

Diffusion-generated Jump Processes:
A framework for efficient pricing with jumps and
stochastic volatility

D. B. R. Derrett

University of Cambridge

June 19, 2015

Abstract

We construct a flexible and numerically tractable class of asset models by firstly choosing a bivariate diffusion process (U, Y) , and then defining the price of the asset at time t to be the value of Y when U first exceeds t . Such price processes will typically have jumps; conventional pricing methodologies would try to solve a PIDE, which can be numerically problematic, but using the fact that the pricing problem is embedded in a two-dimensional diffusion, we are able to exploit efficient methods for two-dimensional diffusion equations to find prices. Models with time dependence (that is, where the bivariate diffusion is U -dependent) are no more difficult in this approach.

Pricing a European option for a model in this class consists of solving a linear second order elliptic PDE. This problem is amenable to highly optimized numerical PDE solving techniques.¹

Keywords: asset price model, finite element method, partial differential equation, jump process, inhomogeneous Markov process.

¹The author is grateful to L.C.G. Rogers for his suggestions on presentation.

1 Introduction

Figure 1.1 shows a sample path of a two-dimensional Brownian motion X with rightwards drift. Those points where the Brownian motion is further right than it has been before are marked in black. These black points may be interpreted as the sample path of a one-dimensional, jumping process Z .




Figure 1.1: A drifting Brownian motion.

Suppose Z represents the path of an asset price. Pricing a European option on Z is a matter of computing an integral with respect to the law of Z_T , where T is fixed and positive.

The value of this integral, given the starting point of X , satisfies a PDE on the half-plane to the left of the vertical line on the right. This is a second order elliptic PDE which may be solved numerically. In this case the homogeneity of the diffusion leads to constant coefficients, so one could obtain an analytical solution using the Fourier transform².

The method we have described associates a jumping process with each rightwards-drifting diffusion. For this class of jumping processes, the prices of European options may be readily computed. We may also add an extra dimension to the diffusion to represent, for example, current market conditions, or the volatility

²Of course, this is no longer possible when the diffusion is fully inhomogenous.

of the asset price process. This allows the asset price to exhibit volatility clustering effects. We consider models where the jumping process Z denotes the discounted asset price process. In section 2, we define this class of models and derive the pricing PDE.

Our asset price model is thus a time-changed diffusion. Unlike many existing models, here we subordinate a diffusion with a *dependent* increasing process. Models generated by subordination of a process with an *independent* subordinator are well-studied. Subordination was first considered by Bochner (1949) and introduced into finance by Clark (1973). Many Lévy process models can be represented as Brownian motions with an appropriate independent subordinator (see Geman, Madan, and Yor (2001)).

The accumulated empirical evidence suggests that there are certain typical properties of asset price processes³, which hold for different asset classes and in different time periods. These are known as stylized facts, and a collection of them may be found in Cont (2001). We present some of these facts and compare the proposed models with well-known models:

1. *The volatility smile.* If we model the rate of return using a Brownian motion, we obtain a lognormal model, under which the European option price may be computed using the formula of Black, Merton and Scholes. It became clear in the wake of the Black Monday crash of 1987 that the lognormal models failed to allow for the heavy tails⁴ seen in the empirical returns of assets. One solution is to model the log asset price instead by a Lévy process. The normal-inverse Gaussian example of Section 3 shows that the models presented here are able to represent this effect.
2. *Volatility clustering.* Volatility clustering is a name given to the autocorrelation of the absolute value of asset returns, as is present, for example, in the Heston model. Models where the asset price is the exponential of a

³The properties given here are with respect to the physical probability measure, rather than an equivalent martingale measure.

⁴This is also known as leptokurtosis – from the Greek leptós ‘thin’ and kurtosis ‘bulging’ – or ‘the volatility smile’.

Lévy process have independent increments and so are unable to represent this effect⁵. Volatility clustering occurs for models in the class we present if there is spatial inhomogeneity or if the model is driven by a diffusion of dimension three or higher⁶.

3. *The term structure of implied volatility.* A reasonable model should achieve a balance between parsimony and the correct pricing of European options. Options traders are aware of planned events which will lead to increased volatility, such as announcements of interest rates, inflation rates and employment rates⁷. The models proposed here allow the volatility of the driving diffusion to depend on the time and asset price. Hence one may construct a sequence of high volatility periods and to introduce level-dependent volatility effects, as in local volatility models. The models presented here are able to represent this structure (see the time-inhomogenous example of section 3.2).
4. *Jumps in price.* Jumps provide a skew to the short end of the term structure of implied volatility which is not present in diffusion models.
5. *Gain/loss asymmetry.* Large drawdowns in asset prices, such as stock market crashes, are observed more frequently than equally large upward movements. The models presented here are able to represent this effect by including a positive correlation between the asset price process and the time process. See the example of section 3.1.3.

The following are desirable properties for an asset pricing model:

1. *Simplicity.* To price an option under the Black-Scholes model requires the solution of a second order parabolic PDE. In the models presented here, one solves an elliptic counterpart to that PDE. This is a classical problem for which fast and robust numerical methods have been developed. Pricing

⁵Mendoza-Arriaga, Carr, and Linetsky (2010) subordinates Lévy processes to obtain volatility clustering.

⁶We only consider two-dimensional diffusions in this paper.

⁷See, for example, <http://www.forexfactory.com/calendar.php>.

a similar option under a Lévy model requires the solution of a partial integro-differential equation (PIDE), which may be more complex.

Pricing an exotic option using Monte Carlo methods is simple under the models presented here, since it just requires the simulation of a diffusion, rather than the direct simulation of a jumping process.

2. *Robustness.* This classical background of this problem in physics and engineering has resulted in robust solution methods. One may use the finite element method (FEM) together with modern techniques such as higher order elements, adaptive meshing and a posteriori error estimates.
3. *Speed.* Fast European option pricing is important, as model calibration may require this to be done repeatedly. For a model in this framework based on a two-dimensional diffusion, European option pricing requires the solution to a two-dimensional PDE on a half-plane. We show this can be done quickly using the FEM. When the underlying diffusion is homogeneous we may use Fourier transform methods.

Our framework includes a wide range of models, ranging from the most parsimonious, with few parameters, to those which can match the observed term structure of implied volatility.

One may price an American option in this framework by approximating it as a Bermudan option and solving a corresponding sequence of European option pricing problems. We do not consider the American option pricing problem in this paper.

We defer most proofs to the appendix, Section 5.

2 Modelling framework

This section introduces the class of models under consideration, and develops their use for pricing European derivatives written on discontinuous underlying assets.

Definition 2.1. A *diffusion-generated jump process* (DJGP) is a process S^* which can be represented as

$$S_t^* = Y(H_t),$$

where

$$H_t \equiv \inf\{s : U_s > t\},$$

and the process $X_t \equiv (U_t, Y_t)$ is a diffusion solving an SDE with Lipschitz coefficients:

$$dX_t = \Sigma(X_t)dW_t + \mu(X_t)dt.$$

REMARKS.

1. The process Y takes values in \mathbb{R}^d for some $d \geq 1$. The examples discussed will all be with $d = 1$.
2. The time index of a DGJP can in principle be any real value, though we may restrict attention to non-negative reals at times.
3. We may speak of a DGJP(μ, Σ) if it is necessary to make explicit the coefficients of the SDE.
4. It will be assumed implicitly that the coefficients of the SDE are such as to guarantee that $H_t < \infty$ almost surely for each t . This has to be checked in any example.
5. The principal focus of this modelling approach is on the pricing of European-style derivatives written on underlying assets which may have jumps. In this context, we think of S^* as the discounted price process, and will accordingly be looking for examples in which S^* is a martingale. The

transformation to the undiscounted price process is trivial in the case of a constant interest rate r :

$$S_t \equiv e^{rt} S_t^*.$$

Lemma 2.1. *Let S^* be a DGJP(μ, Σ), where Σ is bounded and*

$$\mu(x) = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Then S^ is a martingale.*

Proof. See Section 5.1 of the appendix. □

Remark 2.1. Lemma 2.1 gives us a simple way of creating examples, since we may choose any bounded volatility structure Σ which satisfies the Lipschitz condition.

The first and simplest example will be where $d = 1$ and

$$(2.1) \quad \Sigma = \begin{pmatrix} 0 & 0 \\ 0 & \sigma y \end{pmatrix}, \quad \mu(x) = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

for some $\sigma > 0$. This example falls outside the scope of Lemma 2.1, but it is evident that the DGJP(μ, Σ) is the Black-Scholes discounted asset price process.

Lemma 2.2. *The process $(X_{H_t})_{t \geq 0}$ is Strong Markov.*

Proof. This follows from the Strong Markov property of solutions to well-posed martingale problems. See Theorem 5.4.20 of Karatzas and Shreve (1988). □

2.1 Deriving the PDE for pricing European options.

In this section, we show that the price of a European option is given by the solution to a second order elliptic PDE. Consider a European option with expiry

T and bounded⁸ payoff f . Suppose we adopt the notational conventions⁹ $u \equiv x_1 \equiv e_1 \cdot x$ and define

$$(2.2) \quad V(x) = e^{-r(T-u)} E[f(e^{rT} S_T^*) | X_0 = x]$$

for x such that $e_1 \cdot x \leq T$. Now by exploiting the Markov property of the strong solution X to an SDE with Lipschitz coefficients, it is clear that up until the time H_T the process

$$(2.3) \quad M_t \equiv e^{-rU_t} V(X_t) = e^{-rT} E[f(e^{rT} S_T^*) | X_0 = X_t] = e^{-rT} E[f(e^{rT} S_T^*) | \mathcal{F}_t]$$

is a martingale. Moreover, it is also evident that

$$(2.4) \quad V(x) = f(e^{rT} y) \quad \text{if } u \equiv x_1 = T.$$

We intend therefore to find a function $V : (-\infty, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying the boundary condition (2.4), and such that M defined at (2.3) is a martingale. The tool that allows this is of course Itô's formula.

Theorem 2.1. *Suppose that S^* is a DGJP(μ, Σ) which is also a martingale, and suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is bounded. Suppose that r is non-negative. Fix some $T \in \mathbb{R}$ and let $a \equiv \Sigma \Sigma^T$. Suppose that $F : (-\infty, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is bounded, satisfies the boundary condition (2.4), and the PDE*

$$(2.5) \quad 0 = \tfrac{1}{2} \text{tr}(a D^2 F) + (\mu - r a e_1) \cdot DF + r e_1 \cdot (\tfrac{1}{2} r a e_1 - \mu) F.$$

Then $F = V$, where V is the option price function defined at (2.2), at points accessible to the process X .

Proof. Define, for $0 \leq t \leq H_T$, the process N by

$$N_t = e^{-rU_t} F(X_t).$$

⁸The theory survives the relaxation of this assumption, provided we impose suitable technical conditions.

⁹We write e_1 for the $(d+1)$ -vector with 1 in the first entry, zeros elsewhere.

A routine application of Itô's formula shows that dN_t equals, up to driftless terms,

$$(\tfrac{1}{2} \text{tr}(aD^2F) + (\mu - rae_1) \cdot DF + re_1 \cdot (\tfrac{1}{2}rae_1 - \mu)F)dt = 0.$$

Hence N is a continuous local martingale. Since, r is non-negative,

$$N_t \leq e^{-rT} F(U_t, Y_t),$$

and hence N is bounded and thus a true martingale.

$$\begin{aligned} N_{H_T} &= e^{-rT} F(U_{H_T}, Y_{H_T}) \\ &= e^{-rT} f(Y_{H_T}) \\ &= M_{H_T} \end{aligned}$$

In particular, on the event $\{t \leq H_T\}$, $F(U_t, Y_t) = V(U_t, Y_t)$, so F equals V at points accessible to the process X . \square

Remark 2.2. For the Black-Scholes example (Example (2.1)), we see that the PDE (2.5) takes the form

$$0 = \tfrac{1}{2}\sigma^2 y^2 F_{yy} + F_u - rF.$$

Recalling that the variable y is the discounted price, if we try to write the solution in terms of the undiscounted price as $F(u, y) = g(u, e^{ru}y) \equiv g(u, s)$, a few calculations lead to the alternative PDE

$$0 = \tfrac{1}{2}\sigma^2 s^2 g_{ss} + rsg_s + g_u - rg,$$

which is of course the standard Black-Scholes PDE.

3 Examples

In this section, we give examples of DGJP models with and without μ and Σ constant. We give three examples of constant coefficient models: the Black-Scholes model, the normal-inverse Gaussian (NIG) model and a model with

correlation between Y and U . These belong to the family of exponential Lévy models. For these models, the pricing PDE has constant coefficients, so the price may be determined by Fourier transform methods¹⁰.

The constant coefficient models admit a quasi-closed form solution, so we use them to gauge the accuracy of the numerical solution given by the FEM.

Finally, we give an example of a time-dependent volatility model, suitable for short-dated options. In this model there is a single period of high volatility, with low volatility at other times. The period of high volatility could correspond to a volatility event, such as the release of pertinent information to the market.

3.1 Constant coefficient models

The examples studied here are DGJP(μ, Σ) processes of the form

$$\mu = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \Sigma = \begin{pmatrix} \kappa & 0 \\ \tau y & \sigma y \end{pmatrix},$$

where $\kappa, \sigma > 0$, and $\tau \in \mathbb{R}$. This being the case, the process $X = (U, Y)$ has the form

$$\begin{aligned} U_t &= t + \kappa W_t, \\ Y_t &= Y_0 \exp[\sigma W'_t + \tau W_t - \tfrac{1}{2}(\tau^2 + \sigma^2)t], \end{aligned}$$

where W and W' are two independent Brownian motions. Since U drifts upward, all the hitting times H_t are finite. The process Y is of course a martingale, and by Fatou's lemma we know that $(Y_{H_t})_{t \geq 0}$ is a positive supermartingale. But we can say more.

Lemma 3.1. *Assume that*

$$(3.3) \quad 1 + \tau\kappa \geq 0.$$

Then the process $(Y_{H_t})_{t \geq 0}$ is a positive martingale.

¹⁰Of course, this is not possible in the general case.

PROOF. With no loss of generality, we shall assume that $Y_0 = 1$. We shall prove that $E[Y_{H_t}] = 1$ for all $t \geq 0$, which is sufficient. By conditioning firstly on the entire path of W , we develop

$$\begin{aligned} E[Y_{H_t}] &= E \exp[\tau W_{H_t} - \tfrac{1}{2} \tau^2 H_t] \\ &= E \exp[\tau \kappa^{-1} U_{H_t} - (\tfrac{1}{2} \tau^2 + \tau \kappa^{-1}) H_t] \\ &= e^{t\tau/\kappa} E \exp[-(\tfrac{1}{2} \tau^2 + \tau \kappa^{-1}) H_t]. \end{aligned}$$

We see therefore that everything depends on the distribution of the first hitting times of a drifting Brownian motion. As is well known (see, for example, O. E. Barndorff-Nielsen (1997)), the MGF of H_a is given by

$$(3.5) \quad E \exp\{-\lambda H_a\} = \exp\{-a\psi(\lambda)\}$$

where

$$\psi(\lambda) \equiv \frac{\sqrt{1 + 2\lambda\kappa^2} - 1}{\kappa^2},$$

valid for all λ with real part at least $-1/2\kappa^2$. This is now applied to (3.4) with the identification

$$\lambda = \tfrac{1}{2} \tau^2 + \tau \kappa^{-1},$$

which for real τ is easily seen to be at least $-1/2\kappa^2$, and therefore within the range of validity of (3.5). We obtain

$$\psi(\lambda) = \frac{|1 + \tau\kappa| - 1}{\kappa^2} = \tau/\kappa$$

in view of the assumption (3.3). Returning this to (3.4) gives the conclusion $E[Y_{H_t}] = 1$, as required.

□

3.1.1 Black-Scholes model

In Example 2.1, we showed that we can consider the Black-Scholes model as a DGJP model. Under this model $U_t = t$, for $t \geq 0$. The DGJP models with this property are local volatility models, with continuous asset price paths.




Figure 3.1: The price of a put option with parameters $T = 1$, $K = 1$ and $r = 0$ in a normal-inverse Gaussian model (Section 3.1.2) with parameters $\kappa = 0.1$ and $\sigma = 0.3$.

The closed-form expression for the Black-Scholes put price makes the Black-Scholes model a test of the accuracy numerical PDE solution. Table 3.1 compares the prices of European put options given by the formula with those given by the numerical solution of (2.5), using the finite element method of Section 4.

3.1.2 Normal-inverse Gaussian model

In this subsection, we give a class of DGJP models whose stock price processes are exponential normal-inverse Gaussian (NIG) processes. The NIG process is a Lévy process corresponding to the NIG distribution, which was discovered by O. Barndorff-Nielsen and Halgreen (1977). There is a closed-form expression for the PDF of the NIG distribution in terms of a modified Bessel function of the second kind (see O. Barndorff-Nielsen and Halgreen (1977)). By integrating the payoff of the put option against this we are able to check the numerical accuracy of the PDE solver. The NIG process is obtained by subordinating a Brownian

r	T	S_0	K	Analytical price (\$)	Computed price (\$)	Time (s)	Basis point error
0.0000	1.0	1.0	1.0	0.119235	0.119232	0.4	-0.3012
0.0953	1.0	1.0	1.0	0.074009	0.074006	0.4	-0.4327
0.0000	2.0	1.0	1.0	0.167996	0.167993	0.4	-0.1681
0.0000	1.0	1.0	0.8	0.035344	0.035341	0.5	-0.6917
0.0000	1.0	0.8	1.0	0.235344	0.235341	0.4	-0.1043
0.0000	1.0	1.2	1.0	0.054406	0.054402	0.4	-0.5988

Table 3.1: Prices for European put options under the Black-Scholes model, with $\sigma = 0.3$ (see Remark 2.2). r is the risk-free rate, T is the expiry time, S_0 is the initial asset price and K is the strike price. The analytical price is the price given by the closed-form expression for the put price. The option price is computed numerically using the method given in Section 4.

motion with drift by an independent Brownian motion with drift. We note that for the NIG model, the implied volatility surface shows a skew at the short end of the term structure of implied volatility, as is expected from a model with jumps (see figure 3.3).

We start by defining a parameterization of the process, as given in O. Barndorff-Nielsen and Halgreen (1977).

Definition 3.1. Fix parameters α , β and δ in \mathbb{R} satisfying $\alpha \geq |\beta|$ and $\delta > 0$. Let W and W' be independent Brownian motions. Define, for $t \geq 0$, the stopping time

$$H_t = \inf\{u \geq 0 : W_u + \gamma u = \delta t\},$$

where

$$\gamma = \sqrt{\alpha^2 - \beta^2}.$$

Define the process $(Z_t)_{t \geq 0}$, by

$$Z_t = W'_{H_t} + \beta H_t.$$

Then we say Z is the normal-inverse Gaussian (NIG) process corresponding to parameters (α, β, δ) .

The NIG process has stationary and independent increments and is thus a Lévy process. In our framework, the subordinator is given by the hitting times $(H_t)_{t \geq 0}$, and the subordinated process is the log of the stock price.



Figure 3.2: A price path of the NIG model. The grey path is (U, Y) and the black points correspond to the path of $(Y_{H_t})_{t \geq 0}$.

Remark 3.1. Consider the PDE given by (2.5) in the case of the DGJP model of Lemma 3.1 (again writing $x = (u, y)$):

$$0 = \frac{1}{2} ((\sigma y)^2 \partial_{yy} F + \kappa \partial_{uu} F) + (1 - \kappa r) \partial_u F + r (\frac{1}{2} \kappa r - 1) F$$

As in Remark 2.2, we make the change of coordinates

$$s = e^{ru} y$$

and let g be such that $g(t, s) = F(t, y)$. After some calculation, we see that g satisfies the PDE

$$0 = g_u - rg + \frac{1}{2}\sigma^2 s^2 g_{ss} + rsg_s + \kappa \left(\frac{1}{2}r^2 s^2 g_{ss} + \frac{1}{2}g_{uu} + rsg_{su} - \frac{1}{2}r^2 sg_s - rg_u + \frac{1}{2}r^2 g \right).$$

This is the Black-Scholes PDE, but with additional terms (after the first line). In particular, for κ non-zero, this is an elliptic, rather than parabolic, PDE. The Black-Scholes PDE is recovered as κ tends to 0.

r	T	S_0	K	Analytical price (\$)	Computed price (\$)	Time (s)	Basis point error
0.0000	1.0	1.0	1.0	0.117800	0.117799	0.5	-0.0491
0.0953	1.0	1.0	1.0	0.072840	0.072840	0.5	-0.0610
0.0000	2.0	1.0	1.0	0.166944	0.166943	0.8	-0.0554
0.0000	1.0	1.0	0.8	0.034904	0.034904	0.7	-0.0982
0.0000	1.0	0.8	1.0	0.234904	0.234904	0.7	-0.0145
0.0000	1.0	1.2	1.0	0.053580	0.053579	0.8	-0.0854

Table 3.2: Prices for European put options under a NIG model, with parameters $\kappa = 0.1$ and $\sigma = 0.3$. Parameters r , T , S_0 and K are as in Table 3.1. The analytical price is computed by direct integration of the payoff against the PDF of the NIG distribution. The option price is computed numerically using the finite element method (see Section 4). Note that the error is lower than for the Black-Scholes model, potentially due to the added regularity of the elliptic operator.




Figure 3.3: An implied volatility surface for the NIG model with parameters $\kappa = 0.1$ and $\sigma = 0.3$, with the risk-free rate r equal to 0. The implied volatility was computed numerically and interpolated to form a surface. Note the skew at the short end of the term structure.

3.1.3 Model with gain/loss asymmetry

For a model with Σ bounded, it is much easier to show that S^* is a martingale and thus that the model is a DGJP model (Remark 2.1). Thus in practical applications, where the asset price may be modeled as being bounded above by a large constant, the work done in Lemma 3.1 the previous two theorems to show that the martingale property holds is unnecessary.




Figure 3.4: A plot of $\partial V / \partial S$ (the delta) for the NIG model with parameters $\kappa = 0.1$ and $\sigma = 0.3$, with the risk-free r rate equal to 0. This is the derivative of the value function with respect to the price. This was computed by first computing the weak solution to the pricing PDE, and then projecting the derivative onto the finite element space S_h (see Section 4).

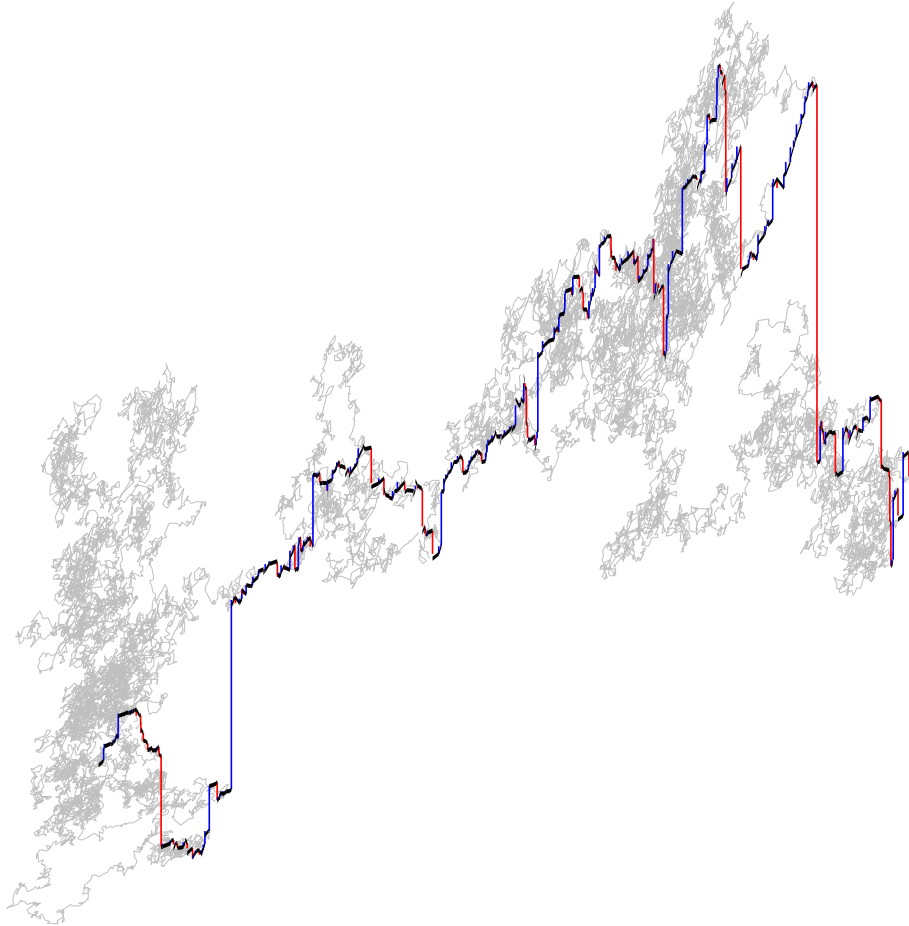


Figure 3.5: A price path of the model with dependence between Y and U . The grey path is a drifting two-dimensional Brownian motion, with positive covariance between the component processes. The black points correspond to the generated price process. The blue and red lines respectively mark upwards and downwards jumps of the asset price. Note that, while the asset price is a martingale, the preponderance of large jumps are downwards, giving a gain/loss asymmetry.




Figure 3.6: An implied volatility surface for the model with gain/loss asymmetry with parameters $\kappa = 0.16$, $\sigma = 0.3$ and $\tau = 0.1$, with the risk-free r rate equal to 0. The implied volatility was computed numerically and interpolated to form a surface.

r	T	S_0	K	Analytical price (\$)	Computed price (\$)	Time (s)	Basis point error
0.0000	1.0	1.0	1.0	0.121913	0.121913	1.2	-0.0252
0.0953	1.0	1.0	1.0	0.078620	0.078620	0.9	-0.0220
0.0000	2.0	1.0	1.0	0.173373	0.173373	0.8	-0.0236
0.0000	1.0	1.0	0.8	0.041379	0.041379	1.3	-0.0302
0.0000	1.0	0.8	1.0	0.234346	0.234346	0.4	-0.0103
0.0000	1.0	1.2	1.0	0.061327	0.061327	0.8	-0.0275

Table 3.3: Prices for European put options under the model with gain/loss asymmetry, with parameters $\kappa = 0.16$, $\sigma = 0.3$ and $\tau = 0.1$. Parameters r , T , S_0 and K are as in Table 3.1. The analytical price is computed using the method given in the appendix (see Section 5.2). The computed option price is computed numerically (see Section 4). Note that accuracy is achieved to the order of a hundredth of a basis point.

3.2 Time-dependent volatility model

In this subsection, we present an example where the volatility of the two-dimensional driving process depends on time. In this model, there is a single period of high volatility, with low volatility at other times. This could correspond, for example, to the case of a short-dated option, with a volatility event occurring before expiry, such as the release of employment figures, or other pertinent economic information.

Define, $\mu(u, y)$, for (u, y) in \mathbb{R}^2 by

$$\mu \begin{pmatrix} u \\ y \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Define, $\Sigma(u, y)$, for (u, y) in \mathbb{R}^2 by

$$\Sigma \begin{pmatrix} u \\ y \end{pmatrix} = \begin{pmatrix} \sqrt{\kappa(u)} & 0 \\ 0 & \sigma(u)y \end{pmatrix},$$

where

$$\sigma(u) = \begin{cases} 0.8 & \text{if } u \in [0.5, 0.6] \\ 0.2 & \text{otherwise,} \end{cases}$$

and

$$\kappa(u) = \begin{cases} 0.2 & \text{if } u \in [0.5, 0.6] \\ 0 & \text{otherwise.} \end{cases}$$

This gives a DGJP model (aside from the fact that Σ is not Lipschitz)¹¹.

¹¹For simplicity of exposition, Σ is not Lipschitz. One could smooth it to introduce this property without changing the behaviour of the model.




Figure 3.7: A price path under the time-dependent volatility model, showing a short period of high volatility.




Figure 3.8: An implied volatility surface for the time-dependent volatility model.

4 Numerical solution to the pricing PDE

In this section, we solve the pricing PDE numerically using the finite element method for the NIG model example, in the case of a European put option, with strike K and expiry T . The same approach works for the general DGJP model. For an introduction to the finite element method, see Iserles (2008). There are three steps required to prepare our PDE for solution by the FEM: first we change variables, second we assign boundary conditions and finally we compute the weak form of the PDE.

Change of variables. Changing of variables to the log price will allow the

resulting PDE to be solved efficiently on an evenly-spaced grid, as in the case of a Black-Scholes model. To compute the price of the option, we compute an approximation to the function \tilde{V} and then change variables to recover V . Define a function V , for $(u, \tilde{y}) \in (-\infty, T] \times \mathbb{R}$ by

$$\tilde{V}(u, \tilde{y}) = V(u, e^{\tilde{y}}).$$

Define $(\tilde{Y}_u)_{u \geq 0}$ by

$$\tilde{Y}_u = \log(Y_u),$$

and note that, by Itô's lemma,

$$d \begin{pmatrix} U_t \\ \tilde{Y}_t \end{pmatrix} = \begin{pmatrix} 1 \\ -\frac{1}{2}\sigma^2 \end{pmatrix} dt + \begin{pmatrix} \sqrt{\kappa} & 0 \\ 0 & \sigma \end{pmatrix} dW_t.$$

Recall that the process given for $u \geq 0$ by

$$e^{-rU_t} V(U_t, Y_t) = e^{-rU_t} \tilde{V}(U_t, \tilde{Y}_t)$$

is a martingale. By Itô's lemma

$$d \left(e^{-rU_t} \tilde{V}(U_t, \tilde{Y}_t) \right) = e^{-rU_t} \begin{pmatrix} -r\tilde{V}(dt + \sqrt{\kappa}dW_t^1) \\ +\frac{1}{2}r^2\tilde{V}\kappa dt \\ +\tilde{V}_{\tilde{y}}(\sigma dW_t^2 - \frac{1}{2}\sigma^2 dt) \\ +\frac{1}{2}\sigma^2\tilde{V}_{\tilde{y}\tilde{y}}dt \\ +\tilde{V}_u(dt + \sqrt{\kappa}dW_t^1) \\ +\frac{1}{2}\kappa\tilde{V}_{uu}dt \\ -r\kappa\tilde{V}_u dt \end{pmatrix},$$

so, since martingales are driftless, \tilde{V} satisfies the PDE

$$0 = (\frac{1}{2}\kappa r^2 - r) \tilde{V} + \begin{pmatrix} 1 - r\kappa \\ -\frac{1}{2}\sigma^2 \end{pmatrix} \cdot \begin{pmatrix} \tilde{V}_u \\ \tilde{V}_{\tilde{y}} \end{pmatrix} + \frac{1}{2}\sigma^2 \tilde{V}_{\tilde{y}\tilde{y}} + \frac{1}{2}\kappa \tilde{V}_{uu}.$$

Boundary conditions. To compute a solution to the PDE, we restrict its domain from a half-plane to a finite domain. The terminal boundary condition for V becomes

$$\tilde{V}(T, y) = (K - e^{\rho T} y)^+,$$

for y in \mathbb{R} . By fixing $m < M$ and stopping \tilde{Y} when it leaves the interval $[m, M]$, we obtain the boundary conditions

$$\tilde{V}(t, \tilde{y}) = (K - \tilde{y})^+,$$

when y equals m or M . Finally, we impose a reflecting boundary condition:

$$\tilde{V}_{\tilde{y}}(t_0, y) = 0,$$

for fixed $t_0 < 0$ and $y \in [m, M]$. These boundary conditions provide a good approximation when the probability that (U, \tilde{Y}) hits $([t_0, T] \times \partial[m, M]) \cup (\{t_0\} \times [m, M])$ is small.

Weak form of the PDE. The FEM computes an approximation to the weak solution to the PDE. To compute the bilinear form associated with the PDE, we integrate by parts¹². Let p be a smooth test function, zero on the boundary of $[m, M] \times [t_0, T]$. Then

$$\begin{aligned} & \int_{t_0}^T \int_m^M \left(\left[\begin{pmatrix} -\frac{1}{2}\sigma^2 \\ 1 \end{pmatrix} \cdot \begin{pmatrix} \tilde{V}_{\tilde{y}} \\ \tilde{V}_u \end{pmatrix} + (\frac{1}{2}\kappa r^2 - r)\tilde{V} \right] p \right) (\tilde{y}, u) d\tilde{y} du \\ &= \int_{t_0}^T \int_m^M -\frac{1}{2} \left([\sigma^2 \tilde{V}_{\tilde{y}\tilde{y}} + \kappa \tilde{V}_{uu}] p \right) (\tilde{y}, u) d\tilde{y} du \\ &= \int_{t_0}^T \int_m^M \frac{1}{2} (\sigma^2 \tilde{V}_{\tilde{y}\tilde{y}} p_{\tilde{y}} + \kappa \tilde{V}_{uu} p_u) (\tilde{y}, u) d\tilde{y} du, \end{aligned}$$

where the second equality follows from integrating by parts and the vanishing of p on the boundary of $[m, M] \times [t_0, T]$.

The finite element space. We first create a uniform triangulation of the rectangle $[t_0, T] \times [m, M]$, where T is the expiry time of the option, m is the minimum value of the log stock price and M is the corresponding maximum

¹²If we were computing the weak form for a DGJP model driven by a higher dimensional diffusion, we would use Gauss's Theorem.

value. We used the following parameters:

$$t_0 = -1$$

$$m = -4$$

$$M = 4$$

For this triangulation, we use a mesh with 400 vertices in the price axis and 8 in the time axis. We denote by S_h the corresponding finite element space, using cubic triangular elements. We implemented this scheme in the Python programming language, using the FEniCS package.

5 Appendix

5.1 The martingale property of the discounted asset price

Lemma 5.1. *Take μ and Σ satisfying the conditions of Lemma 2.1. For $t \geq 0$, $\mathbb{E}(H_t) \leq t$.*

Proof. Fix $t \geq 0$. Define the process $(M_s)_{s \geq 0}$, by

$$M_s = U_s - s.$$

Since $\mu^1 \equiv 1$, M is a local martingale. Thus, as Σ is bounded, M is a true martingale, with $M_0 = 0$. Fix n and apply the Optional Stopping Theorem to M at the stopping time $H_t \wedge n$. Then

$$0 = \mathbb{E}[U_{H_t \wedge n} - H_t \wedge n],$$

so

$$\begin{aligned} \mathbb{E}[t - H_t \wedge n] &= \mathbb{E}[t - U_{H_t \wedge n}] \\ &\geq 0, \end{aligned}$$

thus

$$\mathbb{E}[H_t \wedge n] \leq t,$$

so by Fatou's Lemma, $\mathbb{E}[H_t] \leq t$. □

Proof of Lemma 2.1. Since $\mu^2 \equiv 0$, Y is a local martingale. Thus $(Y_{H_t})_{t \geq 0}$ is a local martingale. As Σ is bounded, there is a constant K such that, for $t \geq 0$, $[Y]_t \leq Kt$. So, by the previous lemma,

$$\begin{aligned}\mathbb{E}([Y]_{H_t}) &\leq \mathbb{E}[KH_t] \\ &\leq Kt \\ &< \infty,\end{aligned}$$

so $(Y_{H_t})_{t \geq 0}$ is a martingale, thus proving Lemma 2.1.

5.2 Analytical price in the model with gain/loss asymmetry

We derive the law of Y_{H_T} for the model with gain/loss asymmetry in terms of a NIG distribution. We follow the notation of that section.

$$H_T + \sqrt{\kappa}W_{H_T}^1 = T,$$

so

$$Y_{H_T} = Y_0 \exp \left[\tau \left(\frac{T - H_T}{\sqrt{\kappa}} \right) + \sigma W_{H_T}^2 - \frac{1}{2}(\sigma^2 + \tau^2)H_T \right],$$

so, the price of a put option may be expressed as

$$\mathbb{E} \left[\left(e^{-rT}K - Y_0 \exp \left[\frac{\tau T}{\sqrt{\kappa}} + \sigma Z \right] \right)^+ \right],$$

where Z has a NIG distribution with parameters¹³

$$\begin{aligned}\beta &= -\frac{\sigma^2 + \tau^2}{2\sigma} - \frac{\tau}{\sqrt{\kappa}\sigma} \\ \gamma &= 1/\sqrt{\kappa} \\ \delta &= T/\sqrt{\kappa}.\end{aligned}$$

¹³That is, has the same law as the NIG process with these parameters at time 1.

5.3 On the boundedness of Σ in Lemma 2.1

Here we give an example showing that if the hypothesis of boundedness of Σ is not satisfied in Lemma 2.1, the result may not pertain.

Define the processes U and Y by

$$\begin{aligned}dU_u &= du + dW_u \\dY_u &= -2Y_u dW_u.\end{aligned}$$

Then Y is a martingale with

$$Y_u = e^{-2U_u}.$$

But $Y_{H_t} = e^{-2t}$, and so Y_{H_t} is not a true martingale, though it is a local martingale.

References

- Barndorff-Nielsen, O. E. (1997). “Normal Inverse Gaussian Distributions and Stochastic Volatility Modelling.” *Scandinavian Journal of Statistics*, 24(1), 1–13.
- Barndorff-Nielsen, O. & Halgreen, C. (1977). “Infinite divisibility of the hyperbolic and generalized inverse Gaussian distributions.” *Probability Theory and Related Fields*, 38(4), 309–311.
- Bochner, S. (1949). “Diffusion Equation and Stochastic Processes.” *Proceedings of the National Academy of Sciences of the United States of America*, 35(7), 368–370.
- Clark, P. (1973). “A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices.” *Econometrica*, 41(1), 135–155.
- Cont, R. (2001). “Empirical properties of asset returns: stylized facts and statistical issues.” *Quantitative Finance*, 1(2), 223–236.
- Geman, H., Madan, D. B., & Yor, M. (2001). “Time Changes for Levy Processes.” *Mathematical Finance*, 11(1), 79–96.

- Iserles, A. (2008). *A First Course in the Numerical Analysis of Differential Equations (Cambridge Texts in Applied Mathematics)*. Cambridge University Press.
- Karatzas & Shreve. (1988). *Brownian Motion and Stochastic Calculus*.
- Mendoza-Arriaga, R., Carr, P., & Linetsky, V. (2010). “Time-changed markov processes in unified credit-equity modeling.” *Mathematical Finance*, 20(4), 527–569.