



Department of Finance

Working Paper Series 1998

FIN-98-057

Pricing Multivariate Contingent Claims using Estimated Risk-neutral Density Functions

Joshua Rosenberg

June 1997

This working paper series has been generously supported by a grant from

CDC Investment Management Corporation

**Pricing Multivariate Contingent Claims using
Estimated Risk-neutral Density Functions**

June 1997

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

Many asset price series exhibit time-varying volatility, jumps, and other features inconsistent with assumptions about the underlying price process made by standard multivariate contingent claims (MVCC) pricing models. This paper develops an interpolative technique for pricing MVCCs — flexible NLS pricing — that involves the estimation of a flexible multivariate risk-neutral density function implied by existing asset prices.

As an application, the flexible NLS pricing technique is used to value several bivariate contingent claims dependent on foreign exchange rates in 1993 and 1994. The bivariate flexible risk-neutral density function more accurately prices existing options than the bivariate lognormal density implied by a multivariate geometric brownian motion. In addition, the bivariate contingent claims analyzed have substantially different prices using the two density functions suggesting flexible NLS pricing may improve accuracy over standard methods.

Multivariate contingent claims are becoming more common as derivative asset use increases. Options on a portfolio of assets, such as options on the S&P500 index, are the most commonly traded multivariate contingent claim. Other multivariate claims include spread options, such as options on the difference between the heating oil and crude oil price, which are traded on the New York Mercantile Exchange. Over-the-counter multivariate claims include options on the minimum or maximum of several assets, dual strike options, and multivariate digital options which are especially common on foreign currencies.

Most existing MVCC pricing techniques rely on the assumption that the underlying asset prices are generated by a multivariate diffusion process, with constant (or deterministic) variance and no jumps. Examples include Margrabe (1978), Stulz (1982), Johnson (1987), Reiner (1992), Shimko (1994). These pricing formulas are derived using a multivariate generalization of the Black-Scholes (1973) technique and imply a multivariate lognormal original price density and multivariate lognormal risk-neutral density. Approximate MVCC pricing formulas which also rely on lognormality include Boyle and Tse (1990), Heenk, Kemna, and Vorst (1990), and Pearson (1995). Multivariate lattices, which represent a discrete-time analog to a multivariate GBM are used by Stapleton and Subrahmanyam (1984a, 1984b), Boyle (1988), Boyle, Evnine, and Gibbs (1989), and Rubinstein (1992, 1994b). In a recent paper, Ho, Stapleton, Subrahmanyam (1995) utilize a binomial approximation with deterministic volatility to price MVCC's allowing for non-stationarity variances.

Departures from lognormality are well documented in many financial time-series. See, for example, Bollerslev, Chou, and Kroner (1992). The lognormal density with constant variance inadequately characterizes the probabilities of large but infrequent market events such as currency devaluations or market crashes. In addition, the "riskless" hedge portfolios used in derivation of existing multivariate contingent claims pricing formulas are not hedged against changes in underlying asset variance. Under stochastic variance, these portfolios will be risky.

A modern interpolative approach that attempts to address deviations from lognormality in the univariate pricing context is that of Sherrick, Irwin, and Foster (1990, 1992), Longstaff (1994), Rubinstein (1994a), Shimko (1993), and Derman and Kani (1994). In these papers, the univariate terminal risk-neutral density or risk-neutral process is estimated by matching fitted and observed option prices. The risk-neutral density is then used to infer prices for assets not included in the estimation procedure.

This paper develops a multivariate contingent claims pricing technique, flexible NLS pricing, that infers a multivariate risk-neutral density from a set of existing prices on multiple assets. Flexible NLS pricing improves upon existing MVCC pricing techniques, because it does not depend on lognormality of the underlying asset returns, and thus may be more accurate for a wider variety of underlying price processes. Flexible NLS pricing introduces a new multivariate parametric family designed to emulate potential characteristics of the risk-neutral density, and estimation is accomplished using a generalization of univariate interpolative pricing techniques.

As an application, several types of bivariate contingent claims that depend on foreign exchange rates are valued using flexible NLS and lognormal NLS techniques. Evidence is found that the empirical risk-neutral density has asymmetries and tail shapes not captured by the lognormal specification. Since a multivariate lognormal risk-neutral density is a property of most existing valuation models, rejection of lognormality suggests potential biases in these models. In several cases, there are substantial differences in estimated contingent claims prices using the two densities, indicating that the flexible NLS technique may improve pricing accuracy.

The paper is organized as follows. In section I, the flexible NLS pricing technique is described. In section II, flexible NLS pricing is applied to bivariate contingent claims that depend on exchange rates, and results are compared to lognormal NLS pricing. Conclusions are presented in section III.

I. Flexible NLS pricing

In this section, the flexible NLS technique for pricing multivariate contingent claims is developed. Flexible NLS pricing advances existing multivariate pricing techniques, because it does not require a multivariate lognormal risk-neutral density. Thus, it is compatible with deviations of the underlying price process from standard conditions such as constant volatility and no jumps. It also utilizes a parametric multivariate risk-neutral density family tailored to this application that allows for shapes associated with a variety of underlying price densities and preferences, but includes the multivariate lognormal as a special case.

Flexible NLS pricing involves estimation of a multivariate risk-neutral density function from a set of existing prices on multiple assets. Estimation is accomplished by non-linear least squares (NLS) minimization of the squared proportional distance between fitted prices using the risk-neutral pricing formula and observed asset prices. The estimated risk-neutral density function is then used to price additional multivariate contingent claims whose prices are not observable. These prices are obtained by inserting the new asset's payoff function into the risk-neutral pricing formula and integrating over the estimated risk-neutral density.

The risk-neutral density function to be estimated is often referred to as an equivalent martingale measure or equilibrium price measure, and it describes the normalized equilibrium state prices. Existence of this density function has been proven by Harrison and Kreps (1979) and Harrison and Pliska (1981) under the assumption of no arbitrage.

The risk-neutral density function, $f^*(\bullet)$ allows all asset prices to be expressed as the present value of their expected payoffs under this density. When states of the world are indexed by the single underlying asset price X and there are no dividend payments, the price of the underlying asset X_t may be represented as:

$$(1) \quad X_t = e^{-rT} E_t^*(X_{t+T}) = e^{-rT} \int X_{t+T} f^*(X_{t+T}) dX_{t+T}$$

The power of the risk-neutral density pricing approach is due to the fact any derivative asset on the underlying may be priced using a generalized version of equation (1). For a derivative asset with payoffs defined by $g(X_{t+T})$ and price D_t ,

$$(2) \quad D_t = e^{-rT} E_t^*(g(X_{t+T})) = e^{-rT} \int g(X_{t+T}) f^*(X_{t+T}) dX_{t+T}$$

A similar formula to equation (2) may be applied to value multivariate contingent claims. Assuming that states of the world are indexed by multiple underlying assets, the only difference between valuing multivariate and univariate claims is that the payoff function for the multivariate claims involves the prices of multiple underlying assets, and the expectation is taken over all the underlying asset prices. Consider an economy with states indexed by two underlying asset prices, X and Y . Then, given a risk-neutral density $f^*(X_{t+T}, Y_{t+T})$, any asset price D_t with payoff function $g(X_{t+T}, Y_{t+T})$ may be written as:

$$(3) \quad D_t = e^{-rT} E_t^*[g(X_{t+T}, Y_{t+T})] = e^{-rT} \iint g(X_{t+T}, Y_{t+T}) f^*(X_{t+T}, Y_{t+T}) dX_{t+T} dY_{t+T}$$

In this paper, estimation of the empirical multivariate risk-neutral density is accomplished by an optimization that matches fitted prices to observed market prices, using a parametric specification for the risk-neutral density. For simplicity, the bivariate case will be described, although greater generality is straightforward. Suppose the true parametric risk-neutral density, $f^*(X_{t+T}, Y_{t+T}; \theta_t)$, is defined over the values of two state variables indexed by two asset prices, X_{t+T} , Y_{t+T} , where θ_t is a parameter vector.

Since the risk-neutral density expresses all asset prices as the present value of their expected payoff under this density, the fitted price of the i^{th} asset ($1 \leq i \leq N$) at date t evaluated at the estimated parameter vector is:

$$(4) \quad \hat{D}_{i,t}(\hat{\theta}_t) = e^{-rT} E_{\hat{\theta}_t}^*[g_i(X_{t+T}, Y_{t+T})]$$

In this case, $g_i(X_{t+T}, Y_{t+T})$ is the payoff function for the i^{th} asset, and r is the riskless rate of interest.

In the absence of arbitrage and correct specification of the risk-neutral density function, the observed asset price should be exactly equal the fitted asset price. However, it is plausible to expect some pricing errors due to potential problems with reported prices, such as non-synchronicity, price-discretization, and bid-ask bounce. This suggests the following moment condition for each observed asset price:

$$(5) \quad E[D_{i,t} - \hat{D}_{i,t}(\hat{\theta}_t)] / D_{i,t} = 0$$

where the expectation is taken over the density of the pricing errors, and $D_{i,t}$ is the observed price of the i^{th} asset on date t . The sample analog to the moment condition in equation (5) is obtained by removing the expectation operator.

With the sample moment conditions implied by equation (5) and a set of observed asset prices, the multivariate risk-neutral density may be estimated. For a parametric risk-neutral density with N free parameters, N risky asset prices are needed for identification. In addition, identification requires at least one moment condition containing each parameter, so that a multivariate option price must be observed for each parameter that does not appear in the marginal risk-neutral densities.

For flexible NLS pricing, it is assumed that the observed assets dynamically complete the market. Selection of a parametric form for the risk-neutral density amounts to restrictions on the asset price processes and investor preferences, although the parametric density family chosen allows for a wide range of possibilities.

When there are more observed asset prices than parameters to estimate, it is reasonable to minimize a fit criterion to estimate the parameter vector. Flexible NLS pricing minimizes the sum of squared proportional pricing errors of observed and fitted asset prices. Compared to a squared pricing error criterion, this criterion increases the importance of option price moment conditions relative to underlying price moment conditions. In addition, this criterion more heavily weights moment conditions for out-of-the-money option prices relative to in-the-money option moment conditions. The optimization program is:

$$(6) \quad \underset{\theta_t}{\text{Min}} \sum_{i=1}^N (\hat{D}_{i,t}(\hat{\theta}_t) - D_{i,t}) / D_{i,t}$$

Estimation is performed by a non-linear least squares optimization algorithm, and the fitted prices are calculated using numerical integration. In the univariate NLS case, this estimation procedure uses prices of European put and call options on a single underlying asset. For the bivariate case, put and call options on underlying assets X and Y are included as are bivariate options, such as a spread option.

At this point, the flexible NLS pricing technique appears similar to the univariate NLS technique. What distinguishes flexible NLS pricing is that the multivariate parametric family selected for estimation is tailored to this particular application. The key elements desired in a multivariate risk-neutral parametric family are threefold. First, the multivariate lognormal should be a special case, since this is used in existing GBM-based pricing formulas. Second, the possibility of tail shapes that deviate from a lognormal specification should be allowed, since these characteristics are observed in many financial asset price densities. And, third, an economical parameterization should be used, since at least as many observed asset prices as parameters are needed for NLS estimation. If substantially more data were available a non-parametric or semi-nonparametric method such as Gallant and Nychka (1987) might be applied.

A multivariate parametric density family developed in this paper that satisfies these three conditions is referred to as the flexible density family. It is based on the idea of approximating a multivariate lognormal density by defining the volatility parameters so that they depend on the evaluation point. This adds flexibility to the density in the sense that the flexible probabilities may be increased or decreased at each point by using the variance parameters to match deviations from lognormality. In particular, characteristics such as excess skewness and kurtosis may be modeled.

Initially, consider the problem of matching the deviations of an empirical univariate density function from a univariate lognormal density function. Since the probability density is monotonically increasing in the σ parameter, the $\sigma(X)$ function may be adjusted to match the empirical density at each point.

In the bivariate case, the lognormal density may be generalized to the flexible density function by allowing all parameters to depend on the evaluation point. Attention will be limited in this paper to the case where the volatility parameters are replaced by exponential functions of x and y , and the rest of the parameter functions are constant. Equation (7) defines the bivariate flexible density function in terms of a bivariate lognormal density and replacement of the volatility parameters with functions that depend on additional parameters and the evaluation point.

$$(7) \quad f_{flex}(x, y; \theta) = \kappa f_{\lognorm}(x, y; \mu_x, \mu_y, \sigma_x(\bullet), \sigma_y(\bullet), \rho)$$

This density function is defined for positive values of x and y over a bounded but large support, which ensures that all moments exist. κ is a scaling factor that ensures that the density function integrates to one, and ρ must lie between one and negative one. The three unchanged parameters from the lognormal density no longer have the interpretation of the mean of the log of x and y and the correlation of the log of x and y . These moments must now be calculated by direct integration.

For the purpose of risk-neutral density estimation in this paper, the sigma functions are defined in terms of powers of log-returns as follows.

$$(8) \quad \sigma_x(X_{t+T}; X_t, \alpha, a) = \exp^{\alpha_1 + \alpha_2 \log(X_{t+T}/X_t) + \alpha_3 \log(X_{t+T}/X_t)^2 + \dots + \alpha_a \log(X_{t+T}/X_t)^{a-1}}$$

$$(9) \quad \sigma_y(Y_{t+T}; Y_t, \beta, b) = \exp^{\beta_1 + \beta_2 \log(Y_{t+T}/Y_t) + \beta_3 \log(Y_{t+T}/Y_t)^2 + \dots + \beta_b \log(Y_{t+T}/Y_t)^{b-1}}$$

These volatility functions defined will be referred to as sigma shape polynomials (SSP's). The flexibility of this parameterization is due to the fact that the SSP's may exhibit a variety of shapes corresponding to deviations from lognormality. For example, when $\sigma_x(\bullet)$ increases with the level of x , the probabilities of larger x events increase relative to smaller x events, since the density at x is positively related to the level of σ_x . This corresponds to an increase in skewness and kurtosis in the marginal density of x . The degrees of the SSP's, given by $a-1$ and $b-1$, determine the types of

deviations from lognormality that are possible. As a special case when a and b are 1, the density is bivariate lognormal.

Deviations from lognormality may be conveniently summarized by plotting the sigma shape polynomials against values of x and y . A flat shape corresponds to a lognormal marginal density, an upward parabolic shape corresponds to heavier tails than lognormal, and a downward parabolic shape indicates the opposite. Similarly, an upward sloping curve is associated with positive skewness in the marginal while a downward sloping curve indicates negative skewness relative to a lognormal.

To illustrate the usefulness of the bivariate flexible density function, it is applied to the problem of estimating the bivariate density of daily Dollar-Yen and Dollar-Deutschemark gross exchange returns. The sigma shape polynomials for the flexible density are specified as quadratic, and 2346 daily observations are used from 1987-1995. The flexible density is estimated using the method of moments and nine moment conditions corresponding to the first four sample central moments of each exchange return and the return correlation. A bivariate lognormal density is also estimated for comparison.

Table 1 compares the sample and estimated moments for a bivariate flexible and bivariate lognormal density. The higher moments are what distinguish the two specifications. The estimated bivariate lognormal density matches the means, standard deviations, and correlation of the exchange returns. However, it is unable to capture the skewness and kurtosis in the marginal densities. The flexible density function matches the means, standard deviations, skewness, kurtosis, and correlation. In addition, it comes close to matching the higher cross-moments including coskewness and cokurtosis.

Figure 1 plots the estimated sigma shape polynomials illustrating the directions of deviation from lognormality. The upward parabolic shape of both curves indicates heavier tails than lognormal in both marginals. The Dollar-Yen SSP is higher for positive returns than negative, while the Dollar-DM SSP has the reverse pattern. These asymmetries correspond to positive skewness in Dollar-Yen exchange returns and negative skewness in Dollar-DM returns.

II. Flexible NLS pricing of bivariate foreign exchange contingent claims

In this section, the flexible NLS pricing technique is applied to several types of bivariate contingent claims that depend on foreign exchange rates in 1993 and 1994. Evidence is found that the empirical risk-neutral density has asymmetries and tail shapes not captured by the lognormal specification. Since a multivariate lognormal risk-neutral density is a property of most existing valuation models, rejection of lognormality suggests potential biases in these models. In several cases, there are substantial differences in estimated contingent claims prices using the two densities, indicating that the flexible NLS technique may improve pricing accuracy.

The difficulty in obtaining multivariate contingent claims prices restricts the current applications of flexible NLS pricing. However, foreign currency cross-rate options, when viewed from a local

country perspective, are equivalent to spread options on the individual currencies. This facilitates estimation of bivariate risk-neutral densities on selected currency pairs for which cross-rate options are traded.

The data used in estimation, provided by Philadelphia Stock Exchange, consists of closing prices for European-style Dollar-Yen, Dollar-Deutschemark, and Yen-Deutschemark currency options on 11/19/93 and 2/18/94 with fourteen trading days until expiration. Contemporaneous exchange rates as well as foreign and domestic riskless rates are obtained from the DataStream database.

There are thirteen closing option prices including two cross-rate option prices available on the first estimation date and fourteen option prices available including five on the cross-rate on the second date. To mitigate the effects of non-synchronous trading, all option prices are converted to implied volatilities using the exchange rate at the time of the trade, and then mapped to synchronous prices using the Garman-Kohlhagen (1983) formula and identically timed exchange rates.

These two estimation dates represent quite different market conditions. While July and August 1993 had considerable currency volatility, and the ERM bands were widened to 30% in August, currency markets were relatively stable in November. In contrast, an increase in short term interest rates by the Federal Reserve in early February 1994 and uncertainty about a rate cut by the Bundesbank caused considerable currency volatility during February. These differences should be reflected in the estimated risk-neutral densities.

The moment conditions corresponding to Dollar-Yen and Dollar-DM currency prices, currency option prices, and cross-currency option prices are listed in listed in the appendix as equations A.1-A.8. These moment conditions express the current price as the present value of the expected payoff under the risk-neutral density using equation (3) and the appropriate payoff function.

Equations A.1-A.6, which involve the currency prices and currency option prices, are only sufficient to identify the marginal risk-neutral densities of the Dollar-Yen and Dollar-DM rates. A bivariate option is required to identify the parameter ρ that does not appear in either marginal risk-neutral density. Interestingly, a DM-Yen option that pays off in yen may be converted to a (bivariate) spread option that pays off in cents. Equations A.7 and A.8 provide the pricing equations for cross-rate options priced in U.S. terms and are necessary to identify the cross-moments of the bivariate risk-neutral density.

The bivariate multivariate lognormal density function is specified based on equations (7)-(9) using quadratic standard deviation shape polynomials ($a=3$, $b=3$). The choice of appropriate order for the shape polynomials involves a tradeoff of a precise fit to the data versus a noisy estimate of the density. A quadratic appears to offer a reasonable balance for this dataset. This means that ten parameters, including the scaling factor κ , are to be estimated. The support is defined over the range of -10 to 10 historical objective standard deviations for the exchange rates.

The fourteen day ahead bivariate flexible density on 11/19/93 and 2/18/94 is estimated using the IMSL nonlinear least squares routine (DUNLSF) and the IMSL bivariate integration routine (DTWODQ). Using the same technique, a lognormal risk-neutral density is fit to the data so that the

underlying assets are priced exactly. This procedure corresponds to choosing the implied standard deviations and correlation that best fit the option prices in the proportional least squares sense.

A fundamental test of any asset pricing model is its ability to match existing asset prices. Table 2 compares of the standard deviation of proportional pricing errors using the flexible density and the lognormal. The flexible density reduces the lognormal pricing error by 35% and 72% for the first and second estimation dates.

A comparison of the moments of the flexible and lognormal densities, also in Table 2, reveals that the empirical risk-neutral density exhibits substantial deviations from lognormality. The flexible density, which fits the existing asset prices much more closely, has higher negative skewness than the lognormal for the Dollar-DM marginals on both dates. Apparently, there is more asymmetry in the empirical risk-neutral density than can be characterized by a lognormal specification. In addition, the flexible density kurtosis of both exchange rates on the second date is significantly higher than the lognormal indicating greater risk-neutral probabilities of tail events.

Figures 2 and 3 plot the estimated sigma shape polynomials for the flexible risk-neutral density on each date. A constant sigma shape polynomial for both exchange rates, i.e. two flat graphs for each date, would be consistent with a lognormal density; however, none of the four curves is flat. The Dollar-Yen SSP has a downward parabolic shape on the first date which shifts to an upward parabolic shape on the second date. This indicates a change from negative excess kurtosis to positive excess kurtosis. The Dollar-DM SSP shifts upward from the first to the second date suggesting an overall increase in risk-neutral uncertainty about the 14-day ahead exchange rates.

These results also reveal that the higher moments of the risk-neutral density change over time. These changes may be attributed to changes in the asset price processes, such as those generated by stochastic volatility. Based on the currency market environment in February 1994, it is reasonable to expect greater uncertainty about fourteen day ahead exchange rates on this date compared to November 1993.

Figures 4 and 5 graph the estimated flexible bivariate densities. It is somewhat difficult to compare these with a lognormal visually. The increase in risk-neutral correlation on the second date is evidenced by the narrower hump in Figure 4 compared to Figure 5. There is some indication of heavier risk-neutral tails in Figure 5.

Once the risk-neutral density estimate has been obtained, its primary use is to price new assets. In order to ensure the risk-neutral density is not changed by the addition of new assets, the span of the market must remain constant. If markets are dynamically complete, any new asset may be replicated by a dynamic trading strategy using the primary assets so that the new asset is redundant and the span is unchanged. If markets are not dynamically complete, it is sufficient that the new asset is redundant relative to the existing primary assets. In this case, the risk-neutral density is not uniquely determined by the existing asset prices, but prices of spanned assets under any valid risk-neutral density are uniquely determined.

An advantage of considering all traded options as primary assets, instead of just the foreign currencies and a riskless bond which are used in the standard MVCC pricing techniques, is that dynamic trading strategies involving all the assets may replicate a much richer variety of payoff patterns. If exact replication of the new asset is not possible using all the traded assets, flexible NLS pricing might be viewed as an approximation assuming only small changes in the risk-neutral density due to increased spanning opportunities.

As an example, consider risk-neutral pricing of a new asset, C1, which is a European option that pays off in cents one-hundred times the maximum of zero and the net return on Dollar-Yen and Dollar-DM rates over the period t until $t+T$. Also, consider another European option, C2, that pays off based on the minimum of the two returns. These options might be considered for use as asset allocation tools. Using an estimated bivariate risk-neutral density, these options are priced in cents using the risk-neutral pricing equations given in the appendix by equations A.9 and A.10.

As a second example, consider two European "double digital" options whose payoffs are dependent on the joint risk-neutral probabilities of large increases in both exchange rates or large decreases in both exchange rates. C3 pays off one dollar if both net exchange returns are greater than 5% over the period t until $t+T$, while C4 pays off one dollar if both net exchange returns are less than -5% over the same period. These options are priced in cents using the risk-neutral pricing equations given in the appendix by equations A.11 and A.12.

Using the estimated flexible and lognormal densities, options of these four types with fourteen days until maturity ($T=14$) are priced on 11/19/93 and 2/18/94. Table 3 reports the differences in estimated prices. Pricing differences of up to 5% might be viewed as being within the range of parameter estimation error and bid-ask spread. Differences over 10% might be considered to be economically significant.

On the first date, the prices are fairly close except for C4, a -5% double digital option, which has a 43% lower price using the flexible density function than the lognormal. Prices for all options increase from the first date to the second date. This reflects the increased uncertainty about fourteen day ahead risk-neutral exchange rate probabilities in February 1994. The most dramatic differences are for C3 and C4 which more than double in price.

On the second date, three of the four options exhibit significant pricing differences using flexible NLS pricing compared to lognormal NLS pricing. These differences reflect the asymmetric reallocation of probability mass in the tails of the flexible density, and its lower risk-neutral correlation than the lognormal. For example, the flexible estimate of the C4 price is about 52% lower than the lognormal price.

It is clear that as the deviations between the flexible and lognormal density increase, the option price differences increase. In the preceding analysis of risk-neutral densities on two dates, there is evidence found for deviations from bivariate lognormality that have a substantial impact on estimated prices for bivariate contingent claims.

III. Conclusions

If underlying asset prices are not lognormally distributed or there are difficulties in constructing a riskless hedge for a multivariate contingent claim using the underlying assets, then there are significant reasons to doubt the accuracy of multivariate contingent claims prices obtained using existing valuation formulas. The flexible NLS pricing technique derived in this paper provides an alternative method for valuing multivariate claims that does not rely on these restrictions on probabilities or hedging opportunities.

The flexible NLS pricing technique involves nonlinear least squares estimation of a flexible multivariate risk-neutral density that best fits observed option and underlying asset prices. The flexible density function developed in this paper is specially designed to model the characteristics of multivariate risk-neutral densities including the possibility of excess skewness and kurtosis compared to a multivariate lognormal density.

Flexible NLS pricing and lognormal NLS pricing are applied to bivariate contingent claims that depend on Dollar-Yen and Dollar-DM exchange rates. In several cases, there are substantial differences in estimated prices indicating that the flexible NLS technique may provide more accurate valuation than existing pricing techniques that rely on a bivariate lognormal risk-neutral density. These price differences are due to the inability of the bivariate lognormal density to characterize the observed asymmetry and tail shapes associated with the empirical risk-neutral density.

Appendix

Moment conditions used in estimation of the bivariate Dollar-Yen, Dollar-DM risk neutral density

$$(A.1) \quad C_{Dollar-Yen,t} = e^{-r_w T} E_{\theta_t}^* \left[\max \{ X_{t+T} - K, 0 \} \right]$$

$$(A.2) \quad P_{Dollar-Yen,t} = e^{-r_w T} E_{\theta_t}^* \left[\max \{ K - X_{t+T}, 0 \} \right]$$

$$(A.3) \quad X_t = e^{-r_w T} E_{\theta_t}^* \left[e^{r_w T} X_{t+T} \right]$$

$$(A.4) \quad C_{Dollar-DM,t} = e^{-r_w T} E_{\theta_t}^* \left[\max \{ Y_{t+T} - K, 0 \} \right]$$

$$(A.5) \quad P_{Dollar-DM,t} = e^{-r_w T} E_{\theta_t}^* \left[\max \{ K - Y_{t+T}, 0 \} \right]$$

$$(A.6) \quad Y_t = e^{-r_w T} E_{\theta_t}^* \left[e^{r_w T} Y_{t+T} \right]$$

$$(A.7) \quad C_{Yen-DM,t}^{cents} = e^{-r_w T} E_{\theta_t}^* \left[\max \{ Y_{t+T} - K X_{t+T}, 0 \} \right]$$

$$(A.8) \quad P_{Yen-DM,t}^{cents} = e^{-r_w T} E_{\theta_t}^* \left[\max \{ K X_{t+T} - Y_{t+T}, 0 \} \right]$$

Pricing formulas for four bivariate contingent claims

$$(A.9) \quad C_t^1 = e^{-r_w T} E_{\theta_t}^* \left[100 * \max \left\{ \frac{X_{t+T}}{X_t} - 1, \frac{Y_{t+T}}{Y_t} - 1, 0 \right\} \right]$$

$$(A.10) \quad C_t^2 = e^{-r_w T} E_{\theta_t}^* \left[100 * \max \left(\min \left\{ \frac{X_{t+T}}{X_t} - 1, \frac{Y_{t+T}}{Y_t} - 1 \right\}, 0 \right) \right]$$

$$(A.11) \quad C_t^3 = e^{-r_w T} E_{\theta_t}^* \left[100 * I_{\left(\frac{X_{t+T}}{X_t} - 1 \right) > .05} * I_{\left(\frac{Y_{t+T}}{Y_t} - 1 \right) > .05} \right]$$

where $I_{(\cdot)}$ is an indicator function

$$(A.12) \quad C_t^4 = e^{-r_w T} E_{\theta_t}^* \left[100 * I_{\left(\frac{X_{t+T}}{X_t} - 1 \right) < -.05} * I_{\left(\frac{Y_{t+T}}{Y_t} - 1 \right) < -.05} \right]$$

where $I_{(\cdot)}$ is an indicator function

Bibliography

- Bollerslev, Tim, Ray Y. Chou and Kenneth F. Kroner, "ARCH Modeling in Finance - A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 1992, 52: 5-59.
- Boyle, Phelim P., "A Lattice Framework for Option Pricing with Two State Variables," *Journal of Financial and Quantitative Analysis*, 1988, 23: 1-12.
- Boyle, Phelim P., Jeremy Evnine and Stephen Gibbs, "Numerical Evaluation of Multivariate Contingent Claims," *Review of Financial Studies*, 1989, 2: 241-250.
- Boyle, Phelim P. and Y. K. Tse, "An Algorithm for Computing Values of Options on the Maximum or Minimum of Several Assets," *Journal of Financial and Quantitative Analysis*, 1990, 25: 215-227.
- Derman, Emanuel and Iraj Kani, "Riding on a Smile," *Risk*, 1994, 7: 32-38.
- Gallant, A. Ronald and Douglas W. Nychka, "Semi-Nonparametric Maximum Likelihood Estimation," *Econometrica*, 1987, 55: 363-390.
- Garman, Mark B. and Steven W. Kohlhagen, "Foreign Currency Option Values," *Journal of International Money and Finance*, 1983, 2: 231-237.
- Harrison, J. Michael and David M. Kreps, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 1979, 20: 381-408.
- Harrison, J. Michael and Stanley R. Pliska, "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and their Applications*, 1981, 11: 215-260.
- Heenk, B. A., Angelien Kemna and Ton C. F. Vorst, "Asian Options on Oil Spreads," *Review of Futures Markets*, 1990, 3: 510-528.
- Ho, Thomas, Richard Stapleton and Marti Subrahmanyam, "Multivariate Binomial Approximations for Asset Prices with Nonstationary Variance and Covariance Characteristics," *Review of Financial Studies*, 1995, 8: 1125-1152.
- Johnson, Herb, "Options on the Maximum or Minimum of Several Assets," *Journal of Financial and Quantitative Analysis*, 1987, 22: 277-283.
- Longstaff, Francis A., "Option Pricing and the Martingale Restriction," *Review of Financial Studies*, 1995, 8: 1091-1124.
- Margrabe, William, "The Value of an Option to Exchange One Asset for Another," *Journal of Finance*, 1978, 33: 177-186.
- Pearson, Neil D., "An Efficient Approach for Pricing Spread Options," *Journal of Derivatives*, 1995, 3: 76-91.
- Reiner, Eric, "Quanto Mechanics," In *From Black-Scholes to Black-Holes: New Frontiers in Option Pricing*, London : Risk Magazine/FINEX, 1992.
- Rubinstein, Mark, "Somewhere Over the Rainbow," *Risk*, 1991: 63-66.
- Rubinstein, Mark, "One for Another," In *From Black-Scholes to Black-Holes: New Frontiers in Option Pricing*, London : Risk Magazine/FINEX, 1992.

- Rubinstein, Mark, "Implied Binomial Trees," *Journal of Finance*, 1994, 49: 771-818.
- Rubinstein, Mark, "Return to Oz," *Risk*, 1994, 7: 67-71.
- Sherrick, Bruce J., Scott H. Irwin and D. Lynn Forster, "Expected Soybean Futures Price Distributions: Option-based Assessments," *Review of Futures Markets*, 1990, 12: 275-290.
- Sherrick, Bruce J., Scott H. Irwin and D. Lynn Forster, "Option-Based Evidence of the Nonstationarity of Expected S&P 500 Futures Price Distributions," *Journal of Futures Markets*, 1992, 12: 275-290.
- Shimko, David C., "Bounds of Probability," *Risk*, 1993, 6: 33-37.
- Shimko, David C., "Options on Futures Spreads - Hedging, Speculation, and Valuation," *Journal of Futures Markets*, 1994, 14: 183-213.
- Stapleton, Richard C. and Marti G. Subrahmanyam, "The Valuation of Multivariate Contingent Claims in Discrete Time Models," *Journal of Finance*, 1984, 39: 207-228.
- Stapleton, Richard C. and Marti G. Subrahmanyam, "The Valuation of Options when Asset Returns are Generated by a Binomial Process," *Journal of Finance*, 1984, 39: 1525-1539.
- Stulz, Rene M., "Options on the Minimum or the Maximum of Two Risky Assets: Analysis and Applications," *Journal of Financial Economics*, 1982, 10: 161-185.

Table 1 - Comparison of estimated bivariate flexible and bivariate lognormal densities

Estimated moment comparison

Currency exchange returns (1987-1995), Dollar-Yen and Dollar-DM
 2346 observations, gross daily exchange returns, $X(t)/X(t-1)$, $Y(t)/Y(t-1)$

	Sample moments	Lognormal estimated moments	Flexible estimated moments
Mean Dollar-Yen	1.0002	1.0002	1.0002
Std. dev. Dollar-Yen	0.0071	0.0071	0.0071
Skewness Dollar-Yen	0.5090	0.0212	0.5090
Kurtosis Dollar-Yen	8.7432	3.0008	8.7432
Mean Dollar-DM	1.0001	1.0001	1.0001
Std. dev. Dollar-DM	0.0072	0.0072	0.0072
Skewness Dollar-DM	-0.0047	0.0216	-0.0047
Kurtosis Dollar-DM	5.1634	3.0008	5.1634
Correlation	0.6401	0.6401	0.6401
Coskewness Yen**2,DM	0.1521	0.0120	0.1520
Coskewness Yen,DM**2	0.0416	0.0121	0.0527
Cokurtosis	3.5078	1.8198	2.5482

Figure 1
 Sigma shape polynomials
 Dollar-Yen, Dollar-DM 1 day objective density




Figure 2
Sigma shape polynomials, 11/19/93
Dollar-Yen, Dollar-DM 14 day risk-neutral density




Figure 3
Sigma shape polynomials, 2/18/94
Dollar-Yen, Dollar-DM 14 day risk-neutral density




Figure 4 -
Estimated 14 day bivariate flexible risk-neutral density
11/19/93




Figure 5 -
Estimated 14 day bivariate flexible risk-neutral density
2/18/94

