A Closed-Form Approximation for Valuing Basket Options

Moshe Arye Milevsky
is assistant professor of finance at the Schulich School of Business at York University in Toronto, Ontario, Canada.

Steven E. Posner
is with ING Baring Securities in New York.

The no-arbitrage valuation of basket options is complicated by the fact that the sum of lognormal random variables is not lognormal. This problem is shared with arithmetic Asian options as well. Various ad hoc approximation techniques have been proposed, none of them very satisfactory or accurate.

In this article we suggest using the reciprocal gamma distribution as an approximation for the state-price density (SPD) function of the underlying basket stochastic variable. This, in turn, allows us to obtain a closed-form expression for the price of a basket option. The technique, when compared against a simple lognormal approximation, performs at its best when the correlation structure of the underlying basket exhibits a specific decaying pattern.

As a by-product, we introduce a formal approach for assessing the goodness of fit of candidate distributions for approximating the SPD. Finally, we present a numerical example in which we apply our formula to value (G-7) index-linked guaranteed investment certificates, which can be decomposed into a zero-coupon bond and a basket option.

Basket options have become extremely popular over the last few years as part of many structured index-linked products offered to institutional and retail investors. A basket option, as its name implies, is an option on a collection, or basket, of assets, typically stocks. A basket call option gives the holder the right, but not the obligation, to purchase a prespecified fixed portfolio of stocks at a fixed strike price.

As a direct consequence of the linear summation, the no-arbitrage valuation of basket options is complicated by the fact that the sum of lognormal distributions is not lognormal. Consequently, there is no known *closed-form* solution to the problem of pricing and hedging these products.

A variety of approximation techniques and ad hoc rules of thumb have arisen to tackle this problem. They include:

1. The lognormal approximations of Hull [1997], Huynh [1994], and Gentle-Vorst [1993].
2. The binomial trees of Hull and White [1993] and Rubinstein [1994].
3. The simulation techniques of Joy, Boyle, and Tang [1996].

Lyden [1996] provides a bibliography of the valuation techniques that have been applied to basket options. A close scrutiny of the literature reveals that, as one would expect in these situations, the more accurate the approximation, the more complicated and computationally expensive its calculation.

This valuation problem is shared with arithmetic Asian options as well. The payoff from an arithmetic Asian option depends on the sum of an underlying stock's prices at prespecified times. An Asian option is essentially a one-stock basket option over time, where the basket of stocks consists of the underlying stock on different days. The difference between an Asian option and a basket option is essentially in the correlation structure of the elements being averaged.

A recent development in arithmetic Asian option pricing has been the introduction of the *reciprocal gamma* distribution as the state-price density function for the sum of contemporaneously correlated lognormal random variables. See Milevsky and Posner [1998] for details. This technique encourages us to try the same approximation for basket options, given the inherent symmetry between the two problems.

We demonstrate how to value basket options by using the reciprocal gamma distribution as the state-price density function for the underlying stochastic variable. This, in turn, allows us to obtain a closed-form expression for the price of a basket option, employing moment matching techniques. The cumulative distribution function (CDF) of the gamma distribution, G(d) in our formula, plays the same role as N(d) in the Black-Scholes formula. Our result should be pedagogically attractive to traders who will understand that we are replacing the normal CDF by an alternative CDF since the basket is not lognormal.

Finally, we present a numerical example in which we apply our formula to value index-linked guaranteed investment certificates (ILGICs) on a collection of international stock indexes, which can be decomposed into a zero-coupon bond and a basket option.

I. RECIPROCAL GAMMA VERSUS LOGNORMAL

If a random variable X is gamma distributed [i.e., $X \sim \text{Gamma}(\alpha, \beta)$], then the probability density function (PDF) of X is

$$g(x,\alpha,\beta) = \frac{e^{-x/\beta}(x/\beta)^{\alpha-1}}{\beta\Gamma(\alpha)}, \quad x \geq 0, \ \alpha,\beta > 0 \quad (1)$$

where $\Gamma(\cdot)$ is the gamma function, which has the property that

$$\Gamma(\alpha + 1) = \alpha\Gamma(\alpha) \quad (2)$$

If a random variable Y is reciprocal gamma distributed [(i.e., $Y \sim \text{R.Gamma}(\alpha, \beta)$)], its inverse, 1/Y, is gamma distributed. We use the notation $G_R(y, \alpha, \beta)$ to denote the CDF of the reciprocal gamma distribution evaluated at y, and $G(x, \alpha, \beta)$ to denote the CDF of the gamma distribution evaluated at x.

By definition:

$$G_R(y, \alpha, \beta) = 1 - G(1/y, \alpha, \beta)$$

It then follows that the probability density functions are related as:

$$g_r(y,\alpha,\beta) = \frac{g(1/y,\alpha,\beta)}{y^2}, \quad y \geq 0, \ \alpha,\beta > 0 \quad (3)$$

One can show using standard calculus techniques that

$$E[Y^i] = \frac{1}{\beta^i(\alpha-1)(\alpha-2)\ldots(\alpha-i)}, \quad i = 1, 2, 3, \ldots, \quad (4)$$

If X is a lognormal random variable [i.e., $X \sim LN(\mu, \sigma)$], however, one can show that

$$E\left[X^i\right] = \exp\left(i\mu + \frac{1}{2}i^2\sigma^2\right), \quad i = 1, 2, 3, \ldots, \quad (5)$$

We will use the notation M_i when referring to the i-th moment of a random variable.

II. THE FORMULA AND JUSTIFICATION

We assume that there are N underlying processes $S_i(t)$ (i = 1, ..., N) that follow the risk-neutral stochastic differential equations:

$$dS_i(t) = (r - q_i)\, S_i(t)\, dt + \sigma_i\, S_i(t) dB_i(t)$$

with a given correlation structure $(\rho)_{i,j}$ between the Brownian motions. Our interest is in the dynamics of a (weighted) arithmetic sum of the underlyings:

$$A(t) = \sum_{i=1}^{N} a_i S_i(t)$$

where a_i are weights.

The Cox-Ross [1976] fundamental theorem of derivative asset pricing states that, in a complete market, the no-arbitrage value of the basket option will be:

$$\text{basket.call} = e^{-rT} E^*[\max(A(T) - K, 0)]$$

The value of the basket option is equal to its expected payoff discounted at the risk-free rate, where E^* denotes the expectation with respect to the state-price density (SPD) function, also known as the risk-neutral probability density function. Consequently:

$$\text{basket.call} = e^{-rT} \int_0^{\infty} \max(A(T) - K, 0) d\Psi[A(T)]$$

where $\Psi[A(T)]$ denotes the SPD. Unfortunately, the SPD of A_T is not known in general, and it appears that closed-form solutions to the problem are not available. A common approach to this problem has been to approximate the SPD $\Psi[A(T)]$ by a lognormal functional form using various moment matching techniques.

We propose moment matching to a reciprocal gamma, as opposed to lognormal, distribution. Milevsky and Posner [1998] demonstrate that, in the context of Asian option pricing, the sum of correlated lognormal random variables is closely approximated by (and converges in the limit to) the reciprocal gamma distribution. We apply the same technique in the basket case, given the intrinsic symmetry between the two problems.

Indeed, Asian options and basket options differ only by the specification of their respective covariance matrices. An Asian option is an arithmetic average of a stock price at N prespecified (equally spaced) times over the life of the option.

The covariance matrix of the logarithms of the stock prices at each of these times is essentially:

$$\frac{\sigma^2}{N}\{\min(i,j)\} = \frac{\sigma^2}{N}\begin{bmatrix} 1 & 1 & \cdots & 1 & 1 \\ 1 & 2 & \cdots & 2 & 2 \\ \vdots & \vdots & \ddots & \cdots & \cdots \\ 1 & 2 & . & N-1 & N-1 \\ 1 & 2 & \cdot & N-1 & N \end{bmatrix} \tag{6}$$

We refer to this as a decaying correlation structure. If the basket of underlying stocks has a decaying correlation structure, it is reasonable to use the Asian option approximation for the basket option as well. Numerical examples will demonstrate that this approximation is more accurate than the lognormal distribution, when compared against results from a Monte Carlo simulation. Of course, when the correlation structure differs significantly from the Asian style payoff, there is no theoretical justification for using the reciprocal gamma distribution compared to any other two-parameter density.

The appendix provides a further discussion of the distance between the reciprocal gamma and lognormal densities, vis-à-vis the sum of lognormals.

Our initial task is to compute the first two moments of the basket's risk-neutral payoff structure at maturity, and then match them to 1) the lognormal distribution, and 2) the reciprocal gamma distribution, to compare the valuations produced by the two methods.

A convenient technique is to define the time T "pseudo-forward" of the basket as

$$F = \sum_{i=1}^{N} a_i F_i$$

where

$$F_i = S_i(0)e^{(r-q_i)T} \tag{7}$$

If we divide A(T) by F and call it $A^*(T)$, it is effectively normalized to have a mean of 1 (i.e., $M_1 = 1$). One can verify that the second moment of $A^*(T)$ is

$$M_2 = \frac{1}{F^2}\sum_i\sum_j a_i a_j S_i(0) S_j(0) e^{(2r-q_i-q_j)T+\rho_{i,j}\sigma_i\sigma_j T}$$

$$= \frac{1}{F^2}\sum_i\sum_j a_i a_j F_i F_j e^{\rho_{i,j}\sigma_i\sigma_j T} \qquad (8)$$

Lognormal Approximation

Assume that the normalized basket $A^*(T)$ is lognormal. Since $M_1 = 1$, we have from Equation (5) that the "variance" of the basket

$$v = \ln(M_2)$$

Then the moment-matched lognormal SPD gives a price for the call, using the usual Black-Scholes formula, as

LN.basket.call =

$$e^{-rT}\left[FN\left(\frac{\ln(F/K)+v/2}{\sqrt{v}}\right) - KN\left(\frac{\ln(F/K)-v/2}{\sqrt{v}}\right)\right]$$

where N(.) is the CDF of a standard normal variate. A similar approach is used by Huynh [1994] and Hull [1997].

Reciprocal Gamma Approximation

We now illustrate our innovation, which is to moment match to a reciprocal gamma distribution instead of a lognormal. To moment match to a reciprocal gamma distribution, we proceed to find the parameters α, β in terms of the first two moments. Since $M_1 = 1$, we have from Equation (4) that

$$\beta = \frac{1}{\alpha - 1} \qquad (9)$$

and

$$M_2 = \frac{1}{\beta^2(\alpha-1)(\alpha-2)}$$

It then follows that:

$$\alpha = \frac{2M_2 - 1}{M_2 - 1}, \quad \beta = 1 - \frac{1}{M_2} \qquad (10)$$

Using the reciprocal gamma distribution with these parameters and using Equation (3), we have that the price of the call is then

$$\text{RG.basket.call} = e^{-rT}\int_0^\infty \max(A(T)-K,0)\,d\Psi(A(T))$$

$$= e^{-rT} F \int_{K/F}^\infty \left(A^*(T) - \frac{K}{F}\right) g_r(A^*(T),\alpha,\beta)\,dA^*(T)$$

$$= e^{-rT}\left[F\int_{K/F}^\infty x g_r(x,\alpha,\beta)dx - K\int_{K/F}^\infty g_r(x,\alpha,\beta)dx\right]$$

$$= e^{-rT}\left[F\int_{K/F}^\infty \frac{g(1/x,\alpha,\beta)}{x}dx - K\int_{K/F}^\infty \frac{g(1/x,\alpha,\beta)}{x^2}dx\right]$$

$$= e^{-rT}\left[F\int_0^{F/K} \frac{g(u,\alpha,\beta)}{u}du - K\int_0^{F/K} g(u,\alpha,\beta)du\right]$$

$$= e^{-rT}\left[F\int_0^{F/K} g(u,\alpha-1,\beta)dx - K\int_0^{F/K} g(u,\alpha,\beta)du\right]$$

The last equality is from the fact that:

$$\frac{g(x,\alpha,\beta)}{x} = \frac{e^{-x/\beta}(x/\beta)^{\alpha-1}}{x\beta\Gamma(\alpha)} = \frac{e^{-x/\beta}(x/\beta)^{\alpha-2}}{\beta^2\Gamma(\alpha)}$$

$$= \frac{e^{-x/\beta}(x/\beta)^{\alpha-2}}{\beta\Gamma(\alpha)/(\alpha-1)}$$

$$= \frac{e^{-x/\beta}(x/\beta)^{\alpha-2}}{\beta\Gamma(\alpha-1)} = g(x,\alpha-1,\beta)$$

[using Equation (1), Equation (2), and Equation (9)].

Finally, we have that

RGbasket.call =

$$e^{-rT}[FG(F/K,\alpha-1,\ \beta) - KG(F/K,\ \alpha,\ \beta)] \qquad (11)$$

III. NUMERICAL EXAMPLE: INDEX-LINKED GICs OR EIA

A numerical example will compare the lognormal and reciprocal gamma techniques for the case of seven stocks in a basket option. The technique performs better than the lognormal when calibrated against a Monte Carlo simulation.

Index-linked guaranteed investment certificates (ILGICs) are a good example of basket options offered to the retail public. With the general decline in the level of interest rates, ILGICs have become extremely popular. Virtually every one of the major Canadian banks offers a variant on this product. An ILGIC (in the U.S. an equity enhanced certificate of deposit or equity indexed annuity) is a retail savings vehicle, similar to a term deposit, with an interest rate that is linked to a collection of diversified stock indexes. Upon maturity of the product, the total return will be determined on the basis of the performance of the underlying indexes over the prespecified term.

An ILGIC is engineered by fusing a zero-coupon bond with a basket option that is struck at the spot rate of the underlying indexes. See Baubonis, Gastineau, and Purcell [1993] for more of the institutional details.

We are interested in pricing the embedded basket option on the product linked to the performance of the G-7 stock markets. (The zero-coupon bond can simply be priced off the term structure.)

Consider a call option that pays off the max of a weighted average of renormalized G-7 stock indexes less 1 (which is the return of principal) and zero. Define

EXHIBIT 1
G-7 ILGIC WEIGHTINGS

COUNTRY	INDEX	WEIGHT (%)	ANN. VOL. (%)	DIV. YIELD (%)
Canada	TSE 100	10	11.55	1.69
France	CAC 40	15	20.68	2.39
Germany	DAX	15	14.53	1.36
Italy	MIB30	5	17.99	1.92
Japan	Nikkei 225	20	15.59	0.81
U.K.	FTSE 100	10	14.62	3.62
U.S.	S&P 500	25	15.68	1.66

$$A(T) = \sum_{i=1}^{7} a_i \frac{S_i(T)}{S_i(0)}$$

where a_i are weights that sum to 1. Our interest is in pricing call options on A(T) struck at 1:

$$\text{basket.call} = e^{-rT}E^*[\max(A[T] - 1, 0)]$$

This is, effectively, a call option on the rate of return of a basket of indexes.

The popular G-7 ILGIC offered by Canada Trust Co. uses the weightings in Exhibit 1 in the seven indexes. The annualized (historical) volatility, dividend yield, and correlation matrix, from a Canadian dollar perspective, were provided by the J.P. Morgan Risk Metrics system on July 17, 1997.

The (log return) correlation structure used is in Exhibit 2. We use a flat (domestic) risk-free rate, r, of

EXHIBIT 2
CORRELATIONS

	CANADA	FRANCE	GERMANY	ITALY	JAPAN	U.K.	U.S.
Canada	1.0000	0.3500	0.1000	0.2700	0.0400	0.1700	0.7100
France	0.3500	1.0000	0.3900	0.2700	0.5000	–0.0800	0.1500
Germany	0.1000	0.3900	1.0000	0.5300	0.7000	–0.2300	0.0900
Italy	0.2700	0.2700	0.5300	1.0000	0.4600	–0.2200	0.3200
Japan	0.0400	0.5000	0.7000	0.4600	1.0000	–0.2900	0.1300
U.K.	0.1700	–0.0800	–0.2300	–0.2200	–0.2900	1.0000	–0.0300
U.S.	0.7100	0.1500	0.0900	0.3200	0.1300	–0.0300	1.0000

EXHIBIT 3
RESULTS

TENOR (YEARS)	MONTE CARLO	STD. ERROR	RECIPROCAL GAMMA	LOGNORMAL
1	0.0587	0.0002	0.0589	0.0591
3	0.1331	0.0004	0.1328	0.1338
5	0.1945	0.0005	0.1942	0.1957
10	0.3104	0.0007	0.3103	0.3119

6.3% (continuously compounded).

Call option values at different tenors are evaluated using these parameters. In reality, the risk-free rate, volatilities, and even correlation structure should vary with tenor. For simplicity and for ease of reproducibility, however, we keep the parameters static.

For example, when pricing a three-year basket option, we plug the weights and the statistical parameters into Equation (7) to obtain that F = 1.1460 and therefore M_2 = 1.0225 [from Equation (8)]. This, in turn, leads to $\alpha = 46.3648$ and $\beta = 0.0220$ from Equation (10). Finally, we plug α, β, K, F, T, and r into Equation (11) to obtain a price of 0.1328 for the basket option using the reciprocal gamma distribution.

The general results are displayed in Exhibit 3. The reciprocal gamma and lognormal approximation results are provided for different tenors (one, three, five, and ten years) along with a Monte Carlo simulation (of length 100,000).

Exhibit 4 displays the reciprocal gamma (RG) approximation and the lognormal (LN) distribution. The Monte Carlo state-price density is our proxy for the true state-price density of the weighted average. The curve seems to indicate that the RG is closer to the true density than the LN. In particular, the right-hand tail of the LN contains a larger mass than the RG, which is consistent with the overpricing of out-of-the-money call options.

IV. CONCLUSION

The lognormal distribution and its properties are very familiar to financial engineers as a result of the Black-Scholes formula and the financial economic justification for geometric Brownian motion. It is no

EXHIBIT 4
TWO PDFS WITH IDENTICAL FIRST TWO MOMENTS

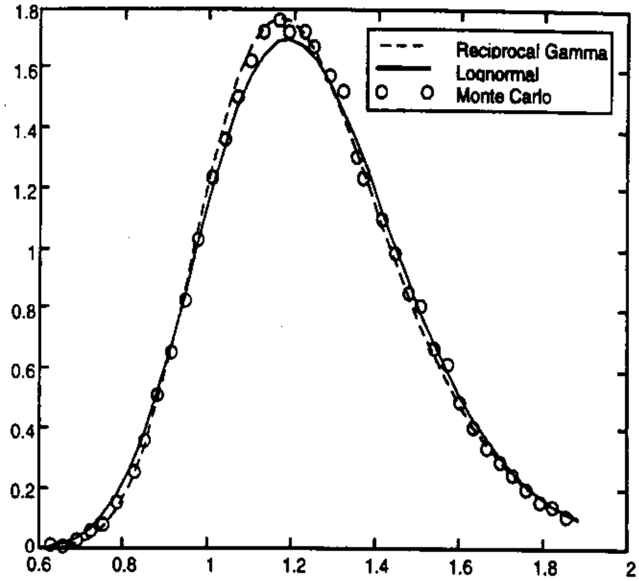

surprise therefore that academics and practitioners alike have attempted to force the lognormal distribution onto situations for which there is little theoretical justification. This is quite evident with basket options, where the finite sum of lognormal variates is conveniently approximated by a lognormal distribution.

Under certain circumstances, basket options can be valued using the reciprocal gamma (instead of the lognormal) distribution as the state-price density function. Our justification comes from a limiting result, originating in Asian option pricing theory, which states that the sum of correlated lognormals converges to the reciprocal gamma distribution in the limit.

Consequently, our basket option pricing formula produces results that are better than the lognormal approximation and are very accurate when compared with Monte Carlo simulations.

A final caveat is in order. After extensive testing, it appears that, compared to the Monte Carlo price, the traditional lognormal method overprices out-of-the-money (relative to the forward) call options, and the reciprocal gamma method slightly underprices them. Perhaps a prudent implementation of this technique would involve averaging the RG and the LN prices to obtain a more reliable estimate when dealing with out-of-the-money basket call options, especially when the correlation structure is non-decaying.

APPENDIX

We introduce a formal (original) approach for assessing the goodness of fit of candidate distributions for approximating the SPD. In particular, we demonstrate that the reciprocal gamma (RG) distribution is closer to the sum of lognormals than the lognormal (LN) distribution, when the correlation structure satisfies the *decaying* criteria. The method we use is motivated by the Kolmogorov-Smirnov notion of distance between random variables.

Specifically, let F(x) denote the cumulative density function (CDF) of the true sum of lognormals, and let H(x) denote the CDF of some known arbitrary random variable, for example, reciprocal gamma or lognormal. Define the Kolmogorov-Smirnov metric:

$$d(F, H) = \max_x \left| F(x) - H(x) \right| \qquad \text{(A-1)}$$

which can be thought of as a measure of how close, or far, the two distributions, F(x) and H(x), are from each other. See Breiman [1993, Chapter 13] for details and justification of this metric as a measure of distance.

We use the notation:

$$d(F, G_R) = \max_x \left| F(x) - G_R(x) \right| \qquad \text{(A-2)}$$

and

$$d(F, L_N) = \max_x \left| F(x) - L_N(x) \right| \qquad \text{(A-3)}$$

for the distance between the true distribution of the sum of lognormals and the two approximating functions: the LN and the RG.

This metric has various applications. In particular, there is a result that states that if a random sample of size m is drawn from one of two distributions, F or H, and if the sample size m, is less than $0.71/d(F, H)^2$, detection at the 95% confidence level is impossible. Likewise, by taking square roots, if $d(F, H) < 0.84/\sqrt{m}$, we cannot detect the difference between the two distributions at a confidence level of 95%.

In our particular case, we make use of this result. If a sample of size m is drawn from a continuous random variable with CDF denoted by F(x), the sample CDF denoted by $F_m(X)$ obeys the inequality:

$$\Pr(d(F, F_m) < \varepsilon_m) = 0.95 \qquad \text{(A-4)}$$

where $\varepsilon_m = \dfrac{1.36}{\sqrt{m}}$.

Intuitively, as the sample size m gets large, the sample CDF gets closer to the true CDF. The practical implication of this statement is that we can simulate samples for the underlying basket of securities, compute the mean, and then collect the terminal values to create a histogram that is compared against the two analytically available alternatives, the LN distribution and the RG distribution.

We now show that if

$$d(G_R, F_m) < d(L_N, F_m) - 2\varepsilon_m \qquad \text{(A-5)}$$

then

$$d(G_R, F) < d(L_N, F) \qquad \text{(A-6)}$$

with 95% confidence. In other words, if the empirical CDF is closer to the RG, by at least an amount $2\varepsilon_m$, then the true CDF is closer to RG than to LN.

The argument is geometric. If the *simulated* CDF of the sum of lognormals is within a prespecified distance of the candidates for the *approximating* CDF of the sum of lognormals, it follows that the *true* CDF is within a prespecified distance as well. More precisely, Equations (A-2) and (A-3) define an uncertainty ball of size ε_m about F.

Accordingly we have the bounds on the true distances:

$$\max_{H \supset d(H,F_m) < \varepsilon_m} d(G_R, H) < d(G_R, F_m) + \varepsilon_m \qquad \text{(A-7)}$$

and

$$\max_{H \supset d(H,F_m) < \varepsilon_m} d(L_N, H) > d(L_N, F_m) - \varepsilon_m \qquad \text{(A-8)}$$

Since these equations hold for the max and the min, respectively, they hold for F as well.

Consequently, if we have that:

$$d(G_R, F_m) < d(L_N, F_m) - 2\varepsilon_m \qquad \text{(A-9)}$$

it follows that:

$$d(G_R, F) < d(G_R, F_m) + \varepsilon_m < d(L_N, F_m) - 2\varepsilon_m + \varepsilon_m =$$

$$d(L_N, F_m) - \varepsilon_m < d(L_N, F) \qquad \text{(A-10)}$$

EXHIBIT 5
SIMULATION RESULTS

$d(G_R, F_m)$	$d(L_N, F_m)$	DIFFERENCE	$2\varepsilon_m$
0.0141	0.0279	0.0138	0.0136

EXHIBIT 6
ATM BASKET OPTION PRICES

TENOR (YEARS)	MONTE CARLO	STD. ERROR	RECIPROCAL GAMMA	LOGNORMAL
1	0.0731	0.0001	0.0730	0.0739
3	0.1111	0.0003	0.1102	0.1140
5	0.1269	0.0005	0.1238	0.1310

To verify that this indeed holds for the sum of lognormals when compared against the LN and RG, we simulate sample paths for the basket price process, where the covariance is structured as

$$\frac{\sigma^2}{N}\{\min(i,j)\} = \frac{\sigma^2}{N}\begin{bmatrix} 1 & 1 & \dots & 1 & 1 \\ 1 & 2 & \dots & 2 & 2 \\ \vdots & \vdots & \ddots & \dots & \dots \\ 1 & 2 & . & N-1 & N-1 \\ 1 & 2 & . & N-1 & N \end{bmatrix}$$

and then construct the sample CDF, $F_m(X)$, using m = 40,000. The results for T = 1 are in Exhibit 5. They indicate that the argument indeed holds: The reciprocal gamma is closer to the true density with probability 95%.

This argument works provided that $d(G_R, F_m) < d(L_N, F_m) - 2\varepsilon_m$ holds in the simulation result. Otherwise we cannot state that $d(G_R, F) < d(L_N, F)$. Of course, to prove this result *always* holds is an open problem.

Pursuant to this example, consider a basket of ten equal-weighted (risk-neutral) forward assets, i.e., r = q = 6%. The covariance structure of the assets is "decaying" with $\sigma_{i,j} = (\sigma^2/N)\{\min(i, j)\}$, where σ = 10%. This means that the actual volatility of asset i is $10\sqrt{i}\%$. Exhibit 6 provides prices for different maturities for ATM call options.

Of course, there is nothing unique about the lognormal density function in our comparison. Indeed, the same experiment can be applied to other candidate distributions.

REFERENCES

Baubonis, Charles, Gary Gastineau, and David Purcell. "The Banker's Guide to Equity-Linked Certificates of Deposit." *Journal of Derivatives*, Vol. 1, 2 (1993), pp. 87-95.

Breiman, Leo. *Probability*. Philadelphia: Society for Industrial and Applied Mathematics, 1993.

Cox, John, and Stephen Ross. "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics*, Vol. 3 (1976), pp. 145-166.

Gentle, David. "Basket Weaving." *RISK*, June 1993, pp. 51-52.

Grannis, Scott. "An Idea Whose Time Has Come." *RISK*, September 1989.

Hull, John C. *Options Futures and Other Derivatives*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall, 1997.

Hull, John C., and Alan White. "Efficient Procedures for Valuing European and American Path-Dependent Options." *Journal of Derivatives*, Vol. 1 (1993), pp. 21-31.

Huynh, Chi Ban. "Back to Baskets." *RISK*, May 1994, pp. 59-61.

Joy, Corwin, Phelim P. Boyle, and Ken Seng Tang. "Quasi-Monte Carlo Methods in Numerical Finance." *Management Science*, Vol. 42 (1996), pp. 926-938.

Lyden, Scott. "Reference Check: A Bibliography of Exotic Options Models." *Journal of Derivatives*, Vol. 4, 1 (1996), pp. 79-91.

Milevsky, Moshe A., and Steven E. Posner. "Asian Options, the Sum of Lognormals and the Reciprocal Gamma Distribution." *Journal of Financial and Quantitative Analysis*, September 1998.

Rubinstein, Mark. "Return to Oz." *RISK*, November 1994.