Quantitative Finance Volume 4 (December 2004) 1–12 Research Paper
Taylor & Francis Ltd tandf.co.uk

Bivariate normal mixture spread option valuation

Carol Alexander[1,3] **and Andrew Scourse**[2]

[1] ISMA Centre, University of Reading, Reading, UK
[2] ABN-Amro Bank, London, UK

E-mail: coalexander@blueyonder.co.uk

Received [redacted], in final form [redacted]
Published [redacted]
Online at www.tandf.co.uk/journals/titles/14697688.asp DOI:10.1080/14697680400016174

Abstract

We discuss the pricing and hedging of European spread options on correlated assets when the marginal distribution each asset return is assumed to be a mixture of normal distributions. Being a straightforward two-dimensional generalization of a normal mixture diffusion model, the prices and hedge ratios have a firm behavioural and theoretical foundation. In this 'bivariate normal mixture' (BNM) model no-arbitrage option values are just weighted sums of different '2GBM' option values that are based on the assumption of two correlated lognormal diffusions, and likewise for their sensitivities. The main advantage of this approach is that BNM option values are consistent with both volatility smiles *and* with the implied correlation 'frown'. No other 'frown consistent' spread option valuation model has such straightforward implementation. We apply analytic approximations to compare BNM valuations of European spread options with those based on the 2GBM assumption and explain the differences between the two as a weighted sum of six second-order 2GBM sensitivities. We also examine BNM option sensitivities, finding that these, like the option values, can sometimes differ substantially from those obtained under the 2GBM model. Finally, we show how the correlation frown that is implied by the BNM model is affected as we change (a) the correlation structure and (b) the tail probabilities in the joint density of the asset returns.

1. Introduction

Spread options are a simple multi-asset derivative. They occur in a wide range of markets, for example credit spread options in fixed income markets, crack spread options in the energy market and index spread options in the equity markets. As spread risk inherently involves correlation risk, spread options are an important tool in hedging and trading correlation. Hedging with spread options can involve buying and selling calls and puts with European, Bermudan or American exercise conditions. Pricing American style spreads is a difficult problem and is not considered here. The interested reader may wish to consider the tree method of Boyle (1988) as a first approximation.

In this paper we consider only European call spread options with pay-off function $[S_1-S_2-K]^+$ where S_1 and

[3] Author to whom any correspondence should be addressed.

T&F (2005 Style) [18.2.2005–7:22pm] [1–12] [Page No. 1] {TandF}Rquf/RQUF-101600.3d (RQUF) First Proof RQUF-101600

S_2 are the prices of the two underlying assets and K is the strike of the option. Extension of the methods developed in this paper to the put option case is straightforward. Valuation of spread options with negative strikes is based on the fact that a call on a negative strike is the same as a put on the reverse spread with a positive strike of the same absolute value. That is, for $K < 0$ and a risk neutral measure Q:

$$\begin{aligned} &E_Q\{\exp(-rT)[S_1 - S_2 - K]_+\} \\ &\quad = E_Q\{\exp(-rT)[|K| - (S_2 - S_1)]^+\}. \end{aligned} \tag{1}$$

In general, no analytic formulae exist for the pricing of multivariate contingent claims such as spreads and, despite considerable research in this area, there is no universally prevalent pricing framework. A very detailed and informed survey of recent research on the valuation of European spread options is given in Carmona and Durrleman (2003a). As well as highlighting the theoretical and computational problems associated with pricing equity and interest rate spread options, Carmona and Durrleman place particular emphasis on commodity spread options and include several example applications to energy markets.

Some of the earliest work on spread options comes from the analysis of 'out-performance' or 'exchange' options. These are just spread options with a strike of zero and, under the assumption of correlated lognormal diffusions for the underlying assets, there is a Black and Scholes (1973) type solution for the option price (Margrabe 1978). Unfortunately, there is no corresponding general closed form for spreads with non-zero strikes, although Kirk (1995) provides a useful analytic approximation. The core issue preventing exact analytic pricing is that a linear combination of correlated lognormals is not lognormal.

One of the earliest publications on spread option valuation is Boyle (1988) who proposes a three-dimensional binomial tree approach. Although computationally burdensome, this approach is flexible enough to be extended to American spread options. For the case of European options, more exact and less challenging methods to program have been developed. It is unrealistic to base spread option values on the assumption that the spread itself has dynamics governed by a univariate diffusion process, because in that case the distribution of the spread would be independent of the correlation between the underlying assets[4]. Hence most of the European spread option valuation models that have been considered in detail in the academic literature have adopted the assumption of two correlated lognormal diffusions (Ravindran 1993, Shimko 1994, Kirk 1995, James, 2002 and others)[5]. To distinguish this approach from modelling the spread itself as a lognormal diffusion, we shall refer to the two correlated lognormal diffusion model as the '2GBM' model. An important extension of the 2GBM model is to include stochastic volatility, so that the individual asset price dynamics are consistent with their implied volatility smiles. Some research provides fast implementation of accurate numerical methods (e.g. Dempster and Hong 2000) while other research focuses on an approximate closed form (e.g. Durrleman 2001, Carmona and Durrleman 2003b).

Dempster and Hong (2000) develop an efficient method to price spread options using the Fast Fourier Transform (FFT). The key is to ensure that the characteristic function of the joint density of asset price movements can be found analytically. Especially in multi-factor models, for example in stochastic volatility and stochastic interest rate environments, this approach can yield great computation time savings over the Monte Carlo approach. They illustrate this by implementing their model in a stochastic volatility environment. The PhD thesis by Hong (2001) provides an extensive analysis of the effect of stochastic volatility on spread option values employing FFT techniques that are sufficiently flexible to admit a stochastic component correlation process[6].

Durrleman (2001) provides upper and lower bounds for spread option prices which may be only a few percent apart for some parameter values, and offers an arbitrarily precise approximation method for spread pricing, based on the value of certain high-order derivatives. Carmona and Durrleman (2003b) developed even more sophisticated price bounds and the numerical results they provide are consistently and impressively accurate.

Few other models have gone beyond the two-factor correlated GBM assumption. A notable exception is the paper by Pliska and Duan (2003) that generalized the discrete time version of the 2GBM model to include the possibility of cointegration between the assets. This is important for pricing longer-term spread options, where asset price cointegration can lead to substantial differences in option value, compared to those based only on correlation. Their results showed that prices with and without cointegration were only significantly different in the presence of stochastic volatility.

Finally, Rosenberg (2001) uses a non-parametric arbitrage-free pricing approach coupled with copulas to price spreads based wholly on observed market data.

[4] But see Carmona and Durrleman (2003a) for a survey of such methods.

[5] For example, Kirk (1995) derives an accurate approximation for the spread price in the special case that spreads are valued on futures or forward contracts within the Black (1976) model, and Ravindran (1993) reduces the pricing problem to taking the expectation of a function of one random variable. This approach yields efficient spread prices under the GBM assumption, but it is difficult to extend this to other distributional assumptions.

[6] Regarding the component correlation process, Hong (2001) finds that if the terminal correlation is negative (resp. positive) the instantaneous correlation rises (resp. falls) with time to maturity.

The results show that lognormal specifications for marginal densities provide very poor descriptions of the higher moments of observed market data, and that the Gaussian copula is an inadequate description of the asymmetric dependence in the data. This non-parametric approach to valuation underlines the importance of accurately modelling the joint distribution to correctly price the spread, and was a motivating factor in the development of the model considered in this paper.

A deficiency of the 2GBM spread option valuation model is the simplicity of the volatility and correlation assumptions. Without adjustment, for instance to include stochastic volatilities for each asset, these models even fail to capture the non-normality of the marginal returns distributions of the underlying assets. Ignoring the volatility smiles of both underlying assets as well as the structure of the implied correlation between these assets can lead to substantial mispricing, because it is likely that the probability of a large spread movement will be under-estimated when using a bivariate normal distribution for the log asset prices.

The aim of this paper is to develop a simple valuation model for European spread options that is consistent with the implied correlation 'frown' that is commonly observed for spread options. Few other frown consistent spread option valuation models have been developed, indeed the authors are only aware of the jump diffusion approach of Carmona and Durrleman (2003b). The approach taken in our paper is not to introduce jump diffusion. Instead we extend the '2GBM' assumption that the two assets follow correlated lognormal processes to the assumption that the two assets follow correlated lognormal *mixture* diffusion processes. This introduces more structure into asset correlations so that, for instance, correlation in the tails of the joint density can be quite different to their correlation in ordinary market circumstances. This bivariate lognormal mixture (or 'BNM') model gives European spread option values that are consistent with both volatility smiles *and* with the correlation frown. The outline of the paper is as follows: section 2 defines the notion of implied correlation and links the existence of a correlation frown to a joint density for asset returns that has heavier tails than the bivariate normal density; section 3 describes the theoretical framework for the model and illustrates the BNM–2GBM value differences with a simple example; section 4 interprets the value difference as a weighted sum of six second-order 2GBM spread option sensitivities and examines the properties of the BNM option sensitivities; the BNM model is a natural framework for capturing trader's uncertainty in correlation and section 5 shows how this uncertainty can influence the shape of the correlation frown; and section 6 summarizes and concludes with a call for further research on this model.

2. The implied correlation 'frown'

To motivate this section let us first fix ideas by considering implied volatility, which is well understood and clearly defined. In a symmetric volatility smile, such as those often observed in currency option markets, implied volatility of in-the-money (ITM) and out-of-the-money (OTM) options is greater than the at-the-money (ATM) implied volatility—hence the term 'volatility smile'. The smile arises because traders perceive a greater probability of large price changes than is assumed in the BS model. The perceived *leptokurtic* asset return density leads to market prices of ITM and OTM options that are greater than BS prices and, all other variables being fixed in the BS model, the only way that the BS model can explain these market prices is to increase the volatility—because of the positive relationship between volatility and option price, increasing the volatility will increase the BS price[7].

In the BS setting there is only one parameter, the volatility, that is not directly observable. So, given an option price, the notion of implied volatility is clearly defined. Even in a stochastic volatility setting we can define implied volatility as the average of instantaneous volatility over the life of the option. However, in the case of spread options there are (at least)[8] three parameters that are not directly observable: the two asset volatilities and the correlation between these assets. So, given a spread option price, how can one 'imply' three distinct parameters? Clearly, defining an implied correlation is not an easy task. But if we are somehow able to fix the values of the other two parameters, i.e. the volatilities, a unique value for implied correlation can be defined. For instance, one might obtain prices on single asset options for each of the two assets in the spread and imply the two volatility values from these. The two implied volatilities would need to be of the same maturity as the spread option. This may require some interpolation or extra-polation of market prices, but that is not necessarily a difficult problem. However, the *strikes* of the two implied volatilities have a more complex relationship with the strike of the spread option. Problems arise because there are an infinite number of price pairs (S_1, S_2) for which $S_1 - S_2 = K$. Some convention is required to fix the strikes of the single asset options. One possibility is to take a strike on S_2 of $(K - S_1)$ and a strike on S_1 of $(K - S_2)$, but clearly other strike conventions are possible. Given the convention we can define the implied correlation of a

[7] Asymmetric smiles (often called 'skews' or 'smirks') are prevalent in equity markets. Here OTM puts often command substantially higher market prices than OTM calls due to demand–supply imbalances: investors seek the downside protection of these options which relatively few institutions write.

[8] There may be more parameters, for instance if one assumes average correlation is higher in the 'tails' of the joint density than it is in the 'core' of the density.

spread option as the correlation that is implicit in the 2GBM model, for the model option value to be equal to an observed market value of the option, *given* that the two volatility parameters are given by the BS implied volatilities defined according to our chosen strike convention. Note that different strike conventions will give different values for the implied correlation, so implied correlation is not a well-defined quantity. It can only be defined conditional on a given strike convention.

Implied correlation will vary with the moneyness of a spread option, just as the Black–Scholes (BS) implied volatility varies with the moneyness of a single asset option. But in contrast to implied volatility, the implied correlation often decreases and the option moves away from at-the-money. That is, implied correlation 'frowns', even though implied volatility smiles!

An argument very similar to that above can be used to link the implied correlation frown to the leptokurtosis of the joint density for asset returns, as follows: if traders perceive a joint density for asset returns that has heavier tails than the bivariate normal, then market prices of ITM and OTM spread options will be greater than those based on the 2GBM assumption. Now, the correlation between asset returns has a negative relationship with the price of a spread option. That is, the lower (or more negative) the correlation, the more valuable is the spread option. Therefore, to match these higher market prices within the 2GBM framework, the implied correlation for OTM and ITM spread options must be *lower* than the implied correlation used for ATM spread options. Hence it is more appropriate to call the variation of implied correlation with moneyness a correlation 'frown' rather than a correlation 'smile'.

The existence of a correlation 'frown' implies that tail probabilities are underestimated in the 2GBM framework. That is, the 2GBM assumption is not consistent with the observed market prices of OTM call and put spread options because asset returns have leptokurtic densities. So, when we speak of a 'frown consistent' valuation model, we mean a model with leptokurtic asset price densities and for which a correlation frown that is similar to the market frown appears in the 2GBM implied correlations that are 'backed-out' from model prices.

3. The bivariate normal mixture (BNM) model

There is a large research literature on the applications of lognormal mixture asset price densities to *single* asset option valuation (see Ritchey 1990, Melick and Thomas 1997, Bingham and Kiesel 2002 and many others). The model described below can be regarded as the two-dimensional generalization of the normal mixture diffusion (NMD) model introduced by Brigo and Mercurio (2001). This is a no-arbitrage single asset option valuation for which there is a unique representation of the local volatility function, consistent with lognormal mixture asset price dynamics. Although the original interpretation of the normal mixture diffusion was one of deterministic local volatility, recent results in financial econometrics show that the normal mixture GARCH model converges to a stochastic normal mixture diffusion in the continuous time limit (Alexander and Lazar 2004)[9]. Here the transition densities and the marginal densities are normal mixtures, which greatly facilitates the pricing of path-dependent options. The BNM spread option valuation model simply extends the stochastic normal mixture diffusion to the case of two, correlated underlying assets. In this paper we only use the model to price European options—our purpose is to understand the determinants of the correlation frown. However, there is no doubt that this model has great potential for pricing American and other path-dependent spread options.

Let $s_1(t)$ and $s_2(t)$ denote the log prices of the two assets at time t. The value of a European spread option with maturity T will depend on the joint density of s_1 and s_2 at time T. In particular, the value at time t of the option will depend on the drift and volatility over the remaining time to expiry, $\tau = T - t$. We assume that each marginal log price density at time T as seen at time t is given by a mixture of two normal components. That is,

$$\begin{aligned} f_{1,t}(s_1) &= \lambda_1 \phi(s_1; \mu_{11}(\tau), v_{11}(\tau)) \\ &\quad + (1-\lambda_1)\phi(s_1; \mu_{12}(\tau), v_{12}(\tau)), \\ f_{2,t}(s_2) &= \lambda_2 \phi(s_2; \mu_{21}(\tau), v_{21}(\tau)) \\ &\quad + (1-\lambda_2)\phi(s_2; \mu_{22}(\tau), v_{22}(\tau)). \end{aligned} \tag{2}$$

Here $v_{ij}(\tau)$ is the τ-period variance of the jth normal component for asset i, and the notation $\varphi(x; \mu_{ij}(\tau), v_{ij}(\tau))$ denotes a normal density function with mean $\mu_{ij}(\tau)$ and variance $v_{ij}(\tau)$[10]. We assume without loss of generality that each $\lambda_i > 0.5$ $\{i = 1, 2\}$ so that each marginal density has

- a 'core' normal density with weight λ_i in the mixture and with the lower volatility, and
- a 'tail' normal density with weight $1 - \lambda_i$ in the mixture and a higher volatility.

[9] Two important points are worth noting about this model. First, it is not equivalent to the uncertain volatility lognormal model of Brigo (2002) and Mercurio (2002) because the transition densities differ, although the marginal densities are the same. Secondly, an interesting corollary of the Alexander and Lazar (2004) result is that the normal GARCH model converges to a simple deterministic local volatility model and NOT a diffusion model, contrary to popular belief. That is, when general limiting assumptions that can be generalized to more than one variance component are made. The original assumptions of Nelson (1990) were very specific and can only be applied in the simple normal GARCH case.

[10] When $\mu_{ij}(\tau) = r_i - v_{ij}(\tau)/2$ where r_i denotes the (constant) risk-free rate for asset i these can be taken as the risk-neutral densities (see Brigo and Mercurio 2001).

Since each marginal log price density has only two normal components, the bivariate normal mixture (BNM) joint density at expiry may be written

$$
\begin{aligned}
f_t(s_1, s_2) = {} & \lambda_1\lambda_2 \boldsymbol{\Phi}(s_1, s_2; \boldsymbol{\mu}_1(\tau), \mathbf{V}_1(\tau)) \\
& + (1-\lambda_1)\lambda_2 \boldsymbol{\Phi}(s_1, s_2; \boldsymbol{\mu}_2(\tau), \mathbf{V}_2(\tau)) \\
& + \lambda_1(1-\lambda_2) \boldsymbol{\Phi}(s_1, s_2; \boldsymbol{\mu}_3(\tau), \mathbf{V}_3(\tau)) \\
& + (1-\lambda_1)(1-\lambda_2) \boldsymbol{\Phi}(s_1, s_2; \boldsymbol{\mu}_4(\tau), \mathbf{V}_4(\tau)),
\end{aligned} \tag{3}
$$

where Φ is the bivariate normal density function. The bivariate normal mixture has four component bivariate normal densities that we can label the 'core', 'tail–core', 'core–tail' and 'tail' components, respectively. From henceforth, for brevity, we shall drop the τ-dependence notation. Hence we denote the mean vectors and covariance matrices in the bivariate normal mixture (3) by

$$
\begin{aligned}
&\boldsymbol{\mu}_1 = (\mu_{11}, \mu_{21})2, \quad \boldsymbol{\mu}_2 = (\mu_{12}, \mu_{21})2, \\
&\boldsymbol{\mu}_3 = (\mu_{11}, \mu_{22})2, \quad \boldsymbol{\mu}_4 = (\mu_{12}, \mu_{22})2, \\
&\mathbf{V}_1 = \begin{pmatrix} v_{11} & \mathrm{cov}_1 \\ \mathrm{cov}_1 & v_{21} \end{pmatrix}, \quad \mathbf{V}_2 = \begin{pmatrix} v_{12} & \mathrm{cov}_2 \\ \mathrm{cov}_2 & v_{21} \end{pmatrix}, \\
&\mathbf{V}_3 = \begin{pmatrix} v_{11} & \mathrm{cov}_3 \\ \mathrm{cov}_3 & v_{22} \end{pmatrix}, \quad \mathbf{V}_4 = \begin{pmatrix} v_{12} & \mathrm{cov}_4 \\ \mathrm{cov}_4 & v_{22} \end{pmatrix}.
\end{aligned}
$$

In subsequent equations, to ease interpretation, we label volatilities and correlations with a mixture number and with C or T to denote 'core' or 'tail'. Thus the core (i.e. the smaller) volatilities are $\sigma_{\mathrm{iC}} = (v_{i1})^{1/2}$ and the tail (i.e. the larger) volatilities are $\sigma_{\mathrm{iT}} = (v_{i2'})^{1/2}$. The core correlation is $\rho_{\mathrm{CC}} = \mathrm{cov}_1/\sigma_{1\mathrm{C}}\sigma_{2\mathrm{C}}$, the correlation between the tail component of asset 1 and the core component of asset 2 is $\rho_{\mathrm{TC}} = \mathrm{cov}_2/\sigma_{1\mathrm{T}}\sigma_{2\mathrm{C}}$, and so forth. The overall volatility of asset i will be denoted σ_{i} and their 'overall' correlation in the bivariate density is denoted ρ.

We write a spread option price as $P(\sigma_1, \sigma_2, \rho)$, dropping the explicit dependence on asset prices, strike, time to expiry, interest rates and so forth. From (3) we have that $P_{\mathrm{BNM}}(\sigma_1, \sigma_2, \rho)$, that is, the price of a European spread option under the BNM assumption, is a probability weighted sum of four different 2GBM prices. That is,

$$
\begin{aligned}
P_{\mathrm{BNM}}(\sigma_1, \sigma_2, \rho) = {} & \lambda_1\lambda_2 P_{\mathrm{2GBM}}(\sigma_{1\mathrm{C}}, \sigma_{2\mathrm{C}}, \rho_{\mathrm{CC}}) \\
& + (1-\lambda_1)\lambda_2 P_{\mathrm{2GBM}}(\sigma_{1\mathrm{T}}, \sigma_{2\mathrm{C}}, \rho_{\mathrm{TC}}) \\
& + \lambda_1(1-\lambda_2) P_{\mathrm{2GBM}}(\sigma_{1\mathrm{C}}, \sigma_{2\mathrm{T}}, \rho_{\mathrm{CT}}) \\
& + (1-\lambda_1)(1-\lambda_2) P_{\mathrm{2GBM}}(\sigma_{1\mathrm{T}}, \sigma_{2\mathrm{T}}, \rho_{\mathrm{TT}}).
\end{aligned} \tag{4}
$$

3.1. Uncertainty in volatility and correlation

The BNM model provides a natural framework for uncertainty in both volatility and correlation. It is well known that uncertainty in volatility is consistent with, and therefore can explain at least part of, an implied volatility smile (see Hull and White 1987, Heston 1993 and many others). In this paper we show that uncertainty in correlation can explain at least part of the implied correlation 'frown'. However, we do not frame this uncertainty in terms of variance or covariance—although these are the basic parameters for the density—because the basic parameters for the option price are the volatility and correlation. We therefore define

$$E(\sigma_i) = \lambda_i \sigma_{i\mathrm{C}} + (1-\lambda_i)\sigma_{i\mathrm{T}} \tag{5}$$

for $i = 1,2$. Thus $E(\sigma_i)$ is the expected value of the volatility of asset i assuming that the uncertainty over volatility is defined by the mixing law of its lognormal mixture marginal density. The mixing law, therefore, captures the trader's Bayesian beliefs about the possible state of volatility. More details about the trader's state of knowledge can be found in Alexander and Lazar (2004).

It is intuitive to hold uncertainty over volatility, not variance, because BS prices of simple European ATM options are approximately linear in volatility. This linearity implies that the expected price of an ATM option, under volatility uncertainty, will equal the BS price. On the other hand, ATM options are a concave increasing function of variance so if the uncertainty were over variance, not volatility, we would conclude that the BS model always over-prices simple European ATM options when volatility is uncertain, and we would hope that this is not the case[11].

However, 2GBM prices of ATM European spread option are not necessarily linear in volatility so this nice property does not generalize to the bivariate case. In general, 2GBM prices of ATM European spread options are convex functions of volatility so the expected price of an ATM spread option under volatility uncertainty may be greater than the 2GBM price.

Figure 1 graphs the 2GBM value of an ATM European call spread option as a function of the volatility of one of the assets. The current price of each asset is taken as 100. In this paper we employ a simple definition of moneyness, as $S_2 - S_1 - K$ where S_1 and S_2 are the prices of the two underlying assets and K is the strike of the option. Hence with our definition of moneyness the ATM strike is zero. The volatility of asset 1 is fixed at 20%, and the volatility of asset 2 varies between 0% and 50%. In figure 1(*a*) the ATM option values are illustrated for three different asset return correlations (−0. 8, 0, and 0. 8). The graphs show that the ATM spread option value is approximately linear in volatility for $\rho = -0.8$, but as ρ increases the value function becomes more convex, and for $\rho = +0.8$ it is a strictly convex, non-monotonic function of volatility. Figure 1(*b*) shows that the convexity of the 2GBM value of an ATM spread option with respect to volatility increases with maturity. Hence the difference

[11] Note that Hull and White (1987) do conclude that BS prices of ATM options are greater than their stochastic volatility model prices because in their framework it is variance, rather than volatility, that is uncertain.

T&F (2005 Style) [18.2.2005–7:22pm] [1–12] [Page No. 5] {TandF}Rquf/RQUF-101600.3d (RQUF) First Proof RQUF-101600

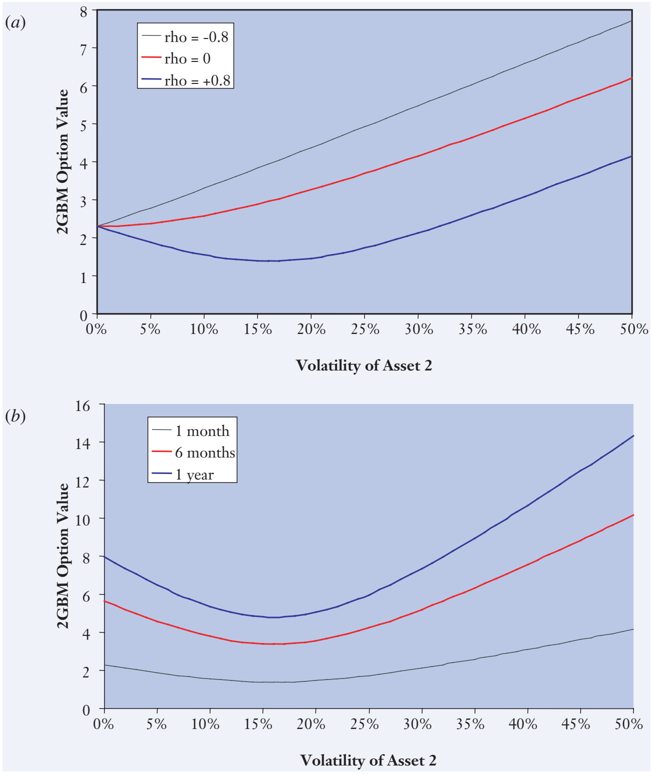

Figure 1. 2GBM values of an ATM European call spread option as a function of the volatility of asset 2. (*a*) For a 1 month option and different asset correlation. (*b*) For asset correlation +0.8 and different option maturities.

between the uncertain volatility spread option price and the 2GBM option price will increase with maturity. We shall return to this point in figure 2 below.

Similarly the mixing law of the bivariate normal mixture density defines uncertainty over correlation, not covariance. That is,

$$E(\rho) = \lambda_1\lambda_2\rho_{CC} + (1-\lambda_1)\lambda_2\rho_{TC} + \lambda_1(1-\lambda_2)\rho_{CT} + (1-\lambda_1)(1-\lambda_2)\rho_{TT}. \quad (6)$$

Note that it does not make sense to equate $E(\rho)$ to the 'ATM correlation' because this correlation is not uniquely defined. There are an infinity of price pairs (S_1, S_2) for which a spread option of strike K will be ATM and these pairs cover a whole 'corridor' in the (S_1, S_2) domain. Some pairs will lie in the extremes of the joint density governed by the 'tail' correlation, others will lie in the 'core' and still others lie in the 'core–tail' regions. Thus, in the BNM framework, or indeed for any joint density that is governed by more than one correlation, a different correlation will apply to different price pairs (S_1, S_2) for which the option is ATM. In any uncertain correlation framework—where by definition correlation is not constant over the domain of (S_1, S_2)—the concept of 'ATM correlation' is difficult to apply.

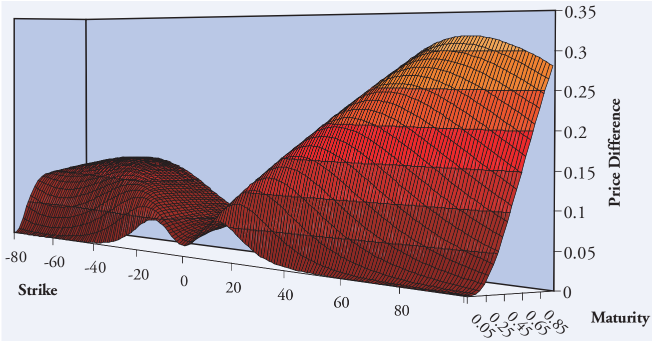

Figure 2. Difference between BNM and 2GBM values of European call spread option as a function of strike. All option parameters as in (7) and (8).

In general, 2GBM European spread option values are a concave monotonic decreasing function of the correlation between asset returns and a convex function of the volatility of each asset return. Figure 1 has shown that this is the case even for ATM options. Thus the expected price of the option will always be less than or equal to the 2GBM price due to uncertainty over correlation, but greater than or equal to the 2GBM price due to uncertainty over volatility. Therefore, the BNM value could theoretically be greater than or less than the 2GBM value based on the expected values of volatilities and correlation, even for ATM spread options.

3.2. An illustration

We investigate the behaviour of BNM spread option prices as the volatility and correlation parameters change. Consider, for example, a 1-year ATM European call spread option with pay-off function $[S_1 - S_2 - K]^+$ based on the following parameters:

$$S_1 = S_2 = 100, \quad \mu_{11} = \mu_{12} = \mu_{21} = \mu_{22} = 0, \quad \sigma_1 = 20\%, \quad \sigma_2 = 25\%, \quad \rho = -0.5. \quad (7)$$

This has a 2GBM price of 15.481[12]. To compute the BNM price, which depends on the excess kurtosis in the marginal densities for each asset as well as their core and tail correlations, we further assume

$$\lambda_1 = \lambda_2 = 0.95, \quad \sigma_{C1} = 18\%, \quad \sigma_{T1} = 58\%, \quad \sigma_{C2} = 24\%, \quad \sigma_{T2} = 44\%. \quad (8)$$

Note that the core and tail volatilities are chosen so that the expected volatility (5) matches the volatility used for the 2GBM price. The two asset returns have no skew but have been chosen so that they have quite different excess kurtosis: 5.81 for asset 1 and 0.64 for asset 2[13].

[12] All 2GBM spread option prices in this paper are based on the analytic approximation given by Kirk (1995).
[13] The excess kurtosis in a mixture $(\lambda, 1-\lambda)$ of two normal densities with volatilities σ_1 and σ_2 under the zero mean assumption is $3[(\lambda\sigma_1^4+(1-\lambda)\sigma_2^4)/((\lambda\sigma_1^2+(1-\lambda)\sigma_2^2)^2-1]$. This is always positive.

Table 1. BNM values of a 1 year ATM European call spread option based on different correlation structures. $S_1=S_2=100$, $\mu_{11}=\mu_{12}=\mu_{21}=\mu_{22}=0$, $\sigma_1=20\%$, $\sigma_2=25\%$, $\rho=-0.5$, $\lambda_1=\lambda_2=0.95$, $\sigma_{C1}=18\%$, $\sigma_{T1}=58\%$, $\sigma_{C2}=24\%$, $\sigma_{T2}=44\%$. 2GBM price 15.312.

ρ_{CC}	ρ_{TC}	ρ_{CT}	ρ_{TT}	Value
−0.51	−0.4	−0.4	−0.9	15.531
−0.52	−0.3	−0.3	−0.9	15.510
−0.53	−0.2	−0.2	−0.9	15.487
−0.541	−0.1	−0.1	−0.9	15.466
−0.552	0	0	−0.9	15.442

Table 2. BNM and 2GBM values of a 1 year call spread option for different strikes. $S_1=S_2=100$, $\mu_{11}=\mu_{12}=\mu_{21}=\mu_{22}=0$, $\sigma_1=20\%$, $\sigma_2=25\%$, $\rho=-0.5$, $\lambda_1=\lambda_2=0.95$, $\sigma_{C1}=18\%$, $\sigma_{T1}=58\%$, $\sigma_{C2}=24\%$, $\sigma_{T2}=44\%$.

Strike	2GBM	BNM	BNM-2GBM
−80	83.034	83.080	0.046
−70	72.344	72.448	0.105
−60	62.376	62.498	0.123
−50	52.878	53.006	0.129
−40	43.866	43.990	0.124
−30	35.462	35.570	0.108
−20	27.830	27.913	0.083
−10	21.130	21.191	0.060
0	15.481	15.531	0.051
10	10.924	10.988	0.064
20	7.419	7.520	0.102
30	4.849	5.005	0.156
40	3.052	3.267	0.215
50	1.853	2.119	0.266
60	1.087	1.388	0.302
70	0.617	0.935	0.318
80	0.340	0.658	0.317

Thus returns on asset 2 are close to normally distributed but asset 1 has a leptokurtic returns density and, whilst less volatile than asset 2, it exhibits more price jumps.

Table 1 reports the BNM values for an ATM option based on the parameters (7) and (8) with different correlation structures in the bivariate normal mixture. Each row of table 1 reports the BNM option value corresponding to feasible values for the core and tail correlations that are consistent with (7) and (8) and when $E(\rho)=-0.5$. In each case the tail–core correlations are set equal and the tail correlation is fixed at $\rho_{TT}=-0.9$. The results show that the BNM values can be greater than or less than the 2GBM value of 15.481 that is based on the expected volatility and the expected correlation. They are not very different from the 2GBM value, but this is because the option is ATM.

To investigate how price differences vary with the moneyness of the spread option we take the correlation structure in the first row of table 1, and table 2 reports the BNM and the 2GBM values of the spread options with different strikes between –80 and +80. As the option moves away from ATM, the differences first increase and then, for far OTM and ITM options, the price differences gradually fall to zero. In fact, the price differences have an asymmetric 'M' shape. The asymmetry arises because the underlying assets have different expected volatilities. In this example, because asset 2 is the more volatile, the largest value differences occur for positive strikes.

Figure 2 illustrates how the difference between the two model values behaves as the strike and maturity changes. Since the convexity of the 2GBM option value as a function of volatility increases with maturity (as shown in figure 1(*b*)), the absolute difference between the BNM and the 2GBM values of European call spread options also increases with maturity. The value differences are not symmetric (by strike) and can become very large for positive strikes, particularly for long maturity options.

Are the value differences large enough to justify the use of the BNM framework? The answer is almost certainly yes, although of course it will depend on the number of spread options in the portfolio and their maturity and strikes, as well as the behaviour of the underlying assets. ATM option value differences are relatively small: for the two assets illustrated above, option value differences (expressed as a percentage of the 2GBM value) are less than 0.5%. But for spread options on highly volatile and strongly correlated assets, value differences can be substantial. Consider, for example, for a call option with strike 30, when the underlying assets are both at 100, having volatilities of 20% and 25% and a negative correlation of –0.5. With 1 year to expiry of the option the BNM value is only 3.22% greater than the 2GBM value, but with only 3 months until expiry the BNM value will be 26.8% greater than the 2GBM value.

4. Spread option sensitivities

We define the first- and second-order sensitivities of a spread option price $P(S_1, S_2, \sigma_1, \sigma_2, \rho)$ as follows:

First-order sensitivities:

$$\begin{aligned}
&\text{Delta}: \delta_i = \partial P/\partial S_i \quad (i=1,2),\\
&\text{Vega}: \nu_i = \partial P/\partial \sigma_i \quad (i=1,2),\\
&\text{Pi}: \pi = \partial P/\partial \rho.
\end{aligned}$$

Second-order sensitivities:

$$\begin{aligned}
&\text{Gamma}: \gamma_i = \partial^2 P/\partial S_i^2 \quad (i=1,2),\\
&\text{Kappa}: \kappa_i = \partial^2 P/\partial \sigma_i^2 \quad (i=1,2),\\
&\text{Phi}: \varphi_{ij} = \partial^2 P/\partial S_i \partial \sigma_j \quad (i,j=1,2),\\
&\text{Zeta}: \zeta_i = \partial^2 P/\partial S_i \partial \rho \quad (i=1,2),\\
&\text{Xi}: \xi_i = \partial^2 P/\partial \sigma_i \partial \rho \quad (i=1,2),\\
&\text{Psi}: \psi = \partial^2 P/\partial \rho^2,\\
&\text{Omega}: \omega = \partial^2 P/\partial \sigma_1 \partial \sigma_2.
\end{aligned}$$

T&F (2005 Style) [18.2.2005–7:22pm] [1–12] [Page No. 7] {TandF}Rquf/RQUF-101600.3d (RQUF) First Proof RQUF-101600

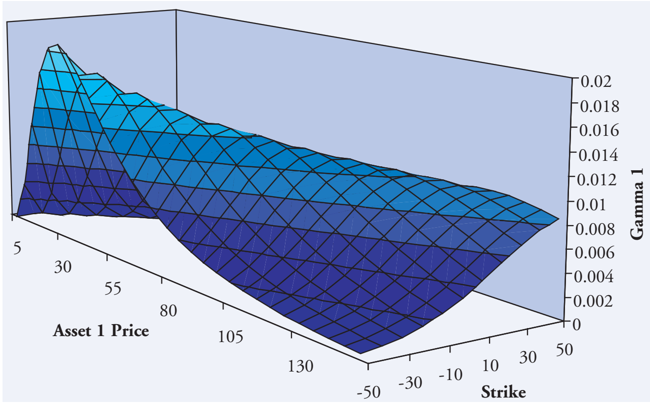

Figure 3. 2GBM gamma for asset 1—all option parameters as in (7).

To illustrate the differences between 2GBM and BNM spread option sensitivities, again we use the asset and option parameters (7) and (8). The 2GBM and BNM spread option sensitivities are based on a 1 unit increase in the underlying parameters (i.e. a 1% absolute increase in volatility and correlation and, since the asset prices are at 100, also a 1% increase in asset price).

Figure 3 shows the 2GBM first-order correlation sensitivity π as a function of the asset correlation, with asset prices and volatilities fixed as in (7). Since spread option prices increase as the correlation decreases, π is always negative and it increases with the option maturity. Note that π is largest (in absolute value) for ATM options, where $\psi = \partial^2 P/\partial\rho^2 = \partial\pi/\partial\rho < 0$, but for OTM calls and puts ψ can be positive. The behaviour of the 2GBM model π and ψ for very highly correlated assets is interesting, as π becomes very large indeed for ATM options, but very small (and ψ can be positive) for OTM options.

The first- and second-order price sensitivities delta and gamma, as a function of their own asset price, and the first- and second-order volatility sensitivities vega and kappa, as a function of their own volatility, have similar characteristics as the 'greeks' of single asset options. For instance, delta 1 is similar to a call option delta and delta 2 is like a put option delta. Figure 4 illustrates the 2GBM gamma 1 as a function of S_1 and K (with $T = 1$ year) and the characteristic asymmetric 'M' shape of kappa 1 as a function of K and T.

Many of the other 2GBM second-order sensitivities also have an asymmetric 'M' shape as a function of moneyness which becomes more pronounced with maturity (although of course all sensitivities become zero in the limits as moneyness increases and decreases towards $\pm\infty$). They often have negatively valued local minima at the ATM strike. This is the case for ψ, the second-order correlation sensitivity, as is evident from figure 3(*b*). Other sensitivities follow an asymmetric 'W' shaped curve with respect to moneyness, with a local maximum around the ATM strike. This is the case for the mixed

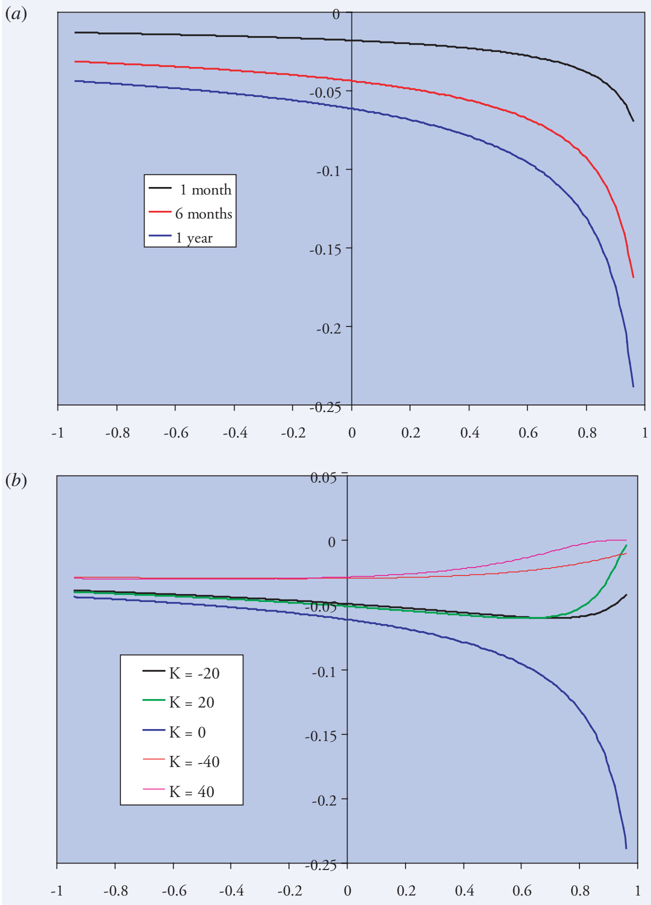

Figure 4. 2GBM first-order correlation sensitivities. (*a*) 2GBM correlation sensitivity of an ATM option for different maturities. (*b*) 2GBM correlation sensitivity of a 1 year option for different strikes.

volatility-correlation sensitivities ξ_1 and ξ_2 and the two price-correlation sensitivities ζ_1 and ζ_2. These are negative because (as seen in figure 1(*a*)) the option value increases as correlation between the two assets decreases. Typically, the mixed price-volatility sensitivities ϕ_{11} and ϕ_{22} will be negative for equity spread options and positive for commodity spread options. Finally, the sign and size of the cross-mixed sensitivities ϕ_{12} and ϕ_{21} depends on the asset correlations. They take the same sign as the own mixed sensitivities when correlation is positive and the opposite sign when correlation is negative.

4.1. Analysis of value differences

Differences between 2GBM and BNM option values arise because of the uncertainty attributed to the volatility and correlation parameters in the bivariate normal mixture density of asset returns. To make this statement more explicit[14], take a second-order Taylor expansion of

[14] To keep notation simple we shall again drop the explicit dependence of the option value on the underlying asset prices, strike, maturity and so forth.

$P(\sigma_1, \sigma_2, \rho)$ about $(E(\sigma_1), E(\sigma_2), E(\rho))$ and take expectations of both sides. All terms in the first-order sensitivities become zero and the result is

$$\begin{aligned} &E(P(\sigma_1,\sigma_2,\rho)) - P(E(\sigma_1),\ E(\sigma_2),\ E(\rho)) \\ &\quad \approx (1/2)[\kappa_1 \mathrm{Var}(\sigma_1) + \kappa_2 \mathrm{Var}(\sigma_2) + \psi \mathrm{Var}(\rho) \\ &\qquad + \omega \mathrm{Cov}(\sigma_1,\sigma_2) + \xi_1 \mathrm{Cov}(\sigma_1,\rho) + \xi_2 \mathrm{Cov}(\sigma_2,\rho)]. \end{aligned} \tag{9}$$

Taking expectation under the probabilities $\{\lambda_1\lambda_2(1-\lambda_1)\lambda_2, \lambda_1(1-\lambda_2), (1-\lambda_1)(1-\lambda_2)\}$ gives

$$P_{\mathrm{BNM}}(\sigma_1,\sigma_2,\rho) = E(P_{\mathrm{2GBM}}(\sigma_1,\sigma_2,\rho)).$$

Putting $P = P_{\mathrm{2GBM}}$ in (9) gives that the price difference $P_{\mathrm{BNM}}(\sigma_1,\sigma_2,\rho)$–$P_{\mathrm{2GBM}}(E(\sigma_1), E(\sigma_2), E(\rho))$ may be approximated as a weighted sum of six 2GBM second-order price sensitivities as follows:

$$\begin{aligned} &P_{\mathrm{BNM}}(\sigma_1,\sigma_2,\rho) - P_{\mathrm{2GBM}}(E(\sigma_1),\ E(\sigma_2),\ E(\rho)) \\ &\quad \approx (1/2)[\kappa_1 \mathrm{Var}(\sigma_1) + \kappa_2 \mathrm{Var}(\sigma_2) + \psi \mathrm{Var}(\rho) \\ &\qquad + \omega \mathrm{Cov}(\sigma_1,\sigma_2) + \xi_1 \mathrm{Cov}(\sigma_1,\rho) + \xi_2 \mathrm{Cov}(\sigma_2,\rho)]. \end{aligned} \tag{10}$$

We have already seen in table 2 and figure 2 that, as the strike of the option varies, the price differences between BNM and 2GBM spread option values typically have an asymmetric 'M' shape that becomes more pronounced with maturity. Given the approximation (10) and our observations above that the second-order volatility and correlation sensitivities typically have an asymmetric 'M' shape with respect to moneyness, the reason for this is now clear.

Some additional features are highlighted by the *weights* in the approximation (10), which represent the degree of uncertainty held over each volatility and the correlation. When the two assets have uncertain volatilities, the covariance between the volatilities affects the BNM option price; likewise, when correlation is uncertain, the covariance of the correlation with both volatilities also affects the BNM price. It is reasonable to suppose that correlation is highly uncertain (i.e. $\mathrm{Var}(\rho)$ is large) and that correlation uncertainty increases with the asset's volatility (i.e. each $\mathrm{Cov}(\sigma_i, \rho)$ is positive) and that there is less uncertainty about volatility than there is about correlation (i.e. both $\mathrm{Var}(\sigma_i)$ and $\mathrm{Cov}(\sigma_1, \sigma_2)$ are relatively low). Hence if the expected correlation is large and negative the BNM price can be less than the 2GBM price for near to ATM options. Negative price differences can also arise if correlation uncertainty decreases as the asset's volatility increases, but this seems unlikely.

It is notable that 'M' shaped price difference surfaces also result when stochastic volatility is introduced to the 2GBM model. Hong (2001) finds that stochastic volatility model prices can also be less than 2GBM prices, without using artificial values for the parameters, and that this difference increases with both the maturity of the option and with the second-order volatility sensitivities.

4.2. Analysis of sensitivity differences

It follows from (4) that each BNM sensitivity is a weighted sum of four different 2GBM sensitivities, each evaluated under different volatility and correlation assumptions. They therefore have similar properties to 2GBM sensitivites, but significant differences may yet arise. For example, table 3 reports the BNM and 2GBM first- and second-order price sensitivity estimates as a function of the price of asset 1. The BNM delta is greater than the 2GBM delta for OTM calls, and less than the 2GBM delta for OTM puts. The BNM gamma is less than the 2GBM gamma for most options, except for OTM options, with differences being more pronounced for short-term, near to ATM options.

Table 4 reports the volatility and correlation sensitivities resulting from the two models[15]. For the spread option parameters of our example, both first- and second-order volatility hedge ratios roughly halve in size when the BNM model is used. Results are only given here for the 1 year maturity option, but even for short-term options when vega is much smaller, the BNM vega will be about half of the 2GBM vega, whatever the volatility level. Likewise, the 2GBM model appears to be substantially overestimating π, the first-order correlation sensitivity. The BNM values for π are much smaller in absolute value at all levels of correlation, and the difference becomes more pronounced as the asset's correlation increases. Hence both volatility and correlation risks are likely to be *over* hedged using the 2GBM model, compared with the BNM model which explicitly accounts for the uncertainty in these parameters.

5. Modelling the correlation frown

The correlation frown arises when 2GBM values of OTM calls and puts are lower than market prices of these spread options. We have seen that BNM values of such options are greater than the 2GBM values. Hence the uncertainty in both volatility and correlation that is implicit in the BNM framework can explain at least part of the correlation frown. We can use the BNM model to analyse how the behaviour of the underlying assets affects the shape of the correlation frown. In this section we replicate a correlation frown using the BNM model, treating the BNM values as if they were market prices, setting implied volatilities equal to the model values and backing-out the implied correlation using the 2GBM model. We then simulate the behaviour of the correlation frown as we change the volatility and correlation structure of the underlying assets.

Continuing with the option and asset parameters in (7) and (8), but changing the 'core' and 'tail' correlations, figure 6(*a*) shows that the principal influence of increasing the core correlation is to increase the height of the frown. The 'tail–core' and 'tail' correlations also alter the height, but to a much lesser degree—they are more for fine

T&F (2005 Style) [18.2.2005–7:22pm] [1–12] [Page No. 9] {TandF}Rquf/RQUF-101600.3d (RQUF) First Proof RQUF-101600

Table 3. First- and second-order sensitivities with respect to the price of asset 1. All other option parameters as in (7) and (8) with $K=0$.

S	Delta 1			Gamma 1		
	BNM (%)	2GBM (%)	Difference (%)	BNM (%)	2GBM (%)	Difference (%)
$T=1$ year						
50	6.60	5.42	1.182	0.56	0.55	0.011
60	14.01	12.84	1.173	0.88	0.89	−0.010
70	24.08	23.07	1.009	1.09	1.11	−0.020
80	35.53	34.75	0.783	1.17	1.19	−0.024
90	47.01	46.46	0.551	1.12	1.14	−0.023
100	57.57	57.24	0.338	1.00	1.02	−0.020
110	66.71	66.55	0.154	0.84	0.86	−0.017
120	74.26	74.26	0.001	0.69	0.70	−0.014
130	80.31	80.43	−0.121	0.54	0.55	−0.011
145	85.05	85.26	−0.215	0.42	0.43	−0.008
150	88.69	88.97	−0.284	0.32	0.33	−0.006
$T=1$ month						
50	0.00	0.00	0.00	0.00	0.00	0.00
60	0.05	0.00	0.047	0.01	0.00	0.009
70	0.37	0.08	0.293	0.07	0.03	0.040
80	3.16	2.40	0.765	0.61	0.57	0.039
90	18.42	17.69	0.734	2.43	2.47	−0.039
100	50.61	50.47	0.131	3.50	3.57	−0.067
110	80.11	80.54	−0.427	2.28	2.32	−0.042
120	94.30	94.91	−0.609	0.83	0.83	0.001
130	98.59	99.06	−0.464	0.21	0.19	0.021
145	99.61	99.87	−0.258	0.05	0.03	0.018
150	99.86	99.99	−0.121	0.01	0.00	0.010

Table 4. Comparison of first- and second-order sensitivities with respect to the volatility of asset 1 and first-order correlation sensitivity. All other option parameters as in (7) and (8) with $K=0$.

$T=1$ year Volatility (%)	Vega 1		Kappa 1			Pi		
	BNM (%)	2GBM (%)	BNM (%)	2GBM (%)	Correlation (%)	BNM (%)	2GBM (%)	Difference (%)
5	11.44	24.40	0.49	0.91	−90	−1.46	−4.43	2.97
10	13.49	28.10	0.35	0.64	−75	−1.53	−4.62	3.10
15	14.92	30.66	0.24	0.44	−50	−1.65	−5.00	3.35
20	15.91	32.43	0.17	0.30	−25	−1.80	−5.49	3.69
25	16.59	33.65	0.11	0.21	−10	−1.91	−5.85	3.94
30	17.06	34.48	0.08	0.14	0	−2.00	−6.14	4.14
35	17.38	35.04	0.05	0.09	10	−2.09	−6.47	4.37
40	17.58	35.40	0.03	0.06	25	−2.27	−7.07	4.81
45	17.71	35.60	0.02	0.03	50	−2.68	−8.61	5.93
50	17.77	35.69	0.01	0.01	75	−3.43	−11.88	8.45
55	17.78	35.68	0.01	−0.01	90	−4.30	−17.47	13.17

tuning—and have only very small effects on the curvature of the frown, being visible only numerically. In figure 6(*b*) the volatilities of each component in the marginal densities and the core and tail correlations remain fixed at their values for the previous example. Now, the excess kurtosis in each marginal density and the two expected volatilities change as we alter λ_1 and λ_2:

- Scenario 1: $\lambda_1=\lambda_2=0.95$, so that $E(\sigma_1)=20\%$, $E(\sigma_2)=25\%$, and the two assets have excess kurtosis 5.81 and 0.64, respectively (this was the scenario for all previous examples);
- Scenario 2: $\lambda_1=\lambda_2=0.75$, so that $E(\sigma_1)=28\%$, $E(\sigma_2)=29\%$, and the two assets have excess kurtosis 4.42 and 1.24, respectively.

With each scenario the expected volatilities of each of the assets increases and so does the height at the centre of the frown. The primary effect of increasing the excess kurtosis is to alter the steepness rather than the height of the frown. For example, when asset 2 has higher excess kurtosis (as in scenario 2) we perceive a much steeper frown. In general, increasing the leptokurtosis in the marginal densities whilst holding expected volatility and

T&F (2005 Style) [18.2.2005–7:22pm] [1–12] [Page No. 10] {TandF}Rquf/RQUF-101600.3d (RQUF) First Proof RQUF-101600

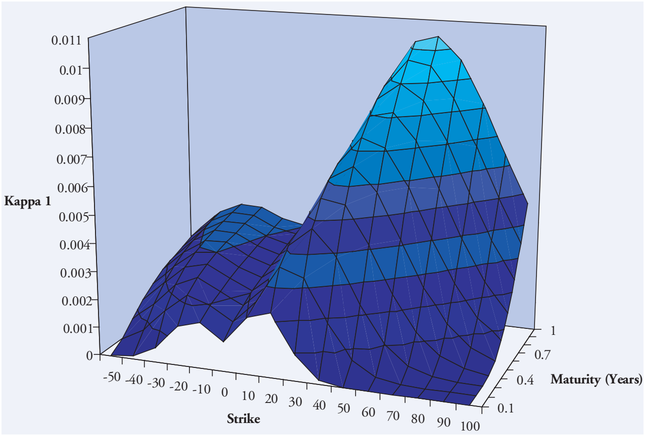

Figure 5. 2GBM second-order volatility sensitivity as a function of strike and maturity.

component correlations constant would increase the convexity but not the height at the centre of the frown.

6. Summary and conclusions

This paper examines the behaviour of European spread option prices and hedge ratios when we account for trader's uncertainty about both volatility and correlation in an intuitive and tractable manner. The simplest possible sample space has been assumed: for each asset there are only two possible states for volatility and four possible states for correlation. Nevertheless, the introduction of volatility and correlation uncertainty makes the structure of the spread option price surface considerably richer than one might naïvely expect. In contrast to the single asset option case, where there are simple rules such as 'the lognormal mixture price is always greater than or equal to the GBM price', here we find complex differences between BNM and 2GBM spread option values. Value differences are influenced by six second-order sensitivities and by the degree of uncertainty that traders hold over both volatility and correlation, an uncertainty that is captured by the bivariate normal mixture joint density of asset returns.

We have shown by example that for a single spread option on moderately volatile assets that have a moderate negative correlation, and when just one of these assets has a fairly leptokurtic returns density, there can be large differences in the prices and hedge ratios that are calculated under the two different models. Clearly, when many spread options are written on volatile and leptokurtic assets, as is the case in the commodity markets in particular, the 2GBM model can lead to substantial under pricing and over hedging of the portfolio!

The BNM model prices are generally greater than the 2GBM prices although, since they are approximated as a weighted sum of six 2GBM second-order option

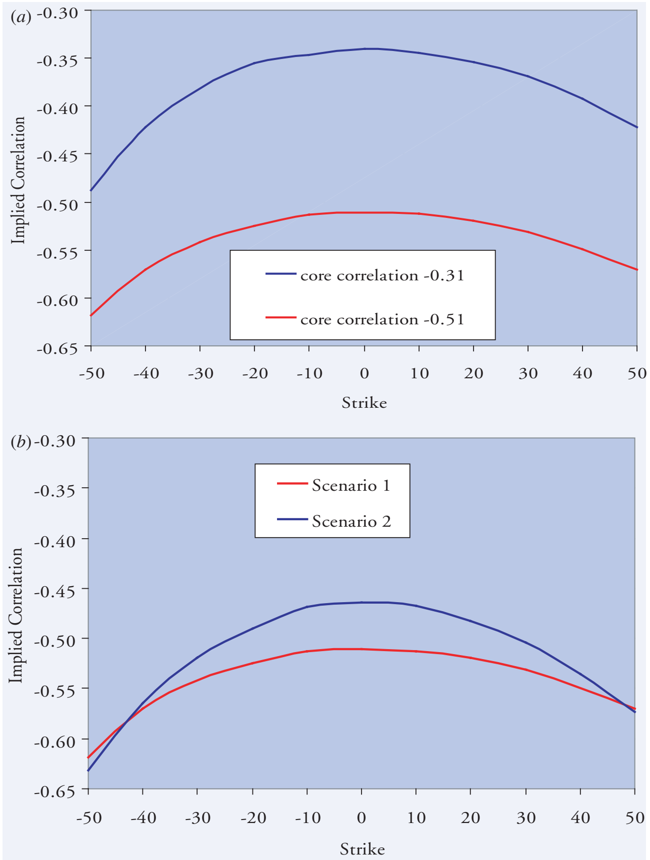

Figure 6. (*a*) Correlation 'frowns' implied by the BNM model. The 'core' correlation determines the height of the frown. (*b*) Correlation 'frowns' implied by the BNM model. The asset's kurtosis determines the steepness of the frown.

sensitivities that can take some fairly strange shapes as a function of strike and maturity, they can, theoretically, be lower. Value differences are small for short-term ATM options, but the differences increase markedly for OTM options and for longer maturity options. Thus BNM option prices can explain at least part of the implied correlation frown.

To our knowledge the only other 'frown consistent' spread option valuation model is the jump-diffusion model of Carmona and Durrleman (2003b). Compared with this approach, the simple mixture paradigm affords ease of understanding and interpretability, a very important attribute in the complicated world of multi-asset derivative pricing. By considering the simplest type of uncertainty in asset volatilities and in their correlation, the arbitrage free spread option values obtained under the BNM assumption can be consistent with implied volatility smiles in both assets and, at the same time, the implied correlation frown.

If there are traded options on each underlying asset, the probability of each state and volatility in each state may be calibrated to each asset's implied volatility smile surface, for example as in Brigo and Mercurio (2001). This leaves only the component correlations to be

T&F (2005 Style) [18.2.2005–7:22pm] [1–12] [Page No. 11] {TandF}Rquf/RQUF-101600.3d (RQUF) First Proof RQUF-101600

calibrated to the market correlation frown. We find that the excess kurtosis of each asset return is the key determinant of the convexity of the frown: indeed, without leptokurtosis in at least one asset, the frown would be flat. Thus, when the marginal density parameters are first calibrated from single asset options, by the time we come to consider the correlation frown its foundations will have already been laid. We have shown that the structure of the asset correlations also affects the convexity of the frown, but to a lesser extent. The principal influence on the height of the frown is the core correlation between the assets. Though tail correlations are often the strongest and most significant, these only play the role of 'fine tuning' to an observed correlation frown in market prices.

We hope that this paper will be the catalyst for further research. For instance, here we have only illustrated the simplest case that each asset return density is a mixture of two normal densities with identical means. This produces a smooth, relatively simple implied correlation frown. If different means are added to the mixtures, more irregular or intricate frowns including highly skewed shapes could be captured in the BNM framework. One of the challenges will be the calibration to a market implied correlation surface so that accurate and robust prices are obtained. Then, since the model has analytically tractable (normal mixture) transition densities, the pricing of American and other path-dependent spread options will be straightforward. Clearly, this analytic approach, which has a rich and intuitive structure, shows considerable promise with the potential for profit through more accurate pricing and hedging of spread options.

Acknowledgement

Many thanks to Dr. Hyungsok Ahn of CommerzBank, London, for valuable discussions about a preliminary version of this paper and to an anonymous referee for very useful comments.

References

Alexander C 2004 Normal mixture diffusion with uncertain volatility: modelling short and long term smile effects *J. Banking Finance* forthcoming

Alexander C and Lazar E 2004 The continuous limit of normal mixture GARCH *ISMA Centre Discussion Papers in Finance 2004-10* Available from www.isma-centre.rdg.ac.uk/dp

Bingham NH and Kiesel R 2002 Semi-parametric modelling in finance: theoretical foundations *Quant. Finance* **2** 241–50

Black F 1976 The pricing of commodity contracts *J. Financial Economics* **3** 167–79

Black F and Scholes M 1973 The pricing of options and corporate liabilities *J. Political Economy* **81** 637–54

Boyle P 1988 A lattice framework for option pricing with two state variables *J. Financial Quant. Anal.* **23** 1–12

Brigo D 2002 The general mixture-diffusion SDE and its relationship with an uncertain-volatility option model with volatility–asset decorrelation, available from www.damianobrigo.it

Brigo D and Mercurio F 2001 Displaced and mixture diffusions for analytically-tractable smile models *Mathematical Finance Bachelier Congress 2000*, eds H Geman, D B Madan, S R Pliska and A C F Vorst (Berlin: Springer)

Brigo D, Mercurio F and Rapisarda F 2003 Connecting univariate smiles and basket dynamics: a new mulitvariate dynamics for basket options *Banca IMI Internal Report*

Carmona R and Durrleman V 2003a Pricing and hedging spread options *SIAM Rev.* **45** 627–87

Carmona R and Durrleman V 2003b Pricing and hedging spread options in a lognormal model *Technical Report* Department of Operations Research and Financial Engineering, Princeton University

Dempster M and Hong G 2000 Spread option valuation and the fast Fourier transform *Research Papers in Management Studies Working Paper* 26/2000

Durrleman V 2001 Implied correlation and spread options *Technical Report* Department of Operations Research and Financial Engineering, Princeton University

Heston S 1993 A closed form solution for options with stochastic volatility with applications to bond and currency options *Rev. Financial Studies* **6** 327–43

Hong S-S G 2001 Pricing and hedging of spread options with stochastic component correlation *PhD Thesis*, Judge Institute of Management Studies, Cambridge, available from www.jims.cam.ac.uk

Hull J and White A 1987 The pricing of options on assets with stochastic volatilities *J. Finance* **42** 281–300

James P 2002 *Option Theory* (Chichester: Wiley)

Kirk E 1995 Correlation in the energy markets *Managing Energy Price Risk* 1st edn (Risk Publications) pp 71–8

Margrabe W 1978 The value of an option to exchange one asset for another *J. Finance* **33** 177–86

Melick WR and Thomas C P 1997 Recovering an assets implied PDF from option prices: an alication to guide oil during the Gulf crisis *J. Financial Quant. Anal.* **32** 91–115

Mercurio F 2002 A multi-stage uncertain volatility model *Banca IMI Working Paper* available from www.fabiomercurio.it/uncertainvol.pdf

Nelson DB 1990 ARCH models as diffusion approximations *J. Econometrics* **45** 7–38

Pliska S and Duan 2003 Option valuation with cointegrated asset prices *J. Economic Dynamics Control* forthcoming

Ravindran K 1993 Low-fat spreads *Risk Mag.* **6** 66–7

Ritchey RJ 1990 Call option valuation for discrete normal mixtures *J. Financial Res.* **13** 285–96

Rosenberg J 2001 Nonparametric pricing of multivariate contingent claims *Unpublished Manuscript* New York University, Stern School Of Business

Shimko D 1994 Options on futures spreads: hedging, speculation and valuation *J. Futures Markets* **14** 182–213