Quanto Pricing with Copulas

MICHAEL N. BENNETT AND JOANNE E. KENNEDY

MICHAEL N. BENNETT is a research student in statistics at the University of Warwick in Coventry, UK.
m.n.bennett@warwick.ac.uk

JOANNE E. KENNEDY is a senior lecturer in statistics at the University of Warwick in Coventry, UK.
j.e.kennedy@warwick.ac.uk

We study the practical problem of pricing a particular multiasset option, a quanto FX option. The Black model, which assumes a jointly lognormal distribution of asset prices at expiration, is inconsistent with the implied volatility smile for the three relevant currency pairs.

We demonstrate a practical methodology for constructing a model for the joint distribution that is calibrated to all relevant implied volatilities. The margins of this distribution are determined separately in an initial stage. To calibrate the joint distribution to the implied volatility smile on the remaining FX rate, we perturb the dependence structure associated with the Black model (the Normal copula) in order to influence the tail-dependence characteristics of the resulting joint distribution.

We calibrate the model to a number of real-life scenarios involving several maturities and currency arrangements. A well-known ad hoc adjustment to the Black pricing formula often gives lower quanto call prices than those calculated under our transformed copula model. The relative difference in quanto prices with strikes furthest away from at-the-money is occasionally large (10-15%).

The price of any European multiasset option may be written as an integral involving the joint density of the asset prices at expiration. Information regarding this joint density is given by the market prices of financial options written on the underlying assets.

Under the Black model, asset prices are jointly lognormal (under the appropriate equivalent martingale measure), and analytical formulas are available. At the same time, both the marginal distributions and the dependence structure associated with this model may not be consistent with information given by the prices of plain vanilla options, since implied volatilities are strike-dependent.

Constructing an alternative model with realistic properties that may be calibrated to all relevant market data is a common problem in finance. Quanto foreign exchange options, spread options, basket options, and so on are all examples of such multiasset pricing problems. It is important to practitioners to find an effective solution.

In our pricing methodology, we employ a copula function as a model for the dependence structure. A copula provides the link between the multivariate joint distribution and the univariate marginals. This allows us to separate the modeling of the marginal distributions from the modeling of the dependence structure, permitting a two-stage calibration.

The idea of obtaining implied marginal distributions from market prices of vanilla options is certainly not new (see, for example, Dupire [1994]), and there is a growing literature in this area. There is so far no literature on the practical extraction of the entire joint distribution of asset prices from implied volatilities.

The application of copulas in finance has recently attracted a great deal of interest.

Lindsog, and McNeil [2003] provide a survey of copulas and dependence concepts in relation to finance, with particular focus on elliptical, Archimedean, and Marshall-Olkin copulas and applications in insurance risk and market risk (VaR). Practical interest has recently focused on credit derivatives, where simple parametric copulas are used to capture the dependence between default and asset prices (see, for example, Schönbucher and Schubert [2001]). In this case, the dependence structure is given exogenously and is not implied from prices of vanilla options. Rosenberg [2003] applies a similar methodology to the FX problem we consider, estimating the dependence function via a non-parametric method based on historical return data.

For the standard quanto FX option, market prices of relevant vanilla options provide both information regarding the marginal distributions of the two relevant FX rates at expiration and partial information on the dependence structure. Calibration of marginal distributions is accomplished in an initial stage. Given these marginal distributions, we find that the dependence structure associated with the Black model (the bivariate Normal copula) is insufficiently flexible to permit calibration to the final set of available implied volatilities.

We accomplish this final calibration step by modifying the tail-dependence structure of the Normal copula using a suitable parametric transformation. This transformation is based on a generic parametric transformation of a bivariate copula first suggested in Genest [2000]. The model is then calibrated to the prices of vanilla options for various maturities and currency arrangements.

We look at only one multiasset option, but this methodology is clearly applicable to other multiasset options such as spread options. The methodology may also have more general application in many real-world statistical problems where the Normal copula does not provide a sufficiently accurate model for the dependence structure under consideration, such as in scenario analysis or stress-testing.

I. MOTIVATION

We first describe the pricing of a quanto call option, and then outline current market practice.

Quanto Call Option Pricing Problem

In a standard quanto FX call option pricing problem, let $D^i_{t,T}$ denote the time t value in currency i of the zero-coupon discount bond with maturity T. Let $\mathbb{Q}^i$ be the equivalent martingale measure associated with this numeraire. Also let $X^{i,j}_t$, $(i \neq j)$, denote the value in currency i of one unit of currency j at time t. For an arbitrage-free economy, we must have $X^{i,j}_t = (X^{j,i}_t)^{-1}$ and $X^{i,j}_t = X^{k,j}_t / X^{k,i}_t$.

For the quanto pricing problem, we need to consider only three currencies. A standard quanto FX call option is a contract that pays the holder a total of $\kappa(X^{i,j}_t - K)_+$ in the "quanto" currency k, where κ is a known constant, the "conversion factor." Note that κ is simply a scaling factor. Currency i is commonly referred to as the "money" currency and currency j as the "dealt" currency.

As we may renumber currencies at our convenience, without loss of generality we choose Currency 1 to be the quanto currency, Currency 2 to be the money currency, and Currency 3 to be the dealt currency. By standard arbitrage pricing theory, the value of this quanto call option at time zero is given by

$$V_0^{\text{quanto}} = \kappa D^1_{0,T} \mathbb{E}_{\mathbb{Q}^1}[(X^{2,3}_T - K)_+] \tag{1}$$

where $\mathbb{Q}^1$ is the equivalent martingale measure associated with the numeraire $D^1_{t,T}$. Therefore the value of the quanto option is completely determined if the distribution of $X^{2,3}_T$ under $\mathbb{Q}^1$ is known.

This quanto FX option must be priced consistently with the prices of plain vanilla options written on the three relevant underlying FX rates, $X^{1,2}_T$, $X^{1,3}_T$, and $X^{2,3}_T$. The prices of vanilla options on each of these FX rates are determined by the market for a finite set of strikes.

The prices of vanilla options on $X^{1,j}_T$ $(j = 2, 3)$ allow us to recover the marginal distributions of $X^{1,j}_T$ under $\mathbb{Q}^1$, under suitable assumptions regarding the shape of the distribution. This subproblem has received some attention recently in the literature, and we proceed by choosing a simple parameterization of the distribution.

Prices of vanilla options on $X^{2,3}_T$ give information about the distribution of this FX rate under $\mathbb{Q}^2$. A change of measure shows these options provide partial information regarding the joint distribution of $(X^{1,2}_T, X^{1,3}_T)$ under $\mathbb{Q}^1$, since $X^{2,3}_T = X^{1,3}_T / X^{1,2}_T$.

Under a suitable model for the joint density, we may calculate the distribution of $X^{2,3}_T$ under $\mathbb{Q}^1$ and hence the price of the quanto FX option. The problem is therefore to formulate a model so that we capture the dependence structure implied by the third set of option prices (on $X^{2,3}_T$) while incorporating the information provided about the marginal distributions given by the first two sets of prices (on $X_T^{1,2}$ and $X_T^{1,3}$, respectively).

We choose to separate the modeling of the marginal distributions and the dependence structure by specifying the dependence structure in terms of a copula function. This allows us to interchange the method we have used to obtain the marginal distributions with other more sophisticated methods. Once we have calibrated the marginal distributions to the first two sets of option prices, the problem remains to choose a suitable parameterization of the dependence structure so that the associated joint density is well calibrated to the third set of vanilla option prices.

Current Market Practice

The standard approach to pricing quanto FX options is based on a Black model that assumes correlated lognormal dynamics for the forward FX rates. Let $M_{t,T}^{i,j}$ denote the forward FX rate at time t with maturity T quoted in units of currency i per unit of currency j. Then by no arbitrage:

$$M_{t,T}^{i,j} = D_{t,T}^{j} X_t^{i,j} / D_{t,T}^{i}$$

and

$$M_{t,T}^{i,j} = (M_{t,T}^{j,i})^{-1}, \quad M_{t,T}^{i,j} = M_{t,T}^{k,j} / M_{t,T}^{k,i}$$

Under the Black model, we assume the forward rates $M_{t,T}^{1,j}$ follow the process

$$dM_{t,T}^{1,2} = \sigma^{(1)} M_{t,T}^{1,2}\, dW_t^1, \quad dM_{t,T}^{1,3} = \sigma^{(2)} M_{t,T}^{1,3}\, dW_t^2$$

under the equivalent martingale measure $\mathbb{Q}^1$ corresponding to the numeraire $D_{t,T}^1$, where $\sigma^{(1)}$ and $\sigma^{(2)}$ are (strike-independent) volatilities of vanilla options on $X_T^{1,2}$ and $X_T^{1,3}$, respectively, and W_t^1 and W_t^2 are standard Brownian motions under $\mathbb{Q}^1$ such that

$$dW_t^1\, dW_t^2 = \rho\, dt$$

It is straightforward to show that under $\mathbb{Q}^2$:

$$dM_{t,T}^{2,3} = \sigma^{(3)} M_{t,T}^{2,3}\, dW_t^3$$

where

$$(\sigma^{(3)})^2 = (\sigma^{(1)})^2 + (\sigma^{(2)})^2 - 2\rho\sigma^{(1)}\sigma^{(2)} \tag{2}$$

and W_t^3 is a standard Brownian motion under $\mathbb{Q}^2$.

Consider a contract that pays $Y_T = \kappa\,(X_T^{2,3} - K)$ at time T in Currency 1. Taking the expectation under $\mathbb{Q}^1$, the usual pricing formula yields

$$Y_0 = \kappa(M_{0,T}^{2,3} e^{((\sigma^{(1)})^2 - \rho\sigma^{(1)}\sigma^{(2)})(T)} - K)$$

which is zero if the strike K is given by the quanto forward rate:

$$M_{0,T}^{2,3*} = M_{0,T}^{2,3} e^{((\sigma^{(1)})^2 - \rho\sigma^{(1)}\sigma^{(2)})T}$$

In the Black model, $X_T^{2,3}$ is lognormal under $\mathbb{Q}^1$, so evaluating the expectation in Equation (1), we find the price of the quanto option is given by

$$V_0^{\text{quanto}} = \kappa D_{0,T}^1 [M_{0,T}^{2,3*}\Phi(d_1^*) - K\Phi(d_2^*)] \tag{3}$$

where

$$d_1^* = \frac{\log(M_{0,T}^{2,3*}/K)}{\sigma^{(3)}\sqrt{T}} + \frac{1}{2}\sigma^{(3)}\sqrt{T}, \quad d_2^* = d_1^* - \sigma^{(3)}\sqrt{T}$$

and $\Phi(\cdot)$ is the standard cumulative Normal distribution function. Under the Black model, $\sigma^{(3)}$ is the volatility of vanilla options on $X_T^{2,3}$, given by Equation (2). Thus, given the (strike-independent) volatilities $\sigma^{(1)}$, $\sigma^{(2)}$, and $\sigma^{(3)}$, we may deduce the implied correlation ρ under the Black model.

This model is a good benchmark against which to compare more sophisticated models. To understand the deficiencies of this model, it is necessary to look at the implied volatilities corresponding to the plain vanilla options written on the relevant underlying assets. Recall that it is standard market convention to quote a vanilla FX option price in units of its implied volatility, that is, the corresponding value of the volatility parameter that gives the correct market price when entered into the Black formula. Thus the Black formula is simply used as a one-to-one mapping of market prices to Black volatilities.[1]

In practice, the Black volatility corresponding to a quoted vanilla option price is dependent on the strike of the option, indicating that the assumptions underlying the Black model do not hold (this well-known feature is termed the *volatility smile*). Therefore the assumption that $X_T^{1,2}$, $X_T^{1,3}$ are jointly lognormal under $\mathbb{Q}^1$ may be inappropriate.

In general, it is not possible to obtain an analytical form for the distribution of $X_T^{2,3}$. To incorporate the effect of the volatility smile, practitioners often adopt an ad hoc approach and modify the Black formula in Equation (3) as follows. The value of ρ may be determined as usual from the at-the-money volatilities of options on $X_T^{1,2}$, $X_T^{1,3}$, and $X_T^{2,3}$ using Equation (2). Then an ad hoc approximation for the price of the quanto option may be found by substituting the actual strike-dependent volatility $\tilde{\sigma}^{(3)}(K)$ for $\sigma^{(3)}$ in Equation (3), where $\tilde{\sigma}^{(3)}(K)$ is the Black volatility corresponding to the price of the vanilla option $X_T^{2,3}$ with the same strike K as the quanto option in question.

Later we compare these Black-adjusted prices with those obtained by the copula approach. Note that this ad hoc modification of the Black formula does not provide a model for the joint distribution of $X_T^{1,2}$ and $X_T^{1,3}$. Therefore, the only comparison with the copula approach available to us is a simple comparison of calculated quanto prices for a range of strikes.

II. METHODOLOGY

In the following, we give an overview of our proposed pricing methodology and describe how an appropriate transformation of the Normal copula provides a suitable model for the dependence structure.

Overview of Pricing Methodology

The price of a standard quanto option may be evaluated directly, given the distribution of $X_T^{2,3}$ under $\mathbb{Q}^1$, which in turn may be computed from the joint distribution function of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$. We model this joint distribution function so that it is well calibrated to all relevant plain vanilla FX options as follows.

In our modeling framework, we choose to specify the dependence structure between $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ using a function known as a copula. This allows us to separate the modeling of the implied marginal distributions (incorporating information given by market quotes for the prices of vanilla options on $X_T^{1,2}$ and $X_T^{1,3}$), from the modeling of the dependence structure (for which we have only partial information given by the prices of vanilla options on $X_T^{2,3}$).

In the first stage of this methodology, we model the implied marginal distributions by parameterizing each distribution as a mixture of lognormals. Since we subsequently model the dependence structure using a copula function, this method may be replaced with more sophisticated methods without affecting the remainder of the methodology.

Once the marginal distributions have been determined, the joint distribution of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ may be computed under the assumption that the dependence structure is modeled appropriately using a given copula. Under this model for the joint distribution, prices of vanilla options on $X_T^{2,3}$ may be calculated by numerical integration. The problem remains to choose a suitable copula so that that the corresponding joint distribution can be calibrated to market quotes for the prices of vanilla options on $X_T^{2,3}$.

There are infinitely many models for the joint distribution that are well calibrated to all relevant market prices, since the available prices of vanilla options on $X_T^{2,3}$ provide only partial information regarding the dependence structure. We model the dependence structure by perturbing the Normal copula, which is the model for the dependence structure associated with the Black pricing formula.

Under the Black pricing model, implied volatilities are independent of the strike. The correlation parameter ρ controls the overall level of association between $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ and therefore the level of (flat) volatilities of vanilla options on $X_T^{2,3}$. In the general case, implied volatilities are strike-dependent, so the Black model can be calibrated only to at-the-money volatilities [by choosing the value of ρ to satisfy (2)]. If we replace the lognormal marginals by those actually implied by vanilla options on $X_T^{1,2}$ and $X_T^{1,3}$, we find that the parameter ρ of the Normal copula still controls the overall level of association and consequently the general level of volatilities of options on $X_T^{2,3}$.

Given the implied marginal distributions, typically it is necessary to alter the tail-dependence characteristics of the normal copula in order to match the volatility smile of vanilla options on $X_T^{2,3}$. We accomplish this by modeling the dependence structure using a specific parametric transformation of the Normal copula (based on the generic copula transformation suggested in Genest [2000]). The parameter ρ of the original Normal copula still controls the overall level of volatilities of options on $X_T^{2,3}$ under this new model for the joint distribution, provided the transformation function exhibits certain qualitative features.

Calibration to implied volatilities of options on $X_T^{2,3}$ is performed by a non-linear weighted least squares optimization procedure. The resulting joint distribution is consistent with distributional information provided by vanilla options on all three FX rates.

Implied marginal distributions. We now consider the problem of recovering the implied marginal distributions of the FX rates $X_T^{1,j}, (j = 2, 3)$ from market quotes of prices of vanilla options written on these FX rates, which are available for a finite set of strikes.[2] One suitable method for constructing the marginal distributions so that they are consistent with these market data is to assume a suitable parametric form such as a mixture of two lognormal distributions.

This method is computationally simple and efficient, and the optimization appears to be relatively stable (see Bahra [1997]). It is often used as a benchmark for more sophisticated models (see, for example, Bliss and Panigirtzoglou [2002], Campa, Chang, and Reider [1998], and Coutant, Jondeau, and Rockinger [2001]). The mixture of lognormals model clearly reduces to the special case of a single lognormal distribution (Black model) when the mixture parameter is one or zero; this is extremely useful in testing the numerical implementation of the copula model. Both density and distribution functions are analytical functions of the cumulative Normal distribution function.

The mixture of lognormals distribution is calibrated to market data as follows. A vanilla call option on the FX rate $X_T^{1,j}$, $(j = 2, 3)$ with strike K gives the holder a payoff of $(X_T^{1,j} - K)_+$ at time T. Hence, by the usual arbitrage pricing argument, the value of this option at time zero is given by

$$V_0^{1,j} = D_{0,T}^1 \mathbb{E}_{\mathbb{Q}^1}[(X_T^{1,j} - K)_+] \tag{4}$$

This expectation may be written as an integral involving the probability density function of $X_T^{1,j}$ under $\mathbb{Q}^1$.

If the distribution of $X_T^{1,j}$ under $\mathbb{Q}^1$ is assumed to be a mixture of lognormals, analytical formulas are obtained for prices of vanilla options in terms of the standard cumulative Normal distribution function (see Appendix A for details). The density may be fitted to market quotes by performing a suitable optimization to find the density parameters that minimize the squared error between model and market prices of the relevant vanilla options. Notice that the forward FX rate $M_{0,T}^{1,j}$ is the mean of the marginal distribution of $X_T^{1,j}$ under $\mathbb{Q}^1$. This is fitted exactly, reducing the number of free distribution parameters by one.

Note that this method always results in proper distribution and density functions that are both smooth and differentiable; this is important in our derivations, and is not necessarily apparent in some methods.[3] Other methods may have superior numerical properties if at the expense of more complex implementation, but it is possible to substitute an alternative to this method of obtaining the marginal distributions without affecting the remainder of the methodology.[4]

Partial information on joint distribution. We now consider the partial information provided about the joint distribution of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ that is given by the prices of vanilla options on $X_T^{2,3}$. The value in Currency 2 of a vanilla call option on $X_T^{2,3}$ with strike K at time zero is given by

$$V_0^{2,3}(K) = D_{0,T}^2 \mathbb{E}_{\mathbb{Q}^2}[(X_T^{2,3} - K)_+]$$

where $\mathbb{Q}^2$ is the equivalent martingale measure corresponding to the numeraire $D_{t,T}^2$. We show that prices of such options provide information on the values of integrals over specific regions of the joint distribution of $(X_T^{1,2}, X_T^{1,3})$ under $\mathbb{Q}^1$, via a change of measure.

The Radon-Nokodym derivative of $\mathbb{Q}^j$ relative to $\mathbb{Q}^i$ is given by

$$\left.\frac{d\mathbb{Q}^j}{d\mathbb{Q}^i}\right|_{\mathcal{F}_t} = \frac{M_{t,T}^{i,j}}{M_{0,T}^{i,j}}$$

So as a consequence of the Radon-Nikodym theorem:

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}^2}[(X_T^{2,3} - K)_+] &= \mathbb{E}_{\mathbb{Q}^1}\left[\left.\frac{d\mathbb{Q}^2}{d\mathbb{Q}^1}\right|_{\mathcal{F}_t}(X_T^{2,3} - K)_+\right] \\ &= X_0^{2,1}\frac{D_{0,T}^1}{D_{0,T}^2}\mathbb{E}_{\mathbb{Q}^1}[X_T^{1,2}(X_T^{2,3} - K)_+]\end{aligned}$$

Therefore

$$\begin{aligned}V_0^{2,3}(K) &= X_0^{2,1} D_{0,T}^1 \mathbb{E}_{\mathbb{Q}^1}[(X_T^{1,3} - K X_T^{1,2})_+] \\ &= X_0^{2,1} D_{0,T}^1 \int_0^\infty \int_{Kv}^\infty (u - Kv)\, f^{\mathbb{Q}^1}_{X_T^{1,2}, X_T^{1,3}}(u, v)\, du\, dv\end{aligned} \tag{5}$$

where $f^{\mathbb{Q}^1}_{X_T^{1,2}, X_T^{1,3}}(\cdot, \cdot)$ is the joint density of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$. If this joint density is known, the integral in Equation (5) may be evaluated numerically, so vanilla call option prices $V_0^{2,3}(K)$ may be calculated for various strikes K and compared against market quotes.

If the marginal distributions of $X_T^{1,3}$ and $X_T^{1,2}$ are already determined, the model of the joint distribution is completed by specifying a suitable model for the dependence structure of $(X_T^{1,3}, X_T^{1,2})$ under $\mathbb{Q}^1$. The remaining practical problem is to choose the copula to match the partial information given by the integrals (5) corresponding to the market prices of vanilla options on $X_T^{2,3}$ for different strikes K.

Dependence Modeling with Copulas

A bivariate copula is a function that allows us to investigate the scale-invariant dependence structure of two variables. It provides the link between the marginal distribution functions and the joint distribution function. By changing the copula we can alter the dependence structure of the joint distribution without altering the marginal distributions.

Here it is convenient to define a bivariate copula as a bivariate distribution function whose one-dimensional margins are uniform on the interval [0, 1]. Assuming we know the margins, Theorem 1 shows we can obtain the copula distribution from the joint distribution and vice versa.

There is a growing literature on copulas and their application in finance. For formal definitions, theoretical properties, and comparisons of common parametric copulas, see Joe [1997] and Nelson [1998]. Embrechts, Lindsog, and McNeil [2003] provide a survey of copulas and dependence concepts in relation to finance, with particular focus on Marshall-Olkin, elliptical, and Archimedean copulas and applications in insurance risk and market risk (VaR).

Recent work has recognized the potential application of copulas in derivatives pricing, especially in the field of credit derivatives, but the copula is generally specified exogeneously. Usually either the copula is assumed to take a particular parametric form, which may have more desirable dependence characteristics than the Normal copula for the particular application (such as the Student-*t* copula), or the copula is estimated using historical data.

Rosenberg [2003] performs the latter, applying non-parametric techniques to estimate the dependence copula between two FX rates, while using implied volatilities to construct the margins, but the implied volatilities on the cross-FX rate $X_T^{2,3}$ do not match market quotes. Practitioners usually need a model that is well calibrated to current market prices. There is at present no literature regarding the practical application of copulas to option pricing when the copula is calibrated to the prices of vanilla options.

Sklar's theorem. Sklar's theorem explains the role of copulas in describing the relationship between a multivariate joint distribution function and its univariate marginal distributions.

Theorem 1 (Existence): Let H be a standard joint distribution function and let F and G be the margins of H. Then there is a copula C such that for all x, y $\in \bar{\mathbb{R}}$:

$$H(x,y) = C(F(x), G(x)) \tag{6}$$

Corollary: If the marginal distribution functions F and G are strictly increasing, then

$$C(u,v) = H(F^{-1}(u), G^{-1}(v)).$$

Since it is a prerequisite that the marginal densities of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ are continuous, if we have in mind a model for the continuous joint density we may calculate the unique copula associated with this joint density, and vice versa.

Parametric copulas. The most common practical example of a copula is probably the Normal copula, which has the distribution function

$$C_N(u, v \mid \rho) = \Phi_\rho(\Phi^{-1}(u), \Phi^{-1}(v))$$

where $\mathbf{\Phi}_\rho$ is the standard bivariate Normal distribution function with correlation ρ, and $\mathbf{\Phi}$ is the standard univariate Normal distribution function. By the corollary above, this is the copula implied by a standard bivariate Normal joint distribution function with standard univariate Normal marginal distributions. It is clear that this elliptical copula exhibits both positive and negative dependence (for positive and negative values of ρ, respectively).

In the following we shall refer informally to the behavior of the copula in the extremes of the upper-right and lower-left quadrants of the domain as upper and lower tail-dependence, respectively. Thus, if we perturb the Normal copula to increase upper tail-dependence, we will observe that the probability of observing very high values in both variables is higher than under the Normal copula.

The Normal copula defined above describes the dependence structure associated with the Black model. We shall see below that in order to calibrate our joint distribution to vanilla options on $X_T^{2,3}$, we require a dependence model with much greater flexibility. Specifically, the copula must exhibit both positive and negative overall dependence structures, and it must be possible to simultaneously alter

Exhibit 1
Original Integration Region

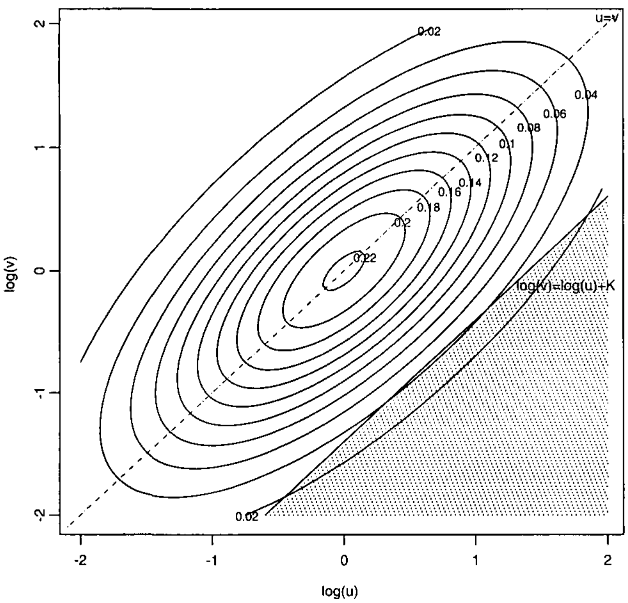

the level of upper-tail and lower-tail dependence.

For most common parametric copulas in the literature, it is not possible to alter the overall level of dependence and the upper and lower tail-dependence characteristics simultaneously. This is a feature of all single-parameter copulas. A few two-parameter copulas have been studied, but these are generally quite inflexible, and it is difficult to construct a multiparameter copula directly that has the desired properties. A copula with all the required characteristics may be constructed by a transformation of the Normal copula, however.

Fitting a Copula Using the Information Available on the Joint Distribution

Calculation of vanilla call prices under a given copula model. Suppose we have in mind a suitable model for the dependence structure of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ in terms of a specific copula. To assess the fit of this copula, we need to be able to calculate the prices of vanilla options on $X_T^{2,3}$.

Using the marginal distributions already determined, it is possible to recover the associated joint density function. From this joint density, prices of vanilla options on $X_T^{2,3}$ may be calculated by numerical integration.

Denote the joint distribution function of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ by

$$F^{\mathbb{Q}^1}_{X_T^{1,2},X_T^{1,3}}(\cdot,\cdot)$$

and the associated marginal distribution functions by $F^{\mathbb{Q}^1}_{X_T^{1,2}}(\cdot)$ and $F^{\mathbb{Q}^1}_{X_T^{1,3}}(\cdot)$, respectively. If the dependence structure between $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ may be modeled appropriately by a known copula with density function $c(\cdot, \cdot)$, then using Sklar's theorem in Equation (6), the joint density is given by

$$f^{\mathbb{Q}^1}_{X_T^{1,2},X_T^{1,3}}(u,v) = f^{\mathbb{Q}^1}_{X_T^{1,2}}(u)\ f^{\mathbb{Q}^1}_{X_T^{1,3}}(v) \times c\left(F^{\mathbb{Q}^1}_{X_T^{1,2}}(u), F^{\mathbb{Q}^1}_{X_T^{1,3}}(v)\right)$$

If the copula density is known, we may now calculate the price of a vanilla option with a given strike by evaluating the integral in Equation (5) numerically. These model prices may then be compared against market quotes.

Reformulating integration region. We now consider the integration regions corresponding to the integral in Equation (5) for different strikes. These integrals correspond to the prices of vanilla options on $X_T^{2,3}$. We need to understand how the Normal copula must be modified so that the resulting joint distribution is well calibrated to market quotes for a given set of strikes. We find that contour plots of copulas versus normal marginals provide a very helpful clue as to how to accomplish this.

Suppose for the moment that the marginal distributions $F^{\mathbb{Q}^1}_{X_T^{1,j}}(\cdot)$ are both lognormal ($j = 1, 2$). Then, for a given strike K, the integration region associated with the integral in Equation (5) corresponds to the region of the copula density below the line $\log(v) = \log(u) + K$ when plotted against normal marginals (*see Exhibit 1*). For high strikes, the integration region is moved further towards the bottom right. Contours shown correspond to a Normal copula with positive correlation.

In the general case, of course, the integration region will not take such a simple form. Since the observed implied marginals will not be too far removed from the single lognormal model, however, any intuition we can gain in this simple setting should carry over to a more realistic model for the marginals.

Suppose now that we wish to modify the Normal copula to fit the strike-dependent volatilities of vanilla call options on $X_T^{2,3}$, but still retain a symmetric structure, i.e., $c(u, v) = c(v, u)$ for all $u, v \in [0, 1]$. If the implied volatility of a vanilla call with a high strike is higher than expected under the Normal copula model, then if we wish our copula model to correctly price this option we must increase the value of the integral (5). This may be accomplished by increasing the mass of the copula density cor-

EXHIBIT 2
Reformulated Integration Region

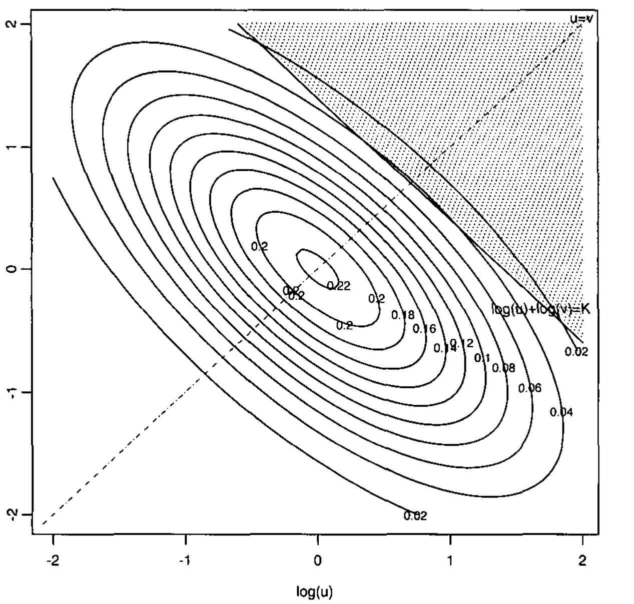

responding to high values of u and low values of v (*see Exhibit 1*). With the symmetry of the copula, this would also increase the mass of the copula density corresponding to low values of u and high values of v, thus also increasing call prices with low strikes simultaneously.

A symmetric structure is a feature of most copulas that are suitable for practical applications. The application of non-symmetric copulas is a much more complex problem. We note that there is very little literature regarding the practical application of non-symmetric copulas. It is therefore desirable to reformulate the model so that symmetric copulas may be used.

If we specify the joint density in terms of $1/X_T^{1,2} = X_T^{2,1}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$, then:

$$V_0^{2,3}(K) = X_0^{2,1} D_{0,T}^1 \times \int_0^\infty \int_{K/v}^\infty (u - K/v)\, f^{\mathbb{Q}^1}_{X_T^{2,1}, X_T^{1,3}}(u,v)\, du\, dv \tag{7}$$

where

$$f^{\mathbb{Q}^1}_{X_T^{2,1}, X_T^{1,3}}(u,v) = f^{\mathbb{Q}^1}_{X_T^{2,1}}(u)\, f^{\mathbb{Q}^1}_{X_T^{1,3}}(v) \times c\left(F^{\mathbb{Q}^1}_{X_T^{2,1}}(u), F^{\mathbb{Q}^1}_{X_T^{1,3}}(v)\right) \tag{8}$$

and both $f^{\mathbb{Q}^1}_{X_T^{2,1}}(\cdot)$ and $F^{\mathbb{Q}^1}_{X_T^{2,1}}(\cdot)$ may be found from $f^{\mathbb{Q}^1}_{X_T^{1,2}}(\cdot)$ and $F^{\mathbb{Q}^1}_{X_T^{1,2}}(\cdot)$ directly (see Appendix A).

If the margins are lognormal, the region of the copula corresponding to the integration region of Equation (7) is illustrated in Exhibit 2 for a given strike. For high strikes, this region is moved further to the upper right. (Contours shown correspond to the same Normal copula show in Exhibit 1.) The values of integrals (7) for different values of K may now be altered independently using a symmetric copula by shifting the mass of the copula density along the line $u = v$.

Thus relatively high call prices for high strikes may be obtained by increasing upper tail-dependence and so forth. By controlling the upper and lower tail-dependence of the copula, it is possible to alter model call prices for different strikes so that they match market quotes.

Equations (7) and (8) allow us to calculate prices of vanilla options under our copula model for any given strike. Of course, under the Black model (corresponding to lognormal margins and a Normal copula with correlation ρ), Equation (7) may be evaluated analytically. The implied volatility associated with these call prices is given by

$$(\sigma^{(3)})^2 = (\sigma^{(1)})^2 + (\sigma^{(2)})^2 + 2\rho\sigma^{(1)}\sigma^{(2)} \tag{9}$$

where $\sigma^{(1)}$ and $\sigma^{(2)}$ are the (flat) implied volatilities of call options on $X_T^{1,2}$, and $X_T^{1,3}$, respectively.[5]

Transformation of Normal copula. The copula that we calibrate to the prices of vanilla options on $X_T^{2,3}$ will not be unique because these prices give only partial information regarding the dependence structure. We would have greater confidence in our inferences if our chosen copula model is close to the Normal copula, which is the dependence model associated with the Black pricing formula. Therefore the copula we choose to fit to market data is a perturbation of the Normal copula.

We construct a new symmetric copula by a suitable transformation of the Normal copula, modifying the tail-dependence structure so that the resulting joint density is well calibrated to prices of vanilla options on $X_T^{2,3}$. Theorem 2 gives sufficient conditions under which the transformed copula is itself a copula (for proof see Appendix B).

Theorem 2 (Genest [2000]): Let φ: $[0, 1] \to [0, 1]$ be a continuous and twice-differentiable concave function such that $\varphi(0) = 0$ and $\varphi(1) = 1$. Then

$$C_\varphi(u,v) := \varphi^{-1}(C(\varphi(u), \varphi(v))) \tag{10}$$

is a copula if $C(\cdot, \cdot)$ is a copula.

The density of the transformed copula C_φ is obtained by taking the partial derivative of (10) with respect to both arguments:

$$c_\varphi(u,v) = \frac{\varphi'(u)\varphi'(v)}{\varphi'(C_\varphi(u,v))}\Bigg(c(\varphi(u),\varphi(v)) - \frac{\varphi''(C_\varphi(u,v))}{[\varphi'(C_\varphi(u,v))]^2}\frac{\partial C}{\partial u}(\varphi(u),\varphi(v))\frac{\partial C}{\partial v}(\varphi(u),\varphi(v))\Bigg) \tag{11}$$

where $c(\cdot, \cdot)$ is the density of the original copula C.

This transformation allows us to construct infinitely many new copulas from any given copula. If our starting point is the normal copula with parameter ρ, then we may use the transformation function to modify the tail-dependence characteristics of the copula as required. Note that the identity transformation $\varphi(x) = x$ does not alter the original copula.

Notice in Equation (10) that the behavior of $\varphi(x)$ near $x = 1$ controls the upper tail-dependence of the transformed copula. To see this, note that for u, v both near one, $C[\varphi(u), \varphi(v)]$ will be near one, so it is the behavior of $\varphi(x)$ and $\varphi^{-1}(x)$ for x close to one that controls the level of $C_\varphi(u, v)$. Similarly, the behavior of $\varphi(x)$ near $x = 0$ controls the lower tail-dependence of the transformed copula.

Henceforth, suppose the original copula C in (10) is the Normal copula. We require that the second derivative $\varphi''(x)$ of the transformation function be near zero around $x = \frac{1}{2}$, or the transformation will significantly alter the overall dependence characteristics of the copula. If this is the case, experimentation shows that the parameter ρ still controls the overall dependence characteristics of the copula. This allows us to use the endpoints of the transformation function to alter the tail-dependence characteristics. Increasing the size of the second derivative of the transformation near zero (one) is seen to increase the lower (upper) tail-dependence of the transformed copula.

Heuristically, how can we see this from Equation (11)? We are performing a slight modification of the normal copula, so the optimal transformation function should be close to $\varphi(x) = x$. For transformation functions with characteristics as described above, increasing the size of $\varphi''(\cdot) < 0$ near zero or one will raise the second term of Equation (11), which will eventually dominate since $\varphi'(\cdot)$ will be close to one and the tails of $c_N[\varphi(u), \varphi(v)]$ decay exponentially. This follows because the quantities $\varphi(\cdot)$, $\varphi(\cdot)$, and $\partial C/\partial u(\cdot, \cdot)$ do not vary as fast as $\varphi''(\cdot)$. The mass of the copula density is shifted toward those regions corresponding to lower or upper tail-dependence.

If the copula C is a Normal copula, then the transformed density function (11) may be evaluated directly because the Normal copula density c_N and the partial derivative $\partial C_N / \partial u$ are analytical (requiring only an approximation to the standard Normal cumulative distribution; see Appendix B), and the distribution function of the bivariate Normal copula $C_N(\cdot, \cdot)$ may be approximated very accurately (see, for example, Genz [2002]).

It remains to parameterize the transformation function φ so that its second derivative may be modified in order to alter the upper and lower tail-dependence as required.

Calibration of transformation function. Given the correlation ρ of the Normal copula that defines the dependence structure of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$, and assuming the transformation function φ has been parameterized and satisfies the conditions of Theorem 2, we can compute the density of the transformed Normal copula. Hence we can find the joint density of $X_T^{1,2}$ and $X_T^{1,3}$ under $\mathbb{Q}^1$ using Equation (8). Vanilla call option prices $V_0^{2,3}(K)$ are computed under this model by evaluating the integral in Equation (7). These prices can then be compared to market quotes.

A non-linear optimization may be performed to find the values of the transformation parameters that give the closest match of model and market vanilla option prices. A standard non-linear least squares algorithm will find the optimal solution $z_* \in \mathbb{R}^n$ that minimizes the objective function:

$$\|r(z)\|_2 = \sum_{i=1}^{m} r_i(z)^2$$

where $n \leq m$.

Let the residuals r_i be given by

$$r_i(z) = w_i(\sigma_i^{\text{observed}} - \sigma_i^{\text{model}}(z)) \text{ for } i = 1, \ldots, m$$

where $\sigma_i^{\text{observed}}$ are the available market implied volatility quotes associated with the prices of vanilla options on $X_T^{2,3}$ with strikes K_i; $\sigma_i^{\text{model}}(z)$ are the corresponding implied volatilities calculated under the transformed copula model with parameter vector $z_* \in \mathbb{R}^n$; and w_i are the given weights.

To modify the tail-dependence structure of the normal copula, it is necessary to construct a transforma-

EXHIBIT 3
Market Data

$x_0^{EUR,JPY} = 0.009781 \quad x_0^{EUR,USD} = 1.175226$

T	EUR	USD	JPY
$\frac{1}{12}$	0.996243	0.996661	0.999948
$\frac{1}{4}$	0.988734	0.990054	0.999805
$\frac{1}{2}$	0.978343	0.980503	0.999566
1	0.958074	0.959798	0.998965
2	0.916243	0.911070	0.997404

Discount factors $D^i_{0,T}$ corresponding to currency i and maturity T.

T	(Puts)		Delta		(Calls)
	10%	20%	50%	20%	10%
$\frac{1}{12}$	0.1289	0.1230	0.1215	0.1260	0.1340
$\frac{1}{4}$	0.1277	0.1210	0.1190	0.1235	0.1319
$\frac{1}{2}$	0.1272	0.1203	0.1180	0.1223	0.1305
1	0.1269	0.1196	0.1170	0.1216	0.1303
2	0.1249	0.1176	0.1150	0.1196	0.1283

EUR/USD implied volatility quotes corresponding to standard values of Black delta and maturity T.

T	(Puts)		Delta		(Calls)
	10%	20%	50%	20%	10%
$\frac{1}{12}$	0.1467	0.1441	0.1480	0.1606	0.1747
$\frac{1}{4}$	0.1466	0.1372	0.1385	0.1487	0.1667
$\frac{1}{2}$	0.1444	0.1329	0.1320	0.1414	0.1597
1	0.1443	0.1320	0.1300	0.1380	0.1557
2	0.1508	0.1358	0.1308	0.1358	0.1508

EUR/JPY implied volatility quotes corresponding to standard values of Black delta and maturity T.

T	(Puts)		Delta		(Calls)
	10%	20%	50%	20%	10%
$\frac{1}{12}$	0.1109	0.1015	0.0950	0.0945	0.0990
$\frac{1}{4}$	0.1156	0.1045	0.0985	0.0995	0.1066
$\frac{1}{2}$	0.1178	0.1055	0.1000	0.1024	0.1122
1	0.1186	0.1065	0.1010	0.1048	0.1166
2	0.1235	0.1105	0.1060	0.1105	0.1235

EUR/USD implied volatility quotes corresponding to standard values of Black delta and maturity T.

tion function with some flexibility, incorporating our prior intuition regarding the shape of a suitable transformation function. Therefore we choose to specify φ as a cubic spline with $k + 1$ suitable predefined knot points $p_i \in [0, 1]$, such that the constraints on φ in Theorem 2 are satisfied.

Since the size of the second derivative gives a good indicator of the increase in tail-dependence, it makes sense to parameterize this directly. Hence φ is found by specifying the second derivative of the spline at each knot point, and then recovering the coefficients of the spline by making use of the requirement that $\varphi(0) = 0$ and $\varphi(1) = 1$. At each optimization step, the first component of $z = (z_1, \ldots, z_n)^T$ may be used to determine the correlation of the normal copula $\rho \in (-1, 1)$ [for example, via $\rho := \sin(z_1)$].

Suppose we choose knot points $p = (p_0, p_1, \ldots, p_k)^T$ with $p_0 = 0$ and $p_k = 1$. Define the cubic spline between the j^{th} and $(j + 1)^{th}$ nodes by

$$S_j(x) := \sum_{i=0}^{3} a_{i,j}(x - p_j)^i \tag{12}$$

for $x \in [p_j, p_{j+1}]$ and coefficients $a_{i,j}$ to be determined. Because $\frac{1}{2} S_j''(x - p_j) = a_{2,j} + 3a_{3,j}(x - p_j)$, setting $a_{2,j} := -z_{j+1}^2 < 0$ $(j = 0, \ldots, k)$ leads to a negative second derivative for all $x \in [0, 1]$.

In this pricing problem, we calibrate to prices of vanilla options for $m = 5$ strikes. Since the objective of the transformation is to alter tail-dependence and not the general dependence structure (which we control by altering ρ), we expect the second derivative of the optimal solution to be near zero around $x = \frac{1}{2}$. Therefore we introduce the condition $\varphi''(\frac{1}{2}) = 0$, reducing the number of free parameters to be fitted by one.

Let knot points be given by $p = (0, 0.1, 0.5, 0.9, 1)^T$ with associated weights $w = (2, 10, 20, 10, 2)^T$ (so volatilities near at-the-money are matched very closely). Then using $n = 5$ parameters, the transformed copula is fully specified by setting $\rho := \sin(z_1)$ and constructing φ as a cubic spline with knot points satisfying:

$$\begin{aligned}
a_{2,0} &:= -z_2^2 \\
a_{2,1} &:= -z_3^2 \\
a_{2,2} &:= 0 \\
a_{2,3} &:= -z_4^2 \\
a_{2,4} &:= -z_5^2
\end{aligned}$$

incorporating the endpoints $\varphi(0) = 0$ and $\varphi(1) = 1$.

At each optimization step, we propose a parameter vector z that corresponds to a value of ρ and a candidate transformation function given by the cubic spline φ. We require that φ satisfy all the conditions in Theorem 2. Since the construction above gives rise to a cubic spline that takes positive values and has negative curvature, it remains to check that $\varphi(x) \leq 1$ for all $x \in [0, 1]$.

But φ is concave and continuous, so this condition must hold unless $\varphi'(1) < 0$. This is trivial to check. If this is the case, a large value of the objective function may be returned to the optimization loop so that the current trial solution z will be discarded. In practice we find this situation does not arise as long we choose the starting point of the optimization $z(0)$ so that this condition holds.[6]

III. RESULTS OF CALIBRATION TO EUROPEAN CALL OPTIONS

Our model of the joint distribution is fitted to actual market quotes for the implied volatilities of plain vanilla options written on exchange rates among the Japaneese yen, the euro, and the U.S. dollar, with strikes corresponding to standard values of the option delta given by 10%, 25%, and 50%, and standard maturities one, three, and six months and one and two years (*see Exhibit 3*). All option price data and corresponding FX rates and discount factors were obtained at the close of trading on July 7, 2001. This market data was kindly provided by West LB.

It is instructive to first consider the joint distribution model provided by the Normal copula with marginal distributions given by market data. Consider as an example the dataset corresponding to maturity one month (1M) and the currency arrangement where (CCY1, CCY2, CCY3) = (USD, EUR, JPY). Implied volatilities of vanilla call options $V_0^{2,3}(K)$ under the Normal copula model with various values of correlation ρ are displayed in Exhibit 4.

As expected, it is observed that changing ρ alters the overall dependence structure (as with the Black model), and therefore results in a parallel shift in model implied volatilities with negligible change in slope or curvature. Therefore it is not possible to find a close match to the market quotes simply by varying ρ.

Exhibit 5 by contrast indicates that in this case the transformed Normal copula offers essentially an exact fit to market quotes (with an objective function of zero to 14 significant places for this dataset), with transformation function $\varphi(\cdot)$ shown in Exhibit 6. This transformation function is constructed by cubic splines, with knot points $p = (0, 0.1, 0.5, 0.9, 1)^T$.

To see the effect of the transformation of the normal copula, compare the contour plot of the optimal transformed copula against that of the Normal copula (*Exhibit 7*) (which is calculated by finding the value of the parameter ρ that minimizes the objective function) against that of the optimal transformed Normal copula (*Exhibit 8*). Although the overall level of association under the optimal transformed copula is similar to that under the

EXHIBIT 4
Implied Volatilities under Normal Copula

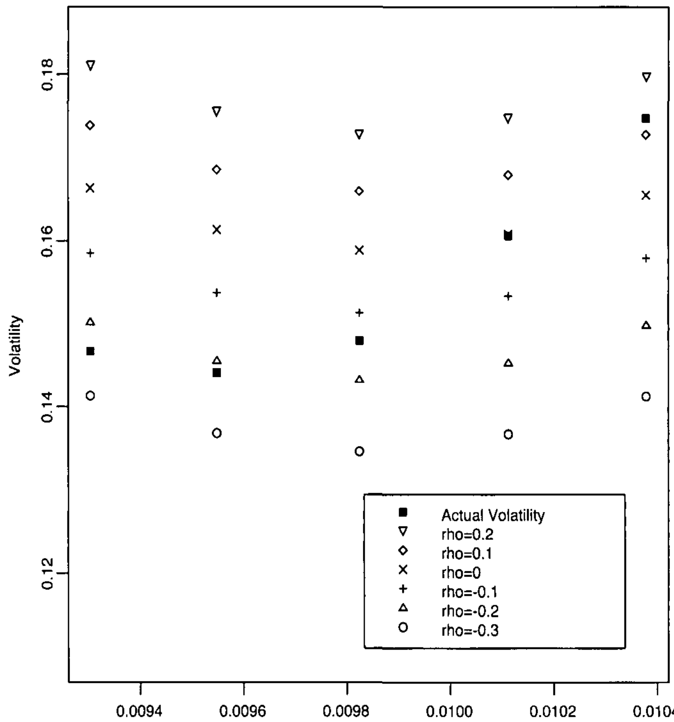

EXHIBIT 5
Implied Volatilities Transformed

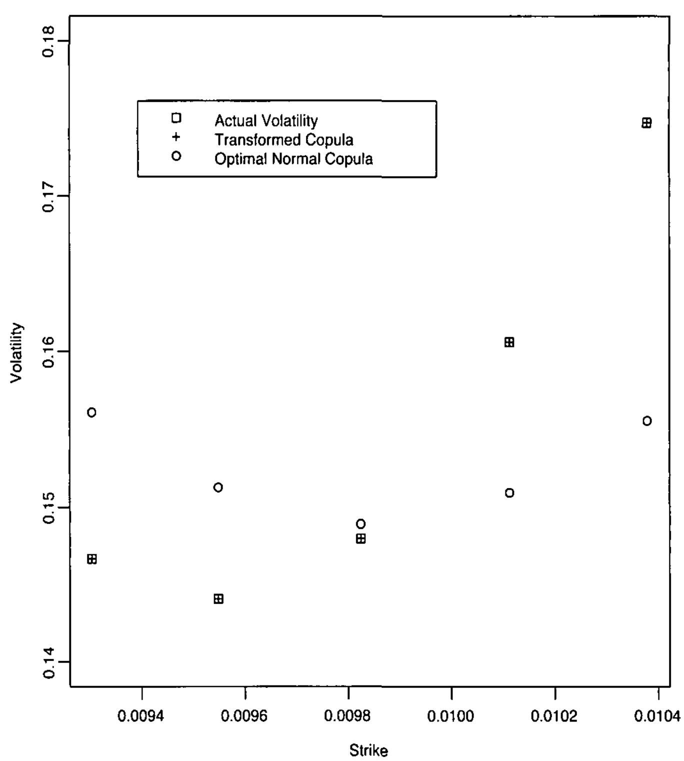

EXHIBIT 6
Optimal Transformation Function φ

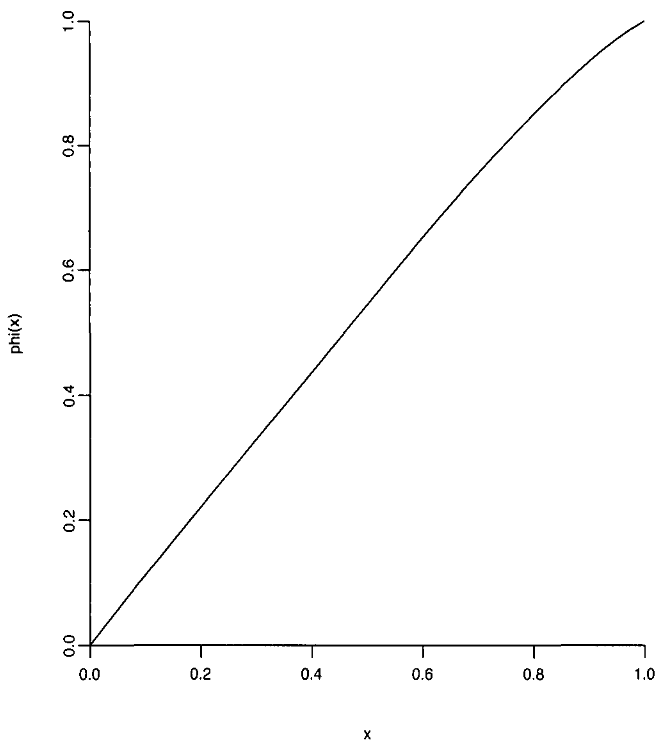

EXHIBIT 7
Optimal Normal Copula

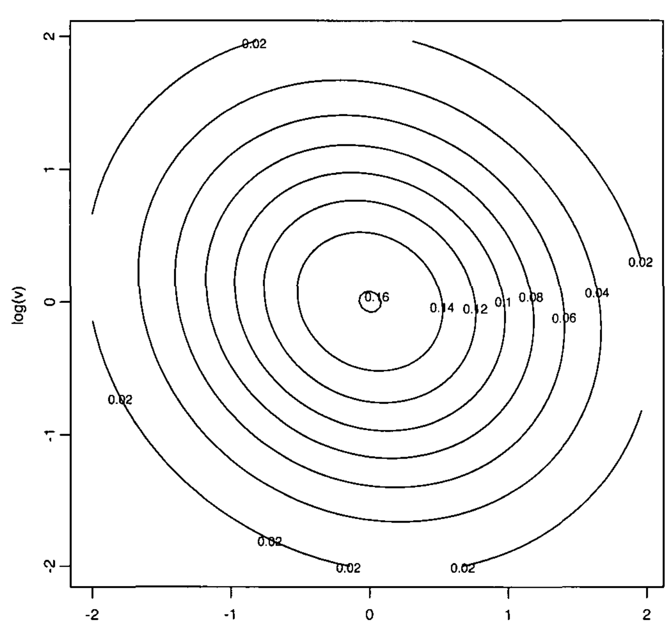

optimal Normal copula, much greater upper tail-dependence is required to fit the implied volatilities of vanilla options on $X_T^{2,3}$ with high strikes. The transformed copula also exhibits increased lower tail-dependence and a much higher central peak. Notice that the optimal transformed copula exhibits much higher upper tail-dependence as well as higher lower tail-dependence than the Normal copula. The resulting upper and lower tail-dependence characteristics of the transformed copula correspond to the graph of $\varphi''(\cdot)$ near one and zero, respectively (*Exhibit 9*). Experimentation shows that in order to attain a good fit to market quotes, a great deal of curvature is required near zero and one, while close to zero near ½.

For all currency set-ups and all five maturities for which we have market implied volatility quotes, the transformed copula model offers a better fit to market data than the Normal copula in all cases. When the smile in implied volatilities is symmetric, the calibration to market data is particularly good (*see Exhibits 10-12*). This is also the case when implied volatilities of options on $X_T^{2,3}$ are skewed but increasing as a function of the strike.

For a given currency set-up, the behavior of φ'' is usually stable across different maturities (*see Exhibit 13*). This is to be expected, since for each currency pair in our dataset the implied volatility smile has the same qualitative features across maturities, so the overall modification of the Normal copula achieved by the calibrated transformation function φ should be the same. In general the dependence structure implied by market quotes of vanilla options on $X_T^{2,3}$ has greater upper and lower tail-dependence than that exhibited by a Normal copula.

When the implied volatilities of options on $X_T^{2,3}$ are heavily skewed and declining, the fit is occasionally not as good. For the currency set-up (EUR, USD, JPY) and maturities 1M and 3M, the market volatility quotes that we calibrate to are almost flat for strikes above ATM rather than the usual more symmetric shape, and it is difficult to fit this skew. The fit of the model may be improved by shifting the knot points appropriately; experimentation shows this works extremely well for the 3M maturity.

Alternatively, notice that the ideal transformation function exhibits a great deal of curvature in its second derivative near zero and one in all cases, and a piecewise linear construction is a poor approximation to this. To allow such curvature and thus permit a better fit in these cases, we may use an alternative parameterization of φ. One possibility is to parameterize φ'' directly as a piecewise polynomial, and then recover φ using the end-point conditions.

EXHIBIT 8
Optimal Transformed Normal Copula

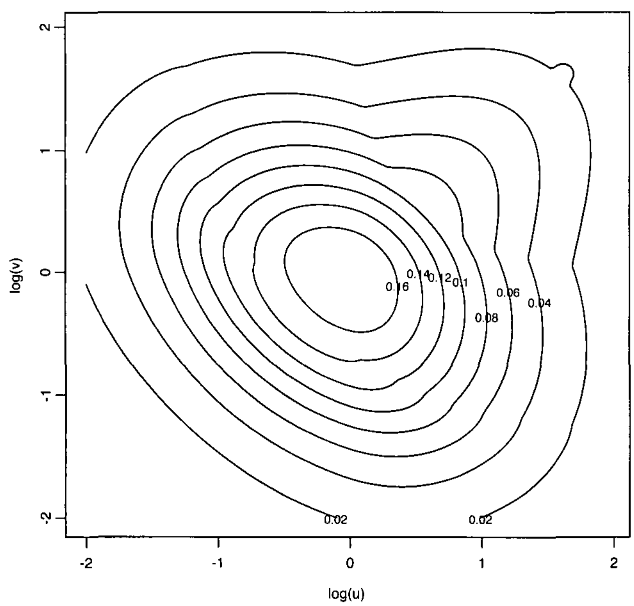

EXHIBIT 9
Second Derivative of Optimal Transformation Function

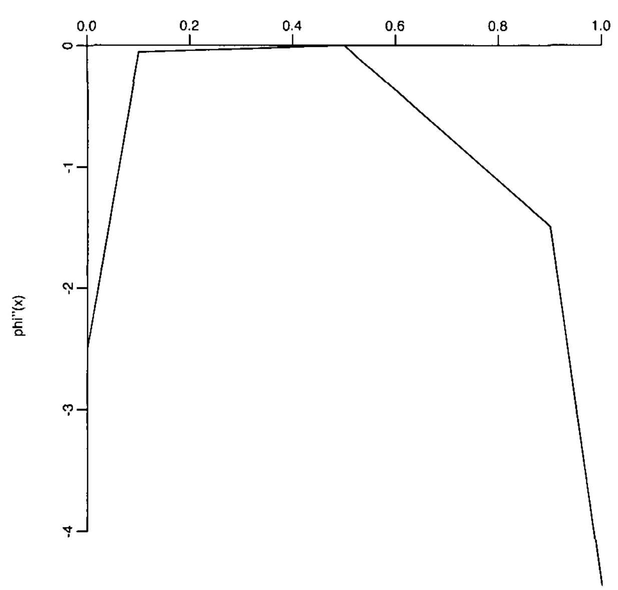

However, the fit to the currency set-up (EUR, JPY, USD) is found to be extremely good for all maturities. Interchanging Currencies 2 and 3 also results in a model that may be used to price any quanto FX option with quanto currency EUR. Therefore for any currency set-up and meturity we are able to find a well-calibrated model for either the joint density $f_{X_T^{2,1},X_T^{1,3}}$ or $f_{X_T^{3,1},\, X_T^{1,2}}$ and we are in a position to value a quanto FX option.

IV. QUANTO OPTION PRICING

We compare prices of a quanto FX option under our model and prices calculated using the Black model and an alternative ad hoc approximation.

Valuation of Quanto FX Options Under a Given Copula Model

The value of a standard quanto FX call option paying out in Currency 1 is computed by evaluating the expectation

$$\mathbb{E}_{\mathbb{Q}^1}[(X_T^{2,3} - K))_+]$$

This is straightforward once the distribution of $X_T^{2,3}$ under $\mathbb{Q}^1$ is known.

Recall $X_T^{2,3} = X_T^{2,1}\, X_T^{1,3}$, so:

$$f^{\mathbb{Q}^1}_{X_T^{2,3}}(z) = \int_0^\infty \frac{1}{x}\, f^{\mathbb{Q}^1}_{X_T^{2,1},X_T^{1,3}}\left(\frac{z}{x}, x\right)\, dx$$

Therefore the price of a standard quanto call option may be evaluated directly by numerical integration:

$$\begin{aligned} V_0^{\text{quanto}} &= \kappa D^1_{0,T} \mathbb{E}_{\mathbb{Q}^1}[(X_T^{2,3} - K)_+] \\ &= \kappa D^1_{0,T} \int_K^\infty \int_0^\infty \frac{1}{x}\, f^{\mathbb{Q}^1}_{X_T^{2,1},X_T^{1,3}}\left(\frac{z}{x}, x\right)(z - K)\, dx\, dz \end{aligned} \quad (13)$$

Standard quanto put prices U_0^{quanto} may be obtained similarly.

Note that the double integral in (13) corresponds to a region of the joint density $f^{\mathbb{Q}^1}_{X_T^{2,1},X_T^{1,3}}(p, q)$ bounded below by the line $p = K/q$. If the margins are lognormal, this corresponds to the region of the contour plot of the copula density above the line $\log(p) + \log(q) = K$ when plotted against normal margins (compare to Exhibit 2). This is the same integration region as a vanilla call option on $X_T^{2,3}$ with strike K. Once the implied marginals have been determined, the value of (13) depends critically on the distribution of the copula mass along the line $u = v$.

In our model we perturb the Normal copula to match information on the dependence structure in vanilla option prices on $X_T^{2,3}$. Given the implied marginal distributions, the price of a quanto option is likely to be highly

EXHIBIT 10
(USD, EUR, JPY) Model Fit

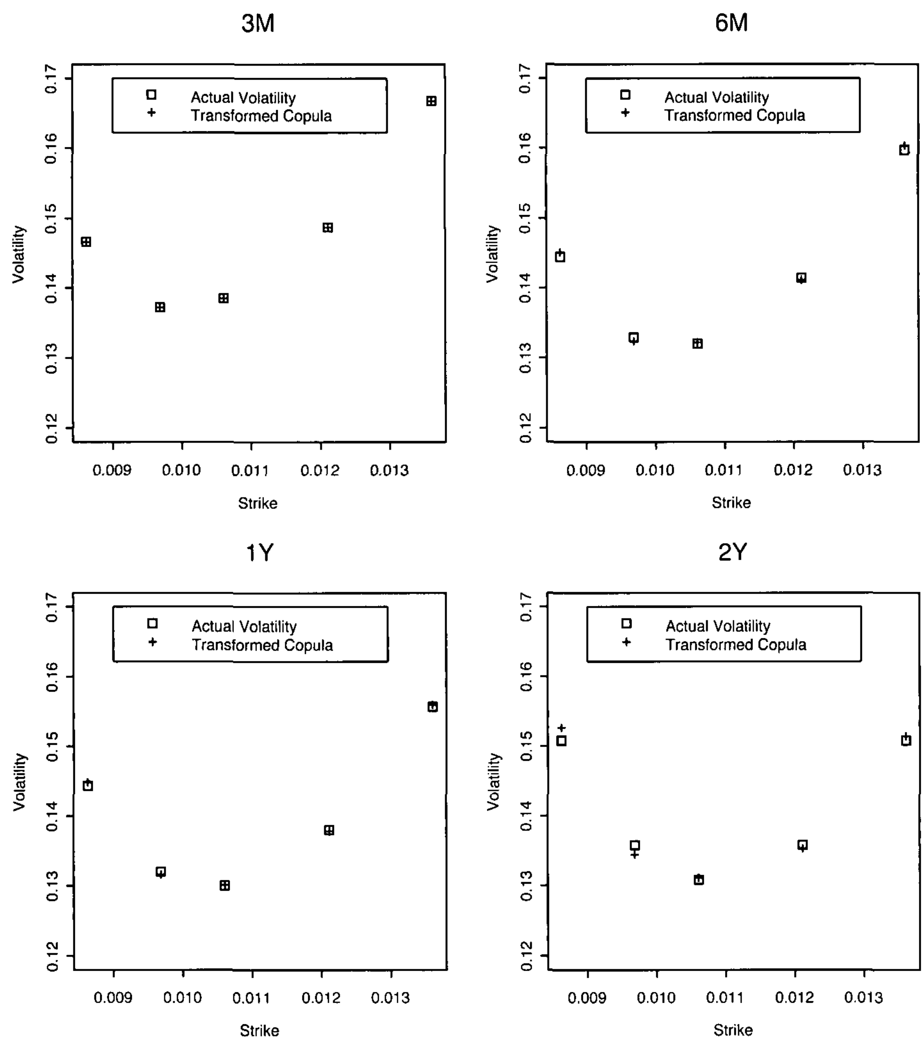

dependent on the fit of the copula to these vanilla option prices. Any distortion of the copula density other than along the line $u = v$ (for instance, any asymmetry) is unlikely to have a material effect on the price of the quanto option. Thus we have more confidence that although our fitted joint distribution model is not the only copula that could be calibrated to market prices, any other copula calibrated to market data is unlikely to result in markedly different quanto prices.

Prices of standard quanto call and put options with the same strike K are linked by the put-call parity relationship:

$$V_0^{\text{quanto}}(K) - U_0^{\text{quanto}}(K) + \kappa D_{0,T}^1 K = \kappa D_{0,T}^1 \mathbb{E}_{\mathbb{Q}^1}[X_T^{2,3}]$$

where the right-hand side is independent of K (but is still dependent on the particular model of the joint distribution). It is a useful numerical check to verify that put-call parity holds for calculated quanto prices under the transformed copula model.

Comparison of Quanto Prices Under Copula Model and Alternatives

We can compare the prices of quanto options calculated under the transformed copula model [evaluating Equation (13)] with both quanto prices calculated under the standard Black model (calibrated to at-the-money volatilities) and prices calculated using the adjusted Black formula. With adjustment of the Black formula, we attempt

EXHIBIT 11
(JPY, EUR, USD) Model Fit

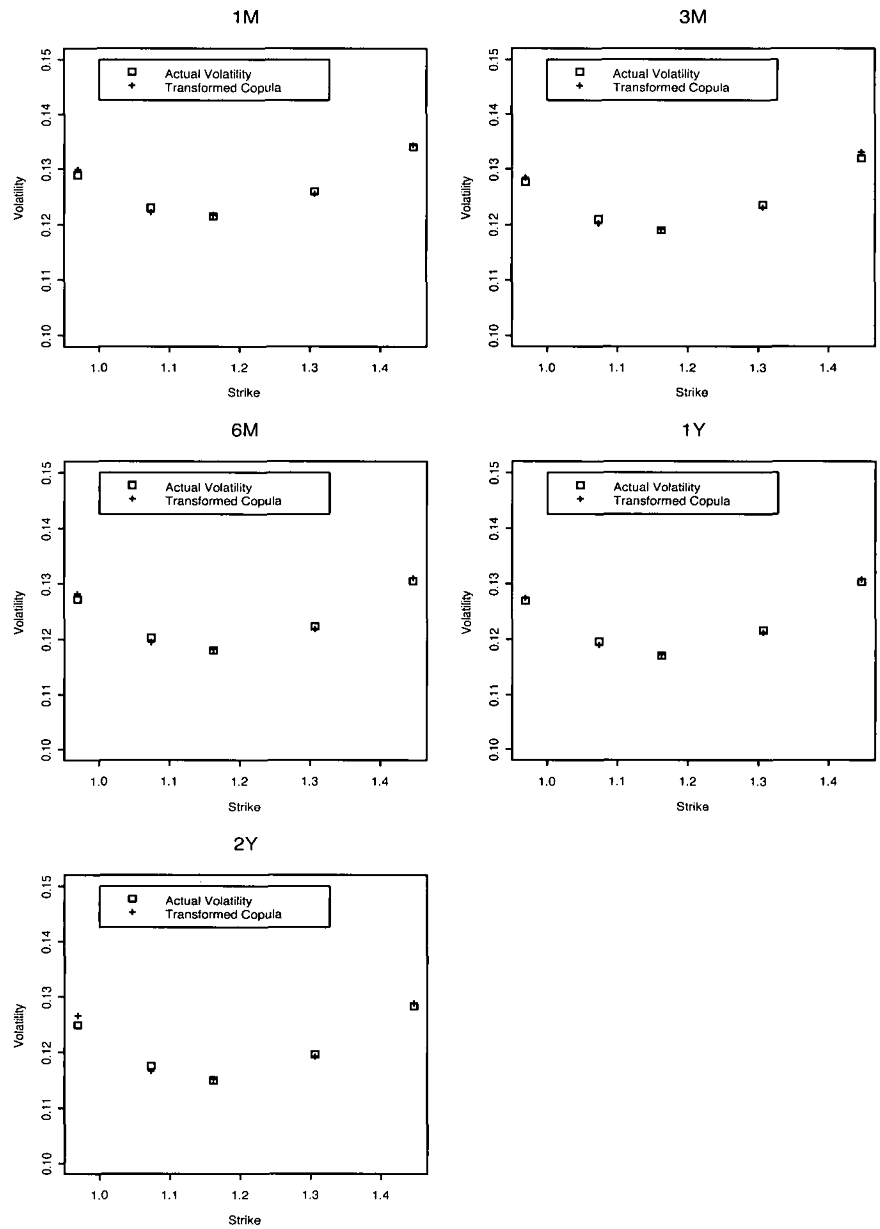

to allow for the smile in implied volatilities on $X_T^{2,3}$, but we ignore the effect of any smile behavior in $X_T^{1,2}$ and $X_T^{1,3}$. This ad hoc formula does not correspond to any joint distribution of $X_T^{1,2}$ and $X_T^{1,3}$, so the only comparison of these approaches available to us is a simple comparison of prices. Note that all prices correspond to given values of the conversion factor κ, but prices for other values of κ are easily obtained since this is just a scaling factor.

The results shown in Exhibits 14-16 indicate that both the adjusted Black formula and the transformed copula model give much higher quanto prices than the simple Black model. This is to be expected. The adjusted Black formula allows the practitioner to compensate somewhat for the increase in implied volatility of $X_T^{2,3}$ for strikes away from at the money (assuming the lowest implied volatility corresponds to the at-the-money strike, which is usually the case). The transformed copula model allows for the smile in implied volatilities of options on all three relevant currency pairs.

We now examine the difference between the prices

EXHIBIT 12
(EUR, JPY, USD) Model Fit

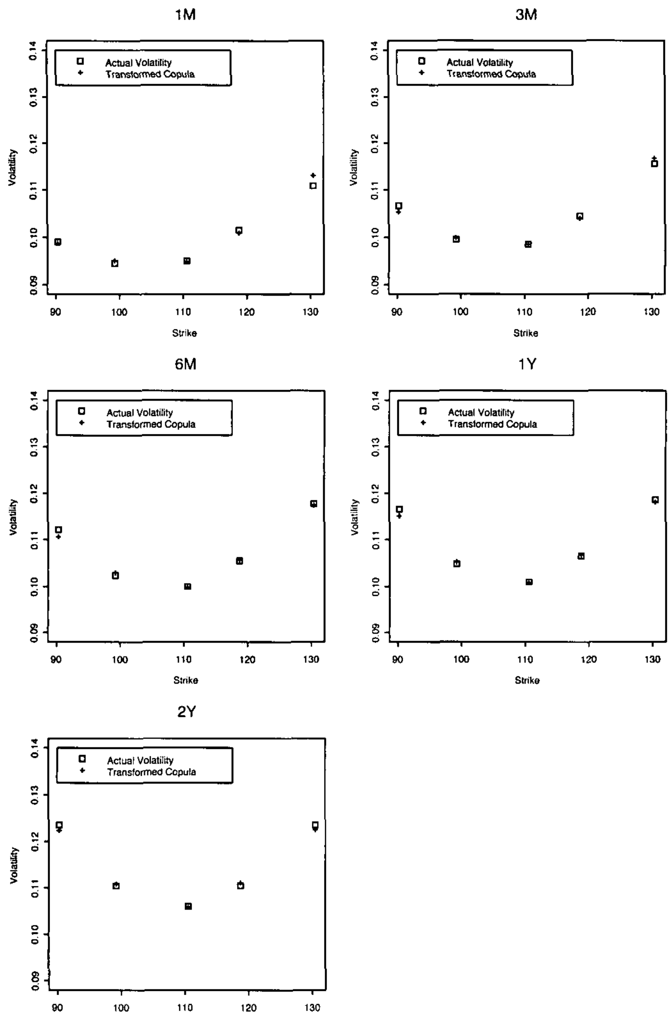

of standard quanto call (or put) options calculated via the adjusted Black formula and the corresponding prices calculated under the transformed copula model. Although call and put quanto prices are generally close, there are sometimes significant differences, usually with strikes farthest from at the money. This difference can be as great as 15% of the option value for standard strikes.

For practitioners familiar with the adjusted Black formula, we also include in the results the value of the adjusted implied volatility parameter $\tilde{\sigma}^{(3)}$ that must be entered into the adjusted Black formula to obtain the model quanto price (calculated under our calibrated transformed copula model). The adjustment required in the size of this parameter is usually in the region of 0.1%–0.3% (in standard units of implied volatility) but is occasionally in the region of 1.0%–1.3% for strikes farthest away from at the money.

For the datasets used in this exercise, the transformed copula model usually gives higher quanto call prices than the adjusted Black formula. The relative difference generally increases with the strike of the option. The absolute difference in prices is usually of similar magnitude across strikes.

The transformed copula model usually leads to lower quanto put prices than the adjusted Black formula. The greatest differences often occur for strikes away from at the money, although there is not such a clear trend as that observed for call prices. In a few scenarios, put prices under the transformed copula model are slightly above those calculated using the adjusted Black formula for some strikes, but on the whole these relative differences are small.

By the results of these scenarios alone, it is not straightforward to determine exactly under what conditions the difference in model and adjusted Black prices will be greatest, because the adjusted Black formula does not correspond to a valid model for the joint distribution of $X_T^{1,2}$ and $X_T^{1,3}$. The transformed copula model provides quanto prices as a result of a calibration to implied volatilities on the three FX rates, and it is not possible to alter the implied volatilities on one FX rate without adjusting the entire calibration.

Further scenario analysis is required to understand the exact relationship between adjusted Black prices and model prices for different patterns of the relevant implied volatilities, but that is beyond the scope of this article.

V. CONCLUSION

We have presented a new methodology for the valuation of quanto FX options that separates the modeling of the dependence structure of the relevant FX rates from the modeling of the implied marginal distributions. Our model for the dependence structure is a copula obtained by a suitable perturbation of the Normal copula. This allows us to modify the upper and lower tail-dependence characteristics appropriately to allow calibration to the smile in implied volatilities on all three relevant foreign exchange rates. Most ad hoc approaches to this problem do not capture all the relevant information embodied in market quotes; indeed, many do not correspond to a proper model for the joint distribution of the relevant FX rates.

Although our model for the dependence structure (the transformed Normal copula) is not the only model that may be calibrated to market prices, it is generally close to the dependence model corresponding to the Black pricing formula. We adjust the dependence struc-

EXHIBIT 13

Second Derivative of Optimal Transformation Function by Currency Arrangement

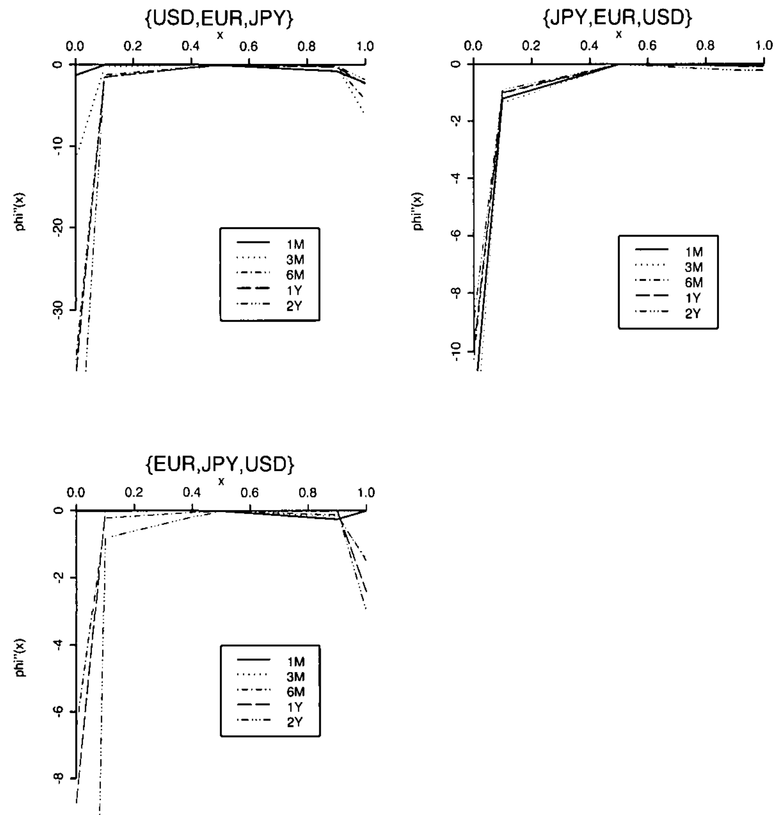

ture so that even though the calibrated dependence model is not unique, we have a degree of confidence that any other model calibrated to all market data will give similar values of the quanto option.

The resulting quanto prices evaluated under a number of real scenarios are generally close to prices calculated using the ad hoc adjusted Black formula, which is certainly a much better proxy for the model price than the standard Black formula. The adjusted Black model often gives lower quanto call prices and higher quanto put prices than those calculated under the transformed copula model. There is occasionally a great difference in quanto prices (10%-15%) for standard strikes farthest away from at the money.

The quanto FX option pricing problem is a particular application of what is a general method for perturbing a normal copula (the dependence model associated with the Black model) to modify the tail-dependence characteristics of the overall dependence structure. A similar methodology may be applied to any two-asset or hybrid option pricing problem if there are plain vanilla option quotes to inform the implied dependence structure. When such option quotes are unavailable (as is the case with equity spread options), historical data may be used to estimate the copula directly. Or, the method we outline could easily be adapted so that the transformed Normal copula is estimated from historical data, allowing construction of a smoother copula that places less reliance on the quality of the historical data and the estimation methods.

This methodology may have other applications in the statistical modeling of real-world dependence where the Normal copula does not provide a realistic enough model for the dependence structure, such as in scenario analysis or stress-testing.

APPENDIX A

Calculation of Implied Marginal Density and Distribution Functions

Initial calibration of marginal distributions to implied volatility quotes is by a weighted non-linear least squares optimization, similarly to Bahra [1997].

Fix the maturity T. Suppose the probability density function $f^{\mathbb{Q}^i}_{X_T^{i,j}}(\cdot)$ of the FX rate $X_T^{i,j}$ under the measure $\mathbb{Q}^i$ is given by a mixture of lognormal distributions

$$f^{\mathbb{Q}^i}_{X_T^{i,j}}(x) = \theta^{i,j} L(\alpha_1^{i,j}, \beta_1^{i,j} \mid x) + (1 - \theta^{i,j}) L(\alpha_2^{i,j}, \beta_2^{i,j} \mid x) \tag{A-1}$$

where $\alpha_k^{i,j} = \mu_k^{i,j} - (\sigma_k^{i,j})^2 T/2$, $\beta_k^{i,j} = \sigma_k^{i,j}\sqrt{T}$, $L(\cdot)$ is the standard lognormal density function where for $k = 1, 2$

$$L(\alpha_k^{i,j}, \beta_k^{i,j} \mid x) = \frac{1}{x\beta_k^{i,j}\sqrt{2\pi}} \exp\left(\frac{-(\log x - \alpha_k^{i,j})^2}{2(\beta_k^{i,j})^2}\right)$$

and $\theta^{i,j}$, $\mu_1^{i,j}$, $\mu_2^{i,j}$, $\sigma_1^{i,j}$ and $\sigma_2^{i,j}$, are parameters to be determined from market quotes of vanilla option prices as follows.

The price of a vanilla FX call option on $X_T^{i,j}$ is given by direct evaluation of the expectation given in Equation (4):

$$\begin{aligned} V_0^{i,j} &= D_{0,T}^i[\theta^{i,j}\{e^{\alpha_1^{i,j}+(\beta_1^{i,j})^2/2}\Phi(d_{1,1}^{i,j}) - K\Phi(d_{2,1}^{i,j}\} + \\ &\quad (1-\theta^{i,j})\{e^{\alpha_2^{i,j}+(\beta_2^{i,j})^2/2}\Phi(d_{1,2}^{i,j}) - K\Phi(d_{2,2}^{i,j}\}] \end{aligned} \tag{A-2}$$

where for $k = 1, 2$

EXHIBIT 14
JPY EUR USD 6M Quanto Prices

Strike	1.0632	1.1189	1.1792	1.2485	1.3139
BS Call	12.53	7.95	4.18	1.62	0.53
Adj BS Call	12.67	8.01	4.18	1.74	0.75
Model Call	12.80	8.10	4.31	1.84	0.85
Call Rel. Err.	1.03%	1.12%	2.87%	5.38%	11.61%
$\Delta\tilde{\sigma}^{(3)}$ (Call)	0.74%	0.30%	0.35%	0.33%	0.50%
BS Put	0.43	1.42	3.69	8.06	13.52
Adj BS Put	0.57	1.48	3.69	8.18	13.73
Model Put	0.55	1.43	3.67	8.13	13.68
Put Rel. Err.	-2.55%	-3.85%	-0.60%	-0.58%	-0.35%
$\Delta\tilde{\sigma}^{(3)}$ (Put)	-0.08%	-0.20%	-0.04%	-0.13%	-0.19%

EXHIBIT 15
EUR USD JPY 6M Quanto Prices

Strike	107.47	112.22	117.78	123.18	128.62
BS Call	0.1087	0.0697	0.0350	0.0146	0.0049
Adj BS Call	0.1104	0.0703	0.0350	0.0162	0.0082
Model Call	0.1112	0.0714	0.0359	0.0172	0.0090
Call Rel. Err.	0.76%	1.53%	2.46%	5.62%	8.66%
$\Delta\tilde{\sigma}^{(3)}$ (Call)	0.55%	0.44%	0.27%	0.33%	0.38%
BS Put	0.0032	0.0106	0.0302	0.0627	0.1062
Adj BS Put	0.0048	0.0112	0.0302	0.0642	0.1095
Model Put	0.0042	0.0108	0.0297	0.0637	0.1088
Put Rel. Err.	-14.43%	-3.42%	-1.92%	-0.78%	-0.63%
$\Delta\tilde{\sigma}^{(3)}$ (Put)	-0.42%	-0.15%	-0.18%	-0.17%	-0.34%

EXHIBIT 16
USD EUR JPY 6M Quanto Prices

Strike	0.008916	0.009440	0.010011	0.010672	0.011299
BS Call	9.8E-04	6.2E-04	3.3E-04	1.3E-04	4.1E-05
Adj BS Call	9.9E-04	6.2E-04	3.3E-04	1.4E-04	7.7E-05
Model Call	1.0E-03	6.3E-04	3.4E-04	1.5E-04	8.7E-05
Call Rel. Err.	1.18%	1.53%	3.33%	6.24%	10.86%
$\Delta\tilde{\sigma}^{(3)}$ (Call)	1.01%	0.52%	0.48%	0.47%	0.63%
BS Put	3.5E-05	1.1E-04	3.0E-04	6.5E-04	1.1E-03
Adj BS Put	4.8E-05	1.2E-04	3.0E-04	6.6E-04	1.1E-03
Model Put	4.7E-05	1.1E-04	2.9E-04	6.6E-04	1.1E-03
Put Rel. Err.	-2.32%	-2.82%	-0.55%	-0.49%	-0.31%
$\Delta\tilde{\sigma}^{(3)}$ (Put)	-0.10%	-0.17%	-0.07%	-0.16%	-0.24%

$$d_{1,k}^{i,j} = \frac{-\log K + \alpha_k^{i,j} + (\beta_k^{i,j})^2}{\beta_k^{i,j}}, \quad d_{2,k}^{i,j} = d_{1,k}^{i,j} - \beta_k^{i,j}$$

and $\Phi(\cdot)$ is the standard Normal cumulative distribution function. Put prices $U_0^{i,j}$ are obtained similarly.

Values of parameters $\theta^{i,j}$, $\mu_1^{i,j}$, $\mu_2^{i,j}$, $\sigma_1^{i,j}$, and $\sigma_2^{i,j}$ are found by a standard least squares optimization procedure that minimizes the weighted squared error between market quotes for implied volatilities of vanilla options and the corresponding values under this model. Since we know the FX forward $M_{0,T}^{i,j}$ is the mean of the distribution of the FX rate $X_T^{i,j}$ under the measure $\mathbb{Q}^i$:

$$M_{0,T}^{i,j} = \theta^{i,j} e^{\alpha_1^{i,j} + (\beta_1^{i,j})^2/2} + (1-\theta^{i,j}) e^{\alpha_2^{i,j} + (\beta_2^{i,j})^2/2} \tag{A-3}$$

hence $\beta_2^{i,j}$ may be found immediately, given the other parameters. A computational requirement of subsequent joint density calculations is that the forward rate implied by the marginal distribution is exact (at least to around 8 significant figures); the optimization procedure is also found to be far more stable if the forward is not allowed to be a free parameter (in contrast with Bahra [1997]).

The distribution function of $X_T^{i,j}$ under $\mathbb{Q}^i$ may be found from the density function (A-1) by direct calculation:

$$F_{X_T^{i,j}}^{\mathbb{Q}^i}(x) = \theta^{i,j} \Phi\left(\frac{\log x - \alpha_1^{i,j}}{\beta_1^{i,j}}\right) + (1-\theta^{i,j}) \Phi\left(\frac{\log x - \alpha_2^{i,j}}{\beta_2^{i,j}}\right)$$

Given the distribution of $X_T^{i,j}$ under $\mathbb{Q}^i$, it is possible to calculate the distribution and density functions of the inverse FX rate $X_T^{j,i}$ under the corresponding measure $\mathbb{Q}^j$, so it is necessary to perform only a single calibration of model parameters for each currency pair $X_T^{i,j}$, $i > j$ under $\mathbb{Q}^i$.

Let us assume the density of $X_T^{i,j}$ under $\mathbb{Q}^i$ is given by a mixture of lognormal distributions as outlined above. By no arbitrage:

$$U_0^{j,i}(K) = K X_0^{j,i} V_0^{i,j}(1/K)$$

Hence, by differentiating both sides with respect to K we see that:

$$F_{X_T^{j,i}}^{\mathbb{Q}^j}(K) = \frac{\theta^{i,j} e^{\alpha_1^{i,j} + (\beta_1^{i,j})^2/2} \Phi(d_{1,1}^{i,j}) + (1-\theta^{i,j}) e^{\alpha_2^{i,j} + (\beta_2^{i,j})^2/2} \Phi(d_{1,2}^{i,j})}{\theta^{i,j} e^{\alpha_1^{i,j} + (\beta_1^{i,j})^2/2} + (1-\theta^{i,j}) e^{\alpha_2^{i,j} + (\beta_2^{i,j})^2/2}}$$

which is the distribution function of a mixture of two lognormal distributions.

APPENDIX B

Copula Calculations

The Normal copula has distribution function

$$C_N(u, v \mid \rho) = \Phi_\rho(\Phi^{-1}(u), \Phi^{-1}(v))$$

where Φ_ρ is the standard bivariate Normal distribution function with correlation ρ, and Φ is the standard univariate Normal distribution function (for which standard approximations are available).

Let $x = \Phi^{-1}(u)$ and $y = \Phi^{-1}(v)$. By standard transformation of variables:

$$c_N(u, v \mid \rho) = (1-\rho^2)^{-1/2} \exp\left(\frac{x^2+y^2}{2} - \frac{1}{2(1-\rho^2)}(x^2+y^2-2\rho xy)\right) \tag{B-1}$$

The partial derivative of the Normal copula with respect to one of its arguments is required in evaluation of the transformed Normal copula density in (11). By integration of (B-1) we obtain:

$$\frac{\partial C}{\partial q}(p, q) = \Phi\left(\frac{\Phi^{-1}(p) - \rho\, \Phi^{-1}(q)}{\sqrt{1-\rho^2}}\right)$$

We transform the Normal copula C_N via

$$C_\varphi(u, v) := \varphi^{-1}(C(\varphi(u), \varphi(v))) \tag{B-2}$$

where φ:[0, 1] → [0,1] is a continuous and twice-differentiable concave function such that $\varphi(0) = 0$ and $\varphi(1) = 1$. This has density given by the partial derivative of (B-2) with respect to both its arguments:

$$\begin{aligned} c_\varphi(u, v) \quad = \quad & \frac{\varphi'(u)\varphi'(v)}{\varphi'(C_\varphi(u,v))}\Bigg(c(\varphi(u), \varphi(v)) \\ & -\frac{\varphi''(C_\varphi(u,v))}{[\varphi'(C_\varphi(u,v))]^2}\frac{\partial C}{\partial u}(\varphi(u), \varphi(v))\frac{\partial C}{\partial v}(\varphi(u), \varphi(v))\Bigg) \end{aligned} \tag{B-3}$$

where $c(\cdot, \cdot)$ is the density of the original copula C.

Following Genest [2000], to show the transformation (B-2) results in a copula, note that if φ is concave then:

$$\frac{\varphi''(\varphi^{-1}(C(u,v)))}{[\varphi'(\varphi^{-1}(C(u,v)))]^2}\frac{\partial C}{\partial u}(u, v)\frac{\partial C}{\partial v}(u, v) \leq 0$$

Hence:

$$c(u, v) - \frac{\varphi''(\varphi^{-1}(C(u,v)))}{[\varphi'(\varphi^{-1}(C(u,v)))]^2}\frac{\partial C}{\partial u}(u, v)\frac{\partial C}{\partial v}(u, v) \geq 0 \tag{B-4}$$

Integrating the transformed density (B-3) between $0 < u_1 \leq u_2 < 1$ and $0 < v_1 \leq v_2 \leq 1$, we obtain:

$$C_\varphi(u_2, v_2) - C_\varphi(u_1, v_2) - C_\varphi(u_2, v_1) + C_\varphi(u_1, v_1) \geq 0 \tag{B-5}$$

Since $\varphi(0) = 0$ and $\varphi(1) = 1$, Equation (B-5) holds for every $0 \leq u_1 \leq u_2 \leq 1$ and $0 \leq v_1 \leq v_2 \leq 1$ with $C_\varphi(u, 0) = C_\varphi(0, v) = 0$, $C_\varphi(u, 1) = u$ and $C_\varphi(1, v) = v$. Thus C_φ is a copula.

Note that from the definition (B-2), if we assume that C_φ is a copula, then the density of the copula (B-3) is positive only if (B-4) holds. Thus Equation (B-4) is actually a necessary and sufficient condition for C_φ to be a copula.

ENDNOTES

[1]Recall that the at-the-money FX option is the option whose strike is given by the current forward FX rate. Options with higher (lower) instrinsic value are known as in-the-money (out-of-the-money) options. Market convention also dictates that the strike of a vanilla FX option is quoted in terms of the corresponding delta, i.e., the value of the delta parameter that gives the correct value of the strike of the option under the Black model, given the implied volatility of the option. By convention, the at-the-money volatility quote is used in this calculation. For a more detailed description of market practice, see, for example, Malz [1997].

[2]Of course, if we knew prices of vanilla options for all strikes, this would completely specify the marginal distributions of $X_T^{1,j}$ (j = 2, 3) (see Dupire [1994]).

[3]We require that the density $f^{1,j}(\cdot)$ be both continuous and differentiable, with $f^{1,j}(0) = 0$, $\int_0^\infty f^{1,j}(u)\, du = 1$ and $f^{1,j}(u) > 0$ for all $u > 0$. In some smoothing methods where call volatilities are interpolated and differentiated with respect to the strike to obtain the distribution (following Dupire [1994]), lognormal tails are pasted onto the central part of the distribution, and it is not always straightforward to smooth the density function across the join while ensuring $\int_0^\infty f^{1,j}(u)\, du = 1$.

[4]The mixture of lognormals model has the disadvantage that the single optimal fit available is not guaranteed to match market quotes with sufficient accuracy in all cases. More sophisticated smoothing methods usually have a goodness-of-fit parameter allowing a closer fit to market quotes at the cost of reduced smoothness. Blitz and Panigirtzoglou [2002] indicate that these methods may offer a more stable fit by observing the effects of perturbing market quotes. Also see Ait-Sahalia and Lo [1998], Malz [1997], and Sherrick, Garcia, and Tirupattur [1996].

[5]Note ρ changes sign; see Equation (2).

[6]For instance, one suitable starting point is the vector $z = (c, 0, \ldots, 0)^T$ for some constant $c \in \mathbb{R}$. This corresponds to a transformation function given by the identity transformation, resulting in a Normal copula with correlation determined by the value of c [e.g., $\rho = \sin(c)$].

REFERENCES

Ait-Sahalia, Y., and A.W. Lo. "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." *Journal of Finance*, 52(5) 1998, pp. 499-548.

Bahra, B. "Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application." Working paper, Bank of England, 1997.

Bliss, R.R., and N. Panigirtzoglou. "Testing the Stability of Implied Probability Density Functions." *Journal of Banking and Finance*, 26(2-3) (2002), pp. 381-422.

Campa, J.M., P.H.K. Chang, and R.L. Reider "Implied Exchange Rate Distributions: Evidence from OTC Option Markets." *Journal of International Money and Finance*, 17(1)(1998), pp. 117-160.

Coutant, S., E. Jondeau, and M. Rockinger. "Reading Interest Rate and Bond Futures Options Smiles: How PIBOR and Notional Operators Appreciated the 1997 French Snap Election." *Journal of Banking and Finance*, 25(11) (2001), pp. 1957-1987.

Dupire, B. "Pricing with a Smile." *Risk Magazine,* 7 (January 1994).

Embrechts, P., F. Lindsog, and A. McNeil. "Modelling Dependence with Copulas and Applications to Risk Management." S. Rachev, ed., *Handbook of Heavy Tailed Distributions in Finance.* Oxford: Elsevier, 2003, pp. 329-384.

Genest C. Presentation at the Fourth International Symposium on Distributions with Given Marginals and Statistical Modelling, Barcelona, 2000.

Genz, A. "Numerical Computation of Rectangular Bivariate and Trivariate Normal and t Probabilities." *Statistics and Computing*, 14(3) (2004), pp. 151-160.

Joe, H. *Multivariate Models and Dependence Concepts.* London: Chapman and Hall, 1997.

Malz, M. "Estimating the Probability Distribution of the Future Exchange Rate from Option Prices." *The Journal of Derivatives,* 5 (Winter 1997), 18-36.

Nelson, R.B. *An Introduction to Copulas: Lecture Notes in Statistics.* Berlin: Springer-Verlag, 1998.

Rosenberg, J.V. "Non-Parametric Pricing of Multivariate Contingent Claims." *The Journal of Derivatives,* 10 (Spring 2003).

Schönbucher, P.J., and D. Schubert. "Copula-Dependent Default Risk in Intensity Models." Working paper, University of Bonn, 2001.

Sherrick, B.J., P. Garcia, and V. Tirupattur. "Recovering Probabalistic Information From Option Markets: Tests of Distributional Assumptions. *Journal of Futures Markets*, 16(5) (1996), pp. 545-560.