Nonparametric pricing of multivariate contingent claims

April 2000

Joshua V. Rosenberg
Department of Finance
NYU - Stern School of Business
44 West 4th Street, Suite 9-190
New York, New York 10012-1126
(212) 998-0311
email: jrosenb0@stern.nyu.edu

Abstract

This paper derives and implements a nonparametric, arbitrage-free technique for multivariate contingent claims (MVCC) pricing. This technique is based on nonparametric estimation of a multivariate risk-neutral density function using data from traded options markets and historical asset returns. "New" multivariate claims are priced using expectations under this measure. An appealing feature of nonparametric arbitrage-free derivative pricing is that fitted prices are obtained that are consistent with traded option prices and are not based on specific restrictions on the underlying asset price process or the functional form of the risk-neutral density.

Nonparametric MVCC pricing utilizes the method of copulas to combine nonparametrically estimated marginal risk-neutral densities (based on options data) into a joint density using a separately estimated nonparametric dependence function (based on historical returns data). This paper provides theory linking objective and risk-neutral dependence functions, and empirically testable conditions that justify the use of historical data for estimation of the risk-neutral dependence function.

The nonparametric MVCC pricing technique is implemented for the valuation of bivariate underperformance and outperformance options on the S&P500 and DAX index. Price deviations are found to be significant in comparisons of fitted prices using nonparametric valuation versus standard multivariate diffusion-based valuation. An analysis of pricing errors indicates that the lognormal marginal density specification poorly approximates the negative skewness and excess kurtosis implied by market data, and the lognormal copula specification poorly approximates the asymmetric return dependence implied by market data. These results suggest that correct specification of the marginal densities and the dependence function that define the multivariate risk-neutral density is essential for accurate MVCC pricing.

The author wishes to thank the Deutsche Borse and Christa Weise for providing the DAX option data used in this paper. The most recent version of this paper is available at http://www.stern.nyu.edu/~jrosenb0/J_wpaper.htm.

I. Introduction

Many recent papers have utilized arbitrage-free techniques for pricing univariate derivative securities. Often, these techniques are based on estimation of a univariate risk-neutral density using a set of observed option prices for a particular underlying asset. The risk-neutral density may then be used to price other derivatives with the same underlying asset. Related papers include Sherrick, Irwin, and Foster (1990, 1992), Shimko (1993), Derman and Kani (1994), Dupire (1994), Rubinstein (1994a), Longstaff (1995), and Ait-Sahalia and Lo (1998).[1]

None of these techniques may be directly applied to the pricing of multivariate contingent claims (MVCCs), because knowledge of the marginal risk-neutral densities is not sufficient to identify the multivariate risk-neutral density that is used for pricing. In general, there are an infinite number of multivariate densities consistent with given marginals. Identification of the multivariate density requires an estimate of the marginal densities and the dependence function that links them together.

Most MVCC pricing papers have used generalizations of the Black-and Scholes (1973) continuous-time Brownian motion framework or the Cox, Ross, Rubinstein (1979) discrete-time binomial framework. Examples of the first type include Margrabe (1978), Stulz (1982), Johnson (1987), Reiner (1992), and Shimko (1994). Examples of the second type include Stapleton and Subrahmanyam (1984a, 1984b), Boyle (1988), Boyle, Evnine, and Gibbs (1989), and Rubinstein (1992, 1994b).

A drawback of these approaches is the restrictive nature of the assumptions on the underlying price processes. To the extent that underlying price processes exhibit features such as jumps or stochastic volatility, the aforementioned pricing formulas will result in pricing inaccuracies and biases. Ho, Stapleton, Subrahmanyam (1995) provide a binomial approximation technique for MVCC pricing that provides more generality than the previously mentioned papers.

Rosenberg (1998) provides an arbitrage-free MVCC valuation technique based on estimation of a flexibly parameterized multivariate risk-neutral density function. The multivariate risk-neutral density function is extracted using the prices of options on individual assets and the prices of options on multiple assets. This technique is well-suited to pricing multivariate currency options for which

[1] Ait-Sahalia and Lo (1998) focus on estimation of the state-price density. The risk-neutral density is obtained from the state-price density by normalizing the state-price density so that it integrates to unity.

there are traded options on individual currencies and cross-currency options.[2] Rosenberg (1999) provides a semi-parametric MVCC pricing technique based on a parametric dependence function that allows the marginal densities to be estimated nonparametrically. Neither of these papers permits nonparametric estimation of the dependence function.

This paper proposes a nonparametric MVCC pricing technique based on the method of copulas (e.g. Nelsen (1998)), in which nonparametric marginal risk-neutral densities are combined with a separately estimated nonparametric risk-neutral dependence function to obtain a nonparametric risk-neutral joint density function.[3] This technique is fully nonparametric, and generalizes univariate nonparametric derivative pricing (Ait-Sahalia and Lo, 1998) to the multivariate case. An appealing feature of *nonparametric* arbitrage-free derivative pricing is that fitted prices are obtained that are consistent with traded option prices and are not based on specific restrictions on the underlying asset price process or the functional form of the risk-neutral density. This paper also provides theory linking objective and risk-neutral dependence functions, and empirically testable conditions that justify the use of historical data for estimation of the risk-neutral dependence function.

Nonparametric MVCC pricing allows the marginal risk-neutral densities to be completely general. Hence, this technique is consistent with non-lognormalities in univariate risk-neutral densities that have been documented in many papers (e.g. Ait-Sahalia and Lo (1998)). Nonparametric MVCC pricing also allows the risk-neutral dependence function that links the behavior of the underlying asset prices to be completely general. In an environment characterized by multivariate stochastic volatility (e.g. Engle and Kroner (1995) or Harvey, Ruiz, and Shepherd (1994)), multivariate extreme value distributions (e.g. Tawn (1990)), asymmetric correlations (e.g. Erb, Harvey, Viskanta (1994)), or multivariate correlated jumps, the dependence function may deviate from the elliptical class of distributions (e.g. multivariate normal or lognormal).

Nonparametric MVCC pricing is implemented for the valuation of bivariate underperformance and outperformance options on the S&P500 and DAX 30 indices.[4] Price deviations are found to be significant in comparisons of fitted prices using nonparametric valuation versus standard multivariate diffusion-based valuation. The pricing errors are attributed to two sources: (1) negative skewness and excess kurtosis in the empirical marginal densities that is not reflected in the

---

[2] As shown in Rosenberg (1998), a cross-currency option is equivalent to a bivariate option on the difference between two exchange rates with the same numeraire currency.

[3] Embrechts, McNeil, and Strauman (1998) discuss the use of copulas in risk management applications where asset returns are non-normal. Hull and White (1998) use a copula-based technique for estimation of value-at-risk.

[4] The DAX index tracks the price of a portfolio of 30 German blue chip stocks including reinvested dividends.

lognormal marginal specification and (2) asymmetric correlation in the empirical dependence function that is not reflected in the lognormal dependence function specification. These results suggest that correct specification of the marginal densities and the dependence function that define the multivariate risk-neutral density is essential for accurate MVCC pricing.

The remainder of the paper is structured as follows. Section II describes the theory underlying the nonparametric multivariate contingent claim pricing. Section III presents the nonparametric multivariate density estimation technique. Section IV utilizes the nonparametric densities for valuation of underperformance and outperformance options on the S&P500 and DAX 30 indices. Section V concludes the paper.

II. Nonparametric pricing of MVCCs: Theory

The "risk-neutral density" representation of the asset pricing formula has been found to be a convenient framework for univariate arbitrage-free pricing. This is because a risk-neutral density representation is guaranteed to exist in the absence of arbitrage, the same risk-neutral density may be used to price all claims dependent on the same underlying asset, and interpolated prices using this representation do not generate arbitrage opportunities.

Equation (1) presents the risk-neutral density representation of the univariate asset pricing formula. $D_{X,t}$ is the price of a European-style derivative asset with payoff function $g_X(X_T)$ and underlying asset price X_T on the expiration date, r_f is the riskless rate of interest, and $f_{X,t}{}^*(X_T)$ is the risk-neutral density of X_T.

(1) $$D_{X,t} = e^{-r_f(T-t)} \int g_X(X_T) f^*_{X,t}(X_T) dX_T$$

Equation (1) shows that knowledge of the risk-neutral density is sufficient to price a European claim with a payoff function that depends on X_T. Equation (1) also implicitly provides an estimation strategy for obtaining the risk-neutral density $f_{X,t}{}^*(X_T)$ given the prices of traded options on X_T.

Theorem 1: Let C(K) be a function that represents the prices of a cross-section of European call options. These options have expiration date T, exercise prices (K) that are continuous on the range $[0,\infty)$, and payoff functions of $\mathrm{Max}[0, X_T - K]$ where X_T is the underlying price on the option expiration date. The riskless rate of interest is denoted rf. Then, the cumulative risk-neutral density and risk-neutral density are uniquely identified as: $Prob^*(X_T < K) = F_{X,t}^*(K) = e^{rf(T-t)}\partial C/\partial K + 1$ and $f_{X,t}^*(K) = e^{rf(T-t)}\partial^2 C/\partial K^2$.

Proof: See Breeden and Litzenberger (1978) or appendix A.

Theorem 1, which was originally proved in Breeden and Litzenberger (1978), shows that if a continuous cross-section of option premia is available, then the risk-neutral density may be obtained by differentiating twice with respect to the exercise price. Ait-Sahalia and Lo (1998) use Theorem 1 as the basis for nonparametric estimation of the S&P500 risk-neutral density using a cross-section of S&P500 option prices.

Equation (2) provides risk-neutral representation of the bivariate contingent claim pricing formula. This is based on a bivariate risk-neutral density $f_{X,Y,t}^*(X_T,Y_T)$. For simplicity, the bivariate case is discussed in the remainder of the paper. Generalization to the n-variate case is straightforward.

(2) $$D_{X,Y,t} = e^{-r_f(T-t)} \iint g_{X,Y}(X_T, Y_T) f^*_{X,Y}(X_T, Y_T) dX_T dY_T$$

Theorem 2: Let C(K,L) be a function that represents the prices of a cross-section of European bivariate product options. These options have expiration date T, exercise prices (K, L) that are continuous on the range $[0,\infty) \times [0,\infty)$, and payoff functions of $Max[0, X_T - K]*Max[0, Y_T - L]$ where X_T and Y_T are the underlying prices on the option expiration date. The cumulative joint risk-neutral density and joint risk-neutral density are uniquely identified as: $Prob^*(X_T < K, Y_T < L) = F_{X,Y,t}^*(K,L) = -e^{rf(T-t)}\partial^2 C/\partial K\partial L + 1$ and $f_{X,Y,t}^*(K,L)= \partial^2 F^*/\partial K\partial L$.

Proof: see appendix A.

If there are traded product options that satisfy the conditions of Theorem 2, then nonparametric estimation of the bivariate risk-neutral density is a straightforward generalization of the univariate case. However, there are rarely liquid markets for multivariate derivatives, and prices for such assets are not publicly disseminated.[5]

[5] The author is unaware of any traded options that have the payoff function required by Theorem 2.

The structure of available options price data motivates the estimation technique developed in this paper. There is sufficient information to nonparametrically estimate $f_{X,t}^{*}(X_T)$ and $f_{Y,t}^{*}(Y_T)$ using Theorem 1, but not to nonparametrically estimate $f_{X,Y,t}^{*}(X_T, Y_T)$ using Theorem 2. Thus, it is natural to separate the estimation of the marginals and the dependence function that combines the marginals into the joint density. The method of copulas is designed to deal with exactly this situation. A copula is a dependence function that links (or couples) separately estimated marginal densities into a joint density function.

Theorem 3 (Sklar's Theorem): If F_{XY} is a joint cumulative distribution function with cumulative marginal densities of F_X and F_Y, then there exists a copula such that $F_{XY}(X_T,Y_T)= Cop(F_X(X_T),F_Y(Y_T))$, where a copula (Cop(u,v)) is any function satisfies the properties of a joint cumulative density function on [0,1]x[0,1]. In this case, $f_{XY}(X_T,Y_T) = \partial^2 F_{XY}(X_T,Y_T)/ \partial X_T \partial Y_T$. Proof: See Nelsen (1998), Theorem 2.3.3.

Theorem 3 states that all bivariate densities have a representation based on their marginal densities and their dependence function (copula). Thus, a bivariate density may be factored into its marginals and dependence function without loss of generality. And, a bivariate density may be constructed using the marginals and the copula of the density function. These results apply to the objective density function that measures the actual probability of state-dependent outcomes and to the risk-neutral density that measures the "risk-adjusted" probability of state-dependent outcomes. The key implication of Sklar's Theorem is that knowledge of the risk-neutral marginals and risk-neutral copula is sufficient to construct the bivariate risk-neutral density and to price bivariate MVCC's.

When there is insufficient option data to directly estimate the risk-neutral copula, another source of data must be used. The following invariance property suggests conditions under which the objective copula may be used to estimate the risk-neutral copula.

Theorem 4: Suppose that X_T and Y_T are continuous random variables with objective cumulative marginal densities of $F_X(X_T)$ and $F_Y(Y_T)$. If the risk-neutral random variables (X_T^{*},Y_T^{*}) are strictly increasing functions of the objective random variables (X_T, Y_T), then the risk-neutral copula is identical to the objective copula, and the risk-neutral joint cumulative density is given by $Cop(F_X^*(X_T), F_Y^*(Y_T))$ where $F_X^*(X_T)$ and $F_Y^*(Y_T)$ are the cumulative marginal risk-neutral densities and Cop is the copula for the objective bivariate density.

Proof: Theorem 2.4.3 of Nelsen (1998).

Theorem 4 states that the copulas for two pairs of random variables are identical as long as on pair is obtained by strictly increasing transformation of the other pair. Hence, the objective and risk-neutral copulas are identical, as long as the risk-neutral random variables (X_T^*, Y_T^*) are strictly increasing functions of the objective random variables (X_T, Y_T).[6] When this theorem is satisfied, the difference between the bivariate risk-neutral and bivariate objective density is embedded in the marginal densities; the dependence function is the same for both bivariate densities.

Lemma 4 provides a sufficient condition for the risk-neutral random variables to be a strictly increasing functions of the objective random variables. This condition must be verified empirically.

Lemma 4: Suppose that $X^*(X)$ and $Y^*(Y)$ represent the risk-neutral random variables as functions of the objective random variables. If $dX^*(X)/dX > 0$ and $dY^*(Y)/dY > 0$, then the risk-neutral random variables are strictly increasing functions of the objective random variables.

It is possible to extract a copula from a given multivariate density function using the "method of inversion." This technique factors out the effects of the marginal densities on the dependence relation by substituting the arguments of the joint density with the quantile functions. The derived function is a copula, since takes a pair of cumulative probabilities as inputs and generates a cumulative probability as an output, i.e. it has the properties of a joint cumulative density function on [0,1]x[0,1].

Theorem 5 ("Method of inversion", Nelsen (1998), Corollary 2.3.7): Suppose that $F_{XY}(X_T, Y_T)$ is a joint density function with marginal quantile functions (inverse marginal cumulative density functions) given by $F_X^{-1}(u)$ and $F_Y^{-1}(v)$ where u and v are cumulative probabilities. Then, the copula of F_{XY} is given by $Cop(u,v) = F_{XY}(F_X^{-1}(u), F_Y^{-1}(v))$.

Proof: See appendix A.

In cases where the objective copula is identical to the risk-neutral copula, the method of inversion may be applied to the objective density (estimated using historical returns data) to obtain the objective copula as an estimate of the risk-neutral copula. Then, the bivariate risk-neutral density function may be obtained by combining the separately estimated marginals and copula.

In some cases, hedge ratios for multivariate derivatives may be estimated directly from the observed cross-section of option premia.[7] These hedge ratios depend on first and second derivatives of the pricing function with respect to the exercise price, so they are closely linked to the risk-neutral density. Applying Bates (1995) to bivariate case: $\Delta_X = \partial D_{X,Y,t}/\partial X_t = (1/X_t)(D_{X,Y,t} + K\partial D_{X,Y,t}/\partial K)$, $\Gamma_X = \partial^2 D_{X,Y,t}/\partial X_t^2 = (K/X_t)^2(\partial^2 D_{X,Y,t}/\partial K^2)$, $\Delta_Y = \partial D_{X,Y,t}/\partial Y_t = (1/Y_t)(D_{X,Y,t} + L\partial D_{X,Y,t}/\partial L)$, and $\Gamma_Y = \partial^2 D_{X,Y,t}/\partial Y_t^2 = (L/Y_t)^2(\partial^2 D_{X,Y,t}/\partial L^2)$.

III. Nonparametric pricing of MVCCs: Estimation

III.i. Risk-neutral marginals

Estimation of the risk-neutral marginal densities is accomplished using the nonparametric estimator developed in Ait-Sahalia and Lo (1998). Specifically, the call option pricing function is estimated using a kernel regression of option prices on pricing formula variables. Then, the first derivative of the nonparametric pricing function with respect to the exercise price defines the cumulative risk-neutral density, as stated in Theorem 1.

The most general estimator proposed by Ait-Sahalia and Lo (1998, equation 11) requires a kernel regression of call prices on five variables (underlying price, exercise price, time-until-expiration, riskless rate, and dividend yield). However, the empirical results reported in Ait-Sahalia and Lo (1998) are based on a dimensionality reduction approach: the call price function is given by the Black-Scholes formula with the volatility parameter replaced by a nonparametrically estimated function (Ait-Sahalia and Lo, 1998, equations 9, 13). This paper utilizes a similar dimensionality reduction approach.

---

[6] A result similar to Theorem 4 holds when either risk-neutral random variable (or both) is a strictly decreasing function of the objective random variable. When the functional relationship is non-monotonic, there is no such result.

[7] In the univariate case, the Bates (1995) hedge ratios require that the call or put pricing formula is homogeneous of degree one in the underlying price and the exercise price. To extend this result to the bivariate case, the derivative pricing formula must be homogeneous of degree one in X_t and K as well as Y_t and L.

Let σ_i be the Black-Scholes implied volatility, K is the option exercise price, T-t is the option time-until-expiration, and F_t is the price of a futures contract with identical expiration and underlying asset as the option contract.[8] Also, let k(·) be the Nadaraya-Watson kernel estimator with a bandwidth of h. Then:

(3) $$\hat{C}_t = BS(F_t, K, T-t, \hat{s}(F_t, K, T-t))$$

(4) $$\hat{s}(F_t, K, T-t) = \frac{\sum_{i=1}^{n} k_F\left((F_t - F_{t_i})/h_F\right) k_K\left((K-K_i)/h_K\right) k_{T-t}\left(((T-t)-(T-t)_i)/h_{T-t}\right) s_i}{\sum_{i=1}^{n} k_F\left((F_t - F_{t_i})/h_F\right) k_K\left((K-K_i)/h_K\right) k_{T-t}\left(((T-t)-(T-t)_i)/h_{T-t}\right)}$$

(5) $$\frac{\partial \hat{C}_t}{\partial K} = \frac{\partial BS}{\partial K} + \frac{\partial BS}{\partial \hat{s}} \frac{\partial \hat{s}}{\partial K}$$

The call price function defined in equation (3) relates the option price to the futures price, exercise price, option maturity, and option implied volatility. The option implied volatility function in equation (4) relates the implied volatility to the futures price, exercise price, and option maturity using a kernel estimator.

The first derivative of the call option pricing formula with respect to the option exercise price is given in equation (5) and obtained using the chain rule. The first term is the change in the call price due to a change in the exercise price, and the second term is the change in the call price due to a change in implied volatility multiplied by the implied volatility change due to a change in the exercise price. The multiplier in the second term depends on the nonparametrically estimated "implied volatility surface."

III.ii. Risk-neutral copula

Theorem 4 shows that the risk-neutral copula may be estimated using the objective copula, as long as the risk-neutral random variables are increasing functions of the objective random variables. This

[8] When the cost-of-carry model holds, $F_t = X_t e^{(r-\delta)(T-t)}$ so that the futures price (F_t) combines the effects of the spot price (X_t), riskless rate (r), and dividend yield (δ) on the call option price. The difference between the riskless rate and the dividend yield (r-δ) is the cost-of-carry.

section provides a technique to verify this condition and a technique for estimation of the objective copula using historical returns data.

The functions $X^*(X)$ and $Y^*(Y)$ define the relationships between the risk-neutral and objective random variables. We are interested in signing the derivatives $dX^*(X)/dX$ and $dY^*(Y)/dY$ to determine whether these functions are strictly increasing. It is not possible to directly evaluate $dX^*(X)/dX$ and $dY^*(Y)/dY$, since the functions $X^*(X)$ and $Y^*(Y)$ are unobserved. However, $X^*(X)$ and $Y^*(Y)$ are implicitly defined by the relations: $E[X^*] = E[X^*(X)]$ and $E[Y^*] = E[Y^*(Y)]$.

First, consider the following local linear approximations and their derivatives:

$$(6) \qquad X^*(X) \cong a_i + b_i X, \quad dX^*(X)/dX \cong b_i \qquad X \in [x_i, x_{i+1}] \quad i=1\ldots n$$

$$Y^*(Y) \cong c_i + d_i Y, \quad dY^*(Y)/dY \cong d_i \qquad Y \in [y_i, y_{i+1}] \quad i=1\ldots n$$

The b_i's and d_i's in equation (6) represent the slope of the approximating function over each interval. If $b_1 \ldots b_n > 0$ and $d_1 \ldots d_n > 0$, then the approximating functions are strictly increasing, and we will infer that the original functions are also strictly increasing.

To estimate the b_i's and d_i's, it must be noted that the risk-neutral expected value within each interval has a linear representation in terms of the objective expected value. Define EX_i^{*+} and EX_i^{+} as the risk-neutral and objective expectations taken over the first-half of interval i, i.e. over the interval $[x_i, x_i + (x_i + x_{i+1})/2]$. Define EX_i^{*-} and EX_i^{-} as the risk-neutral and objective expectations taken over the second-half of interval i, i.e. over intervals of $[x_i + (x_i + x_{i+1})/2, x_{i+1}]$. The analogous definitions apply to EY_i^{*+}, EY_i^{+}, EY_i^{*-}, and EY_i^{-}.

Taking expectations of equation (6) over each subinterval generates two equations for each random variable over each interval.

$$(7) \qquad EX_i^{*+} = a_i + b_i \, EX_i^{+} \qquad EX_i^{*-} = a_i + b_i \, EX_i^{-} \qquad i = 1...n$$

$$EY_i^{*+} = c_i + d_i \, EY_i^{+} \qquad EY_i^{*-} = c_i + d_i \, EY_i^{-} \qquad i = 1...n$$

The risk-neutral expectations on the left-hand sides of equation (7) are evaluated by numerical integration over the estimated risk-neutral densities, $f(X^*)$ and $f(Y^*)$. The objective expectations on the right-hand sides of these equations are evaluated by numerical integration over the estimated objective densities, $f(X)$ and $f(Y)$. Solving these equations, the b_i's and d_i' are estimated as the change in the risk-neutral expectation divided by the change in the objective expectation.

(8) $b_i = (EX_i^{*+} - EX_i^{*-}) / (EX_i^{+} - EX_i^{-})$ $\quad$ $i = 1...n$

$d_i = (EY_i^{*+} - EY_i^{*-}) / (EY_i^{+} - EY_i^{-})$ $\quad$ $i = 1...n$

In cases where $b_1...b_n > 0$ and $d_1...d_n > 0$, it is appropriate to estimate the risk-neutral copula using the objective copula. The method of inversion (Theorem 5) requires an estimate of the joint and marginal densities to obtain the copula. In practice, these densities are not known "in population," but instead must be estimated using sample data.

This paper proposes a sample estimator for the objective copula estimator based on a continuous nonparametric representation of the sample cumulative density function.[9] Consider the Nadaraya-Watson kernel estimators of the marginal and joint objective densities, where $k(\cdot)$ is the kernel, h_x, and h_y are the kernel bandwidths and X_t and Y_t are historical returns observed from period 1…T. The kernel estimates may be interpreted as "smoothed histograms," where the degree of smoothness is determined by the magnitude of the kernel bandwidths.

(9) $$f_X(x) = (Th_x)^{-1}\sum_{t=1}^{T} k\big((X_t - x)/h_x\big) \qquad f_Y(y) = (Th_y)^{-1}\sum_{t=1}^{T} k\big((Y_t - y)/h_y\big)$$

(10) $$f_{XY}(x,y) = (T^2 h_x h_y)^{-1}\sum_{t=1}^{T} k\big((X_t - x)/h_x\big) * k\big((Y_t - y)/h_y\big)$$

The required functions to obtain the objective copula are the quantile functions and the joint cumulative density function. Using numerical integration, F_X and F_Y are obtained from the kernel estimates of f_X and f_Y in (9) and (10). Then, the quantile functions are obtained by numerical inversion of F_X and F_Y.

(11) $$F_X^{-1}(u) = x \mid F_X(x) = u \qquad F_Y^{-1}(v) = y \mid F_Y(y) = v$$

[9] Deheuvels (1979) and Nelsen (1998, pp. 176-179) provide a sample estimator for the objective copula using the sample cumulative density functions, e.g. $\text{Prob}(X \le x)$ = (number of X_T such that $X_T \le x$) / T. These cumulative density functions are discontinuous, exhibiting a "staircase" pattern of constant cumulative probability in-between observed return points and jumps at observed return points.

The joint cumulative density function, F_{XY}, is obtained by numerical integration of the kernel-estimated joint density function in (10). Then, the sample copula estimator is identical to the population copula using the method of inversion with the sample quantile functions and the joint density function replacing their population counterparts, i.e. $Cop = Cop(u,v) = F_{XY}(F_X^{-1}(u), F_Y^{-1}(v))$.

IV. Pricing underperformance and outperformance options on the S&P500 and DAX 30 indices

IV.i Data

S&P500 futures option and S&P500 futures data is obtained from the Chicago Mercantile Exchange for the month of December 1999. Daily settlement prices for each futures option contract and corresponding maturity futures contract are extracted along with contract type, maturity, and daily trading volume. Futures options and futures with expirations in the March expiration cycle (March, June, September, and December) with expiration dates through December 2000 are utilized.[10] A riskless interest rate for each contract is calculated using the corresponding maturity interest rate from the daily term structure of British Banker's Association LIBOR rates, based on Datastream data.

The risk-neutral density estimation technique described in the previous section is based on implied volatilities for a cross-section of European-style options. Index futures options, which have American exercise style, exhibit a price differential compared with index options due to the early exercise premium and the effect of convergence of index and index futures prices, as shown in Brenner, Courtadon, Subrahmanyam (1985). To infer the price of a European-style index option corresponding to the observed price of an American-style futures option, it is necessary to measure and then subtract off the price differential.

---

[10] S&P500 futures option and futures data offers superior price synchronicity compared with S&P500 index option and S&P500 index data. Under most circumstances, the futures and futures option settlement prices are the average of the highest and lowest transaction prices in the last 30 seconds of trading. To ensure that settlement prices reflect market conditions at the end of the day, alternative calculation techniques are used if the futures option or futures does not trade near the close. For additional details, see Rule 813 of the Chicago Mercantile Exchange Rulebook.

The CBOE does not report option settlement prices, so tick-by-tick quote or trade data must be used to infer the price at the close of trading. In addition, the S&P500 index price (as an average of closing prices of the constituent stocks) may not reflect closing market conditions due to the fact that some stocks may not trade near the close. Various techniques to ameliorate these problems have been proposed and implemented, see e.g. Dumas, Fleming, and Whaley (1998).

The price differential for each option contract is calculated as the difference between the fitted price of American and European futures option contract with the same terms as the desired contract. This paper uses a modified version of the BBSR pricing algorithm described in Broadie and Detemple (1996) to obtain these fitted option prices and the price differential (see Appendix B). The adjusted price (the observed price minus the price differential) is used as an input to calculate the implied volatility for each contract. Estimation of the implied volatility is accomplished by numerically inverting the Black (1976) futures option pricing formula, using the data described above.

DAX 30 index option and futures settlement data is obtained from the Deutsche Borse for the month of December 1999.[11] Daily settlement prices for each option contract and corresponding maturity futures contract are extracted along with contract type, maturity, and daily trading volume.[12] Options and futures with expirations in the March expiration cycle and expirations through December 2000 are utilized. A riskless interest rate for each contract is calculated using the corresponding maturity interest rate from the daily term structure of FIBOR (Frankfurt Interbank Offer) rates, based on Datastream data. An implied volatility for each contract is obtained by numerically inverting the Black (1976) futures option pricing formula, using the data described above.[13] Since these contracts have European exercise-style, no price adjustment is necessary.

Several screening conditions are used to eliminate illiquid options. Options with less than two weeks until expiration or less than five contracts traded over the day are excluded from the sample. Options that are deeply in-the-money or out-of-the-money as measured by a proportional moneyness, ln(exercise price/futures price), less than -40% or greater than 30% are excluded from the sample. Options with unreasonably high or low implied volatilities (less than 5% or greater than 90%) are also excluded.

The first panel of Table 1 describes the characteristics of the S&P500 futures option data and the second panel of Table 1 describes the DAX 30 options data. There are 711 S&P500 contracts and 1680 DAX 30 contracts in December 1999 that satisfy the screening criteria. The average implied

---

[11] The DAX index option is an equity index contract based on the DAX 30 price index of thirty blue chip German stocks. These options have similar specifications to the S&P500 index options. In particular, the DAX options expire on the third Friday of the month, they have European exercise style, and they are cash settled based on a final calculation of the index level.
[12] The settlement price is the last option transaction price of the day, unless this price is more than fifteen minutes before the close of trading or it “does not reasonably reflect market conditions.” In these cases, the Eurex determines the settlement price. These rules are described in the DAX option contract specifications reported on the Eurex website.

volatility across S&P500 contracts is 13.6% with a range from 6.5% to 61.2%, and the average across DAX 30 contracts is 17.2% with a range of 6.5% to 70.8%. S&P500 contracts range in moneyness from –40.0% to 22.1%, and DAX 30 contracts range in moneyness from –40.0% to 20.1%. The average maturity for S&P500 and DAX 30 contracts is about four months.

The last row of the first panel describes the characteristics of the price differential calculated for the S&P500 options. The price differential captures the effect of the early exercise feature and price convergence for an American futures option versus a European futures option. Typically, the price differential as a percentage of the option price is small; it averages .2% of the option price with a standard deviation of .5%. The largest price differential, which is for a deep in-the-money option, is 4.6%.

In addition to the option data, monthly S&P500 returns and DAX 30 returns are obtained for the period from January 1995 through December 1999 from Datastream. The third panel of Table 1 presents the sample moments of these time series. During this period, average annual returns were high for in both markets: 25.9% for the S&P500 and 26.7% for the DAX 30. The annualized S&P500 return volatility was 14.1% and DAX 30 volatility was 21.3%. Both series exhibited negative skewness and excess kurtosis relative to a normal density, and the return correlation over this period was .66.

IV.ii. Estimation of objective densities

The objective S&P500 and DAX 30 joint and marginal densities are estimated nonparametrically using the data and techniques described in the previous sections.[14] The estimated marginal objective densities are graphed in Figure 1, and the moments of these densities are reported in the Panel A of Table 2. The negative skewness of both marginals is apparent in Figure 1. Panel A indicates that the

[13] It is equivalent to use the Black (1976) formula with the DAX 30 futures price as the underlying price or the Black-Scholes (1973) formula with the DAX 30 index level as the underlying price, as long as the cost-of-carry model holds.
[14] All kernel estimation in this paper is implemented using a Gaussian kernel and Silverman (1986, pp. 45-48) bandwidth. The Silverman bandwidth is defined by $h = .9*N^{-1/5}*\min(std, iqr/1.34)$ where h is the bandwidth, N is the number of data points in the series, std is the standard deviation of the series, and iqr is the interquartile range of the series.

skewness and kurtosis of the S&P500 return density are –1.3 and 6.2, and the skewness and kurtosis of the DAX 30 return density are -.6 and 4.4.[15]

The estimated joint density is graphed in Figure 2. It is somewhat difficult to determine the dependence relationship by visual inspection of this figure, so this paper uses correlation and semicorrelation as summary statistics to characterize multivariate dependence. The semicorrelation statistic is proposed by Erb, Harvey, and Viskanta (1994) to measure differences in the linkage between returns during simultaneous up moves or simultaneous down moves. They find evidence for asymmetric correlation in cross-country equity index returns; cross-country returns are much more closely linked during down moves than up moves.

In this paper, the semicorrelation[+,+] statistic is defined as the Pearson correlation coefficient estimated by numerical integration over the first quadrant of returns (positive returns only). The semicorrelation[-,-] statistic is the Pearson correlation coefficient estimated by numerical integration over the third quadrant of returns (negative returns only).

The last column Panel A (Table 2) reports that the correlation between S&P500 and DAX 30 returns, using the estimated objective joint density function, is .57. However, the semicorrelation[+,+] is .64 compared with a semicorrelation[-,-] of .89. Thus, the objective joint density exhibits asymmetric correlation, with a stronger linkage between returns during joint down moves than up moves.

IV.iii. Estimation of risk-neutral densities

The risk-neutral S&P500 and DAX 30 joint and marginal densities are estimated nonparametrically using the data and techniques described in the previous sections. This paper focuses on one-month risk-neutral densities estimated on the last trading day of December 1999.[16]

---

[15] Notice that these moments are slightly different than the sample moments reported in the third panel of Table 1. The Table 2 moments are calculated by numerical integration of the kernel density function estimate, rather than as sample averages.

[16] Evaluation of the nonparametric density functions (equations 3-5) requires a time-horizon to define the state variable T-t (in this case, T-t = 1/12 year or one month) and one-month S&P500 and DAX 30 futures prices to define the state variable F_t. There are no traded S&P500 or DAX 30 futures that expire at the end of January; so estimated futures prices are obtained using the cost-of-carry model for one-month futures as of the last trading day of December. Using the closing S&P500 level (1469.25), the annualized one-month U.S. riskless rate (6.0%), and the dividend yield (1.1%), the estimated one-month S&P futures price is F_t=1475.26. Using DAX 30 level (6958.14), the annualized one-month German riskless rate (3.2%), and the dividend yield (0%, due to reinvestment of dividends in the index), the estimated one-month DAX 30 futures price is F_t = 6976.55.

Figure 3 graphs the nonparametrically estimated one-month “implied volatility functions” against proportional moneyness, as of the last trading day of December 1999. Both functions exhibit pattern of decreasing implied volatility as a function of moneyness. This result is consistent with Ait-Sahalia and Lo (1998), and it indicates that the implied risk-neutral density has higher probability of large negative events than a lognormal density.

Figures 4 and 5 graph the estimated nonparametric marginal densities. For comparison, plots of lognormal densities that have identical mean and standard deviation as the nonparametric estimate are included. These figures illustrate that the implied risk-neutral densities have higher probabilities for large negative returns than a comparable lognormal density and lower probabilities for large positive returns than a lognormal.

These results are confirmed by comparison of moments of the risk-neutral marginal densities, as reported in the second two columns of Panel B of Table 2. The nonparametric S&P500 and DAX 30 risk-neutral densities have negative skewness (-.6 and -.4), while the lognormal densities have positive skewness (.2 and .3). Also, the nonparametric S&P500 and DAX 30 densities have greater kurtosis (3.8 and 3.4) than their lognormal counterparts (3.1 and 3.1). These results illustrate two types of error caused by use of a lognormal specification for the risk-neutral marginals.

To construct the joint risk-neutral density function, it is necessary to specify the marginal densities and the copula. This paper considers the four bivariate risk-neutral densities generated by the choice of lognormal or nonparametric marginals combined with the choice of a lognormal or nonparametric copula. These comparison densities are constructed to separately identify the effects of misspecification of the marginals or copula on the joint density.

In order to justify the use of the objective copula as an estimate of the risk-neutral copula, it is necessary to verify that the risk-neutral random returns are increasing functions of the objective returns. Using equation (8) and expectations under the nonparametric objective and risk-neutral measures, the slope of the function that transforms the objective return into the risk-neutral return is estimated for five return intervals.

Return interval	$dX^*(X)/dX$	$dY^*(Y)/dY$
[-20%,10%]	3.85	2.58
[10%,0%]	2.90	1.85
[0%, 10%]	0.99	0.78
[10%, 20%]	13.85	1.63

The derivatives are positive in all cases, indicating that the conditions of Theorem 4 are satisfied. Hence, the objective copula — obtained by the method of inversion applied to the nonparametric joint objective density — may be used to estimate the risk-neutral copula.

The lognormal copula is defined using the method of inversion as $Cop(u,v) = ProbBN(\Theta^{-1}(u), \Theta^{-1}(v))$, where $ProbBN(\cdot,\cdot)$ is the cumulative standard bivariate normal density function, $\Theta^{-1}(\cdot)$ is the standard normal quantile (inverse standard normal cumulative density) function, and u and v are cumulative probabilities. Notice that the copulas for normal and lognormal variates are identical, since lognormal variates are obtained by an increasing function of normal variates. The correlation coefficient in the cumulative standard bivariate normal density function is set equal to the nonparametric objective correlation estimate (.57).

Figures 6 – 9 graph the joint risk-neutral densities obtained by combining nonparametric or lognormal risk-neutral marginals with a nonparametric or risk-neutral copula. Figure 6 presents the risk-neutral density (a joint lognormal density) obtained by combining lognormal marginals with a lognormal copula. The joint lognormal is unimodal with a narrow hump that corresponds to a relatively high positive risk-neutral correlation between S&P500 and DAX returns.

Figure 7 generalizes the lognormal density to allow the dependence structure to reflect the empirical data, but preserves the lognormality of the marginals. While this density exhibits positive correlation, there appears to be significant asymmetry in the density with the large joint negative returns more tightly clustered than the large joint positive returns.

Figure 8 generalizes the lognormal density to allow the marginals to reflect the empirical data, but preserves the lognormality of the dependence structure. This is a joint density based on nonparametric marginals and a lognormal copula. Figure 8 is more similar to a joint lognormal than Figure 7, but the negative skewness and excess kurtosis of the marginals result in a larger amount of density weight in the region of joint negative returns than for a joint lognormal.

Figure 9 completely generalizes the lognormal density, so that both the marginals and the dependence structure reflect the empirical data. This is a fully-nonparametric joint density based on nonparametric marginals and a nonparametric copula. It has the greater density weight in the joint negative return region of Figure 8, and the tighter clustering of large joint negative returns of Figure 7.

An alternative comparison of these joint densities, based on joint moments, is presented in Panel C of Table 2. Both densities that use a lognormal copula (Columns 1 and 3) have correlations (.57

and.56) close to that of the objective density (.57). Neither of these densities exhibits substantial differences in semicorrelation[+,+] compared with semicorrelation[-,-]. These results indicate that a lognormal copula imposes nearly symmetrical correlation on the risk-neutral density, and preserves the correlation coefficient specified in the lognormal copula.

The joint densities constructed using the nonparametric copula (Columns 2 and 4) inherit the asymmetric correlation of the objective density. The semicorrelation[+,+] for these densities (.60 and .63) is close to that of the objective density (.64), and the semicorrelation[-,-] for these densities (.85 and .85) is close to that of the objective density (.89). These densities have lower correlation (.46 and .52) than the objective correlation (.56).

IV.iv. Estimation of underperformance and outperformance option prices

To measure the importance of nonparametric estimation of the risk-neutral density function on MVCC prices, this section compares estimated underperformance and outperformance option prices using lognormal and nonparametric bivariate risk-neutral densities. A joint lognormal risk-neutral density is an implication of standard MVCC pricing formulas.

Consider a bivariate underperformance option that pays in dollars the lesser of the one-month S&P500 index return ($r_X = X_T/X_t - 1$) in percent and the fully-hedged DAX 30 index return ($r_Y = Y_T/Y_t - 1$) in percent, with a lowest possible payoff of zero. Applying the equation (2) pricing formula using the appropriate payoff function, the price of such an option is:[17]

$$D_{r_X, r_Y, t} = e^{-r(T-t)} \iint 100 Max[0, Min(r_X, r_Y)] f^*_{r_X, r_Y}(r_X, r_Y) dr_X dr_Y \tag{12}$$

A bivariate outperformance option pays the better of the percentage return on the S&P500 index and the fully-hedged percentage return on the DAX 30 index, with a lowest possible payoff of zero. The pricing formula for this type of option is:

$$D_{r_X, r_Y, t} = e^{-r(T-t)} \iint 100 Max[0, r_X, r_Y] f^*_{r_X, r_Y}(r_X, r_Y) dr_X dr_Y \tag{13}$$

---

[17] Stulz (1982) and Johnson (1987) derive closed-form pricing formulas for pricing options on minimum and maximum asset prices. These formulas, which are based on a multivariate geometric Brownian motion for the asset prices, cannot be directly applied to options on minimum or maximum returns.

Inspection of these pricing formulas indicates the importance of correct specification of the risk-neutral density. To the extent that a lognormal density incorrectly specifies the characteristics of the actual risk-neutral density, fitted prices based on the lognormal density will deviate from the correct values. The types of biases generated by pricing under a misspecified lognormal risk-neutral density will depend on the types of error in the density estimate (e.g. underestimation of negative skewness or positive kurtosis of the marginals, imposition of symmetry on the dependence function) and the particular payoff function of the option.

Underperformance and outperformance option prices are particularly sensitive to the dependence structure (especially the correlation) of the risk-neutral density. An underperformance option only pays off if both returns are positive (quadrant I, see Figure 10). In the lognormal case, the risk-neutral probability that both returns are positive is increasing as the risk-neutral return correlation increases. Hence, the value of an underperformance option is increasing as the return correlation increases.[18]

For an underperformance option, only dependence over quadrant I (e.g. semicorrelation[+,+]) is relevant for pricing, since the payoff in all other return regions is zero. A misspecified risk-neutral density that overestimates (underestimates) positive dependence in this region will tend to overprice (underprice) underperformance options.

An outperformance option pays off as long at least one of the returns is positive (quadrants I, II, IV, see Figure 11). In the lognormal case, the risk-neutral probability that both returns are negative (and the payoff is zero) is increasing as correlation increases. In fact, when the asset return correlation is equal to one, the diversification benefit is lost, and an outperformance option is equivalent to an option on a single asset return with an exercise price of zero. Thus, the value of an outperformance option is decreasing as the return correlation increases.

For an outperformance option, dependence over all quadrants except quadrant III is relevant for the pricing, since the payoff in quadrant III is zero. A misspecified risk-neutral density that overestimates (underestimates) positive dependence in these regions will tend to underprice (overprice) outperformance options.

Panel D of Table 2 presents a pricing comparison of one-month underperformance and outperformance options on the S&P500 and DAX 30 indices, valued on the last trading day of

December 1999. The four risk-neutral densities analyzed in the previous section are used for valuation. Column 4 reports that the nonparametric value of an underperformance option is $1.52 and the value of an outperformance option is $4.86. Using these nonparametric prices as a benchmark, we can separately examine the importance of correct specification of the marginals and the copula, as well as the relationship between semicorrelation and pricing biases.

For example, column 3 reports the value of underperformance and outperformance options when the marginals are nonparametric but the copula is lognormal ($1.76 and $4.62). This density has correct marginals but incorrect dependence structure. The misspecified density overestimates the underperformance option value by about 16% and underestimates of the outperformance option value by about 5%. These are exactly the biases that would be expected when the misspecified semicorrelation[+,+] is higher than the true semicorrelation[+,+], which is the case for these densities (.71 versus .63).

An analysis of underperformance option prices using the three misspecified density functions indicates that that misspecification of the copula results in greater pricing error than misspecification of the marginals. Replacing the lognormal with a nonparametric copula (while retaining lognormal marginals) results in a substantial decrease in pricing error (from 16% to 8%). Replacing parametric with nonparametric marginals (while retaining a lognormal copula) also reduces pricing error, but not by as much (from 16% to 10%).

An analysis of outperformance option prices using the three misspecified density functions shows that these prices are much less sensitive to density misspecification. To some extent, outperformance option pricing accuracy appears to depend more on correct specification of the marginals than correct specification of the copula. Replacing the lognormal with a nonparametric copula (while retaining lognormal marginals) results in an increase in pricing error (from 3% to 5%). Replacing parametric with nonparametric marginals (while retaining a lognormal copula) slightly reduces pricing error (from 3% to 2%).

The observed pricing errors may be attributed to misspecification of the marginal densities and misspecification of the dependence function. The nonparametrically estimated risk-neutral marginal densities exhibit negative skewness and excess kurtosis that is not captured in a lognormal specification. The nonparametrically estimated dependence function exhibits asymmetric correlation that is not reflected in the lognormal copula specification.

---

[18] Stulz (1982) shows that the value of an option on the minimum asset price is increasing in correlation, while the value

In all cases, the pricing biases are exactly what would be expected based on an analysis of semicorrelation[+,+]. When the misspecified risk-neutral density overestimates (underestimates) semicorrelation[+,+], there is overpricing (underpricing) of underperformance options and underpricing (overpricing) of outperformance options. It is clear that correct specification of the marginals and dependence function that define the multivariate risk-neutral density function is essential for accurate MVCC pricing.

V. Conclusion

This paper has developed a nonparametric technique for multivariate contingent claim valuation with several attractive features. First, this technique does not impose restrictions on the asset return process or on the functional form of the risk-neutral density. The characteristics of the risk-neutral density are estimated using market data. Second, MVCC prices obtained using this procedure are consistent with the market prices of existing options; and, the marginal risk-neutral densities are identical to the marginals used for risk-neutral univariate contingent claim pricing. Third, the estimation methodology does not require data on traded multivariate claims, since the dependence function is estimated using historical returns data.

The nonparametric pricing technique is applied to value underperformance and outperformance options on the DAX 30 and S&P500 indices. Nonparametric prices are compared with prices based on a joint lognormal risk-neutral density, which is an implication of standard MVCC pricing formulas. The results of the pricing comparisons indicate that a bivariate lognormal density is an inadequate approximation to the bivariate risk-neutral density implied by market data. Important features of the estimated risk-neutral density that are not reflected in a joint lognormal density include negative skewness and excess kurtosis in the marginals and asymmetric correlation in the dependence function. Accurate MVCC pricing requires correct specification of both the marginal densities and the dependence function that define the multivariate risk-neutral density function.

of an option on the maximum asset price is decreasing in correlation.

Appendix A

Proof of Theorem 1

$$C(K) = e^{-r(T-t)}\int_K^{\infty}(X_T - K)\, f^*_{X,t}(X_T)dX_T \Rightarrow \partial C/\partial K = e^{-r(T-t)}\int_K^{\infty} - f^*_{X,t}(X_T)dX_T = -e^{-r(T-t)}(1 - F^*_{X,t}(K))$$

$$F^*_{X,t}(K) = e^{r(T-t)}\partial C/\partial K + 1$$

$$f^*_{X,t}(K) = e^{r(T-t)}\partial^2 C/\partial K^2$$

Proof of Theorem 2

$$C(K,L) = e^{-r(T-t)}\int_K^{\infty}\int_L^{\infty}(X_T - K)(Y_T - L)\, f^*_{X,Y,t}(X_T, Y_T)dX_T dY_T \Rightarrow$$

$$\partial C/\partial K = e^{-r(T-t)}\int_K^{\infty}\int_L^{\infty} -(Y_T - L)\, f^*_{X,Y,t}(X_T, Y_T)dX_T dY_T \Rightarrow$$

$$\partial^2 C/\partial K \partial L = e^{-r(T-t)}\int_K^{\infty}\int_L^{\infty} f^*_{X,Y,t}(X_T, Y_T)dX_T dY_T = e^{-r(T-t)}(1 - F^*_{X,Y,t}(K,L))$$

$$F^*_{X,Y,t}(K,L) = \text{Prob}(X_T < K, Y_T < L) = -e^{r(T-t)}\partial^2 C/\partial K \partial L + 1$$

$$f^*_{X,Y,t}(K,L) = -e^{r(T-t)}\partial^2 F^* / \partial K \partial L$$

Proof of Theorem 5

Suppose that $F_{XY}(X_T,Y_T)$ is a joint density function with marginal quantile functions (inverse marginal cumulative density functions) given by $F_X^{-1}(u)$ and $F_Y^{-1}(v)$ where u and v are cumulative probabilities. Using Sklar's Theorem, $F_{XY}(X_T,Y_T) = Cop(F_X(X_T),F_Y(Y_T))$ where Cop(u,v) is a copula. $Cop(u,v) = F_{XY}(F_X^{-1}(u), F_Y^{-1}(v))$ is the copula of F_{XY}, since $Cop(F_X^{-1}(F_X(X_T)), F_Y^{-1}(F_Y(Y_T)) = F_{XY}(X_T,Y_T)$.

Appendix B

The BBSR algorithm (Broadie and Detemple (1996), pp. 1243-1245) is an enhancement to the Cox, Ross, Rubinstein (1979) binomial tree approach for pricing American or European options on the underlying asset (e.g. options on an equity index). Broadie and Detemple (1996) find the BBSR algorithm offers significant improvements in pricing accuracy over alternative algorithms.

The BBSR algorithm modifies the CRR algorithm in two ways. First, the option price tree is modified so that the continuation values at the nodes just prior to expiration are replaced with values from the Black-Scholes formula. This is referred to as the BBS algorithm (Binomial-Black-Scholes). Second, the option price is calculated using the desired number of time steps and one-half the number of time steps using the BBS algorithm. The BBSR (Binomial-Black-Scholes-Richardson) price is calculated using the Richardson extrapolation, which sets the BBSR price equal to twice the BBS price estimated using the desired number of time-steps minus the BBS price using one-half the number of time-steps. In this implementation of the algorithm, the number of time-steps is set to 100, which offers maximum precision and is the maximum number tested in Broadie and Detemple (1996).

The BBSR algorithm is designed to price options on the underlying asset, so several minor adjustments must be made to the BBSR algorithm so that it is appropriate for pricing futures options. To price futures options, the up and down parameters must reflect futures returns (u_F, d_F) rather than the underlying asset returns (u and d). The appropriate formulas are $u_F = u*e^{-c\Delta}$ and $d_F = d*e^{-c\Delta}$, where Δ is the amount of time represented by one-time-step, and u and d are the parameters from the BBSR algorithm.

These parameters are derived as follows. The cost-of-carry model states that $F_t = e^{c(T-t)}X_t$, where F_t is the futures price, X_t is the spot price, T-t is the number of years until the futures contract expires, and c is the continuously compounded annual cost of carry. For equity index options, the cost of carry equals the difference between the continuously compounded annual riskless rate and the dividend yield.

$$u_F = F_u / F_t = e^{c(T-t-\Delta)}X_u / e^{c(T-t)}X_t = u*e^{-c\Delta}$$

$$d_F = F_d / F_t = e^{c(T-t-\Delta)}X_d / e^{c(T-t)}X_t = d*e^{-c\Delta}$$

To price futures options, we can also write the risk-neutral probabilities (p, 1-p) in terms of the futures returns rather than the underlying returns to obtain p_F and $1-p_F$. The appropriate formula is $p_F = (1-d_F)/(u_F - d_F)$. This parameter is derived as follows. The BBSR algorithm defines p as $p= (e^{c\Delta} - d) / (u - d)$. Using the definitions above, $u=u_F e^{c\Delta}$ and $d=d_F e^{c\Delta}$. Substituting into p to define p_F, we have $p_F = (1-d_F)/(u_F - d_F)$.

The following data is used to implement the BBSR algorithm. The cost-of-carry (c) is calculated as the difference between the riskless rate and S&P500 dividend yield. The riskless rate corresponding to the option maturity is extracted from the daily term structure of BBA LIBOR rates as reported in the Datastream database. The S&P500 dividend yield is the predicted one-year dividend yield for the S&P500 as reported by Standard and Poor's (1.14% based on the 1999 closing level of the S&P500). The volatility parameter (σ), which is used to calculate u_F and d_F, is the annualized historical standard deviation of monthly S&P500 returns over previous five years (14.07%).

Bibliography

Ait-Sahalia, Y., Lo, A.W., 1998. Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices. Journal of Finance 53, 499-547.

Bates, D., 1995. Hedging the Smile. Manuscript.

Black, F., 1976. The Pricing of Commodity Contracts. Journal of Financial Economics 3, 167-179.

Black, F., Scholes, M., 1973. The Pricing of Options and Corporate Liabilities. Journal of Political Economy 81, 637-654.

Boyle, P.P., 1988. A Lattice Framework for Option Pricing with Two State Variables. Journal of Financial and Quantitative Analysis 23, 1-12.

Boyle, P.P., Evnine, J., Gibbs, S., 1989. Numerical Evaluation of Multivariate Contingent Claims. Review of Financial Studies 2, 241-250.

Breeden, D.T., Litzenberger, R.H., 1978. Prices of State-Contingent Claims Implicit in Option Prices. Journal of Business 51, 621-651.

Brenner, M., Courtadon, G., Subrahmanyam, M., 1985. Options on the Spot and Options on Futures. Journal of Finance 40, 1303-1317.

Broadie, M., Detemple, J., 1996. American Option Valuation: New Bounds, Approximations, and a Comparison of Existing Methods. Review of Financial Studies 9, 1211 - 1250.

Cox, J.C., Ross, S.A., Rubinstein, M., 1979. Option Pricing: A Simplified Approach. Journal of Financial Economics 7, 229-263.

Deheuvels, P., 1979. La fonction de dependence empirique et ses proprietes: Un test non parametrique d'independence. Acad. Roy. Belg. Bull. Cl. Sci. 5, 274-292.

Derman, E., Kani, I., 1994. Riding on a Smile. Risk 7, 32-38.

Dumas, B., Fleming, J., Whaley, R.E., 1998. Implied Volatility Functions: Empirical Tests. Journal of Finance , 2059-2106.

Dupire, B., 1994. Pricing with a Smile. Risk 7, 18-20.

Embrechts, P., McNeil, A., Straumann, D., 1998. Correlation and Dependency in Risk Management: Properties and Pitfalls. Manuscript.

Engle, R.F., Kroner, K.F., 1995. Multivariate Simultaneous Generalized ARCH. Econometric Theory 11, 122-150.

Erb, C.B., Harvey, C.R., Viskanta, T.E., 1994. Forecasting International Equity Correlations. Financial Analysts Journal 50, 32-45.

Harvey, A., Ruiz, E., Shephard, N., 1994. Multivariate Stochastic Variance Models. Review Of Economic Studies 61, 247-264.

Ho, T., Stapleton, R., Subrahmanyam, M., 1995. Multivariate Binomial Approximations for Asset Prices with Nonstationary Variance and Covariance Characteristics. Review of Financial Studies 8, 1125-1152.

Hull, J., White, A., 1998. Value at Risk when Daily Changes in Market Variables are not Normally Distributed. Journal of Derivatives , 9-19.

Johnson, H., 1987. Options on the Maximum or Minimum of Several Assets. Journal of Financial and Quantitative Analysis 22, 277-283.

Longstaff, F.A., 1995. Option Pricing and the Martingale Restriction. Review of Financial Studies 8, 1091-1124.

Margrabe, W., 1978. The Value of an Option to Exchange One Asset for Another. Journal of Finance 33, 177-186.

Nelsen, R., 1998. An Introduction to Copulas. Springer-Verlag, New York.

Reiner, E., 1992. Quanto Mechanics. In From Black-Scholes to Black-Holes: New Frontiers in Option Pricing. Quanto Mechanics147-154. Risk Magazine/FINEX, London.

Rosenberg, J., 1998. Pricing Multivariate Contingent Claims using Estimated Risk-Neutral Density Functions. Journal of International Money and Finance 17, 229-247.

Rosenberg, J., 1999. Semiparametric Pricing of Multivariate Contingent Claims. Manuscript.

Rubinstein, M., 1992. One for Another. In From Black-Scholes to Black-Holes: New Frontiers in Option Pricing. One for Another147-154. Risk Magazine/FINEX, London.

Rubinstein, M., 1994a. Implied Binomial Trees. Journal of Finance 49, 771-818.

Rubinstein, M., 1994b. Return to Oz. Risk 7, 67-71.

Rubinstein, M., 1991. Somewhere Over the Rainbow. Risk , 63-66.

Sherrick, B.J., Irwin, S.H., Forster, D.L., 1990. Expected Soybean Futures Price Distributions: Option-based Assessments. Review of Futures Markets 12, 275-290.

Sherrick, B.J., Irwin, S.H., Forster, D.L., 1992. Option-Based Evidence of the Nonstationarity of Expected S&P 500 Futures Price Distributions. Journal of Futures Markets 12, 275-290.

Shimko, D.C., 1993. Bounds of Probability. Risk 6, 33-37.

Shimko, D.C., 1994. Options on Futures Spreads - Hedging, Speculation, and Valuation. Journal of Futures Markets 14, 183-213.

Silverman, B.W., 1986. Density Estimation for Statistics and Data Analysis. Chapman and Hall, London.

Stapleton, R.C., Subrahmanyam, M.G., 1984. The Valuation of Multivariate Contingent Claims in Discrete Time Models. Journal of Finance 39, 207-228.

Stapleton, R.C., Subrahmanyam, M.G., 1984. The Valuation of Options when Asset Returns are Generated by a Binomial Process. Journal of Finance 39, 1525-1539.

Stulz, R.M., 1982. Options on the Minimum or the Maximum of Two Risky Assets: Analysis and Applications. Journal of Financial Economics 10, 161-185.

Tawn, J., 1990. Modelling Multivariate Extreme Value Distributions. Biometrika 77, 245-253.

Table 1
Data description

S&P500 futures option data
Daily settlements, December 1999

	N	Mean	Std. dev.	Minimum	Maximum
Implied volatility (annual %)	711	23.89%	6.46%	13.56%	61.21%
Proportional moneyness (%)	711	-7.42%	14.67%	-39.99%	22.14%
Time until expiration (years)	711	0.36	0.22	0.04	0.99
Trading volume (contracts)	711	64.42	115.49	5	1030
Option price (in dollars)	711	41.30	56.07	0.10	478.70
Early exercise premium (% of option price)	711	0.21%	0.47%	0.00%	4.64%

DAX 30 option data
Daily settlements, December 1999

	N	Mean	Std. dev.	Minimum	Maximum
Implied volatility (annual %)	1680	29.77%	6.51%	17.16%	70.82%
Proportional moneyness (%)	1680	-9.94%	13.29%	-39.96%	20.05%
Time until expiration (years)	1680	0.36	0.17	0.04	0.74
Trading volume (contracts)	1680	519.42	998.59	5	13902
Option price (in deustchemarks)	1680	374.58	411.59	0.30	2289.70

Returns data
Monthly, 1995.01 -1999.12

	S&P500	DAX 30
N	60	60
Mean (annualized)	25.90%	26.70%
Std. dev. (annualized)	14.07%	21.31%
Skewness	-1.56	-0.74
Kurtosis	7.30	4.82

*Return correlation: (S&P500,DAX 30) = .6573

This table reports characteristics of the option and returns data used in the paper. The option data covers all traded contracts for the month of December 1999, subject to exclusion criteria for illiquid contracts, as detailed in the text of the paper. The returns data covers the five-year period from 1995 through 1999. Proportional moneyness is defined as the natural log of the ratio of the option exercise price and the futures price.

The S&P500 options data is extracted from a futures options database distributed by the Chicago Mercantile Exchange that reports settlement prices and contract terms for all traded contracts. For each futures option contract, the corresponding futures contract settlement price also obtained from a CME database. Option implied volatilities are calculated by numerically inverting the Black (1976) formula with an adjustment for the early-exercise premium, using a dollar-denominated riskless rate with identical maturity as the option contract extracted from the contemporaneous LIBOR term structure.

The DAX 30 options data is extracted from an options and futures database distributed by the Deutsche Borse (the German option and futures exchange, which is now known as EUREX) that reports settlement prices and contract terms for all traded contracts. Option implied volatilities are calculated by numerically inverting the Black-Scholes (1976) formula, using a German Mark denominated riskless rate with identical maturity as the option contract extracted from the contemporaneous FIBOR (Frankfurt Interbank Offer Rate) term structure.

The S&P500 return and DAX 30 return (in local terms) are calculated as the daily proportional change in the S&P500 index as reported in Datastream database. Moments are for the raw data are calculated using the standard formulas. Kurtosis is reported in total (rather than excess terms).

Table 2

Density estimation and pricing

Panel A: Characteristics of the objective density (kernel estimate)

	Marginal densities	
	S&P500	DAX 30
Mean	27.89%	28.70%
Standard deviation	15.03%	22.76%
Skewness	-1.25	-0.60
Kurtosis	6.19	4.35

Joint density	
S&P500,DAX 30	
Correlation	0.5665
Covariance	0.0016
Semicorrelation[+,+]	0.6375
Semicorrelation[-,-]	0.8897

Panel B: Characteristics of the marginal risk-neutral densities

	S&P 500 risk-neutral density		DAX 30 risk-neutral density	
S&P500 one-month return	*Lognormal*	*Nonparametric*	*Lognormal*	*Nonparametric*
Mean	7.30%	7.30%	4.28%	4.28%
Standard deviation	22.34%	22.34%	29.82%	29.82%
Skewness	0.19	-0.58	0.26	-0.36
Kurtosis	3.07	3.80	3.11	3.38

Panel C: Characteristics of the joint risk-neutral density

Marginals	Lognormal		Nonparametric	
Copula	*Lognormal*	*Nonparametric*	*Lognormal*	*Nonparametric*
Correlation	0.5656	0.4646	0.5627	0.5192
Covariance	0.0031	0.0026	0.0031	0.0029
Semicorrelation[+,+]	0.6933	0.6016	0.7148	0.6322
Semicorrelation[-,-]	0.7151	0.8466	0.6723	0.8548

Panel D: Comparison of bivariate option prices

Marginals	Lognormal		Nonparametric	
Copula	*Lognormal*	*Nonparametric*	*Lognormal*	*Nonparametric*
Underperformance option value	$1.68	$1.40	$1.76	$1.52
Outperformance option value	$4.75	$5.03	$4.62	$4.86

This table reports characteristics of the objective density (Panel A), risk-neutral density (Panels B & C), and bivariate option prices estimated using the risk-neutral densities (Panel D). Two types of marginal densities (lognormal and nonparametric) are paired with two types of copulas (lognormal and nonparametric) to obtain four bivariate risk-neutral densities. These densities are estimated using the data described in Table 1.

Marginal and joint moments are obtained by numerical integration of the four joint density functions. Semicorrelation[+,+] is defined as the Pearson correlation coefficient where the integration range is exclusively over positive returns. Semicorrelation[-,-] is defined similarly over negative returns. The payoff function for the underperformance option is $100*Max[0,Min[0, X_T,Y_T]]$, and the payoff for the outperformance option is $100*Max[0,X_T,Y_T]$ where X_T, and Y_T are one-month returns on the S&P500 and DAX 30 indices (in local terms). The options are valued on the last trading day of December 1999.

Figure 1
Marginal objective return density functions
S&P500 and DAX 30 one-month returns, 1995.1 - 1999.12

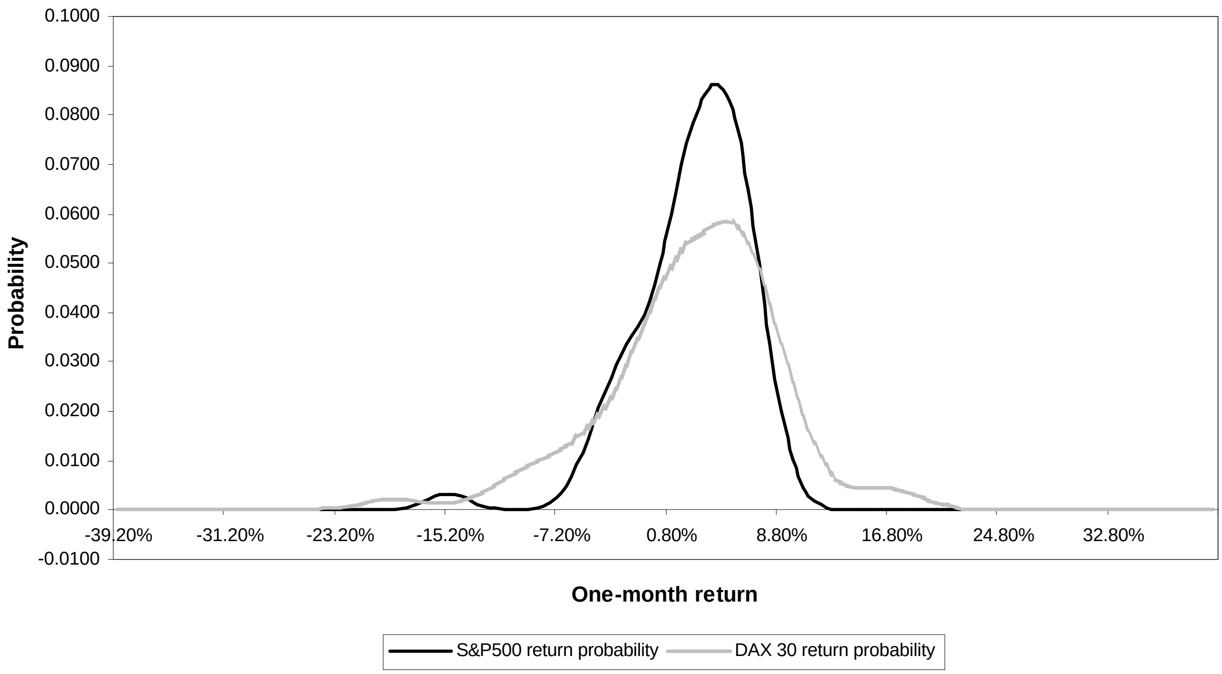

Figure 2
Joint objective return density function
S&P500 and DAX 30 one-month returns
1995.1 - 1999.12

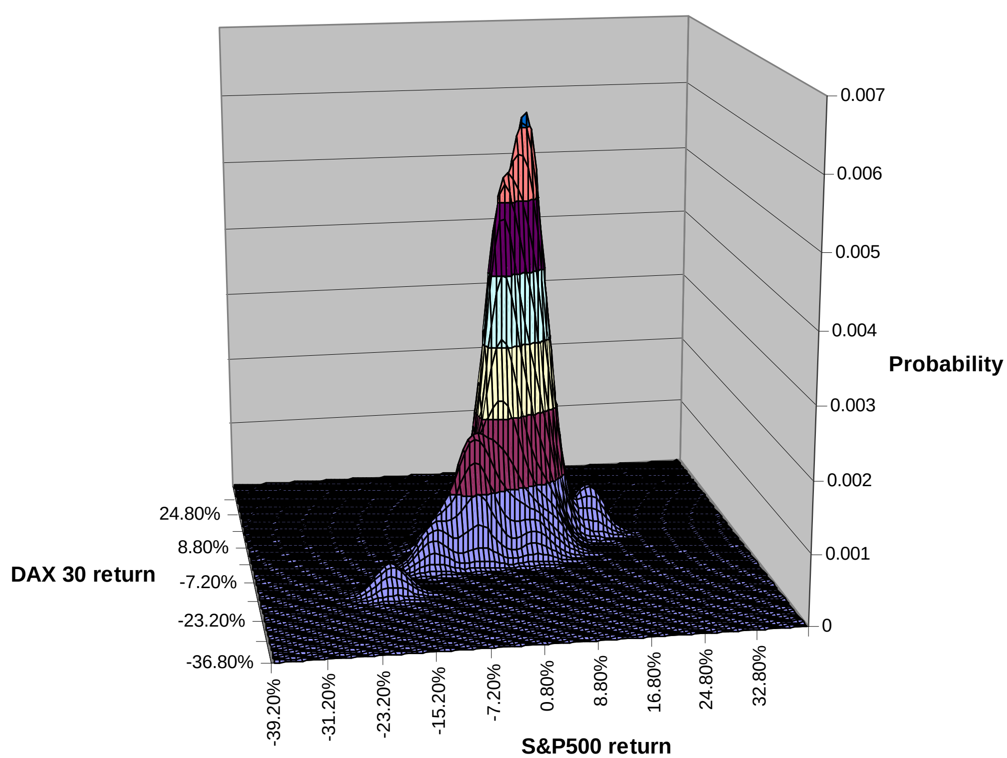

Figure 3
Fitted one-month implied volatility functions
S&P500 and DAX 30 options
December 1999

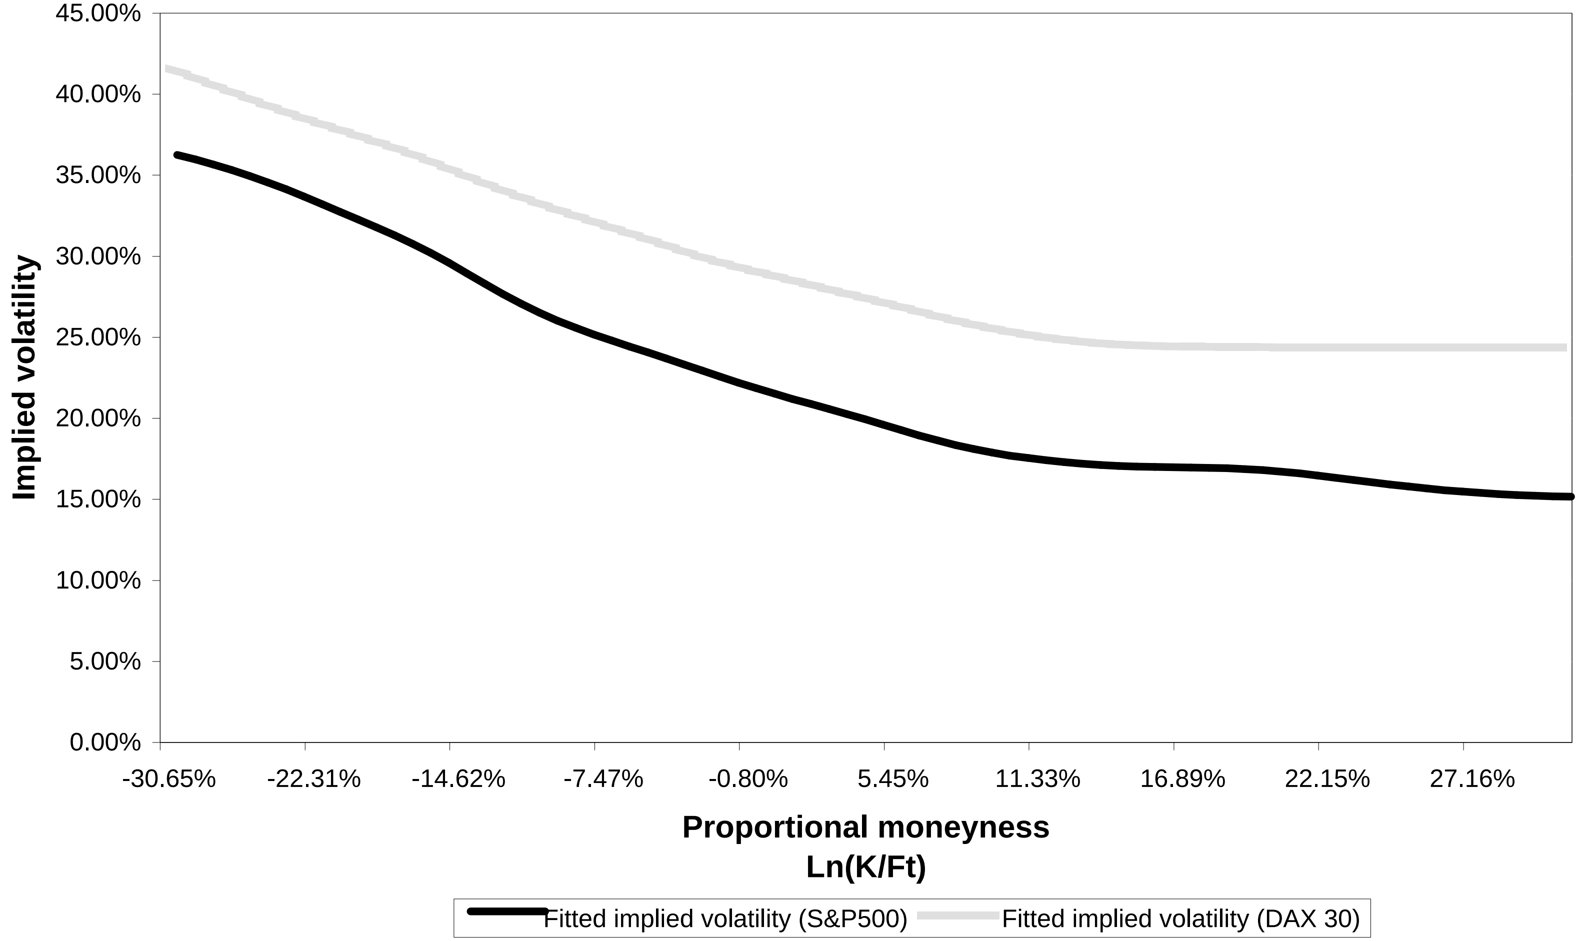

Figure 3 plots the estimated implied volatility functions for S&P500 and DAX 30 options. These functions are estimated using a kernel regression of option-specific implied volatilities over the month of December 1999 on the corresponding underlying price, strike price, and time until expiration. The fitted functions are obtained using the end of December 1999 underlying price and a maturity of one-month. Implied volatilities are reported in annualized percentage terms.

Figure 4
Marginal risk-neutral density
S&P500 one-month returns, December 31, 1999

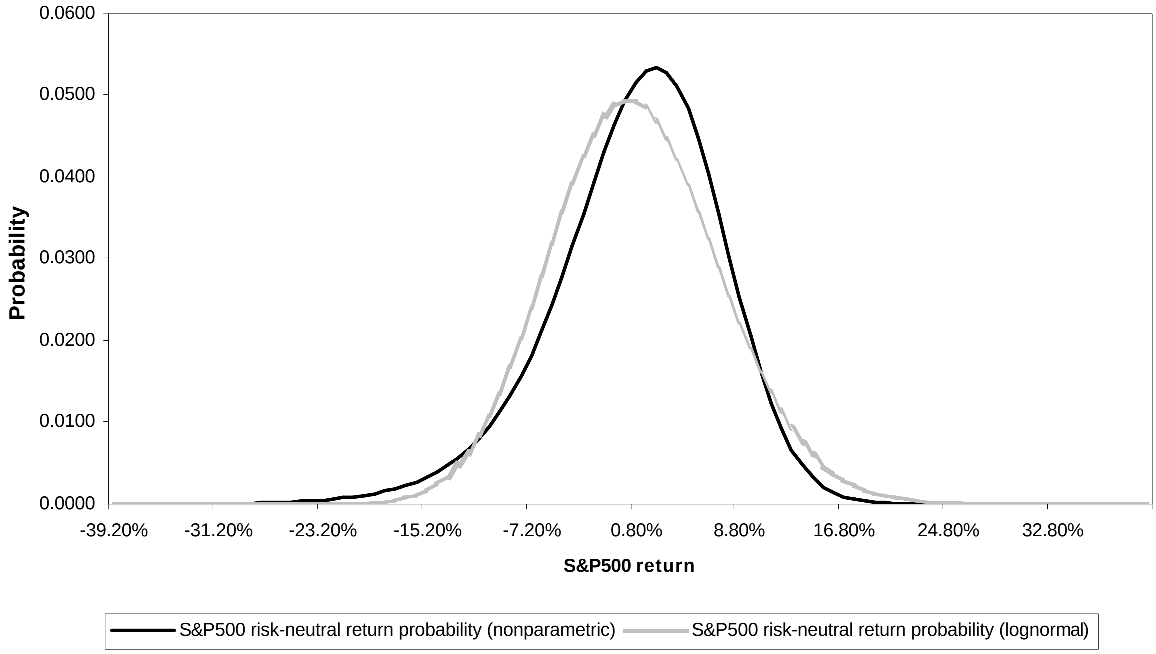

Figure 5
Marginal risk-neutral density
DAX 30 one-month returns, December 31, 1999

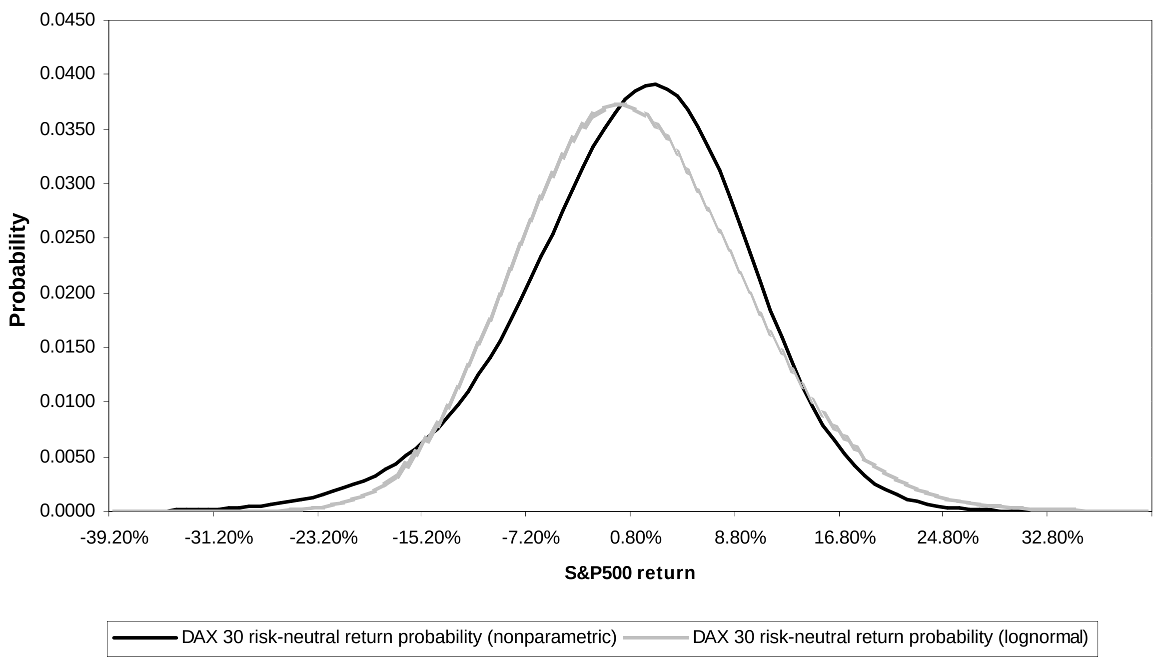

Figure 6
Joint risk-neutral return density
Lognormal marginals and lognormal copula
One-month S&P500 and DAX 30 returns, December 31, 1999

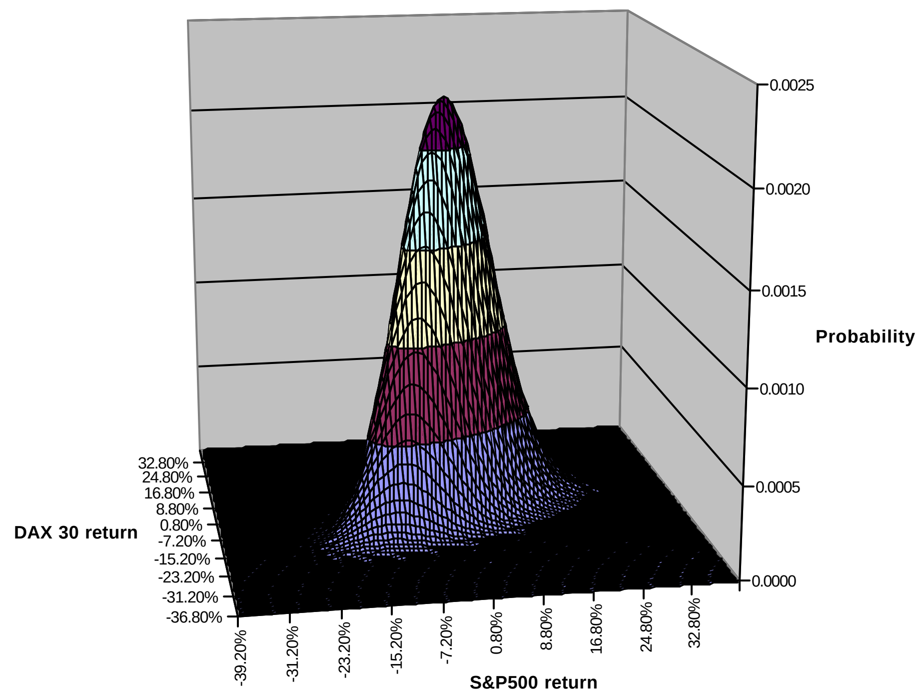

Figure 7
Joint risk-neutral return density
Lognormal marginals and nonparametric copula
One-month S&P500 and DAX 30 returns, December 31, 1999

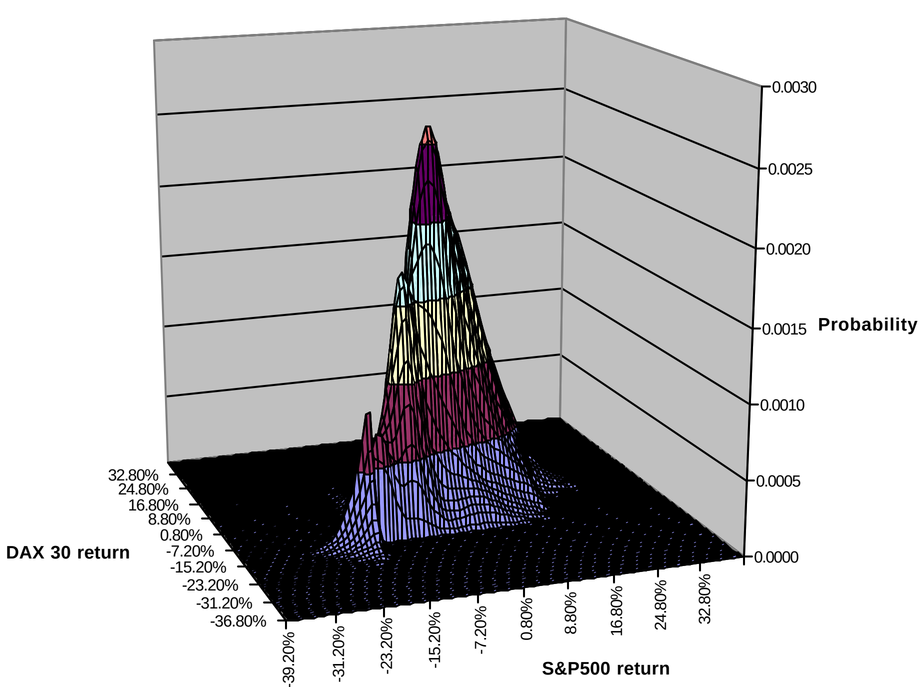

Figure 8
Joint risk-neutral return density
Nonparametric marginals and lognormal copula
One-month S&P500 and DAX 30 returns, December 31, 1999

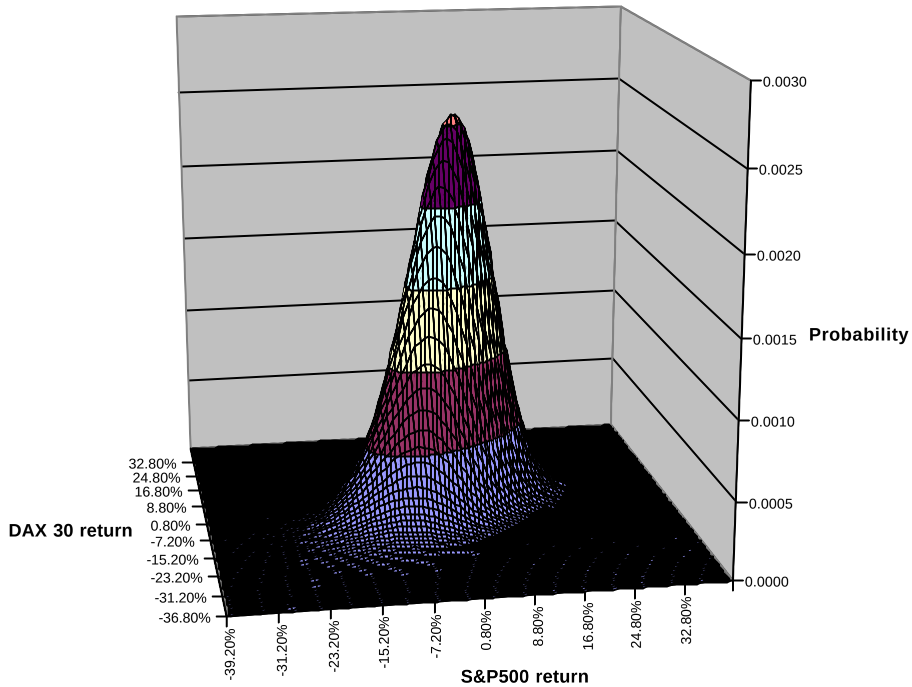

Figure 9
Joint risk-neutral return density
Nonparametric marginals and nonparametric copula
One-month S&P500 and DAX 30 returns, December 31, 1999

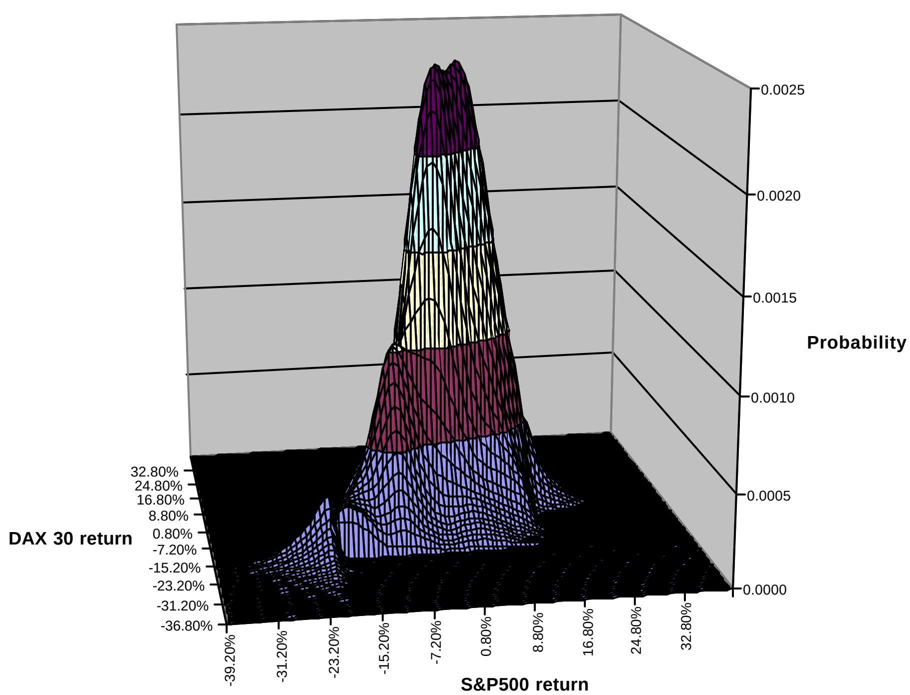

Figure 10
Payoff surface for an underperformance option

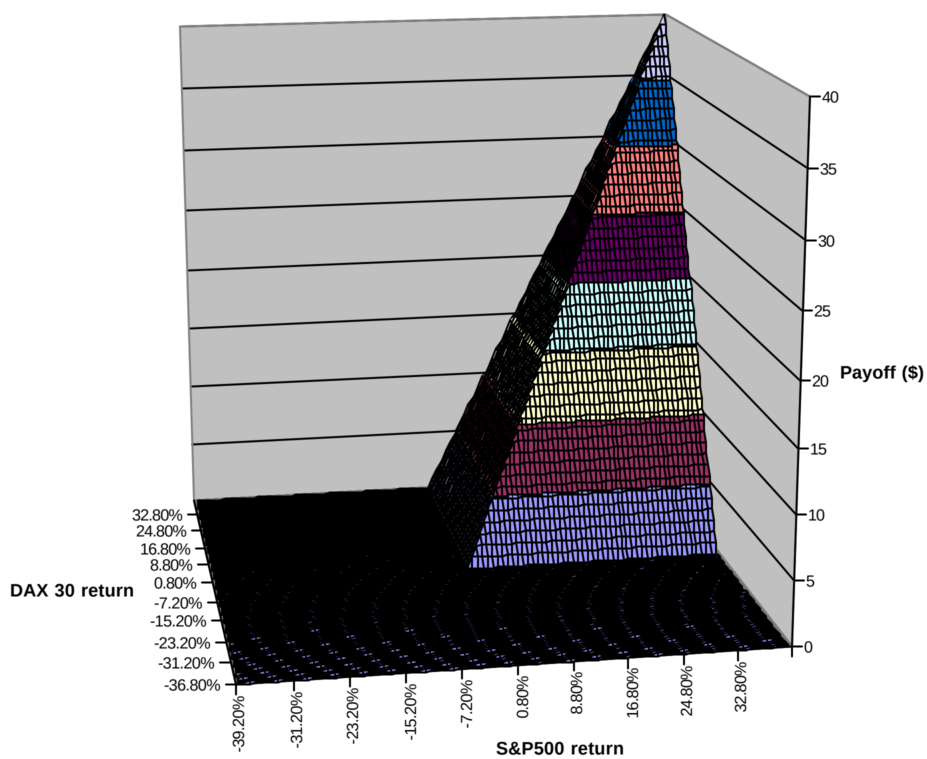

Figure 11
Payoff surface for an outperformance option