Semiparametric pricing of multivariate contingent claims

August 1999

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

Abstract

This paper develops and implements a methodology for pricing multivariate contingent claims (MVCC’s) based on semiparametric estimation of the multivariate risk-neutral density function. This methodology generates MVCC prices which are consistent with current market prices of univariate contingent claims.

This method allows for completely general marginal risk-neutral densities and is compatible with all univariate risk-neutral density estimation techniques. The univariate risk-neutral densities are related by their risk-neutral correlation, which is estimated using time-series data on asset returns and an empirical pricing kernel (Rosenberg and Engle, 1999). This permits the multivariate risk-neutral density to be identified without requiring observation of multivariate contingent claims prices.

The semiparametric MVCC pricing technique is used for valuation of one-month options on the better of two equity index returns. Option contracts with payoffs dependent on are four equity index pairs are considered: S&P500 - CAC40, S&P500 - NK225, S&P500 - FTSE100, and S&P500 - DAX30.

Five marginal risk-neutral densities (S&P500, CAC40, NK225, FTSE100, and DAX30) are estimated semiparametrically using a cross-section of contemporaneously measured equity index option prices in each market. A bivariate risk-neutral Plackett (1965) density is constructed using the given marginals and risk-neutral correlation derived using an empirical pricing kernel and the historical joint density of the index returns. Price differences from a lognormal pricing formula using historical and risk-neutral return moments are found to be significant.

This paper has benefited from the suggestions of David Bates, Stephen Blyth, Frank Diebold, as well as participants at the Risk 99 conference and the 1999 Derivatives Research Project Research Retreat. The most recent version of this paper is available at http://www.stern.nyu.edu/~jrosenb0/J_wpaper.htm.

I. Introduction

Recently, the literature in univariate contingent claims valuation has focused on pricing by interpolation. A wide variety of estimation techniques have been developed — including parametric, semiparametric, and nonparametric — which rationalize an observed set of option prices, and may then be used to interpolate prices of unobserved assets.

Most of these techniques are based on estimation of a univariate risk-neutral (or state price) density which is fitted to a set of observed option prices for a particular underlying asset. The risk-neutral density may then be used to price other derivatives which are based on the same underlying asset. Related papers include Sherrick, Irwin, and Foster (1990, 1992), Shimko (1993), Derman and Kani (1994), Rubinstein (1994a), Longstaff (1995), and Ait-Sahalia and Lo (1998).

None of these techniques may be directly applied to the pricing of multivariate contingent claims (MVCC's) because knowledge of the marginal risk-neutral densities is not sufficient to identify the multivariate risk-neutral density which is necessary for pricing. In general, there are an infinite number of multivariate densities consistent with given marginals.

The majority of previous 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. These techniques do not usually replicate the prices of traded univariate options, and result in MVCC prices which are not reflective of current market conditions.

Rosenberg (1998) provides a methodology for MVCC valuation using a general parametric multivariate risk-neutral density function — the multivariate flexible density function — and 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 there are traded options on individual currencies and traded cross-currency options. Ho, Stapleton, Subrahmanyam (1995) provide a binomial approximation technique for MVCC pricing.

This paper generalizes univariate interpolation-based pricing methods to the multivariate case by developing a multivariate contingent claims pricing technique based on semiparametric estimation of a multivariate risk-neutral Plackett (1965) density. This method allows for completely general marginal risk-neutral densities and is compatible with all univariate risk-neutral density estimation techniques. MVCC prices using this method are consistent with current market prices of univariate contingent claims.

This technique has several advantages over existing techniques. First, any desired method may be used to estimate the marginal densities including non-parametric, semiparametric, or fully parametric techniques, depending on the quantity and quality of observed prices of options traded on individual assets.

Second, prices of options whose payoff depends on multiple assets do not need to be observed. A technique is introduced to estimate the risk-neutral correlation using time-series individual asset returns and an empirical pricing kernel (Rosenberg and Engle, 1999). This means that prices of multivariate contingent claims may be estimated as long as there are options traded on the individual assets.

The semiparametric MVCC pricing technique is used for valuation of one-month options on the better of two equity index returns. Option contracts with payoffs dependent on are four equity index pairs are considered: S&P500 - CAC40, S&P500 - NK225, S&P500 - FTSE100, and S&P500 - DAX30.

Five marginal risk-neutral densities (S&P500, CAC40, NK225, FTSE100, and DAX30) are estimated semiparametrically using a cross-section of contemporaneously measured equity index option prices in each market. A bivariate risk-neutral Plackett (1965) density is constructed using the given marginals and an estimated risk-neutral correlation based on an empirical pricing kernel and the historical joint density of the index returns. Price differences from a lognormal pricing formula using historical and risk-neutral return moments are found to be significant.

The remainder of the paper is structured as follows. Section II describes the semiparametric multivariate risk-neutral density estimation technique. Section III presents an application of this technique for the pricing of options on the better of two equity index returns. Section IV concludes the paper.

II. Semiparametric pricing of MVCCs

Univariate interpolative pricing techniques are designed to provide the price of a contingent claim on a particular underlying asset, which is consistent with the observed prices of contingent claims on the same underlying asset. These techniques identify a pricing function based on the condition that the pricing function should reprice existing assets accurately.

The “risk-neutral density” representation of the asset pricing problem has been found to be a convenient framework for univariate interpolative 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 “risk-neutral density” representation of the asset pricing problem which prices all contingent claims dependent on an underlying asset with price X_T. Any contingent claim which has its date T payoff entirely determined by the date T underlying asset price may be priced using this equation. Let the payoff function be written as $g_X(X_T)$ and the univariate risk-neutral density function be written as $f_{X,t}^*(X_T)$.

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

From a standard asset pricing perspective, equation (1) is a pricing formula which generates the price of a contingent claim ($D_{X,t}$), given its payoff function and an estimated risk-neutral density. Equation (1) also implicitly provides the estimation strategy for obtaining the risk-neutral density $f_{X,t}^*(X_T)$. If there are traded contingent claims dependent an underlying asset with price X_T, the observed prices of these assets ($D_{X,t,1}$, $D_{X,t,2}$, …, $D_{X,t,N}$) transform equation (1) into an identifying condition for the risk-neutral density. Estimation of the risk-neutral density entails solution of the inverse asset pricing problem: finding a risk-neutral density which is consistent with observed prices.

The multivariate contingent claim pricing problem may be viewed in terms of finding a multivariate risk-neutral density $f^*_{X,Y,\ldots,t}$ (X_T,Y_T,Z_T, …) which is consistent with observed option prices for several different underlying assets. Equation (2) provides the special case of a bivariate pricing equation which depends on a bivariate risk-neutral density $f_{X,Y,t}^*(X_T,Y_T)$.

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

If there are traded contingent claims dependent on an underlying asset with price Y_T, then the univariate risk-neutral density $f_{Y,t}^*(Y_T)$ may be obtained by an of the existing univariate risk-neutral density estimation techniques.

However, it is clear that knowledge of two marginals is not sufficient to uniquely determine their joint density. For example, in the bivariate normal case, knowledge that both X_T and Y_T have standard normal densities (μ_X=0, σ_X=1, μ_Y=0, σ_Y=1) does not uniquely determine their joint density. In particular, the correlation coefficient (ρ) is unknown.

The estimation problem in this paper is to find a multivariate risk-neutral density which has given marginal risk-neutral densities. For simplicity, the bivariate case is discussed in the remainder of the paper. Thus, the estimation problem is to determine $f_{X,Y,t}^*(X_T,Y_T)$ given $f_{X,t}^*(X_T)$ and $f_{Y,t}^*(Y_T)$.

There are two facets to this problem. First, an technique is required to construct a bivariate density given two previously estimated marginal densities. This paper utilizes the multivariate density due to Plackett (1965). Second, a technique is required to describe the relationship between the marginal densities. This paper estimates the risk-neutral correlation using a time-series of asset returns and an estimated pricing kernel.

II.i. Estimation of a multivariate risk-neutral densities with given marginals

While several approaches have been proposed to construct multivariate densities with completely general marginals, this paper utilizes the bivariate density due to Plackett (1965), which is further developed by Mardia (1967), and generalized to higher orders by Molenberghs and LeSaffre (1994) as well as Chakak and Koehler (1995).

There are several advantages to the multivariate Plackett density family. First, this density family is compatible with completely general marginals. Second, this family allows for correlations spanning the range of -1 to 1 and includes a member corresponding to independence. Third, association between pairs of random variables is measured using a single parameter. This is ideal for modeling situations in which data on the cross-moments of the multivariate density is limited.

Alternative density families which allow completely general marginals include that of Morgenstern (1956), Farlie (1960), and Gumbel (1960). Unfortunately, this density family severely restricts the degree of allowable correlation — typically, to less than 1/3 in absolute value — limiting its usefulness in practice. The technique of copulas (see, e.g.) is another possible methodology for combining specified marginals into a joint distribution. Copulas require an inverse probability transformation and are typically used when the marginals are specified in closed form; however, extension to the semi-parametric and non-parametric case is possible.

The bivariate Plackett density is defined as follows. Given univariate densities $f_{X,t}^{*}(X_T)$ and $f_{Y,t}^{*}(Y_T)$ with cumulative density functions $F_{X,t}^{*}(X_T)$ and $F_{Y,t}^{*}(Y_T)$, the cumulative Plackett joint density function $F_{X,Y,t}^{*}(X_T,Y_T)$ is given by:

(3)
$$F_{X,Y,t}(X_T,Y_T) = \left[\{S^2 + 4y(1-y)F_{X,t}(X_T)F_{Y,t}(Y_T)\}^{\frac{1}{2}} - S\right] / \{2(1-y)\} \qquad y \neq 1$$
$$\text{with } S = 1 - \left[F_{X,t}(X_T) + F_{Y,t}(Y_T)\right](1-y)$$
$$F_{X,Y,t}(X_T,Y_T) = F_{X,t}(X_T)F_{Y,t}(Y_T) \qquad y = 1$$

The degree of association between X_T and Y_T is determined by the parameter ψ, which is non-negative. When ψ is zero, the correlation is -1; when ψ is one, the correlation is 0; and as ψ approaches positive infinity, the correlation approaches 1. In the continuous case, the bivariate density function is obtained by taking the second cross-partial derivative with respect to X_T and Y_T. The discrete density function is obtained in the usual manner.

II.ii. Estimation of risk-neutral correlation using time-series data

In the bivariate Plackett density, given the two marginals, only one parameter (ψ) needs to be estimated for identification. For some densities (e.g. Gaussian), correlation is a sufficient statistic to characterize dependence, and provides a natural identifying condition: choose ψ to match the sample correlation of X and Y.

However, the joint densities of some financial time series exhibit asymmetric correlation (different correlations in up and down markets) suggesting that another criterion (e.g. semi-correlation) should be used. See, e.g., Erb, Harvey, Viskanta (1994), Longin and Solnik (1995), and De Santis and Gerard (1997). For the purposes of multivariate contingent claim pricing, only the dependence structure over in-the-money states is relevant. Hence, the suggested criterion would be a truncated correlation measure (e.g. semi-correlation in the case of a standard bivariate at-the-money option) over these states.

This paper develops an estimator based on risk-neutral correlation; the extension to risk-neutral semi-correlation or truncated correlation is straightforward. In particular, the time-series correlation between the asset returns and an empirical pricing kernel to estimate the risk-neutral correlation. Then, an optimization procedure is used to select the value of ψ which replicates the estimated risk-neutral correlation in the bivariate risk-neutral density with the given marginals.

The objective and risk-neutral densities are linked to each other through the risk-neutral pricing equation. This relationship is used to transform data on underlying asset density into data on the risk-neutral density.

Consider the ratio of the discounted risk-neutral density and the original density, which is referred to as the pricing kernel. The pricing kernel measures investor risk aversion by defining the probability normalized value of the payoff in each state, i.e. the state price per unit probability. Notice that the discounted risk-neutral density is the state price density.

$$\begin{aligned} &\text{Objective density: } f_{X,Y,t}(X_T, Y_T) \\ (4) \qquad &\text{Pricing kernel: } K_{t,T}(X_T, Y_T) = e^{-r(T-t)} f^*_{X,Y,t}(X_T, Y_T) \,/\, f_{X,Y,t}(X_T, Y_T) \\ &\text{Risk - neutral density: } f^*_{X,Y,t}(X_T, Y_T) = e^{r(T-t)} f_{X,Y,t}(X_T, Y_T)\, K_{t,T}(X_T, Y_T) \end{aligned}$$

It is clear that knowledge of the pricing kernel and objective density fully determines the risk-neutral density. A more limited goal is to determine the risk-neutral correlation, given data on the objective density and the pricing kernel.

Consider the risk-neutral correlation between X_T and Y_T, where $E^*[\bullet]$ denotes the expectation taken under the risk-neutral measure $f_{X,Y,t}{}^*(X_T,Y_T)$ and $\sigma^*[X_T]$ and $\sigma^*[Y_T]$ are the risk-neutral standard deviations of X_T and Y_T.

$$(4) \qquad r_{X,Y} = \left(E^*[X_T Y_T] - E^*[X_T]E^*[Y_T]\right) / \left(s^*[X_T]s^*[Y_T]\right)$$

There are five risk-neutral moments to be estimated to evaluate equation (4). The following describes the estimator for risk-neutral expectation of the product of X_T and Y_T: $E^*[X_T Y_T]$. The remaining moments estimators are determined in an analogous manner. Using equation (4) to rewrite the moment condition in terms of the objective density and the pricing kernel:

$$(5) \qquad E^*[X_T Y_T] = \iint X_T Y_T f^*_{X,Y,t}(X_T, Y_T) dX_T dY = e^{r(T-t)} \iint X_T Y_T f_{X,Y,t}(X_T, Y_T) K_{t,T}(X_T, Y_T) dX_T dY_T$$

Equation (5) states that the risk-neutral moment may also be interpreted as a scaled pricing-kernel weighted objective moment. If the objective density $f_{X,Y,t}(X_T,Y_T)$ is known in closed form, the double integral may be evaluated analytically or numerically to obtain the desired moment. If the objective density is obtainable by simulation conditional on the current state variables, then equation (5) may be evaluated by replacing the integral with a sample average over simulated realizations of the random variables. If the objective density is time-invariant, then equation (5) may be evaluated by replacing the integral with an average over observed realizations using time-series data on the asset prices (or returns). Equation (6) provides the estimator applicable to the second and third cases.

$$(6) \qquad \hat{E}^*[X_T Y_T] = e^{r(T-t)} \frac{1}{N} \sum_{i=1}^{N} X_{i,T} Y_{i,T} K_{t,T}(X_{i,T}, Y_{i,T})$$

The third case is compatible with stochastic volatility models in which the correlation is constant but the individual variances are time varying. See, for example, Bollerslev (1990). In this case, the returns should be normalized by their conditional variances prior to calculating the sample averages.

III. Semiparametric pricing of options on the better of two equity index returns

To illustrate the use of this estimation and pricing technique, it is applied to the case of pricing one-month options on the better of two equity index returns. In this case, X_T and Y_T represent one-month fully-hedged returns on one of five equity index portfolios: the S&P500 (U.S), the CAC40 (France), the Nikkei 225 (Japan), the FTSE 100 (England), and the DAX30 (Germany).

III.i. Estimation of marginal equity index risk-neutral densities

The marginal equity index risk-neutral densities are estimated semiparametrically using the method of Rosenberg and Engle (1999) which is closely related to the semi-parametric technique of Ait-Sahalia and Lo (1998) and Shimko (1993). For this purpose, a general pricing formula is proposed which is fit to a cross-section of option prices by allowing a different implied volatility ($\sigma_t(K)$) for each option contract.

(7) $$C_t = BS(X_t, K, T-t, r_f, s_t(K))$$

It is well known — see, e.g. Breeden and Litzenberger (1978) —that the first derivative of the call price formula with respect to the exercise price provides an estimate of the cumulative state price density (and cumulative risk-neutral density). Let n(•) be the standard normal density function, let N(•) be the cumulative standard normal density function, and let $d_2 = [\ln(X_t/K) + ½(r - \sigma^2(T\text{-}t))] / \sigma\sqrt{T\text{-}t}$, and let $\sigma_{K,t}(K)$ be the first derivative of the implied volatility function with respect to the option exercise price.

(8) $$e^{-r(T-t)}\text{Prob}^*(X_T \leq K) = e^{-r(T-t)}[1 + K * n(d_2) * s_{K,t}(K) - N(d_2)]$$

A discrete risk-neutral marginal density is defined over 1000 states corresponding to asset returns (X_T/X_t - 1) between -30% and 30%. The cumulative risk-neutral density function is evaluated using equation (8) with the first derivative numerically evaluated based on a B-spline fit of the implied volatility function. For additional details, see Rosenberg and Engle (1999).

This procedure is applied to estimate marginal risk-neutral densities defined over equity index returns for the S&P500 index (United States), the CAC40 index (France), the Nikkei 225 index

(Japan), the FTSE 100 index (England), and the DAX30 index (Germany). For the estimation, end-of-day prices of April European index put option contracts are collected on March 19, 1999. Implied volatilities for each contract are calculated by numerically inverting the Black-Scholes (1973) formula adjusted for dividends.

For the implied volatility estimation, the dividend yield used is the implied dividend yield based on the put-call parity relationship and the prices of at-the-money put and call index options in each market. The underlying price used is the closing price of the index, and the riskless interest rate is the continuously compounded yield to maturity of the national treasury bill contract with expiration closest to the option expiration. The time until expiration used is the number of calendar days until the option expiration.

Ideally, the option expiration dates would be identical across markets, so that the implied volatilities reflect the same time-span. In practice, there are some differences in option expiration dates, e.g. the S&P500, FTSE100, and NK225 option expire mid-month, while the DAX30 and CAC40 options expire at the end of the month. For simplicity, it is assumed that option implied volatilities across maturities are equal (i.e. the implied volatility term structure is flat), so that the calculated implied volatility is a proxy for the implied volatility for an option with one-month until expiration.

The first panel of Table 1 describes the data used to estimate the implied volatilities, and Figure 1 plots the annualized implied volatilities for each of the option classes against option moneyness (K/S_t - 1). The implied volatility skew, i.e. the downward sloping curve of implied volatility against option moneyness, is present in all of the markets. This has been attributed to crash fears — Bates (1991) — as well as portfolio insurance demand — Rosenberg and Engle (1999).

Figure 2 plots the one-month marginal risk-neutral densities defined over local equity index returns. While the densities are drawn as if they were continuous, the actual densities are discrete with a support defined over 1000 one-month returns ranging from -30% to 30%. The second panel of Table 1 presents estimates of risk-neutral standard deviation, skewness, and kurtosis for these densities. All of the risk-neutral densities exhibit negative skewness except for the NK225, and all of the densities exhibit positive excess kurtosis. Negative skewness and positive excess kurtosis in the S&P500 risk-neutral density have been also been found by Rubinstein (1994), Ait-Sahalia and Lo (1998), and Rosenberg and Engle (1999).

III.ii. Estimation of the empirical pricing kernel

The transformation of time-series data on asset returns into an estimate of risk-neutral correlation requires a pricing kernel. In its most general form, the pricing kernel — $K_{t,T}(X_T,Y_T)$ — depends on both asset price outcomes. This paper uses the special case of a univariate pricing kernel which only depends on S&P500 index return states: $K_{t,T}(X_T)$.

In this paper, the pricing kernel is estimated using the empirical pricing kernel (EPK) technique of Rosenberg and Engle (1999). A short summary of this technique is presented below. For additional details, see the previously mentioned paper.

The pricing kernel, which measures the state price per unit probability across states, is naturally estimated by separately obtaining the state price density and the state probability density and taking their ratio. An innovation of the EPK technique is that the state price density is estimated using a current cross-section of option prices so that it reflects current investor risk aversion and probability beliefs. This is paired with a conditional state probability density function which is time-varying reflecting current market conditions.

The state price density is estimated using the risk-neutral density estimation technique described in section III.i. using the S&P500 index option data on March 19, 1999. Note that the state price density is the discounted risk-neutral density. One minor difference from section III.i. is that the state space is discretized into 100 discrete return states ranging from -10% to 5%.

The state probability density is estimated by monte-carlo simulation (200,000 replications) of an estimated dynamic state probability model. The state probability model used is a GJR-GARCH model (Engle and Ng, 1993) estimated over the period from January 1, 1994 through March 18, 1999. Standardized residuals from the model are saved and used as an empirical innovation density.

Figure 3 graphs the empirical pricing kernel which provides a measure of investor risk aversion over S&P500 return states. The EPK has a downward slope and is steeply rising for large negative return states. This indicates that investors are willing to pay more for a probability normalized dollar payoff when the S&P500 returns are large and negative, i.e. crash states, than for positive return states. In other words, investors are willing to pay up (probability normalized) for hedging assets which payoff in bad states of the world.

III.iii. Estimation of the bivariate risk-neutral densities

The risk-neutral correlations for the equity index return pairs (S&P500 - CAC40, S&P500 - NK225, S&P500 - FTSE100, and S&P500 - DAX30) are obtained using the methodology described in section II.ii. Equation (6) and its analogs are calculated using historical data on the asset returns and the empirical pricing kernel described in section III.ii.

For the risk-neutral correlation estimation, monthly local currency (i.e. fully-hedged) returns for each of the equity indices are obtained for the period from January 1994 - February 1999. The first panel of Table 2 describes this data. Four of five return series exhibit positive skewness and positive excess kurtosis.

The second panel of Table 2 presents the objective correlations and the risk-neutral correlations. The risk-neutral correlations range from .07 to .21 higher than the objective correlations. This indicates a serious potential problem using objective correlations instead of risk-neutral correlations for MVCC pricing.

The estimated association parameters (ψ) are reported in the third panel of Table 2. These are calculated by an optimization which matches the risk-neutral correlation of the estimated bivariate density to the risk-neutral correlation estimated above. As noted previously, ψ is monotonically increasing in correlation. The equity index pair with the lowest risk-neutral correlation S&P500-NK225 (.617) has the lowest ψ (11.33); the equity index pair with the highest risk-neutral correlation S&P500 - FTSE100 (.833) has the highest ψ (52.09).

Figures 4-7 plot the estimated joint risk-neutral density functions. The characteristics of these joint densities which are most easily identified are: (1) strong positive correlations for all return pairs as indicated by the narrow right-facing diagonal hump in each graph (2) different risk-neutral marginal standard deviations reflected in the length and width of the hump in each graph.

III.iv. Pricing options on the better of two equity index returns

MVCC pricing using equation (2) is straightforward given an estimated bivariate risk-neutral density and a known payoff function. This paper is concerned with the pricing of equity index options on the better of two equity index returns.

The options of interest for this paper have payoffs based one-month S&P500 return and the return of a selected fully-hedged foreign equity index. These options have payoffs scaled by 100. Thus, the equation (2) risk-neutral pricing formula specializes to:

$$(9) \qquad D_{r_X, r_Y, t} = e^{-r(T-t)} \iint Max[0, 100 r_X, 100 r_Y] f^{*}_{r_X, r_Y, t}(r_X, r_Y) dr_X \, dr_Y$$

It is important to note that the foreign equity index return risk-neutral density is identical to the fully-hedged foreign index return risk-neutral density for valuation of dollar denominated payoffs when a single international pricing kernel exists. Equation (4) shows that the risk-neutral density depends on the objective payoff probabilities, which are identical for foreign index returns in foreign currency and fully-hedged foreign index returns. The risk-neutral density also depends on the pricing kernel. If there is a single international pricing kernel, then the foreign and fully-hedged foreign risk-neutral densities are the same.

The fourth panel of Table 2 presents a comparison of prices for options on the better of two returns using the bivariate Plackett densities estimated above and prices of options using a bivariate lognormal risk-neutral density. The bivariate lognormal risk-neutral density is a reasonable pricing benchmark, because it has been frequently used in the literature for MVCC valuation.

Two comparison bivariate lognormal densities are used. The first has the standard deviation and correlation parameters set to the sample values for the historical return density (Table 2, panels 1 and 2). The second has the standard deviation and correlation parameters set to the risk-neutral values (Table 1 - panel 2 and Table 2 panel 2).

The Plackett - semiparametric prices are higher than the lognormal prices in seven of eight cases. And, the lognormal prices using the risk-neutral parameters are closer to the Plackett - semiparametric prices in three of four cases. The largest pricing differences between the Plackett - semiparametric and lognormal prices is for the S&P500 - NK225 option with differences of 57.78% and 55.21%. Overall, the pricing differences range from -14.55% to 57.78%. These results indicate that a bivariate lognormal density is an inadequate approximation to the bivariate risk-neutral density implied by option price data.

IV. Conclusion

This paper has developed a semiparametric technique for multivariate contingent claim valuation. A multivariate risk-neutral density function estimated using market data on the prices of options on individual assets and a risk-neutral correlation estimated using the time-series of asset returns and an empirical pricing kernel. The estimated risk-neutral density is used to infer prices of additional MVCC's.

This technique has several attractive features. First, the 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. Second, the estimation methodology is compatible with existing nonparametric, semiparametric, and parametric techniques for univariate risk-neutral density estimation. Third, estimation is possible without observed prices for related multivariate contingent claims.

This pricing technique is used for valuation of options on the better of two equity index returns. These prices are compared to those generated by bivariate lognormal risk-neutral density, which is often used in this context. The results of the pricing comparisons indicate that a bivariate lognormal density is an inadequate approximation to the bivariate risk-neutral density implied by option price data.

Bibliography

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

Bates, D. S., 1991, "The Crash of 87 - Was It Expected? - the Evidence from Options Markets," *Journal of Finance*, 46, 1009-1044.

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

Bollerslev, T., 1990, "Modelling the Coherence in Short-Run Nominal Exchange Rates - a Multivariate Generalized ARCH Model," *Review of Economics and Statistics*, 72, 498-505.

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

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

Chakak, A. and K. J. Koehler, 1995, "A Strategy for Constructing Multivariate Distributions," *Communications in Statistics - Simulation and Computation*, 24, 537-550.

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

De Santis, G. and B. Gerard, 1997, "International Asset Pricing and Portfolio Diversification with Time-Varying Risk," *Journal of Finance,* 52, 1881-1912.

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

Engle, R. F. and V. K. Ng, 1993, "Measuring and Testing the Impact of News on Volatility," *Journal of Finance*, 48, 1749-1778.

Erb, C. B., C. R. Harvey and T. E. Viskanta, 1994, "Forecasting International Correlation," *Financial Analysts Journal*, 50, 32-45.

Farlie, D. J. G., 1960, "The Performance of Some Correlation Coefficients for a General Bivariate Distribution," *Biometrika*, 47, 307-322.

Gumbel, E. J., 1960, "Bivariate Exponential Distributions," *Journal of the American Statistical Association*, 55, 698-707.

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

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

Longin, F. and B. Solnik, 1995, "Is the Correlation in International Equity Returns Constant: 1960-1990?," *Journal of International Money and Finance*, 14, 3-26.

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

Mardia, K. V., 1967, "Some Contributions to Contingency-type Bivariate Distributions," *Biometrika*, 54, 239-249.

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

Molenberghs, G. and E. Lesaffre, 1994, "Marginal Modelling of Correlated Ordinal Data using a Multivariate Plackett Distribution," *Journal of the American Statistical Association*, 89, 633-644.

Morgenstern, D., 1956, "Einfache Beispiele Zweidimesnsionaler Verteilungen," *Mitteilingsblatt fur Mathematische Statistik*, 8, 234-235.

Nelsen, R. B. "An Introduction to Copulas." . New York: Springer-Verlag, 1999.

Plackett, R. L., 1965, "A Class of Bivariate Distributions," *Journal of the American Statistical Association*, 60, 516-522.

Reiner, E. "Quanto Mechanics." In *From Black-Scholes to Black-Holes: New Frontiers in Option Pricing*, 147-154. London: Risk Magazine/FINEX, 1992.

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

Rosenberg, J. and R. F. Engle, 1999, "Empirical Pricing Kernels," Manuscript.

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

Rubinstein, M. "One for Another." In *From Black-Scholes to Black-Holes: New Frontiers in Option Pricing*, 147-154. London: Risk Magazine/FINEX, 1992.

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

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

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

Sherrick, B. J., S. H. Irwin and D. L. Forster, 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.

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

Stapleton, R. C. and M. G. Subrahmanyam, 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.

Table 1
Analysis of marginal risk neutral densities
Defined over equity index returns

Option data used to estimate marginal risk-neutral densities
April European equity index puts on March 19, 1999

	S&P500	CAC40	NK225	FTSE100	DAX30
Num. of obs.	24	19	11	22	22
Minimum moneyness	-30.73%	-26.53%	-26.73%	-21.71%	-25.48%
Maximum moneyness	15.45%	0.72%	6.85%	5.87%	1.97%
Minimum implied standard deviation (annualized)	15.86%	21.73%	23.75%	12.18%	15.15%
Maximum implied standard deviation (annualized)	39.07%	33.66%	44.45%	33.37%	29.33%
Days until expiration	28	42	21	28	42

Properties of the risk-neutral marginal densities

	S&P500	CAC40	NK225	FTSE100	DAX30
Risk-neutral standard deviation (annualized)	23.03%	24.55%	33.71%	19.88%	19.36%
Risk-neutral skewness	-0.29	-0.25	0.48	-1.07	-0.42
Risk-neutral excess kurtosis	1.65	0.41	0.35	1.88	0.79

Table 2
Analysis of joint risk neutral densities
Defined over equity index returns

Objective return moments (monthly returns: Jan. 1994 - Feb. 1999)

	S&P500	CAC40	NK225	FTSE100	DAX30
Num. obs.	62	62	62	62	62
Mean (annualized)	21.82%	14.20%	-1.55%	12.80%	17.94%
Standard dev. (annualized)	13.14%	21.92%	22.00%	13.95%	20.39%
Skewness	-0.02	-0.12	-0.31	0.39	-0.64
Excess kurtosis	0.13	0.42	1.61	-0.35	0.93

Correlation with S&P500 returns

	S&P500	CAC40	NK225	FTSE100	DAX30
Objective	1.000	0.699	0.406	0.722	0.682
Risk-neutral	1.000	0.770	0.617	0.833	0.806

Estimation of joint risk-neutral density

		S&P500 - CAC40	S&P500 - NK225	S&P500 - FTSE100	S&P500 - DAX30
ψ		26.76	11.13	52.09	36.38

Estimated outperformance option prices

Bivariate risk-neutral density		S&P500 - CAC40	S&P500 - NK225	S&P500 - FTSE100	S&P500 - DAX30
Plackett - semiparametric		3.68	5.43	3.00	3.18
Lognormal (historical moments)		3.24	3.44	2.46	3.09
Lognormal (risk-neutral moments)		3.50	3.50	2.88	3.72
Percent price difference (Plackett vs. lognormal-historical)		13.86%	57.78%	21.65%	2.79%
Percent price difference (Plackett vs. lognormal-risk-neutral)		5.19%	55.21%	4.08%	-14.55%

Figure 1
Equity index options
Implied standard deviations (March 19, 1999)

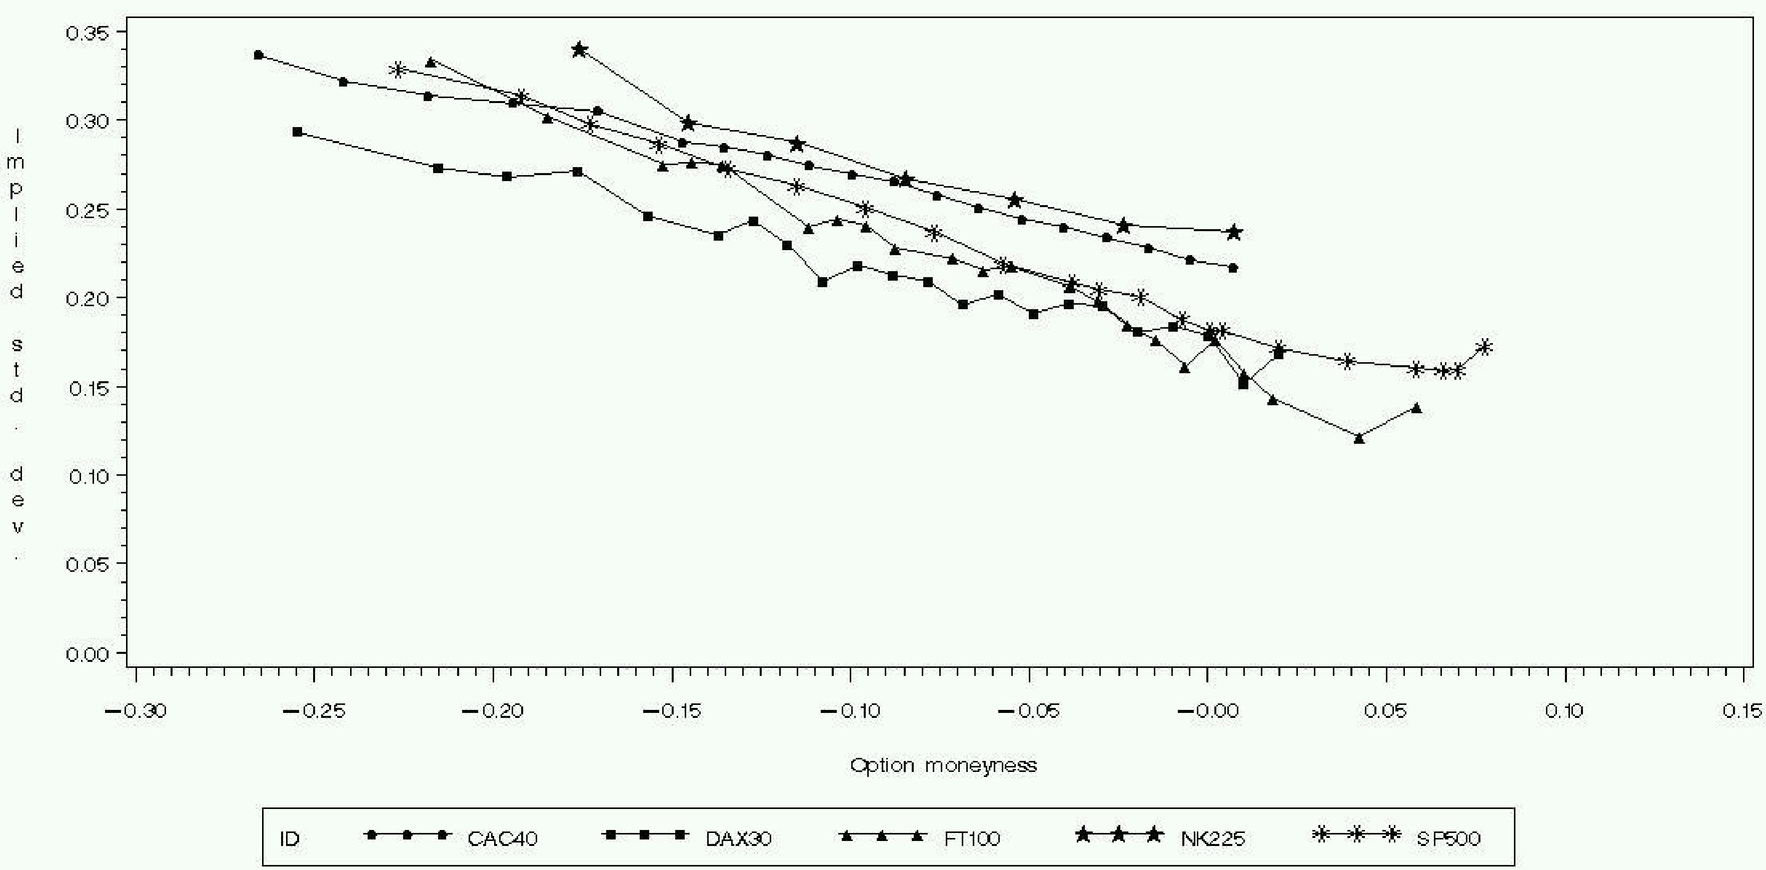

Figure 2
One-month marginal risk-neutral densities
Defined over equity index returns
March 19, 1999

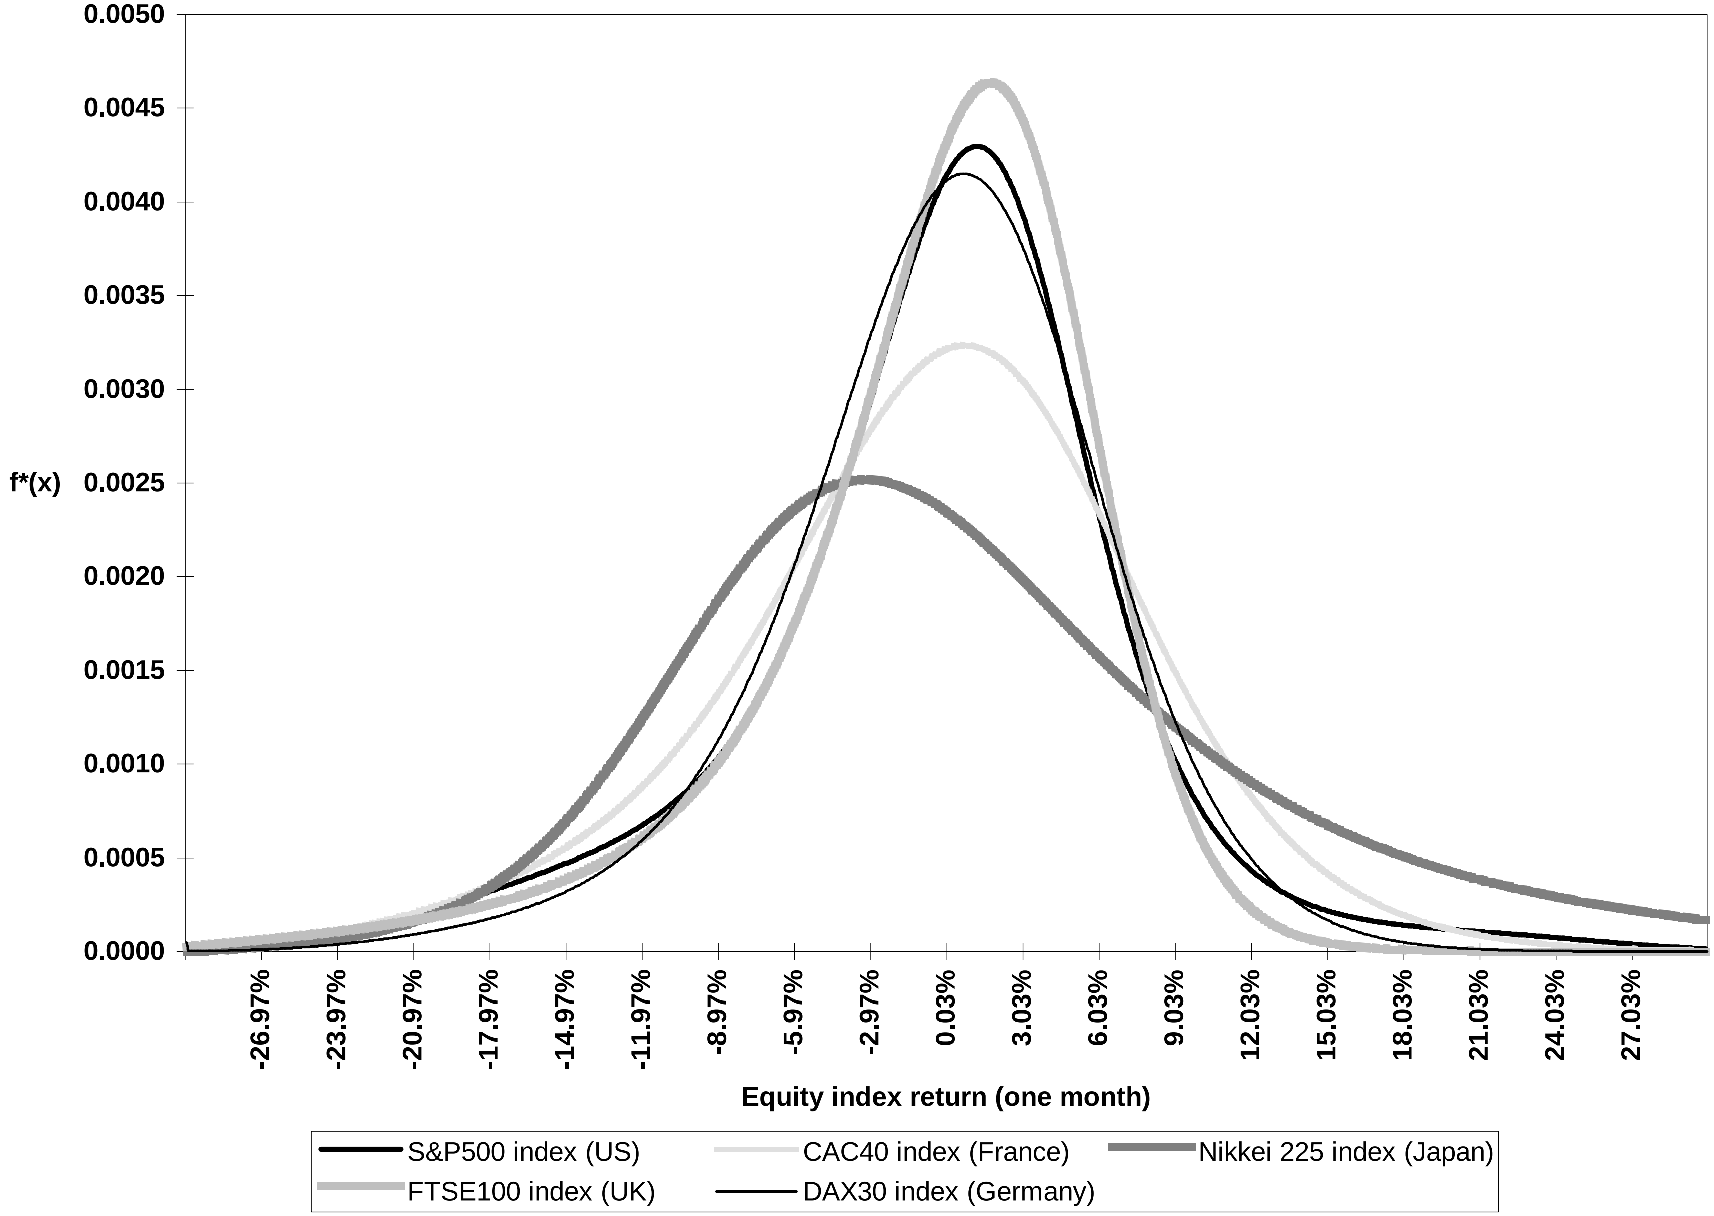

Figure 3
S&P500 empirical pricing kernel
March 19, 1999

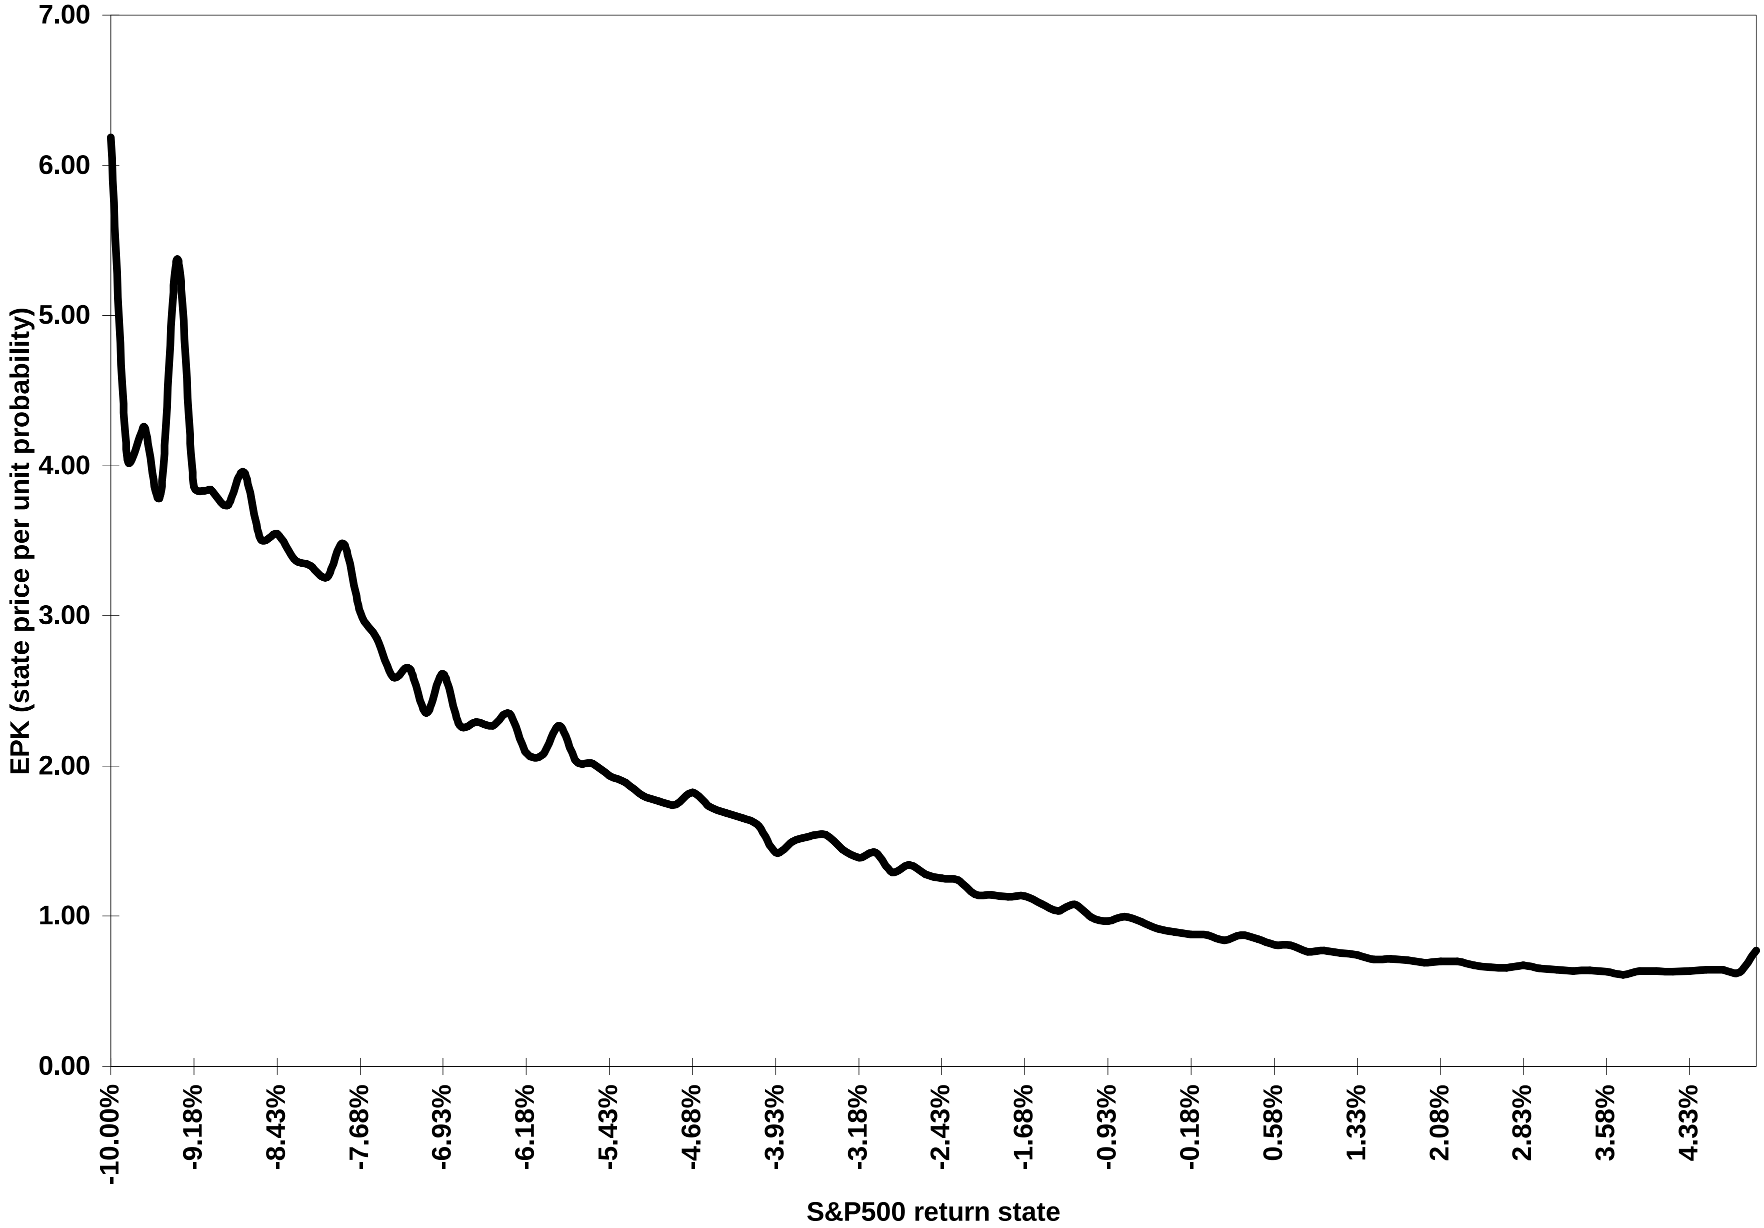

Figure 4
One-month joint risk-neutral density
S&P500 returns and CAC 40 returns
March 19, 1999

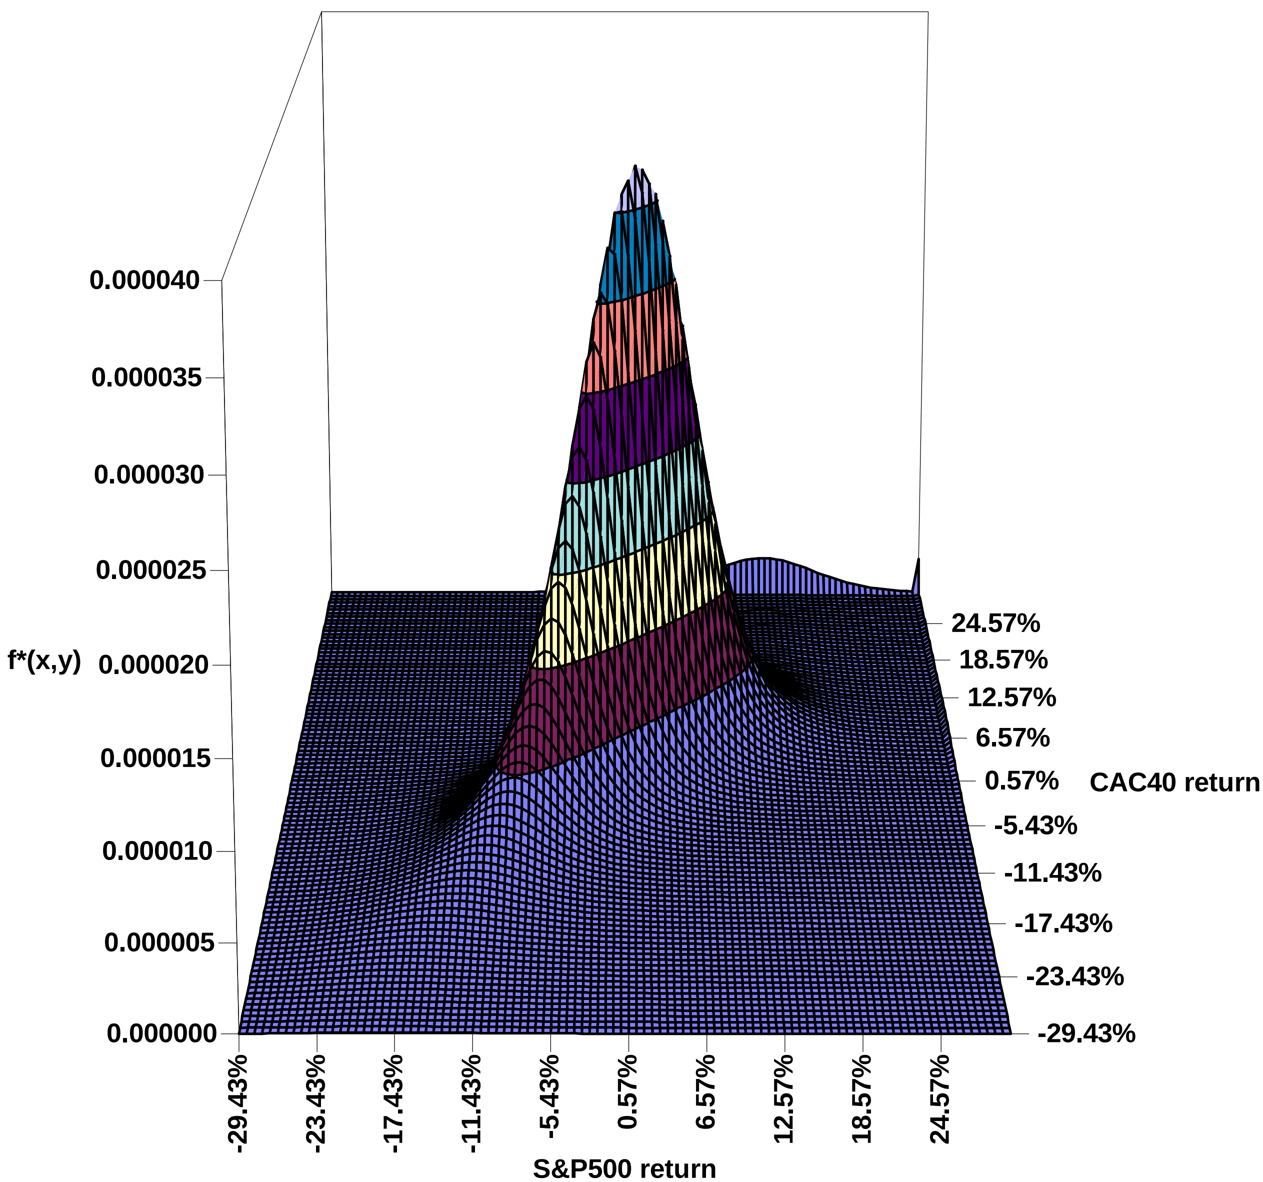

Figure 5
One-month joint risk-neutral density
S&P500 returns and NK 225 returns
March 19, 1999

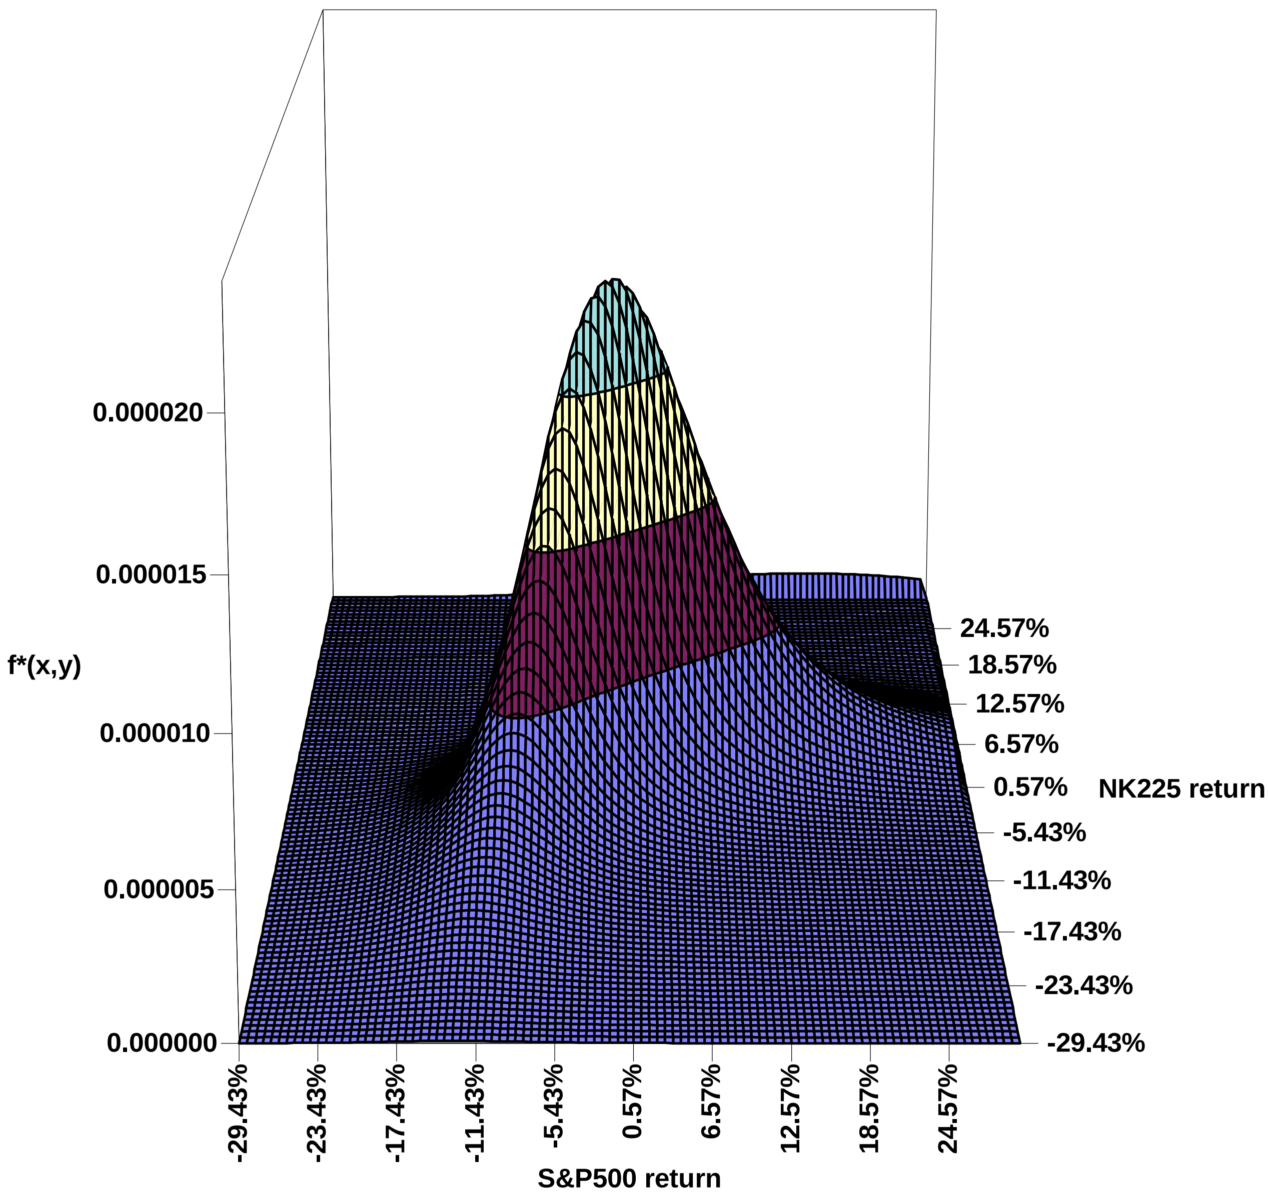

Figure 6
One-month joint risk-neutral density
S&P500 returns and FTSE100 returns
March 19, 1999

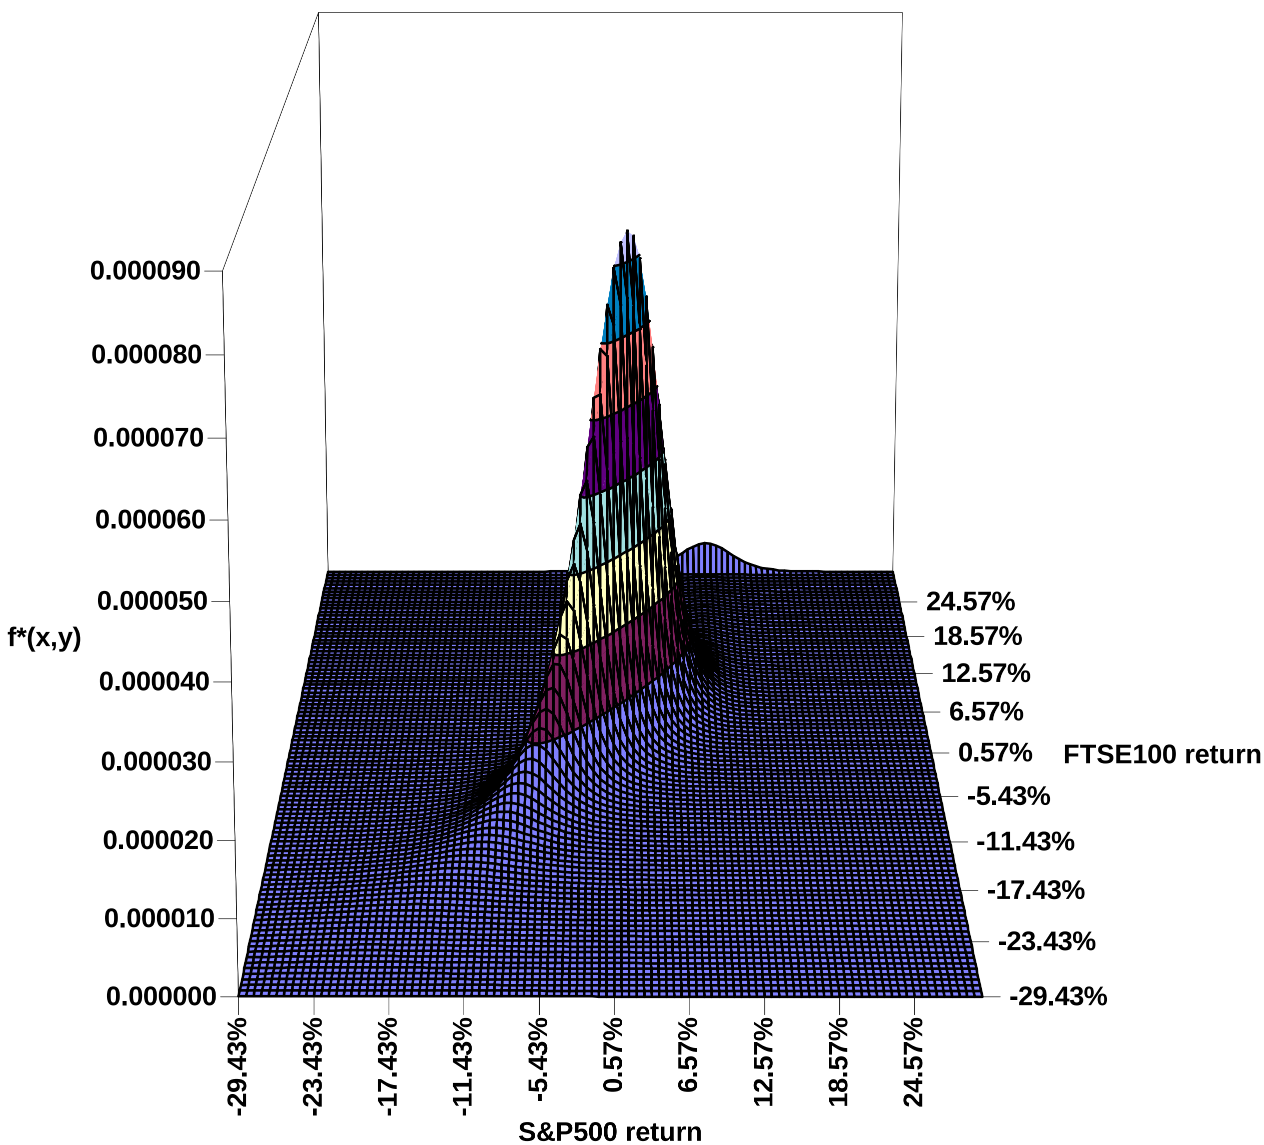

Figure 7
One-month joint risk-neutral density
S&P500 returns and DAX 30 returns
March 19, 1999