Summary of the Research Project Proposal

'Improving Portfolio Selection Using Option-Implied Volatility and Skewness'

Victor DeMiguel
Yuliya Plyakhax
RamanUppal
GrigoryVilkov

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. Several approaches have been proposed in the literature for improving the performance of portfolios based on historical data.

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. The main contribution of our work is to demonstrate empirically how one can use option-implied information to improve the selection of a portfolio 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, where performance is measured in terms of the out-of-sample portfolio variance, portfolio Sharpe ratio, certainty equivalent, and turnover.

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.
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 or to an increase in the Sharpe ratio. This is surprising because there is a large literature that documents that implied volatility can predict stock-return volatility better than sample volatility (e.g., 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 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. 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.

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. 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). 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.