

ANALYTIC APPROXIMATIONS FOR MULTI-ASSET OPTION PRICING

CAROL ALEXANDER AND AANAND VENKATRAMANAN

Reading University

We derive general analytic approximations for pricing European basket and rainbow options on N assets. The key idea is to express the option's price as a sum of prices of various compound exchange options, each with different pairs of subordinate multi- or single-asset options. The underlying asset prices are assumed to follow lognormal processes, although our results can be extended to certain other price processes for the underlying. For some multi-asset options a strong condition holds, whereby each compound exchange option is equivalent to a standard single-asset option under a modified measure, and in such cases an almost exact analytic price exists. More generally, approximate analytic prices for multi-asset options are derived using a weak lognormality condition, where the approximation stems from making constant volatility assumptions on the price processes that drive the prices of the subordinate basket options. The analytic formulae for multi-asset option prices, and their Greeks, are defined in a recursive framework. For instance, the option delta is defined in terms of the delta relative to subordinate multi-asset options, and the deltas of these subordinate options with respect to the underlying assets. Simulations test the accuracy of our approximations, given some assumed values for the asset volatilities and correlations. Finally, a calibration algorithm is proposed and illustrated.

KEY WORDS: basket options, rainbow options, best-of and worst-of options, compound exchange options, analytic approximation.

1. INTRODUCTION

This paper presents a recursive procedure for pricing European basket and rainbow options on N assets. The pay-off to a basket option at its maturity T is $\omega[\Theta_N \mathbf{S}'_T - K]^+$, and that of a rainbow option is $\omega[\max\{S_{1T}, \dots, S_{NT}\} - K]^+$, where $\mathbf{S}_T = (S_{1T}, S_{2T}, \dots, S_{NT})$ are the N asset prices, K is the option strike, $\omega = 1$ (call) or -1 (put) and the weights $\Theta_N = (\theta_1, \theta_2, \dots, \theta_N)$ are real constants. Zero-strike rainbow options are commonly termed best-of- N -assets or worst-of- N -assets options. Since $\min\{S_{1T}, \dots, S_{NT}\} = -\max\{-S_{1T}, \dots, -S_{NT}\}$, a pricing model for best-of options also serves for worst-of options. The most commonly traded two-asset options are exchange options (rainbow options with zero strike) and spread options (basket options with weights 1 and -1). Margrabe (1978) derived an exact solution for the price of an exchange option, under the assumption that the two asset prices follow correlated lognormal processes. However, a straightforward generalization of this formula to any two-asset option with non-zero strike is impossible. A linear combination of lognormal processes is no longer lognormal, and it is only when the strike of a two-asset option is zero that one may circumvent this problem by reducing the dimension of the correlated lognormal processes to one. Hence,

Manuscript received July 2009; final revision received September 2010.

Address correspondence to Aanand Venkatramanan, ICMA Centre, Henley Business School at Reading, Reading University, RG6 6BA, UK; e-mail: aanand.venkat@gmail.com.

DOI: 10.1111/j.1467-9965.2011.00481.x

© 2011 Wiley Periodicals, Inc.

academic research has focussed on deriving good analytic approximations for pricing two-asset options with non-zero strike, as well as more general multi-asset options.

Analytic approximations for pricing rainbow options include that of Johnson (1987), who extends the two-asset rainbow option pricing formula of Stulz (1982) to the general case of N assets. Ouwehand and West (2006) verify the results of Johnson (1987) and explain how to prove them using a multivariate normal density approximation derived by West (2005). Then, they explain how to price N asset rainbow options using this approach, providing an explicit approximation for the case $N = 4$. For options on two assets with different pay-off profiles, Topper (2001) uses a finite element scheme to solve nonlinear parabolic price PDEs.

Approximations for basket options began with Levy (1992), who approximated a basket price distribution with that of a single lognormal variable, matching the first and second moments. Then Gentle (1993) derived the basket option price by approximating the arithmetic average by a geometric average. Milevsky and Posner (1998a) used the reciprocal gamma distribution and Milevsky and Posner (1998b) used the Johnson (1949) family of distributions to approximate the distribution of the basket price. Extending the Asian option pricing approach of Rogers and Shi (1995), Beißer (1999) expressed a basket option price as a weighted sum of single-asset Black-Scholes prices, with adjusted forward price and adjusted strike for every constituent asset. Finally, Ju (2002) used a Taylor expansion to approximate the ratio of the characteristic function of the average of correlated lognormal variables, which is approximately lognormal for short maturities. Krekel et al. (2004) compare the performance of these models, concluding that the approximations of Ju (2002) and Beißer (1999) are most accurate, although they tend to slightly over- and underprice, respectively. However, many of these methods have limited validity or scope. They may require a basket value that is always positive, or they may not identify the effect of each individual volatility or pair-wise correlation on the multi-asset option price or its hedge ratios.

This paper develops an analytic approximation that has none of these limitations. It is based on the novel idea of writing a basket or rainbow pay-off as a sum of pay-offs to compound exchange options (CEOs), where the “assets” in these exchange options are pairs of subordinate basket options, and then applying a recursive procedure to obtain a decomposition of the multi-asset option’s pay-off into a sum of pay-offs to standard, single-asset European calls and puts (henceforth, vanilla options) and options to exchange vanilla options on the assets. Each asset price is assumed to follow a standard geometric Brownian motion (GBM) process, but this assumption may be relaxed to allow more general drift and local volatility components. The price of a vanilla option follows an Itô’s process, which we approximate with a lognormal process and hence the prices of the exchange options on vanilla options are obtained using the formula of Margrabe (1978). Then an analytic approximation to the multi-asset option price is given by a recursive application.

In the following, Section 2 presents these recursive pay-off decompositions and derives the prices of European basket and rainbow options as linear combinations of prices of vanilla options and options to exchange such options. Section 3 analyzes CEOs on vanilla options, deriving conditions under which (1) the price processes for the vanilla options are approximately lognormal, and (2) their relative prices are almost exactly lognormal. Section 4 derives approximate basket option prices based on both these conditions. Section 5 presents simulations that illustrate the accuracy of our approximations and Section 6 describes a general calibration algorithm. Section 7 summarizes and concludes.

2. PRICING FRAMEWORK

The pay-off to a basket option on N assets with maturity T and (any real) strike K is

$$(2.1) \quad V_{NT} = [\omega(B_T - K)]^+ = [\omega(\Theta_N \mathbf{S}'_T - K)]^+ = [\omega \Theta_N (\mathbf{S}'_T - \mathbf{K}')]^+,$$

where $B_T = \Theta_N \mathbf{S}'_T$ and $\mathbf{K} = (K_1, K_2, \dots, K_N)$ such that $\Theta_N \mathbf{K}' = K$, other notation being previously defined. Now set $\Theta_N = (\Theta_m, \Theta_n)$, $\mathbf{S}_T = (\mathbf{S}_{mT}, \mathbf{S}_{nT})$, and $\mathbf{K} = (\mathbf{K}_m, \mathbf{K}_n)$ for some positive integers m and n such that $m + n = N$. Define the sub-basket call and put pay-offs $C_{mT} = [\Theta_m (\mathbf{S}'_{mT} - \mathbf{K}'_m)]^+$, and $P_{mT} = [-\Theta_m (\mathbf{S}'_{mT} - \mathbf{K}'_m)]^+$, and similarly for n . Then (2.1) becomes

$$(2.2) \quad V_{NT} = [C_{mT} - P_{nT}]^+ + [C_{nT} - P_{mT}]^+.$$

Alternatively, setting $\Theta_N = (\Theta_m, -\Theta_n)$,

$$(2.3) \quad V_{NT} = [C_{mT} - C_{nT}]^+ + [P_{nT} - P_{mT}]^+.$$

The above representations follow on noting that, if $f^+ = \max\{f, 0\}$ and $f^- = \max\{-f, 0\}$ and f and g are any real-valued functions, then $[f + g]^+ = [f^+ - g^-]^+ + [g^+ - f^-]^+$ and $[f - g]^+ = [f^+ - g^+]^+ + [g^+ - f^+]^+$. For $0 \leq t < T$, V_{Nt} may now be computed as the discounted sum of the risk-neutral expectations of the two replicating CEO pay-offs which appear on the right-hand side: in case (2.2), they are pay-offs to exchange options on a basket call and a basket put; and in case (2.3) they are pay-offs to exchange options on two basket calls and two basket puts with a different number of assets in each basket.

Decompositions of the form (2.2) or (2.3) are then applied to each of C_{mT} , P_{mT} , C_{nT} , and P_{nT} in turn, choosing suitable partitions for m and n which determine the number of assets in the subordinate basket calls and puts. By applying the pay-off decomposition repeatedly, each time decreasing the number of assets in the subordinate options, one eventually expresses the pay-off to an N -asset basket option as a sum of pay-offs to CEOs in which the subordinate options are standard single-asset calls and puts, and ordinary exchange options. The price of the original N -asset basket option is then computed as the sum of the prices of CEOs and standard exchange options. We explain how to price CEOs in terms of standard single-asset option prices in the next section.

Figure 2.1 illustrates this recursive decomposition for pricing a four-asset basket option (we have only shown one leg of the tree, but the other leg can be priced similarly). The four-asset basket option has price equal to the sum of the two CEO prices on the penultimate level of the tree. To price these CEOs we have first to compute the prices of the 4 two-asset basket options on the level below, but to price these we must first price eight CEOs—and to price the CEOs we need to price the standard calls on the four assets (we only need call prices, because we can deduce the corresponding put prices using put-call parity).

For a general N -asset basket option our approach requires the evaluation of $2(N - 1)$ CEO prices and $2N$ standard option prices. When N is odd, terminal nodes will contain vanilla options. For instance, for $N = 3$ and the pay-off $[S_{1T} - S_{2T} + S_{3T} - K]^+$, the decomposition may be written $[[S_{1T} - S_{2T}]^+ - [K - S_{3T}]^+]^+ + [[S_{3T} - K]^+ - [S_{2T} - S_{1T}]^+]^+$, so we choose a pair of exchange options on assets one and two and a pair of vanilla options on asset three. The weighted sum of the strikes of the exchange or vanilla options that appear in the terminal nodes of the tree is equal to K , the strike of the basket option.

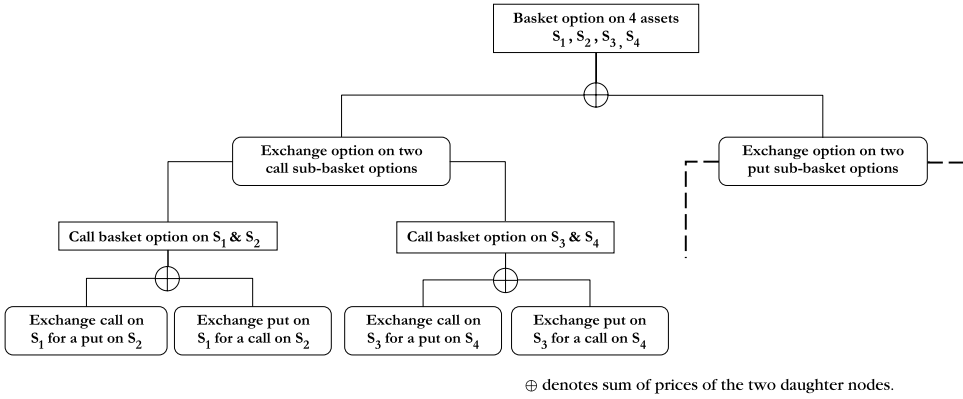


FIGURE 2.1. Pricing tree for four-asset basket option.

Now consider a general rainbow option with pay-off $\max\{S_{1T} - K, \dots, S_{NT} - K\}$. Its price may be expressed in terms of an N -asset basket option price and exchange option prices. To see this, let (n_1, n_2, \dots, n_N) be a permutation of $(1, 2, \dots, N)$ and choose k to be some integer between 1 and N . By splitting the basket of N assets into two sub-baskets, where k determines the size of the sub-baskets and the permutation (n_1, n_2, \dots, n_N) determines the assets in these sub-baskets, the rainbow option's pay-off may be written as the sum of two pay-offs, one to a best-of option on a sub-basket and the other to a compound option, as

$$\begin{aligned} & \max\{S_{n_{k+1}T}, S_{n_{k+2}T}, \dots, S_{n_NT}\} - K \\ & + [\max\{S_{n_1T}, S_{n_2T}, \dots, S_{n_kT}\} - \max\{S_{n_{k+1}T}, S_{n_{k+2}T}, \dots, S_{n_NT}\}]^+. \end{aligned}$$

The best-of option pay-off terms here may themselves be represented as the sum of two such pay-offs, until all sub-baskets are on one or two assets. Once the sub-basket size is eventually reduced to two, we use the identity $\max\{S_{iT}, S_{jT}\} = S_{jT} + [S_{iT} - S_{jT}]^+$.

For every permutation (n_1, n_2, \dots, n_N) and index k we obtain a different pay-off decomposition for the rainbow option. Obviously, the value of the pay-off will be the same in each case, so the model should be calibrated in such a way that the option price is invariant to the choice of (n_1, n_2, \dots, n_N) and k . For illustration, consider a rainbow option on four assets. Here, it is convenient to use the notation $X_{n_{it}}$ for the price of an option to receive asset n_i in exchange for selling asset n_{i+1} . Choosing $(n_1, n_2, n_3, n_4) = (1, 2, 3, 4)$ and $k = 2$, the rainbow option's pay-off P_{4T} may be written $P_{4T} = S_{4T} + X_{3T} + [B_T]^+ - K$. Thus, the price of the rainbow option is $P_{4t} = S_{4t} + X_{3t} + V_t - Ke^{-r(T-t)}$, where $V_t = e^{-r(T-t)}\mathbb{E}_Q\{[B_T]^+\}$ is the price of a zero-strike basket option with four assets whose prices are $\{X_{1t}, S_{2t}, X_{3t}, S_{4t}\}$ and with weights $\{1, 1, -1, -1\}$. Recall that, under the correlated lognormal assumption, an analytic solution exists for X_{it} . Hence, P_{4t} may be evaluated because we have already derived the price V_t of the basket option; it may be expressed in terms of CEO prices.

We have chosen $k = 2$ above because this choice leads to the simplest form of pay-off decomposition for a four-asset rainbow option. In fact, for any permutation (n_1, n_2, n_3, n_4) , the pay-off may be written $P_{4T} = S_{n_4T} + X_{n_3T} + [S_{n_2T} - X_{n_1T} - S_{n_4T} + X_{n_3T}]^+ - K$. Hence, a general expression for the price of a four-asset rainbow option is $P_{4t} = S_{n_4t} + X_{n_3t} + V_{4t} - Ke^{-r(T-t)}$, where $V_{4t} = e^{-r(T-t)}\mathbb{E}_Q\{[B_T]^+\}$ denotes the price of a zero-strike

basket option with four assets whose prices are $\{X_{n_1t}, S_{n_2t}, X_{n_3t}, S_{n_4t}\}$ and with weights $\{1, 1, -1, -1\}$. This argument can be extended to rainbow options on more than four assets. For example, the pay-off to a rainbow option on three assets, having prices S_5 , S_6 , and S_7 , may be written $P_{3T} = S_{7T} + X_{6T} + [S_{5T} - S_{7T} + X_{6T}]^+$, and the price of a seven-asset rainbow option is $P_{7t} = P_{3t} + \mathbb{E}_{\mathbb{Q}}\{[P_{4T} - P_{3T}]^+\}$.

3. PRICING CEO'S

The pay-off decompositions illustrated above have employed options on a basket which may contain the assets themselves, options written on these assets, and options to exchange these assets. A suitable pay-off decomposition will express the price of these basket options as a sum of CEO prices. The two options exchanged in the CEO (which may be on different assets or may themselves be compound options) always have the same maturity as the CEO. We now derive an analytic approximation for the CEO price, first assuming the underlying assets follow correlated lognormal processes, and then under more general price processes. Our key idea is to express the price of such a CEO as the price of a single-asset option that can be easily derived.

Let $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$ be a filtered probability space, where Ω is the set of all possible events such that $S_{1t}, S_{2t} \in (0, \infty)$, $(\mathcal{F}_t)_{t \geq 0}$ is the filtration produced by the sigma algebra of the price pair $(S_{1t}, S_{2t})_{t \geq 0}$ and \mathbb{Q} is a bivariate risk-neutral probability measure. Assume the risk-neutral price dynamics are governed by $dS_{it} = rS_{it}dt + \sigma_i S_{it}dW_{it}$ with $\langle dW_{1t}, dW_{2t} \rangle = \rho dt$ for $i = 1, 2$, where W_1 and W_2 are Wiener processes under the risk-neutral measure \mathbb{Q} , σ_i is the constant volatility of asset i , and ρ is the constant correlation between dW_1 and dW_2 . Consider an option to exchange a vanilla option on asset one with a vanilla option on asset two, where all options have the same maturity T . Let U_{it} and V_{it} denote the prices of the vanilla call and put on asset i with common strike K_i , for $i = 1, 2$. If the CEO is on two calls the pay-off is $[\omega(U_{1T} - U_{2T})]^+$, if the CEO is on two puts the pay-off is $[\omega(V_{1T} - V_{2T})]^+$, and the pay-off is either $[\omega(U_{1T} - V_{2T})]^+$ or $[\omega(V_{1T} - U_{2T})]^+$ if the CEO is on a call and a put, where $\omega = 1$ for a call CEO and -1 for a put CEO.

The price f_i of a CEO can be obtained as a risk-neutral expectation of the terminal pay-off, for example, for a CEO on two calls

$$f_i = e^{-r(T-t)} \mathbb{E}_{\mathbb{Q}}\{[\omega(U_{1T} - U_{2T})]^+ | \mathcal{F}_t\}.$$

But the application of risk-neutral valuation requires a description of the price processes U_{it} and V_{it} , for $i = 1, 2$. To this end we apply Itô's Lemma to f_i and use the Black-Scholes differential equation to obtain

$$(3.1) \quad dU_{it} = \left(\frac{\partial U_{it}}{\partial t} + r S_{it} \frac{\partial U_{it}}{\partial S_{it}} + \frac{1}{2} \sigma_i^2 S_{it}^2 \frac{\partial^2 U_{it}}{\partial S_{it}^2} \right) dt + \sigma_i S_{it} \frac{\partial U_{it}}{\partial S_{it}} dW_{it} \\ = r U_{it} dt + \xi_{it} U_{it} dW_{it},$$

where $\xi_{it} = \sigma_i \frac{S_{it}}{U_{it}} \frac{\partial U_{it}}{\partial S_{it}}$. Similarly, setting $\eta_{it} = \sigma_i \frac{S_{it}}{V_{it}} \left| \frac{\partial V_{it}}{\partial S_{it}} \right|$ yields

$$(3.2) \quad dV_{it} = r V_{it} dt + \eta_{it} V_{it} dW_{it}.$$

We restrict our subsequent analysis to pricing a CEO on two calls (the following derivations are very similar when one or both of the standard options are puts). We first

solve for the prices U_{it} of the underlying options and their volatilities ξ_{it} , for $i = 1, 2$. Then we derive “weak” and “strong” lognormality conditions, under which the CEO price is approximately equal to a single-asset option price under an equivalent measure.

LEMMA 3.1. *When the asset price S_{it} follows a GBM with Wiener process W_{it} , a standard call option on asset i has price process described by (3.1) with volatility ξ_{it} following the process described by*

$$(3.3) \quad d\xi_{it} = \xi_{it} \left(\sigma_i - \xi_{it} + \sigma_i S_{it} \frac{\Gamma_{it}}{\Delta_{it}} \right) [-\xi_{it} dt + dW_{it}],$$

where $\Delta_{it} = \frac{\partial U_{it}}{\partial S_{it}}$ and $\Gamma_{it} = \frac{\partial \Delta_{it}}{\partial S_{it}}$ are the delta and gamma of the call option.

Proof. Let $\theta_{it} = \frac{\partial U_{it}}{\partial t}$ and $X_{it} = \frac{\Delta_{it}}{U_{it}}$. Then, dropping the subscripts, $d\xi = \sigma(XdS + SdX + dSdX)$. By Itô's Lemma

$$\begin{aligned} d\Delta &= \frac{\partial \theta}{\partial S} dt + \frac{\partial \Delta}{\partial S} (rSdt + \sigma SdW) + \frac{1}{2} \frac{\partial^2 \Delta}{\partial S^2} \sigma^2 S^2 dt \\ &= \frac{\partial}{\partial S} (rU) dt - (r\Delta + \Gamma\sigma^2 S) dt + \sigma S\Gamma dW \\ &= -\Gamma\sigma^2 Sdt + \sigma S\Gamma dW, \end{aligned}$$

and

$$\begin{aligned} dX &= U^{-1} d\Delta - U^{-2} \Delta dU + U^{-3} \Delta dU^2 - U^{-2} d\Delta dU \\ &= U^{-1} (\Delta\xi^2 - \sigma S\Gamma\xi - r\Delta - \Gamma\sigma^2 S) dt + U^{-1} (\sigma S\Gamma - \xi\Delta) dW. \end{aligned}$$

Hence $d\xi = (\sigma\xi - \xi^2 + \frac{\sigma^2 S}{U} \Gamma)[- \xi dt + dW]$, which can be rewritten as (3.3). \square

The only approximation we need to make is that $\sigma_i S_{it} \Gamma_{it} / \Delta_{it}$ is constant in equation (3.3), that is, set $c_i = \sigma_i S_{it} \Gamma_{it} / \Delta_{it}$ and write $\tilde{\sigma}_i = \sigma_i + c_i$. Then (3.3) becomes

$$(3.4) \quad d\xi_{it} = \xi_{it} (\xi_{it} - \tilde{\sigma}_i) (\xi_{it} dt - dW_{it}).$$

LEMMA 3.2. *Let $k_i = \tilde{\sigma}_i / \xi_0 - 1$ and $W_{it}^* = -\frac{1}{2} \tilde{\sigma}_i t + W_{it}$, for $i = 1, 2$. Then (3.4) has solution*

$$(3.5) \quad \xi_{it} = \tilde{\sigma}_i (1 + k_i e^{-\tilde{\sigma}_i W_{it}^*})^{-1}.$$

For ξ_{it} to remain finite, we must have $W_{it}^* > \frac{1}{2} \tilde{\sigma}_i t + \tilde{\sigma}_i^{-1} \ln |k_i|$ for all $t \in [0, T]$. When $\xi_{i0} > \tilde{\sigma}_i$ the option volatility process explodes in finite time and the boundary at ∞ is an exit boundary.

Proof. The proof is the same for $i = 1$ and 2 , so we can drop the subscript i for convenience. If $\xi_0 = \tilde{\sigma}$ then $\xi_t = \tilde{\sigma}$ for all $t > 0$. So in the following we consider two separate cases, according as $\xi_0 < \tilde{\sigma}$ and $\xi_0 > \tilde{\sigma}$.

It follows from (3.4) that $d\xi_t \rightarrow 0$ whenever $\xi_t \rightarrow 0$ or $\tilde{\sigma}$. So if the process is started with a value $\xi_0 < \tilde{\sigma}$ then ξ_t will remain bounded between 0 and $\tilde{\sigma}$ ($\tilde{\sigma} \geq \xi_t \geq 0$) for all $t \in [0, T]$. On the other hand, when $\xi_0 > \tilde{\sigma}$, ξ_t is bounded below by $\tilde{\sigma}$ but not bounded

above, and $\xi_t - \tilde{\sigma} \geq 0$. Setting $x_t = \frac{1}{\tilde{\sigma}} \ln \left| \frac{\xi_t - \tilde{\sigma}}{\xi_t} \right|$ and applying Itô's Lemma, we can show that $dx_t = \frac{1}{2} \tilde{\sigma} dt - dW_t$ and $x_t = x_0 + \frac{1}{2} \tilde{\sigma} t - W_t$. Substituting x_t back into the above equation yields $\xi_t = \tilde{\sigma} (1 + ke^{-\tilde{\sigma} W_t^*})^{-1}$.

Now we show that the option volatility process explodes when $\xi_{i0} > \tilde{\sigma}_i$ and the boundary at ∞ is an exit boundary. That is, ξ_{it} reaches ∞ in finite time and once it reaches ∞ , it stays there. From the above equation, we see that $\xi_t \rightarrow \infty$ when $W_t^* \rightarrow \tilde{\sigma}^{-1} \ln |k|$. So the volatility process could reach ∞ in finite time. However, when $W_t^* < \tilde{\sigma}^{-1} \ln |k|$ the above equation implies that ξ_t could become negative, which cannot be true. So to prove that ξ_t remains strictly positive, we need to know more about the boundary at ∞ . If the boundary is an "exit" boundary, then ξ_t cannot return back once it enters the region. That is, if $\xi_\tau = \infty$ for some stopping time $0 \leq \tau \leq T$, then $\xi_s = \infty$ for all $s > \tau$.

In fact, we can indeed classify ∞ as an exit boundary, and to show this we perform the test described in Lewis (2000), Durrett (1996), and Karlin and Taylor (1981). To this end, let $s(y)$, $m(y)$ be functions such that, for $0 < y < \infty$, $s(y) = \exp\{-\int^y \frac{2\alpha(x)}{\beta(x)^2} dx\}$, $m(y) = \beta(y)^2 s(y)$. Define $S(c, d)$, $M(c, d)$, and $N(d)$ as $S(c, d) = \int_c^d s(y) dy$, $M(c, d) = \int_c^d \frac{1}{m(y)} dy$, $N(d) = \lim_{c \downarrow 0} \int_c^d \frac{S(c, y)}{m(y)} dy$. Then, the Feller test states that the boundary at ∞ is an exit boundary of the process if $M(0, d) = \infty$ and $N(0) < \infty$. In our case, we have $\alpha(x) = x^2(x - \tilde{\sigma})$ and $\beta(x) = -x(x - \tilde{\sigma})$, so

$$\begin{aligned} m(y) &= \tilde{\sigma}^2 y^2; & s(y) &= \exp\left\{-\int^y \frac{2}{x - \tilde{\sigma}} dx\right\} = \frac{\tilde{\sigma}^2}{(y - \tilde{\sigma})^2} \\ M(c, d) &= \frac{1}{\tilde{\sigma}^2} \left[\frac{1}{c} - \frac{1}{d}\right]; & S(c, d) &= \tilde{\sigma}^2 \left[\frac{1}{\tilde{\sigma} - d} - \frac{1}{\tilde{\sigma} - c}\right] \\ N(d) &= \int_0^d S(0, x) m(x) dx = -\frac{d}{\tilde{\sigma}} - \ln |d - \tilde{\sigma}| + \ln \tilde{\sigma}. \end{aligned}$$

This shows that $M(0, d) = \infty$ and $N(0) < \infty$, hence ξ_{it} explodes and ∞ is an exit boundary. \square

Now that we have characterized the option price volatility, it is straightforward to find the option price under our lognormal approximation, as follows:

LEMMA 3.3. *When the option volatilities are given by (3.5) the call option price at time t is*

$$(3.6) \quad U_{it} = U_{i0} e^{rt} \frac{(e^{\tilde{\sigma}_i W_{it}^*} + k_i)}{1 + k_i},$$

where $k_i = (\frac{\tilde{\sigma}_i}{\xi_{i0}} - 1)$ and $W_{it}^* = -\frac{1}{2} \tilde{\sigma} t + W_{it}$. Moreover, when $\xi_{i0} > \tilde{\sigma}_i$, $U_{it} \rightarrow 0$ as $\xi_{it} \rightarrow \infty$.

Proof. Given the option price SDE (3.1), dropping the subscript i and solving for U , we get

$$(3.7) \quad U_t = U_0 \exp\left(rt - \frac{1}{2} \int_0^t \xi_t^2 dt + \int_0^t \xi_t dW_t\right).$$

Substituting $d(\ln |\tilde{\sigma} - \xi_t|) = \frac{1}{2} \xi_t^2 dt - \xi_t dW_t$ in the above equation gives

$$\begin{aligned}
(3.8) \quad U_t &= U_0 \exp\left(rt - \int_0^t d(\ln|\tilde{\sigma} - \xi_t|)\right) \\
&= U_0 \exp(rt - \ln|\tilde{\sigma} - \xi_t| + \ln|\tilde{\sigma} - \xi_0|) \\
&= U_0 e^{rt} \left(\frac{\tilde{\sigma} - \xi_0}{\tilde{\sigma} - \xi_t}\right),
\end{aligned}$$

which can be rewritten as equation (3.6). \square

From Lemma 3.3, since ξ_t is bounded between 0 and $\tilde{\sigma}$ when $\xi_0 < \tilde{\sigma}$ and bounded below by $\tilde{\sigma}$ when $\xi_0 > \tilde{\sigma}$, we can conclude that U_t will remain strictly positive for all time $t \in [0, T]$. However, for $\xi_0 > \tilde{\sigma}$, if the volatility explodes ($\xi_t \rightarrow \infty$), equation (3.8) shows that $U_t \rightarrow 0$. Moreover, once U_t reaches zero, it will stay there.

The option price process (3.6) will follow an approximately lognormal process if $\xi_{i0} \approx \tilde{\sigma}_i$, and we call this the weak lognormality condition. But when is $\xi_{i0} \approx \tilde{\sigma}_i$? By definition, $\xi_{it} \rightarrow \sigma_i$ if $\frac{\partial U_{it}}{\partial S_{it}} \rightarrow 1$ and $\frac{S_{it}}{U_{it}} \rightarrow 1$ as $t \rightarrow T$. Then, $\sigma_i \approx \tilde{\sigma}_i$ and therefore $\xi_{i0} \approx \tilde{\sigma}_i$. Moreover, from equation (3.5), $\xi_{it} \approx \tilde{\sigma}_i$, for all $t \in [0, T]$. Under the weak lognormality condition, the option price volatilities ξ_{it} and η_{it} are directly approximated as constants in equations (3.1) and (3.2), respectively. This allows us to approximate the option price processes as lognormal processes. Hence, we can change the numeraire to be one of the option prices, so that the price of a CEO may be expressed as the price of a single-asset option, and we can price the CEO using the formula given by Margrabe (1978).

The weak lognormality condition is used to derive analytic approximations to basket option prices in Section 4. However, for some multi-asset options it is possible to obtain an almost exact price. The following result provides a strong lognormality condition, under which the relative option price follows a lognormal process almost exactly.

THEOREM 3.4. *The CEO on calls has the same price as a standard single-asset option under a modified yet equivalent measure if the following condition holds*

$$(3.9) \quad U_{10} \frac{k_1}{1+k_1} - U_{20} \frac{k_2}{1+k_2} = 0.$$

Proof. The call CEO on two calls has time 0 price given by

$$\begin{aligned}
(3.10) \quad f_0 &= e^{-rT} \mathbb{E}_{\mathbb{Q}} \left\{ \left[\frac{U_{10} e^{rT}}{1+k_1} (e^{\tilde{\sigma}_1 W_{1T}^*} - k_1) - \frac{U_{20} e^{rT}}{1+k_2} (e^{\tilde{\sigma}_2 W_{2T}^*} - k_2) \right]^+ \right\} \\
&= \mathbb{E}_{\mathbb{Q}} \left\{ \left[\frac{U_{10}}{1+k_1} e^{\tilde{\sigma}_1 W_{1T}^*} - \frac{U_{20}}{1+k_2} e^{\tilde{\sigma}_2 W_{2T}^*} - \left(U_{10} \frac{k_1}{1+k_1} - U_{20} \frac{k_2}{1+k_2} \right) \right]^+ \right\} \\
&= \mathbb{E}_{\mathbb{Q}} \left\{ \left[\frac{U_{10}}{1+k_1} \exp\left(\int_0^T -\frac{1}{2} \tilde{\sigma}_1^2 ds + \tilde{\sigma}_1 dW_{1s}\right) \right. \right. \\
&\quad \left. \left. - \frac{U_{20}}{1+k_2} \exp\left(\int_0^T -\frac{1}{2} \tilde{\sigma}_2^2 ds + \tilde{\sigma}_2 dW_{2s}\right) \right]^+ \right\}.
\end{aligned}$$

Let $dW_{1t} = \rho dW_{2t} + \rho' dZ_{1t}$, where W_2 and Z_1 are independent Wiener processes, $\rho' = \sqrt{1 - \rho^2}$ and \mathbb{P} is a probability measure whose Radon-Nikodym derivative with respect to \mathbb{Q} is given by

$$\frac{d\mathbb{P}}{d\mathbb{Q}} = \exp\left(-\frac{1}{2}\tilde{\sigma}_2^2 t + \tilde{\sigma}_2 W_{2t}\right).$$

Let $Y_t = \frac{U_{1t}(1+k_2)}{U_{2t}(1+k_1)}$ and $Z_{2t} = W_{2t} - \tilde{\sigma}_2 t$. Then Z_1 and Z_2 are independent Brownian motions under the measure \mathbb{P} and the dynamics of Y can be described by

$$d(\ln Y_t) = \left(-\frac{1}{2}\tilde{\sigma}_1^2 + \frac{1}{2}\tilde{\sigma}_2^2\right) dt + (\rho\tilde{\sigma}_1 - \tilde{\sigma}_2) dW_{2t} - \rho'\tilde{\sigma}_1 dZ_{1t} = -\frac{1}{2}\tilde{\sigma}^2 dt + \tilde{\sigma}_1 dW_{1t},$$

where W is a Brownian motion under \mathbb{P} and $\tilde{\sigma}^2 = (\rho\tilde{\sigma}_1 - \tilde{\sigma}_2)^2 + (\rho'\tilde{\sigma}_1)^2 = \tilde{\sigma}_1^2 + \tilde{\sigma}_2^2 - 2\rho\tilde{\sigma}_1\tilde{\sigma}_2$. Now the CEO price can be written as the price of a single-asset option written on Y , as

$$f_0 = \frac{U_{20}}{1+k_2} \mathbb{E}_{\mathbb{P}} \left\{ \left[Y_t \exp\left(-\frac{1}{2} \int_0^T \tilde{\sigma}_s^2 ds + \int_0^T \tilde{\sigma}_s dW_s\right) - 1 \right]^+ \right\} = \frac{U_{20}}{1+k_2} \mathbb{E}_{\mathbb{P}} \{ [Y_T - 1]^+ \}.$$

When $\tilde{\sigma}_i > \xi_{i0}$, for $i = 1, 2$, the above expectation simply yields the Black-Scholes price of a European option on Y_t with strike equal to one. But when $\tilde{\sigma}_i < \xi_{i0}$, ξ_{it} could reach ∞ for some $\tau \leq T$. However, when the volatility explodes $U_{is} = 0$ for $s \geq \tau$, since the boundary at ∞ is an exit boundary, and the expectation in equation (3.10) need only be computed over the paths for which the individual option volatilities remain finite. Now by Lemma 3.2,

$$Z_{1t} > \rho'^{-1} \left(\frac{1}{2}(\tilde{\sigma}_1 - \tilde{\sigma}_2)t + \tilde{\sigma}_1^{-1} \ln |k_1| - \tilde{\sigma}_2^{-1} \ln |k_2| \right) = \mu_1$$

$$\text{and } Z_{2t} > -\frac{1}{2}\tilde{\sigma}_2 t + \tilde{\sigma}_2^{-1} \ln |k_2| = \mu_2.$$

Hence, setting $m_i = \min \{Z_{is}; 0 \leq s \leq T\}$, the price of the CEO may be written

$$(3.11) \quad f_0 = \frac{U_{20}}{1+k_2} \mathbb{E}_{\mathbb{P}} \{ \mathbf{1}_{m_1 > \mu_1; m_2 > \mu_2} [Y_T - 1]^+ \} + \frac{U_{10}}{1+k_1} \mathbb{E}_{\mathbb{Q}} \{ \mathbf{1}_{m_1 > \mu_1; m_2 < \mu_2} [e^{\tilde{\sigma}_1 W_{1T}^*} - k_1]^+ \},$$

where $\mathbf{1}$ is the indicator function. The first term on the right-hand side gives the expected value of the pay-off when neither of the volatilities explode. This is equal to the price of a down-and-out barrier exchange option which expires if either of the asset prices crosses the barrier. The second term is equivalent to a single-asset external barrier option when only ξ_2 explodes. Then U_{2T} becomes zero and the CEO pay-off reduces to U_{1T} .

Single-asset barrier options can be priced by an application of the reflection principle (see Karatzas and Shreve 1991) and the case of a two-asset barrier exchange option is an extension of that.¹ Using the reflection principle, the first term in the right-hand side of equation (3.11) may be written

$$\mathbb{E}_{\mathbb{P}} \{ \mathbf{1}_{Y_T > 1; m_1 > \mu_1; m_2 > \mu_2} \} = \mathbb{E}_{\mathbb{P}} \{ \mathbf{1}_{Y_T > 1; m_1 > \mu_1} \} - \mathbb{E}_{\mathbb{P}} \{ \mathbf{1}_{\ln Y_T < 2\mu_2; m_1 > \mu_1} \}.$$

¹ Carr (1995), Banerjee (2003), and Kwok, Wu, and Yu (1998) discuss the pricing of two-asset and multi-asset external barrier options that knock out if an external process crosses the barrier. Lindset and Persson (2006) discuss the pricing of two-asset barrier exchange options, where the option knocks out if the price of one asset equals the other.




FIGURE 3.1. Plot of $U_{1t}k_1/(1+k_1) - U_{2t}k_2/(1+k_2)$ against strike ($S_1 = 75$, $S_2 = 65$, $\sigma_1 = 0.25$, $\sigma_2 = 0.25$, $r = 4\%$, $T = 3$ months).

The two expectation terms on the right are equivalent to the ITM probabilities of a call option ($Y_T > 1$) and a put option ($Y_T < \exp(2\mu_2)$) with external barriers. These may be computed using the results of Carr (1995) and Kwok, Wu, and Yu (1998). Similarly, the second term in the right-hand side of equation (3.11) may also be evaluated as a combination of external barrier options. \square

Whether we price the multi-asset option under the weak or the strong lognormality condition, we must calibrate the strikes K_i of the vanilla options in the CEOs that replicate the pay-off so that $\sum \theta_i K_i = K$. Furthermore, when applying the strong lognormality condition we aim to find particular strikes K_i for which the relative price of these vanilla options is lognormal. When the underlying option prices satisfy condition (3.9) almost exactly, Theorem 3.4 allows one to price CEOs almost exactly. As discussed earlier, the approximation error can be extremely small for certain strikes of the vanilla options and, since we are calibrating these strikes, our approximation can be justified. For instance, Figure 3.1 plots the behavior of condition (3.9) for two sample vanilla options and shows that the condition holds when the strikes of the options are equal to 60.8.

In fact, we can price a CEO by dimension reduction under both the weak and the strong lognormality conditions. In both cases a CEO becomes equivalent to a simple lognormal exchange option, and the price of such an option can be found by change of numeraire, as in Margrabe (1978). Under the weak lognormality condition the vanilla option prices follow approximate lognormal processes, and under the strong lognormality condition their relative prices follow lognormal processes, because the displacement terms cancel out under the condition (3.9).

Finally, we remark on extending our results to asset price processes that are more general than GBM. When we apply the weak lognormality condition σ_i must be constant, but many choices of drift are possible. For instance, for non-traded assets such as commodity spots, we may choose $\mu_{it} = \kappa(\theta(t) - \ln S_{it})$ as in Schwartz (1997) and Pilipovic (2007). Theorem 3.4 may also be extended to cases where the underlying asset prices follow certain nonlognormal processes. To see this, suppose the risk-neutral price dynamics are governed by a general two-factor model: $dS_{it} = \mu_i(S_{it}, t)dt + \sigma_i(S_{it}, t)dW_{it}$, $\langle dW_{1t}, dW_{2t} \rangle = \rho dt$ for $i = 1, 2$, where ρ is constant. It is easy to show that the option

price processes will still be given by (3.6) whenever $\mu_i(S_{it}, t)$ and $\sigma_i(S_{it}, t)$ satisfy

$$\left(\frac{\partial \sigma_{it}}{\partial t} + \mu_{it} \frac{\partial \sigma_{it}}{\partial S_{it}} + \frac{1}{2} \sigma_{it}^2 \frac{\partial^2 \sigma_{it}}{\partial S_{it}^2} \right) = \sigma_{it} \frac{\partial \mu_{it}}{\partial S_{it}}.$$

4. PRICING BASKET OPTIONS

First, we apply the strong lognormality condition and use the formula of Margrabe (1978) recursively to derive almost exact prices for specific examples of basket options with two or three assets. However, when $N \geq 4$ the condition in Theorem 3.4 is too strong. In that case, we derive analytic approximate basket option prices under the weak lognormality condition.

Consider a CEO written on two lognormal exchange options, both having a common asset which is used as numeraire in the method of Margrabe (1978). Then the exchange option price processes may be described by equations (3.1) or (3.2), and the CEO can be priced by applying Theorem 3.4. For example, the pay-off to a two-asset basket option can be written as a sum of pay-offs to two CEOs on single-asset call and put options, as in Section 3. A three-asset basket option with zero strike, when the signs of the asset weights Θ are a permutation of $(1, 1, -1)$ or $(-1, -1, 1)$, is just an extension of the two-asset case where we have an additional asset instead of the strike. The option can be priced as a CEO either to exchange a two-asset exchange option for the third asset or to exchange 2 two-asset exchange options with a common asset. For example, a 3:2:1 spread option has pay-off decomposition²

$$\begin{aligned} P_T &= [3S_{1T} - 2S_{2T} - S_{3T}]^+ \\ &= [3[S_{1T} - S_{2T}]^+ - [S_{3T} - S_{2T}]^+]^+ + [[S_{2T} - S_{3T}]^+ - 3[S_{2T} - S_{1T}]^+]^+. \end{aligned}$$

In the general case of basket options on N underlying assets, except for the ones discussed above, the two replicating CEOs are no longer written on plain vanilla or lognormal exchange options, but on sub-basket options. For instance, consider a four-asset basket option with zero strike and pay-off $P_T = [S_{1T} - S_{2T} - S_{3T} + S_{4T}]^+$. We have

$$P_T = [[S_{1T} - S_{2T}]^+ - [S_{3T} - S_{4T}]^+]^+ + [[S_{4T} - S_{3T}]^+ - [S_{2T} - S_{1T}]^+]^+,$$

and since the two replicating CEOs are written on lognormal exchange options with no common asset, the CEOs cannot be priced using Theorem 3.4. But we can adjust the volatilities of the CEOs using the weak lognormality condition, so that the sub-basket option price processes are approximately lognormal processes. Then the relative sub-basket option prices also follow approximate lognormal processes and the two replicating CEO prices can be computed by using the formula of Margrabe (1978). Thus, exact pricing under the strong lognormality condition is only possible in special cases, and in the general case we must use the weak lognormality condition to find an approximate price.

² Note that there is a closed form formula for the price of this option of the form

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\ln \frac{3}{2} + x} \int_{-\infty}^{\ln(3e^x - 2e^y)} (3e^x - 2e^y - e^x) f(x, y, z) dz dy dx,$$

where f is the trivariate normal density function and x, y, z are the log stock price processes. However, it is not easy to evaluate the triple integral.

Let $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$ be a filtered probability space, where Ω is the set of all possible events such that $(S_{1t}, S_{2t}, \dots, S_{Nt}) \in (0, \infty)^N$, $(\mathcal{F}_t)_{t \geq 0}$ is the filtration produced by the sigma algebra of the N -tuple $(S_{1t}, S_{2t}, \dots, S_{Nt})_{t \geq 0}$ of asset prices, and \mathbb{Q} is a multivariate risk-neutral probability measure. We assume that the underlying asset prices processes S_i are described by

$$dS_{it} = \mu_i(S_{it}, t)S_{it} dt + \sigma_i S_{it} dW_{it}, \quad \langle dW_{it}, dW_{jt} \rangle = \rho_{ij} dt, \quad 1 \leq i, j \leq N,$$

where W_i are Wiener processes under the risk-neutral measure \mathbb{Q} , σ_i is the volatility of i th asset (assumed constant), $\mu_i(\cdot)$ is a well-defined function of S_{it} and t , and ρ_{ij} is the correlation between i th and j th assets (assumed constant).

We now describe the price process V_{Nt} of the basket option on N assets. Using a recursive argument, we begin by assuming that the prices of the call and put sub-basket options on m and n assets follow lognormal processes. Then we show that when the basket option volatility is approximated as a constant the basket option price process V_{Nt} can be expressed as a lognormal process. Since C_{mt} , C_{nt} , P_{mt} , and P_{nt} are prices of basket options themselves, we may also express their processes as lognormal process, assuming their sub-basket option prices follow lognormal processes. In the end, these assumptions yield an approximate lognormal process for the price of a basket option on N assets.

As before, let $\Theta_N = (\Theta_m, -\Theta_n)$. Then the basket option price may be computed as a sum of the price E_{1t} of a CEO on two sub-basket calls and the price E_{2t} of a CEO on two sub-basket puts, that is, $V_{Nt} = E_{1t} + E_{2t}$. Each sub-basket option follows a price process with a nonconstant volatility, but we shall express it using a constant. For instance, the sub-basket call option price processes are written, for $i = m$ and n

$$dC_{it} = r C_{it} dt + \bar{\sigma}_{ci} C_{it} d\tilde{W}_{it},$$

where \tilde{W}_i is a Wiener process and $\bar{\sigma}_{ci}$ is a constant.³ Then, by Itô's Lemma

$$(4.1) \quad dE_{1t} = r E_{1t} dt + \sum_{i=m,n} \bar{\sigma}_{ci} C_{it} \Delta_{C_{it}} d\tilde{W}_{it}.$$

Similarly, we assume the price E_{2t} of the CEO on puts follows a process analogous to (4.1) with C_i replaced by P_i ; and when the decomposition gives two CEOs written on call and put sub-basket options, so their price processes have a call and a put option component, we have

$$dE_{1t} = r E_{1t} dt + \bar{\sigma}_{cm} C_{mt} \Delta_{C_{mt}} d\tilde{W}_{mt} + \bar{\sigma}_{pn} P_{nt} \Delta_{P_{nt}} d\tilde{W}_{nt}.$$

Now write $d\tilde{W}_{nt} = \gamma_{mn} d\tilde{W}_{mt} + \sqrt{1 - \gamma_{mn}^2} d\tilde{W}_t$, where \tilde{W} is a Wiener process, independent of \tilde{W}_m , and γ_{mn} is the correlation between the options on sub-baskets of size m and n . Further, define

$$\tilde{\sigma}_{it} = \bar{\sigma}_{ci} \frac{C_{it}}{V_{Nt}} \frac{\partial E_{1t}}{\partial C_{it}} - \bar{\sigma}_{pi} \frac{P_{it}}{V_{Nt}} \frac{\partial E_{2t}}{\partial P_{it}},$$

for $i = m$ and n .

³ For brevity, we suppress the dependence of $\bar{\sigma}_{ci}$ on the option's strike and maturity, the discount rate etc.

Then

$$\begin{aligned}
dV_{Nt} &= r(E_{1t} + E_{2t}) dt + \sum_{i=m,n} \bar{\sigma}_{Ci} C_{it} \Delta_{C_{it}} d\tilde{W}_{it} - \sum_{i=m,n} \bar{\sigma}_{Pi} P_{it} \Delta_{P_{it}} d\tilde{W}_{it}, \\
&= r V_{Nt} dt + V_{Nt} \left(\left(\bar{\sigma}_{Cm} \frac{C_{mt}}{V_{Nt}} \frac{\partial E_{1t}}{\partial C_{mt}} - \bar{\sigma}_{Pm} \frac{P_{mt}}{V_{Nt}} \frac{\partial E_{2t}}{\partial P_{mt}} \right) d\tilde{W}_{mt} \right. \\
&\quad \left. - \left(\bar{\sigma}_{Cn} \frac{C_{nt}}{V_{Nt}} \frac{\partial E_{1t}}{\partial C_{nt}} - \bar{\sigma}_{Pn} \frac{P_{nt}}{V_{Nt}} \frac{\partial E_{2t}}{\partial P_{nt}} \right) d\tilde{W}_{nt} \right), \\
&= r V_{Nt} dt + V_{Nt} (\tilde{\sigma}_{mt} d\tilde{W}_{mt} - \tilde{\sigma}_{nt} d\tilde{W}_{nt}).
\end{aligned}$$

Now setting

$$\begin{aligned}
\tilde{\sigma}_t^2 &= \tilde{\sigma}_{mt}^2 + \tilde{\sigma}_{nt}^2 - 2\gamma_{mn} \tilde{\sigma}_{mt} \tilde{\sigma}_{nt}, \\
\tilde{\sigma}_{E_{1t}}^2 &= \tilde{\sigma}_{Cmt}^2 + \tilde{\sigma}_{Cnt}^2 - 2\gamma_{mn} \tilde{\sigma}_{Cmt} \tilde{\sigma}_{Cnt}, \\
\tilde{\sigma}_{E_{2t}}^2 &= \tilde{\sigma}_{Pmt}^2 + \tilde{\sigma}_{Pnt}^2 - 2\gamma_{mn} \tilde{\sigma}_{Pmt} \tilde{\sigma}_{Pnt}
\end{aligned}$$

yields

$$(4.2) \quad dV_{Nt} = r V_{Nt} dt + \tilde{\sigma}_t V_{Nt} d\tilde{W}_t.$$

So that its price is given by Margrabe's formula, the basket option price V_{Nt} is approximated by a lognormal process, that is, we replace the volatility in the SDE (4.2) by a constant, $\bar{\sigma}$. Thus, in place of (4.2) we assume $dV_{Nt} = r V_{Nt} dt + \bar{\sigma} V_{Nt} d\tilde{W}_t$.

Let \tilde{v} be a positive real such that $\bar{\sigma} = \tilde{v} \tilde{\sigma}_0$, where $\tilde{\sigma}_0$ is the value of $\tilde{\sigma}_t$ at time zero. Then $dV_{Nt} = r V_{Nt} dt + \tilde{v} \tilde{\sigma}_0 V_{Nt} d\tilde{W}_t$. We have also assumed that the underlying call and put sub-basket options follow lognormal processes, with average volatilities $\bar{\sigma}_{Ci}$ and $\bar{\sigma}_{Pi}$, respectively. Thus, we could likewise set $\bar{\sigma}_{Ci} = \tilde{v}_{Ci} \tilde{\sigma}_{Ci,0}$ and $\bar{\sigma}_{Pi} = \tilde{v}_{Pi} \tilde{\sigma}_{Pi,0}$. However, for ease of calibration we shall use only one parameter for the volatilities of all the options above the base level of the pricing tree. That is, we set all \tilde{v} 's equal to a constant, ν , which we call the "volatility correction factor" (VCF).⁴ Thus, $\bar{\sigma} \approx \nu \tilde{\sigma}_0$, $\bar{\sigma}_{Ci} \approx \nu \tilde{\sigma}_{Ci,0}$, $\bar{\sigma}_{Pi} \approx \nu \tilde{\sigma}_{Pi,0}$, and so on. Note that $\tilde{\sigma}_{E_{1t}}$ is also replaced by a constant $\bar{\sigma}_{E_{1t}}$, and with a single VCF ν , $\bar{\sigma}_{E_{1t}} \approx \nu (\tilde{\sigma}_{Cm0}^2 + \tilde{\sigma}_{Cn0}^2 - 2\gamma_{mn} \tilde{\sigma}_{Cm0} \tilde{\sigma}_{Cn0})^{1/2}$ and similarly for $\bar{\sigma}_{E_{2t}}$.

THEOREM 4.1. *The price of a general N -asset basket option may be approximated by the recursive formula*

$$(4.3) \quad V_{Nt}(\Theta, \mathbf{S}_t, \mathbf{K}, T, \nu, \omega) = E_{1t}(\Theta, \mathbf{S}_t, \mathbf{K}, T, \nu, \omega) + E_{2t}(\Theta, \mathbf{S}_t, \mathbf{K}, T, \nu, \omega),$$

⁴The VCF does not apply to single-asset option volatilities. When the sub-basket size reduces to one, $\nu = 1$ and the option has price process (3.1) if it is a call, or (3.2) if it is a put. The sole function of the VCF is to approximate the process volatilities of the options above the base level of the tree.

where

$$\begin{aligned}
E_{1t}(\Theta, \mathbf{S}_t, \mathbf{K}, T, \nu, \omega) &= \omega(V_{mt}(\Theta_m, \mathbf{S}_{mT}, \mathbf{K}_m, \nu, +1)\Phi(\omega d_{11}) \\
&\quad - V_{nt}(\Theta_n, \mathbf{S}_{nT}, \mathbf{K}_n, \nu, -\chi)\Phi(\omega d_{12})) \\
&= \omega(V_{mt}^1\Phi(\omega d_{11}) - V_{nt}^1\Phi(\omega d_{12})), \text{ say} \\
E_{2t}(\Theta, \mathbf{S}_t, \mathbf{K}, T, \omega) &= \omega(V_{nt}(\Theta_n, \mathbf{S}_{nT}, \mathbf{K}_n, \nu, \chi)\Phi(\omega d_{21}) \\
&\quad - V_{mt}(\Theta_m, \mathbf{S}_{mT}, \mathbf{K}_m, \nu, -1)\Phi(\omega d_{22})) \\
&= \omega(V_{nt}^2\Phi(\omega d_{21}) - V_{mt}^2\Phi(\omega d_{22})), \text{ say}
\end{aligned}$$

with $\chi = +1$ for a call and -1 for a put, and

$$d_{i1} = \frac{\ln(V_{mt}^i/V_{nt}^i) + \frac{1}{2}\nu^2\bar{\sigma}_{Ei}^2(T-t)}{\nu\bar{\sigma}_{Ei}\sqrt{T-t}}; \quad d_{i2} = d_{i1} - \nu\bar{\sigma}_{Ei}\sqrt{T-t}.$$

Proof. Since C_m, P_m and C_n, P_n are themselves prices of options on baskets of sizes m and n , respectively, they can be computed by applying equation (4.3) recursively until the size of a sub-basket reaches one and $\mathbf{S}_t = (S_{it}), \mathbf{K}_1 = (K_i)$, and $\Theta_1 = \theta_i$ for some $1 \leq i \leq N$. Then

$$E_1(\Theta, \mathbf{S}_t, \mathbf{K}, T, 1, \omega) = \omega e^{-r(T-t)}\theta_i(F_{it,T}\Phi(\omega d_1) - K_i\Phi(\omega d_2)), \quad E_2(\Theta, \mathbf{S}_t, \mathbf{K}, T, 1, \omega) = 0,$$

where $F_{it,T}$ is the i th asset forward price and

$$d_1 = \left[\ln(F_{it,T}/K_i) + \left(r + \frac{1}{2}\Sigma_i^2 \right) (T-t) \right] / [\Sigma_i\sqrt{T-t}], \quad d_2 = d_1 - \Sigma_i\sqrt{T-t}.$$

For instance, when $\mu_i = (r - q_i)$ in equation (4.1), $F_{it,T} = S_{it}e^{(r-q_i)(T-t)}$ and $\Sigma_i = \sigma_i$. Or more generally, when $\mu_i = \kappa(\theta(t) - \ln S_{it})$, as in Pilipovic (2007)

$$\begin{aligned}
F_{it,T} &= \exp\left(e^{-\kappa(T-t)} \ln S_{it} + \int_t^T e^{-\kappa(T-s)}\theta(s) ds + \frac{\sigma_i^2}{2\kappa}(1 - e^{-2\kappa(T-t)}) \right), \\
\Sigma_i &= \sigma_i \sqrt{\frac{1 - e^{-2\kappa(T-t)}}{2\kappa}}.
\end{aligned}$$

□

One of the main advantages of our approximation is that we can derive analytic formulae for the multi-asset option Greeks which, unlike most other approximations, capture the effects that individual asset's volatilities and correlations have on the hedge ratios. Differentiating the basket option price given in Theorem 4.1, using the chain rule, yields the following

COROLLARY 4.2. *The deltas, gammas, and vegas of our basket option price f are*

$$\begin{aligned}
(4.4) \quad \Delta_{S_i}^f &= \Delta_{C_j}^f \Delta_{S_i}^{C_j} + \Delta_{P_j}^f \Delta_{S_i}^{P_j} \\
\Gamma_{S_i}^f &= \Gamma_{C_j}^f (\Delta_{S_i}^{C_j})^2 + \Gamma_{S_i}^{C_j} \Delta_{C_j}^f + \Gamma_{P_j}^f (\Delta_{S_i}^{P_j})^2 + \Gamma_{S_i}^{P_j} \Delta_{P_j}^f \\
\mathcal{V}_{\sigma_i}^f &= \mathcal{V}_{\bar{\sigma}_{E1}}^f \frac{\partial \bar{\sigma}_{E1}}{\partial \sigma_i} + \mathcal{V}_{\bar{\sigma}_{E2}}^f \frac{\partial \bar{\sigma}_{E2}}{\partial \sigma_i} + \mathcal{V}_{\sigma_i}^{C_j} \Delta_{C_j}^f + \mathcal{V}_{\sigma_i}^{P_j} \Delta_{P_j}^f,
\end{aligned}$$

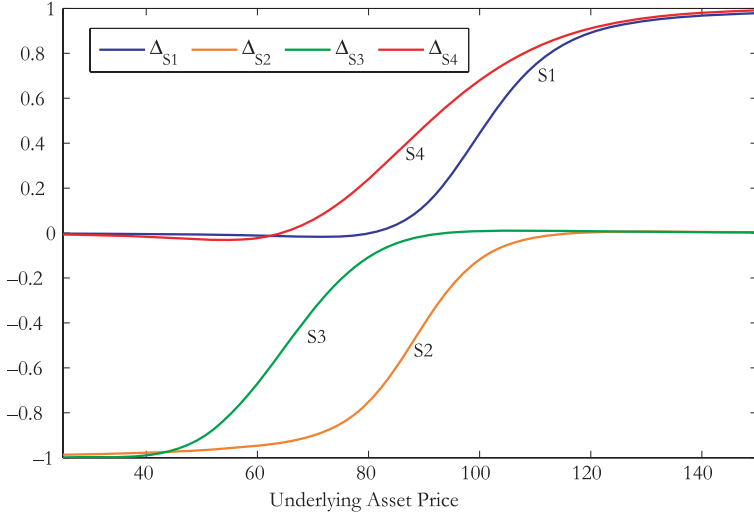


FIGURE 4.1. Deltas with respect to all four underlying asset prices.

where j is equal to m when $1 \leq i \leq m$ and equal to n when $m + 1 \leq i \leq N$. Here Δ_x^z , Γ_x^z , and \mathcal{V}_x^z denote the delta, gamma, and vega of z with respect to x , respectively. The corresponding formulae for a rainbow option are obtained from its basket-option representation.

For example, Figure 4.1 depicts the deltas of a four-asset basket option with pay-off given by $[S_1 - S_2 - S_3 + S_4]^+$.⁵ Here, we assumed the assets pay no dividends and that their current prices are $S_1 = 100$, $S_2 = 90$, $S_3 = 85$, and $S_4 = 75$, so the basket price is currently at zero. The discount rate is 4% and the volatilities and correlations of the asset prices are

$$(4.5) \quad \Sigma = \begin{pmatrix} 0.10 \\ 0.15 \\ 0.18 \\ 0.20 \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} 1 & 0.8 & 0.6 & 0.2 \\ 0.8 & 1 & 0.55 & 0.65 \\ 0.6 & 0.55 & 1 & 0.57 \\ 0.2 & 0.65 & 0.57 & 1 \end{pmatrix}.$$

The basket option is ATM with respect to every underlying asset (i.e., the price of every asset is equal to the weighted sum of prices of the other three assets). The basket has positive weights on assets one and four, so the deltas with respect to S_1 and S_4 resemble the deltas of vanilla call options, whereas it has negative weights on assets two and three so the deltas with respect to S_2 and S_3 resemble the deltas of vanilla put options. Due to differences in the asset's volatilities and correlations, the deltas differ. For instance, $\Delta_1 \approx 0.5$ but $\Delta_2 \approx -0.4$. This property is not captured by any other existing approaches to analytic approximations for multi-asset options, because they ignore the effects of asset price volatilities and correlations on the basket option deltas.

⁵ To derive the deltas in Figure 4.1 we have chosen the strikes of the single-asset options so that our approximate four-asset basket price matches its simulated price under GBM (further details are given in Section 5). But in practice we can also calibrate these strikes using the algorithm described in Section 6. Either way, an advantage of our approach is that analytic Greeks are obtainable, using (4.4), and it is not necessary to derive them using time-consuming simulations.

5. SIMULATION RESULTS

This section tests the accuracy of the weak lognormal approximation by comparing the simulated price of various options based on standard correlated GBM processes with that based on the approximate process (4.2). We have seen that we price a multi-asset option as a sum of CEOs on sub-basket options. These sub-basket options are in turn priced as sums of CEOs of smaller baskets until we arrive at a sum of CEOs on vanilla call and put options. By rewriting the option price processes as approximate lognormal processes we are able to apply Margrabe's formula to compute the price of the CEO, and hence also the price of the multi-asset option. Therefore, the accuracy of our weak lognormal approximation for pricing CEOs is crucial for the accuracy of our multi-asset option price approximations.

Consider the compound call exchange option that is an option to receive a vanilla call on asset one in exchange for a vanilla call on asset two, which has pay-off

$$P_T = [U_{1T} - U_{2T}]^+, \text{ with } U_{1t} = [S_{1t} - K_1]^+ \text{ and } U_{2t} = [S_{2t} - K_2]^+.$$

We now simulate the price of this option in two different ways, by: (1) simulating the underlying asset prices themselves, using correlated GBMs; and (2) simulating the CEO price P_t directly using our lognormal approximation, viz:

$$(5.1) \quad \frac{dP_t}{P_t} = r dt + \xi_1 \frac{\partial P_t}{\partial U_{2t}} \frac{U_{1t}}{P_t} dW_{1t} - \xi_2 \frac{\partial P_t}{\partial U_{2t}} \frac{U_{2t}}{P_t} dW_{2t}.$$

Note that the deltas used here are those of the CEO with respect to the vanilla options, and ξ_1 and ξ_2 are the option volatilities, which are approximated using (3.6). Hence, in this case the weak lognormal approximation is done twice, once for each ξ_i ($i = 1, 2$).

The difference between these two CEO prices depends on moneyness (of the vanilla options and the CEO) and the option's characteristics. We therefore change moneyness by fixing $S_1 = 100$, letting $S_2 = 80, 90, 100, 110$ and setting $K_1 = K_2 = 60, 80, 100, 120$. To further limit the number of options considered we simulate only two maturities (1 and 6 months), set discount rates to 0% or 4%, let the asset price volatilities be 20% or 50%, and set their correlation to be 0.5 or 0.8.

Table 5.1 presents the results. To compute each option price we employed two million simulations, including antithetic sampling to reduce the sampling variance. The GBM price is obtained by simulating U_{1t} and U_{2t} and then computing $[U_{1T} - U_{2T}]^+$ and taking the average over all two million simulations. The standard deviation is reported next to the GBM price. In the last column, we report our approximate prices, obtained by simulating (5.1). The approximation errors are very small indeed, especially for ITM options, as expected.⁶

6. MODEL CALIBRATION

To express the price process of a sub-basket option as an approximate log-normal process we have approximated the option volatility and assumed it is constant. Hence, any error

⁶ Results for other values of volatilities and correlations yield similar results. We also ran similar experiments to examine the effect of weak lognormal approximation on the price processes of vanilla options and three types of two-asset options: standard exchange options, call spread options, and two-asset basket calls. In every case the errors between prices from approximate and GBM processes were less than 10^{-2} ; in fact for ITM options they were usually much smaller than this.

TABLE 5.1
 CEO Prices: Correlated GBM Process versus Weak Lognormal Approximate Process
 ($S_1 = 100$)

S_2	K_1, K_2	r	T	σ_1	σ_2	ρ	Sim Price	Std Dev	Appr Price	Error
80	60	0	0.08	0.2	0.2	0.50	20.0001	0.0003	20.0007	-0.0006
80	60	0.04	0.08	0.2	0.2	0.50	20.0669	0.0003	20.0665	0.0004
80	60	0	0.08	0.2	0.2	0.80	20.0000	0.0002	20.0003	-0.0003
80	60	0.04	0.08	0.2	0.2	0.80	20.0667	0.0002	20.0690	-0.0023
80	100	0	0.50	0.2	0.2	0.50	5.3946	0.0048	5.3948	-0.0002
80	100	0.04	0.08	0.5	0.5	0.50	5.6498	0.0049	5.6496	0.0001
80	100	0.04	0.08	0.2	0.2	0.80	2.4787	0.0019	2.4782	0.0005
80	100	0	0.50	0.5	0.5	0.80	9.7367	0.0109	9.7370	-0.0003
80	120	0	0.08	0.2	0.5	0.50	0.0012	0.0000	0.0014	-0.0002
80	120	0.04	0.08	0.2	0.5	0.50	0.0015	0.0001	0.0019	-0.0005
80	120	0	0.08	0.5	0.5	0.80	0.7668	0.0023	0.7649	0.0019
80	120	0.04	0.08	0.5	0.5	0.80	0.8058	0.0024	0.8039	0.0019
90	60	0	0.08	0.2	0.2	0.80	10.0019	0.0002	10.0053	-0.0034
90	60	0.04	0.08	0.2	0.2	0.80	10.0353	0.0002	10.0407	-0.0054
90	100	0.04	0.08	0.2	0.2	0.50	2.4068	0.0018	2.4108	-0.0040
90	100	0	0.08	0.2	0.2	0.80	2.2308	0.0017	2.2347	-0.0039
90	100	0.04	0.08	0.2	0.2	0.80	2.3938	0.0017	2.3954	-0.0016
90	120	0	0.08	0.2	0.2	0.50	0.0013	0.0000	0.0011	0.0002
90	120	0.04	0.08	0.2	0.2	0.50	0.0016	0.0001	0.0014	0.0002
90	120	0	0.08	0.2	0.5	0.50	0.0009	0.0000	0.0008	0.0001
90	120	0.04	0.08	0.2	0.5	0.50	0.0011	0.0000	0.0013	-0.0002
90	120	0	0.08	0.2	0.2	0.80	0.0013	0.0000	0.0012	0.0001
90	120	0.04	0.08	0.2	0.2	0.80	0.0017	0.0001	0.0015	0.0001
90	120	0	0.50	0.2	0.2	0.80	0.6314	0.0020	0.6307	0.0008
100	60	0	0.08	0.2	0.2	0.50	2.3022	0.0017	2.2984	0.0038
100	60	0.04	0.08	0.2	0.2	0.50	2.3099	0.0018	2.3031	0.0068
100	60	0	0.08	0.2	0.5	0.50	5.0126	0.0032	5.0133	-0.0007
100	60	0.04	0.08	0.2	0.5	0.50	5.0294	0.0033	5.0304	-0.0010
100	60	0	0.08	0.2	0.5	0.80	4.1516	0.0026	4.1505	0.0011
100	60	0.04	0.08	0.2	0.5	0.80	4.1655	0.0026	4.1646	0.0009
100	120	0	0.08	0.2	0.2	0.50	0.0013	0.0000	0.0010	0.0003
100	120	0.04	0.08	0.2	0.2	0.50	0.0016	0.0001	0.0014	0.0002
100	120	0	0.08	0.2	0.5	0.50	0.0004	0.0000	0.0006	-0.0002
100	120	0.04	0.08	0.2	0.5	0.50	0.0005	0.0000	0.0006	-0.0001
110	60	0	0.08	0.2	0.2	0.80	0.0054	0.0001	0.0052	0.0002
110	60	0.04	0.08	0.2	0.2	0.80	0.0054	0.0001	0.0051	0.0003
110	80	0	0.08	0.2	0.2	0.50	0.1245	0.0005	0.1242	0.0003
110	80	0.04	0.08	0.2	0.2	0.50	0.1249	0.0005	0.1247	0.0002
110	80	0	0.08	0.2	0.2	0.80	0.0054	0.0001	0.0042	0.0013
110	120	0	0.08	0.2	0.2	0.50	0.0006	0.0000	0.0004	0.0002
110	120	0	0.08	0.2	0.5	0.50	0.0001	0.0000	0.0002	-0.0001
110	120	0.04	0.08	0.2	0.5	0.50	0.0002	0.0000	0.0003	-0.0001

in the pricing of a basket option arises from our approximation of the volatility of the sub-basket options in the pricing tree. Quantifying this error theoretically is difficult, given the dimensionality of the problem, but the previous section has circumvented this by quantifying it computationally. There we showed that the option prices obtained by simulating our approximate (weak) lognormal processes were very close to prices obtained using GBM processes.

Now we focus on the approximate price recursion, describing how this may be applied to basket option pricing. To this end we develop a general algorithm for choosing the strikes of the vanilla options and for approximating the volatilities of all the options in the pricing tree.

There are two stages to the algorithm: (1) construct the pricing tree; and (2) approximate the volatilities of the sub-basket options at each level in the tree. Stage 1 consists of three steps

- (1) Fix a partition of k , where k is the number of assets in the basket to be priced;
- (2) Use this partition to decompose the basket option pay-offs into sub-basket option pay-offs at every level of the pricing tree; and
- (3) Fix the strikes of the sub-basket options at every level of the tree and choose an appropriate permutation of asset prices. This should be consistent with the strike of the original basket option and the decomposition structure of the pricing tree.

For step (1) we advocate setting $[1: k - 1]$ as the partition, for $k = N, N - 1, \dots, 2$, because it will allow the volatilities in stage 2 to be approximated with minimal error (given a suitable choice of strikes at step (3), as discussed below). Thus, we choose a pay-off decomposition in step (2) where, at every stage in the pricing tree, the option on k assets is split into two sub-basket options: on $(k - 1)$ assets and a single asset. At step (3), the single-asset option will always be ATM, and the sub-basket option on $(k - 1)$ assets will be of very similar moneyness (defined as the ratio between the asset price and strike) as the k -asset option. The ATM option is written on the asset with the highest volatility among k assets while keeping the moneyness of the $k - 1$ -asset option as close to that of k -asset option's moneyness. The reason for choosing this convention is that the sub-basket option on $(k - 1)$ assets is the closest we can get to the option on k assets. Moreover, the resulting permutation would minimize the cascading of error due to volatility approximation in stage 2 from a lower level to a higher level of the tree.

Hence, in stage 1, the pricing tree for an N -asset option is constructed so that the terminal nodes contain $2(N - 1)$ ATM vanilla options and a pair of vanilla call and put options with the same moneyness as the N -asset option. A simple illustration is provided by a three-asset basket call with pay-off

$$\begin{aligned} [S_1 + S_2 + S_3 - K]^+ &= [[S_1 + S_2 - K_{12}]^+ - [K_3 - S_3]^+]^+ \\ &\quad + [[S_3 - K_3]^+ - [K_{12} - (S_1 + S_2)]^+]^+. \end{aligned}$$

The two-asset options have similar moneyness as the three-asset basket and we set $K_3 = S_3$. Similarly,

$$[[S_1 + S_2 - K_{12}]^+]^+ = [[S_1 - K_1]^+ - [K_2 - S_2]^+]^+ + [[S_2 - K_2]^+ - [K_1 - S_1]^+]^+,$$

and we set $K_2 = S_2$, which leaves $K_1 = K - S_2 - S_3$. Thus the vanilla options on S_2 and S_3 are ATM while the vanilla options on S_1 are of the similar moneyness as the basket option. If the chosen permutation does not permit the option on S_1 to be of similar

moneyness as of the basket option, then we rearrange the order picking the next higher asset price with similar volatility to S_1 .

Stage 2 consists of calibrating the VCF ν , defined in Section 4. Any basket option price could be used for its calibration, but we recommend using an ATM basket option price (simulated, or better, if available, its market price) because using an ITM (OTM) basket option for calibration leads to larger errors on OTM (ITM) basket options than we obtain using the ATM ν value. Having calibrated ν , the prices of other basket options and all hedge ratios are derived analytically, using (4.3) and (4.4).

To illustrate the calibration algorithm, consider the five-asset basket calls examined by Ju (2002), for which $S_1 = S_2 = S_3 = S_4 = S_5 = 100$, $\theta_1 = 35$, $\theta_2 = 25$, $\theta_3 = 20$, $\theta_4 = 15$, $\theta_5 = 5$, and $T = 1$ or 3 years. As in Ju (2002) the asset price volatilities are equal ($\sigma = 20\%$ or 50%), their correlations are also equal ($\rho = 0$ or 0.5), the discount rate is either 0% or 5% , and the five-asset basket call has strike 90, 100 or 110.⁷ Following step (1), we construct the pricing tree so that the strikes of the four ATM vanilla options at the base have $K_i = \theta_i$ for $i = 2, \dots, 5$ and the other vanilla option has $K_1 = 25, 35$, or 45 according as the five-asset basket call has strike 90, 100, or 110.

Table 6.1 reports the results from step (2), that is, the calibration of the VCF for each set of options. The column headed “simulation results” reports the basket option prices obtained by simulating the five correlated GBM processes for the underlying assets and taking the average price over two million simulations. The next column reports the standard error of these simulations and the column headed “approximate price” is the price obtained using our weak lognormal approximation in the calibration algorithm described above. The last two columns report the calibrated VCF and the difference between the two prices.

The VCF is sensitive to asset price volatilities and correlation. A VCF equal to or close to one implies that the approximation naturally yields accurate prices and there is very minimal or no correction required to the sub-basket option volatilities. However, a VCF different to one implies that the true sub-basket option volatilities may be different to the constant approximate volatility $\tilde{\sigma}$ and hence requires some correction.

Table 6.2 shows that the VCF is close to one for uncorrelated assets with relatively low volatility ($\sigma = 20\%$) implying that our approximate formula naturally works well for ATM options under such cases. But for uncorrelated assets with relatively high volatility ($\sigma = 50\%$) we have a VCF less than one (about 0.83–0.9) suggesting a downward correction to the approximate sub-basket volatilities. The largest VCFs (about 1.5–1.7) are for correlated assets with relatively low volatility ($\rho = 0.5$, $\sigma = 20\%$) in which case the sub-basket option volatilities are scaled up.

By construction, the error is very small for ATM options. Using the ATM-calibrated values for the VCF also produces relatively small errors for ITM options, and when pricing 3-year options the errors are generally less than when pricing 1-year options. The largest errors are for 1-year OTM options. Here errors can exceed 10% of the simulated price for options on uncorrelated assets. When the asset prices are correlated and the basket option is a 3-year OTM basket call, our approximate prices are slightly greater than the simulated prices, but in all other cases our approximation tends to over-price slightly.

⁷ In practice, we could set the volatilities of the vanilla options that are ATM to be equal to the ATM implied volatilities and the volatility of the vanilla option that has the same moneyness as the basket option to a value implied from the skew. Similarly, we could use implied correlation matrix, if available from market prices, or set it equal to some statistical estimate.

TABLE 6.1
Five-Asset Basket Prices: Correlated GBM versus Weak Lognormal Approximation

T	K	r	σ	ρ	Sim Price	Std Error	Approx Price	VCF
1	90	0.05	0.2	0	14.6254	0.0011	15.6434	0.9896
1	90	0.05	0.2	0.5	15.6479	0.0005	17.0599	1.7002
1	90	0.05	0.5	0	18.3388	0.0062	18.3719	0.8950
1	90	0.05	0.5	0.5	22.8694	0.0029	24.1031	1.3935
1	90	0.1	0.2	0	18.6285	0.0012	19.4295	1.0136
1	90	0.1	0.2	0.5	19.2149	0.0006	20.2970	1.4921
1	90	0.1	0.5	0	21.2996	0.0065	21.3933	0.9017
1	90	0.1	0.5	0.5	25.3757	0.0031	26.6393	1.3593
1	100	0.05	0.2	0	6.8143	0.0009	6.8150	0.9896
1	100	0.05	0.2	0.5	8.8929	0.0004	8.8929	1.7002
1	100	0.05	0.5	0	12.6438	0.0054	12.6433	0.8950
1	100	0.05	0.5	0.5	17.8991	0.0028	17.8990	1.3935
1	100	0.1	0.2	0	10.307	0.0011	10.3068	1.0136
1	100	0.1	0.2	0.5	11.9199	0.0005	11.9199	1.4921
1	100	0.1	0.5	0	15.2241	0.0058	15.2236	0.9017
1	100	0.1	0.5	0.5	20.2037	0.0028	20.2036	1.3593
1	110	0.05	0.2	0	2.2074	0.0007	4.1564	0.9896
1	110	0.05	0.2	0.5	4.3969	0.0004	5.0478	1.7002
1	110	0.05	0.5	0	8.426	0.005	9.6427	0.8950
1	110	0.05	0.5	0.5	13.8766	0.0028	13.9065	1.3935
1	110	0.1	0.2	0	4.2398	0.0007	5.8384	1.0136
1	110	0.1	0.2	0.5	6.5267	0.0003	6.5727	1.4921
1	110	0.1	0.5	0	10.5148	0.0052	11.6391	0.9017
1	110	0.1	0.5	0.5	15.9268	0.0028	15.7634	1.3593
3	90	0.05	0.2	0	23.0121	0.0017	23.4937	0.95928
3	90	0.05	0.2	0.5	24.8104	0.0008	26.0281	1.38662
3	90	0.05	0.5	0	29.9998	0.0103	29.7121	0.83373
3	90	0.05	0.5	0.5	36.828	0.0057	37.7945	1.14902
3	90	0.1	0.2	0	33.3711	0.0018	33.7350	1.00713
3	90	0.1	0.2	0.5	34.0088	0.0008	34.7613	1.23057
3	90	0.1	0.5	0	37.2287	0.0107	36.9049	0.84556
3	90	0.1	0.5	0.5	42.7673	0.0058	43.6977	1.12109
3	100	0.05	0.2	0	15.678	0.0016	15.6782	0.95928
3	100	0.05	0.2	0.5	18.5798	0.0007	18.5799	1.38662
3	100	0.05	0.5	0	25.161	0.01	25.1607	0.83373
3	100	0.05	0.5	0.5	32.7054	0.0058	32.7049	1.14902
3	100	0.1	0.2	0	26.1671	0.0017	26.1664	1.00713
3	100	0.1	0.2	0.5	27.5462	0.0008	27.5463	1.23057
3	100	0.1	0.5	0	32.1145	0.0104	32.1141	0.84556
3	100	0.1	0.5	0.5	38.5906	0.0057	38.5907	1.12109
3	110	0.05	0.2	0	9.8016	0.0013	10.9059	0.95928
3	110	0.05	0.2	0.5	13.4902	0.0006	13.0938	1.38662
3	110	0.05	0.5	0	21.0184	0.0098	21.9056	0.83373
3	110	0.05	0.5	0.5	29.1034	0.0059	28.8132	1.14902
3	110	0.1	0.2	0	19.4368	0.0016	20.0164	1.00713
3	110	0.1	0.2	0.5	21.7596	0.0008	21.0829	1.23057
3	110	0.1	0.5	0	27.6233	0.0101	28.4547	0.84556
3	110	0.1	0.5	0.5	34.8364	0.0057	34.4549	1.12109

The calibration algorithm described and illustrated above is just one possible model calibration procedure. We recommended always adhering to our convention for choosing the single-asset strikes in stage 1, as this will minimize the average calibration error over all option strikes. However, for users seeking to increase pricing accuracy for OTM (ITM) basket options, more than one VCF could be calibrated at stage 2, or a single VCF could be calibrated to an OTM (ITM) basket option. Of course, our analytic formulae for the Greeks will apply regardless of the calibration procedure chosen.

7. CONCLUSION

Most of the existing approaches to pricing basket options are based on approximating the distribution of the basket price, or they are limited to pricing average price basket options, or they apply only to options on a small number of assets. This paper develops a recursive framework for pricing and hedging European basket options which has no such constraints, and which may also be extended to rainbow options. Our key idea is to write the option pay-off as a sum of pay-offs to CEOs on sub-basket options. By writing the pay-offs to these sub-basket options in terms of pay-offs to CEOs on smaller sub-baskets, and repeating, we can draw a pricing tree that applies to any given basket option. This yields an approximate pricing formula for a general, N -asset basket option, which expresses the basket option price in terms of the prices of $2(N - 1)$ CEOs and $2N$ standard option prices. Rainbow options may also be priced in this framework, as the price of a rainbow option may be expressed in terms of a basket option price and exchange option prices.

Our recursive procedure provides an almost exact price for certain options on baskets containing no more than three assets, because they satisfy what we call the “strong” lognormality condition where exact lognormal option price processes may be applied, under a change of measure. For general N asset options the error stems from our approximation of the price process of the options in the pricing tree as lognormal processes with constant volatility, which is possible under a “weak” lognormality condition. Simulations have tested the accuracy of our approximations for pricing various vanilla options, exchange options, CEOs, spread options, and two-asset basket options. The results show that the approximation errors are very small.

Then, we described how one can apply our option price recursion to compute basket option prices. In general, the model calibration problem takes two stages: defining a pricing tree and approximating the sub-basket option volatilities at each level in the tree. We advocate the use of a tree based on $2(N - 1)$ ATM vanilla options and a pair of vanilla call and put options with the same moneyness as the N -asset option. Then the volatility approximation reduces to calibrating a single parameter to a given (simulated or market) price for an ATM N -asset option. This way, the approximation error is negligible for ATM options. Empirical examples have quantified the error for OTM and ITM five-asset options.

Our recursive approach is quite novel, and has several advantages over those already developed in the literature. First, the underlying asset prices may follow heterogeneous lognormal processes. For instance, some asset prices could follow mean-reverting processes whilst others follow standard lognormal processes. Second, our framework provides a convention for selecting the strikes of vanilla options in the pricing tree, and the asset volatilities may therefore be consistent with the volatility skew. Third, model calibration requires only one parameter to be calibrated to a single simulated ATM

basket option price, which can be done in a few seconds; then simple analytic formulae yield prices of ITM and OTM baskets and all deltas, gammas, and vegas. Most other approaches are much more time consuming, requiring either many more simulations or computations of multidimensional integrals.

REFERENCES

- BANERJEE, P. (2003): Close from Pricing on Plain and Partial Outside Double Barrier Options, *Wilmott* 6, 46–49.
- BEIBER, J. (1999): Another Way to Value Basket Options, Working Paper, Johannes Gutenberg-Universität Mainz.
- CARR, P. (1995): Two Extensions to Barrier Option Valuation, *Appl. Math. Finance* 2(3), 173–209.
- DURRETT, R. (1996): *Stochastic Calculus: A Practical Introduction*. Boca Raton, FL: CRC Press.
- GENTLE, D. (1993): Basket Weaving, *Risk* 6(6), 51–52.
- JOHNSON, H. (1987): Options on the Maximum or the Minimum of Several Assets, *J. Financ. Quant. Anal.* 22(3), 277–283.
- JOHNSON, N. L. (1949): Systems of Frequency Curves Generated by Methods of Translation, *Biometrika* 36(1/2), 149–176.
- JU, N. (2002): Pricing Asian and Basket Options Via Taylor Expansion, *J. Comput. Finance* 5(3), 79–103.
- KARATZAS, I., and S. E. SHREVE (1991): *Brownian Motion and Stochastic Calculus*. New York: Springer-Verlag.
- KARLIN, S., and H. M. TAYLOR (1981): *A Second Course in Stochastic Processes*. New York: Academic Press.
- KREKEL, M., J. DE KOCK, R. KORN, and T. K. MAN (2004): An Analysis of Pricing Methods for Baskets Options, *Wilmott* 2004(3), 82–89.
- KWOK, Y. K., L. WU, and H. YU (1998): Pricing Multi-asset Options with an External Barrier, *Int. J. Theor. Appl. Finance* 1(4), 523–541.
- LEVY, E. (1992): The Valuation of Average Rate Currency Options, *J. Int. Money Finance* 11, 474–491.
- LEWIS, A. L. (2000): *Option Valuation under Stochastic Volatility*. Newport Beach, CA: Finance Press.
- LINDSET, S., and S. A. PERSSON (2006): A Note on a Barrier Exchange Option: The World's Simplest Option Formula? *Financ. Res. Lett.* 3(3), 207–211.
- MARGRABE, W. (1978): The Value of an Option to Exchange One Asset for Another, *J. Finance* 33(1), 177–186.
- MILEVSKY, M. A., and S. E. POSNER (1998a): Asian Options, the Sum of Lognormals, and the Reciprocal Gamma Distribution, *J. Financ. Quant. Anal.* 33(3), 409–422.
- MILEVSKY, M. A., and S. E. POSNER (1998b): A Closed-form Approximation for Valuing Basket Options, *J. Derivatives* 5(4), 54–61.
- OUWEHAND, P., and G. WEST (2006): Pricing Rainbow Options, *Wilmott Mag.* 74–80.
- PILIPOVIC, D. (2007): *Energy Risk: Valuing and Managing Energy Derivatives*. New York: McGraw-Hill.
- ROGERS, L. C. G., and Z. SHI (1995): The Value of an Asian Option, *J. Appl. Probab.* 32(4), 1077–1102.

- SCHWARTZ, EDUARDO S. (1997): The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging, *J. Finance* 52(3), 923–973.
- STULZ, R. M. (1982): Options on the Minimum or the Maximum of Two Risky Assets—Analysis and Applications, *J. Financ. Econ.* 10(2), 161–185.
- TOPPER, J. (2001): *Worst Case Pricing of Rainbow Options*, SSRN eLibrary.
- WEST, G. (2005): Better Approximations to Cumulative Normal Functions, *Wilmott Mag.* 9, 70–76.