Stepping Through Fourier Space

Article · October 2008

Vladimir Surkov
Royal Bank of Canada

Stepping Through Fourier Space

Sebastian Jaimungal
Department of Statistics and
Mathematical Finance Program,
University of Toronto, Toronto, ON, Canada
sebastian.jaimungal@utoronto.ca

Vladimir Surkov
Department of Computer Science,
University of Toronto, Toronto, ON, Canada
vsurkov@cs.toronto.edu

October 9, 2008

ABSTRACT

Diverse finite-difference schemes for solving pricing problems with Lévy underliers have been used in the literature. Invariably, the integral and diffusive terms are treated asymmetrically, large jumps are truncated, the methods are difficult to extend to higher dimensions and cannot easily incorporate regime switching or stochastic volatility. We present a new efficient approach which switches between Fourier and real space as time propagates backwards. We dub this method *Fourier Space Time-Stepping* (FST). The FST method applies to regime switching Lévy models and is applicable to a wide class of path-dependent options (such as Bermudan, barrier, shout and catastrophe linked options) and options on multiple assets.

KEY WORDS

Option Pricing, Lévy Processes, Regime Switching, Fourier Methods, American Options, Catastrophe Options.

Jump-diffusion, more generally Lévy models, have been extensively applied in practice to partially correct the defects in the Black and Scholes (1973) and Merton (1973) (BSM) model and explain the implied volatility surface's shape and dynamics. Popular models include the variance gamma (VG) model (Madan, Carr, and Chang 1998) and normal inverse Gaussian (NIG) model (Barndorff-Nielson 1997). Under such models, the pricing PDE transforms into a partial-integro differential equation (PIDE) with a non-local integral term. By writing the price as convolution and performing a transform on log-strike (Carr and Madan 1999) introduced an FFT method which prices European calls/puts with a series of strikes in "one shot".

For path dependent options or options with early exercise clauses, the full PIDE must be solved. A plethora of finite difference methods for solving these PIDEs have been proposed in the recent literature (see e.g. Andersen and Andreasen (2000), Briani, Natalini, and Russo (2004), Cont and Tankov (2004), and d'Halluin, Forsyth, and Vetzal (2005)), yet they all have many points in common: Firstly, the integral and diffusion terms of the PIDE are often treated separately – invariably, the integral term is evaluated explicitly in order to avoid solving a dense system of linear equations. Secondly, within infinite activity processes, the small jumps are approximated by a diffusion and incorporated into the diffusion term. Thirdly, the integral term is localized to the bounded domain of the diffusion term, i.e. large jumps are truncated. Finally, the separate treatment of diffusion and integral components requires that function values are interpolated and extrapolated between the diffusion and integral grids.

These factors together make finite difference methods for option pricing under jump models quite complex, and potentially prone to accuracy and stability problems, especially for path dependent claims or claims with early exercise clauses. In this work, we present a new Fourier Space Time-stepping (FST) algorithm which avoids the problems associated with finite difference methods by switching back and forth between Fourier and real space. Working in Fourier space transforms the PIDE into a system of easily solvable ordinary differential equations (ODE). The characteristic exponent of the underlying Lévy process naturally appears and since it is available, through the Lévy-Khintchine formula, in closed form for all independent increment processes, makes the FST method quite flexible and generic. The FST algorithm naturally leads to a symmetric treatment of the diffusion and jump terms, avoids stepping through time between observation dates, easily generalizes to higher dimensions and can incorporate regime switching rather simply.

The method, in its first incarnation together with extensive numerical experiments and further examples can be found in Jackson, Jaimungal, and Surkov (2007). Recently, Lord, Fang, Bervoets, and Oosterlee (2008) have independently worked along similar lines of thought without making the connection to the PIDE or extending to regime switching models and instead carried out an interesting analysis of the method's convergence.

1 The Model

Lévy models on their own do not fully capture the stylized features observed in stock returns nor are they capable of fitting the entire implied volatility surface. It is well known that stochastic volatility effects are necessary to model the mid/long term dynamics while jump components assist in fitting the short term dynamics. Stochastic time changes alá Carr and Wu (2002) is one approach for integrating Lévy processes and stochastic volatility models. They demonstrate that the Laplace transform of the time-changed process can still be obtained in closed form via a clever complex valued measure change. Regime switching models provide an enticing alternative for combining stochastic volatility and Lévy models. Elliott and Osakwe (2006) employ a Markov modulated pure jump process, determine its Laplace transform and proceed to value European styled options on a single underlier.

Motivated by these works, we are interested in modeling a vector of returns when the world switches between a finite number of regimes. Within each regime, returns are driven by a vector valued Lévy process incorporating both diffusive and jump components, while the regime changes may induce:

- a new covariance structure in the underlying diffusions – preserving variances but changing correlations, or changing variances while preserving correlations, or changing both variances and correlations;
- a new jump measure to incorporate regimes of high uncertain and a series of small market crashes – perhaps in one state the jump measure may be absent all together and in the other state the jump process is turned on and has infinite activity with large negative jumps;
- a new co-dependence in the jump structure – Lévy copulas have recently been introduced by Kallsen and Tankov (2006) which allow non-trivial co-dependence between jump components of two or more processes;
- new interest rate levels – to mimic stochastic interest rates;
- any combination of the above.

A typical sample path using a two-state regime switching VG model together with the underlying Lévy paths are depicted in Figure 1. In regime 1 (panel (a)) the jumps are of moderate size and symmetric, while in regime 2 (panel (b)) jumps are biased downwards and are larger. In the regime switching case (panel (c)) volatility and downward jumps cluster – non-stationary features that cannot be captured by Lévy models alone.

In our framework, the state of the world, enumerated $1, \ldots, K$, is driven by a Markov chain denoted Z_t with generator matrix $\mathbf{A}$. Conditional on the

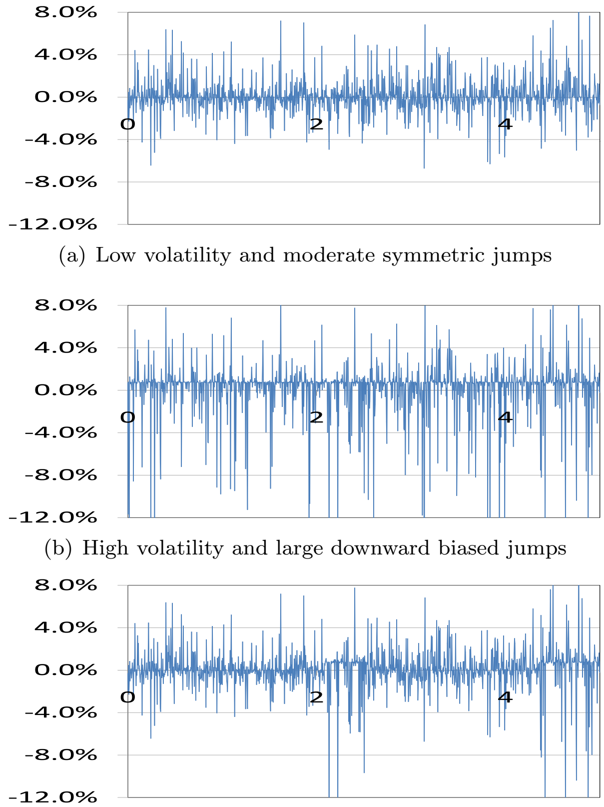

(a) Low volatility and moderate symmetric jumps

(b) High volatility and large downward biased jumps

(c) Regime Switching between models (a) and (b)

Figure 1: Sample returns from VG models both with and without regime switching. The top and middle panels are produced by the underlying VG models, while the bottom panel is produced by a Markov chain inducing switching between the two models.

world state, the d-dimensional vector of stock price log returns accumulate according to a d-dimensional Lévy process $\mathbf{X}_t^{(k)}$ (when state k prevails). Letting $(X_t)_i = \ln(S_t)_i$ denote the logarithm of the i^{th} stock price, the model can succinctly be written $d\mathbf{X}_t = d\mathbf{X}_t^{(Z_t)}$. In this manner, the logarithmic returns, rather than the prices themselves, are modulated by the Markov chain.

In the remainder of this article, the Lévy triple (which characterizes the diffusive and jump components) for state k is denoted $(\boldsymbol{\gamma}^{(k)}, \boldsymbol{\Sigma}^{(k)}, \boldsymbol{\nu}^{(k)})$. Recall that for a Lévy process, $\boldsymbol{\gamma}$ is the uncompensated drift vector, $\boldsymbol{\Sigma}$ is the covariance matrix of the Brownian pieces of the Lévy process, and $\boldsymbol{\nu}(d\mathbf{y} \times ds)$ is the Lévy density representing the rate of arrival of jumps of size $(y_1, y_1 + dy_1] \times \cdots \times (y_d, y_d + dy_d]$. It is sometimes useful to incorporate a time-dependent Lévy density – to mimic seasonality effects for example – for brevity, such simple generalizations are left for the interested reader to explore.

For derivative valuation, the drift vector $\boldsymbol{\gamma}$ must be pinned to risk-neutralize the price processes by setting $\Psi^{(k)}(-i\mathbf{1}) = r^{(k)} \times \mathbf{1}$ for each state k, where $r^{(k)}$ denotes the risk-free rate prevailing in world state k and $\Psi^{(k)}(\boldsymbol{\omega}) := \mathbb{E}^{\mathbb{Q}}[\exp\{i\boldsymbol{\omega} \cdot \mathbf{X}_t^{(k)}\}|\mathcal{F}_t]$ denotes the characteristic exponent of the Lévy process. Fortunately, the Lévy-Khintchine formula provides the characteristic exponent for any Lévy model:

$$\Psi^{(k)}(\boldsymbol{\omega}) = i\,\boldsymbol{\gamma}^{(k)} \cdot \boldsymbol{\omega} - \tfrac{1}{2}\,\boldsymbol{\omega} \cdot \mathbf{C}^{(k)} \cdot \boldsymbol{\omega} + \int_{\mathbb{R}^n/\{0\}} \left(e^{i\boldsymbol{\omega}\cdot\mathbf{y}} - 1 - i\mathbf{y} \cdot \boldsymbol{\omega}\ \mathbb{1}_{|\mathbf{y}|<1}\right) \boldsymbol{\nu}^{(k)}(d\mathbf{y}) \,.$$

This model framework is quite general and can produce an extremely wide variety of effects. Here are a few examples:

- Modeling defaultable equities by associating one of the states as a default state and making it absorbing;
- Modeling catastrophic losses which have both moderate and extreme epochs;
- Modeling equity price dynamics in which volatility / jumps cluster;
- Modeling currencies in which the positive and negative jumps are emphasized in different regimes.

It is not difficult to imagine other applications. We now move on to valuing derivatives based on underlying assets following such regime switching Lévy models.

2 The FST Algorithm

We now develop an algorithm for valuing various derivatives based on the above modeling assumptions. To this end, let $v^{(k)}(\mathbf{X}_t, t)$ denote the discounted price of a derivative instrument given that state k is prevailing at time t. The discounted price process is a martingale under the risk-neutral measure $\mathbb{Q}$; consequently, after applying Ito's lemma for

Lévy processes together with regime changes, the prices $v^{(k)}$ satisfy the system of *coupled* PIDEs:

$$\partial_t v^{(k)}(\mathbf{x},t) + (A_{kk} + \mathcal{L}^{(k)})v^{(k)}(\mathbf{x},t) + \sum_{j\neq k} A_{jk} v^{(j)}(\mathbf{x},t) = 0 , \quad (1)$$

where $\mathcal{L}^{(k)}$ is the infinitesimal generator of the multi-dimensional Lévy process in state k and acts on twice differentiable functions $f(\mathbf{x})$ as follows

$$\mathcal{L}^{(k)} f(\mathbf{x}) = \left(\boldsymbol{\gamma}^{(k)} \cdot \partial_{\mathbf{x}} + \tfrac{1}{2}\,\partial_{\mathbf{x}} \cdot \boldsymbol{\Sigma}^{(k)} \cdot \partial_{\mathbf{x}}\right) f(\mathbf{x}) + \int_{\mathbb{R}^n/\{\mathbf{0}\}} \big(f(\mathbf{x}+\mathbf{y}) - f(\mathbf{x}) - \mathbf{y}\cdot\partial_{\mathbf{x}} f(\mathbf{x})\, \mathbb{1}_{|\mathbf{y}|<1}\big)\, \boldsymbol{\nu}^{(k)}(d\mathbf{y}) \,.$$

As is well known, the various derivative instruments are valued by augmenting the PIDE with appropriate boundary conditions. For example, European options require specifying the terminal boundary only, Bermudan options in addition require an optimal exercise constraint at the exercise dates, knock-out Barrier options require truncating or modifying the value on the barrier monitoring dates, etc. . . .

It is important to note that (1) is a *coupled* system of PIDEs in multiple dimensions. In principle, any of the usual jump-modified finite-difference schemes can be adapted to this case. However, as discussed earlier, this is quite difficult due to the non-local integral terms and especially so for the present coupled multi-dimensional problem. Instead, we use transform methods to simplify the problem dramatically.

Define the Fourier transform as follows:

$$\mathcal{F}[h](\boldsymbol{\omega}) := \int_{\mathbb{R}^d} h(\mathbf{x})\, e^{i\boldsymbol{\omega}\cdot\mathbf{x}} d\mathbf{x} \quad (2)$$

Applying this operation to both sides of (1) reduces the coupled PIDEs to a simple coupled system of ODEs indexed by the vector of frequencies $\boldsymbol{\omega}$:

$$\partial_t \mathcal{F}[\mathbf{v}](\boldsymbol{\omega},t) + \boldsymbol{\Psi}(\boldsymbol{\omega})\mathcal{F}[\mathbf{v}](\boldsymbol{\omega},t) = \mathbf{0} \,.$$

where $\mathbf{v}$ denotes the vector of world state prices $(v^{(1)}, \ldots, v^{(K)})$ and the elements of the matrix $\boldsymbol{\Psi}(\boldsymbol{\omega})$ are $(\boldsymbol{\Psi}(\boldsymbol{\omega}))_{kl} = A_{kl} + \Psi^{(k)}(\boldsymbol{\omega})\,\delta_{kl}$. The solutions of these ODEs are straightforward, and taking the inverse transform leads to the simple exact algorithm between any two monitoring dates

$$\mathbf{v}(\mathbf{x}, t_{n-1}) = \mathcal{F}^{-1}\left[e^{\boldsymbol{\Psi}(\boldsymbol{\omega})\,\Delta t_n}\, \mathcal{F}[\mathbf{v}(\cdot, t_n)](\boldsymbol{\omega})\right](\mathbf{x}) \,.$$

Note that $e^{\boldsymbol{\Psi}(\boldsymbol{\omega})\,\Delta t_n}$ is a $K \times K$ matrix exponential and $\mathcal{F}[\mathbf{v}(\cdot, t_n)](\boldsymbol{\omega})$ is a K-dimensional vector for each d-dimensional frequency $\boldsymbol{\omega}$.

Introducing a grid in real space and Fourier space, letting $v^{(n)}$ denote the vector of prices at time t_n on the real grid, replacing the Fourier transform by the discrete Fourier transform and making use of the FFT algorithm to perform the forward and inverse transforms leads to the **FST method**:

$$\mathbf{v}^{(n-1)} = \mathsf{F}^{-1}\left[\exp\{\boldsymbol{\Psi}(\cdot)\,\Delta t_n\}\, \mathsf{F}[\mathbf{v}^{(n)}(\cdot)]\right] . \quad (3)$$

For Bermudan styled claims, the above price is viewed as the continuation (holding) value at time step t_{n-1}. Before iterating backwards, at each node in real space, it must be compared with the intrinsic value of the option, if the intrinsic value is larger, then that node is replaced by the intrinsic value. For American option, a penalty method can be used to augment the system of PIDEs which results in a modified version of the FST algorithm – this work is currently being written up. For barrier options, a simple application of the reflection principal neatly accounts for the barrier condition. Double barrier options can be treated by simply truncating the payoff, or alternatively, applying several iterations of the reflection scheme.

The FST scheme is related to the method employed by Carr and Madan (1999) for valuing European options. In their work, the authors performed a transform in the log-strike price to obtain prices of options of all strikes using an FFT operation. Here, however, we employ the FFT to obtain prices for all *log-spot levels* at intermediate decision dates. Another difference is that the adjustment factor which pushes the pole singularity for calls and puts off of the real axis is absent. This issue can be treated in two ways: (i) shift the integration path appropriately; or (ii) truncate the payoff for sufficiently large spot values. We find that truncating is a very simple and stable procedure which avoids the singularity problems; furthermore, we find that the results are insensitive to the truncation point for sufficiently large truncation.

The FST method is far superior to any finite difference inspired scheme for Bermudan or discretely monitored barriers, shout styled options, options on multiple asset among others. It can be applied to a wide class of Lévy processes – even incorporating Lévy copulas in a trivial manner, can deal with regime switches efficiently and we find it to be very stable to model parameters[1] and option payoffs.

It is more efficient precisely because between monitoring dates the FST method is exact (up to the truncation of very large frequencies) and requires only a single step from one date to the next. For example, a Bermudan monitored on a weekly basis will only require 50 time steps using the FST method, while any finite-difference scheme will require at least 250 time steps (assuming 1 step per day provides sufficient accuracy) – even without jumps. Furthermore, the jump measure is treated exactly, there is no necessity to approximate small jumps with a diffusion term, nor is it necessary to cut off large jumps.

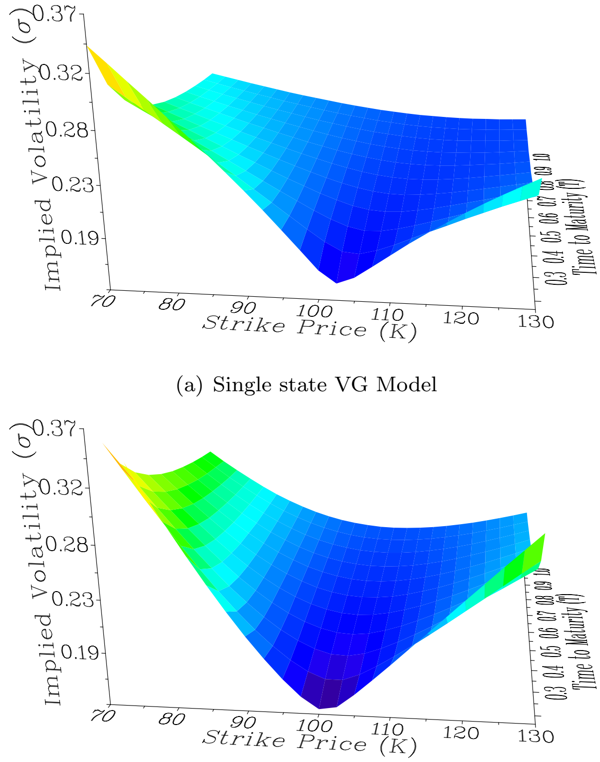

(a) Single state VG Model

(b) Two state VG model

Figure 2: Implied volatility surface using (a) a single state VG model and (b) a two-state regime switching VG model.

3 Examples

Lévy processes on their own have a difficult time capturing longer termed options – allowing regime switches rectifies this problem quite well. In Figure 2 we present the implied volatility surface obtained by using (a) a single-sate VG model (b) a two-state regime switching VG model which matches the short term smile of the single state model. The characteristic exponent within world-state $i = 1, 2$ is

$$\Psi^{(i)}(\omega) = -\frac{1}{\kappa^{(i)}} \log(1 - i\mu^{(i)}\kappa^{(i)}\omega + \frac{(\sigma^{(i)})^2\kappa^{(i)}\omega^2}{2})$$

Notice that the surface flattens out much more slowly in the two-state model than in the single state model. A feature which is difficult to capture using single state models.

As our next application, we turn to an American spread option with underliers driven by two-regime Merton jump-diffusion process. Specifically, the intrinsic value of the option is $(S_t^2 - S_t^1 - K)_+$ and the characteristic exponent is

$$\begin{aligned}\Psi^{(i)}(\omega_1, \omega_2) &= i(\mu_1 - \frac{\sigma_1^2}{2})\omega_1 + i(\mu_2 - \frac{\sigma_2^2}{2})\omega_2 \\ &\quad - \frac{\sigma_1^2\omega_1^2}{2} - \rho^{(i)}\sigma_1\sigma_2\omega_1\omega_2 - \frac{\sigma_2^2\omega_2^2}{2} \\ &\quad + \lambda_1 \, (e^{i\tilde{\mu}_1\omega_1 - \tilde{\sigma}_1^2\omega_2^2/2} - 1) \\ &\quad + \lambda_2 \, (e^{i\tilde{\mu}_2\omega_2 - \tilde{\sigma}_2^2\omega_2^2/2} - 1).\end{aligned}$$

Here, $\mu_{1,2}$ and $\sigma_{1,2}$ represent the mean and volatility of the diffusive components of each asset, $\tilde{\mu}_{1,2}$ and $\tilde{\sigma}_{1,2}$ represent the mean and variance of the jumps for each asset, and $\rho^{(i)}$ is the instantaneous correlation in regime-i ($i = 1, 2$). The parameter values for the spread option are $S_0^1 = S_0^2 = 100, K = 3, T = 1$ and $\sigma_1 = 0.19, \sigma_2 = 0.27, \lambda_1 = 10.0, \lambda_2 = 15.0, \tilde{\mu}_1 = 0.01, \tilde{\mu}_2 = -0.01, \tilde{\sigma}_1 = 0.017, \tilde{\sigma}_2 = 0.013, q_1 = q_2 = 0.01, r = 0.05$ for the stock price process. We allow for ρ to be in two different states, 0.3 and 0.8, where the regime-switching dynamics are driven by genera-

[1] Several finite difference schemes are tuned to specific models and fail under certain model parameters.

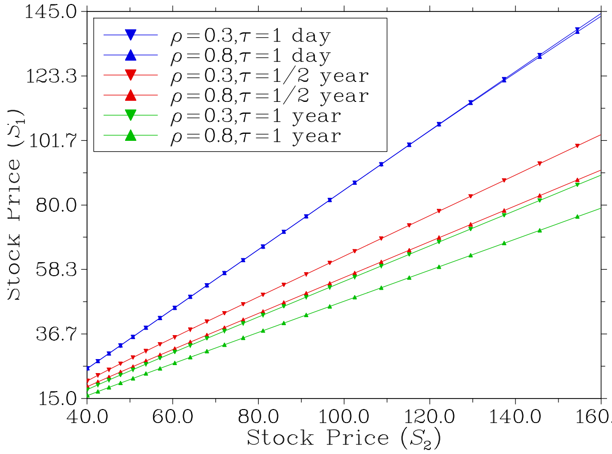

(a) Exercise policy

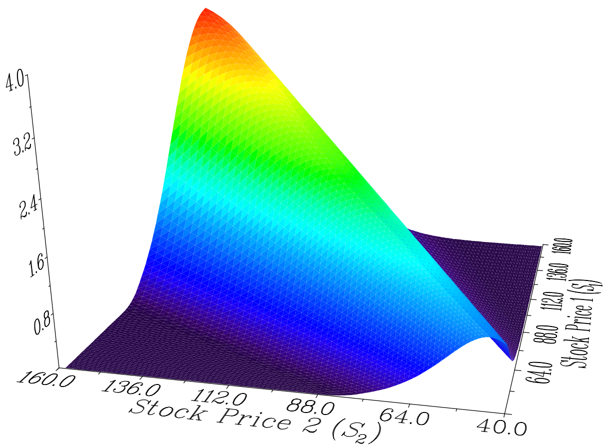

(b) Price difference in two regimes

Figure 3: Exercise slices (a) and price difference (b).

tor matrix $\mathbf{A} = [-0.2, 0.2; 0.9 - 0.9]$.

Figure 3 shows the effect of introduction of regime switching to stock price process. The exercise policy is dependent on the current state, as shown by exercise boundaries at $\tau = T/2$ and $\tau = T$ in panel (a). Near maturity, the exercise policy is independent of the state since the payoff depends only on terminal stock price. In panel (b), the difference in option prices between the low correlation state and high correlation state are provided. Notice that the spread option prices in high correlation state are lower since the difference between the two stock prices will typically be lower.

To further illustrate the versatility of the FST method, we value a non-standard path-dependent option known as a double-trigger stop-loss (DTSL) contract. The DTSL issues a stop-loss contract in the event that the insurers own share value falls below a critical level. Such options are issued as Bermudan contracts which can be exercised on a quarterly, monthly, weekly or daily basis. The exercise value of the contract can therefore be written $\varphi = \mathbb{I}(S < S^*)\,[(L - L_a)_+ - (L - L_d)_+]$ where S and L are the the share price and absorbed losses upon exercise.

Since large losses will drive share value down, the codependence between loss jumps and share value must be incorporated. Jaimungal and Wang (2006) introduced a model which jointly models interest rates, share value and losses and value a related European contract called a catastrophe equity put option. Here, we value the Bermudan DTSL product as well as its European counterpart. To keep things simple, interest rates are assumed constant. The joint stock price and loss dynamics is then

$$\begin{aligned} S_t &= S_0 \exp\{-\alpha L_t + \gamma\, t + \sigma\, W_t\}\ , \\ L_t &= \sum_{i=1}^{N_t} l_i\ . \end{aligned}$$

Here, N_t is a Poisson process which drives the arrival of losses and the random variables $l_1, l_2 \ldots$ are *iid* and gamma distributed with mean m and variance v. In this case, the characteristic exponent is

$$\begin{aligned} \Psi(\omega_1, \omega_2) &= i\,\gamma\,\omega_1 - \tfrac{1}{2}\sigma^2\,\omega_1^2 \\ &\quad + \lambda\left[\left[1 - \frac{i\,v}{m}(-\alpha\,\omega_1 + \omega_2)\right]^{-\frac{m^2}{v}} - 1\right]\ . \end{aligned}$$

In Figure 4, we provide the pricing surfaces for the European case, together with the daily exercisable Bermudan cases. The contract parameters are $S^* = 100$, $L_a = 5$ and $L_d = 40$. The joint stock price and loss parameters are $\sigma = 0.2, r = 0.05, m = 2, v = 5, \lambda = 1.25, \alpha = 0.005$.

Notice that the European case is a smoothed version of the payoff function. Contrastingly, the Bermudan price surfaces have smooth behavior when the stock price is above the trigger level of 100 while there is a distinct kink across this spot price level for losses above 5. Within this region, the DTSL

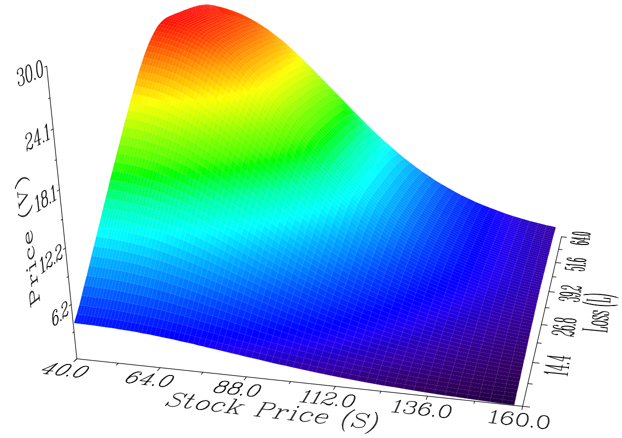

(a) European claim

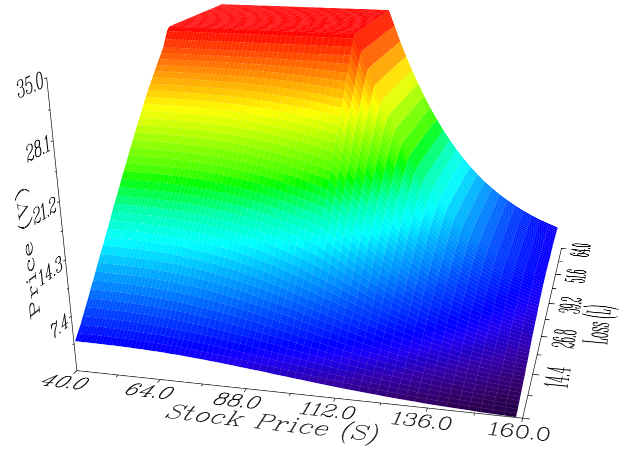

(b) Daily exercise

Figure 4: Price surfaces for three year double trigger stop-loss contracts assuming (a) exercise at maturity only and (b) daily exercise.

option is in-the-money. The behavior along the line $(S = 100, L > 5)$ reflects the discontinuous behavior of the intrinsic value of the option. These features are most easily explained and observed by investigating the optimal exercise policies.

Figure 5 depicts the optimal exercise behavior at different points in time (exercise occurs in the area above the exercise curve and to the left of the line $S = 100$). The exercise policy is governed by two competing factors. An investors is inclined to delay exercise to allow for the losses to accumulate and thus raise the option value at exercise. However, by waiting the investor is running the risk that the stock price will cross the threshold S^* and never venture below before maturity, thus rendering the option worthless. Since the latter risk is far greater than the former upside potential, the exercise boundary trends away from the threshold S^* as maturity approaches. Furthermore, the exercise curve flattens at around $L = 38$ for all maturities when $S < 95$. This can be explained by the discrete nature of loss arrival – since the maximal payoff is achieved at $L = 40$, it is optimal to exercise at $L = 38$ rather than waiting and possibly exercising at the same loss level if no additional losses arrive. In a separate experiment, as we increased the arrival rate of losses, this upper threshold approached $L = 40$.

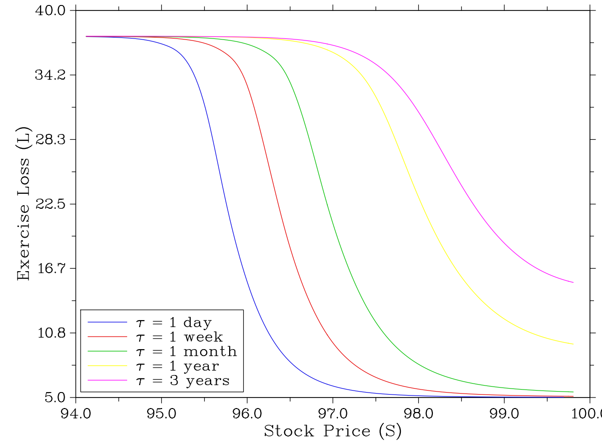

Figure 5: Exercise surfaces for three year double trigger stop-loss contracts assuming daily exercise rights.

4 Conclusions

We introduced a new pricing method for regime switching Levy models coined Fourier Space Time-Stepping and demonstrated its usefulness in several applications. There are many doors still open for exploration of these techniques such as currency options, interest rates, and higher dimensional options. We are also currently developing a modification appropriate for mean-reverting processes such as those appearing in Energy markets which will be useful in valuing swing and/or storage options with mean-reversion, jumps and regime changes without resorting to least-square Monte-Carlo.

5 Acknowledgements

The authors would like to thank the participants of the workshop on Actuarial Science and Financial Engineering at U. Michigan, the Computational Finance Conference at U. Waterloo, and the 2007 Daiwa International workshop on Financial Engineering in Tokyo for providing interesting feedback on various stages of this work. The Natural Sciences and Engineering Research Council of Canada partially funded this work.

References

Andersen, L. and J. Andreasen (2000). Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing. *Review of Derivatives Research 4*, 231–262.

Barndorff-Nielson, O. (1997). Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics 24*, 1–13.

Black, F. and M. S. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy 81*(3), 637–654.

Briani, M., R. Natalini, and G. Russo (2004). Implicit-explicit numerical schemes for jump-diffusion processes. IAC Report 38, Istituto per le Applicazioni del Calcolo IAC-CNR.

Carr, P. and D. B. Madan (1999). Option valuation using the fast Fourier transform. *The Journal of Computational Finance 2*(4), 61–73.

Carr, P. and L. Wu (2002). Time changed levy processes and option pricing. *Journal of Financial Economics forthcoming*.

Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes*. Chapman & Hall / CRC.

d'Halluin, Y., P. A. Forsyth, and K. R. Vetzal (2005). Robust numerical methods for contingent claims under jump diffusion processes. *IMA Journal of Numerical Analysis 25*, 87–112.

Elliott, R. and C. Osakwe (2006). Option pricing for pure jump processes with markov switching compensators. *Finance and Stochastics 10*, 250–275.

Jackson, K., S. Jaimungal, and V. Surkov (2007). Fourier space time-stepping for option pricing with lévy models. *Forthcoming in Special Issue of Journal of Computational Finance*.

Jaimungal, S. and T. Wang (2006). Catastrophe options with stochastic interest rates and compound Poisson losses. *Insurance: Mathematics and Economics 38(3)*, 469–483.

Kallsen, J. and P. Tankov (2006). Characterization of dependence of multidimensional Lévy processes using Lévy copulas. *Journal of Multivaraite Analysis 97*, 1551–1572.

Lord, R., F. Fang, F. Bervoets, and K. Oosterlee (2008). A fast and accurate FFT-based method for pricing early-exercise options under Lévy processes. *SIAM Journal on Scientific Computing 30*(4), 1678–1705.

Madan, D., P. Carr, and E. Chang (1998). The Variance Gamma process and option pricing. *European Finance Review 2(1)*, 79–105.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science 4*(1), 141–183.