A Fourier transform method for pricing options on mean-reverting Lévy-driven assets

Tony Ware
University of Calgary
(Joint work with Xu Li)

November 12th, 2008

Overview

- A quick look again at Lévy processes.
- A review of Fourier transform methods for option pricing.
- A discussion of the extension to Lévy-driven processes and the use of semi-Lagrangian time-stepping for computation.
- A numerical example - put option pricing for a mean-reverting asset.

Introduction

Lévy processes

are the continuous-time version of random walks, i.e. cumulative sums of i.i.d. random variables, which must be infinitely divisible.
Given an infinitely-divisible distribution, with characteristic function $\phi(u) = \mathrm{e}^{\psi(u)}$, a Lévy process $\{X(t)\}_{t\geq 0}$ constructed from this satisfies

- $X(0) = 0$,
- the increments are independent and stationary,
- for $t > 0$ and $s \geq 0$ the distribution of the increment $X(t+s) - X(s)$ has characteristic function $\mathrm{e}^{t\psi(u)}$.
- $\{X(t)\}_{t\geq 0}$ is *càdlàg*.

Introduction

Lévy-Khintchine formula

The characteristic exponent $\psi(u)$ of a Lévy process can be expressed in the form

$$\psi(u) = \underbrace{\mathrm{i}\gamma u - \frac{1}{2}\sigma^2 u^2}_{diffusion} + \underbrace{\int_{\mathbb{R}} \mathrm{e}^{\mathrm{i}ux} - 1 - \mathrm{i}ux\mathbf{1}_{\{|x|<1\}}\nu(\mathrm{d}x)}_{jump(compensated)}.$$

Examples

- Normal distributions: $X(1) \sim N(\mu, \sigma^2)$, $\psi(u) = \mathrm{i}\mu u - \frac{1}{2}\sigma^2 u^2$.
- Poisson distributions: $\mathbb{P}[X_1 = n] = \mathrm{e}^{-\lambda}\frac{\lambda^n}{n!}$, $\psi(u) = \lambda(\mathrm{e}^{\mathrm{i}u} - 1)$.
- Variance-Gamma: $\psi(u) = -\frac{1}{\nu}\log\left(1 - \mathrm{i}\theta\nu u + \frac{1}{2}\sigma^2\nu u^2\right)$.

Variance gamma process simulations (Matthias Winkel)

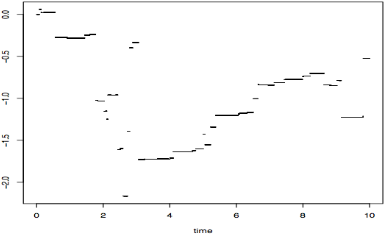

Variance Gamma process with shape parameter 0.5 and scale parameter 1

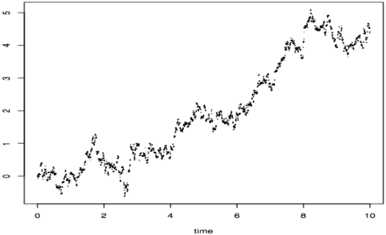

Variance Gamma process with shape parameter 50 and scale parameter 10

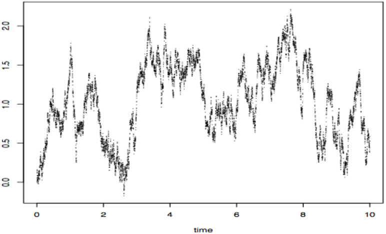

Variance Gamma process with shape parameter 5000 and scale parameter 100

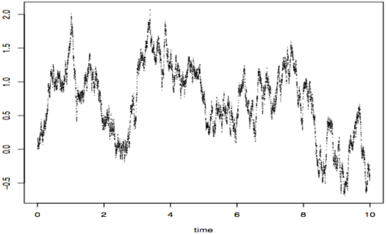

Variance Gamma process with shape parameter 5e+05 and scale parameter 1000

Fourier transforms

The Fourier transform of a sufficiently well-behaved function f on $\mathbb{R}$ is given by

$$\mathcal{F}(\omega) = \int_{-\infty}^{\infty} f(x) e^{-\mathrm{i}\omega x} dx,$$

and the inverse transform is then

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathcal{F}(\omega) e^{\mathrm{i}\omega x} d\omega.$$

If the function f is zero outside $[-R, R]$, then the first integral can be approximated by the compound trapezium rule:

$$\mathcal{F}(\omega) \approx \sum_{-N}^{N}{}'' e^{-\mathrm{i}\omega x_j} f(x_j) \Delta x, \qquad \Delta x = \frac{R}{N},$$

where $x_j = j\Delta x$, and the dashes indicate that the first and last terms are halved.

Fourier transforms

If we set $\omega_k = \frac{k\pi}{R}$, $k = -N, \ldots, N$, we can define

$$\widehat{f}_k = \sum_{-N}^{N}{}'' e^{-\mathrm{i}\omega_k x_j} f(x_j) \Delta x.$$

It turns out then that

$$f(x_j) = \sum_{-N}^{N}{}'' e^{\mathrm{i}\omega_k x_j} \widehat{f}_k.$$

These relations describe the forward and inverse **discrete** Fourier transforms. Note that $\widehat{f}_k$ can be thought of as approximations to $\mathcal{F}(\omega_k)$. However, if

$$f(x) = \sum_{-N}^{N}{}'' e^{\mathrm{i}\omega_k x} \widehat{f}_k$$

for all $x \in [-R, R]$ (and is zero outside), then they are *exact*.

The Fast Fourier Transform (FFT)

The most valuable numerical algorithm of our lifetime (Gil Strang)

The discrete Fourier transform can be written (in aliased form) as

$$f_k = \sum_{n=0}^{2N-1} F_n \omega^{nk}, \qquad j = 0, \ldots, 2N-1, \quad (\omega = \mathrm{e}^{\mathrm{i}\Delta})$$
$$= \sum_{n=0}^{N-1} F_{2n} (\omega^2)^{nk} + \omega^k \sum_{n=0}^{N-1} F_{2n+1} (\omega^2)^{nk}.$$

Each of these sums is a FFT of size N if we let $k = 0, \ldots, N-1$.
To produce the values at the other sample points we have

$$F_{k+N} = \sum_{n=0}^{N-1} F_{2n} (\omega^2)^{nk} - \omega^k \sum_{n=0}^{N-1} F_{2n+1} (\omega^2)^{nk},$$

because $\omega^N = -1$ and $\omega^{2N} = 1$.

The Fast Fourier Transform

We write the discrete Fourier Transform as $\mathbf{f} = A_{2N}\mathbf{F}$, with

$$A_{2N} = \begin{pmatrix} 1 & 1 & \dots & 1 \\ 1 & \omega & \dots & \omega^{2N-1} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & \omega^{2N-1} & \dots & \omega^{(2N-1)^2} \end{pmatrix}$$

Then we can factorize

$$A_{2N} = \begin{pmatrix} I & D_N \\ I & -D_N \end{pmatrix} \begin{pmatrix} A_N & 0 \\ 0 & A_N \end{pmatrix} \Big(\text{even-odd shuffle} \Big),$$

where D_n has the nth roots of -1 along the diagonal.
By continuing recursively, A_{2N} can be expressed as a product of $1 + \log_2 N$ sparse matrices (with only 2 non-zero elements in each row), and a single permutation matrix (bit-reversal).

The resulting operations count is $O(N \log_2 N)$.

Carr and Madan (1999)

Let $C_T(k)$ be the value of a European call option with maturity T and strike price K; set $k = \log K$, and let q_T be the risk-neutral density of $s_T = \log S_T$. Then

$$C_T(k) = \int_k^\infty e^{-rT}(e^s - e^k)q_T(s)ds.$$

Note that $C_T(k) \to S_0$ as $k \to -\infty$.
For $\alpha > 0$ define

$$c_T(k) = e^{\alpha k} C_T(k).$$

The Fourier transform $\psi_T(\omega)$ of $c_T(k)$ is given by

$$\psi_T(\omega) = \int_{-\infty}^\infty e^{-\mathrm{i}\omega k} c_T(k)ds,$$

and it can be computed as

Carr and Madan (1999)

$$
\begin{aligned}
\psi_T(\omega) &= \int_{-\infty}^{\infty} e^{-\mathrm{i}\omega k} \int_k^{\infty} e^{\alpha k} e^{-rT}(e^s - e^k) q_T(s) ds\, dk \\
&= \int_{-\infty}^{\infty} e^{-rT} q_T(s) \int_{-\infty}^{s} (e^{\alpha k + s} - e^{(\alpha+1)k}) e^{-\mathrm{i}\omega k} dk\, ds \\
&= \int_{-\infty}^{\infty} e^{-rT} q_T(s) \left(\frac{e^{(\alpha+1-\mathrm{i}\omega)s}}{\alpha - \mathrm{i}\omega} - \frac{e^{(\alpha+1-\mathrm{i}\omega)s}}{\alpha + 1 - \mathrm{i}\omega} \right) ds \\
&= \frac{e^{-rT} \phi_T(\omega + (\alpha+1)\mathrm{i})}{\alpha^2 + \alpha - \omega^2 - \mathrm{i}(2\alpha+1)\omega},
\end{aligned}
$$

where $\phi_T(\omega)$ is the characteristic function of s_T:

$$
\phi_T(\omega) = \mathbb{E}[e^{\mathrm{i}\omega s_T}].
$$

Carr and Madan (1999)

Then the inverse transform gives the option price:

$$C_T(k) = \frac{e^{-\alpha k}}{2\pi} \int_{-\infty}^{\infty} e^{\mathrm{i}\omega k} \psi_T(\omega) d\omega = \frac{e^{-\alpha k}}{\pi} \int_{0}^{\infty} e^{\mathrm{i}\omega k} \psi_T(\omega) d\omega.$$

If $\psi_T(\omega) \to 0$ as $\omega \to \infty$ sufficiently quickly, we may approximate this by

$$C_T(k) \approx \frac{\exp(-\alpha k)}{\pi} \sum_{j=0}^{N-1} e^{\mathrm{i}\omega_j k} \psi_T(\omega_j) \eta,$$

where $\omega_j = \eta j, j = 0, \cdots, N-1$, so that η is the distance between the grid points, and we are hoping that the integral beyond $N\eta$ is small.

The FFT then allows approximate option prices for a range of values of k to be computed efficiently.

Fourier-space time-stepping

Jackson, Jaimungal and Surkov (2007)

- Suppose that, under the pricing measure, $X(t) = \ln S(t)$ follows a Lévy process with triplet (γ, σ^2, ν).
- Let $V(t, S(t))$ be the price of the option at time t, with payoff $\phi(S)$ at time T, and define $v(t, x)$ via

$$v(t, x) = e^{-r(T-t)} V(t, S(0)e^x).$$

- Then v satisfies the PIDE

$$\begin{aligned}(\partial_t + \mathcal{L})v &= 0 \\ v(T, x) &= \phi(S(0)e^x),\end{aligned}$$

where $\mathcal{L}$ is the infinitesimal generator of the Lévy process:

$$\mathcal{L}f(x) = \lim_{t \to 0+} \frac{\mathbb{E}[f(x + X(t))] - f(x)}{t}.$$

Fourier-space time-stepping

- The advantages of Fourier methods in this context arise from the fact that if we take the Fourier transform of $\mathcal{L}v$ the characteristic exponent Ψ of the process $X(t)$ factors out:

$$\int_{-\infty}^{\infty} \mathcal{L}v(t,x)e^{\mathrm{i}\omega x}dx = \Psi(\omega)\hat{v}(t,\omega),$$

  where $\hat{v}(t,\omega)$ is the Fourier transform of $v(t,x)$.

- The PIDE reduces to

$$\frac{\partial \hat{v}(t,\omega)}{\partial t} + \Psi(\omega)\hat{v}(t,\omega) = 0,$$

  which is really a set of decoupled ODEs, so that

$$\hat{v}(t_1,\omega) = e^{(t_2-t_1)\Psi(\omega)}\hat{v}(t_2,\omega).$$

Fourier-space time-stepping

Initialization

- $\hat{v}(T,\omega)$ is given by the Fourier transform of the final condition:

$$\hat{v}(T,\omega) = \int_{-\infty}^{\infty} v(T,x) e^{\mathrm{i}\omega x} dx.$$

- If the payoff does not decay sufficiently, a device similar to that employed by Carr and Madan can be used: the payoff is multiplied by $e^{\epsilon x}$ for some suitably-chosen ϵ.
- Given sufficient decay, the integral may be truncated (to $[-R, R]$, say) and computed, approximately, for a range of equally-spaced $\omega_k = \frac{k\pi}{R}$ with the help of the FFT:

$$\hat{v}(T,\omega_k) = \frac{R}{N} \sum_{-N}^{N}{}'' e^{-\mathrm{i}\omega_k x_j}\, v(T,x_j), \quad x_j = \frac{jR}{N}.$$

Fourier-space time-stepping

- For a European-style option, we can obtain the option price at any earlier time by means of another Fourier transform. If we can assume the appropriate decay, then all of these transforms may be approximately computed, efficiently, using the FFT, whenever the underlying processes is of exponential Lévy form.
- For other styles of options, the above procedure can be used simply to evolve the option value function between timesteps. In this way, for example, American options may be valued (taking sufficiently small time steps).
- The authors apply the technique to barrier and spread options (since the FFT can be used in higher dimensions), and have also incorporated regime switching into the underlying asset model.

Lévy-driven processes

We are interested in extending the approach to cases where the underlying asset is *driven* by a Lévy process, but where neither the asset nor its logarithm are themselves Lévy. We consider processes of the form

$$dx(t) = a(t, x(t))dt + dL(t),$$

where $L(t)$ is a Lévy process. If a is constant, then $x(t)$ will again be a Lévy process. In general, however, neither $x(t)$ nor any function of it will be Lévy. Suppose that the above SDE represents the dynamics of a function $x(t) = f(t, S(t))$ of an asset $S(t)$ (with inverse $S(t) = g(t, x(t))$, and that—given an option value function $V(t, S)$—we define

$$v(t, x) = e^{-r(T-t)} V(t, g(t, x)).$$

Then v satisfies

$$\frac{\partial v}{\partial t} + a(t, x)\frac{\partial v}{\partial x} + \mathcal{L}v = 0$$
$$v(T, x) = \phi(g(T, x)).$$

Semi-Lagrangian time stepping

In order to deal with the non-constant coefficients in the drift term, we use semi-Lagrangian time stepping. We define $X(\tau;t,x)$ to be the solution to

$$\frac{\mathrm{d}}{\mathrm{d}\tau}X(\tau;t,x) = a\big(\tau, X(\tau;t,x)\big); \qquad X(t;t,x) = x.$$

Then

$$\frac{\partial v}{\partial t} + a(t,x)\frac{\partial v}{\partial x} = \frac{\mathrm{d}}{\mathrm{d}\tau}v(\tau, X(\tau;t,x))\bigg|_{\tau=t}.$$

Discretisation

We lay down a set of discrete times $T = t_N > t_{N-1} > \cdots > t_0 = 0$. This **directional derivative** may be approximated at t_n:

$$\frac{\mathrm{d}}{\mathrm{d}\tau}v(\tau, X(\tau;t,x))\bigg|_{\tau=t_n} \approx \frac{v(t_n,x) - v(t_{n+1}, X(t_{n+1};t_n,x))}{t_n - t_{n+1}}.$$

Semi-Lagrangian time stepping

Time-discrete equations

We use the shorthand $X_n(x) = X(t_{n+1}; t_n, x)$, and define the functions $v_n(x)$ recursively via $v_N(x) = v(t_N, x)$ and, with $\Delta t_n = t_{n+1} - t_n$,

$$\big(I - \Delta t_n \mathcal{L}\big) v_n(x) = v_{n+1}(X^n(x)).$$

Now we take Fourier transforms. We find, as before, that the Lévy characteristic exponent factors out and we obtain

$$\big(1 - \Delta t \Psi(\omega)\big)\widehat{v_n}(\omega) = \widehat{\phi_{n+1}}(\omega),$$

where we have defined $\phi_{n+1}(x) := v_{n+1}(X^n(x))$, and we assume for the moment we can evaluate this function straightforwardly. In this case, we have

$$\widehat{v_n}(\omega) = \frac{\widehat{\phi_{n+1}}(\omega)}{1 - \Delta t \Psi(\omega)}.$$

What is required, in order for this recursion to proceed, is a way of computing $\widehat{\phi_{n+1}}(\omega)$ from $\widehat{v_{n+1}}(\omega)$.

Semi-Lagrangian time stepping

Finding $\widehat{\phi_{n+1}}(\omega)$ from $\widehat{v_{n+1}}(\omega)$

Once $\phi_{n+1}(x_j)$ is known at equally-spaced samples x_j, $\widehat{\phi_{n+1}}$ may be calculated at equally-spaced samples of ω via the FFT. However, $\phi_{n+1}(x_j)$ involves evaluating v_{n+1} at the **unequally-spaced** points $X^n(x_j)$.

If $v_{n+1}(x) = \sum_{-N}^{N}{}'' e^{i\omega_k x}\, \widehat{v_{n+1}}(\omega_k)$ for all $x \in [-R, R]$, then

$$\phi_{n+1}(x_j) = \sum_{-N}^{N}{}'' e^{i\omega_k X^n(x_j)}\, \widehat{v_{n+1}}(\omega_k)$$

The problem here is that the points $X^n(x_j)$ are **not** equally-spaced, and computing this sum is going to be prohibitively expensive.
Fortunately, the non-uniform FFT[a] is available, making it possible to compute this term arbitrarily accurately at the cost of a small number of regular FFTs.

[a] See http://www-user.tu-chemnitz.de/~potts/nfft/ for more information.

Mean-reverting Variance-Gamma

We consider the situation where (under the pricing measure being used) the logarithm of the asset price follows the following SDE:

$$dx(t) = (a - bx(t))dt + dL(t),$$

where $L(t)$ is (here) a Variance-Gamma process, with characteristic exponent

$$\Phi(\omega) = -\frac{1}{\nu}\log\left(1 - \mathrm{i}\theta\nu\omega + \frac{1}{2}\sigma^2\nu\omega^2\right).$$

The parameters used for $L(t)$ are $\theta = -0.22898$, $\sigma = 0.20722$, $\nu = 0.50215$ (see Hirsa & Madan 2004). For the asset $S(t) = S(0)e^{x(t)}$ we have $S_0 = 1369.41$. The mean-reversion parameters are $a = 0.2$ and $b = 0.5$, and the risk-free interest rate $r = 0.0541$.

Mean-reverting Variance-Gamma

We price a put option (both American and European styles) with strike price $K = 1200$ and expiry time $T = 0.56164$. We use grids with M time steps and $2N$ sample points in x, between -10 and 10. The prices of European and American puts in this OU-VG case are below. The errors are computed with reference to a high-resolution computation.

Table: Prices of European puts under OU-VG

M	N	Price	Error	CPU time
250	256	30.62785840	1.002E-3	2.07
500	512	30.58442947	4.165E-4	4.99
1000	1024	30.59469280	8.105E-5	19.86
2000	2048	30.59565782	4.951E-5	94.25
4000	4096	30.59630475	2.837E-5	378.57

Mean-reverting Variance-Gamma

We price a put option (both American and European styles) with strike price $K = 1200$ and expiry time $T = 0.56164$. We use grids with M time steps and $2N$ sample points in x, between -10 and 10. The prices of European and American puts in this OU-VG case are below. The errors are computed with reference to a high-resolution computation.

Table: Prices of American puts under OU-VG

M	N	Price	Error	CPU time
250	256	37.05025901	7.035E-3	1.36
500	512	37.11888322	5.196E-3	5.08
1000	1024	37.28919849	6.319E-4	20.28
2000	2048	37.30675911	1.612E-4	83.27
4000	4096	37.31020341	6.894E-5	378.38

Mean-reverting Variance-Gamma

The relative errors in both cases are shown below.

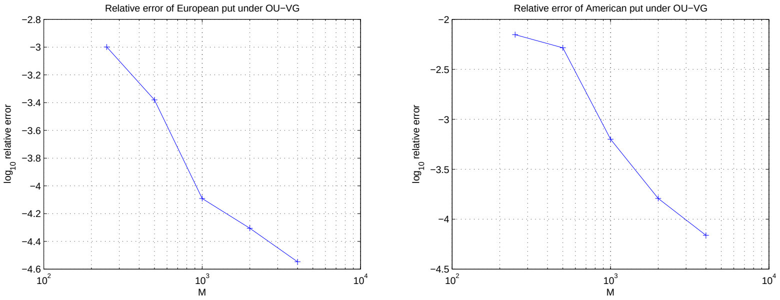

Figure: Relative error of European puts (left) and American puts (right) under OU-VG.

Mean-reverting Variance-Gamma

Below, we compare the price of the European puts and American puts in the left plot, from which we find that the price of an American put is always above the payoff. In the right plot, we compare the price of American puts with different maturities, and it shows that the price increases as the maturity increases.

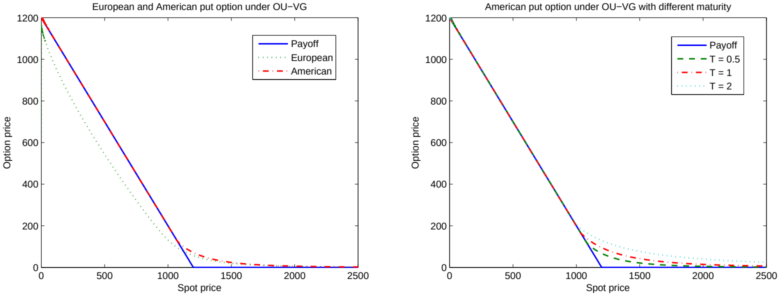

Figure: Left: Comparison of European and American puts under OU-VG.