Centre for Computational Finance and Economic Agents

Working Paper Series

WP003-06

Kyriakos Chourdakis

Switching Lévy models in continuous time: Finite distributions and option pricing

March 2006

Switching Lévy models in continuous time: Finite distributions and option pricing

Kyriakos Chourdakis

CCFEA, University of Essex, Colchester CO4 3SQ, United Kingdom
e-mail: kchour@essex.ac.uk
URL address: http://www.theponytail.net/

Abstract

This paper introduces a general regime switching Lévy process, and constructs the characteristic function in closed form. Correlations between the underlying Markov chain and the asset returns are also allowed, by imposing asset price jumps whenever a regime change takes place. Based on the characteristic function the conditional densities and vanilla option prices can be rapidly computed using the FFT. It is shown that the regime switching model has the potential to capture a wide variety of implied volatility skews. The paper also discusses the pricing of exotic contracts, like barrier, Bermudan and American options, by implementation of a quadrature method. A detailed numerical experiment illustrates the application of the regime switching framework.

1 Introduction

The implied volatility skew (or smile) has been long recognized as the pattern that summarizes the failures of the Black-Merton-Scholes (henceforth BMS) option pricing formula.[1] Using a set of observed option prices, one can invert the BMS formula and retrieve the volatility that would price each one of these contracts correctly. Typically, these volatilities vary substantially across different strike prices and maturity horizons, indicating that the simple geometric Brownian motion might not be sufficient to capture all features of options markets. A typical textbook argument would link the volatility skew with (risk neutral) return distributions that are substantially skewed and leptokurtic, contrasting the normal returns assumed in the BMS framework. Therefore, it appears that options markets and the study of the volatility skew could reveal significant information regarding the true dynamics of the underlying asset price. This information can be valuable, not only from a vanilla pricing and hedging point of view, but for the more daunting task of pricing and hedging exotic contracts.

The unique BMS price is an outcome of market completeness, which renders derivative contracts redundant securities. That is to say, in frictionless markets their payoff structure can be replicated using primitive securities, such as the underlying asset and the riskless bond. When attempting to capture the implied volatility patterns, research on extensions of the BMS framework has taken two distinctive approaches regarding completeness. In most cases, sacrificing market completeness would be equivalent to introducing extra stochastic factors, which cannot be dynamically hedged using the primitive securities alone.

The "local volatility" models of Derman and Kani (1994) and Dupire (1994)

[1] See for example Ghysels, Harvey, and Renault (1996) for an overview of the empirical stylized facts and their relationships to implied volatilities.

retain the market completeness and assume that the underlying price follows a diffusion similar to the BMS one but allowing for the volatility to depend explicitly on time and the underlying price

$$\frac{\mathrm{d}S}{S} = \mu \mathrm{d}t + \sigma(S, t)\mathrm{d}B,$$

where S denotes the underlying asset price, μ is the constant mean return and σ is a volatility function of the asset and time. It can be shown that such a specification leads naturally to a (binomial or trinomial) tree that exhibits inhomogeneous transition probabilities across its nodes. This approach permits flexible volatility structures, and, in principle, can be accurately calibrated to any volatility surface. Since the assumption of one source of uncertainty is maintained, all derivative contracts are redundant and can therefore be perfectly replicated and hedged. Unfortunately, the volatility function $\sigma(S, t)$ is not always sufficiently smooth, and often takes counterintuitive forms. In addition, the large number of parameters can easily lead to overfitting and instabilities. Therefore, the local volatility models have not been successful in producing forward curves, and thus do not offer acceptable prices for exotic contracts.

A different approach has been taken in the stochastic volatility and/or jump diffusion frameworks (see for example Merton 1976; Hull and White 1987; Heston 1993; Bates 1998; and Duffie, Pan, and Singleton 2000, inter alia). The assumption of a single source of randomness is now dropped and either the volatility assumes a stochastic form of its own, or price discontinuities in the form of randomly arriving jumps are introduced. In both cases market completeness is sacrificed, although in the case of stochastic volatility the market can be completed using a number of derivative contracts. Although such models do not offer the perfect calibration fit of local volatility models, they can compensate by offering a robust alternative, where all parameter values admit an intuitive interpretation.

In this paper we will consider a class of models that belong in the second class. The instantaneous asset log-returns are generated by one of N candidate Lévy processes. A Markov chain with N states will specify the process that generates the data at any point in time. Thus, the analysis takes place in a "regime switching" framework. Since the seminal papers of Hamilton (1989, 1990), regime switching models have been applied to virtually every economic and financial time series. They offer a robust, yet parsimonious, methodology to model variables with conditional distributions that evolve and change through time. The business cycle, the collapse of speculative bubbles, interest and exchange rate fluctuations, asset pricing with fundamental uncertainty, and the analysis of volatility regimes are but a small number of areas that have benefited from the regime switching approach.[2]

Despite their popularity as an econometric tool, regime switching models have not been used for the purpose of option pricing. This stems from the unavailability of closed form solutions for such specifications, except for the very restrictive two-state case.[3] This paper attempts to bridge this gap, by constructing the characteristic function of a general regime switching Lévy process. Based on the characteristic function, vanilla option prices can be rapidly computed using the FFT procedure of Carr and Madan (1999). Therefore, a regime switching model can be calibrated to a whole implied volatility surface in a

[2] For examples of the regime switching framework see Hamilton (1990), Hamilton and Susmel (1994), Driffill and Sola (1998), Veronesi (1999, 2001) and Ang and Bekaert (2002) inter alia.

[3] See Naik (1993) for the two-state model with Brownian innovations, and Konikov and Madan (2001) for the two-state model with Lévy innovations.

matter of minutes.

One of the well documented stylized facts in the stochastic volatility literature is the substantial degree of negative correlation between the asset returns and the volatility process. To accommodate for a correlation structure between the Markov chain and the asset log-returns, jumps are introduced whenever the chain switches states. The paper derives the characteristic function in this general case, allowing these jumps to be deterministic or stochastic. Thus, the regime switching model can be successfully calibrated to a wide range of asymmetric volatility skews.

There has been a surge of research in the are of pricing and hedging exotic contracts.[4] Typically, finite difference methods are employed, which numerically solve a partial (integro-) differential equation (PIDE), under the appropriate boundary conditions. In the presence of stochastic volatility, these PIDEs will be two-dimensional, putting a heavy burden on computing resources. In a regime switching framework a system of one-dimensional PIDEs has to be solved instead. This can substantially simplify the numerical complexity. In this paper we also discuss an alternative based on the QUAD method, recently proposed in Andricopoulos, Widdicks, Duck, and Newton (2003). The QUAD method uses the probability density rather than solving the PIDEs and is well suited for the pricing of barrier, Bermudan and American options.

The paper is organized in the following way: Section 2 lays out the underlying assumptions, presents the regime switching Lévy model and computes the characteristic function. Section 3 explains how the conditional moments and densities can be retrieved and discusses the pricing of vanilla and exotic contracts. A numerical example is presented in section 4 showing examples of

[4] A bird's eye view of a number of exotic contracts and their pricing can be found in Wilmott (2004).

volatility structures and exotic prices. Section 5 concludes.

2 Assumptions and the main results

This section lays down the underlying process for the asset price, and gives the main theoretical results of the paper. In particular, the asset returns are modelled as a regime switching Lévy process. A correlation structure between the asset and the regime process is installed, and the characteristic function of the log-returns is derived.

2.1 The regime structure

Consider a Markov chain s_t, taking values in $E_0 = \{1, 2, \ldots, N\}$. Denote the generator of s_t with $\mathbf{Q} = \{q(j, i)\}$, where $i, j \in E_0$. Also denote with $E = \{(j, i) : i, j \in E_0, i \neq j\}$ the set of all possible state transitions. The generator of the Markov chain will define the infinitesimal transition probabilities,

$$\begin{cases} \mathrm{P}(s_{t+\Delta} = j | s_t = i) = q(j, i)\Delta + o(\Delta), & \text{for } i \neq j \\ \mathrm{P}(s_{t+\Delta} = i | s_t = i) = 1 + q(i, i)\Delta + o(\Delta), & \text{otherwise} \end{cases}$$

The above structure implies the relationship $q(i, i) = -\sum_{j \neq i} q(j, i)$.

Also denote with $T^{(j,i)}$ the set of the (stopping) times, where these regime shifts take place, that is to say

$$T^{(j,i)} = \{t > 0 : s_{t-} = i, s_t = j\}$$

for all $(j, i) \in E$.

2.2 The Lévy structure

We will now turn into the mechanism that generates the underlying asset price, which we assume takes the form

$$S_t = \exp X_t$$

The log-price process X_t will be constructed from a collection of Lévy processes, as described below.

Consider a collection of independent Lévy processes, indexed by $i \in E_0$, say Y_t^i. The increments of the log-price process will switch between the N Lévy processes, depending on the state s_t

$$\mathrm{d}X_t = \mathrm{d}Y_t^{s_t}$$

Each one of the Lévy processes Y_t^i, is assumed to have a Lévy-Itô decomposition of the form

$$\mathrm{d}Y_t^i = \mu^i \mathrm{d}t + \sigma^i \mathrm{d}B_t + \int_{\mathbb{R}/\{0\}} z N^i(\mathrm{d}z, \mathrm{d}t)$$

In the above expression, μ^i is the drift and σ^i is the diffusion coefficient of the continuous path. For any Borel set $\Lambda \in \mathcal{B}(\mathbb{R}/\{0\})$, $N^i(\Lambda, t)$ is a Poisson random measure, with $\nu^i(\cdot)$ its associated Lévy measure, describing the discontinuities.

Now consider the conditional characteristic function of each process Y_i, as a function of the initial value y, and the time t

$$\psi^i(y, t) = \mathrm{E}^y \exp(\mathrm{i}uY_t^i)$$

with initial condition $\psi^i(y,0) = \exp(\mathrm{i}uy)$. By construction, this function will satisfy the forward Kolmogorov equation, given for instance in Øksendal (1991)

$$\partial_t \psi^i(y,t) = \mathscr{A}^i \psi^i(y,t)$$

The generator $\mathscr{A}^i$ is applied to the function $y \mapsto \psi^i(y,t)$ and is given by (see for example Applebaum, 2004, p. 139)

$$\begin{aligned}\mathscr{A}^i f(y) = \mu^i \partial_y f(y) + \frac{(\sigma^i)^2}{2} \partial_{yy} f(y) \\ + \int_{\mathbb{R}/\{0\}} (f(y+z) - f(y) - z\partial_y f(y) \mathrm{I}(|z| < 1)) \nu^i(\mathrm{d}z)\end{aligned}$$

Assuming a separation $\psi^i(y,t) = \exp(\mathrm{i}uy)\psi^i(t)$, with $\psi^i(0) = 1$, and applying the generator $\mathscr{A}^i$ yields

$$\partial_t \psi^i(t) = \psi^i(t)\phi^i(u)$$

or $\psi^i(y,t) = \exp(\mathrm{i}uy + t\phi^i(u))$, where $\phi^i(u)$ is the *characteristic exponent* or *Lévy symbol* associated with each process Y_t^i (Applebaum, 2004, p. 30). This is given by

$$\phi^i(u) = \mathrm{i}u\mu^i - \frac{(u\sigma^i)^2}{2} + \int_{\mathbb{R}/\{0\}} (e^{\mathrm{i}uz} - 1 - \mathrm{i}uz\mathrm{I}(|z| < 1))\nu^i(\mathrm{d}z)$$

Since $\mathrm{d}Y_t^i$ will represent an asset log-price increment, it is intuitive to re-parameterize

$$\mu^i = \tilde{\mu}^i - \frac{(\sigma^i)^2}{2} - \int_{\mathbb{R}/\{0\}} (e^z - 1 - z\mathrm{I}(|z| < 1))\nu^i(\mathrm{d}z)$$

With this parametrization, the expected instantaneous asset price growth will be equal to $\tilde{\mu}$. We assume that this integral is finite and that all Lévy processes possess at least the first two moments, or equivalently that $\int_{\mathbb{R}/\{0\}} (z^2 \vee z)\nu^i(\mathrm{d}z) <$

∞. Under the above substitution the Lévy generator can be written as

$$\mathscr{A}^i f(y) = \left(\tilde{\mu}^i - \frac{(\sigma^i)^2}{2}\right) \partial_y f(y) + \frac{(\sigma^i)^2}{2} \partial_{yy} f(y) \\ + \int_{\mathbb{R}/\{0\}} (f(y+z) - f(y) - (e^z - 1)\partial_y f(y))\nu^i(\mathrm{d}z) \quad (2.1)$$

2.3 The regime switching Lévy process

Recall that log-price increases follow the i-th Lévy process $\mathrm{d}X_t = \mathrm{d}Y_t^i$, given the regime $s_t = i$. We will now proceed in determining the characteristic function of X_t. Given $X_0 = x$ and $s_0 = i$, we denote the characteristic function

$$\psi(x,i,t) = \mathrm{E}^{(x,i)} \exp(\mathrm{i}uX_t)$$

For the regime switching model we have the following

Lemma 2.1. *The generator of the process X_t, conditional on $X_0 = x$ and $s_0 = i$, will be equal to*

$$\mathscr{A}^x f(x,i) = (q(i,i) + \mathscr{A}^i) f(x,i) + \sum_{j \neq i} q(j,i) f(x,j)$$

Proof. By definition, the generator of X will be equal to

$$\mathscr{A}^x f(x,i) = \lim_{\Delta \downarrow 0} \frac{\mathrm{E}^{(x,i)} f(X_\Delta, s_\Delta) - f(x,i)}{\Delta}$$

The expectation can be written, by conditioning on the future regime s_Δ, as

$$\mathrm{E}^{(x,i)} f(X_\Delta, s_\Delta) = \sum_{j\neq i} q(j,i)\mathrm{E}^{(x,i)} f(X_\Delta, j)\Delta \\ + (1 + q(i,i)\Delta)\mathrm{E}^{(x,i)} f(X_\Delta, i) + o(\Delta)$$

Now the conditional quantities $\mathrm{E}^{(x,i)} f(X_\Delta, j)$, for all $j = 1, \ldots, N$, can be retrieved by applying the generator $\mathscr{A}^i$, since the increment (over Δ) will follow the i-th Lévy process. Therefore,

$$\mathrm{E}^{(x,i)} f(X_\Delta, j) = f(x,j) + \mathscr{A}^i f(x,j)\Delta + o(\Delta)$$

Substituting, and collecting the $o(\Delta)$ terms gives

$$\mathrm{E}^{(x,i)} f(X_\Delta, s_\Delta) = f(x,i) + \left(\sum_{j\neq i} q(j,i) f(x,j) + \mathscr{A}^i f(x,i) \right) \Delta + o(\Delta)$$

By the definition of the generator, and passing to the limit $\Delta \downarrow 0$, yields the result. □

Having established the generator, we now turn to determining the characteristic function of the regime switching log-price. The characteristic function has to satisfy a *system* of forward Kolmogorov equations. We shall show that the solution of this system is given by a simple matrix exponential.

Theorem 2.2. *Assume that X_t follows the regime switching Lévy specification, as described above. The characteristic function of X_t is given by*

$$\mathrm{E}^{(x,\mathbf{\Pi})} \exp(\mathrm{i}uX_t) = \exp(\mathrm{i}ux) \cdot [\mathbf{1}' \cdot \exp(t \cdot \mathbf{\Phi}(u)) \cdot \mathbf{\Pi}] \tag{2.2}$$

where the matrix $\mathbf{\Phi}(u)$ has elements given by

$$[\mathbf{\Phi}(u)]_{i,j} = \begin{cases} q(i,i) + \phi^i(u), & \textit{if } j = i \\ q(j,i), & \textit{otherwise} \end{cases}$$

and $\mathbf{\Pi}$ is the initial regime distribution. Recall that the functions $\phi^i(u)$ are the conditional characteristic exponents of the i-th Lévy process.

Proof. Since we have N states, the functions $\{\psi(x,i,t), i = 1,\dots,N\}$ solve a system of forward Kolmogorov equations, namely

$$\partial_t \psi(x,i,t) = (q(i,i) + \mathscr{A}^i)\psi(x,i,t) + \sum_{j \neq i} q(j,i)\psi(x,j,t)$$

with initial conditions $\psi(x,i,0) = \exp(\mathrm{i}ux)$, for all $i \in E_0$. We conjecture a solution of the form $\psi(x,i,t) = \exp(\mathrm{i}ux)\psi(i,t)$.[5] Applying the corresponding generators produces the system

$$\partial_t \psi(i,t) = (q(i,i) + \phi^i(u))\psi(i,t) + \sum_{j \neq i} q(j,i)\psi(j,t)$$

This system can be written in matrix form as

$$\partial_t \psi(t) = \mathbf{\Phi}(u) \cdot \psi(t)$$

where $\psi(t) = (\psi(1,t) \cdots \psi(N,t))'$ and the matrix $\mathbf{\Phi}(u)$ has elements given by

$$[\mathbf{\Phi}(u)]_{i,j} = \begin{cases} q(i,i) + \phi^i(u), & \text{if } j = i \\ q(j,i), & \text{otherwise} \end{cases}$$

[5] Please note that we slightly abuse the notation here, by using ψ for both functions.

The initial condition is once more $\psi(i, 0) = 1$, for all $i \in E_0$. The solution of this differential equation is given by the matrix exponential,

$$\psi(t) = \mathbf{1}' \cdot \exp(t \cdot \mathbf{\Phi}(u))$$

Thus, given an initial regime distribution, $\mathbf{\Pi}$, and an initial log-price, $X_0 = x$, the characteristic function of the log-price process X_t will be given by

$$\mathrm{E}^{(x,\mathbf{s})} \exp(\mathrm{i}uX_t) = \exp(\mathrm{i}ux) \cdot [\mathbf{1}' \cdot \exp(t \cdot \mathbf{\Phi}(u)) \cdot \mathbf{\Pi}]$$

$\square$

2.4 The correlation structure

A number of studies indicate that asset prices can be strongly correlated with some underlying, possibly unobserved, state processes.[6] The most prominent example is the popular stochastic volatility model of Heston (1993), where a strong spot/volatility correlation appears to be necessary, in order to capture the familiar implied volatility skew. In Heston's framework, the implied volatility asymmetries are attributed to the pronounced skewness of asset returns. The *leverage effect* induced by this correlation is responsible for controlling the skewness. Of course, in the more general Lévy framework, skewness can be controlled by the appropriate Lévy measure.

In the analysis above, it was maintained that the driving Markov chain, s_t, is independent of the log-price process, X_t. In this subsection we relax this

[6] For instance, the typical asset/volatility correlation is documented in Pan (2002) within a standard stochastic volatility jump diffusion framework. In a detailed study Chernov, Gallant, Ghysels, and Tauchen (2003) investigate a number of different volatility specifications, some of them multi-factor. A stochastic volatility model with jumps, where the spot price and volatility jump together, is estimated in Eraker (2004).

assumption, in an attempt to introduce a correlation structure.

One approach of introducing correlations can be found in Naik (1993): every time the chain switches, that is to say on the stopping times $t \in \bigcup_{(i,j)\in E_0} T^{(j,i)}$, there is an associated deterministic jump in the log-price process. A regime shift from state i to j will induce a price jump of $J^{(j,i)} < \infty$.[7] Thus, changes of the chain will be accompanied by changes of the underlying asset.[8]

Here we will generalize this approach, and assume the jump to be of a random magnitude. We will assume that $\mathrm{E}J^{(j,i)} < \infty$, for all jumps $(j,i) \in E$

Once again, it is intuitive to re-parameterize the drift of the Lévy processes, by setting

$$\mu^i = \tilde{\mu}^i - \frac{(\sigma^i)^2}{2} - \int_{\mathbb{R}/\{0\}} (e^z - 1 - \mathrm{I}(|z|<1))\nu^i(\mathrm{d}z) \\ - \sum_{j\neq i} q(j,i)(\mathrm{E}\exp(J^{(j,i)}) - 1)$$

This will ensure that the coefficient $\tilde{\mu}$ will be equal to the instantaneous growth of the asset price.

The following lemma gives the generator of the regime switching log-price process described above.[9]

Lemma 2.3. *The generator of the process X_t, conditional on $X_0 = x$ and $s_0 = i$,*

[7] Although Naik describes this procedure of inducing a spot/volatility correlation, he only provides option prices for a simplified two-state model with zero correlations.

[8] This behavior is consistent with an equilibrium model where dividend growth is regime dependent. Switches of the price dividend ratio will cause the equilibrium prices to jump whenever the regime changes.

[9] The analysis of the corresponding jumps follows to some extend the methodology of Merton (1992, Ch. 5).

is given by

$$\mathscr{A}^x f(x,i) = (q(i,i) + \mathscr{A}^i) f(x,i) + \sum_{j \neq i} q(j,i) \mathrm{E} f(x + J^{(j,i)}, j)$$

where the expectation is taken with respect to each $J^{(j,i)}$.

Proof. Once again we use the definition of the generator

$$\mathscr{A}^x = \lim_{\Delta \downarrow 0} \frac{\mathrm{E}^{(x,i)} f(X_\Delta, s_\Delta) - f(x,i)}{\Delta}$$

and expand over the future regime s_Δ

$$\mathrm{E}^{(x,i)} f(X_\Delta, s_\Delta) = \sum_{j \neq i} q(j,i) \mathrm{E}^{(x,i)} f(X_\Delta, j) \Delta + (1 + q(i,i)\Delta) \mathrm{E}^{(x,i)} f(X_\Delta, i) + o(\Delta)$$

The quantity $\mathrm{E}^{(x,i)} f(X_\Delta, i)$, where no regime change takes place, follows from the definition of the generator, and is equal to

$$\mathrm{E}^{(x,i)} f(X_\Delta, i) = f(x,i) + \mathscr{A}^i f(x,i) \cdot \Delta + o(\Delta)$$

The quantities that incorporate regime switches, $\mathrm{E}^{(x,i)} f(X_\Delta, j)$ for $j \in E_0$, $j \neq i$, will be dependent on the associated jump $J^{(j,i)}$. In particular, we can write

$$\mathrm{E}^{(x,i)} f(X_\Delta, j) = \mathrm{E}^{(x,i)} f(x + J^{(j,i)}, j) + o(\Delta) = \mathrm{E} f(x + J^{(j,i)}, j) + o(\Delta)$$

where the last expectation integrates the jump size. Substituting in the definition of the generator and passing to the limit yields the result. □

The following theorem gives the characteristic function of the augmented model. Once again matrix exponentiation is sufficient to produce the characteristic function. Observe that the only difference occurs in the off-diagonal elements of the matrix function $\mathbf{\Phi}(u)$ which are now non-zero.

Theorem 2.4. *Assume that X_t follows the regime switching Lévy specification, as described above, together with its correlation structure. Then, the characteristic function of X_t is given by*

$$\mathrm{E}^{(x,\mathbf{\Pi})}\exp(\mathrm{i}uX_t) = \exp(\mathrm{i}ux)\cdot[\mathbf{1}'\cdot\exp(t\cdot\mathbf{\Phi}(u))\cdot\mathbf{\Pi}] \tag{2.3}$$

where the matrix $\mathbf{\Phi}(u)$ has elements given by

$$[\mathbf{\Phi}(u)]_{i,j} = \begin{cases} q(i,i)+\phi^i(u), & \text{if } j=i \\ q(j,i)\mathrm{E}\exp(\mathrm{i}uJ^{(j,i)}), & \text{otherwise} \end{cases}$$

and $\mathbf{\Pi}$ is the initial regime distribution. The quantities $\mathrm{E}\exp(\mathrm{i}uJ^{(j,i)})$ are the characteristic functions of the jump sizes.

Proof. Denote with $\{\psi(x,i,t), i\in E_0\}$ the conditional characteristic functions

$$\psi(x,i,t) = \mathrm{E}^{(x,i)}\exp(\mathrm{i}uX_t)$$

These satisfy the system of forward Kolmogorov equations

$$\begin{aligned}\partial_t\psi(x,i,t) &= \mathscr{A}^i\psi(x,i,t) + \sum_{j\neq i} q(j,i)(\psi(x,j,t)-\psi(x,i,t)) \\ &= (q(i,i)+\mathscr{A}^i)\psi(x,i,t) + \sum_{j\neq i} q(j,i)\psi(x,j,t)\end{aligned}$$

with initial conditions $\psi(x, i, 0) = \exp(\mathrm{i}ux)$, for all $i \in E_0$. We conjecture a solution of the form $\psi(x, i, t) = \exp(\mathrm{i}ux)\psi(i, t)$. Applying the corresponding generators, derived above, produces the system

$$\partial_t \psi(i, t) = (q(i, i) + \phi^i(u))\psi(i, t) + \sum_{j \neq i} q(j, i)\psi(j, t)\mathrm{E}\exp(\mathrm{i}uJ^{(j,i)})$$

This system can be written in matrix form as

$$\partial_t \psi(t) = \mathbf{\Phi}(u) \cdot \psi(t)$$

where $\psi(t) = (\psi(1, t) \cdots \psi(N, t))'$, and the matrix $\mathbf{\Phi}(u)$ has elements given by

$$[\mathbf{\Phi}(u)]_{i,j} = \begin{cases} q(i, i) + \phi^i(u), & \text{if } j = i \\ q(j, i)\mathrm{E}\exp(\mathrm{i}uJ^{(j,i)}), & \text{otherwise} \end{cases}$$

The result follows. □

Remark 2.5. Note that the (i, j)-th element of the matrix $\exp(t \cdot \mathbf{\Phi}(u))$ in equations (2.2) and (2.3) is the Fourier transform of the conditional quantity

$$\mathrm{P}(s_t = j | s_0 = i) f(y | X_0 = 0, s_0 = i, s_t = j) \tag{2.4}$$

where $f(y|\cdot) = \mathrm{P}(X_t \in \mathrm{d}y|\cdot)$ is the probability density of the log-return over a time interval t, conditional on the initial and final regimes.

3 Pricing under regime switching

Thus far, the processes of the previous sections describe the price evolution under the statistical (or objective) probability measure. In order to price derivative contracts, the risk neutral (or pricing) measure has to be established. We will assume, from this point onwards, that the parameter set of the regime switching process is specified under the risk neutral measure directly, and to this end we impose the restriction $\tilde{\mu} = r$. This will not affect the generality of the results, and it reflects the common practice of calibrating pricing models using derivative contracts alone. The appendix illustrates a setup based on the *Esscher transform* method of Gerber and Shiu (1994), which formally establishes the pricing measure, quantifies the various prices of risk, and shows how the parameter values are affected when switching between the two probability measures.

Having established the characteristic function of the log-return in the previous section, we now investigate a number of applications. In this section we first derive the moments of the general regime switching Lévy model in closed form and show how the conditional probability functions can be retrieved. Next, we set up the system of partial integro-differential equations that derivative prices will satisfy. The pricing of vanilla and exotic contracts is the topic of the last two subsections.

3.1 Moments and the risk neutral density

Simple differentiation of the characteristic function yields the t-period moments of any order. The main obstacle in our setting is differentiating the matrix exponential function. Mathias (1997) gives a simple procedure to carry out this differentiation analytically. The advantage of this method is that we can retrieve the derivatives of order $\{1, \ldots, k\}$ *simultaneously*, by computing a single $N(k+1) \times N(k+1)$ matrix exponential.

The following definition of the *block-upper triangular-block Toeplitz matrix* is needed: Given a sequence of $N \times N$ matrices $\mathbf{A}_0, \mathbf{A}_1, \ldots, \mathbf{A}_k$, denote with

$\mathscr{T}(\mathbf{A}_0, \mathbf{A}_1, \ldots, \mathbf{A}_k)$ the $N(k+1)\times N(k+1)$ block-upper triangular-block Toeplitz matrix, with (i, j)-block equal to the matrix $\mathbf{A}_{j-i}$ for $j \geq i$. As an example,

$$\mathscr{T}(\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2) = \begin{pmatrix} \mathbf{A}_0 & \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{0} & \mathbf{A}_0 & \mathbf{A}_1 \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_0 \end{pmatrix}$$

In our framework, given the matrix $\mathbf{\Phi}(u)$, in order to produce the k derivatives of the matrix exponential $\exp(t \cdot \mathbf{\Phi}(u))$, evaluated at $u = 0$, we need to take the following steps:

1. Compute the first k (element-wise) derivatives of the matrix $\mathbf{\Phi}(u)$, evaluated at $u = 0$, that is to say

$$\mathbf{\Phi}_j = \frac{\mathrm{d}^j}{\mathrm{d}u^j}\mathbf{\Phi}(u)\bigg|_{u=0}, \text{ for } j = 0, \ldots, k$$

2. Construct the Toeplitz matrix

$$\mathbf{\Phi}_{\mathscr{T}} = \mathscr{T}\left(\mathbf{\Phi}_0, \frac{1}{1!}\mathbf{\Phi}_1, \ldots, \frac{1}{k!}\mathbf{\Phi}_k\right)$$

as prescribed above.

3. Compute the matrix exponential

$$\mathbf{\Phi}_t^e = \exp(t \cdot \mathbf{\Phi}_{\mathscr{T}})$$

4. The $(1, j)$ block of $\mathbf{\Phi}_t^e$, denoted $\mathbf{\Phi}_{t,j}^e$, is equal to $\frac{1}{j!}$ times the derivative of the matrix exponential. Thus, to retrieve the matrix exponential derivatives,

compute

$$\frac{\mathrm{d}^j}{\mathrm{d}u^j}\exp(t\cdot\mathbf{\Phi}(u))\bigg|_{u=0} = j!\cdot\mathbf{\Phi}^e_{t,j},\ \text{for } j=0,\dots,k$$

In many practical applications the matrix $\mathbf{\Phi}_{\mathscr{T}}$ will be large but sparse. In such cases the matrix exponential can be computed rapidly using the ExpoKit routines[10] of Sidje (1998).

The raw moments (of order j) are therefore easily obtained as the derivatives of the characteristic function. Assuming, without loss of generality, that $X_0 = 0$, the moments are computed as

$$\mathrm{E}^{(0,\pi)}X_t^j = \frac{\mathrm{d}^j}{\mathrm{d}u^j}\mathbf{1}'\cdot\exp(t\cdot\mathbf{\Phi}(u))\cdot\pi\bigg|_{u=0} = j!\cdot\left[\mathbf{1}'\cdot\Phi^e_{j,t}\cdot\pi\right]$$

The centered moments follow in a straightforward fashion from the raw ones.

The conditional moments can be subsequently used to approximate the probability density function of the conditional log-returns. For example, Gram-Charlier and Edgeworth expansions, or members of the Pearson Type IV density family, can be used to match the theoretical moments (for details on the approximation methods see Kendal and Stuart 1977; implementations in a general option pricing framework are given in Jarrow and Rudd 1982, while density approximations for Garch models can be found in Duan, Gauthier, and Simonato 1999).

Alternatively, the risk neutral density can be retrieved by inverting the conditional characteristic function numerically. Numerical inversion of Laplace and Fourier transforms is treated extensively in Abate and Whitt (1992) and

[10] The ExpoKit routines provide Krylov subspace projection approximations of vectors of the form $\exp(\mathbf{W})\cdot\mathbf{u}$. These can be easily adapted to compute the characteristic function or the exponential of the Toeplitz matrix.

Abate, Choudhury, and Whitt (1999). Typically, a discretization of the characteristic function is employed, with the trapezoidal rule utilized to approximate the Fourier integral. Using the FFT or the fractional FFT (FRFT) can substantially speed up the computations, retrieving all density points in a single F(R)FT run. Details on the FRFT implementation can be found in Bailey and Swarztrauber (1991).

Note that following remark 2.5, if we invert each element of the matrix $\exp(t \cdot \mathbf{\Phi}(u))$ separately, we will obtain the weighted conditional probability densities of the log-returns

$$f(y,i,j) = \mathrm{P}(s_t = j | s_0 = i) \cdot \mathrm{P}(X_t - X_0 \in \mathrm{d}y | s_0 = i, s_t = j) \qquad (3.1)$$

for all regimes i and j. These conditional densities will be used extensively in section 3.4, where the QUAD method is introduced for the pricing of exotic options.

3.2 The system of partial integro-differential equations

The price of any derivative contract, $V(x,t)$, satisfies the Feynman-Kac formula, that is to say

$$\partial_t V(x,t) + \mathscr{A} V(x,t) - rV(x,t) = 0$$

In this relationship x denotes the log-price, t denotes the time, and $\mathscr{A}$ is the appropriate generator (under risk neutrality). r is the risk free rate of return. Therefore, for the one-state Lévy process (where the generator is given in equation 2.1), the derivative price satisfies the following partial integro-differential equation (see for example Chan 1999; and Cont and Voltchkova 2004)

$$\partial_t V(x,t) + \left(r - \frac{\sigma^2}{2}\right)\partial_x V(x,t) + \frac{\sigma^2}{2}\partial_{xx}V(x,t) + \int_{\mathbb{R}/\{0\}} (V(x+z,t) - V(x,t) - (e^z - 1)\partial_x V(x,t))\nu(\mathrm{d}z) = rV(x,t)$$

Partial integro-differential equations can be solved by finite difference or finite element methods, in a fashion similar to the numerical solution of PDEs.[11] Of course, the integral term complicates the analysis, since a quadrature will have to be implemented at every node of the grid.

Under the regime switching structure, a system of PIDEs will have to be solved. In particular, the Feynman-Kac formula takes the form

$$\begin{aligned}\partial_t V(x,t,i) &+ q(i,i)V(x,t,i) \\ &+ \left(r - \frac{\sigma^2}{2}\right)\partial_x V(x,t,i) + \frac{\sigma^2}{2}\partial_{xx}V(x,t,i) + \sum_{j\neq i} q(j,i)\mathrm{E}V(x+J^{(j,i)},t,j) \\ &+ \int_{\mathbb{R}/\{0\}} (V(x+z,t,i) - V(x,t,i) - (e^z-1)\partial_x V(x,t,i))\nu^i(\mathrm{d}z) \\ &= rV(x,t,i) \qquad (3.2)\end{aligned}$$

Although the numerical solution of the above system can be a computationally intensive task, it should not be as demanding as the numerical solution of the two dimensional PIDEs arising from a standard Lévy model with diffusive volatility. Nevertheless, in section 3.4 we investigate an alternative that is based on an integral recursive representation of (3.2), which offers a fast and robust method for the pricing of a variety of exotic option contracts.

[11] For a brief exposition of the numerical methods involved see Duffy (2005).

3.3 Vanilla option pricing

European plain vanilla calls and puts can be priced by inverting the characteristic function. Carr and Madan (1999) show that the (time-0) price of a call option with strike log-price k can be computed as the (one-sided) Fourier integral

$$C(k) = \frac{\exp(-\xi k)}{\pi} \int_0^\infty \exp(-\mathrm{i}ku)\eta(u;\xi)\mathrm{d}u \tag{3.3}$$

In the above expression, ξ is a parameter that controls the speed of the decay of the integrand, while the function $\eta(\cdot;\cdot)$ is given in terms of the characteristic function of the log-price $\phi(\cdot)$ as follows:

$$\eta(u;\xi) = \frac{\exp(-rt)\phi(u-\mathrm{i}(\xi+1))}{\xi^2+\xi-u^2+\mathrm{i}(2\xi+1)u}$$

Having established the characteristic function of the regime switching Lévy models in equations (2.2) and (2.3), computing the corresponding option prices is a straightforward exercise. Based on a set of equidistant abscissas $\mathscr{U} = \{u_1,\ldots,u_\ell\}$ one can produce a set of options for a set of equidistant strike log-prices $\mathscr{K} = \{k_1,\ldots,k_\ell\}$ by approximating the integral using a quadrature of the form:

$$\int_0^\infty \exp(-\mathrm{i}k_j u)\eta(u;\xi)\mathrm{d}u \approx \sum_{i=1}^{\ell} w_i \exp(-\mathrm{i}k_j u_i)\eta(u_i;\xi)\triangle u$$

The weights w_i implement the appropriate quadrature[12], while $\triangle u$ is the grid spacing.

Details on the derivation of (3.3) can be found in Carr and Madan (1999).

[12] For example a set of weights $\mathscr{W} = \{\frac{1}{2}, 1, 1, \ldots, 1, \frac{1}{2}\}$ implements the trapezoidal rule.

Choudakis (2005) shows how the above summation can be rewritten as

$$\sum_{i=1}^{\ell} \exp(-\mathrm{i} ij\bar{\alpha})\bar{\eta}_i$$

for some $\bar{\alpha}$ and $\bar{\eta}_i$, with $j = 1, \ldots, \ell$, and thus it can be rapidly computed using the fractional FFT (FRFT) procedure.

3.4 The QUAD procedure and pricing of exotics

The QUAD procedure of Andricopoulos et al. (2003) can be used to rapidly produce prices for exotic contracts, such as discretely monitored barrier or Bermudan options. This is achieved by numerically integrating the option payoffs over the risk neutral density. As pointed out in Andricopoulos et al. (2003), this strategy can be thought of as a "perfect" multinomial tree method. Since, in our regime switching setting, inversion of the characteristic function yields the risk neutral density, the QUAD procedure is particularly suited for the computation of exotic prices.

Denote with $V(x,t)$ the value of the derivative contract at time t, conditional on the underlying log-price being equal to x. The QUAD method is based on the recursive relationship

$$V(x,t) = \exp(-r\triangle t) \cdot \mathrm{E}\left(V(X_{t+\triangle t}, t+\triangle t) | X_t = x\right)$$

If we denote with $f(y)$ the log-return density over a time interval of $\triangle t$, that is $f(y) = \mathrm{P}(X_{\triangle t} - X_0 \in \mathrm{d}y)$, then we can write the above relationship as the integral

$$V(x,t) = \exp(-r\triangle t) \cdot \int_{\mathbb{R}} f(y-x) V(y, t+\triangle t) \mathrm{d}y$$

After constructing a log-price grid, the above integral is numerically computed using the trapezoidal rule over the grid.

When compared to the standard finite difference methods, QUAD typically evaluates the option values at a substantially smaller number of points. Figure 1 illustrates this important difference. A discretely monitored up-and-out call option is priced using finite differences (Crank-Nicolson) and the QUAD method. The underlying spot and strike prices are set to \$100, while the up-and-out barrier is equal to \$120. The maturity of option is one year, and the barrier is monitored four times over the life of the option (every 0.20 years). In order to retrieve the option price, the Crank-Nicolson method reconstructs the price surface at a large number of points between the monitoring dates, while the QUAD method updates the function on these dates alone. In addition, QUAD does not attempt to solve the sometimes cumbersome partial (integro) differential equation; instead it utilizes the risk neutral probability function.

In a regime switching setting we will have to keep track of the conditional (on the regime i) option values $V(x,t,i)$. If a model with N regimes is used, we can write the recursive relationship as

$$V(x,t,i) = \exp(-r\triangle t)\cdot \mathrm{E}\left(V(X_{t+\triangle t}, t+\triangle t, s_{t+\triangle t})|X_t = x, s_t = i\right)$$

which gives

$$\begin{aligned} V(x,t,i) = \exp(-r\triangle t)\cdot \sum_{j=1}^{N}\int_{\mathbb{R}} & V(y,t+\triangle t,j) \\ & \times \mathrm{P}(X_{t+\triangle t} - x \in \mathrm{d}y | s_t = i, s_{t+\triangle t} = j)\mathrm{P}(s_{t+\triangle t} = j|s_t = i) \end{aligned}$$

Fig. 1: Comparison of the QUAD and the finite difference methods. An up-and-out barrier call is considered, with maturity one year. The spot and strike prices are set to \$100, and the barrier is \$120. The option is monitored discretely every 0.20 years. The stock price follows a geometric Brownian motion with volatility 20% and the interest rate is 4%. The top graph illustrates part of the grid and the solution using finite differences; the bottom graph shows part of the values computed using the QUAD method.

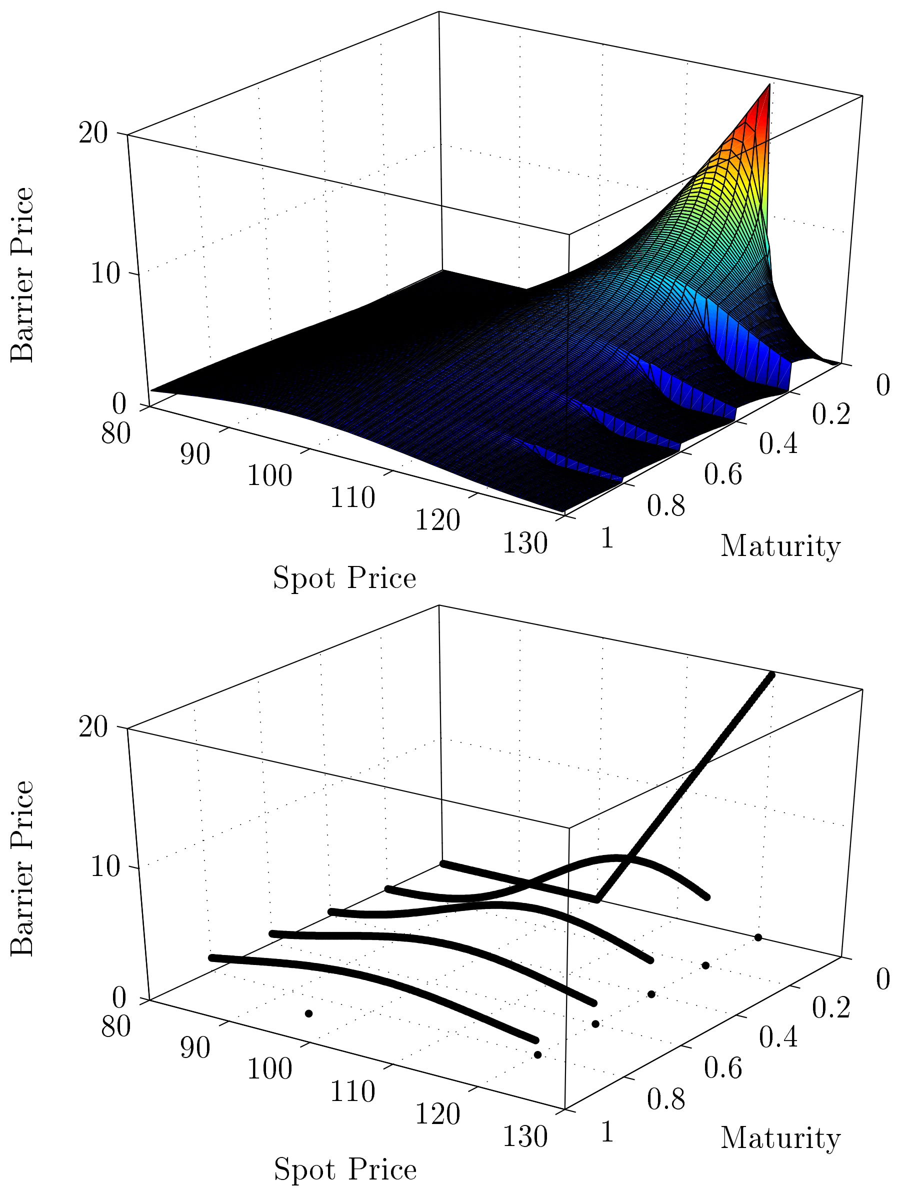

Then, using (2.4) yields

$$V(x,t,i) = \exp(-r\triangle t) \cdot \sum_{j=1}^{N} \int_{\mathbb{R}} V(y, t+\triangle t, j) f(y-x,i,j) \tag{3.4}$$

As pointed out in remark 2.5, the weighted conditional densities $f(y,i,j)$ can be retrieved by taking the inverse Fourier transform of the matrix exponential $\exp(t \cdot \mathbf{\Phi}(u))$ element-wise.[13]

4 A numerical example

In the numerical example that follows we will assume a two-state regime switching model, alternating between two Brownian motions. The interest rate is set at 4% and the volatilities for the two Brownian motions are set to 10% and 40%, respectively. The rate matrix of the underlying Markov chain is given by

$$\mathbf{Q} = \begin{pmatrix} -0.5 & 2.5 \\ 0.5 & -2.5 \end{pmatrix}$$

to reflect the higher persistence of the low volatility regime. Correlations between the volatility process and the return process are introduced via the jumps that occur, conditional on a regime change. For simplicity, we consider deterministic jumps, with magnitudes

$$\mathbf{J} = \begin{pmatrix} 0 & +2\% \\ -5\% & 0 \end{pmatrix}$$

[13] Note that since the conditional densities $f(y,i,j)$ are time-invariant, the Fourier inversions have to be carried out only once and the conditional densities can be stored.

Tab. 1: Conditional moments of the regime switching model. A two-state Brownian motion model is considered, with correlations introduced via conditional jumps. The time horizon is 0.25 years. The moments conditional on the high (40%) and low (10%) volatility regimes are presented.

	High Vol	Low Vol
Vol:	39.16%	23.12%
Skew:	-0.0275	-0.9053
Kurt:	+3.0645	+5.8631

This indicates, for example, that a volatility switch from 10% to 40% is accompanied by a return jump of -5%. For this experiment we consider a time horizon of 0.25 years.

4.1 The risk neutral density and moments

Table 1 presents the conditional moments of the above specification, while Figure 2 gives the conditional densities. The conditional moments were computed by differentiating the characteristic function, using the procedure outlined in section 3.1. The conditional densities were constructed using the fractional FFT procedure, as described in detail in Bailey and Swarztrauber (1991). An adaptive 128-point FRFT was used and the densities were retrieved in under a tenth of a second on a standard notebook computer.

It is apparent that the low volatility density is heavily skewed and leptokurtic. This observation is verified by the entries of Table 1: while conditional on the high volatility regime the return distribution appears to be close to normal, the distribution which is conditional on low volatility has a skewness value of -0.9, and kurtosis equal to 5.9. This is a result not only of the impact of the conditional jumps but also of the asymmetries in the rate matrix of the Markov chain.

Fig. 2: Conditional densities of the regime switching model. See Table 1 for details. The solid (dashed) line gives the probability density conditional on the low (high) initial volatility regime.

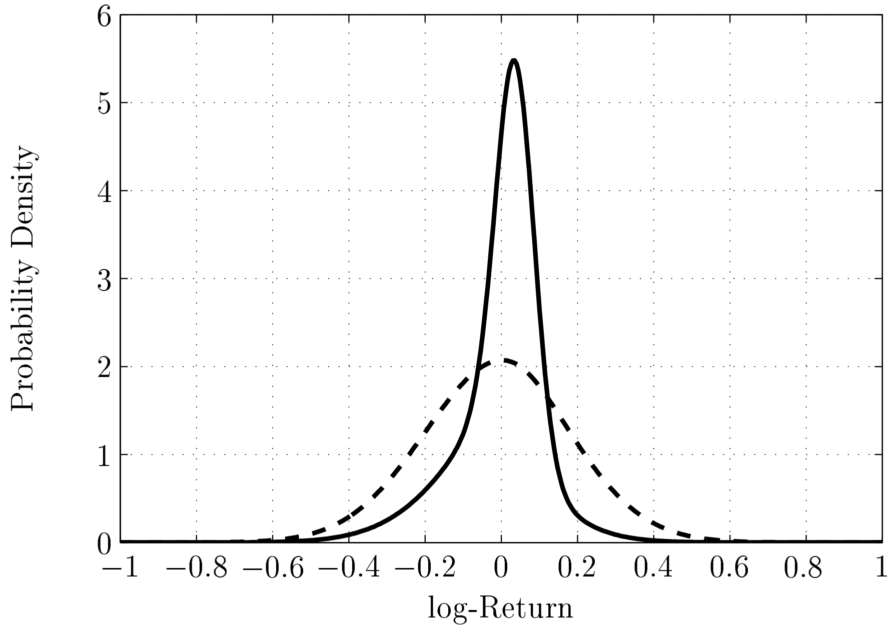

4.2 Vanilla prices

Based on the closed form characteristic function, we apply equation (3.3) to retrieve a set of call option prices, for different strike prices ranging from \$70 to \$130 and for maturities up to six months.

Figure 3 presents the implied volatility surface, conditional on a current low volatility regime. The volatility skews resemble the patterns encountered in equity and index vanilla option markets. In particular, the smile asymmetry is apparent, courtesy of the leverage effect which is captured by the regime-change dependent jumps. Deep in-the-money calls (and equivalently out-of-the-money puts) exhibit substantially higher implied volatilities. These results are in line with the highly skewed and leptokurtic conditional density of figure 2.[14]

[14] The implied volatility smile conditional on the high initial volatility regime, which is not presented here, appears to be fairly flat, ranging from 35% to 40%. This is in line with the

Fig. 3: The conditional implied volatility surface, based on the regime switching model described in section 4. The surface is conditional on the low volatility regime.

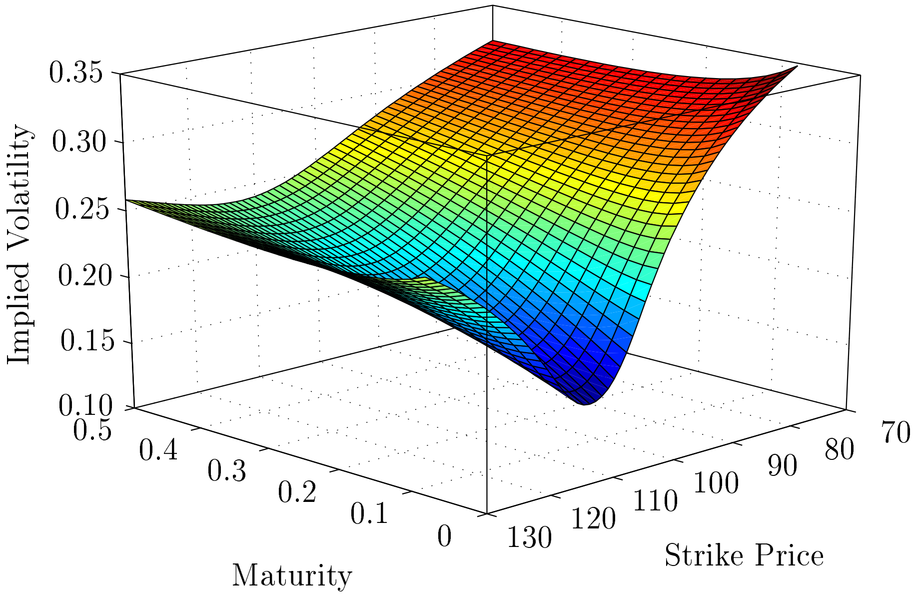

The prices were computed using an adaptive 128-point FRFT procedure, which sequentially integrated the support of the modified characteristic function (3.3).[15] A dampening parameter $\xi = 1.5$ was used throughout.

The regime switching model is able to produce flexible implied volatility surfaces and thus one can calibrate it accurately to a given set of option contracts. This flexibility can be significantly extended, if one augments the regime depended processes with a Lévy measure, incorporating a jump structure. This can offer a simple and robust alternative to stochastic volatility models and a more intuitive approach to the local volatility framework.

nearly normal conditional density, given in figure 2.

[15] The support of the modified characteristic function was split in subintervals of the form $[0, 50]$, $[50, 100]$, etc. The function was integrated over successive subintervals using FRFT, until the contribution to the option prices was insignificant.

Fig. 4: Barrier pricing under regime switching. For the model parameters described see section 4. The prices conditional on the high (resp. low) volatility regime are given in white (resp. black).

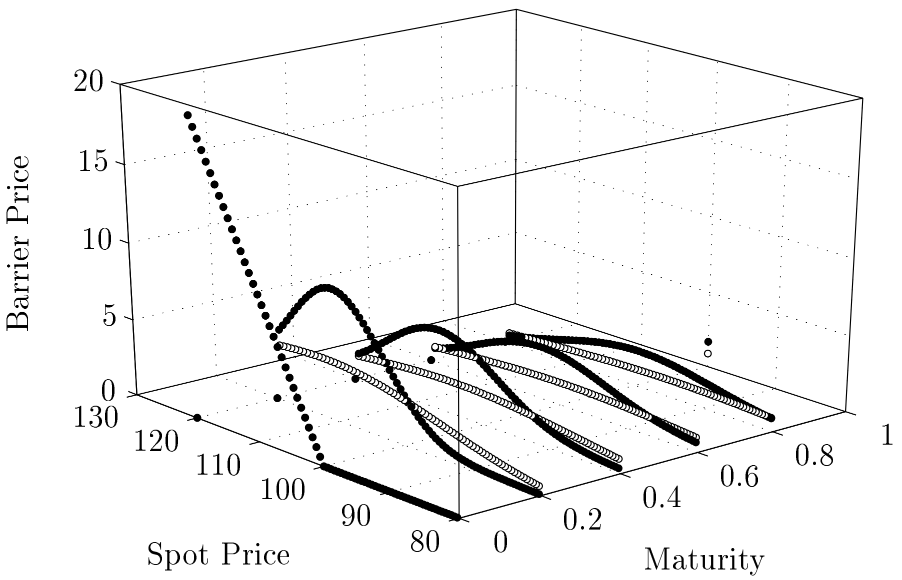

4.3 Exotic prices

In this subsection we will implement the QUAD method, as described in section 3.4, in order to price an up-and-out call with maturity of one year. We assume that the barrier is set at \$120 and that we monitor the barrier four times over the life of the option, that is to say every 0.20 years. The spot and strike prices are set at \$100.

Thus, in order to implement the QUAD method we need to numerically invert the weighted densities $f(y, i, j)$ over a set grid. We assume a dense log-price grid, with spacing $\frac{\sqrt{\triangle t}}{80} = 0.55\%$.[16] The trapezoidal rule was used for the numerical integration.

Barrier options were computed in a fraction of a second on a notebook. The

[16] Although Andricopoulos et al. (2003) suggest a lower value of $\frac{\sqrt{\triangle t}}{4}$ for the log-price grid, we found that such coarse discretization did not offer a high degree of accuracy.

resulting patterns are given in figure 4 and are intuitive. The QUAD method only computes the option prices at the monitoring points: prices given the high volatility regime are shown in red, while prices given the low volatility regime are shown in black.

One can observe that during high volatility periods the barrier prices are lower around the barrier, since the probability of breaching it is substantially higher. On the other hand, they are higher away from the barrier (out-of-the-money), since at these price levels the probability of returning in-the-money and exercising is relatively higher. Overall, the barrier option is more valuable in the low volatility regime.

With one year to maturity, the barrier options are \$0.90 for the high and \$1.70 for the low volatility regime. For comparison, if regime switches were not present, the barrier prices would have been \$0.75 and \$4.20, respectively. The impact of possible volatility changes is obvious, especially if we are currently in a low volatility period. This observation verifies the sensitivity of barrier contracts to the time variation of the underlying volatility.

Overall, the QUAD procedure appears to deliver fast and accurate results for the pricing of discretely monitored options. Andricopoulos et al. (2003) discuss in detail the QUAD implementation for the pricing of Bermudan options. They also show how one can extrapolate, in order to compute the prices of continuously monitored contracts, for example American options. Hedge parameters can be retrieved using finite differences to approximate the appropriate derivatives.

5 Conclusion

This paper introduces a regime switching Lévy model for the purpose of option pricing. The analysis is largely focused on the by-products of the conditional characteristic function, which is derived in closed form. Moments, the conditional densities, and derivative prices are all evaluated based on these conditional characteristic functions.

In particular, the paper illustrates how moments of all orders can be retrieved by numerically computing a simple matrix exponential, while densities are readily available via an application of an FFT. A regime switching approach offers not only an intuitive model and parameter interpretation but also densities that can accommodate a very wide range of skewness and kurtosis, providing the potential to fit a regime switching structure to the market at any point in time.

Vanilla call (and put) prices for a whole array of strike prices can be computed simultaneously, following the Carr and Madan (1999) procedure. This is a potentially important feature, since it allows one to calibrate the regime switching model to market prices capturing the observed volatility skew. The paper also shows how the QUAD method of Andricopoulos et al. (2003) can be employed for the pricing of standard exotic contracts, such as barrier or American options.

The purpose of this paper is largely to introduce the regime switching specification and to lay down the potential of this model. Detailed calibration exercises and the numerical analysis of the stability of the pricing methods are left for future research.

A Appendix: The pricing process and the market prices of risk

A.1 The Esscher transform

Given a process X_t that serves as a source of uncertainty or risk, one can define the function

$$M_t(x;\vartheta) = \frac{\exp(\vartheta x)}{\mathrm{E}\exp(\vartheta X_t)}$$

If we denote with $f_t(x)$ the pdf of X_t, then the function

$$f_t(x;\vartheta) = M_t(x;\vartheta) f_t(x)$$

is also a pdf and is call the *Esscher transform* of the original distribution with parameter ϑ. In that way equivalent probability measures to the original one can be constructed. Denote with $\mathbb{P}$ the statistical probability measure. Then, the Radon-Nikodym derivative of an equivalent measure $\mathbb{P}_\vartheta$, with respect to the statistical one is given by the martingale

$$\left.\frac{\mathrm{d}\mathbb{P}_\vartheta}{\mathrm{d}\mathbb{P}}\right|_t = M_t(X_t;\vartheta)$$

Thus, the Esscher transform generalizes the popular Girsanov's formula for general stochastic processes.

Consider an asset price that depends on this source of risk, say $S_t = S_t(X_t)$.[17] The *risk neutral Esscher transform* has a parameter ϑ^* such that

$$S_0 = \exp(-rt)\mathrm{E}_* S_t(X_t) = \exp(-rt)\mathrm{E}M_t(X_t;\vartheta)S_t(X_t) \tag{A.1}$$

[17] If $S_t = \exp X_t$, it can be easily shown that the function $M_t(x;\vartheta)$ will be the stochastic discount factor associated with a power utility function.

where E_* is the expectation taken under the probability measure $\mathrm{I\!P}_{\vartheta^*}$. The parameter ϑ^* that satisfies the above relationship can be though of as the price of risk associated with factor X_t.

If multiple *independent* risk factors are present, say $X_{1,t}, \ldots, X_{\ell,t}$, we can construct equivalent measures by sequentially applying the Esscher transform:

$$\left.\frac{\mathrm{d}\mathrm{I\!P}_{\vartheta_1,\ldots,\vartheta_\ell}}{\mathrm{d}\mathrm{I\!P}}\right|_t = \left.\frac{\mathrm{d}\mathrm{I\!P}_{\vartheta_1,\ldots,\vartheta_\ell}}{\mathrm{d}\mathrm{I\!P}_{\vartheta_1,\ldots,\vartheta_{\ell-1}}}\right|_t \cdot \left.\frac{\mathrm{d}\mathrm{I\!P}_{\vartheta_1,\ldots,\vartheta_{\ell-1}}}{\mathrm{d}\mathrm{I\!P}_{\vartheta_1,\ldots,\vartheta_{\ell-2}}}\right|_t \cdots \left.\frac{\mathrm{d}\mathrm{I\!P}_{\vartheta_1,\vartheta_2}}{\mathrm{d}\mathrm{I\!P}_{\vartheta_1}}\right|_t \cdot \left.\frac{\mathrm{d}\mathrm{I\!P}_{\vartheta_1}}{\mathrm{d}\mathrm{I\!P}}\right|_t$$
$$= M_t(X_{\ell,t};\vartheta_\ell) \cdot M_t(X_{\ell-1,t};\vartheta_{\ell-1}) \cdots M_t(X_{2,t};\vartheta_2) \cdot M_t(X_{1,t};\vartheta_1)$$

We will employ this strategy in order to characterize the pricing measure in our regime switching Lévy setting. In particular, we have the following sources of risk to consider:

- The regime risk associated with the Markov chain s_t.
- The jump risk associated with the correlation inducing jumps $J^{(j,i)}$.
- The Lévy risk, associated with the processes Y_t^i.

In the following subsections we separately consider each source of randomness and investigate its implications on the risk neutral process.

A.2 Regime risk

Exponential changes of measure for continuous time Markov chains are extensively discussed in Rolski, Schmidli, Schmidt, and Teugels (1999, 12.3). In particular, given a set of Esscher parameters $\{\vartheta_1, \ldots, \vartheta_N\}$, each associated with a regime, the rate matrix under the equivalent measure will have elements given by

$$q_*(j,i) = q(j,i)\exp(\vartheta_j - \vartheta_i)$$

It follows from the above relationship that $q_*(j,i) > q(j,i)$, that is to say a transition from i to j is more likely under risk neutrality, iff $\vartheta_j > \vartheta_i$. This is an intuitive result: If we consider ϑ_k as a proxy of the risk aversion of regime k, the relationship $\vartheta_j > \vartheta_i$ would imply that market participants "dislike" regime j more than regime i. Therefore, they would behave *as if* regime switches towards j are more likely than they truly are.

A.3 Jump risk

If the jumps associated with a regime shift from i to j, that is $J^{(j,i)}$, are of random magnitude, then their Esscher transform will imply a Radon-Nikodym derivative of the form

$$\frac{\exp(\vartheta^{(j,i)} J^{(j,i)})}{\mathrm{E}\exp(\vartheta^{(j,i)} J^{(j,i)})}$$

under the assumption that the expectation exists. If the jumps are deterministic, as for example in Naik (1993), then their sizes will be identical under the two measures. Nevertheless, their frequency will not be the same, since the intensities of regime switches under risk neutrality can be higher or lower, as discussed in the previous section.

A.4 Lévy risk

Esscher transforms of Lévy processes are discussed in detail in Chan (1999). In particular, we fix the Esscher parameter ϑ^i associated with the Lévy process Y_t^i. Then, the Radon-Nikodym derivative that corresponds to this process is given by

$$\exp(-\vartheta^i Y_t^i + t\varphi^i(\vartheta^i))$$

where $\varphi^i(\vartheta^i) = -\log \mathrm{E}\exp(-\vartheta^i Y_1^i)$. Based on the characteristic function of Y_t^i we can express $\varphi^i(\cdot)$ as

$$\varphi^i(\vartheta^i) = \vartheta^i \mu^i - \frac{(\vartheta^i \sigma^i)^2}{2} - \int_{\mathbb{R}/\{0\}} (e^{-\vartheta^i z} - 1 + \vartheta^i z \mathrm{I}(|z| < 1))\nu^i(\mathrm{d}z)$$

Chan (1999) gives the derivation details and the appropriate parameter conditions. The parameter ϑ^i can be thought of as a proxy of the risk aversion of the market participants, associated with price shocks that are due to the corresponding Lévy process, under a power utility framework.

Having established the Radon-Nikodym derivatives for all risk sources, we identify all equivalent probability measures in terms of the Esscher parameters $\vartheta_* = \{\vartheta_1, \ldots, \vartheta_N, \vartheta^1, \ldots, \vartheta^N, \vartheta^{(1,2)}, \ldots, \vartheta^{(N-1,N)}\}$. Pricing measures will satisfy the martingale restriction (A.1), and thus we have to select the Esscher parameters that ensure this.

If we do not have any particular reason to assume that risk aversion itself is regime dependent, we can set the Esscher parameters for the Lévy processes equal, that is to say $\vartheta^1 = \vartheta^2 = \ldots = \vartheta^N = \vartheta$. In addition, we can also equate the Esscher parameters that are associated with the jump amplitudes, $\vartheta^{(1,2)} = \ldots = \vartheta^{(N-1,N)} = \vartheta$. This can substantially simplify the conversion from the objective to the pricing probability measure.

References

Abate, J., G. L. Choudhury, and W. Whitt (1999). An introduction to numerical transform inversion and its application to probability models. In W. Grassman (Ed.), *Computational Probability*, pp. 257–323.

Abate, J. and W. Whitt (1992). The Fourier-series method for inverting transforms of probability distributions. *Queueing Systems 10*(1), 5–88.

Andricopoulos, A. D., M. Widdicks, P. W. Duck, and D. P. Newton (2003). Universal option valuation using quadrature methods. *Journal of Financial Economics 67*, 447–471.

Ang, A. and G. Bekaert (2002). Regime switches in interest rates. *Journal of Business and Economic Statistics 20*(2), 163–182.

Applebaum, D. (2004). *Lévy Processes and Stochastic Calculus.* Cambridge Studies in Advanced Mathematics. Cambridge, UK: Cambridge University Press.

Bailey, D. H. and P. N. Swarztrauber (1991). The fractional Fourier transform and applications. *SIAM Review 33*(3), 389–404.

Bates, D. S. (1998). Pricing options under jump diffusion processes. Technical Report 37/88, The Wharton School, University of Pennsylvania.

Carr, P. and D. Madan (1999). Option valuation using the Fast Fourier Transform. *Journal of Computational Finance 3*, 463–520.

Chan, T. (1999). Pricing contingent claims on stocks driven by Lévy rocesses. *The Annals of Applied Probability 9*(2), 504–528.

Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen (2003). Alternative models for stock price dynamics. *Journal of Econometrics 116*, 225–257.

Chourdakis, K. (2005). Option pricing using the Fractional FFT. *Journal of Computational Finance 8*(2), 1–18.

Cont, R. and E. Voltchkova (2004). Integro-differential equations for option prices in exponential Lévy models. Technical Report R.I. N^0 547, Centre de Mathématiques Appliquées, Ecole Polytechnique.

Derman, E. and I. Kani (1994). Riding on a smile. *RISK 7*, 32–39.

Driffill, J. and M. Sola (1998). Intrinsic bubbles and regime switching. *Journal of Monetary Economics 42*(2), 357–73.

Duan, J.-C., G. Gauthier, and J.-G. Simonato (1999). An analytical approximation for the Garch option pricing model. *Journal of Computational Finance 2*(4), 75–116.

Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump–diffusions. *Econometrica 68*, 1343–1376.

Duffy, D. J. (2005). Numerical analysis of jump diffusion models: A partial differential equation approach. Technical report, Datasim, http://www.datasim-component.com.

Dupire, B. (1994). Pricing with a smile. *RISK 7*, 18–20.

Eraker, B. (2004). Do stock prices and volatility jump? reconciling evidence from spot and option prices. *Journal of Finance 54*(3), 1367–1403.

Gerber, H. U. and E. S. W. Shiu (1994). Option pricing by Esscher transforms (with discussion). *Transactions of Society of Actuaries 46*, 99–191.

Ghysels, E., A. Harvey, and E. Renault (1996). Stochastic volatility. In G. Maddala and C. Rao (Eds.), *Handbook of Statistics, 14, Statistical Methods in Finance*. North Holland.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica 57*, 357–384.

Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics 45*, 39–70.

Hamilton, J. D. and R. Susmel (1994). Autoregressive conditional heteroscedasticity and changes in regime. *Journal of Econometrics 64*, 307–333.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies 6*, 327–344.

Hull, J. C. and A. White (1987). The pricing of options with stochastic volatilities. *The Journal of Finance 42*, 281–300.

Jarrow, R. and A. Rudd (1982). Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics 10*, 347–349.

Kendal, M. and A. Stuart (1977). *The Advanced Theory of Statistics.* (4th ed.), Volume I. London, U.K.: Charles Griffin and Co.

Konikov, M. and D. Madan (2001). Stochastic volatility via Markov chains. Technical report, Robert H. Smith School of Business, University of Maryland.

Mathias, R. (1997). A chain rule for matrix functions and applications. SIAM *Journal on Matrix Analysis and Applications 17*(2), 610 – 620.

Merton, R. (1976). Option pricing when the underlying stock returns are discontinuous. *Journal of Financial Economics 4*, 125–144.

Merton, R. C. (1992). *Continuous time finance.* Oxford, U.K.: Blackwell Publishing.

Naik, V. (1993). Option valuation and hedging strategies with jumps in the volatility of asset returns. *The Journal of Finance 48*, 1969–1984.

Øksendal, B. (1991). *Stochastic Differential Equations.* (3rd ed.). New York, NY: Springer-Verlag.

Pan, J. (2002). The jump risk-premia implicit in options: Evidence from an integrated time series study. *Journal of Financial Economics 63*, 3–50.

Rolski, T., H. Schmidli, V. Schmidt, and J. Teugels (1999). *Stochastic Processes for Insurance and Finance.* New York, NY: John Wiley and Sons.

Sidje, R. B. (1998). ExpoKit: A software package for computing matrix exponentials. *ACM Transactions of Mathematical Software 24*(1), 130–156.

Veronesi, P. (1999). Stock market overreaction to bad news in good times. a rational expectations equilibrium model. *Review of Financial Studies 12*(5), 975–1007.

Veronesi, P. (2001). Belief-dependent utilities, aversion to state-uncertainty and asset prices. University of Chicago, Manuscript.

Wilmott, P. (2004). *Paul Wilmott on quantitative finance.* Chichester, UK: John Wiley and Sons.