

Correlation Risk Premia for Multi-Asset Equity Options

Matthias R. Fengler*

Institut für Ökonometrie und Statistik, Wirtschaftswissenschaftliche Fakultät
Center for Applied Statistics and Economics (CASE)
Humboldt Universität zu Berlin,
Spandauer Straße 1, 10178 Berlin, Germany
fengler@wiwi.hu-berlin.de
TEL ++49 30 2093 5654
FAX ++49 30 2093 5649

Peter Schwendner[†]

Equity Trading & Products
Sal. Oppenheim jr. & Cie.
Königsberger Straße 29, 60487 Frankfurt am Main, Germany
peter.schwendner@oppenheim.de
TEL ++49 69 7134 5460
FAX ++49 69 7134 9 5460

March 10, 2003

*I gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft and the Sonderforschungsbereich 373 “Quantifikation und Simulation ökonomischer Prozesse”.

[†]This paper represents the author’s personal opinion and does not necessarily reflect the views of Sal. Oppenheim.

Correlation Risk Premia for Multi-Asset Equity Options

Abstract

The lack of a liquid market for implied correlations requires traders to estimate correlation matrices for pricing multi-asset equity options from historical data. To quantify the precision of these correlation estimates, we devise a block bootstrap procedure. The resulting bootstrap distributions are mapped on price distributions of three standard types of multi-asset options. ‘Minimal’ bid-ask spreads that reflect the risk from estimating the unknown correlations are quoted as quantiles of the price distributions. We discuss the influence of different market regimes and different payoff structures on the price distributions and on the the size of the resulting bid-ask spreads.

Keywords: Multi-Asset Options, Correlation Derivatives, Correlation Risk, Bid-Ask Spreads, Block Bootstrapping, Market Making, Equity Derivatives

JEL Classification: G12, G11

1 Introduction

In recent years, various types of multi-asset equity options emerged in the markets. They are either sold separately over-the-counter or as an ‘equity kicker’ of bond-like structures, where they usually offer a certain participation in equity performance or a large coupon conditionally on a defined performance of a basket of stocks. The latest kinds of these products are heavily path-dependent options with a very large lifetime (up to 10-15 years), and contain intrinsic barriers, or even some of their underlyings may be withdrawn at certain fixing dates, Overhaus [2002] and Quessette [2002].

Apart from the sophisticated payout structures, the inherent challenge of pricing and hedging multi-asset equity options is the illiquidity of implied correlations due to the lack of standardized multi-asset contracts. Consequently, equity correlation risk cannot be hedged as precisely as volatility risk. This is unlike to foreign exchange (FX) markets: Here, hedging correlation risk is possible, since volatilities and correlations of currency pairs are linked together via the exchange rate mechanism, as has been shown in a geometric interpretation by Wystup [2002]. Unfortunately, this does not hold for equity markets, as stocks are traded for cash and not in pairs like currencies. Therefore, at equity derivatives desks, traders monitor the statistics of their correlation exposure and try to avoid risk peaks in certain correlation pairs by managing the product flow via dynamic price margins. Spread positions between index options and baskets of single stock options may hedge the ‘average’ correlation risk within the index basket, but specific correlation risk in individual stock pairs remains.

Correlation risk stems particularly from two sources: First, correlations cannot directly be observed, but must be estimated. Second, as has recently been documented by Walter and Lopez [2000] and Neftci and Genberg [2002] for the FX markets, correlation may change

over time in the sequel of different market regimes. To cope with these facts, an obvious way to proceed would be to directly model time dependent correlations via multivariate generalizations of general autoregressive conditionally heteroscedastic (GARCH) processes such as the BEKK-model suggested by Baba et al. [1990]. However, multivariate GARCH modelling is a challenge, since the number of parameters even in very moderate dimensions is high. Furthermore, the log-likelihood function is often found to be flat, thus rendering parameter identification difficult, Fengler and Herwartz [2002]. Finally, since these types of processes typically converge very quickly to their unconditional means, forecasting performance of the correlations will be low.

In this study, we are not concerned about time-dependent modelling of correlations, but about the *quality* of the correlation estimates. We propose to ‘price’ correlation risk by constructing ‘minimal’ bid-ask-spreads derived from the asymptotic distribution of the correlation estimates. These spreads are minimal in the sense that they cover the risk from estimating the unknown correlations up to a certain degree of confidence. Thus, when quoting multi-asset option prices, one can translate the statistical confidence in the precision of the correlation estimate directly into the price margins. Since the distribution of correlation estimates is nonstandard, we devise a bootstrap methodology for its derivation. We apply this procedure to three prominent examples of multi-asset equity options such as a simple basket option, and a maximum and a minimum option, where the option writer either delivers the best or the worst performing asset, if exercised. We also discuss the specific risk exposure of these options and hedging.

This paper is organized as follows: in Section (2) the underlying asset price model is presented, and pricing and hedging of the options is discussed. Section (3) explains the bootstrap procedure and Section (4) discusses the results. Section (5) concludes.

2 Multi-Asset Options

2.1 Pricing model for European multi-asset equity options

We treat European multi-asset equity options in the standard Black/Scholes framework, i.e. we model n correlated Brownian motions W_i for the n underlying assets with spot prices $S_i(t)$ and constant parameters for the correlations ρ_{ij} , volatilities σ_i , dividend yields q_i and the risk-free rate r :

$$dS_i(t) = (r - q_i)S_i(t)dt + \sigma_i S_i(t)dW_i(t) \quad (1)$$

$$\rho_{ij}dt = dW_i(t)dW_j(t) \quad (2)$$

The literature suggests various pricing methods for European multi-asset options. For specific payout structures, analytical formulae are known, see Margrabe [1978], Stulz [1982], and Johnson [1987]. For arbitrary continuous payouts in lower dimensions ($n \leq 2$), numerical PDE solvers like Fast Fourier Transforms are advantageous, Engelmann and Schwendner [1998]. In higher dimensions, usually Monte Carlo integration or low-discrepancy series are used, Engelmann [2002]. For our numerical examples, we choose the standard Monte Carlo scheme.

We discuss three different option types with the following payout structures, Table (1): An at-the-money (ATM) call option on an equally weighted basket of n stocks, an option on the maximum performance of n assets and an option on the minimum performance of n assets. We define our payouts in percent of the underlying price to make them better comparable. Payouts at some terminal date T are defined relative to $S_i(0)$, i.e. the asset price at time $t = 0$ when the option is issued.

Insert Table (1) here

The risk management of multi-asset options at the trading desk usually involves monitoring the first and second order spot and volatility risks: the delta vector $\partial C/\partial S_i$, the gamma matrix $\partial^2 C/\partial S_i \partial S_j$, the vega vector $\partial C/\partial \sigma_i$ and the volga matrix $\partial^2 C/\partial \sigma_i \partial \sigma_j$. First-order correlation risk ('correlation vega') can be calculated as a triangular matrix $\partial C/\partial \rho_{ij}$ with $i < j$. The standard hedge instruments, single stocks and plain vanilla options on single stocks, hedge only the delta and vega risks and the diagonal elements of the gamma and volga matrix. The remaining risks, i.e. the nondiagonal elements of the gamma and volga matrix ('cross gamma' and 'cross volga') and the correlation vega, can only be hedged by involving other multi-asset options.

2.2 Drivers of correlation risk

To get an intuitive understanding of the influence of correlation on the price of our selected option types, we discuss two different drivers for the correlation risk:

- the influence of correlation on the volatility of the whole basket,
- the influence of correlation on the dispersion of individual assets in the basket against each other.

The impact of these drivers on the different option types can be summarized as follows, Table (2):

1. Basket option: The basket option is affected by the basket volatility, only. The dispersion of individual assets does not influence the option price, because the payout does

only depend on the sum of the asset prices at maturity. As higher correlations increase basket volatility and thus the option price, basket options are *long* in correlation.

2. Max option: The max option is affected by both drivers. Increasing dispersion of individual assets increases the chance of any asset to reach particularly high values at maturity. This effect grows with *declining* correlations. So for max options, the two drivers work in opposite directions. For the parameter spectrum in our application, i.e. from moderate to high correlation, we observe option prices to decrease with increasing correlation: hence, the dispersion effect outpaces the basket volatility effect and the holder of our max options is *short* in correlation. However, this is the initial exposure when the $S_i(0)$ are fixed. Situations are possible where the max option is both short and long in correlation depending on the specific levels of correlation and spot prices.
3. Min option: The min option is affected by both drivers as well, but both take the same direction. The dispersion effect increases option prices as correlations become higher, since the holder of the option is interested in a small dispersion of the single assets in the basket. This minimizes the probability that any asset reaches particularly low levels at maturity and maximizes the value of the min option contract. Thus, the net position of min options is long correlation. As both drivers reinforce each other, the correlation sensitivity is especially high for this option type.

Insert Table (2) here

2.3 Stability of vega hedges

The mechanisms that affect correlation sensitivity also have an impact on the quality and stability of the vega hedges. However, it is difficult to give a coherent picture, since for

all products the quality of vega hedges hinges in a complex manner on current volatility, correlation and spot price levels. Hence we only discuss the vega hedge shortly after fixing the $S_i(0)$, and not a hedge over the entire life-time of the option.

Buying a basket and max option, one is vega, volga and cross volgas long: Thus when implied volatility rises, additional profits accrue due to increased basket volatility. For the max option this effect is even stronger as the dispersion between the stocks also increases with increasing volatility.

The min option is more involved: One is vega, volga and cross volga long in implied volatilities of all underlyings, but vega and volga short in the one belonging to the underlying with highest implied volatility. For the former, the basket effect exceeds the dispersion effect. For the latter, due to the one-sided risk profile, one expects the underlying with highest implied volatility to be the worst performing. Thus increasing implied volatility in this asset reduces the value of the min option and the dispersion effect tends to outperform the basket effect, which works in the opposite direction. Since highest implied volatility rises relatively to the other implied volatilities, the effect is even aggravated: hence the volga short position. On the other hand, one is cross volga long with respect to the other underlyings, since increasing volatility reduces the dominance of the ‘highest implied volatility’ position.

However, it needs to be mentioned that this picture may change considerably as stocks move away from their initial values, and implied volatility or even correlation changes during the holding period of the option. E.g. in the min option it may happen that one becomes vega short in two implied volatilities. Also for the case of the max option correlation may generate ‘surprising effects’ on the vega hedges through volga and cross volga. As our numerical simulations show, although being of second order magnitude, through cross volga

effects, which are not necessarily smaller than the volga itself, the vega hedge in one asset can be significantly distorted even if another implied volatility moves. Thus, to ensure the quality of the vega hedge, in multi-asset options, volga and cross volga need to be monitored carefully.

3 Bootstrapping Correlations

The bootstrap is a Monte Carlo ‘resampling’ method for estimating the distribution of an estimator or a test statistic, when the distribution is known to exist, but is difficult to compute or explicit approximations are poor. This situation is particularly obvious in the case of the correlation coefficient ρ . Being bounded between minus and plus one, the estimator of the correlation coefficient has a nonstandard distribution, and can well be approximated using the bootstrap.

The bootstrap has been introduced by Efron [1979]. Here, we present the technique only. For the necessary theoretical background, we refer the reader to the literature, such as the monographs by Hall [1992] and Mammen [1992]. Härdle et al. [2002] give a lucid review on the most recent developments of the bootstrap theory, especially focussing on the case of dependent data. Bootstrapping the correlation coefficient for independently and identically distributed (*iid*) random variables seems to be introduced by Lunneborg [1985] and Rasmussen [1987], however we are unaware of work on bootstrapping the correlation coefficient in dependent data.

Among practitioners, who are concerned about portfolio allocation and optimization in the Markowitz sense, a ‘portfolio resampling’ technique is popular, Michaud [1998] and Scherer

[2002]. Portfolio resampling means sampling the mean and the covariance matrix for a relatively small number of draws. For these draws optimal portfolio weights are calculated. The average of the latter is then used to derive a ‘resampled frontier’. This procedure seems to produce more stable portfolio weights for the computation of the true efficient frontier, than only one optimization step, and is thus similar to our procedure. However, it seems that dependence in the data is not explicitly taken into account.

For presentation of our bootstrap procedure, suppose for a moment that asset returns be *iid* random variables. The case for dependent data is explained afterwards. We denote by $\mathcal{X} \stackrel{\text{def}}{=} \{X_t\}_{t=1}^T$ a sample of return vectors $X_t \in \mathbb{R}^n$ of the n underlying assets, on which the multi-asset options are based. The X_t belong to some common population distribution P_0 . Furthermore, since we wish to estimate the correlation ρ_{ij} between asset i and j , we assume that the first and the second moments exist. We are interested in estimating the distribution of the correlation estimator ρ_T , which will be denoted by $L_T(P_0) \stackrel{\text{def}}{=} \mathcal{L}\{\rho_T(X_1, \dots, X_T | P_0)\}$. Let \tilde{P}_T be a consistent estimate of P_0 . The key insight of the bootstrap procedure is that the bootstrap estimate of $L_T(P)$ simply is $L_T(\tilde{P}_T)$. An approximate confidence interval with significance level $(1 - \alpha)$ can be based on the $(1 - \alpha)$ -quantile of $L_T(\tilde{P}_T)$.

Since one can hardly compute $L_T(\tilde{P}_T)$ explicitly, it is approximated numerically. The method works as follows:

STEP 1. Generate a *resample* $\mathcal{X}^* = \{X_t^*\}_{t=1}^T$ from \mathcal{X} as an unordered collection of T elements drawn randomly and with replacement from \mathcal{X} , such that the probability of each X_t^* to be equal to any of the $X_{t'}$ is T^{-1} : $P(X_t^* = X_{t'} | \mathcal{X}) = T^{-1}$, $1 \leq t, t' \leq T$. This way the X_t^* are *iid* distributed, conditional on \mathcal{X} . Thus, the resample \mathcal{X}^* may contain repeats or may not contain some of the X_t at all.

STEP 2. Compute ρ_T^* from \mathcal{X}^* .

STEP 3. Repeat the steps one and two for a large number of times M .

STEP 4. Approximate $L_T(\tilde{P}_T)$ by the empirical distribution of the pseudo random variables $\rho_{T,1}^*$ to $\rho_{T,M}^*$.

The challenge in dependent data is that the bootstrap must be carried out in a way which preserves the dependence structure present in the data. One possible solution is the *block* bootstrap, Hall [1985] and Carlstein [1986]. Instead of drawing single values from \mathcal{X} , the idea is to draw entire blocks or subseries with replacement and to compute a pseudo time series from these blocks. More precisely, split the data \mathcal{X} into B blocks $\mathcal{Y}_t \stackrel{\text{def}}{=} \{X_t, \dots, X_{t+\ell-1}\}$ such that block length $\ell = T/B$ (we assume for the sake of exposition that $B\ell = T$ exactly). Blocks may be overlapping or non-overlapping, so either $\{X_1, \dots, X_\ell\}, \{X_2, \dots, X_{\ell+1}\}$ etc., or $\{X_1, \dots, X_\ell\}, \{X_{\ell+1}, \dots, X_{2\ell}\}$ and so forth. Under suitable assumptions on the underlying stochastic process generating X_t , the approximate bootstrap distribution $L_T(\tilde{P}_T)$ is again given by the empirical distribution of $\rho_{T,1}^*$ to $\rho_{T,M}^*$, Härdle et al. [2002]. Note however that overlapping block bootstrap ensues a bias since observations at both far ends of the original sample have less probability in entering the bootstrap samples than those in the middle, which needs to be corrected for. This is why we prefer the non-overlapping bootstrap.

Choice of the block length ℓ is an intricate question: The optimal ℓ intimately depends on the underlying data generating process and must increase with sample size T in order to achieve asymptotically correct coverage probabilities of the confidence intervals. Depending on the particular assumptions on the stochastic process and on the object to be estimated, a number of asymptotically optimal block lengths have been derived. Typically, they range between $\ell \propto T^{1/5}$ to $\ell \propto T^{1/3}$ (Härdle et al. [2002] and references therein), which – with a sample size

of around 250 days in our application – amounts to 3 to 6 days. We experimented between these values and did not find any strong deviations among the bootstrapped distributions. Since our returns are not, or at most mildly serially correlated, we believe a short block length to be sufficient to capture the local dependence of the data, and decided for $\ell = 3$.

4 Empirical Results

4.1 General Set-up

In our empirical analysis we took times series of six major stocks in the German DAX index, adjusted for stock splits: Allianz (ALV), BMW, Commerzbank (CBK), Deutsche Bank (DBK), DaimlerChrysler (DCX), Deutsche Telekom (DTE). The data was kindly provided by Sal. Oppenheim jr. & Cie.

We choose two time periods with entirely different market regimes to investigate the pricing results. As an example of a bullish trending market, where correlations are relatively low, we take the year 1999, whereas 2002 is an example of a bearish market, where correlation tends to be higher. In each case we price multi-asset options with a maturity of one year. Returns are calculated as log differences. We compute correlation from lagged returns, where the time window of past returns corresponds to the life time of the option, i.e. around 250. This is in line with market conventions. Statistically speaking, this convention implies that the market assumes stationarity of the model over the life time of the option. As spelled out in Section (3) we bootstrap in blocks of $\ell = 3$ days $M = 20\,000$ times. Proceeding this way, we generate in the case of the three asset model ($n = 3$) a trivariate distribution of correlations,

in the case of the six-asset model ($n = 6$) an $\frac{n(n-1)}{2} = 15$ -dimensional distribution.

Correlation distributions are mapped via the multi-asset pricers into an option price distribution. The first basket contains the three stocks CBK, DBK and DTE, the second basket contains all six stocks from Table (3). In Table (3) we also display the implied volatility and dividend yields under which we price. Since we like to focus on correlation risk only, they were not altered in the 1999 and 2002 scenarios. Interest rate is fixed at $r = 5\%$ throughout. All options are assumed to be of European type and priced via the discounted mean of the payout function after 50 000 simulations. The maturity is $T = 1$ year. Bid and ask spreads are computed from the empirical price distribution as the 90% confidence around the mean.

The resulting price distributions are exploited for the derivation of bid-ask prices. In our application we stipulate a confidence level of $\alpha = 10\%$ percent. A 90%-confidence interval around the (fair) mean value is hence given by computing the $\alpha/2$ and $1 - \alpha/2$ quantiles of the price distribution. Note however, although the statistical coverage of this interval is 90%, from a trading point view it is 95%, since the trader faces only a one-tailed risk for his bid and ask price.

Insert Table (3) here

4.2 Discussion

Tables (4) and (5) summarize general statistics of the bootstrapped distributions. For the sake of clarity, we focus in our discussion on the three asset case, only. (Statistics on the six asset case may be obtained from the authors upon request). Large differences between the two different market scenarios are visible: During 1999 mean correlation between the

bank stocks DBK/CBK and DTE is low, around 0.25–0.27, and within the two bank stocks moderately, 0.54. However, as is expected, in 2002 correlations sharply rise: All correlations range from 0.56 (DTE/CBK) up to 0.74 (DBK/CBK). At the same time standard deviation of correlations drops almost by one half. Thus precision of correlation estimates becomes more accurate in the bad market, whereas in the stable upward trending market the precision is lower. For instance, from the distribution of the correlation estimates one can infer that the hypothesis of a correlation for DBK/DTE of only $\rho = 0.15$, i.e. of 10 basis points lower than the mean, is well within the 95% confidence, whereas in the 2002 market for a mean of 0.67, the hypothesis of $\rho = 0.57$ is safely rejected at the 5% level. Thus, hedges taking a correlation estimate of $\rho_T = 0.25$ for granted, can be misled, even if the data generating process is stationary.

Figures (1) and (2) display bivariate contour plots of the (marginal) bivariate correlation distributions. From the general position of the ellipses it is visible that the correlation structure between the estimates varies among the different assets: Whereas for DBK/DTE and DBK/CBK cross correlation between estimates is low, as is seen from the concentric circles (this holds also for DBK/CBK and CBK/DTE), in the case of DBK/DTE and DTE/CBK correlation estimates are positively correlated, see also Table (5). During the bearish market in 2002 the general orientation of the ellipses from ‘correlation of correlation point of view’ is almost unchanged, but due to the dropped standard deviation, it is more aggravated (note that the width between contour lines is 5 in both plots). This is also economically plausible: One would expect correlation pairs between two different sectors (banking sector versus a cross sector) to be not or only little correlated, whereas pairs of common cross sector correlations should be correlated, if common, sector specific shocks exist that hit the market. Interestingly, the correlation of correlation remains relatively stable across the different mar-

ket regimes. This evidence corresponds with the findings of Neftci and Genberg [2002] who investigate the correlation of implied volatilities between each other.

Stable correlations of correlations have also important implications for hedging: In this case, it may be sufficient to hedge the ‘average’ correlation risk, e.g. to hedge a short correlation position by selling single stock vega and buying index vega. If correlations rise, we gain on the index vega position which in turn offsets the losses in the short correlation position.

Insert Table (4) here

Insert Table (5) here

Insert Figure (1) about here

Insert Figure (2) about here

In Tables (6) and (7), we present descriptive statistics of the price distributions obtained from mapping the correlation distribution via the option pricers. Since in the bootstrap the mean of the price distribution is preserved, the mean corresponds to the fair price directly obtained from the correlation estimate in the time period under consideration. In order to remove level effects in our dispersion measure, we also report the coefficient of variation of the distributions, which is defined as standard deviation divided by mean. In our discussions, we do not discuss the absolute difference between bid and ask. Rather we focus on the bid-ask spreads relative to the mean, since one typically is interested in the spread relative to the fair price of the option. Figures (3) and (4) also display the discussed price distributions.

Basket option

Comparing the three versus the six asset cases, regardless of the correlation scenarios, one observes that the basket option becomes cheaper, the more assets are included into the payout. Clearly, this is a risk diversification effect: The more assets are included in the

basket, the less its volatility, the lower its price. On the other hand, increasing the number of assets should introduce additional risk of the unknown correlations. Standard deviation and the bid-ask spreads increase only moderately. Coefficient of variation also indicates a very slight increase in dispersion when one moves from three to six assets. Thus diversification seems to almost compensate the risk from correlations. As is to be expected, with increasing correlations between different market scenarios the diversification effect deteriorates, and prices increase. However, since estimates of correlations become more precise minimal bid-ask spreads can be reduced.

Maximum option

Max options are by far the most expensive ones accounting for up to 60% of the spot. Since the payout is directed only at the performance of the best asset, prices strongly increase, the more assets we include in the basket. So do bid-ask spreads. Level adjusted coefficients of variation indicate an increase of dispersions also. As is explained in Section (2), for max options two competing drivers influence option prices: a dominating dispersion and a basket volatility effect. When correlation is low it is more probable that some stock is performing well when others do not. Thus a max option holder is short in correlation. Observe that this effect can be very strong: option prices are 15% lower in the three asset case and 20% in the six asset case. Bid-ask spreads (and coefficients of variation also) stay approximately of same size in both scenarios. Thus the effect of higher estimation precision of the correlation, which diminishes the spread, is dominated by the fact that the correlation sensitivity increases as correlations rise.

Minimum option

Since the buyer of the option receives always the worst performing asset, the price of min option falls the more assets are included. Due to the dispersion effect bid-ask spreads increase

with a growing number of assets. This is particularly evident in the 6 asset minimum option for 1999. Increasing correlation reduces both dispersion and volatility of the entire basket resulting in increasing prices for 2002 compared with 1999. At the same time higher precision of correlations reduces bid-ask spreads relative to the low correlation scenario. This effect is stronger than the increased correlation sensitivity in high correlation situations.

Insert Table (6) here

Insert Table (7) here

Insert Figure (3) about here

Insert Figure (4) about here

Although we do not compute formal tests for normality, our figures for skewness and kurtosis seem to confirm that using a normal variate for the computation of minimal risk premia for multi-asset basket options, such as the well known ‘two-times-sigma’ rule will not be too bad an approximation to the quantiles of the price distribution. Under the normal distribution, coefficient of variation and spread per mean are multiples of the same quantity and thus always move in the same direction. In case for the min options the normal distribution may not be good, because non-negligible skewness is emerging, the closer the price approaches zero. However, when prices are close to zero, using a normal approximation may not be advisable anyway.

5 Conclusions

Pricing and hedging multi-asset equity options requires estimating correlations. Even when time series of stock prices are assumed to be stationary, correlation estimates from historical time series involve estimation errors. We quantify this error by approximating the asymptotic distribution of the correlation estimate via block bootstrapping. The bootstrapped correlation distributions are mapped on prices of three standard types of multi-asset options (basket option, minimum and maximum option). Bid-ask spreads of the prices are computed as statistical quantiles of the resulting price distributions. The resulting bid-ask spreads can be considered as ‘minimal’ spreads since they accommodate the risk from estimating unknown correlations only.

We discuss hedging the correlation products and present numerical results for two different market regimes (1999 and 2002), two different numbers of stocks (three and six) and three payout types. The shown differences in the bid-ask spreads can be attributed to two drivers of correlation risk acting differently for the three payout types: basket volatility and dispersion. These drivers influence both the absolute level of the prices and bid-ask spreads of the three individual options in a significantly different manner.

Also, the two market regimes affect the bid-ask spreads: typically, in a bearish market regime correlations and statistical confidence in the estimates are high, whereas in the bullish market correlations may be moderate, while having low statistical confidence at the same time. Thus bid-ask spreads are influenced in two ways: first via the correlation sensitivity of the option and second via the confidence that can be attributed to correlation estimates in the current regime.

We point out that traders quoting bid-ask spreads of multi-asset equity options need the same thorough correlation analysis as do investors considering these products for their portfolios.

References

- Y. Baba, R. F. Engle, D. F. Kraft, and K. F. Kroner. “Multivariate simultaneous generalized arch models.” Technical report, Department of Economics, University of California, San Diego, 1990.
- E. Carlstein. “The use of subseries methods for estimating the variance of a general statistic from stationary time series.” *Annals of Statistics*, 14 (1986), pp. 1171–1179.
- B. Efron. “Bootstrap methods: Another look at the jackknife.” *Annals of Statistics*, 7 (1979), pp. 1–26.
- B. Engelmann. “Quasi-monte carlo techniques for the valuation of contingent claims on several assets.” In J. Hakala and U. Wystup, editors, *Foreign Exchange Risk*, pages 209–225. Risk Books, 2002.
- B. Engelmann and P. Schwendner. “The pricing of multi-asset options using a fourier grid method.” *Journal of Computational Finance*, 1(4) (1998), pp. 63–72.
- M. Fengler and H. Herwartz. “Multivariate volatility models.” In W. Härdle, T. Kleinow, and G. Stahl, editors, *Applied Quantitative Finance*, chapter 10. Springer, Heidelberg, 2002.
- P. Hall. “Resampling a coverage process.” *Stochastic Process Applications*, 19 (1985), pp. 259–269.

- P. Hall. “The Bootstrap and the Edgeworth Expansion.” Springer Verlag, New York, Heidelberg, 1992.
- W. Härdle, J. Horowitz, and J.-P. Kreiss. “Bootstrap methods for time series.” Technical report, Humboldt-University, Berlin, 2002.
- H. Johnson. “Options on the maximum or the minimum of several assets.” *Journal of Financial and Quantitative Analysis*, 22 (1987), pp. 277–283.
- C. E. Lunneborg. “Estimating the correlation coefficient: The bootstrap approach.” *Psychological Bulletin*, 98 (1985), pp. 209–215.
- E. Mammen. “When Does Bootstrap Work? Asymptotic Results and Simulations.” Number 77 in *Lecture Notes in Statistics*. Springer Verlag, New York, Heidelberg, 1992.
- W. Margrabe. “The value of an option to exchange one asset for another.” *Journal of Finance*, 33 (1978), pp. 177–186.
- R. Michaud. “Efficient Asset Management.” Harvard Business School Press, Boston, MA, 1998.
- S. N. Neftci and H. Genberg. “Excessive variation of risk factor correlations and volatilities.” *Journal of Futures Markets*, 22 (2002), pp. 1119–1146.
- M. Overhaus. “Himalaya options.” *Risk*, 15(12) (2002).
- R. Quessette. “New products, new risks.” *Risk*, 15(12) (2002).
- J. L. Rasmussen. “Estimating the correlation coefficient: Bootstrap and parametric approaches.” *Psychological Bulletin*, 101 (1987), pp. 136–215.

- B. Scherer. "Portfolio resampling: Review and critique." *Financial Analysts Journal*, 58(6) (2002), pp. 98–109.
- R. Stulz. "Options on the minimum or the maximum of two risky assets." *Journal of Financial Economics*, 10 (1982), pp. 161–181.
- C. A. Walter and J. A. Lopez. "Is implied correlation worth calculating? Evidence from foreign exchange options." *Journal of Derivatives*, 22 (2000), pp. 65–81.
- U. Wystup. "How the greeks would have hedged correlation risk of foreign exchange option." In J. Hakala and U. Wystup, editors, *Foreign Exchange Risk*, pages 143–146. Risk Books, 2002.

Table 1: Payout structures of the multi-asset options

Option Type		Payout
ATM call on a basket of n assets	basket option	$\max\{\frac{1}{n} \sum_{i=1}^n \frac{S_i(T)}{S_i(0)} - 1, 0\}$
ATM call on the maximum of n assets	max option	$\max\{(\max_{i=1}^n \frac{S_i(T)}{S_i(0)}) - 1, 0\}$
ATM call on the minimum of n assets	min option	$\max\{(\min_{i=1}^n \frac{S_i(T)}{S_i(0)}) - 1, 0\}$

$S_i(T)$ denotes the (uncertain) spot value of asset i at some terminal date T , $S_i(0)$ denotes its value at time $t = 0$ when the payouts for the at-the-money options (ATM) are fixed.

Table 2: Influence of correlation on option prices

	Correlation driver		net effect
	basket volatility	asset dispersion	
Basket option	+	0	+
Max option	+	-	-
Min option	+	+	++

Influence of correlation on the option prices via the volatility and dispersion driver. ‘+’ indicates a positive influence of the respective correlation driver when correlation levels rise, ‘-’ a negative influence. Reported net effects as observed from our numerical examples.

Table 3: Underlying assets: assumed properties

Sector	stock	Reuters code	implied vol.	div. yield
insurance	Allianz	ALV	52	1.5
automobiles	BMW	BMW	42	1.7
banking	Commerzbank	CBK	60	4.9
banking	Deutsche Bank	DBK	48	2.8
automobiles	DaimlerChrysler	DCX	47	3.3
telecommunication	Deutsche Telekom	DTE	51	2.9

Implied volatilities and dividend yield quoted in percent. Numbers are fixed in all calculations and supposed to reflect average market conditions, not the specific 1999 and 2002 scenarios, since the focus is on correlation risk, only.

Table 4: Summary statistics of bootstrapped correlations in the three asset case

year	correlation pair	min.	max.	mean	std.
1999	DBK/DTE	0.01	0.49	0.25	0.06
	DBK/CBK	0.33	0.67	0.53	0.04
	DTE/CBK	0.03	0.45	0.27	0.06
2002	DBK/DTE	0.49	0.79	0.67	0.03
	DBK/CBK	0.61	0.84	0.74	0.03
	DTE/CBK	0.40	0.74	0.56	0.04

Reported variables: minimum (min.), maximum (max.), mean and standard deviation (std.). Bootstrap distributions obtained after 20 000 block draws from the original asset returns sample, block length $\ell = 3$ days.

Table 5: Correlation of correlations in bootstrap distributions

year	correlation pair	DBK/DTE	DBK/CBK	DTE/CBK
1999	DBK/DTE	1	0.01	0.54
	DBK/CBK		1	0.17
	DTE/CBK			1
2002	DBK/DTE	1	0.23	0.64
	DBK/CBK		1	0.37
	DTE/CBK			1

Correlations between correlation pairs computed from the bootstrap distributions as described in Table (4).

Table 6: Descriptive statistics of option price distributions, 3 asset case

Option type	year	mean	std.	coeff. var.	skew.	kurt.	bid	ask	spread/mean
Basket option	1999	16.03	0.31	0.019	-0.262	3.159	15.51	16.51	0.06
	2002	18.31	0.20	0.011	-0.111	3.060	17.99	18.60	0.04
Max option	1999	44.35	0.68	0.015	0.121	3.073	43.26	45.49	0.05
	2002	37.70	0.71	0.019	-0.008	3.032	36.52	38.87	0.06
Min option	1999	3.43	0.33	0.097	-0.002	3.034	2.89	3.98	0.32
	2002	7.14	0.45	0.063	0.092	3.032	6.42	7.91	0.21

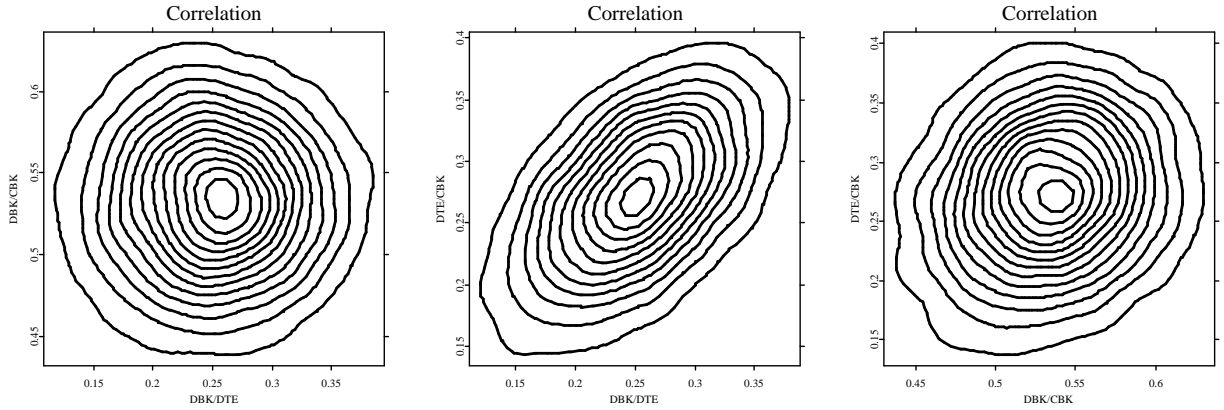
Reported variables are fair price (mean), standard deviation (std.), coefficient of variation (coeff. var.), skewness (skew.), kurtosis (kurt.), bid-ask spread (5% and 95% quantiles), spread per fair price (spread/mean). DBK, DTE and CBK are included in the basket. Prices obtained by Monte Carlo simulation.

Table 7: Descriptive statistics of option price distributions, 6 asset case

Option type	year	mean	std.	coeff. var.	skew.	kurt.	bid	ask	spread/mean
Basket option	1999	14.35	0.36	0.025	-0.144	3.001	13.75	14.92	0.08
	2002	17.23	0.20	0.012	-0.129	2.989	16.89	17.55	0.04
Max option	1999	60.68	1.38	0.022	0.028	2.969	58.42	62.96	0.07
	2002	47.67	1.09	0.023	0.022	2.977	45.89	49.52	0.08
Min option	1999	1.07	0.19	0.175	0.298	3.081	0.79	1.40	0.57
	2002	3.90	0.34	0.086	0.124	3.005	3.35	4.47	0.29


Reported variables are fair price (mean), standard deviation (std.), coefficient of variation (coeff. var.), skewness (skew.), kurtosis (kurt.), bid-ask spread (5% and 95% quantiles), spread per fair price (spread/mean). All assets available from Table (3) included in the basket. Prices obtained by Monte Carlo simulation.

Figure 1: Contour plots, 1999 data



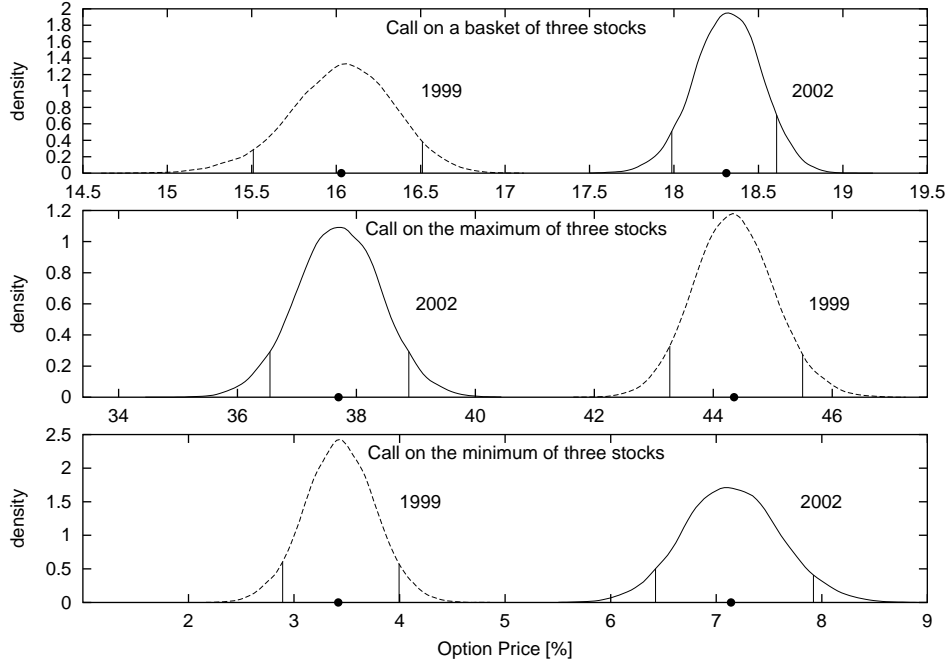
Bivariate contour plots of the (marginal) bootstrap distributions of estimated correlations in the three asset case. Obtained via a kernel smoothing procedure. Initial contour line starts at 5, step width between contour lines is also 5.

Figure 2: Contour plots, 2002 data



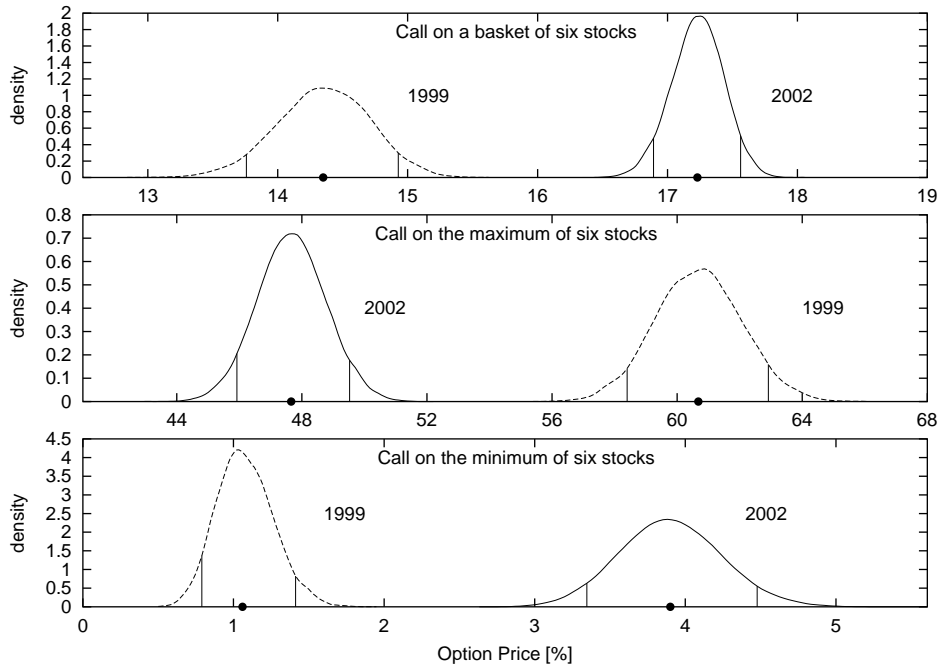
Bivariate contour plots of the (marginal) bootstrap distributions of estimated correlations in the three asset case. Obtained via a kernel smoothing procedure. Initial contour line starts at 5, step width between contour lines is also 5.

Figure 3: Price Densities and Bid-Ask Spreads, three asset case



Price densities obtained from a kernel smoothing procedure. Fair price denoted by a black bullet, bid-ask spreads, quoted as the 5% and the 95% interval, indicated by vertical lines. DBK, DTE and CBK are included in the basket. Prices obtained by Monte Carlo simulation.

Figure 4: Price Densities and Bid-Ask Spreads, six asset case



Price densities obtained from a kernel smoothing procedure. Fair price denoted by a black bullet, bid-ask spreads, quoted as the 5% and the 95% interval, indicated by vertical lines. All assets available from Table (3) included in the basket. Prices obtained by Monte Carlo simulation.