Portfolio Optimization Using Forward-Looking Information [1]

by

Alexander Kempf[2], Olaf Korn[3], and Sven Saßning[4]

[1]We thank Rainer Baule, Felix Brinkmann, Gjergji Cici, Alexander Pütz, and Monika Trapp for their helpful comments and suggestions. All remaining errors are ours.

[2]Alexander Kempf, Department of Finance and Centre for Financial Research Cologne (CFR), University of Cologne, D-50923 Cologne, Germany, Phone +49 221 470 2741, Fax + 49 221 470 3992, Email kempf@wiso.uni-koeln.de

[3]Olaf Korn, Chair of Finance, Georg-August-Universität Göttingen and Centre for Financial Research Cologne (CFR), Platz der Göttinger Sieben 3, D-37073 Göttingen, Germany, Phone +49 551 39 7265, Fax +49 551 39 7665, Email okorn@uni-goettingen.de

[4]Sven Saßning, Chair of Finance, Georg-August-Universität Göttingen, Platz der Göttinger Sieben 3, D-37073 Göttingen, Germany, Phone +49 551 39 8305, Fax +49 551 39 7665, Email ssassni@uni-goettingen.de

Portfolio Optimization Using Forward-Looking Information

Abstract

In this paper we develop the first estimator of the covariance matrix that relies solely on forward-looking information. This estimator only uses price information from a cross-section of plain-vanilla options. In an out-of-sample study for US blue-chip stocks we show that a minimum-variance strategy based on this fully implied estimator consistently outperforms a wide range of benchmark strategies, including strategies based on historical estimates, index investing, and investing according to the 1/N rule. The outperformance is strong in periods of high information asymmetry, whereas in quiet periods all strategies lead to similar results. The outperformance can only be reached using a fully implied approach; partially implied approaches that combine implied moments with historical ones might even perform worse than purely historical approaches.

JEL Classification: G11, G13, G17

1 Introduction

Selecting an optimal portfolio is a classical problem in finance. Although the solution for the mean-variance investor is well known since the seminal work by Markowitz (1952), implementation remains a challenging task. Generally, estimation errors make a simple implementation based on historical sample moments very unstable (see, for example, Best and Grauer (1991), Chopra and Ziemba (1993), and Michaud (1989)), and specifically, expected returns have proved hard to estimate (see Merton (1980)). Facing these implementation limitations, researchers are paying growing attention to the global minimum variance portfolio (GMVP), the only efficient portfolio that doesn't depend on expected returns. The GMVP often leads to a better out-of-sample performance than a mean-variance optimized portfolio (see, for example, Ledoit and Wolf (2003) and Jagannathan and Ma (2003)).

Despite the attractiveness of the GMVP, an estimation risk with respect to the covariance matrix still remains, and several recent papers have suggested ways to reduce estimation errors. For example, Ledoit and Wolf (2004) impose restrictions on the covariance matrix, whereas Jagannathan and Ma (2003) as well as DeMiguel, Garlappi, Nogales, and Uppal (2009) put restrictions on the portfolio weights. Despite all the progress that has been made, a recent study by DeMiguel, Garlappi, and Uppal (2009) concludes that " there are still many miles to go before the gains promised by optimal portfolio selection can actually be realized out of sample." Specifically, DeMiguel, Garlappi, and Uppal (2009) base their conclusion on the finding that the GMVP is unable to beat simple benchmark strategies like investing in an equally weighted portfolio (1/N-strategy).

Echoing these concerns and the limitations raised by previous research, we suggest a new approach to estimate the covariance matrix and show that the GMVP based on these estimates is able to beat the 1/N-strategy and other benchmark strategies in an out-of-sample test. The core idea of our new approach is to rely solely on current option prices when estimating the covariance matrix instead of using historical return information. Since option prices reflect the expectations of market participants about risk, our approach - unlike the backward-looking approaches used so far - is inherently forward-looking.

Our paper makes two major contributions to the literature. First, we develop a new estimator of the covariance matrix that uses exclusively forward-looking information from a cross-section of option prices. This estimator requires implied volatilities as well as implied correlations. While implied volatilities can be easily derived from plain-vanilla options, implied correlations cannot be derived in a similar way since cross-correlation derivatives, like exchange options and quantos, are usually not available.[1] Therefore, we suggest a different route by developing a model which allows us to derive implied correlations from a cross-section of plain-vanilla options.

Our second main contribution is to show that the GMVP based on the implied estimates of the covariance performs extremely well in an out-of-sample study for the US stock market. In particular, we obtain the following main empirical results: (i) The implied approach consistently beats all the benchmark strategies: GMVP based on historical estimates of the covariance matrix, 1/N-strategy, and index investments. (ii) This result is robust and holds with and without short sales restrictions, with portfolios being rebalanced at different frequencies, and with transactions costs taken into account. (iii) The superiority of the implied approach is mostly pronounced in crisis periods, while in quiet periods all strategies lead to similar results. (iv) The implied approach is superior only if implied variances and implied covariances are used at the same time. A partially implied approach that combines implied moments with historical ones can lead to even worse results than the exclusive use of historical moments.

Our paper relates to two strands of literature. First, we extend the scarce literature on implied estimation of covariances. Buss and Vilkov (2010) use cross-sectional information from option prices combined with time series information to estimate correlations. Thus, in contrast to our paper, their approach is not a fully implied one. Chang, Christoffersen, Jacobs, and Vainberg (2011) develop a beta estimator based on implied skewness that could be combined with an implied index variance to obtain fully implied covariances. However, positive definiteness of the resulting covariance matrix is not guaranteed.[2] In contrast, our estimator leads to a positive definite matrix by construction, which is essential for an application to portfolio

[1]Siegel (1995) suggests the use of cross-correlation derivatives to estimate implied betas. Campa and Chang (1998) and Walter and Lopez (2000) analyze the predictive power of implied correlations based on cross-correlation derivatives.

[2]We provide empirical evidence in Subsection III 3.3 that the problem is severe.

problems.

The second related strand of literature consists of the few papers that have used forward-looking information from option prices to solve portfolio problems. Kostakis, Panigirtzoglou, and Skiadopoulos (2011) show that implied distributions can be useful to solve the problem of how to allocate wealth between a market index and a risk-free asset. However, in this problem, correlations do not matter and Kostakis, Panigirtzoglou, and Skiadopoulos (2011) consequently make no attempt to estimate implied correlations. The same holds for the portfolio allocation problem studied by Jabbour, Peña, Vera, and Zuluaga (2008), who seek to find a portfolio's worst case Conditional Value-at-Risk. Aït-Sahalia and Brandt (2008) study the dynamic consumption and portfolio choice problem of an investor who can invest in the stock market, the bond market, and in a risk-free asset. They characterize the properties of consumption and portfolio rules using implied marginal distributions. Since they do not derive the joint implied distribution, they have to estimate the correlation between the bond and stock market from historical returns. DeMiguel, Plyakha, Uppal, and Vilkov (2010) analyze the portfolio selection problem among a set of stocks and provide evidence on minimum-variance portfolios, but they either combine implied variances with historical correlations or historical variances with partially implied correlations derived as in Buss and Vilkov (2010). Thus, DeMiguel, Plyakha, Uppal, and Vilkov (2010) do not consider a fully implied approach. This is the main difference of our paper, which is the first that presents and tests a fully implied approach to find the GMVP.

The remainder of the paper is organized as follows. In Section 2 we develop our fully implied estimator of the covariance matrix. In Section 3 we present our empirical study. Section 4 concludes.

2 A Fully Implied Covariance Estimator

To derive implied covariances from a cross-section of plain-vanilla options, we have to make two main assumptions. First, we assume that asset returns follow a generalized version of the Sharpe (1963) market model with time-varying coefficients:

$$R_{it} = \alpha_{it} + \beta_{it} R_{mt} + \epsilon_{it}, \qquad \forall\, i = 1, \ldots, N. \tag{1}$$

R_{it} and R_{mt} denote the returns of the ith asset and the market, respectively. ϵ_{it} is a zero mean idiosyncratic error term with positive variance that is independent of the market return. In addition, ϵ_{it} and ϵ_{jt} are independent for all $i \neq j$. α_{it} and β_{it} are model coefficients. The index t indicates that these coefficients can be time-varying. In the market model the return covariances depend only on the beta coefficients and the variance of the market return:

$$Cov(R_{it}, R_{jt}) = \beta_{it}\beta_{jt} Var(R_{mt}), \qquad \forall\, i \neq j\,. \tag{2}$$

Since the variance of the market return can be implied from traded index options, we are left with the problem of identifying betas from the prices of plain-vanilla options written on the individual assets. To solve this problem, we have to make our second assumption, a cross-sectional restriction on one moment of the return distribution. We choose the second moment and assume that the same proportion c_t of total variance is systematic for all assets. Note that the proportion c_t can vary over time. Under this assumption, $\beta_{it}^2 Var(R_{mt}) = c_t Var(R_{it})$ and $Var(\epsilon_{it}) = (1 - c_t) Var(R_{it})$ with $0 \leq c_t < 1$, i.e., stocks with a higher total variance have both a higher beta and a higher idiosyncratic variance. This assumption fits well with empirical evidence well known since Fama and MacBeth (1973). Our cross-sectional restriction guarantees that the covariance matrix is positive definite. The return variance of the ith asset equals

$$Var(R_{it}) = \beta_{it}^2 Var(R_{mt}) + (1 - c_t)\, Var(R_{it}). \tag{3}$$

Solving for β_{it} leads to

$$\beta_{it} = c_t^{1/2} \left(\frac{Var(R_{it})}{Var(R_{mt})} \right)^{1/2}. \tag{4}$$

Since the market beta equals one, we can use the weights w_{itm} of the different assets $i =$

$1, \ldots, N$ in the market portfolio to identify the parameter c_t:

$$\sum_{i=1}^{N} w_{itm}\, \beta_{it} = \sum_{i=1}^{N} w_{itm}\, c_t^{1/2} \left(\frac{Var(R_{it})}{Var(R_{mt})} \right)^{1/2} = 1$$

$$\Leftrightarrow \quad c_t = \frac{Var(R_{mt})}{\left(\sum_{i=1}^{N} w_{itm}\, Var(R_{it})^{1/2} \right)^2} \quad . \tag{5}$$

Substitution of c_t from equation (5) into equation (4) and substitution of the resulting betas into equation (2) finally leads to

$$Cov(R_{it}, R_{jt}) = \frac{Var(R_{it})^{1/2}\, Var(R_{jt})^{1/2}}{\left(\sum_{i=1}^{N} w_{itm}\, Var(R_{it})^{1/2} \right)^2} Var(R_{mt}), \qquad \forall\, i \neq j. \tag{6}$$

Equation (6) shows that the covariances are functions of individual asset variances and the variance of the market only. No cross-moments appear. Therefore, we can use implied volatilities from plain-vanilla options written on individual assets and the market index to obtain a fully implied covariance estimate.

The estimator (6) is just one version of a whole class of implied estimators. Estimators based on higher moments can also be constructed. In Section III 3.3 we present two alternative versions of the estimator, a skewness-based and a kurtosis-based one, and analyze their performance. Compared to the variance-based estimator, skewness- and kurtosis-based estimators have the disadvantage that they do not guarantee a positive definite covariance matrix. However, they do not imply a constant correlation of c_t between all assets, as does the variance-based estimator according to equations (5) and (6).

Options provide moments under the risk-neutral measure, whereas portfolio selection requires moments under the physical measure. However, we make no attempt to risk-adjust implied moments in our study. Rubinstein (1994) provides examples of how the risk-neutral and physical distributions of the market are related. He concludes that the difference lies mainly in a mean shift and the distributions are similar in shape. Moreover, even if the shape of the distribution is changed, the GMVP is not affected as long as the proportional variance risk premia are the same for all assets. Ultimately, it is an empirical question of how valuable implied moments from option prices are for investment strategies[3] - a question which is addressed in the section to come.

3 Empirical Study

3.1 Data and Portfolio Strategies

Our sample consists of stocks included in the Dow Jones Industrial Average. We use daily prices (adjusted for dividends and stock splits) of all stocks that had been index components over the period from January 1998 to October 2009, which is our out-of-sample period. In addition, we use stock prices from January 1993 to December 1997 to estimate covariance matrices for historical benchmark strategies. The data source is Datastream. Table 1 provides information on the stock price data for the out-of-sample period. It shows annualized average excess log-returns, return volatilities, and the period for which the stocks have been included in the index.[4]

[Insert Table 1 about here]

Table 1 shows that excess returns of stocks included over the whole period range from -4.8% to 9.5% p.a. and volatilities from 20% to 40%. The Dow Jones Index itself has an excess return of 1%. The index volatility is 17.7%, providing evidence for potential diversification benefits.

To estimate the covariance matrix we use model-free implied moments, i.e., we do not rely on a particular valuation model. The idea goes back to the seminal paper by Breeden and Litzenberger (1978), who show that the whole risk-neutral return distribution can be recovered from option prices if a continuum of strike prices is available.[5] We apply the approach of Bakshi, Kapadia, and Madan (2003) to derive model-free implied variances

[3] A large literature shows that implied volatilities are useful to predict future realized volatilities. See Poon and Granger (2003) and Poon and Granger (2005) for surveys and Busch, Christensen, and Ørregaard Nielsen (2011) for a recent study that uses state-of-the-art benchmark predictors.

[4] For the calculation of excess returns, we used the zero rates provided by the IvyDB data base. Interest rates provided by IvyDB are derived from BBA LIBOR rates and settlement prices of CME Eurodollar futures. We interpolate them with a cubic spline to get the appropriate yield.

[5] See, for example, Britten-Jones and Neuberger (2000), Carr and Madan (2001), Jiang and Tian (2005), and Vanden (2008) for more recent theoretical and empirical research on model-free implied moments.

which we then put into equation (6) to obtain implied covariance estimates. As our data source we use the volatility surfaces provided by IvyDB. They deliver Black-Scholes implied volatilities for a variety of strike prices and maturity buckets. These surfaces are available for all individual stocks and the Dow Jones Index.

When implementing our trading strategy, we use a monthly rebalancing frequency for the portfolio in the base case. We choose the first trading day after the expiration day of options contracts at CBOE since liquid options with a time to maturity of about 30 days exist then. For this maturity bucket, we select all out-of-the-money put and call options and fit a cubic spline to obtain a smooth volatility curve for each stock and the index.[6] Outside the available range of strike prices, we assume that the volatility curve is flat. Then, we select 1000 equally spaced strike prices on the interval $[0.003 \cdot S_i, 3 \cdot S_i]$, where S_i denotes the current spot price of the ith asset. For these 1000 strike prices we convert back the implied volatilities into European option prices via the Black-Scholes formula. To do so, we use the matching spot prices and the risk-free interest rates provided by IvyDB. Based on these option prices we calculate the model-free implied moments following Bakshi, Kapadia, and Madan (2003).

For every month in the out-of-sample period, we set up GMVPs which differ in the way the covariance matrices are estimated. The first GMVP strategy is based on fully implied estimates of the covariance matrix (Implied). Another two GMVP strategies serve as historical benchmarks. The first one uses daily returns of the preceding 60 days to estimate the covariance matrix (Hist_60D). The second strategy uses monthly returns of the preceding 60 months (Hist_60M). Since both strategies employ the same number of observations, they allow us to look at the consequences of using more or less recent information. A natural hypothesis is that the Hist_60D strategy performs more like the implied one, which is based on current market information only. As further benchmarks, we consider three passive strategies for which no estimates are needed. The first one follows the construction of the Dow Jones Industrial Average and uses a price weighting of the 30 stocks (Pass_DJ). The second strategy uses a capital weighting (Pass_CW) and the third gives equal weights to all stocks (Pass_1/N).

[6]Using out-of-the-money options diminishes the price effect of a potential early exercise of American options, which reduces model risk with respect to the early exercise premium. Moreover, out-of-the money options are usually much more liquid than in-the-money options.

For each strategy and every month in our out-of-sample period we calculate the excess log-return and the realized volatility. Realized volatility is calculated as the volatility of the daily returns within the particular month.

3.2 Main Results

3.2.1 The Unrestricted GMVP and the Role of Short-sales Constraints

Table 2 shows the average realized volatilities ($\overline{\sigma}$) and average realized excess returns ($\overline{R}$) of the different portfolio strategies. For ease of interpretation, annualized values are provided. The p-values in brackets refer to tests of significant differences between the implied strategy and the five benchmark strategies.

When looking at the unrestricted GMVP (Short Sales Allowed) we see that the implied strategy has a volatility of 16.5%. This value is clearly smaller than the ones obtained from the benchmark strategies. The differences in volatility are statistically significant at least at a five percent level. Since no statistically significant difference exists for mean returns, the implied strategy beats all benchmark strategies with respect to the risk-return trade off. Interestingly, the historical strategies do not outperform the passive ones in terms of volatility reduction. Therefore, variance minimization based on historical covariance matrices would be less attractive out of sample, which is in line with previous findings in DeMiguel, Garlappi, and Uppal (2009).

[Insert Table 2 about here]

It is well documented that short-sales constraints often improve the out-of-sample performance of portfolio strategies based on historical parameter estimates since they prevent extreme stock positions.[7] This happens in our sample as well. The volatilities decline from 19.1% to 14.9% (Hist_60D) and from 20% to 15.7% (Hist_60M) when short-sales restrictions are introduced. Nevertheless, the implied estimator remains superior. With respect to average returns, there is still no significant difference between any of the strategies.

[7] See, for example, Frost and Savarino (1988), Grauer and Shen (2000), and Jagannathan and Ma (2003).

3.2.2 The Effects of Crisis

Implied moments are derived from current option prices and are not based on historical information. Therefore, one expects the implied strategy to perform particularly well in crisis periods for two reasons. First, the information flow is high in crisis periods. This could make historical information less valuable leading to superiority of the implied approach, which relies solely on current data. Second, informed investors can better exploit their private information in options markets than in spot markets.[8] Therefore, one expects to see a higher fraction of informed investors in the options market during periods of high information asymmetry. This would make option prices particular informative, again suggesting that the implied strategy performs particularly well in crisis periods. In quiet market periods, however, when no major events occur and the information asymmetry is low, implied moments might not have a big advantage over historical ones. We test these hypotheses now. In doing so, we concentrate on strategies with short-sales constraints because of their superior performance.[9]

The sample period between 1998 and 2009 contains two major stock market crisis. The first crisis is the burst of the Dot-com bubble; the second one the global financial crisis. Therefore, we divide the whole out-of-sample period into "crisis periods" and "non-crisis periods". The burst of the Dot-com bubble began in March 2000, when the NASDAQ index lost almost nine percent of its value in just six days. As the end of the Dot-com crisis we choose April 2003, the month when the stock market started its recovery. As the starting point of the global financial crisis, we use June 2007, the month when the problems of two of Bear Stearns' hedge funds became public, and we consider the whole remaining sample period until October 2009 a crisis period. This classification leaves us with 67 observations in crisis periods and 75 observations in non-crisis periods. Table 3 provides the average realized volatilities and excess returns of all strategies for crisis and non-crisis periods.

[Insert Table 3 about here]

Table 3 clearly shows that volatility is higher in the crisis periods than in the non-crisis periods, which indicates that our definition of crisis periods is sensible. When looking at crisis

[8] See, for example, Kumar, Sarin, and Shastri (1992), Easley, O'Hara, and Srinivas (1998), Chakravarty, Gulen, and Mayhew (2004), Cao, Chen, and Griffin (2005), and Pan and Poteshman (2006).

[9] Allowing for short sales leads to qualitatively identical results.

periods, we clearly see that the implied strategy is best in terms of realized volatility and the passive strategies are worst. The historical strategies are in-between, with the strategy using 60 previous trading days delivering a significantly lower (1% significance level) volatility than the one that uses the last 60 months. This ordering of the strategies suggests that it is crucial to adapt quickly to new information in times of crisis. Therefore, the implied estimator performs best.

The non-crisis periods provide a different picture of the performance of various strategies. Implied and historical strategies lead to almost identical volatilities. Only the passive strategies have a somehow higher volatility than the other ones, but the disadvantage is much smaller than in the crisis periods. Thus, in quiet periods it makes no big difference how one estimates the covariance matrix.

3.2.3 The Contribution of Implied Variances and Implied Correlations

Using current market information to estimate the covariance matrix leads to a lower portfolio volatility. We now analyze the contribution of implied variances and implied correlations to this volatility reduction. In Table 4 we compare realized portfolio volatilities for the fully implied approach with two partially implied approaches. The first is based on implied variances and historical correlations, and the second approach combines historical variances with implied correlations. For the historical estimates, we use the returns of the previous 60 days and 60 months, respectively.

The upper part of Table 4 partly repeats results from Tables 2 and 3 for the implied strategy. In addition, it includes p-values which refer to tests of significant differences between the numbers obtained from the fully implied approach and the partially implied approaches which are shown in the lower part of the table.

[Insert Table 4 about here]

Most importantly, Table 4 shows that the fully implied strategy delivers the lowest volatilities, most pronounced in crisis periods. The differences in volatility between the fully implied approach and the partially implied approaches are significant at the 1%-level.

The strategy that combines historical variances with implied correlations generally delivers the highest realized volatilities. They are in some cases even higher than the ones resulting from a fully historical approach.[10] This result shows that using implied moments can even be harmful if we combine them with historical ones.

If we use implied variances together with historical correlations, we see that some of the improvement of the implied strategies can be attributed to implied variances. However, for crisis periods, this partially implied approach still leads to worse results than the fully implied one. Only the joint use of implied variances and implied correlations allows implied information to exploit their full potential.

3.3 Robustness Checks

In this section, we run three robustness tests: First, we consider alternative implied covariance estimators based on higher moments. Second, we test whether our results still hold if we change the rebalancing frequency. Third, we check whether our main conclusions remain valid if we take transactions costs into account.

3.3.1 Alternative Implied Covariance Estimators

The implied covariance estimator of Section 2 was derived from a linear factor model and a cross-sectional restriction on the proportion of systematic variance. Within the same framework, we can derive alternative implied estimators based on higher moments. We just have to replace the assumption concerning the proportion of systematic variance with a corresponding assumption about how systematic risk affects higher moments. We will demonstrate the approach based on the third and fourth moments.

(i) Skewness-Based Estimator of Covariances:[11] We now assume that the proportion of systematic return skewness is equal for all assets.[12] Denote this proportion by c_t^{Skew}. Then,

[10] Compare the results for the Hist_60D and the Hist_60M strategy in Tables 2 and 3.

[11] When referring to skewness and kurtosis, we mean the third and fourth central moment without any re-scaling or normalization.

[12] Note that the assumption by Chang, Christoffersen, Jacobs, and Vainberg (2011) that the proportion of systematic skewness equals 100% is a special case.

the return skewness of the ith asset is

$$Skew(R_{it}) = \beta_{it}^3 Skew(R_{mt}) + (1 - c_t^{Skew})\, Skew(R_{it}), \tag{7}$$

and solving for β_{it} leads to

$$\beta_{it} = (c_t^{Skew})^{1/3} \left(\frac{Skew(R_{it})}{Skew(R_{mt})} \right)^{1/3}. \tag{8}$$

Again, the proportion c_t^{Skew} can be identified from the condition that the market beta equals one. Solving for c_t^{Skew} and substitution of the solution into equation (8) provides the beta coefficients and finally leads to the following covariances:

$$Cov(R_{it}, R_{jt}) = \frac{Skew(R_{it})^{1/3} Skew(R_{jt})^{1/3}}{\left(\sum_{i=1}^{N} w_{itm} Skew(R_{it})^{1/3} \right)^2} Var(R_{mt}), \qquad \forall\, i \neq j\,. \tag{9}$$

Equation (9) provides our first alternative implied estimator of covariances.

(ii) Kurtosis-Based Estimator of Covariances: Another way to estimate covariances is to use fourth moments. Assume that the proportion of systematic kurtosis is equal for all N assets and denote this proportion by c_t^{Kurt}. Based on this assumption we can derive our second alternative implied estimator of covariances as

$$Cov(R_{it}, R_{jt}) = \frac{Kurt(R_{it})^{1/4} Kurt(R_{jt})^{1/4}}{\left(\sum_{i=1}^{N} w_{itm} Kurt(R_{it})^{1/4} \right)^2} Var(R_{mt}), \qquad \forall\, i \neq j\,. \tag{10}$$

However, note that the skewness- and kurtosis-based estimators might not lead to a positive definite covariance matrix since the implied beta estimates together with the implied stock and index volatilities can induce a negative idiosyncratic variance. This is the main disadvantage of using higher moments as compared to the variance-based estimator of Section 2, which guarantees the positive definiteness of the resulting covariance matrix.

Using the model-free implied skewness and kurtosis from Bakshi, Kapadia, and Madan (2003), we provide empirical evidence on the severity of this problem for our data set. The skewness-based estimator leads to a negative or very small positive (less than 5% of total variance) idiosyncratic variance 259 times (out of 4260 times).[13] The problem is much less severe for the kurtosis-based estimator. In our sample, this estimator does not lead to a negative or very small idiosyncratic variance in a single case. This finding suggests that the variance-based and kurtosis-based estimators are more promising than the skewness-based estimator. Martellini and Ziemann (2010) provide complementary evidence on this issue for historical moment estimates by showing that odd moments are much more difficult to estimate than even ones. Ultimately, however, the crucial question is how the GMVPs resulting from different estimators perform out of sample.

Table 5 presents results for the portfolio strategies using implied covariance estimators based on the second (Imp_Var), third (Imp_Skew), and fourth moment (Imp_Kurt), respectively.

[Insert Table 5 about here]

For all periods (full sample, periods of crisis, periods of no crisis), the variance-based and the kurtosis-based estimator show almost identical results. In contrast, the skewness-based estimator always leads to a higher realized volatility. The difference is statistically significant. These results indicate that the implementation problems of the skewness-based estimator translate into a poor performance of the corresponding portfolio strategy.

3.3.2 Quarterly Rebalancing

As a second robustness check we repeat our analyses using quarterly rebalancing. We now extract implied moments from options with a remaining time to maturity of 91 days. Table 6 shows the results. They are based on 142 overlapping three-months periods. Again, we report annualized values of the average realized volatilities and excess returns. The p-values of the significance tests are based on the heteroscedasticity and autocorrelation consistent estimator by Newey and West (1987) with two lags, which accounts for the autocorrelation caused by the overlapping intervals.

[13]In these cases, we construct the implied covariance matrix by using an adjusted beta that ensures an idiosyncratic variance equal to 5% of total variance. This adjustment is likely to reduce the precision of the covariance estimates based on skewness. If we make the assumption of zero idiosyncratic skewness, as in Chang, Christoffersen, Jacobs, and Vainberg (2011), the situation becomes even worse. The implied idiosyncratic variance would be negative or very small in 1389 cases. Due to the great number of necessary adjustments, we do not consider this case any further.

[Insert Table 6 about here]

Table 6 shows that our main results are not changed when using quarterly rebalancing. The implied strategies perform better than the benchmark strategies, particularly in times of crisis.[14] Thus, the performance of the implied approach is robust to a change from monthly to quarterly rebalancing.

3.3.3 Transactions Costs

As our third robustness check we repeat our analysis, but now take transactions costs into account. Since the implied approach adapts quickly to new information, implied strategies might cause a higher turnover and, thus, higher transactions costs than historical ones. This might deteriorate the performance of the implied strategies.

To get a first impression of the trade intensity of the strategies we calculate the average turnovers. We measure average turnover as the average fraction of the portfolio that has to be traded (purchased and sold) at a rebalancing date. This fraction is within the interval $[0, 2]$ when short sales are not allowed. Table 7 shows the average turnover for different strategies and periods.

[Insert Table 7 about here]

As expected, the passive strategies and the historical strategy based on 60 months have the lowest turnover. The strategies using more recent information lead to higher turnovers, but the implied strategies do not cause a higher turnover than the historical strategy based on the previous 60 days.

We now quantify the effect of transactions costs on the expected returns and volatilities of different portfolio strategies. To do so, we need to specify the level of transactions costs. Venkataraman (2001) reports average quoted and realized spreads for a group of 40 major US stocks that were traded on NYSE between July 1997 and March 1998, which includes the beginning of our sample period. The quoted spread equals 24 bp and the realized spread 16 bp. Jain (2003) provides spread information for the period from January to April 2000.

[14]Again, the skewness based estimator is the worst within the group of implied estimators. It fails to beat the historical estimator based on 60 days, but performs better than the other benchmark strategies.

He reports an average quoted spread of 20 bp and a realized spread of 16 bp for the 25 NYSE stocks with the highest market capitalization. Roll, Schwartz, and Subrahmanyam (2007) construct a liquidity index and show that there was a further decrease in transactions costs after the reduction of the minimum tick-size from eights to cents in January 2001. Based on the evidence from the literature, we consider 20 bp a conservative estimate of the spread during our sample period and choose a halfspread of 10 bp as percentage transactions costs in our study.

As expected, the transactions costs mainly influence the average return. Given the transactions costs estimate of 10 bp we find that the volatilities of all trading strategies change by less than one basis point. None of the reported realized volatilities from previous tables would have to be changed after adjusting for transactions costs. Therefore, we report in Table 8 only the effect of transactions costs on average returns.

[Insert Table 8 about here]

Obviously, the mean returns are lower now, but the differences between the strategies remain insignificant. Thus, our main results remain valid even when considering transactions costs.

4 Conclusions

In this paper, we develop the first fully implied estimator of the covariance matrix. Our basic idea is to obtain estimates of the covariance matrix solely from current prices of plain-vanilla options. These prices reflect market expectations about the return distributions of the underlying assets. Therefore, the approach is forward-looking and differs fundamentally from the backward-looking approach that uses time series of returns. We show that one can set up different versions of the new estimator. They differ with respect to the moment of the implied distribution they use to obtain the covariance matrix. We implement versions based on the second to fourth moment, but recommend use of the variance based one. This is the only one which guarantees that the resulting implied covariance matrix is positive definite.

We test the quality of this new estimator by analyzing the out-of-sample performance of a corresponding global minimum variance (GMVP) strategy. We compare its performance with several benchmark strategies, namely GMVP strategies based on historical estimates, 1/N-strategy, and index investments. Using a sample of US blue chips we get the following main results: (i) Our implied approach consistently beats all the benchmark strategies. (ii) This superior result is very robust. It holds no matter whether short sales are allowed or not, it does not depend on the rebalancing frequency of the portfolios, and it does not change when reasonable transactions costs are taken into account. (iii) The superiority of the implied approach results from its superiority in crisis periods, i.e., adapting to current information is particularly useful when the speed of information flow is high. (iv) The superiority of the implied approach can only be reached when implied variances and implied covariances are used at the same time. A partially implied approach that combines implied moments with historical ones can lead to even worse results than the exclusive use of historical moments.

Although our paper is only a first step, it suggests that we can do better than the simple 1/N-strategy by using forward-looking instead of historical information. The natural next step would be to combine the fully implied covariance matrix suggested in this paper with option-implied information about expected returns. This would allow us to go beyond the GMVP and possibly reach an even better out-of-sample performance. We believe that this is a promising avenue for further research since several recent papers like Ang, Bali, and Cakici (2010), Bali and Hovakimian (2009), Conrad, Dittmar, and Ghysels (2009), Cremers and Weinbaum (2010), DeMiguel, Plyakha, Uppal, and Vilkov (2010), Rehman and Vilkov (2010), and Xing, Zhang, and Zhao (2010) have shown that option-implied information has predictive power for expected returns.

References

Ang, A., T. G. Bali, and N. Cakici, 2010, "The Joint Cross Section of Stocks and Options," *Working Paper, NBER.*

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

Bakshi, G., N. Kapadia, and D. Madan, 2003, "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options," *The Review of Financial Studies*, 16(1), 101–143.

Bali, T. G., and A. Hovakimian, 2009, "Volatility Spreads and Expected Stock Returns," *Management Science*, 55, 1797–1812.

Best, M. J., and R. R. Grauer, 1991, "On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results," *Review of Financial Studies*, 4(2), 315–342.

Breeden, D. T., and R. H. Litzenberger, 1978, "Prices of State-Contingent Claims Implicit in Option Prices," *The Journal of Business*, 51(4), 621–651.

Britten-Jones, M., and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *The Journal of Finance*, 55(2), 839–866.

Busch, T., B. J. Christensen, and M. Ørregaard Nielsen, 2011, "The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets," *Journal of Econometrics*, 160(1), 48–57.

Buss, A., and G. Vilkov, 2010, "Option-Implied Correlation and Factor Betas Revisited," *Working Paper, Goethe University Frankfurt.*

Campa, J. M., and P. H. K. Chang, 1998, "The forecasting ability of correlations implied in foreign exchange options," *Journal of International Money and Finance*, 17(6), 855–880.

Cao, C., Z. Chen, and J. M. Griffin, 2005, "Informational Content of Option Volume Prior to Takeovers," *The Journal of Business*, 78(3), 1073–1109.

Carr, P., and D. Madan, 2001, "Optimal positioning in derivative securities," *Quantitative Finance*, 1(1), 19–37.

Chakravarty, S., H. Gulen, and S. Mayhew, 2004, "Informed Trading in Stock and Option Markets," *The Journal of Finance*, 59(3), 1235–1258.

Chang, B. Y., P. Christoffersen, K. Jacobs, and G. Vainberg, 2011, "Option-Implied Measures of Equity Risk," *Review of Finance*.

Chopra, V. K., and W. T. Ziemba, 1993, "The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice," *Journal of Portfolio Management*, 19(2), 6–11.

Conrad, J. S., R. F. Dittmar, and E. Ghysels, 2009, "Ex Ante Skewness and Expected Stock Returns," *Working Paper, University of North Carolina.*

Cremers, M., and D. Weinbaum, 2010, "Deviations from Put-Call Parity and Stock Return Predictability," *Journal of Financial and Quantitative Analysis*, 45(02), 335–367.

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?," *Review of Financial Studies*, 22(5), 1915–1953.

DeMiguel, V., Y. Plyakha, R. Uppal, and G. Vilkov, 2010, "Improving Portfolio Selection Using Option-Implied Volatility and Skewness," *Working Paper, Goethe University Frankfurt.*

Easley, D., M. O'Hara, and P. Srinivas, 1998, "Option Volume and Stock Prices: Evidence on Where Informed Traders Trade," *The Journal of Finance*, 53(2), 431–465.

Fama, E. F., and J. D. MacBeth, 1973, "Risk, Return, and Equilibrium: Empirical Tests," *The Journal of Political Economy*, 81(3), 607–636.

Frost, P. A., and J. E. Savarino, 1988, "For better performance: Constrain portfolio weights," *The Journal of Portfolio Management*, 15(1), 29–34.

Grauer, R. R., and F. C. Shen, 2000, "Do constraints improve portfolio performance?," *Journal of Banking & Finance*, 24(8), 1253–1274.

Jabbour, C., J. F. Peña, J. C. Vera, and L. F. Zuluaga, 2008, "An estimation-free, robust conditional value-at-risk allocation model," *Journal of Risk*, 11, 57–78.

Jagannathan, R., and T. Ma, 2003, "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps," *The Journal of Finance*, 58(4), 1651–1684.

Jain, P. K., 2003, "Institutional Design and Liquidity at Stock Exchanges around the World," *Working Paper, University of Memphis.*

Jiang, G. J., and Y. S. Tian, 2005, "The Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, 18(4), 1305–1342.

Kostakis, A., N. Panigirtzoglou, and G. Skiadopoulos, 2011, "Market timing with option-implied distributions: A forward looking approach," *Management Science*, forthcoming.

Kumar, R., A. Sarin, and K. Shastri, 1992, "The Behavior of Option Price Around Large Block Transactions in the Underlying Security," *The Journal of Finance*, 47(3), 879–889.

Ledoit, O., and M. Wolf, 2003, "Improved estimation of the covariance matrix of stock returns with an application to portfolio selection," *Journal of Empirical Finance*, 10(5), 603–621.

——— , 2004, "Honey, I Shrunk the Sample Covariance Matrix," *Journal of Portfolio Management*, 30(4), 110–119.

Markowitz, H., 1952, "Portfolio Selection," *The Journal of Finance*, 7(1), 77–91.

Martellini, L., and V. Ziemann, 2010, "Improved Estimates of Higher-Order Comoments and Implications for Portfolio Selection," *Review of Financial Studies*, 23(4), 1467–1502.

Merton, R. C., 1980, "On estimating the expected return on the market: An exploratory investigation," *Journal of Financial Economics*, 8(4), 323–361.

Michaud, R. O., 1989, “The Markowitz Optimization Enigma: Is ’Optimized’ Optimal?,” *Financial Analysts Journal*, 45(1), 31–42.

Newey, W. K., and K. D. West, 1987, “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55(3), 703–708.

Pan, J., and A. M. Poteshman, 2006, “The Information in Option Volume for Future Stock Prices,” *Review of Financial Studies*, 19(3), 871–908.

Poon, S.-H., and C. W. J. Granger, 2003, “Forecasting Volatility in Financial Markets: A Review,” *Journal of Economic Literature*, 41(2), 478–539.

——— , 2005, “Practical issues in forecasting volatility,” *Financial Analysts Journal*, 61(1), 45–56.

Rehman, Z. U., and G. Vilkov, 2010, “Risk-neutral skewness: Return predictability and its sources,” *Working Paper, Goethe University Frankfurt.*

Roll, R., E. Schwartz, and A. Subrahmanyam, 2007, “Liquidity and the Law of One Price: The Case of the Futures-Cash Basis,” *The Journal of Finance*, 62(5), 2201–2234.

Rubinstein, M., 1994, “Implied Binomial Trees,” *The Journal of Finance*, 49(3), 771–818.

Sharpe, W. F., 1963, “A Simplified Model for Portfolio Analysis,” *Management Science*, 9(2), 277–293.

Siegel, A. F., 1995, “Measuring Systematic Risk Using Implicit Beta,” *Management Science*, 41(1), 124–128.

Vanden, J. M., 2008, “Information Quality and Options,” *Review of Financial Studies*, 21(6), 2635–2676.

Venkataraman, K., 2001, “Automated Versus Floor Trading: An Analysis of Execution Costs on the Paris and New York Exchanges,” *The Journal of Finance*, 56(4), 1445–1485.

Walter, C. A., and J. A. Lopez, 2000, “Is Implied Correlation Worth Calculating?,” *The Journal of Derivatives*, 7(3), 65–81.

Xing, Y., X. Zhang, and R. Zhao, 2010, "What Does the Individual Option Volatility Smirk Tell Us About Future Equity Returns?," *Journal of Financial and Quantitative Analysis*, 45(03), 641–662.

Table 1: Average Excess Returns and Realized Volatilities of Individual Stocks and the Dow Jones Index.

Tickers	$\overline{R}$	$\overline{\sigma}$	Period
UTX	0.095	0.282	01/98 - 10/09
CAT	0.07	0.337	01/98 - 10/09
MCD	0.063	0.270	01/98 - 10/09
XOM	0.063	0.250	01/98 - 10/09
WMT	0.056	0.275	01/98 - 10/09
IBM	0.047	0.280	01/98 - 10/09
MMM	0.041	0.246	01/98 - 10/09
HPQ	0.035	0.390	01/98 - 10/09
JNJ	0.034	0.206	01/98 - 10/09
PG	0.022	0.230	01/98 - 10/09
AXP	0.022	0.367	01/98 - 10/09
JPM	0.012	0.389	01/98 - 10/09
BA	-0.002	0.321	01/98 - 10/09
KO	-0.026	0.228	01/98 - 10/09
DIS	-0.033	0.328	01/98 - 10/09
AA	-0.038	0.399	01/98 - 10/09
MRK	-0.039	0.291	01/98 - 10/09
DD	-0.043	0.295	01/98 - 10/09
GE	-0.048	0.295	01/98 - 10/09
DJ Index	0.010	0.177	01/98 - 10/09
TRV	0.227	0.208	06/09 - 10/09
UK	0.181	0.393	01/98 - 11/99
MO	0.090	0.270	01/98 - 02/08
CVX	0.063	0.282	01/98 - 11/99
T	0.047	0.219	11/05 - 10/09
CSCO	0.040	0.273	06/09 - 10/09
HON	0.026	0.323	01/98 - 02/08
VZ	0.018	0.205	04/04 - 10/09
KFT	-0.004	0.241	09/08 - 10/09
CVX	-0.015	0.360	02/08 - 10/09
IP	-0.027	0.340	01/98 - 04/04
MSFT	-0.061	0.304	11/99 - 10/09
HD	-0.064	0.326	11/99 - 10/09
INTC	-0.073	0.395	11/99 - 10/09
PFE	-0.108	0.236	04/04 - 10/09
SBC	-0.136	0.331	11/99 - 11/05
EK	-0.144	0.333	01/98 - 04/04
C	-0.192	0.406	01/98 - 06/09
T	-0.205	0.412	01/98 - 04/04
S	-0.253	0.370	01/98 - 11/99
MTLQQ	-0.326	0.462	01/98 - 06/09
GT	-0.342	0.322	01/98 - 11/99
BAC	-0.463	0.779	02/08 - 10/09
AIG	-0.653	0.332	04/04 - 09/08

This table shows average excess returns and realized volatilities of individual stocks and the Dow Jones Index for the period from January 1998 to October 2009. Calculations are based on daily data and annualized values are reported. All stocks that are in the index during the full period are included in the upper part of the table. The lower part of the table lists all other stocks together with the periods when they have been in the index.

Table 2: Out-of-Sample Volatilities and Excess Returns of the Portfolio Strategies.

	Short Sales Allowed		Short Sales not Allowed	
	$\overline{\sigma}$	$\overline{R}$	$\overline{\sigma}$	$\overline{R}$
Implied	**0.165**	**-0.019**	**0.143**	**0.014**
	(0.000)	(0.368)	(0.010)	(0.328)
	(0.000)	(0.556)	(0.000)	(0.731)
	(0.017)	(0.630)	(0.000)	(0.904)
	(0.000)	(0.876)	(0.000)	(0.515)
	(0.000)	(0.676)	(0.000)	(0.906)
Hist_60D	**0.191**	**0.025**	**0.149**	**-0.009**
Hist_60M	**0.200**	**0.013**	**0.157**	**0.005**
Pass_DJ	**0.177**	**0.010**	**0.177**	**0.010**
Pass_CW	**0.185**	**-0.009**	**0.185**	**-0.009**
Pass_1/N	**0.190**	**0.009**	**0.190**	**0.009**

This table shows average realized out-of-sample volatilities ($\overline{\sigma}$) and excess returns ($\overline{R}$) of different portfolio strategies. Portfolios are rebalanced monthly and annualized values of volatilities and returns are reported. The out-of-sample period starts in January 1998 and ends in October 2009, i.e., the number of observations is 142. Columns 2 and 3 refer to the case when short sales are allowed, Columns 4 and 5 to the case with short-sales constraints. The numbers in brackets are p-values for tests of significant differences between the implied strategy (Implied) and the benchmark strategies: GMVP using historical estimates over the last 60 days (Hist_60D), GMVP using historical estimates over the last 60 months (Hist_60M), and passive investments using a price weighting (Pass_DJ), capital weighting (Pass_CW), and equal weighting (Pass_1/N), respectively.

Table 3: Out-of-Sample Volatilities and Excess Returns of the Portfolio Strategies in Periods of Crisis and Periods of no Crisis.

	Periods of Crisis		Periods of no Crisis	
	$\overline{\sigma}$	$\overline{R}$	$\overline{\sigma}$	$\overline{R}$
Implied	**0.167**	**-0.036**	**0.120**	**0.059**
	(0.003)	(0.262)	(0.747)	(0.885)
	(0.000)	(0.513)	(0.599)	(0.718)
	(0.000)	(0.335)	(0.015)	(0.187)
	(0.000)	(0.174)	(0.001)	(0.318)
	(0.000)	(0.355)	(0.008)	(0.123)
Hist_60D	**0.178**	**-0.078**	**0.121**	**0.055**
Hist_60M	**0.195**	**-0.065**	**0.122**	**0.069**
Pass_DJ	**0.231**	**-0.093**	**0.128**	**0.104**
Pass_CW	**0.243**	**-0.122**	**0.131**	**0.094**
Pass_1/N	**0.256**	**-0.104**	**0.129**	**0.112**

This table shows average realized out-of-sample volatilities ($\overline{\sigma}$) and excess returns ($\overline{R}$) of different portfolio strategies in periods of crisis and periods of no crisis. Crisis periods are March 2000 to April 2003 and June 2007 to October 2009 (67 observations). Non-crisis periods are January 1998 to February 2000 and May 2003 to May 2007 (75 observations). Portfolios are rebalanced monthly and annualized values of volatilities and returns are reported. The numbers in brackets are p-values for tests of significant differences between the implied strategy (Implied) and the benchmark strategies: GMVP using historical estimates over the last 60 days (Hist_60D), GMVP using historical estimates over the last 60 months (Hist_60M), and passive investments using a price weighting (Pass_DJ), capital weighting (Pass_CW), and equal weighting (Pass_1/N), respectively.

Table 4: Performance of GMVPs Based on Fully Implied and Partially Implied Estimates of the Covariance Matrix.

	Full Sample		Periods of Crisis		Periods of no Crisis	
	$\overline{\sigma}$	$\overline{R}$	$\overline{\sigma}$	$\overline{R}$	$\overline{\sigma}$	$\overline{R}$
	Implied Correlations - Implied Variances					
	0.143	**0.014**	**0.167**	**-0.036**	**0.120**	**0.059**
	(0.055)	(0.313)	(0.010)	(0.075)	(0.783)	(0.477)
Implied	(0.003)	(0.517)	(0.000)	(0.146)	(0.518)	(0.394)
	(0.000)	(0.971)	(0.000)	(0.833)	(0.004)	(0.949)
	(0.000)	(0.888)	(0.000)	(0.778)	(0.000)	(0.894)
	Historical Correlations - Implied Variances					
60D	**0.147**	**-0.008**	**0.176**	**-0.100**	**0.120**	**0.077**
60M	**0.150**	**-0.001**	**0.183**	**-0.092**	**0.119**	**0.082**
	Implied Correlations - Historical Variances					
60D	**0.150**	**0.013**	**0.177**	**-0.041**	**0.126**	**0.062**
60M	**0.164**	**0.010**	**0.192**	**-0.047**	**0.138**	**0.063**

This table shows average realized out-of-sample volatilities ($\overline{\sigma}$) and excess returns ($\overline{R}$) of portfolio strategies based on fully implied and partially implied estimates of the covariance matrix. The out-of-sample period starts in January 1998 and ends in October 2009, i.e., the number of observations is 142 for the full sample. Crisis periods are March 2000 to April 2003 and June 2007 to October 2009 (67 observations). Non-crisis periods are January 1998 to February 2000 and May 2003 to May 2007 (75 observations). Historical variances and correlations are estimated from daily returns of the preceding 60 days (60D) or from monthly returns of the preceding 60 months (60M). Portfolios are rebalanced monthly and annualized values of volatilities and returns are reported. The numbers in brackets are p-values for tests of significant differences between the fully implied strategy (implied correlations - implied variances) and the partially implied strategies (historical correlations - implied variances; implied correlations - historical variances).

Table 5: Out-of-Sample Volatilities and Excess Returns of the Portfolio Strategies for Different Implied Covariance Estimators.

	Full Sample		Periods of Crisis		Periods of no Crisis	
	$\overline{\sigma}$	$\overline{R}$	$\overline{\sigma}$	$\overline{R}$	$\overline{\sigma}$	$\overline{R}$
	0.143	**0.014**	**0.167**	**-0.036**	**0.120**	**0.059**
Imp_Var	(0.000)	(0.761)	(0.001)	(0.804)	(0.014)	(0.511)
	(0.107)	(0.275)	(0.155)	(0.904)	(0.338)	(0.177)
Imp_Skew	**0.150**	**0.019**	**0.176**	**-0.043**	**0.127**	**0.076**
Imp_Kurt	**0.143**	**0.016**	**0.167**	**-0.036**	**0.121**	**0.063**

This table shows average realized out-of-sample volatilities ($\overline{\sigma}$) and excess returns ($\overline{R}$) of portfolio strategies based on different implied covariance estimators. Annualized values of volatilities and returns are reported. The out-of-sample period starts in January 1998 and ends in October 2009, i.e., the number of observations is 142 for the full sample. Crisis periods are March 2000 to April 2003 and June 2007 to October 2009 (67 observations). Non-crisis periods are January 1998 to February 2000 and May 2003 to May 2007 (75 observations). The numbers in brackets are p-values for tests of significant differences between the Imp_Var strategy and the alternative implied strategies (Imp_Skew, Imp_Kurt).

Table 6: Out-of-Sample Volatilities and Excess Returns of the Portfolio Strategies for Quarterly Rebalancing.

	Full Sample		Periods of Crisis		Periods of no Crisis	
	$\overline{\sigma}$	$\overline{R}$	$\overline{\sigma}$	$\overline{R}$	$\overline{\sigma}$	$\overline{R}$
Imp_Var	**0.150**	**0.006**	**0.175**	**-0.045**	**0.126**	**0.052**
	(0.077)	(0.190)	(0.053)	(0.114)	(0.595)	(0.694)
	(0.014)	(0.693)	(0.000)	(0.620)	(0.972)	(0.899)
	(0.006)	(0.981)	(0.000)	(0.574)	(0.134)	(0.230)
	(0.007)	(0.631)	(0.000)	(0.317)	(0.036)	(0.362)
	(0.061)	(0.926)	(0.001)	(0.544)	(0.163)	(0.254)
Imp_Skew	**0.156**	**0.015**	**0.185**	**-0.049**	**0.129**	**0.074**
	(0.871)	(0.097)	(0.926)	(0.192)	(0.597)	(0.283)
	(0.094)	(0.374)	(0.005)	(0.635)	(0.243)	(0.378)
	(0.003)	(0.765)	(0.000)	(0.554)	(0.304)	(0.475)
	(0.001)	(0.433)	(0.000)	(0.312)	(0.057)	(0.774)
	(0.046)	(0.735)	(0.002)	(0.521)	(0.326)	(0.445)
Imp_Kurt	**0.150**	**0.007**	**0.176**	**-0.044**	**0.126**	**0.053**
	(0.084)	(0.168)	(0.059)	(0.093)	(0.604)	(0.673)
	(0.013)	(0.654)	(0.000)	(0.592)	(0.979)	(0.935)
	(0.006)	(0.958)	(0.000)	(0.559)	(0.138)	(0.236)
	(0.006)	(0.609)	(0.000)	(0.305)	(0.038)	(0.374)
	(0.057)	(0.907)	(0.001)	(0.531)	(0.165)	(0.261)
Hist_60D	**0.155**	**-0.019**	**0.185**	**-0.085**	**0.128**	**0.042**
Hist_60M	**0.161**	**-0.004**	**0.199**	**-0.069**	**0.126**	**0.055**
Pass_DJ	**0.182**	**0.004**	**0.235**	**-0.093**	**0.132**	**0.094**
Pass_CW	**0.189**	**-0.016**	**0.248**	**-0.125**	**0.135**	**0.084**
Pass_1/N	**0.193**	**0.001**	**0.258**	**-0.104**	**0.133**	**0.096**

This table shows average realized out-of-sample volatilities ($\overline{\sigma}$) and excess returns ($\overline{R}$) of different portfolio strategies for quarterly rebalancing. Annualized values of volatilities and returns are reported. The out-of-sample period starts in January 1998 and ends in October 2009, i.e., the number of observations is 142 for the full sample. Crisis periods are March 2000 to April 2003 and June 2007 to October 2009 (67 observations). Non-crisis periods are January 1998 to February 2000 and May 2003 to May 2007 (75 observations). The numbers in brackets are p-values for tests of significant differences between the respective implied strategy (Imp_Var, Imp_Skew, or Imp_Kurt) and the benchmark strategies: GMVP using historical estimates over the last 60 days (Hist_60D), GMVP using historical estimates over the last 60 months (Hist_60M), and passive investments using a price weighting (Pass_DJ), capital weighting (Pass_CW), and equal weighting (Pass_1/N), respectively. Quarterly returns are calculated for every month of the out-of-sample period leading to overlapping returns. Therefore, p-values are based on the heteroscedasticity and autocorrelation consistent estimator by Newey and West (1987) with two lags.

Table 7: Average Turnover of Different Strategies

	Monthly Rebalancing			Quarterly Rebalancing		
	Full Sample	Periods of Crisis	Periods of no Crisis	Full Sample	Periods of Crisis	Periods of no Crisis
Imp_Var	0.672	0.655	0.688	0.618	0.573	0.657
Imp_Skew	0.956	0.865	1.039	0.835	0.756	0.904
Imp_Kurt	0.709	0.671	0.743	0.623	0.580	0.661
Hist_60D	0.745	0.685	0.800	1.208	1.097	1.306
Hist_60M	0.187	0.211	0.165	0.333	0.354	0.315
Pass_DJ	0.016	0.022	0.010	0.050	0.071	0.031
Pass_CW	0.023	0.028	0.018	0.072	0.089	0.057
Pass_1/N	0.060	0.068	0.053	0.110	0.124	0.098

This table shows the average turnover of different portfolio strategies for different periods and both monthly and quarterly rebalancing. Average turnover is the average fraction of the portfolio that has to be traded (purchased and sold) at a rebalancing date. The full sample covers the period from January 1998 to October 2009. Crisis periods are March 2000 to April 2003 and June 2007 to October 2009. Non-crisis periods are January 1998 to February 2000 and May 2003 to May 2007.

Table 8: Out-of-Sample Excess Returns of the Portfolio Strategies Adjusted for Transactions Costs

	Monthly Rebalancing			**Quarterly Rebalancing**		
	Full Sample	Periods of Crisis	Periods of no Crisis	Full Sample	Periods of Crisis	Periods of no Crisis
Imp_Var	**0.006**	**-0.044**	**0.051**	**0.003**	**-0.047**	**0.050**
	(0.311)	(0.259)	(0.847)	(0.152)	(0.099)	(0.621)
	(0.905)	(0.591)	(0.557)	(0.728)	(0.633)	(0.849)
	(0.908)	(0.402)	(0.120)	(0.981)	(0.590)	(0.204)
	(0.666)	(0.214)	(0.220)	(0.666)	(0.329)	(0.328)
	(0.932)	(0.413)	(0.076)	(0.962)	(0.559)	(0.226)
Imp_Skew	**0.008**	**-0.053**	**0.064**	**0.012**	**-0.051**	**0.070**
	(0.309)	(0.440)	(0.505)	(0.083)	(0.177)	(0.257)
	(0.823)	(0.721)	(0.865)	(0.424)	(0.662)	(0.439)
	(0.957)	(0.473)	(0.172)	(0.830)	(0.577)	(0.400)
	(0.609)	(0.256)	(0.359)	(0.478)	(0.329)	(0.700)
	(0.973)	(0.466)	(0.078)	(0.792)	(0.543)	(0.373)
Imp_Kurt	**0.007**	**-0.044**	**0.054**	**0.004**	**-0.046**	**0.051**
	(0.278)	(0.261)	(0.757)	(0.134)	(0.080)	(0.603)
	(0.853)	(0.595)	(0.627)	(0.689)	(0.606)	(0.884)
	(0.949)	(0.407)	(0.140)	(0.997)	(0.575)	(0.209)
	(0.634)	(0.219)	(0.255)	(0.644)	(0.317)	(0.339)
	(0.966)	(0.418)	(0.089)	(0.942)	(0.546)	(0.233)
Hist_60D	**-0.018**	**-0.087**	**0.046**	**-0.024**	**-0.089**	**0.037**
Hist_60M	**0.003**	**-0.068**	**0.067**	**-0.006**	**-0.070**	**0.054**
Pass_DJ	**0.009**	**-0.093**	**0.104**	**0.004**	**-0.093**	**0.094**
Pass_CW	**-0.010**	**-0.122**	**0.094**	**-0.016**	**-0.125**	**0.084**
Pass_1/N	**0.009**	**-0.104**	**0.112**	**0.001**	**-0.104**	**0.096**

This table shows average excess returns adjusted for transaction costs ($\overline{R}$) of different portfolio strategies for different periods and both monthly and quarterly rebalancing. Percentage transactions costs are 10 bp. The full sample covers the period from January 1998 to October 2009. Crisis periods are March 2000 to April 2003 and June 2007 to October 2009. Non-crisis periods are January 1998 to February 2000 and May 2003 to May 2007. The numbers in brackets are p-values for tests of significant differences between the respective implied strategy (Imp_Var, Imp_Skew, or Imp_Kurt) and the benchmark strategies: GMVP using historical estimates over the last 60 days (Hist_60D), GMVP using historical estimates over the last 60 months (Hist_60M), and passive investments using a price weighting (Pass_DJ), capital weighting (Pass_CW), and equal weighting (Pass_1/N), respectively. Quarterly returns are calculated for every month of the out-of-sample period leading to overlapping returns. Therefore, the p-values for quarterly returns are based on the heteroscedasticity and autocorrelation consistent estimator by Newey and West (1987) with two lags.