, but this is accompanied by higher portfolio turnover.

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

JEL: G11, G12, G13, G17

Contents

1	Introduction	1
2	Portfolio Selection	5
3	Data	5
3.1	Data on Stock Returns	5
3.2	Data on Stock Options	7
4	Description of Portfolio-Performance Metrics	7
5	Benchmark Portfolios that Do Not Use Option Prices	9
5.1	The equally-weighted portfolio	10
5.2	Three minimum-variance portfolios based on historical returns	10
5.3	The mean-variance portfolio based on historical returns	12
6	Option-Implied Information	12
6.1	Option-Implied Volatility and Volatility Risk Premium	12
6.2	Option-Implied Correlation	15
6.3	Option-Implied Skewness	16
7	Portfolios that Use Option-Implied Information	16
7.1	Portfolios Using Option-Implied Volatilities	16
7.2	Portfolios Using Option-Implied Correlations	18
7.3	Portfolios Using Volatility Risk Premium	19
7.4	Portfolios Using Option-Implied Skewness	20
8	Robustness Tests	21
8.1	Different Datasets	21
8.2	Different Data Frequencies	21
8.3	Different Rebalancing Frequencies	22
8.4	Different Benchmark Portfolios	22
8.5	Different Objective Functions	23
8.6	Risk exposures of the option-implied portfolios	23
9	Conclusion	23
A	Computing (Co)variances from Intraday Data	25
B	Shrinkage and Regularization of Covariance Matrix	26
C	The Construction of the Risk-Neutral Implied Moments	27
	Tables	29
	Figures	36
	References	38

1 Introduction

To determine the optimal portfolio of an investor, one needs to estimate the moments of asset returns, such as the 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 short sales are constrained), but this is accompanied by an increase in portfolio turnover. Our analysis is carried out in a comprehensive fashion: we consider four benchmark portfolios (1/ N , sample-covariance-based minimum-variance, short-sale-constrained minimum-variance, and minimum-variance with shrinkage of the covariance matrix); four performance metrics (portfolio volatility, Sharpe ratio, certainty-equivalent return, and turnover); two data sets (with 100 assets and with 561 assets); two data frequencies (daily and intraday); two portfolio rebalancing periods (daily and monthly); and, two objective functions (minimum-variance optimization and utility maximization).

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, Jagan-

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

²These 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).

³For 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.

nathan and Ma (2003, pp. 1652–3) 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, even though the methodology we develop applies also to mean-variance portfolios and to portfolios obtained from the maximization of more general utility functions.⁴

To determine minimum-variance portfolio, one needs to estimate for each stock its volatility and correlations with all the other stocks. We undertake our analysis in four 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 the implied volatilities to compute the optimal portfolio does *not* lead to a substantial reduction in the out-of-sample portfolio volatility (standard deviation) 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 for this. 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 model-free 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

⁴A discussion of the poor performance of the mean-variance portfolio for our data sets is given in Section 5.3, and the performance of portfolio weights obtained from maximizing log utility is discussed in Section 8.5.

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 appears 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 premiums 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.⁵ 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 short-sale constraints), but these gains are accompanied with higher portfolio turnover. Moreover, we observe that the use of option-implied information for portfolio selection does not lead to a substantial change in the systematic risk profile of the portfolios with respect to the three Fama and French (1993) factors and the Carhart (1997) momentum factor.

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).⁶ The focus of our work is to investigate how the information implied in 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

⁵For 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).

⁶For 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 obtain also implied correlations from the prices of options on individual stocks and on the index, while Bali and Hovakimian (2009) and 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.

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 order to help the reader understand each step of our analysis. In Section 2, we provide a brief background to portfolio selection. In Section 3, we describe the data on stock returns and options that we use in our empirical work. In Section 4, we explain the various measures we use to evaluate portfolio performance. The construction and performance of our benchmark portfolios that do *not* use option-implied information is 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 explains the method used for the shrinkage and regularization of the covariance matrix; and, Appendix C explains the construction of model-free option-implied moments.

2 Portfolio Selection

The minimum-variance problem is

$$\min_w w^\top \Sigma w, \quad (1)$$

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

where $w \in \mathbb{R}^N$ is the vector of portfolio weights invested in stocks, $\Sigma \in \mathbb{R}^{N \times N}$ is the covariance matrix, and $e \in \mathbb{R}^N$ is the vector of ones. The objective in (1) is to minimize the variance of the portfolio return, $w^\top \Sigma w$, subject to the constraints that the portfolio weights sum to one. The constraint that the portfolio weights sum to one is required because we consider the case without the risk-free asset, given our objective is to explore how to use option-implied information to select the portfolio of risky stocks. The solution to the above problem is:

$$w_{min} = \frac{\Sigma^{-1}e}{e^\top \Sigma^{-1}e}. \quad (3)$$

Note that the covariance matrix Σ in (3) can be decomposed into volatility and correlation matrices,

$$\Sigma = \text{diag}(\sigma) \Omega \text{diag}(\sigma), \quad (4)$$

where $\text{diag}(\sigma)$ denotes the diagonal matrix with volatilities of the stocks on the diagonal, and Ω is the correlation matrix. Thus, to obtain the optimal portfolio weights in (3) there are two quantities that need to be estimated: volatilities (σ) and correlations (Ω). Information from prices of options can be used 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. 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 3146 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.⁷

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, delete entries with a bid, ask or transaction price equal to zero, delete entries with corrected trades (trades with a correction indicator, “corr” \neq 0), delete entries with an abnormal sale condition (trades where the variable “cond” has a letter code, except for “E” and “F”)⁸ and delete entries with 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.⁹ 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,¹⁰ and the second consists of 100 stocks out of the 219 for which data are available for all dates; to select these 100 stocks, we first order the 219 stocks with respect to the security identifier code of the IvyDB data base, and then select the first 100.¹¹

⁷For a survey of the literature on using high-frequency data to estimate moments of asset returns, see Andersen, Bollerslev, and Diebold (2009).

⁸See the TAQ 3 Users Guide for additional details about sale conditions.

⁹Rules P1, P2, T1, T2 and an adjusted version of T4.

¹⁰At 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.

¹¹In 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 robustness check are provided in Section 8.1.

3.2 Data on Stock Options

For stock options we use the IvyDB that contains data on all U.S.-listed index and equity options. We use data from January 4, 1996 to June 29, 2007.¹²

We do not use option *prices* directly in our analysis, but 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.¹³

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.¹⁴ For each date, each underlying stock, and each 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.¹⁵

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.

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 considering the certainty-equivalent return, in addition to the Sharpe

¹²Note 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.

¹³We calculated implied moments also from the raw data on option prices, and the results are similar.

¹⁴The 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.

¹⁵There are 13 implied volatilities given for standard delta points for each call and put. For puts, these 13 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.

ratio, is that the Sharpe ratio considers only the mean and volatility of returns, while the certainty-equivalent 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, which for daily data corresponds to three years.¹⁶ Two, using the return data over the estimation window τ , we compute the various portfolios we wish to compare. Three, 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 $t = \tau, \dots, T - 1$ and for each strategy.

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 ($\hat{\sigma}$), Sharpe ratio of returns (SR), and certainty-equivalent return (ce) are:

$$\hat{\mu}^{strategy} = \frac{1}{T - \tau} \sum_{t=\tau}^{T-1} r_{t+1}^{strategy}, \quad (5)$$

$$\hat{\sigma}^{strategy} = \left(\frac{1}{T - \tau - 1} \sum_{t=\tau}^{T-1} \left(r_{t+1}^{strategy} - \hat{\mu}^{strategy} \right)^2 \right)^{1/2}, \quad (6)$$

$$\widehat{\text{SR}}^{strategy} = \frac{\hat{\mu}^{strategy}}{\hat{\sigma}^{strategy}}, \quad (7)$$

$$\widehat{\text{ce}}^{strategy} = u^{-1} \left(\frac{1}{T - \tau} \sum_{t=\tau}^{T-1} u \left(r_{t+1}^{strategy} \right) \right), \quad (8)$$

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 benchmark, we report also the p-values for these differences. We are interested in the finite-

¹⁶Because our sample consists of 561 stocks, estimation window lengths shorter than $\tau = 750$ often give singularities in the covariance matrix.

sample properties of the portfolios. Therefore, when 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* 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:

$$\text{Turnover} = \frac{1}{T - \tau - 1} \sum_{t=\tau}^{T-1} \sum_{j=1}^N \left(\left| w_{j,t+1}^{strategy} - w_{j,t^+}^{strategy} \right| \right). \quad (9)$$

5 Benchmark Portfolios that Do Not Use Option Prices

For robustness, we consider four 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 short-sale-constrained minimum-variance portfolio; and (iv) the unconstrained minimum-variance portfolio with shrinkage of the covariance matrix. The construction of these four 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 our four benchmark portfolios listed above, and so we do not report the results for the mean-variance portfolio in the tables.

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

¹⁷Specifically, consider two portfolios i and n , with $\mu_i, \mu_n, \sigma_i, \sigma_n$ as their true means and volatilities. We wish to test the hypothesis that the Sharpe ratio (or certainty-equivalent return) of portfolio i is worse (smaller) than that of the benchmark portfolio n , that is, $H_0 : \mu_i/\sigma_i - \mu_n/\sigma_n \leq 0$. To do this, we obtain B pairs of size $T - \tau$ of the portfolio returns i and n by simple resampling with replacement for daily returns, and by blockwise resampling with replacement for overlapping 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 $\hat{\mu}_i/\hat{\sigma}_i - \hat{\mu}_n/\hat{\sigma}_n$, then a one-sided P-value for the previous null hypothesis is given by $\hat{p} = \hat{F}(0)$, and we will reject it for a small \hat{p} . In a similar way, to test the hypothesis that the variance of the portfolio i is greater (worse) than the variance of the benchmark portfolio n , $H_0 : \sigma_i^2/\sigma_n^2 \geq 1$, if \hat{F} denotes the empirical distribution function of the B bootstrap pairs corresponding to: $\hat{\sigma}_i^2/\hat{\sigma}_n^2$, then, a one-sided P-value for this null hypothesis is given by $\hat{p} = 1 - \hat{F}(1)$, and we will reject the null for a small \hat{p} . For a nice discussion of the application of other bootstrapping methods to tests of differences in portfolio performance, see Ledoit and Wolf (2008).

trading days); therefore, this corresponds to the average of 21 daily returns. The advantage of this approach is that it is not sensitive to the particular day on which the portfolio is formed.

5.1 The equally-weighted portfolio

For the “*equally-weighted*” ($1/N$) portfolio, one invests an equal amount of wealth across all N available stocks each period. 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.

5.2 Three minimum-variance portfolios based on historical returns

The second benchmark portfolio that we consider is the minimum-variance portfolio, w_{min} , which is given in (3). To identify the portfolio weights for this strategy, we need to estimate only the “*sample covariance*” matrix, that is, the realized volatilities and correlations. For *daily data* we compute the conventional sample estimators of (co)variances using data over the past 750 days.¹⁸

In the existing literature, two methods have been adopted to improve the out-of-sample performance of the minimum-variance portfolio based on the sample (co)variances. One approach is to impose constraints on the portfolio weights, which Jagannathan and Ma (2003) show can lead to substantial gains in performance. Thus, our third benchmark is the “*constrained*” portfolio, where we compute the short-sale-constrained minimum-variance portfolio weights. To compute these portfolio weights, we solve the problem in (1)–(2), after imposing the additional constraint that all weights have to be non-negative.

Another approach to improve the out-of-sample performance of the minimum-variance portfolio based on the sample covariance matrix is to use “shrinkage.” Our fourth benchmark consists of “*shrinkage*” portfolios, where we compute the minimum-variance portfolio weights *after* shrinking the covariance matrix.¹⁹ The sample covariance matrix for daily data and intraday data is computed using the same approach that is described above. To shrink the covari-

¹⁸For the realized (co)variance estimators based on *intraday data* we use the filtered and calendar-time aligned transaction prices over the last 30 trading days to estimate the (co)variances. 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 Ait-Sahalia (2005), and the details of our procedure are provided in Appendix A.

¹⁹We 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.

ance matrix for daily returns, we then 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.²⁰

In **Table 1** we report the performance of these four benchmark portfolios, 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 ($\hat{\sigma}$) 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 0.1347, for the minimum-variance portfolio with constraints is 0.1240, and for the minimum-variance portfolio with shrinkage is 0.1180. The first set of p-values indicate that the volatilities of the three 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) and the certainty-equivalent return (ce) are typically higher for the $1/N$ portfolio. The one-sided p-values again indicate that we cannot reject the null that the Sharpe ratio of the minimum-variance portfolios is worse than that of $1/N$. These results are similar for “Monthly rebalancing” reported in Panel B for the data with 100 stocks; but for the data with 561 stocks and monthly rebalancing, the Sharpe ratio and certainty-equivalent return are highest for the constrained-minimum-variance portfolio rather than the $1/N$ portfolio.

Of the three minimum-variance portfolios, the short-sale-constrained portfolio has the lowest turnover (which is true also in the tables that follow, where we use option-implied information). However, the turnover of the $1/N$ portfolio is almost always lower than even that of the constrained minimum-variance portfolio.

²⁰For intraday data, instead of shrinkage, we use the regularization approach of Zumbach (2009) because the distribution of intraday returns is different from that of daily returns and does not satisfy the assumptions of Ledoit and Wolf (2004a,b). 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.

5.3 The mean-variance portfolio based on historical returns

For completeness, we also discuss briefly the results for the *mean-variance* portfolio portfolios. Of the three variants of the mean-variance portfolio we consider, the first is based on the sample covariance matrix, the second has short-sale 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. For example, while the volatility of the three minimum-variance portfolios is less than 0.1350, the volatility of the corresponding three mean-variance portfolios is at least 3 times higher. Similarly, the Sharpe ratio of the short-sale-constrained mean-variance portfolios is less than half of that of the short-sale-constrained minimum-variance portfolios, and it is negative for the other two mean-variance portfolios. Finally, the turnover of the short-sale-constrained mean-variance portfolio is five times that of the minimum-variance portfolio, and that for the other two mean-variance portfolios is about fifteen times higher. 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.

6 Option-Implied Information

In this section, we explain how we compute the quantities implied by option prices that we use in our selection of portfolios. These quantities are: (i) option-implied volatility and the volatility risk premium; (ii) option-implied correlation; and, (iii) option-implied skewness. We also compare the ability of option-implied volatility and option-implied correlation to predict realized volatility and realized correlation, respectively. Our objective here 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.²¹

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

²¹There is a very large literature on forecasting stock-return volatility and correlations; for a discussion of this, see for instance, Engle (1982), Bollerslev (1986), and 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).

the *model-free implied volatility* (MFIV), which represents a nonparametric estimate of the risk-neutral expected stock-return volatility until the option’s expiration.

Model-free implied volatility is given by a single number and it subsumes information in the whole Black-Scholes implied volatility smile. Theoretical and empirical research (see Jiang and Tian (2005) and Vanden (2008)) finds that model-free implied volatility is better at predicting the future realized volatility than the Black-Scholes implied volatility, and it is used by the CBOE to compute VIX, which is the ticker symbol for the CBOE Volatility Index that gives the implied volatility of S&P500 index options. To compute the model-free implied volatility, we first calculate the option prices from the interpolated volatility surface data. We then use these prices to find the value of the “variance contract,” following the approach in Bakshi, Kapadia, and Madan (2003); the formula for the variance contract and the procedure used to compute it is provided in Appendix C.²² The square root of the variance contract then gives us the model-free implied volatility.

To verify whether the intuition that the model-free implied volatility is a useful predictor of realized volatility in the future relative to using the volatility estimate based on historical returns, we regress realized variances on the model-free implied variances and compare the R^2 with that when variances based on historical data are used as a predictor. We see from Panel A of **Table 2** that when regressing the 30-day realized variances in the future on (i) 750-day historical daily variances, (ii) 30-day intraday historical variances, and (iii) model-free implied variances, the R^2 for the model-free implied variances is higher than that for intraday historical variances, which is higher than for daily historical variances. This is true for both the dataset with 100 stocks and that with 561 stocks.²³ For example, in the case of the data with 561 stocks, the R^2 for historical daily variances is 19.66%, for intraday historical variances is 22.99%, and for model-free implied variances is 33.12%.

However, what we need for portfolio selection is not the risk-neutral implied volatilities of stock returns but the expected volatilities under the *objective* distribution. We now explain how to make a correction to the model-free implied volatilities in order to get the volatilities 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 (instead of the expected volatility) to estimate

²²For 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).

²³The results are similar also when we regress the 30-day intraday (high-frequency) volatilities on the same regressors.

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) as the square root of the average variance risk premium for each stock on each day for the past period of T trading days:²⁴

$$\text{HVRP}_t^2 = \frac{1}{T - \Delta t} \sum_{i=t-T+1}^{t-\Delta t} \frac{\text{MFIV}_{i,i+\Delta t}^2}{\text{RV}_{i,i+\Delta t}^2}. \quad (10)$$

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

$$\widehat{\text{RV}}_{t,t+\Delta t} = \frac{\text{MFIV}_{t,t+\Delta t}}{\text{HVRP}_t}. \quad (11)$$

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 31.55%, which is about the same as the R^2 for the model-free implied volatility (33.12%), 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 by comparing the time series for the different volatility measures in **Figure 1**, where we plot the 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. Note also that the variability of each of these volatility series is quite different.

²⁴Note that because HVRP_t^2 is calculated as the average of the ratio of $\text{MFIV}_{i,i+\Delta t}^2$ and $\text{RV}_{i,i+\Delta t}^2$, both of which are calculated over Δt days, as a result we will have only $T - \Delta t$ observations when computing the average.

6.2 Option-Implied Correlation

The second piece of option-implied information that we consider is implied correlation. It has been shown by Driessen, Maenhout, and Vilkov (2009) that the risk-neutral expectation of the correlation is higher than the expectation under the objective probability measure. Therefore, one approach for computing implied correlation is to take the historical correlation matrix as the proxy for the future realized correlation matrix and increase each pairwise historical correlation by the correlation risk premium. We use the same methodology as in Buss and Vilkov (2008) to obtain implied correlation. That is, we assume that the size of the adjustment to each pairwise historical correlation is proportional to the distance between the calculated historical correlation and the perfect correlation of one, which ensures that the implied correlation matrix satisfies the necessary positive-definiteness restriction.

Having computed implied correlations using the above approach, we examine whether they are superior at predicting realized correlations. To do this, we regress realized correlations on (i) 750-day historical daily correlations, (ii) 750-day historical daily correlations after shrinkage, (iii) 30-day intraday correlations, (iv) 30-day intraday correlations after regularization, and (v) implied correlations. We report in Panel B of **Table 2** that for the data with 561 stocks, the R^2 for historical daily correlations is 3.97%, for historical daily correlations with shrinkage is 3.28%, for intraday correlations is 6.26%, for intraday correlations with shrinkage is 4.79%, and for implied correlation is 9.32%, with the results being similar for the data with 100 stocks. Thus, while high-frequency intraday correlations do not seem to predict realized correlations better than historical correlations based on daily data, implied correlations are better, though the improvement in R^2 is smaller than it was for predicting realized volatilities.

In **Figure 2**, we plot the historical correlation based on the last 750 days (dashed blue line), implied correlation (solid red line), and 30-day realized correlations (thick black line). Just as in the figure for volatilities, the plot is based on the cross-sectional equally-weighted average of average correlations across 561 stocks. There are two observations about these series: first, implied correlation follows the level of realized correlation much more closely than historical correlation; two, implied correlation is much more volatile (that is, contains more noise) than realized correlation, while historical correlation is almost a smooth function (but contains much less of the 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 and at-the-money put options (Xing, Zhang, and Zhao, 2009). However, that measure does not take into account the whole distribution, but rather just the left tale. 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 the recent research (Rehman and Vilkov (2009)) that risk-neutral skewness contains information about the future stock returns above and beyond 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 model-free implied skewness as described by 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 option-implied forward-looking information 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 select portfolios with superior out-of-sample performance.

In each of these sections, we use only one source of option-implied information, 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 (4); that is, we

use as the covariance matrix: $\hat{\Sigma} = \text{diag}(\text{MFIV}) \hat{\Omega} \text{diag}(\text{MFIV})$. Using this adjusted covariance matrix, we then determine the minimum-variance portfolio in (3), along with the portfolios where short sales are constrained, and where the shrinkage is applied to this adjusted covariance matrix. In computing these portfolios, we continue to use historical correlations. Effectively, we are computing the minimum-variance portfolio but with implied volatilities instead of historical volatilities in the covariance matrix.

The results are reported in **Table 3**. This table and also all the remaining tables 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 optimal 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 always helps to reduce the out-of-sample volatility of the portfolio, although the reduction itself is relatively small (an average of 1.6% for daily rebalancing and 4.7% for monthly rebalancing) and the p-values for this reduction are significant only for the data with 100 stocks for monthly rebalancing.

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 damaging, 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 ratios and certainty-equivalent returns and turnovers when using the risk-premium-corrected implied volatilities, which confirms our motivation for correcting the implied-volatilities for the volatility risk premium. However, the Sharpe ratios and certainty-equivalent returns in Table 4, are still lower than those for the sample-based benchmark portfolios in Table 1. And, 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 exception is the shortsale-constrained portfolio for 561 stocks with monthly rebalancing, 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 instability 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 use of 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.

Table 5 shows that of the three minimum-variance portfolios we consider with the implied correlations, the constrained portfolio does best; but the portfolios relying on implied correlations typically perform worse than the corresponding minimum-variance benchmark portfolios in Table 1 in terms of all four performance metrics: volatility, Sharpe ratio, certainty-equivalent return, and turnover. One possible explanation for this poor performance is that by replacing historical correlations by the implied correlations, we are essentially increasing the magnitude of the off-diagonal elements of the covariance matrix, making the covariance matrix less stable,

which then leads to poor performance. Moreover, as can be seen in Figure 2, the implied correlations are also much more variable than the other series; this variability is reflected also in Panel B of Table 2 where the p-values of the β estimated in the predictive regressions is very low, even though the R^2 is the highest when implied correlation is the predictor.²⁵

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.²⁶ 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 volatilities of the top decile to $\hat{\sigma} \cdot (1 - \delta)$ and of the bottom decile to $\hat{\sigma} \cdot (1 + \delta)$. This scaling allows us to increase the weights of assets that are expected to outperform other stocks and decrease the weights of assets that are expected to underperform other stocks in the portfolio. In our empirical implementation, to give the reader a sense of how the results change with δ , we consider two values of the scaling factor: $\delta_1 = 0.50$ and $\delta_2 = 0.70$. To compute the minimum-variance portfolio, we use the scaled volatilities, along with correlations estimated from historical data.

Table 6 gives the results when the estimates of the historical volatilities have been scaled, based on the volatility risk premium for each stock. The table shows 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 (the only exception are the constrained portfolios for 561 stocks). Moreover, in all cases but that of the sample covariance portfolio for 561 stocks, the Sharpe ratios exceed that of the $1/N$ portfolio (with several p-values being below 0.10). We also see that increasing δ *reduces* turnover, increases portfolio volatility, and has a mixed effect on Sharpe ratios and certainty-equivalent returns.

²⁵In order to conserve space, we do not present the table where the portfolios are computed using implied correlations *and* risk-premium-corrected implied volatilities. The table with the results for this case is available from the authors. As one would expect, the results are a combination of the insights from Tables 4 and 5. What one can conclude from these results is that using implied correlations and risk-premium-corrected implied volatilities has (i) a mixed effect on portfolio volatility, reducing it in some cases but increasing it in others; (ii) reduces Sharpe ratios and certainty-equivalent returns considerably; and (iii) usually increases turnover. Compared to the $1/N$ portfolio, the portfolio using option-implied volatilities *and* correlations is better in terms of portfolio volatility, but worse in terms of Sharpe ratio, certainty-equivalent return, and turnover.

²⁶See 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.

We conclude from the comparison of the results in Table 6 and 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 the section above, but in this case adjust volatilities based on the characteristic “option-implied skewness.” We first sort all the stocks in the data set by the characteristic “model-free implied skewness” into deciles, and then scale the volatilities of the top decile to $\hat{\sigma} \cdot (1 - \delta)$ and for the bottom decile to $\hat{\sigma} \cdot (1 + \delta)$. We consider two values of the scaling factor: $\delta_1 = 0.30$ and $\delta_2 = 0.50$.²⁷

Comparing Table 7 to the benchmark portfolios in Table 1, we see that using the option-implied skewness leads to a significant increase in the out-of-sample Sharpe ratio, which ranges between 30% and 400% for daily rebalancing, and between 11% and 87% for monthly rebalancing. However, this is accompanied by a considerable increase in turnover. Comparing Table 7 to Table 6, we see that the Sharpe ratios obtained from using implied skewness are higher than those from using the volatility risk premium. The main difference between the results obtained with the scaling based on the model-free implied skewness and that based on the historical volatility risk premium is the effect on portfolio turnover. While the adjustment based on implied skewness increases Sharpe ratio but also leads to a much higher turnover, the adjustment based on the volatility risk premium achieves improvement in Sharpe ratios while often *reducing* the portfolio turnover, compared to the adjustment based on option-implied skewness, and the use of implied volatility or implied correlation. 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,²⁸ and therefore is based on the market’s expectations about the future, while the historical volatility risk premium is computed from historical return data as well as historical model-free implied volatilities, and therefore, is more stable.

²⁷Note that we use smaller values of δ than for the scaling based on the volatility risk premium in the previous section. This is because the model-free implied skewness has a stronger effect on the portfolio’s out-of-sample performance.

²⁸In order to reduce variability of implied skewness, we use its average value over the last five days but it is still quite variable.

Overall, the empirical evidence demonstrates that using information in the volatility risk premium and model-free implied skewness can lead to a substantial improvement in the out-of-sample portfolio Sharpe ratios and certainty-equivalent returns, 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 from our empirical analysis.

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 the second dataset is a subset of the first. The first 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. The second 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 the IvyDB data base, and then selecting the first 100. The results for this dataset are reported in Columns 2–4 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 set of 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 ratios 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, and 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 short-sale 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.

8.4 Different Benchmark Portfolios

We consider four benchmark portfolios, all of which 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 short-selling 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.

Note that 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 reported in Section 5.3.

8.5 Different Objective Functions

In our analysis, our objective has been 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 6 and 7. We find a distinct improvement in portfolio performance, and the scaling works in the same direction as reported in these tables.

8.6 Risk exposures of the option-implied portfolios

To investigate the exposures of the option-implied portfolios to different risk factors, we estimated 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.²⁹

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

²⁹The tables with details of these risk exposures are available from the authors.

benchmarks being the $1/N$ portfolio and three types of minimum-variance portfolios based on historical data: the first based on the sample covariance matrix, the second with short sale constraints, and the third with shrinkage applied to the covariance matrix.

We find that using *option-implied correlations* and *option-implied volatilities*, even after correcting for the volatility-risk premium, do *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 adjust our estimates of historical volatilities on stocks 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 mean-variance 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):

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

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; Jacod and Protter, 1998). However, it is inconsistent under real market conditions where there is noise and asynchronous trading (Barndorff-Nielsen and Shephard, 2002; Zhang, Mykland, and Ait-Sahalia, 2005).³⁰ To mitigate this problem, we use the “second-best” estimator from Zhang, Mykland, and Ait-Sahalia (originally derived to estimate the 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 and to overcome this problem Zhang, Mykland, and Ait-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:

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

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. As our sample also includes less frequently traded stocks, especially early in the sample period, we choose relatively low sampling frequencies from 240 to 390 minutes (which corresponds to the number of minutes in a typical trading day) with a step size of 10 minutes for the estimator

³⁰The 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.

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

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

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):

$$\hat{\Sigma}_{Shrunk} = (1 - \phi)\hat{\Sigma} + \phi S, \quad (\text{B1})$$

where we shrink the sample estimate of the covariance matrix $\hat{\Sigma}$ toward a diagonal matrix with the cross-sectional average variance on the diagonal, defined as the target S .³¹ 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 the asymptotic properties 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 $\hat{\Sigma}$, which is:

$$\hat{\Sigma} = \sum_{n=1}^N \lambda_n V_n V_n', \quad (\text{B2})$$

where $\{\lambda_1, \dots, \lambda_N\}$ are eigenvalues and $\{V_1, \dots, V_N\}$ 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:

$$\hat{\Sigma}^{-1/2} = \sum_{n=1}^N \frac{1}{\sqrt{\lambda_n}} V_n V_n', \quad (\text{B3})$$

³¹We 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

$$\tilde{\Sigma}^{-1/2} = \sum_{n=1}^k \frac{1}{\sqrt{\lambda_n}} V_n V_n' + \frac{1}{\sqrt{\lambda_{k+1}}} \sum_{n=k+1}^N V_n V_n', \quad (\text{B4})$$

and use $\tilde{\Sigma}$ as the estimator of the covariance matrix. This approach allows us to maintain the covariance structure by keeping all eigenvectors $\{V_1, \dots, V_N\}$ 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 , $R(t, \tau)$ be the τ -period return (seen at time $t + \tau$) given by the log-price relative:

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

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

$$\begin{aligned} V(t, \tau) = & \int_{S(t)}^{\infty} \frac{2 \left(1 - \log \left(\frac{K}{S(t)}\right)\right)}{K^2} \cdot C(t, \tau; K) dK \\ & + \int_0^{S(t)} \frac{2 \left(1 - \log \left(\frac{K}{S(t)}\right)\right)}{K^2} \cdot P(t, \tau; K) dK, \end{aligned} \quad (\text{C2})$$

the price of the cubic contract is

$$\begin{aligned} W(t, \tau) = & \int_{S(t)}^{\infty} \frac{6 \log \left(\frac{K}{S(t)}\right) - 3 \left(\log \left(\frac{K}{S(t)}\right)\right)^2}{K^2} \cdot C(t, \tau; K) dK \\ & - \int_0^{S(t)} \frac{6 \log \left(\frac{K}{S(t)}\right) + 3 \left(\log \left(\frac{K}{S(t)}\right)\right)^2}{K^2} \cdot P(t, \tau; K) dK, \end{aligned} \quad (\text{C3})$$

and the price of the quartic contract is

$$\begin{aligned}
X(t, \tau) = & \int_{S(t)}^{\infty} \frac{12 \left(\ln\left[\frac{K}{S(t)}\right] \right)^2 - 4 \left(\ln\left[\frac{K}{S(t)}\right] \right)^3}{K^2} \cdot C(t, \tau; K) dK \\
& + \int_0^{S(t)} \frac{12 \left(\ln\left[\frac{S(t)}{K}\right] \right)^2 + 4 \left(\ln\left[\frac{S(t)}{K}\right] \right)^3}{K^2} \cdot P(t, \tau; K) dK.
\end{aligned} \tag{C4}$$

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

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

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 3.³² 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.

³²The 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 portfolio based on the sample covariance matrix but with short sales constrained; and, “Shrinkage” is the portfolio where shrinkage has been applied to the sample covariance matrix using the Ledoit and Wolf (2004a,b) methodology. 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).

Strategy	100 stocks				561 stocks			
	$\hat{\sigma}$	sr	ce	trn	$\hat{\sigma}$	sr	ce	trn
<i>Panel A: Daily rebalancing</i>								
1/ N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7489	0.1155	0.0144
<i>Minimum-variance portfolios</i>								
Sample cov	0.1217 (0.00) (0.50)	0.8089 (0.67) (0.50)	0.0911 (0.87) (0.50)	0.0849	0.1347 (0.00) (0.50)	0.4737 (0.78) (0.50)	0.0547 (0.86) (0.50)	0.5360
Constrained	0.1242 (0.00) (0.93)	0.7557 (0.83) (0.63)	0.0862 (0.96) (0.60)	0.0262	0.1240 (0.00) (0.00)	0.9701 (0.15) (0.03)	0.1126 (0.53) (0.05)	0.0376
Shrinkage	0.1205 (0.00) (0.00)	0.8514 (0.61) (0.05)	0.0953 (0.86) (0.09)	0.0745	0.1180 (0.00) (0.00)	0.6142 (0.66) (0.06)	0.0655 (0.84) (0.20)	0.3046
<i>Panel B: Monthly rebalancing</i>								
1/ N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7339	0.1107	0.0673
<i>Minimum-variance portfolios</i>								
Sample cov	0.1221 (0.00) (0.50)	0.7903 (0.74) (0.50)	0.0891 (0.91) (0.50)	0.1762	0.1296 (0.00) (0.50)	0.4768 (0.82) (0.50)	0.0534 (0.91) (0.50)	0.7785
Constrained	0.1207 (0.00) (0.37)	0.7632 (0.84) (0.58)	0.0848 (0.97) (0.62)	0.0545	0.1180 (0.00) (0.01)	1.0135 (0.07) (0.00)	0.1126 (0.49) (0.01)	0.0690
Shrinkage	0.1211 (0.00) (0.13)	0.8280 (0.69) (0.03)	0.0930 (0.89) (0.05)	0.1602	0.1196 (0.00) (0.00)	0.5842 (0.73) (0.09)	0.0627 (0.90) (0.17)	0.4979

Table 2: Prediction of variances and correlations

In Panel A of this table, we report the results of variance prediction regressions $RV = \alpha + \beta \widehat{RV}$, where we regress the 30-days realized variance (RV) on the variance predictors (\widehat{RV}). 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}). We use 750 days of data to calculate the historical predictors. For covariances based on daily data, shrinkage is applied using the Ledoit and Wolf (2004a,b) methodology; for covariances based on intraday data, regularization is applied using the methodology in Zumbach (2009). In both cases, we infer the correlations from the covariances to which shrinkage/regularization has been applied. In parenthesis we report the p-values for the two-sided null hypotheses: $\alpha = 0$ and $\beta = 1$.

Predictor	100 stocks			561 stocks		
	α	β	R^2	α	β	R^2
<i>Panel A: Prediction of variances</i>						
Historical daily	0.0369 (0.03)	0.6795 (0.05)	0.2031	0.0520 (0.10)	0.6515 (0.09)	0.1966
Historical intraday	0.0616 (0.95)	0.5075 (0.01)	0.2308	0.0717 (0.94)	0.5667 (0.06)	0.2299
Implied variance	0.0214 (0.98)	0.7460 (0.19)	0.3334	0.0164 (0.97)	0.7793 (0.21)	0.3312
Implied variance (HVRP corrected)	0.0372 (0.97)	0.8189 (0.31)	0.3222	0.0435 (0.96)	0.8735 (0.31)	0.3155
<i>Panel B: Prediction of correlations</i>						
Historical daily	0.2007 (0.84)	0.0821 (0.02)	0.0238	0.2548 (0.80)	-0.1761 (0.07)	0.0397
Shrinkage daily	0.1987 (0.84)	0.0988 (0.03)	0.0245	0.2501 (0.80)	-0.0814 (0.04)	0.0328
Historical intraday	0.1703 (0.87)	0.2312 (0.00)	0.0680	0.1776 (0.86)	0.1880 (0.00)	0.0626
Shrinkage intraday	0.1899 (0.85)	1.4419 (0.38)	0.0419	0.2014 (0.84)	1.2203 (0.37)	0.0479
Implied correlation	0.1130 (0.91)	0.3528 (0.00)	0.0826	0.1116 (0.91)	0.3503 (0.01)	0.0932

Table 3: Portfolios using implied volatilities

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 “Sample cov” portfolio is the minimum-variance portfolio based on the sample covariance matrix but where historical volatility is replaced by the option-implied volatility; “Constrained” is the portfolio based on the same covariance matrix as above, but where short sales are constrained; and, “Shrinkage” is the portfolio based on the same covariance matrix as above, but with shrinkage applied to the covariance matrix using the Ledoit and Wolf (2004a,b) 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).

Strategy	100 stocks				561 stocks			
	$\hat{\sigma}$	sr	ce	trn	$\hat{\sigma}$	sr	ce	trn
<i>Panel A: Daily rebalancing</i>								
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7489	0.1155	0.0144
Sample cov	0.1212 (0.00) (0.39)	0.3724 (0.98) (0.99)	0.0378 (0.99) (0.99)	0.5801	0.1307 (0.00) (0.19)	0.4792 (0.81) (0.49)	0.0541 (0.89) (0.50)	1.5504
Constrained	0.1210 (0.00) (0.02)	0.6339 (0.94) (0.81)	0.0694 (0.99) (0.83)	0.2312	0.1240 (0.00) (0.46)	0.6458 (0.65) (0.92)	0.0724 (0.85) (0.92)	0.2743
Shrinkage	0.1188 (0.00) (0.14)	0.4174 (0.98) (0.99)	0.0425 (0.99) (1.00)	0.5154	0.1156 (0.00) (0.17)	0.3953 (0.90) (0.84)	0.0390 (0.96) (0.85)	1.1032
<i>Panel B: Monthly rebalancing</i>								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7339	0.1107	0.0673
Sample cov	0.1139 (0.00) (0.02)	0.5566 (0.96) (0.94)	0.0569 (0.99) (0.96)	0.6168	0.1251 (0.00) (0.23)	0.5051 (0.85) (0.46)	0.0555 (0.94) (0.49)	1.6359
Constrained	0.1143 (0.00) (0.01)	0.6745 (0.96) (0.80)	0.0705 (1.00) (0.87)	0.2402	0.1159 (0.00) (0.36)	0.7628 (0.41) (0.92)	0.0817 (0.84) (0.93)	0.2832
Shrinkage	0.1132 (0.00) (0.02)	0.5930 (0.94) (0.96)	0.0607 (0.99) (0.97)	0.5508	0.1161 (0.00) (0.24)	0.4546 (0.92) (0.79)	0.0461 (0.98) (0.80)	1.1763

Table 4: Portfolios using risk-premium-corrected implied volatilities

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 “Sample cov” portfolio is the minimum-variance portfolio based on the sample covariance matrix but where historical volatility is replaced by the risk-premium-corrected option-implied volatility; “Constrained” is the portfolio based on the same covariance matrix as above, but where short sales are constrained; and, “Shrinkage” is the portfolio based on the same covariance matrix as above, but with shrinkage applied to the covariance matrix using the Ledoit and Wolf (2004a,b) 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).

Strategy	100 stocks				561 stocks			
	$\hat{\sigma}$	sr	ce	trn	$\hat{\sigma}$	sr	ce	trn
<i>Panel A: Daily rebalancing</i>								
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7489	0.1155	0.0144
Sample cov	0.1215 (0.00) (0.46)	0.7735 (0.72) (0.58)	0.0866 (0.90) (0.58)	0.4782	0.1246 (0.00) (0.03)	0.6837 (0.58) (0.23)	0.0775 (0.78) (0.27)	1.1788
Constrained	0.1226 (0.00) (0.17)	0.7990 (0.73) (0.39)	0.0904 (0.93) (0.42)	0.2139	0.1180 (0.00) (0.16)	0.6903 (0.58) (0.86)	0.0745 (0.82) (0.89)	0.2433
Shrinkage	0.1185 (0.00) (0.12)	0.7787 (0.72) (0.67)	0.0853 (0.91) (0.69)	0.4183	0.1107 (0.00) (0.00)	0.6047 (0.72) (0.53)	0.0608 (0.91) (0.58)	0.8444
<i>Panel B: Monthly rebalancing</i>								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7339	0.1107	0.0673
Sample cov	0.1136 (0.00) (0.04)	0.9282 (0.53) (0.20)	0.0990 (0.85) (0.31)	0.5133	0.1191 (0.00) (0.07)	0.7368 (0.50) (0.11)	0.0807 (0.79) (0.16)	1.2480
Constrained	0.1155 (0.00) (0.10)	0.8509 (0.69) (0.26)	0.0916 (0.94) (0.34)	0.2231	0.1089 (0.00) (0.12)	0.8863 (0.20) (0.74)	0.0907 (0.73) (0.83)	0.2496
Shrinkage	0.1119 (0.00) (0.01)	0.9104 (0.56) (0.27)	0.0956 (0.89) (0.44)	0.4529	0.1126 (0.00) (0.09)	0.6549 (0.67) (0.31)	0.0674 (0.92) (0.40)	0.9062

Table 5: Portfolios using implied correlations

In this table, we evaluate the performance of various portfolios that use option-implied correlations, as computed in Buss and Vilkov (2008), while volatilities are estimated from historical data. The “Sample cov” portfolio is the minimum-variance portfolio based on the sample covariance matrix but where option-implied correlations are used; “Constrained” is the portfolio based on the same covariance matrix as above, but where short sales are constrained; and, “Shrinkage” is the portfolio based on the same covariance matrix as above, but with shrinkage applied to the covariance matrix using the Ledoit and Wolf (2004a,b) 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).

Strategy	100 stocks				561 stocks			
	$\hat{\sigma}$	sr	ce	trn	$\hat{\sigma}$	sr	ce	trn
<i>Panel A: Daily rebalancing</i>								
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7489	0.1155	0.0144
Sample cov	0.1462 (0.00) (1.00)	0.2161 (0.99) (0.99)	0.0209 (0.99) (0.98)	0.4535	0.1302 (0.00) (0.07)	0.3780 (0.86) (0.61)	0.0407 (0.91) (0.63)	0.4957
Constrained	0.1311 (0.00) (1.00)	0.9001 (0.55) (0.15)	0.1094 (0.83) (0.10)	0.2109	0.1525 (0.00) (1.00)	0.8086 (0.28) (0.79)	0.1117 (0.59) (0.52)	0.1044
Shrinkage	0.1357 (0.00) (1.00)	0.3509 (0.98) (0.99)	0.0384 (0.99) (0.97)	0.2286	0.1344 (0.00) (1.00)	0.2743 (0.92) (0.88)	0.0278 (0.95) (0.85)	0.3354
<i>Panel B: Monthly rebalancing</i>								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7339	0.1107	0.0673
Sample cov	0.1455 (0.20) (1.00)	0.1660 (0.99) (1.00)	0.0135 (1.00) (1.00)	0.5098	0.1448 (0.02) (0.95)	0.2666 (0.90) (0.80)	0.0280 (0.93) (0.78)	0.5825
Constrained	0.1231 (0.00) (0.71)	0.8009 (0.81) (0.41)	0.0910 (0.96) (0.36)	0.2284	0.1450 (0.00) (1.00)	0.8488 (0.07) (0.81)	0.1125 (0.46) (0.50)	0.1445
Shrinkage	0.1428 (0.12) (1.00)	0.2545 (1.00) (1.00)	0.0261 (1.00) (1.00)	0.2820	0.1504 (0.06) (1.00)	0.1441 (0.95) (0.97)	0.0102 (0.96) (0.95)	0.4188

Table 6: Portfolios using historical volatilities scaled by volatility risk premium

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 of the top decile to $(1 - \delta)$ and of the bottom decile to $(1 + \delta)$, with $\delta_1 = 0.50$ and $\delta_2 = 0.70$. Correlations are estimated from historical data. The “Sample cov” portfolio is the minimum-variance portfolio based on the adjusted sample covariance matrix with volatilities scaled as described above; “Constrained” is based on the same adjusted covariance matrix but where short sales are constrained; and, “Shrinkage” is based on the same adjusted covariance matrix but with shrinkage applied to the sample covariance matrix. 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.

Strategy	100 stocks				561 stocks			
	$\hat{\sigma}$	sr	ce	trn	$\hat{\sigma}$	sr	ce	trn
<i>Panel A: Daily rebalancing</i>								
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7489	0.1155	0.0144
Sample cov: δ_1	0.1400 (0.00) (1.00)	1.2109 (0.13) (0.03)	0.1598 (0.26) (0.01)	0.1213	0.1343 (0.00) (0.44)	0.7288 (0.54) (0.18)	0.0889 (0.76) (0.18)	0.4248
Sample cov: δ_2	0.1442 (0.00) (1.00)	1.1922 (0.13) (0.06)	0.1616 (0.24) (0.02)	0.0876	0.1354 (0.00) (0.57)	0.7624 (0.47) (0.18)	0.0941 (0.74) (0.17)	0.2492
Constrained: δ_1	0.1384 (0.00) (1.00)	0.9465 (0.45) (0.09)	0.1214 (0.72) (0.04)	0.0470	0.1480 (0.00) (1.00)	0.9401 (0.03) (0.55)	0.1282 (0.25) (0.29)	0.0358
Constrained: δ_2	0.1398 (0.00) (1.00)	1.0049 (0.32) (0.05)	0.1307 (0.58) (0.01)	0.0326	0.1492 (0.00) (1.00)	0.8680 (0.08) (0.70)	0.1184 (0.43) (0.41)	0.0343
Shrinkage: δ_1	0.1378 (0.00) (1.00)	1.2269 (0.11) (0.03)	0.1596 (0.26) (0.01)	0.1092	0.1299 (0.00) (1.00)	0.7648 (0.47) (0.27)	0.0909 (0.76) (0.21)	0.2750
Shrinkage: δ_2	0.1431 (0.00) (1.00)	1.2049 (0.11) (0.06)	0.1622 (0.23) (0.02)	0.0808	0.1344 (0.00) (1.00)	0.7797 (0.43) (0.28)	0.0958 (0.74) (0.20)	0.1673
<i>Panel B: Monthly rebalancing</i>								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7339	0.1107	0.0673
Sample cov: δ_1	0.1381 (0.04) (0.99)	1.0961 (0.28) (0.07)	0.1419 (0.39) (0.02)	0.1841	0.1371 (0.00) (0.82)	0.7447 (0.48) (0.09)	0.0928 (0.71) (0.07)	0.5157
Sample cov: δ_2	0.1416 (0.08) (1.00)	1.0993 (0.26) (0.08)	0.1457 (0.35) (0.03)	0.1368	0.1370 (0.00) (0.83)	0.7903 (0.38) (0.07)	0.0990 (0.65) (0.06)	0.3068
Constrained: δ_1	0.1323 (0.00) (1.00)	0.9444 (0.48) (0.12)	0.1162 (0.77) (0.04)	0.0809	0.1469 (0.00) (1.00)	0.8935 (0.01) (0.75)	0.1204 (0.22) (0.37)	0.0819
Constrained: δ_2	0.1340 (0.00) (1.00)	1.0177 (0.31) (0.07)	0.1274 (0.59) (0.02)	0.0677	0.1492 (0.00) (1.00)	0.8708 (0.01) (0.79)	0.1187 (0.24) (0.40)	0.0813
Shrinkage: δ_1	0.1359 (0.02) (0.99)	1.1119 (0.25) (0.08)	0.1419 (0.39) (0.03)	0.1699	0.1371 (0.00) (1.00)	0.7507 (0.46) (0.16)	0.0936 (0.71) (0.09)	0.3585
Shrinkage: δ_2	0.1406 (0.07) (1.00)	1.1088 (0.24) (0.11)	0.1461 (0.33) (0.04)	0.1295	0.1393 (0.00) (1.00)	0.7803 (0.41) (0.14)	0.0990 (0.66) (0.07)	0.2245

Table 7: Portfolios using historical volatilities scaled by implied skewness

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 of the top decile to $(1 - \delta)$ and of the bottom decile to $(1 + \delta)$, with $\delta_1 = 0.30$ and $\delta_2 = 0.50$. Correlations are estimated from historical data. The “Sample cov” portfolio is the minimum-variance portfolio based on the adjusted sample covariance matrix with volatilities scaled as described above; “Constrained” is based on the same adjusted covariance matrix but where short sales are constrained; and, “Shrinkage” is based on the same adjusted covariance matrix but with shrinkage applied to the sample covariance matrix. 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.

Strategy	100 stocks				561 stocks			
	$\hat{\sigma}$	sr	ce	trn	$\hat{\sigma}$	sr	ce	trn
<i>Panel A: Daily rebalancing</i>								
1/N	0.1609	0.9286	0.1365	0.0129	0.1745	0.7489	0.1155	0.0144
Sample cov: δ_1	0.1417 (0.00) (1.00)	1.6164 (0.00) (0.00)	0.2190 (0.01) (0.00)	0.6574	0.1410 (0.00) (0.99)	1.4079 (0.01) (0.00)	0.1886 (0.04) (0.00)	1.8965
Sample cov: δ_2	0.1663 (0.96) (1.00)	1.7304 (0.00) (0.00)	0.2740 (0.00) (0.00)	0.7509	0.1543 (0.00) (1.00)	1.4576 (0.00) (0.00)	0.2130 (0.00) (0.00)	1.5829
Constrained: δ_1	0.1439 (0.00) (1.00)	1.1196 (0.05) (0.01)	0.1507 (0.22) (0.00)	0.2108	0.1651 (0.00) (1.00)	0.8075 (0.11) (0.79)	0.1197 (0.32) (0.42)	0.0300
Constrained: δ_2	0.1517 (0.00) (1.00)	1.2464 (0.02) (0.00)	0.1776 (0.05) (0.00)	0.2683	0.1653 (0.00) (1.00)	0.7923 (0.22) (0.80)	0.1173 (0.42) (0.44)	0.0491
Shrinkage: δ_1	0.1394 (0.00) (1.00)	1.6156 (0.00) (0.00)	0.2155 (0.01) (0.00)	0.5788	0.1319 (0.00) (1.00)	1.5254 (0.00) (0.00)	0.1926 (0.02) (0.00)	1.2677
Shrinkage: δ_2	0.1647 (0.89) (1.00)	1.7379 (0.00) (0.00)	0.2727 (0.00) (0.00)	0.6838	0.1527 (0.00) (1.00)	1.5133 (0.00) (0.00)	0.2194 (0.00) (0.00)	1.1407
<i>Panel B: Monthly rebalancing</i>								
1/N	0.1531	0.9402	0.1322	0.0595	0.1708	0.7339	0.1107	0.0673
Sample cov: δ_1	0.1454 (0.18) (1.00)	1.1664 (0.13) (0.00)	0.1592 (0.17) (0.00)	0.6949	0.1448 (0.00) (1.00)	1.0174 (0.08) (0.00)	0.1370 (0.20) (0.00)	1.9612
Sample cov: δ_2	0.1682 (0.97) (1.00)	1.1906 (0.07) (0.01)	0.1864 (0.02) (0.00)	0.7791	0.1653 (0.24) (1.00)	1.0349 (0.04) (0.00)	0.1576 (0.04) (0.00)	1.6204
Constrained: δ_1	0.1401 (0.00) (1.00)	1.0323 (0.15) (0.03)	0.1348 (0.42) (0.00)	0.2381	0.1628 (0.00) (1.00)	0.7662 (0.12) (0.91)	0.1114 (0.43) (0.51)	0.0806
Constrained: δ_2	0.1499 (0.24) (1.00)	1.0601 (0.13) (0.03)	0.1477 (0.16) (0.00)	0.2919	0.1617 (0.00) (1.00)	0.8249 (0.01) (0.85)	0.1203 (0.09) (0.39)	0.0982
Shrinkage: δ_1	0.1427 (0.08) (1.00)	1.1856 (0.09) (0.00)	0.1591 (0.15) (0.00)	0.6159	0.1432 (0.00) (1.00)	1.0670 (0.04) (0.00)	0.1426 (0.13) (0.00)	1.3281
Shrinkage: δ_2	0.1662 (0.96) (1.00)	1.2104 (0.04) (0.01)	0.1876 (0.01) (0.00)	0.7124	0.1705 (0.49) (1.00)	1.0499 (0.02) (0.01)	0.1647 (0.02) (0.00)	1.1787

Figure 1: Volatilities: Realized, historical, implied, and risk-premium-corrected implied

In this figure, we plot the historical volatility based on the 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. And, the historical 750-day volatility does not track realized volatility closely. 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 historical volatility is the smoothest.





Figure 2: Correlations: Realized, historical, and implied

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



References

- Aït-Sahalia, Y., and M. W. Brandt, 2008, "Consumption and Portfolio Choice With Option-Implied State Prices," NBER Working Paper.
- Aït-Sahalia, Y., and A. W. Lo, 1998, "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices," *Journal of Finance*, 53, 499–547.
- Andersen, T. G., T. Bollerslev, P. Christoffersen, and F. X. Diebold, 2006, "Volatility and Correlation Forecasting," in Graham Elliott, Clive W. J. Granger, and Allan G. Timmermann (ed.), *Handbook of Economic Forecasting*, North Holland.
- Andersen, T. G., T. Bollerslev, and F. X. Diebold, 2009, "Parametric and Nonparametric Measurement of Volatility," in Yacine Aït-Sahalia, and Lars Peter Hansen (ed.), *Handbook of Financial Econometrics*, North Holland, Amsterdam (forthcoming).
- Bakshi, G. S., N. Kapadia, and D. B. Madan, 2003, "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options," *The Review of Financial Studies*, 16, 101–143.
- Bali, T. G., and A. Hovakimian, 2009, "Volatility Spreads and Expected Stock Returns," *Management Science*, 55, 1797–1812.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard, 2009, "Realised Kernels in Practice: Trades and Quotes," *Econometrics Journal*.
- Barndorff-Nielsen, O. E., and N. Shephard, 2002, "Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models," *Journal of The Royal Statistical Society Series B*, 64, 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, 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 Heteroscedasticity," *Journal of Econometrics*, 31, 34–105.
- Bollerslev, T., R. Y. Chou, and K. Kroner, 1992, "ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 307–327.
- Bollerslev, T., M. Gibson, and H. Zhou, 2004, "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities," *Finance and Economics Discussion Series*, 2004-56, Board of Governors of the Federal Reserve System, Washington.
- 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, 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 Stephen Figlewski, William L. Silber, and Martin G. Subrahmanyam (ed.), *Financial Options: From Theory to Practice*, Business One-Irwin, Homewood, IL.
- Buss, A., and G. Vilkov, 2008, "Option-Implied Correlation and Factor Betas Revisited," Working Paper, Goethe University, Frankfurt.
- Carhart, M. M., 1997, "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52, 57–82.
- Carr, P., and L. Wu, 2009, "Variance Risk Premiums," *Review of Financial Studies*, 22, 1311–1341.
- Carr, P. P., and D. B. Madan, 1998, "Towards a Theory of Volatility Trading," in Robert A. Jarrow (ed.), *Volatility: New Estimation Techniques for Pricing Derivatives*, RISK Publications, London.
- Carr, P. P., and D. B. Madan, 2001, "Optimal Positioning in Derivative Securities," *Quantitative Finance*, 1, 19–37.
- Chan, L. K. C., J. Karceski, and J. Lakonishok, 1999, "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model," *Review of Financial Studies*, 12, 937–974.
- Chernov, M., 2007, "On the Role of Risk Premia in Volatility Forecasting," *Journal of Business & Economic Statistics*, 25, 411–426.
- Chernov, M., and E. Ghysels, 2000, "A Study towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation," *Journal of Financial Economics*, 56, 407–58.
- Christensen, B. J., and N. R. Prabhala, 1998, "The Relation between Implied and Realized Volatility," *Journal of Financial Economics*, 50, 125–50.
- Christoffersen, P., and B.-Y. Chang, 2009, "Option-Implied Measures of Equity Risk," Working Paper, McGill University.
- Corsi, F., G. Zumbach, U. Müller, and M. Dacorogna, 2001, "Consistent High-precision Volatility From High-Frequency Data," *Economic Notes*, 30, 183–204.
- Cox, J. C., and C.-f. Huang, 1989, "Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process," *Journal of Economic Theory*, 49, 33–83.
- Cremers, M., and D. Weinbaum, 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, 798–812.

- DeMiguel, V., L. Garlappi, and R. Uppal, 2009, "Optimal Versus Naive Diversification: How Inefficient Is the $1/N$ Portfolio Strategy? Portfolio Strategy?," *Review of Financial Studies*, 22, 1915–1953.
- Driessen, J., P. Maenhout, and G. Vilkov, 2009, "The Price of Correlation Risk: Evidence from Equity Options," *Journal of Finance*, 64, 1375–1404.
- Dumas, B., 1995, "The Meaning of the Implicit Volatility Function in Case of Stochastic Volatility," Working Paper, HEC, Paris.
- Efron, B., and R. Tibshirani, 1993, *An Introduction to the Bootstrap*, Chapman and Hall, Boca Raton, FL.
- Engle, R. F., 1982, "Autoregressive conditional heteroscedasticity with estimates of the variance of U.K. inflation," *Econometrica*, 50, 987–1008.
- Engle, R. F., 1993, "Statistical Models For Financial Volatility," *Financial Analysts Journal*, 49, 72–78.
- Engle, R. F., 2002, "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business and Economic Statistics*, 20, 339–350.
- Epps, T. W., 1979, "Comovements in Stock Prices in the Very Short Run," *Journal of the American Statistical Association*, 74, 291–298.
- Fama, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns on Stock and Bonds," *Journal of Financial Economics*, 33, 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, 1263–95.
- Heston, S., 1993, "A Closed-form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327–343.
- Jackwerth, J. C., 2000, "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, 13, 433–51.
- Jackwerth, J. C., and M. Rubinstein, 1996, "Recovering Probability Distributions from Option Prices," *Journal of Finance*, 51, 1611–32.
- Jacod, J., 1994, "Limit of Random Measures Associated with the Increments of a Brownian Semimartingale," *Preprint University Paris VI*.
- Jacod, J., and P. Protter, 1998, "Asymptotic Error Distributions for the Euler Method for Stochastic Differential Equations," *Annals of Probability*, 26, 267–307.
- Jagannathan, R., and T. Ma, 2003, "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps," *Journal of Finance*, 58, 1651–1684.
- Jiang, G. J., and Y. S. Tian, 2005, "The Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, 18, 1305–1342.
- Jobson, J. D., R. Korkie, and V. Ratti, 1979, "Improved Estimation for Markowitz Portfolios Using James-Stein Type Estimators," *Proceedings of the American Statistical Association*, 41, 279–292.

- Jorion, P., 1986, "Bayes-Stein Estimation for Portfolio Analysis," *Journal of Financial and Quantitative Analysis*, 21, 279–292.
- Kostakis, A., N. Panigirtzoglou, and G. Skiadopoulos, 2009, "Asset Allocation with Option-Implied Distributions: A Forward-Looking Approach," SSRN-id1288103.
- Kraus, A., and R. H. Litzenberger, 1976, "Skewness Preference and the Valuation of Risk Assets," *Journal of Finance*, 31, 1085–1100.
- Lamoureux, C. G., and W. D. Lastrapes, 1993, "Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities," *Review of Financial Studies*, 6, 293–326.
- Latane, H. A., and J. Rendleman, Richard J., 1976, "Standard Deviations of Stock Price Ratios Implied in Option Prices," *Journal of Finance*, 31, 369–381.
- Ledoit, O., and M. Wolf, 2004a, "Honey, I Shrank the Sample Covariance Matrix," *Journal of Portfolio Management*, 30, 110–119.
- Ledoit, O., and M. Wolf, 2004b, "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices," *Journal of Multivariate Analysis*, 88, 365–411.
- Ledoit, O., and M. Wolf, 2008, "Robust Performance Hypothesis Testing with the Sharpe Ratio," *Journal of Empirical Finance*, 15, 850–859.
- MacKinlay, A. C., and Ľ. Pástor, 2000, "Asset Pricing Models: Implications for Expected Returns and Portfolio Selection," *Review of Financial Studies*, 13, 883–916.
- Merton, R. C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.
- Merton, R. C., 1973, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.
- Merton, R. C., 1980, "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323–361.
- 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.
- Pástor, Ľ., 2000, "Portfolio Selection and Asset Pricing Models," *Journal of Finance*, 55, 179–223.
- Pástor, Ľ., and R. F. Stambaugh, 2000, "Comparing Asset Pricing Models: An Investment Perspective," *Journal of Financial Economics*, 56, 335–81.
- Politis, D., and J. Romano, 1994, "The Stationary Bootstrap," *Journal of the American Statistical Association*, 89, 1303–1313.
- Poon, S.-H., and C. W. J. Granger, 2005, "Practical Issues in Forecasting Volatility," *Financial Analysts Journal*, 61, 45–56.
- Rehman, Z., and G. Vilkov, 2009, "Option-Implied Skewness as a Stock-Specific Sentiment Measure," Working Paper, Goethe University, Frankfurt and SSRN id=1301648.
- Rubinstein, M. E., 1973, "The Fundamental Theorem of Parameter-Preference Security Valuation," *Journal of Financial and Quantitative Analysis*, 8, 61–69.
- Vanden, J. M., 2008, "Information Quality and Options," *Review of Financial Studies*, 21, 2635–2676.

- Xing, Y., X. Zhang, and R. Zhao, 2009, “What Does Individual Option Volatility Smirk Tell Us About Future Equity Returns?,” forthcoming in *Journal of Financial and Quantitative Analysis*.
- Zhang, L., P. A. Mykland, and Y. Aït-Sahalia, 2005, “A Tale of Two Time Scales: Determining Integrated Volatility With Noisy High-Frequency Data,” *Journal of the American Statistical Association*, 100, 1394–1411.
- Zhou, B., 1996, “High-Frequency Data and Volatility in Foreign-Exchange Rates,” *Journal of Business and Economic Statistics*, 14, 45–52.
- Zumbach, G., 2009, “Inference on Multivariate ARCH Processes with Large Sizes,” *SSRN eLibrary*.