Unifying volatility models

This article introduces a method for building analytically tractable option pricing models that combine state-dependent volatility, stochastic volatility and jumps. The eigenfunction expansion method is used to add jumps and stochastic volatility to hypergeometric Brownian motions. Claudio Albanese and Alexey Kuznetsov conclude that such comprehensive unified models are not only able to reflect the complexities of exotic option prices, but are also analytically tractable

The success of pricing models is often measured by the extent to which closed-form solutions of the Black-Scholes type are available for the basic payouts. Analytic tractability is often crucial for the calibration to market data and is helpful for implementing numerical algorithms for exotics. Extensions of the Black-Scholes formula have been directed towards three main model classes: local volatility models, which postulate a deterministic relationship between the underlying state variable, time and volatility; stochastic volatility models, which assume that the volatility follows a distinct but correlated process; and jump models, which can be regarded as limits of stochastic volatility models whereby the volatility can occasionally be singularly large at some points in time, enough to cause the underlying sample path to have discontinuous jumps.

As Dupire (1994) demonstrated, state-dependent volatility models are able to reproduce arbitrage-free implied volatility surfaces. However, robust estimations require either regularisations or settling on a parametric form for the local volatility such as the constant elasticity of variance model in Cox & Ross (1976), the quadratic volatility models in Rady (1997) and the more comprehensive hypergeometric Brownian motions in Albanese *et al* (2001). A widely adopted jump model is Madan's variance-gamma model (see Madan, Carr & Chang, 1998, and references therein). Stochastic volatility models such as Heston's (1993) have also found widespread industry applications.

The three volatility models are all able to reproduce most of the features of market-implied volatility skews for European-style options at a fixed maturity, although sometimes the implied process departs too radically from the one estimated historically, does not reflect properly the price of exotic options and/or does not produce correct hedge ratios. An important feature of price processes is the so-called leverage effect first considered by Black (1976), according to which price levels are negatively correlated with spot volatility. Models that incorporate this effect produce hedge ratios that reflect the ability to vega-hedge using the stock. From the modelling perspective, the leverage effect can be captured either by local volatility models or by stochastic volatility models in which equity returns are negatively correlated to volatility changes. Local volatility models have the advantage of being analytically much simpler as the volatility is tied directly to the value of the underlying instead of being driven by an independent process. However, pure state-dependent volatility models run into dynamic inconsistencies, as the implied volatility of call options of the same strike tend to stay constant as the underlying price changes, a phenomenon sometimes called sticky-strike dynamics. Empirical evidence shows that this happens only occasionally, while in normal market conditions, options of the same delta (not the same strike) tend to retain a roughly constant implied volatility as time evolves (Derman, 1999). Market prices of forward starting calls can be explained interchangeably by either a stochastic volatility or a jump component or a mixture of the two effects. Prices of other exotics, such as barrier and American-style options, are instead quite sensitive to the degree of mixture between the jump and the stochastic volatility components. Furthermore, pure stochastic volatility models encounter difficulties capturing the very short-term behaviour of the implied volatility skew, unless one assumes that the instantaneous volatility can become unusually large. Jumps instead readily justify short-term skews. Pure jump models also run into difficulties as the implied volatility surface flattens out too rapidly unless jump amplitudes quickly increase as a function of time. As a rule, the better hedge ratios and the most consistent pricing schemes across exotic options result from the more inclusive models, as long as they are properly calibrated.

The practice of using different pricing models for different classes of exotics is driven by the desire to capture the most relevant effects while keeping models as simple as possible. However, this practice gives rise to pricing inconsistencies across asset classes and exposes users to the potential of subtle arbitrage opportunities. In addition, price discrepancies often translate into even larger discrepancies between hedge ratios. The desire to bridge the gaps between the model classes in use and to construct a more consistent pricing framework is driving much of the current research in option pricing theory. This complex technical challenge includes three steps: (i) add flexibility to the option pricing formalism in such a way to make calibrations and statistical estimations possible; (ii) construct lattice models to add further flexibility and price exotics; and (iii) carry out systematic empirical studies to find the right mixture of volatility processes that best fits data across the various asset classes.

In this article, we address (i) by constructing solvable models in a continuous pricing variable that combines the three volatility models. Within these models, one can price European-style and some barrier options (with barriers struck at constant forward). Also, these models can be discretised in a natural way while preserving solvability, thus yielding efficient lattice models for a broad range of exotics. For this, we refer the interested reader to our more technical paper (Albanese & Kuznetsov, 2003), where we show how a slight modification of the derivation in this paper yields models in a discrete price variable for which one can calculate transition probabilities in analytically closed form, thus giving lattices where the distance between the time nodes can be adjusted at will without altering final results. Similar lattice models also apply to the case with stochastic volatility, which we study in a forthcoming paper on interest rate models. Regarding (iii), our group worked out applications to credit derivatives (Albanese *et al*, 2003) and is about to release two more papers on convertible bonds and on interest rate derivatives with stochastic volatility. The purpose of this article is to outline in a stylised framework what the main technical ideas are by following the simplest, most direct route we know, thus providing a pedagogical introduction to our other more application-specific articles and perhaps motivating the technical reader to adapt these general ideas to her specific domain of interest.

The recent derivatives literature on unified models is mostly aimed at combining jumps with stochastic volatility. For instance, Carr *et al* (2001) discuss time-changed Levy processes and model the leverage effect by correlating jumps with the stochastic volatility factor. In this article, we pursue a different direction and combine the three volatility models together. Although we assume that the stochastic volatility driver is independent, the presence of a local volatility component accounts for the leverage effect. Our method is also related to the recent work by Carr & Wu (2004) on Fourier methods in option pricing theory. When using ordinary Fourier series, one expands payouts and pricing kernels as integrals over a continuous basis of exponential functions with imaginary arguments. The Fourier basis is suitable for processes that are homogeneous in the price of the underlying, the most far-reaching example being represented by models combining jumps with correlated stochastic volatility. The pricing kernels of the state-dependent volatility models in this article also admit a generalised form of Fourier representation. Our approach differs from the one of Carr & Wu (2004) in that our basis is given by a discrete family of hypergeometric functions and the integrals in the Fourier representation are replaced by infinite series. This gives rise to distinct advantages from the numerical viewpoint, as we find that the series can be evaluated by truncating them to manageable finite sums over a few hundred terms.

The article is organised as follows. In the next section, we give a new and simplified derivation of the hypergeometric Brownian models in Albanese *et al* (2001) and recast the pricing kernel and pricing formulas for European-style calls as series expansions. In the following section, we discuss time changes associated with jumps or stochastic volatility and the case of barrier options. We then summarise the results by discussing an implementation example. We then conclude.

Hypergeometric Brownian motions

Here, we give a new simplified derivation of the main results about pricing models built upon the hypergeometric Brownian motions in Albanese *et al* (2001). That article contains a general procedure to construct a driftless process starting from a generic diffusion process X_t with drift $m(x)$ and volatility $\sigma(x)$, that is, satisfying the equation:

$$dX_t = m(X_t)dt + \sigma(X_t)dW \tag{1}$$

To complete the formal definition of the process X_t, we add that X_t is bound to take values on a (possibly infinite) interval $D = (a, b)$ where $\infty \le a < b \le \infty$ and the process terminates upon hitting the boundary.

Solutions to the Black-Scholes equation with a generic payout function contingent on the underlying X_t can be expressed in terms of the so-called pricing kernel $p^X(t, x_0, x_1)$ solving the following equation for $t > 0$:

$$\frac{d}{dt}p^X(t,x_0,x_1) = \mathcal{L}^X p^X(t,x_0,x_1) \quad \text{and} \quad \lim_{t\to 0} p^X(t,x_0,x_1) = \delta(x_0 - x_1)$$

Here, the operator $\mathcal{L}^X$ is the so-called Markov generator for the process X_t, and is given by:

$$\mathcal{L}^X = m(x)\frac{d}{dx} + \frac{\sigma(x)^2}{2}\frac{d^2}{dx^2} \tag{2}$$

and acts on the variable x_0 in the first argument. From the probabilistic point of view, the pricing kernel $p^X(t, x_0, x_1)$ as a function of the end-point x_1 is the conditional probability density of the variable X_t given the initial condition $X_0 = x_0$.

A family of drift-less processes can be constructed based on the following ingredients: a parameter $\rho > 0$; two linearly independent solutions $f_1(x)$, $f_2(x)$ to the differential equation $\mathcal{L}^X f = \rho f$; and two constants c_1, c_2 such that the function $g(x) = c_1 f_1(x) + c_2 f_2(x)$ is strictly positive, that is, $g(x) > 0$ for all $x \in D$.

Since the process $e^{-\rho t}g(X_t)$ is positive and drift-less, it can be interpreted as a new numeraire and used to define a measure change. Let Q denote the measure such that:

$$\frac{dQ_t}{dP_t} = e^{-\rho t} g(X_t) \tag{3}$$

where P is the measure corresponding to equation (1).

Consider the function:

$$Y(x) = \frac{c_3 f_1(x) + c_4 f_2(x)}{c_1 f_1(x) + c_2 f_2(x)} \tag{4}$$

Since the transformed process $e^{-\rho t}(c_3 f_1(X_t) + c_4 f_2(X_t))$ is also drift-less under P, the process $Y(X_t)$ is drift-less under Q. It turns out that the function $Y(x)$ is invertible for all choices of the constants c_3, c_4 such that $c_1c_4 - c_2c_3 \neq 0$. In fact, the derivative $Y'(x)$ can be expressed through the Wronskian between the independent solutions f_1 and f_2, namely:

$$Y'(x) = (c_1c_4 - c_2c_3)\frac{W_{f_1,f_2}(x)}{g^2(x)} \quad \text{where} \quad W_{f_1,f_2}(x) = \det\begin{vmatrix} f_1(x) & f_2(x) \\ f_1'(x) & f_2'(x) \end{vmatrix}$$

Notice that:

$$W_{f_1,f_2}(x) = W(x_0)\exp\left(-2\int_{x_0}^{x}\frac{m(u)}{\sigma(u)^2}du\right)$$

Hence, $Y'(x)$ is non-zero and $Y(x)$ is invertible. Let $X(y)$ denote the inverse of the function $Y(x)$.

The process $Y_t = Y(X_t)$ is drift-less and satisfies the stochastic differential equation $dY_t = \nu(Y_t)dW_t^Q$ with volatility function:

$$\nu(y) = Y'(X(y))\sigma(X(y)) \tag{5}$$

The probability density function $p^Y(t, y_0, y_1)$ for the process $Y_t = Y(X_t)$ can be expressed in terms of the corresponding probability density function $p^X(t, x_0, x_1)$ for the process X_t as follows:

$$\begin{aligned} p^Y(t,y_0,y_1) &= E_t^Q\left[\delta\left(Y(X_t) - Y(x_1)\right) \middle| Y(X_0) = Y(x_0)\right] \\ &= E_t^P\left[\frac{g(X_t)}{g(X_0)}\frac{e^{-\rho t}}{Y'(x_1)}\delta(X_t - x_1)\middle| X_0 = x_0\right] \\ &= \frac{g(x_1)}{g(x_0)}\frac{1}{Y'(x_1)}e^{-\rho t}p^X(t,x_0,x_1) \end{aligned}$$

Here and throughout the article, we set $y_0 = Y(x_0)$, $y_1 = Y(x_1)$ and:

$$H(x_0,x_1) = \frac{g(x_1)}{g(x_0)}\frac{1}{Y'(x_1)} \tag{6}$$

Next, we specialise to the case whereby the Markov generator $\mathcal{L}^X$ has discrete spectrum. This happens when the eigenvalue equation:

$$\mathcal{L}^X \psi_n(x) = \lambda_n \psi_n(x) \tag{7}$$

admits a complete set of eigenvalues λ_n and eigenfunctions $\psi_n(x)$, $n = 0, 1, \ldots$, which are orthogonal with respect to some measure $\mu(x)dx$, that is:

$$\int \psi_n(x)\psi_m(x)\mu(x)dx = \delta_{mn} \tag{8}$$

In the implementation section below, we consider the generator of the Cox-Ingersoll-Ross (CIR) process as an example. If the eigenfunctions are known, the pricing kernel can be represented as the following series:

$$p^X(t, x_0, x_1) = e^{t\mathcal{L}^X}(x_0, x_1) = \sum_{n=0}^{\infty} e^{\lambda_n t}\psi_n(x_0)\psi_n(x_1)\mu(x_1) \tag{9}$$

The pricing kernel for the martingale process Y_t is then given by:

$$p^Y(t, y_0, y_1) = H(x_0, x_1)\mu(x_1)\left[\sum_{n=0}^{\infty} e^{-(\rho-\lambda_n)t}\psi_n(x_0)\psi_n(x_1)\right]$$

Notice that the time and space variables are separated in each term of the series expansion, a useful property that – as we discuss below – allows one to readily include jumps and stochastic volatility into the process.

An interesting expansion also applies for European-style call option prices defined as follows:

$$C^Y(t, y_0, K) = E^Q\left((Y_t - K)^+ | Y_0 = y_0\right) \tag{10}$$

where $(y)^+ = \max(y, 0)$. If $k = X(K)$ is the strike price in the X-variable, then:

$$\begin{aligned} C^Y(t, y_0, K) &= (y_0 - K)^+ \\ &+ \frac{1}{2}\nu^2(K)H(x_0, k)\left[G(\rho, x_0, k) + \mu(k)\sum_{n=0}^{\infty}\frac{1}{\rho-\lambda_n}e^{-(\rho-\lambda_n)t}\psi_n(x_0)\psi_n(k)\right] \end{aligned} \tag{11}$$

where $G(\rho, x_0, x_1)$ is the Green's function defined either as:

$$\int_0^{\infty} e^{-\rho s} p^X(s, x_0, x_1)ds$$

or as the integral kernel of the operator $(\rho - \mathcal{L}^X)^{-1}$. This formula is useful in practice as it gives a computationally efficient method for evaluating European-style call option prices and thus calibrating the model to market data. A derivation of this formula is given in Albanese & Kuznetsov (2003).

Jumps and stochastic volatility

Here, we show how to build jumps and stochastic volatility features into the local martingale Y_t constructed in the previous section. The strategy is to make use of stochastic time changes and leverage on the fact that variables are separated in formulas (9) and (11).

A stochastic time change process T_t is defined as a right-continuous non-decreasing process starting from zero and with values in $[0, \infty)$. The time changed version of the process Y_t is the process $\tilde{Y}_t = Y_{T_t}$. We shall refer to t as the calendar time and call $s = T_t$ the economic time co-ordinate. As discussed in the introduction, we make the simplifying assumption that the time change process T_t is independent of the underlying process Y_t. In this case, if $\rho_t(ds)$ is the distribution of T_t, we have that:

$$E\left[f(\tilde{Y}_t)\right] = \int E\left[f(Y_s)\right]\rho_t(ds) \tag{12}$$

for all continuous functions $f(x)$. A stochastic time change is also characterised by the Laplace transform:

$$L(t, \lambda) = E\left[e^{-\lambda T_t}\right] \tag{13}$$

Formula (12) allows one to express as follows the probability kernel for the time-changed process $\tilde{Y}_t$:

$$p^{\tilde{Y}}(t, y_0, y_1) = \int p^Y(s, y_0, y_1)\rho_t(ds)$$

In conjunction with equation (10) and given that in (9) the time variable and space variables are separated, we find that:

$$p^{\tilde{Y}}(t, y_0, y_1) = H(x_0, x_1)\mu(x_1)\left[\sum_{n=0}^{\infty} L(t, \rho-\lambda_n)\psi_n(x_0)\psi_n(x_1)\right] \tag{14}$$

Similarly, European-style option prices under the time-changed process are calculated as follows:

$$\begin{aligned} C^{\tilde{Y}}(t, y_0, K) &= (y_0 - K)^+ \\ &+ \frac{1}{2}\nu^2(K)H(x_0, k)\left[G(\rho, x_0, k) + \mu(k)\sum_{n=0}^{\infty}\frac{L(t, \rho-\lambda_n)}{\rho-\lambda_n}\psi_n(x_0)\psi_n(k)\right] \end{aligned} \tag{15}$$

A rich variety of time changes can be constructed by combining time changes of three different types. The first and most elementary one is given by a deterministic time change process of the form $T_t^d = f(t)$, where $f(t)$ is an increasing right continuous function from $[0, +\infty)$ to $[0, +\infty)$. In this case the Laplace transform is simply:

$$L(t, \lambda) = e^{-\lambda f(t)} \tag{16}$$

The second building block is a non-decreasing jump process T_t^j. If the increments $T_{t_1}^j - T_{t_0}^j$ are independent and their distribution is time-homogeneous, the stochastic time change process is a non-decreasing Levy process called Bochner subordinator. It turns out that Bochner subordinators are characterised by a so called Bernstein function $\phi(\lambda)$, such that:

$$L(t, \lambda) = e^{-t\phi(\lambda)} \tag{17}$$

Perhaps the most widely used example of Bochner subordinator encountered in pricing theory is given by the gamma process Γ_t of density:

$$\rho_t^{\Gamma}(ds) = \frac{\left(\frac{\mu}{\nu}\right)^{\frac{\mu^2 t}{\nu}}}{\Gamma\left(\frac{\mu^2 t}{\nu}\right)} s^{\frac{\mu^2 t}{\nu}-1} e^{-\frac{\mu}{\nu}s} ds \tag{18}$$

In this case, the Laplace transform is given by:

$$L^{\Gamma}(t, \lambda) = \left(1 + \frac{\lambda\mu}{\nu}\right)^{-\frac{\mu^2}{\nu}t} \tag{19}$$

and the Bernstein function is:

$$\phi(\lambda) = \frac{\mu^2}{\nu}\ln\left(1 + \frac{\lambda\mu}{\nu}\right)$$

The third example of stochastic time change we build upon is given by a process T_t^c with continuous paths that can be represented as follows:

$$T_t^c = \int_0^t u_s ds$$

where u_s is a positive stochastic process. This is the kind of stochastic time change that arises in stochastic volatility models such as Heston's. The integrand u_s can be interpreted as the process for local variance. In Heston's model, the variable u_s follows the CIR process. This choice is advantageous as in this case the Laplace transform of T_t^c can be evaluated in analytically closed form and is structurally similar to the pricing function of zero-coupon bonds in the CIR model. More precisely, if u_t is a CIR process with $du_t = (c - du_t)dt + \zeta dW$, then the Laplace transform has the form:

$$L(t, \lambda) = e^{m(t)u_0 + n(t)} \tag{20}$$

where:

$$\begin{aligned} n(t) &= \frac{2c}{\zeta_0^2}\log\left(\frac{\gamma\exp\left(\frac{1}{2}dt\right)}{\gamma\cosh(\gamma t) + \frac{1}{2}d\sinh(\gamma t)}\right) \\ m(t) &= -\frac{\lambda\sinh(\gamma t)}{\gamma\cosh(\gamma t) + \frac{1}{2}d\sinh(\gamma t)} \end{aligned} \tag{21}$$

where:

$$\gamma = \frac{1}{2}\sqrt{d^2 - 2\lambda\zeta^2}$$

Notice that in this case the Laplace transform T_t^c depends on the initial value of u_0, which is intuitively clear since:

$$T_t^c = \int_0^t u_s ds$$

clearly depends on u_0. To be completely precise in our notations, because of this dependency of the Laplace transform on u_0, we may have defined the Laplace transform with a time subscript as $L_{t_0}(t, \lambda) = E[e^{-\lambda T_t}|\mathcal{F}_{t_0}]$. We opted otherwise for the sake of having less encumbered notations.

Our strategy to build a unified volatility model is to combine time changes of the three types above into one, that is, we consider the process:

$$T_t = T_t^d + T_t^j + T_t^c$$

where T_t^d is a deterministic function of time, T_t^j is a Bochner subordinator and the third process has the form $T_t^c = \int_0^t u_s ds$ where u_s follows a CIR process. We also assume these processes are independent, so that the Laplace transform can be calculated in analytically closed form and is given by the product:

$$L^T(t,\lambda) = L^{T^d}(t,\lambda) L^{T^j}(t,\lambda) L^{T^c}(t,\lambda) \tag{22}$$

Barrier options and exotics are somewhat more difficult to price in the presence of jumps and stochastic volatility. In this case, it is possible to resort to the lattice approximations in Albanese & Kuznetsov (2003), which are more flexible and cover the cases of barriers of arbitrary shape, American-style options, one-touch options, forward starting calls, compound options, etc. The analytical framework we presented is, however, still quite powerful in cases such as that of barrier options struck at constant forward. As a working example, consider a double barrier option, with an up-and-out barrier at the forward value $Y = U$ and a down-and-out barrier at $Y = L$. Stochastic volatility is fairly easy to incorporate by directly subordinating pricing formulas. This is possible because the time change process in this case has continuous paths.

In case there are jumps instead, the Markov generator in the X-representation for the underlying diffusion process with jumps and absorption outside the interval $B = [X(L), X(U)]$ has the integral kernel:

$$\mathcal{L}^{XBJ}(x_0, x_1) = 1_B(x_0) 1_B(x_1) \sum_{n=0}^{\infty} \lambda_n^J \psi_n(x_0) \psi_n(x_1) \mu(x_1) \tag{23}$$

Here, $1_B(F)$ is the indicator function for B, $\lambda_n^J = -\phi(-\lambda_n)$ and ϕ is the Bernstein function that characterises the jump process. To find the pricing kernel, we need to exponentiate the operator $\mathcal{L}^{XBJ}$. To do that, one has to diagonalise it. A possible strategy is to carry out the diagonalisation in the basis $\psi_n^B(x)$ of the eigenfunctions of the operator $\mathcal{L}^X$ with Dirichlet boundary conditions $\psi_n^B(L) = \psi_n^B(U) = 0$. Let λ_n^B be the corresponding eigenvalues. The matrix elements of the operator $\mathcal{L}^{XBJ}$ in this basis are:

$$\mathcal{L}^{XBJ}(m,n) = \sum_{j=0}^{\infty} \lambda_j^J c_{mj} c_{nj} \quad \text{where} \quad c_{mn} = \int_B \psi_m^B(x) \psi_n(x) \mu(dx) \tag{24}$$

Let $a_n^{BJ}(j)$ be the eigenvectors of $\mathcal{L}^{BJ}(m, n)$ and let λ_n^{XBJ} be the corresponding eigenvalues. Furthermore, let $\psi_n^{BJ}(x) = \Sigma_j a_n^{BJ}(j) \psi_j^B(x)$ be the eigenvectors in the X-representation. Then the pricing kernel in the Y-representation is given by:

$$p^{\tilde{Y}BJ}(t, y_0, y_1) = H(x_0, x_1) \sum_{n=0}^{\infty} e^{t\lambda_n^{BJ}} \psi_n^{BJ}(x_0) \psi_n^{BJ}(x_1) \mu(x_1) \tag{25}$$

Implementation example

Here, we give a step-by-step construction to illustrate the theory in the article. As for X_t, we choose the CIR process of equation $dX = (a - bX)dt + \sqrt{\sigma} dW$. The Markov generator $\mathcal{L}^X$ in the new variable $\xi = \theta x$ is given by:

$$\mathcal{L}^X = \frac{1}{b}\left((\alpha + 1 - \xi) \frac{d}{d\xi} + \xi \frac{d^2}{d\xi^2} \right) \tag{26}$$

where $\theta = 2b/\sigma^2$ and $\alpha = (2a/\sigma^2) - 1$. Two linearly independent solutions f_1 and f_2 to the equation $\mathcal{L}^X f = \rho f$ are given by the following formulas:

$$f_1(x) = {}_1F_1\left(\frac{\rho}{b}, \alpha + 1, \theta x \right), \quad f_2(x) = (\theta x)^{-\alpha} {}_1F_1\left(\frac{\rho}{b} - \alpha, 1 - \alpha, \theta x \right) \tag{27}$$

Here, ${}_1F_1(A, B, z)$ is a confluent hypergeometric function and can be calculated by means of its Taylor expansion:

$${}_1F_1(A, B, z) = 1 + \frac{A}{B} z + \frac{A(A+1)}{B(B+1)} \frac{z^2}{2!} + \ldots$$

The eigenfunctions ψ_n can be expressed as follows:

$$\psi_n(x) = \sqrt{\frac{n!}{\Gamma(n + \alpha + 1)}} L_n^{(\alpha)}(\theta x) \tag{28}$$

where the $L_n^{(\alpha)}$ are Laguerre polynomials of order α.

The corresponding eigenvalues are $\lambda_n = -bn$ and the orthogonality measure $\mu(x)dx$ is equal to $\theta^{\alpha+1} x^{\alpha} e^{-\theta x} dx$. Laguerre polynomials can be efficiently calculated by means of the recurrence relation:

$$(n+1) L_{n+1}^{(\alpha)}(z) = (2n + \alpha + 1 - z) L_n^{(\alpha)}(z) - (n + \alpha) L_n^{(\alpha)}(z) \tag{29}$$

with initial conditions $L_{-1}^{(\alpha)} = 0$, $L_0^{(\alpha)} = 1$.

Once the constants c_1, c_2, c_3, c_4 are specified, we can construct the diffeomorphism $Y(x)$ and invert it to find the function $X(y)$. The derivative $Y'(x)$ can be explicitly expressed in terms of confluent hypergeometric functions using equation (27) and the identity:

$$\frac{d}{dz} {}_1F_1(A, B, z) = \frac{A}{B} {}_1F_1(A + 1, B + 1, z)$$

This yields the local volatility function $\nu(y)$ in equation (5). The function $H(x_0, x_1)$ in equation (6) can also be expressed in terms of the derivative $Y'(x)$.

To calculate the Green's function defined after equation (11), it is convenient to start from the following equation:

$$(\mathcal{L}^X - \rho) G(\rho, x_0, x_1) = \delta(x_1 - x_0)$$

The solution can be expressed in terms of two linearly independent solutions of the differential equation $\mathcal{L}^X f = \rho f$, namely:

1. Convergence of series for option prices

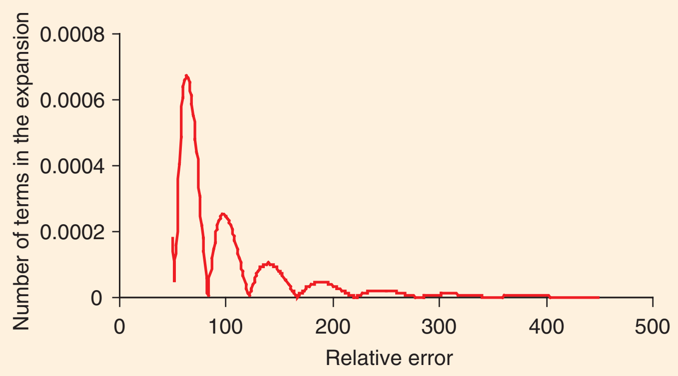

2. Implied volatility surface for the example

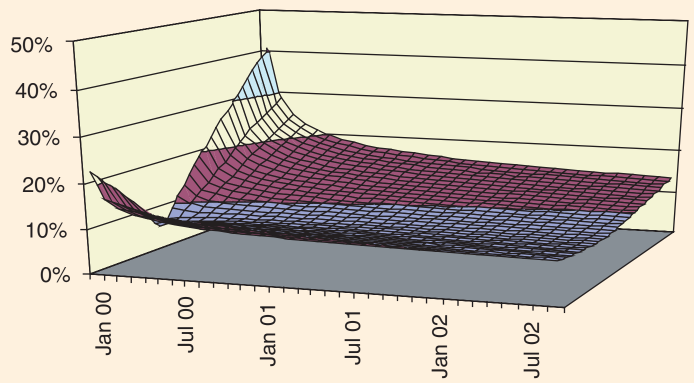

$$G(\rho, x_0, x_1) = \begin{cases} \dfrac{\Gamma\left(\frac{\rho}{b}\right)}{b\Gamma(\alpha+1)} M(x_0) U(x_1) \mu(x_1) & \text{in case } x_0 < x_1 \\ \dfrac{\Gamma\left(\frac{\rho}{b}\right)}{b\Gamma(\alpha+1)} M(x_1) U(x_0) \mu(x_1) & \text{otherwise} \end{cases}$$

Here $M(x) = f_1(x)$ and:

$$U(x) = \frac{\pi}{\sin(\pi(\alpha+1))} \left(\frac{f_1(x)}{\Gamma\left(\frac{\rho}{b} - \alpha\right)} - \frac{f_2(x)}{\Gamma\left(\frac{\rho}{b}\right)\Gamma(1-\alpha)} \right)$$

To construct a realistic example, we choose the following parameters: $a = 0.1$, $b = 0.06$, $\sigma = 0.42$, $X_0 = 1$, $\rho = 0.05$, $c_1 = 5$, $c_2 = -2$, $c_3 = -2$ while c_4 is determined out of the condition $g(X_0) = 1$. The gamma process for the Bochner subordinator has parameters $\mu = 1$ and $\nu = 0.1$ while the stochastic volatility portion is given by $c = 0.1$, $d = 0.1$, $\zeta = 0.3$ with initial condition $u_0 = 1$. To assess the speed of convergence, we estimated the number N of terms in the series expansion for European-style call prices that are needed to achieve penny accuracy, that is, errors smaller than 10^{-4}. We find that six-month maturities require only 100 terms, one-month maturities about 500 terms (see figure 1), and one-week maturities 2,000 terms. When applied to the lattice discretisations in Albanese & Kuznetsov (2003), convergence is noticeably faster and one-week time horizons require about 50 terms to calculate the node-to-node transition probabilities.

Figures 2 and 3 show sample shapes for the implied volatility surface and the terms structure of at-the-money implied volatilities corresponding to the particular choice of parameters in our examples. Obviously, all the shapes corresponding to the standard variance-gamma, Heston and hypergeometric local volatility models can also be reproduced as particular limiting cases.

3. Term structure of implied ATM volatility

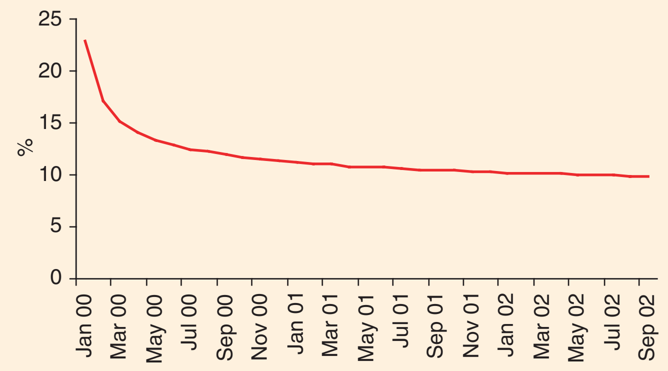

REFERENCES

Albanese C, G Campolieti, P Carr and A Lipton, 2001
Black-Scholes goes hypergeometric
Risk December, pages 99–103

Albanese C, G Campolieti, O Chen and A Zavidonov, 2003
Credit barrier models
Risk June, pages 109–113

Albanese C and A Kuznetsov, 2003
Discretization schemes for subordinated processes
Preprint, available at *www.ma.ic.ac.uk/_calban*

Black F, 1976
The pricing of commodity contracts
Journal of Financial Economics 3(1), pages 167–179

Carr P, H Geman, D Madan and M Yor, 2001
Stochastic volatility for Levy processes
Working paper

Carr P and L Wu, 2004
Time changed Levy processes and option pricing
Journal of Financial Economics 71(1), January, pages 113–141

Cox J and S Ross, 1976
The valuation of options for alternative stochastic processes
Journal of Financial Economics 3, pages 145–166

Derman E, 1999
Regimes of volatility
Quantitative Strategies Research Notes, Goldman Sachs

Dupire B, 1994
Pricing with a smile
Risk January, pages 18–20

Heston S, 1993
A closed-form solution for options with stochastic volatility with applications to bond and currency options
Review of Financial Studies 6, pages 327–343

Madan D, P Carr and E Chang, 1998
The variance-gamma process and option pricing
European Financial Review 2, pages 79–105

Rady S, 1997
Option pricing in the presence of natural boundaries and a quadratic diffusion term
Finance and Stochastics 1, pages 331–344

Concluding remarks

We have built a family of analytically tractable option pricing models combining state-dependent volatility, stochastic volatility and jumps. The construction is based on the method of eigenfunction expansions, and accounts for jumps and stochastic volatility by altering the time-dependent coefficients of a series expansion. This technique can be applied to numerous sub-fields of option pricing theory. Lattice versions of these models are discussed in separate papers such as Albanese & Kuznetsov (2003). ■

Claudio Albanese recently took up a chair in mathematical finance at Imperial College, University of London. Alexey Kuznetsov is a graduate student at the University of Toronto. **This work was written while the authors were visiting the National University of Singapore, whose hospitality is gratefully acknowledged. We thank Oliver Chen, Giuseppe Campolieti, Peter Carr, Pierre Hauvillier and Stephan Lawi for discussions. Remaining errors are our own.** **Email: claudio.albanese@imperial.ac.uk, alexey.kuznetsov@utoronto.ca**