Quoting Multiasset Equity Options in the Presence of Errors from Estimating Correlations

MATTHIAS R. FENGLER AND PETER SCHWENDNER

MATTHIAS R. FENGLER is a quantitative analyst at Sal. Oppenheim jr. & Cie. in Frankfurt am Main, Germany.
matthias.fengler@oppenheim.de

PETER SCHWENDNER is a senior vice president at Sal. Oppenheim jr. & Cie.
peter.schwendner@oppenheim.de

Because there is no liquid market for implied correlations, traders must estimate correlation matrices for pricing multiasset 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 multiasset options. Minimum 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 size of the resulting bid-ask spreads.

Various types of multiasset equity options have emerged in the markets. They are either sold separately over-the-counter or issued as "equity kickers" in bond-like structures. Usually, they offer a certain participation in the equity performance or a large coupon conditional on some defined performance of a basket of stocks. The newest kinds of these products are heavily path-dependent long-dated options (up to 10-15 years) that include intrinsic barriers or from which one or more of their underlyings can be withdrawn at certain fixing dates (see Overhaus [2002] and Quessette [2002]).

The inherent challenge in pricing and hedging multiasset equity options, besides their sophisticated payout structures, lies in the illiquidity of implied correlations because there are no standardized multiasset contracts.[1] Consequently, equity correlation risk cannot be hedged as precisely as volatility risk. This is different from foreign exchange (FX) markets, where hedging correlation risk is possible, since volatilities and correlations of currency pairs are linked via the exchange rate mechanism (see Wystup [2002]). This does not hold for equity markets, as stocks are traded for cash and not in pairs like currencies.

At equity derivatives desks, therefore, 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 from two sources. First, correlations cannot be observed directly, but must be estimated. Second, as has been documented by Walter and Lopez [2000] and Neftci and Genberg [2002] for the FX markets, correlations may change over time in different market regimes.

To cope with these facts, an obvious way to proceed would be to model time-dependent correlations directly via multivariate general autoregressive conditionally heteroscedastic (GARCH) processes such as the BEKK model suggested by Baba et al. [1990]. Multivariate GARCH modeling is a challenge, however, as the number of parameters even in very moderate dimensions is high. Furthermore, the

EXHIBIT 1
Payout Structures of Multiasset 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.

log-likelihood function is often found to be flat, which renders parameter identification difficult; see Fengler and Herwartz [2002]. Finally, since these types of processes typically converge very quickly to their unconditional means, the forecasting usefulness of the correlations will be limited.

In this research, we are concerned not about modeling of time-dependent correlations, but about the *quality* of the correlation estimates. We propose to price correlation risk by constructing minimum bid-ask-spreads derived from the asymptotic distributions of the correlation estimates. These spreads are minimum in the sense that they cover the risk from estimating the unknown correlations up to a certain degree of confidence. That is, in quoting multiasset option prices, one can translate the statistical confidence in the precision of the correlation estimates directly into 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 multiasset equity options such as a simple basket option and maximum and minimum options, where the option writer delivers either the best- or the worst-performing asset if the option is exercised. We also discuss hedging and the specific risk exposure of these options.

I. MULTIASSET EQUITY OPTIONS

We first describe the underlying asset pricing model, and then discuss the pricing and hedging of options.

Pricing Model for European Multiasset Equity Options

We treat European multiasset equity options in the standard Black-Scholes framework. A stochastic volatility or a local volatility model could also be used (see Avellaneda et al. [2002]). We model n correlated Brownian motions W_i for the n underlying assets. Denote the spot price of asset i by $S_i(t)$; the risk-free rate by r; the dividend yield by q_i; asset price volatility by σ_i; and the constant parameter for the correlation between asset i and j by ρ_{ij}:

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

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

for $i, j = 1, \ldots, n$ and $i \neq j$.

The literature suggests various pricing methods for European multiasset options. For specific payout structures, analytical formulas are known; see Margrabe [1978], Stulz [1982], and Johnson [1987]. For arbitrary continuous payouts in lower dimensions ($n \leq 2$), numerical partial differential equation solvers like fast Fourier transforms are advantageous; see Engelmann and Schwendner [1998]. In higher dimensions, usually Monte Carlo integration or low-discrepancy series is used; see Engelmann [2002]. We choose the standard Monte Carlo scheme for our numerical examples, as performance is not an issue in our analysis.

We discuss three different option types with the payout structures as described in Exhibit 1: an at-the-money (ATM) call option on an equally weighted basket of n stocks (a *basket option*); an option on the maximum performance of n assets (*max option*); and an option on the minimum performance of n assets (*min option*). We define our payouts as a percentage of the underlying price to make them 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.

The risk management of multiasset 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 $\partial C/\partial \sigma_i$; the volga matrix $\partial^2 C/\partial \sigma_i \partial \sigma_j$; and the vanna matrix $\partial^2 C/\partial \sigma_i \partial S_j$. First-order correlation risk (correlation vega) can be calculated as a triangular matrix $\partial C/\partial \rho_{ij}$ with $i < j$.[2]

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 non-diagonal elements of the gamma and volga matrix (cross-gamma and cross-volga) and the correlation vega, can be hedged only by other multiasset options. Therefore, at equity derivatives desks, traders monitor the statistics of their correlation exposure. If the relationship among correlations proves to be sufficiently stable, one may take spread positions between options on some stock index, which itself is a basket, and a basket of individual stock options.

Such a position can hedge the average correlation risk within the index basket, although the specific correlation risk in individual stock pairs remains. Alternatively, one can control the risk in certain correlation pairs by managing the product flow via dynamic price margins (this issue is beyond the scope of this research).

The Monte Carlo pricing scheme requires a covariance matrix as input. For the diagonal elements of the covariance matrix, we use the (squared) implied volatilities of ATM plain vanilla single stock options that are currently traded; only if market-implied volatilities are used will the vega hedge be successful. Historical volatilities are not of importance in this context. We would like to do the same with the off-diagonal elements of the covariance or the correlation matrix, but there is no standardized market for implied correlations outside the foreign exchange market.

In the FX market, the exchange rate mechanism imposes restrictions on the covariance matrix that allow exact identification of implied correlations from observed implied volatilities of plain vanilla currency options, as in Walter and Lopez [2000]. Wystup [2002] shows via a geometric interpretation how this mechanism works.

Implied correlations need to be proxied by estimates from historical return data. When the Black-Scholes model is used for pricing multiasset equity options, it is common practice to assume that correlations are stationary and to estimate them as constants over a price history that corresponds to the lifetime of the option. Inevitably, estimation involves some estimation error, and we are concerned with quoting multiasset equity options in a way that takes this estimation error into account. In this sense, our bid-ask spreads will be the minimum.

Naturally, if correlations are time-dependent, the bid-ask spreads may underestimate this error. From a practical point of view, this effect can be remedied by hedging average correlation risk via spread positions in an index and individual stock options.

EXHIBIT 2
Influence of Correlation on Option Prices

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

+ Indicates a positive influence of the correlation driver when correlation levels rise; – indicates a negative influence. Reported net effects as observed from our numerical examples.

Drivers of Correlation Risk

To provide for an intuitive understanding of the influence of correlation on the prices 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, as laid out in Exhibit 2:

1. The basket option is affected by the basket volatility only. The dispersion among the individual assets does not influence the option price, because the payout depends only on the sum of the asset prices at maturity. As rising correlations increase the basket volatility and thereby the option price, basket options are *long* in correlation.
2. The max option is affected by both drivers. An increasing dispersion of individual assets increases the odds that any asset may 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 correlations, we observe that option prices decline with increasing correlation; hence, the dispersion effect outweighs the basket volatility effect, and the holder of the max options is *short* in correlation. This is the initial exposure when the $S_i(0)$ are fixed. In some cases, the max option is both short and long in correlation, depending on the specific levels of correlation and the spot prices.

3. The min option is affected by both drivers as well, but here both influences are in the same direction. The dispersion effect raises option prices as correlations become greater, since the holder of the option is interested in low dispersion of the assets in the basket. This minimizes the probability that any asset will reach particularly low levels at maturity, and maximizes the value of the min option contract. Thus, the net position of a min option is *long* in correlation. As both drivers reinforce each other, the correlation sensitivity is especially high for this option type.

Stability of Vega Hedges

The mechanisms that affect correlation sensitivity also have an impact on the quality and stability of the vega hedges. It is difficult to describe a coherent picture here, however, because for all products the quality of the vega hedge hinges in a complex manner on current volatility, correlation, and spot price levels. Hence we discuss the vega hedge only shortly after fixing the $S_i(0)$, and not the hedge over the entire lifetime of the option.

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

The min option is more involved. The option holder is vega, volga, and cross-volga long in the volatilities of all underlyings, but vega and volga short in the one for the underlying with the greatest implied volatility. For the former, the basket effect exceeds the dispersion effect. For the latter, the dispersion effect dominates. The reason is that the underlying with the greatest implied volatility can be expected to wander the farthest from its initial value; consequently, it is the most likely one to be delivered at maturity.

To put this more generally, because of the one-sided risk profile of the min option, it is the most dangerous asset from the holder's point of view. That is, an increase in implied volatility of this asset reduces the value of the min option, and the dispersion effect tends to override the basket effect, which works in the opposite direction. Since the greatest implied volatility also rises relative to the other implied volatilities, the effect is exacerbated, hence the volga short position. Yet one is cross-volga long with respect to the other underlyings, since increasing volatility reduces the dominance of the greatest implied volatility position.

Note that this picture may change considerably when 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 one might become vega short in two implied volatilities. In the max option, correlation may have a severe impact on the vega hedge through volga and cross-volga, although these are second-order effects. For instance, the vega hedge for asset i can be distorted through a cross-volga effect, when the implied volatility of asset j moves. Thus, to ensure the quality of the vega hedge for multiasset options, volga and cross-volga have to be monitored very carefully.

II. BOOTSTRAPPING CORRELATIONS

The (implied) correlations among single assets that are necessary for pricing the multiasset options are not directly observable. Therefore, they are approximated by consistent estimates from historical data. In our multidimensional framework, we estimate the entries in the correlation matrix $\hat{\rho}_T$ via:

$$\hat{\rho}_{ij} = \frac{\sum_{t=1}^{T}(x_{i,t} - \bar{x}_i)(x_{j,t} - \bar{x}_j)}{\sqrt{\sum_{t=1}^{T}(x_{i,t} - \bar{x}_i)^2 \sum_{t=1}^{T}(x_{j,t} - \bar{x}_j)^2}} \tag{3}$$

for $i, j = 1, \ldots, n$ and $i \neq j$. Here, $x_{i,t}$ and $x_{j,t}$ denote the daily return on assets i and j at date t. Returns are computed by log-differences from the asset price history of T days, which corresponds to the lifetime T of the option. The sample means of the returns are denoted by $\bar{x}$.

The value at risk literature often assumes that the mean of the returns is zero; see, for instance, Hendricks [1996] or Alexander and Leigh [1997]. The argument is that for financial time series the mean is close to zero and prone to estimation error. Despite this well-founded practice, we choose to include the estimated mean in our computations of the correlation matrix. Our main concern is consistency of the correlation estimator.

If we are sure that the population mean is indeed zero, then imposing the constraint yields a consistent and also more efficient estimator. Yet, because we sample from the physical distribution and not from the risk-neutral one, we cannot be sure that the mean is zero. The best we can say by economic reasoning is that in equilibrium it is non-negative for normal assets. Thus, we might bootstrap the distribution of an inconsistent estimator, which one wants to avoid.

We nevertheless test the bootstrap procedure using the constrained estimator. The differences both in terms of the correlation coefficients and the resulting prices are negligible. We attribute this finding to an averaging effect that ensues from resampling and that remedies the estimation error problem of the mean. We take this as additional support for our preference for using the unconstrained estimator.

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. As it is bounded between minus one and plus one, the estimator of the correlation coefficient has a non-standard distribution, and can be approximated well using the bootstrap.

The bootstrap was developed by Efron [1979]. Here, we present the technique only. For the theoretical background, the monographs by Hall [1992] and Mammen [1992] are useful. Härdle, Horowitz, and Kreiss [2003] review the most recent developments in bootstrap theory, focusing on the case of dependent data. Bootstrapping the correlation coefficient for independent and identically distributed (i.i.d.) random variables appears to have been introduced by Lunneborg [1985] and Rasmussen [1987], but we are unaware of work on bootstrapping the correlation coefficient in dependent data.

Among portfolio managers who allocate and optimize portfolios in the Markowitz sense, a portfolio resampling technique is popular (see 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, and their average is used to derive a resampled frontier. This procedure appears to produce more stable portfolio weights for the computation of the efficient frontier than a single optimization step. In this respect, portfolio resampling is similar to our procedure, although it does not explicitly take into account the dependence present in financial data.

Suppose for a moment that the asset returns are i.i.d. random variables. We denote by $\chi \stackrel{\text{def}}{=} \{X_t\}_{t=1}^{T}$ a sample of return vectors $X_t \in \mathbb{R}^n$ of the n underlying assets on which the multiasset options are based. The X_t belong to some common population distribution P_0. As we 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 estimated correlation matrix $\hat{\rho}_T$.

The distribution of the estimator is denoted by $L_T(P_0) \stackrel{\text{def}}{=} \mathcal{L}\{\hat{\rho}_T(X_1, \ldots, X_T | P_0)\}$. Let $\widetilde{P}_T$ be a consistent estimate of P_0. The key insight of the bootstrap procedure is that the bootstrap estimate of $L_T(P_0)$ is simply $L_T(\widetilde{P}_T)$. An approximate confidence interval with significance level $(1-\alpha)$ can be based on the $(1-\alpha)$-quantile of $L_T(\widetilde{P}_T)$.

As we cannot compute $L_T(\widetilde{P}_T)$ explicitly, it is approximated numerically. The method works as follows:

Step 1. Generate a *resample* $\chi^* = \{X_t^*\}_{t=1}^{T}$ from χ as an unordered collection of T elements drawn randomly and with replacement from χ, such that the probability each X_t^* is equal to any of the $X_{t'}$ is T^{-1}: $P(X_t^* = X_{t'} | \chi) = T^{-1}$, $1 \le t, t' \le T$. This way, the X_t^* are i.i.d. distributed, conditional on χ. Thus, the resample χ^* may include repeats or may not include any of the X_t at all.
Step 2. Compute $\hat{\rho}_T^*$ from χ^*.
Step 3. Repeat steps one and two for many times M.
Step 4. Approximate $L_T(\widetilde{P}_T)$ by the empirical distribution of the pseudo-random variables $\hat{\rho}_{T,1}^*$ to $\hat{\rho}_{T,M}^*$.

The challenge in dependent data is that the bootstrap must be carried out in a way that preserves the dependence structure present in the data. One possible solution is the *block bootstrap*; see Hall [1985] and Carlstein [1986]. Instead of drawing single values from χ, 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 χ into B blocks $\mathcal{Y}_t \stackrel{\text{def}}{=} \{X_t, \ldots, X_{t+\ell-1}\}$ with 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 $\{X_1, \ldots, X_\ell\}$, $\{X_2, \ldots, X_{\ell+1}\}$, and so on, or $\{X_1, \ldots, X_\ell\}$, $\{X_{\ell+1}, \ldots, X_{2\ell}\}$, and so on. Now draw with replacement B blocks $\mathcal{Y}_t$ from χ. In our case the $\mathcal{Y}_t$ are $\ell \times n$ matrices of returns. These blocks are put together to form the resample χ^*. One may interpret the resample as an artificial time series of returns.

EXHIBIT 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 annualized percentages. Numbers are fixed in all calculations and supposed to reflect average market conditions.

Finally, from the entire resample compute the bootstrap correlation matrix $\hat{\rho}^*_T$. The procedure is repeated M times.

Under suitable assumptions as to the underlying stochastic process generating X_t, the approximate bootstrap distribution $L_T(\widetilde{P}_T)$ is again given by the empirical distribution of $\hat{\rho}^*_{T,1}$ to $\hat{\rho}^*_{T,M}$, (Härdle, Horowitz, and Kreiss [2003]). Note, however, that a bias results from an overlapping block bootstrap since the observations at both ends of the original sample have a lower probability of getting into the bootstrap samples than those in the middle; this needs to be corrected for. Therefore we prefer the non-overlapping bootstrap.

Choice of the block length ℓ is a delicate question. The optimal ℓ depends intrinsically on the underlying data-generating process, and must increase with the sample size T in order to achieve asymptotically correct coverage probabilities of the confidence intervals. Depending on the particular assumptions about the stochastic process and on the object to be estimated, a number of asymptotically optimal block lengths are derived. They typically range between $\ell \propto T^{1/5}$ to $\ell \propto T^{1/3}$ (see Härdle et al. [2003]). With a sample size of around 250 days in our application, this amounts to three to six days.

We experiment using these values, and do not find any strong deviations among the bootstrapped distributions. As our returns are not, or are at most mildly, serially correlated, we believe a short block length to be sufficient to capture the local dependence of the data, and decided on $\ell = 3$.

III. EMPIRICAL RESULTS

The empirical analysis uses time series of six major stocks in the German DAX index, adjusted for stock splits: Allianz (ALV), BMW, Commerzbank (CBK), Deutsche Bank (DBK), DaimlerChrysler (DCX), and Deutsche Telekom (DTE). The data were 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 select the year 1999. The year 2002 is an example of a bearish market, where correlation tends to be higher. In each case, we price multiasset options with a maturity of one year. Returns are calculated as log differences.

We compute correlation from daily returns, where the time window of past returns corresponds to the lifetime of the option, i.e., around 250 days. This is in line with market conventions. Statistically speaking, this convention implies that the market assumes stationarity of the model over the lifetime of the option. We bootstrap in blocks of $\ell = 3$ days $M = 20{,}000$ times. In this setting, we generate in the case of the three-asset model ($n = 3$) a trivariate distribution of correlations, and 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 multiasset pricers into an option price distribution. The first basket includes the three stocks CBK, DBK, and DTE; the second basket includes all six.

Exhibit 3 shows the implied volatility and dividend yields on which pricing is based. As we want to focus on correlation risk only, we keep them unchanged in the

EXHIBIT 4
Summary Statistics of Bootstrapped Correlations in Three-Asset Case

year	correlation pair	min.	max.	mean	std.
1999	DBK/DTE	0.01	0.50	0.25	0.06
	CBK/DBK	0.33	0.67	0.53	0.04
	CBK/DTE	0.03	0.45	0.27	0.06
2002	DBK/DTE	0.49	0.79	0.67	0.04
	CBK/DBK	0.62	0.84	0.75	0.03
	CBK/DTE	0.40	0.74	0.56	0.04

Bootstrap distributions obtained after 20,000 block draws from original asset returns sample; block length $\ell = 3$ days.

EXHIBIT 5
Correlation of Correlations in Bootstrap Distributions

year	correlation pair	DBK/DTE	CBK/DBK	CBK/DTE
1999	DBK/DTE	1	0.00	0.54
	CBK/DBK		1	0.17
	CBK/DTE			1
2002	DBK/DTE	1	0.23	0.64
	CBK/DBK		1	0.38
	CBK/DTE			1

1999 and 2002 scenarios. The interest rate is fixed at $r = 5\%$ throughout.

All options are assumed to be European and priced via the discounted mean of the payout function after 50,000 simulations. The maturity is $T = 1$ year. The resulting price distributions are exploited for the derivation of the bid-ask prices. In our application we stipulate a confidence level of $\alpha = 10\%$. 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 trader's point of view it is 95%, as the trader faces only a one-tailed risk for the bid and ask price.

Exhibits 4 and 5 summarize the general statistics of the bootstrapped distributions. For the sake of clarity, we focus in our discussion on the three-asset case only.[3]

There are clearly great differences between the two different market scenarios. In 1999, the mean of the two cross-sectional correlations CBK/DTE and DBK/DTE is low (around 0.25–0.27), and moderate between the two bank stocks (0.54). As expected, in 2002, correlations are higher; they range from 0.56 (CBK/DTE) to 0.75 (CBK/DBK). At the same time, the standard deviations of the correlations drop by almost one half.

Thus the precision of the correlation estimates becomes more accurate in the bear market, while in the stable upward-trending market the estimates are less precise. 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., 10 percentage points lower than the estimated mean, is well within the 95% confidence bounds, while 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 building on a correlation estimate of $\hat{\rho}_T = 0.25$ can be wrong, even if the data-generating process is stationary.

Exhibits 6 and 7 display bivariate contour plots of the (marginal) bivariate correlation distributions. Both plots are constructed so that moving from one contour line to the next corresponds to the *same* increase in altitude. Thus, the greater accuracy of the correlation estimates in the 2002 case is apparent from the much closer contour lines.

All bivariate contour plots are repeated in Exhibit 8. This gives us a better idea about their position in the unit square. The correlations in the 1999 and 2002 regimes cluster at very distinct locations in the space. As the within-sector correlations in the banking sector are much higher than the cross-sectional correlations, the DBK/DTE versus CBK/DBK ellipses always appear in the top position of each cluster.

From the general shape of the ellipses, we can deduce that the correlation structure between the estimates varies among the different assets. While for the pairs DBK/DTE and CBK/DBK and for the pairs CBK/DBK and CBK/DTE, the cross-correlation between estimates is low, which can be seen from the concentric circles in the exhibits and also from Exhibit 5, for the pairs DBK/DTE and CBK/DTE the correlation estimates are positively correlated. This is also economically plausible. One would expect correlation pairs between two different sectors (banking sector versus a cross-sector) to be not correlated or only slightly correlated, while pairs of common cross-sector correlations should be correlated, if there are common sector-specific shocks.

Interestingly, the correlation between the correlations remains relatively stable across the different market regimes

Exhibit 6
Contour Plots—1999 Data

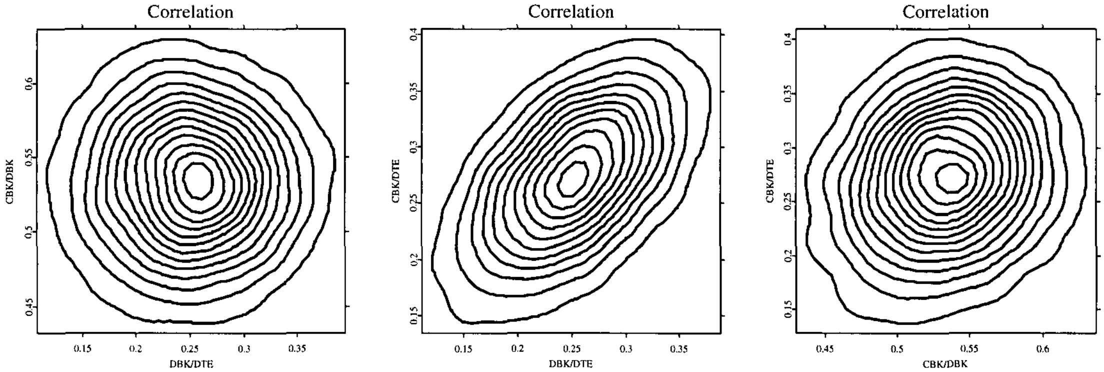

Bivariate contour plots of the (marginal) bootstrap distributions of estimated correlations in the three-asset case obtained via a kernel smoothing procedure. Initial contour line starts at 5; step width between contour lines is also 5.

Exhibit 7
Contour Plots—2002 Data

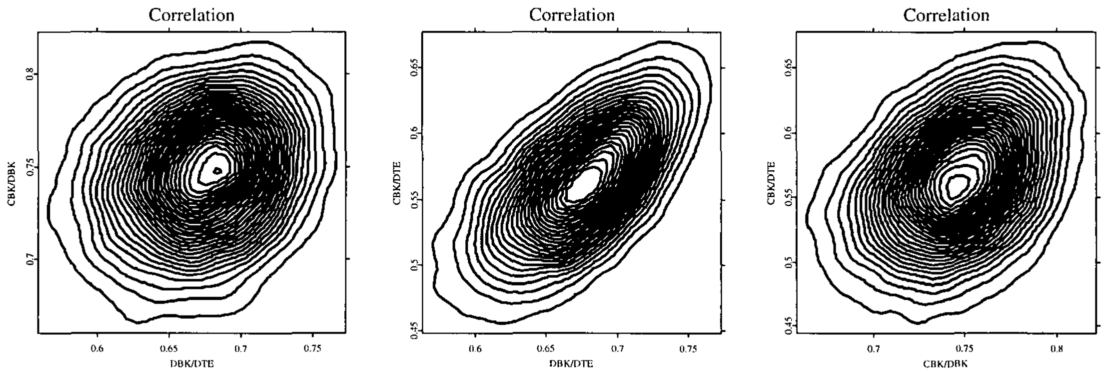

Bivariate contour plots of the (marginal) bootstrap distributions of estimated correlations in the three-asset case obtained via a kernel smoothing procedure. Initial contour line starts at 5; step width between contour lines is also 5.

(Exhibit 5). During the bearish market in 2002, the correlation of correlations increases somewhat, but in Exhibit 8 we see that the relative position of the marginal distributions remains roughly the same during the two regimes. This evidence corresponds with the findings of Neftci and Genberg [2002]; they report that the correlation between implied volatilities is relatively stable over time as well.

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

Exhibits 9 and 10 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 particular time period. In order to remove level effects in our dispersion measure, we also report the coefficient of variation of the distributions, which is defined as the standard deviation divided by the mean.

Our analysis does not discuss the absolute difference between bid and ask, but rather focuses on the bid-ask spreads relative to the mean, since typically one is interested in the spread relative to the fair price of the option. Exhibits 11 and 12 also display the price distributions.

Basket Option

Regardless of the correlation scenarios, a comparison of the three- versus the six-asset case indicates that the basket

EXHIBIT 8

All Contour Plots—1999 and 2002 Data

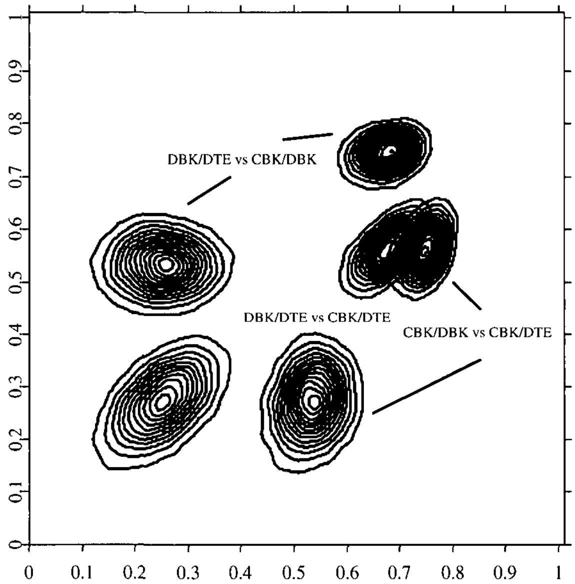

Bivariate contour plots of the (marginal) bootstrap distributions of estimated correlations in the three-asset case as plotted separately in Exhibits 6 and 7. The lower-left cluster is the 1999 data; upper-right cluster is the 2002 data.

option becomes cheaper, the more assets are included in the payout. Clearly, this is a risk diversification effect; the more assets included in the basket, the lower its volatility, and the lower its price. On the other hand, increasing the number of assets should introduce additional risk of the unknown correlations. The standard deviation and the bid-ask spreads increase only moderately. The 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 for the risk from correlations.

As is to be expected, with increasing correlations between the different market scenarios, the diversification effect is overcome, and prices rise. Yet, since the estimates of the correlations become more precise, the minimum bid-ask spreads can be reduced.

Maximum Option

Max options are by far the most expensive options; their prices can be up to 60% of the spot price. Since the payout is determined only by the performance of the best asset, prices increase more, the more assets we include in the basket. So do bid-ask spreads. Level-adjusted coefficients of variation indicate an increase of dispersion also.

As we noted earlier, 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 while the others are 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% lower in the six-asset case.

Bid-ask spreads (and coefficients of variation also) are approximately the same size in both scenarios. Thus, the effect of higher precision in estimating the correlations, 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 always receives the worst-performing asset, the price of the 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 for the six-asset minimum

EXHIBIT 9

Descriptive Statistics of Option Price Distributions—Three-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

EXHIBIT 10
Descriptive Statistics of Option Price Distributions—Six-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

EXHIBIT 11
Price Densities and Bid-Ask Spreads—Three-Asset Case

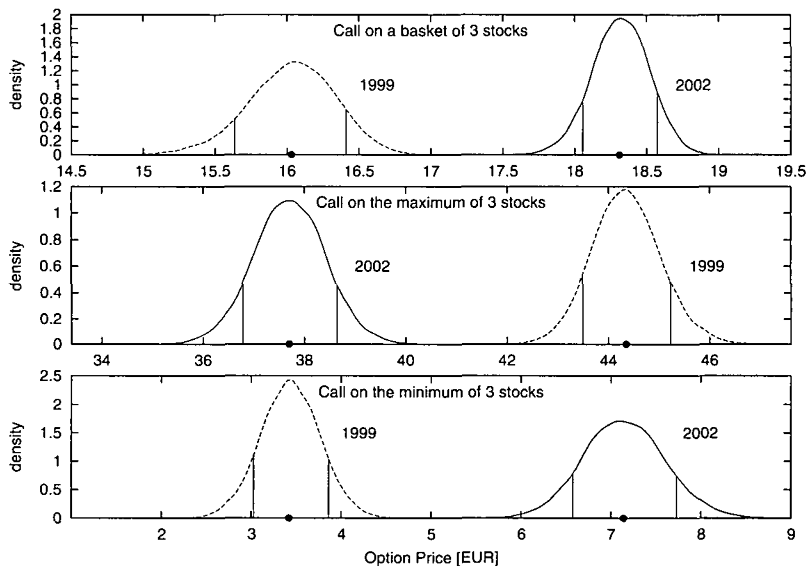

Price densities obtained from a kernel smoothing procedure. Fair price denoted by the black dot. Bid-ask spreads, quoted as the 5% and the 95% interval, indicated by vertical lines. Prices obtained by Monte Carlo simulation.

option for 1999. A higher correlation reduces both dispersion and volatility of the entire basket and results in higher prices for 2002 than for 1999. At the same time, the higher precision of the correlation estimates reduces the bid-ask spreads relative to the low-correlation scenario. This effect is stronger than the increased correlation sensitivity in high-correlation situations.

Although we do not compute formal tests for normality, our exhibits for skewness and kurtosis confirm that a normal variate for computing the price quantiles of multiasset basket options, such as the well known two-times-sigma rule, will be an acceptable approximation. Under the normal distribution, the coefficient of variation and the spread per mean are multiples of the same quantity, and thus always move in the same direction.

In the case of the min options, the normal distribution may not be well-suited, because non-negligible skewness emerges, the closer the price is to zero. When prices are close to zero, using the normal distribution as an approximation may not be advisable anyway.

IV. CONCLUSIONS

Pricing and hedging multiasset equity options requires estimating correlations. Even when the time series of stock prices is assumed to be stationary, the correlation estimates

EXHIBIT 12
Price Densities and Bid-Ask Spreads—Six-Asset Case

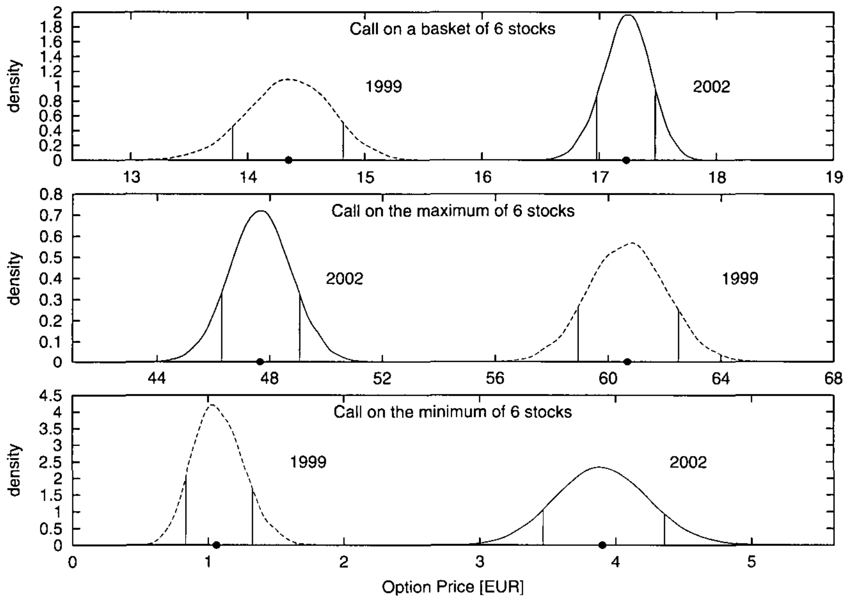

Price densities obtained from a kernel smoothing procedure. Fair price denoted by the black dot. Bid-ask spreads, quoted as the 5% and the 95% interval, indicated by vertical lines. Prices obtained by Monte Carlo simulation.

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 multiasset options (a basket option, and minimum and maximum options). Bid-ask spreads of the prices are computed as statistical quantiles of the resulting price distributions. The resulting bid-ask spreads can be considered minimum 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 differences shown 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 the 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, while in a bullish market correlations are more moderate, with low statistical confidence. 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 the correlation estimates in the current regime.

Our results imply that traders quoting bid-ask spreads of multiasset equity options need to carry out the same thorough correlation analysis as do investors considering these products for their portfolios. Finally, our block bootstrap procedure may prove to be useful in similar contexts, such as in an analysis of the impact of estimation errors on the quality of value at risk computations.

ENDNOTES

Fengler acknowledges financial support by the Deutsche Forschungsgemeinschaft and the Sonderforschungsbereich 373 *Quantifikation und Simulation ökonomischer Prozesse.* An earlier version of this article circulated under the title "Correlation Risk Premia for Multi-Asset Equity Options." This article represents the authors' personal opinions and does not necessarily reflect the views of their employers.

[1]An exception are the liquid OTC quanto options on the so-called world basket consisting of the EuroStoxx 50, the Nikkei, and the S&P 500 index.

[2]Note that in the Black-Scholes framework, correlation vegas and cross-gammas are linked via the relationship: $\partial C/\partial \rho_{ij} = S_i(t)S_j(t)\sigma_i\sigma_j(T-t)\,\partial^2 C/\partial S_i \partial S_j$ (see Reiß and Wystup [2002]). Therefore, the arguments we develop to explain the sign of a

particular correlation vega apply also to the corresponding cross-gamma.

[3]Statistics on the six-asset case may be obtained from the authors upon request.

REFERENCES

Alexander, C., and C.T. Leigh. "On the Covariance Matrices Used in Value at Risk Models." *The Journal of Derivatives*, 4 (1997), pp. 50–62.

Avellaneda, M., D. Boyer-Olson, J. Busca, and P. Friz. "Reconstructing Volatility." *Risk*, 15 (2002), pp. 87-91.

Baba, Y., R.F. Engle, D.F. Kraft, and K.F. Kroner. "Multivariate Simultaneous Generalized ARCH Models." Working paper, Department of Economics, University of California, San Diego, 1990.

Carlstein, E. "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.

Efron, B. "Bootstrap Methods: Another Look at the Jackknife." *Annals of Statistics*, 7 (1979), pp. 1–26.

Engelmann, B. "Quasi-Monte Carlo Techniques for the Valuation of Contingent Claims on Several Assets." In J. Hakala and U. Wystup, eds., *Foreign Exchange Risk*. New York: Risk Books, 2002, pp. 209–225.

Engelmann, B., and P. Schwendner. "The Pricing of Multi-Asset Options Using a Fourier Grid Method." *Journal of Computational Finance*, 1 (1998), pp. 63–72.

Fengler, M.R., and H. Herwartz. "Multivariate Volatility models." In W. Härdle, T. Kleinow, and G. Stahl, eds., *Applied Quantitative Finance*. Heidelberg: Springer, 2002.

Hall, P. *The Bootstrap and the Edgeworth Expansion*. New York and Heidelberg: Springer Verlag, 1992.

——. "Resampling a Coverage Process." *Stochastic Process Applications*, 19 (1985), pp. 259–269.

Härdle, W., J. Horowitz, and J.-P. Kreiss. "Bootstrap Methods for Time Series." *International Statistical Review*, 71 (2003), pp. 435–359.

Hendricks, D. "Evaluation of Value-at-Risk Models Using Historical Data." *Federal Reserve Bank of New York Economic Policy Review*, 2 (1996), pp. 39–69.

Johnson, H. "Options on the Maximum or the Minimum of Several Assets." *Journal of Financial and Quantitative Analysis*, 22 (1987), pp. 277–283.

Lunneborg, C.E. "Estimating the Correlation Coefficient: The Bootstrap Approach." *Psychological Bulletin*, 98 (1985), pp. 209–215.

Mammen, E. *When Does Bootstrap Work? Asymptotic Results and Simulations*. Number 77 in Lecture Notes in Statistics. New York and Heidelberg: Springer Verlag, 1992.

Margrabe, W. "The Value of an Option to Exchange One Asset for Another." *Journal of Finance*, 33 (1978), pp. 177–186.

Michaud, R. *Efficient Asset Management*. Boston: Harvard Business School Press, 1998.

Neftci, S.N., and H. Genberg. "Excessive Variation of Risk Factor Correlations and Volatilities." *Journal of Futures Markets*, 22 (2002), pp. 1119–1146.

Overhaus, M. "Himalaya Options." *Risk*, 15 (2002), pp. 101–104.

Quessette, R. "New Products, New Risks." *Risk*, 15 (2002), pp. 97–100.

Rasmussen, J.L. "Estimating the Correlation Coefficient: Bootstrap and Parametric Approaches." *Psychological Bulletin*, 101 (1987), pp. 136–139.

Reiß, O., and U. Wystup. "Efficient Computation of Option price Sensitivities Using Homogeneity and Other Tricks." In J. Hakala and U. Wystup, eds., *Foreign Exchange Risk*. New York: Risk Books, 2002, pp. 127–142.

Scherer, B. "Portfolio Resampling: Review and Critique." *Financial Analysts Journal*, 58 (2002), pp. 98–109.

Stulz, R. "Options on the Minimum or the Maximum of Two Risky Assets." *Journal of Financial Economics*, 10 (1982), pp. 161–181.

Walter, C.A., and J.A. Lopez. "Is Implied Correlation Worth Calculating? Evidence from Foreign Exchange Options." *The Journal of Derivatives*, 7 (2000), pp. 65–82.

Wystup, U. "How the Greeks Would Have Hedged Correlation Risk of Foreign Exchange Options." In J. Hakala and U. Wystup, eds., *Foreign Exchange Risk*. New York: Risk Books, 2002, pp. 143–146.