

Modelling spikes and pricing swing options in electricity markets

Ben Hambly

Sam Howison

Tino Kluge

April 24, 2007

Abstract

Most electricity markets exhibit high volatilities and occasional distinctive price spikes, which result in demand for derivative products which protect the holder against high prices. In this paper we examine a simple spot price model that is the exponential of the sum of an Ornstein-Uhlenbeck and an independent mean reverting pure jump process. We derive the moment generating function as well as various approximations to the probability density function of the logarithm of the spot price process at maturity T . Hence we are able to calibrate the model to the observed forward curve and present semi-analytic formulae for premia of path-independent options as well as approximations to call and put options on forward contracts with and without a delivery period. In order to price path-dependent options with multiple exercise rights like swing contracts a grid method is utilised which in turn uses approximations to the conditional density of the spot process.

1 Introduction

A distinctive feature of electricity markets is the formation of price spikes which are caused when the maximum supply and current demand are close, often when a generator or part of the distribution network fails unexpectedly. The occurrence of spikes has far reaching consequences for risk management and pricing purposes. In this context a parsimonious model with some degree of analytic tractability has clear advantages, and in this paper we propose and examine in detail a simple mean-reverting spot price process exhibiting spikes.

In our model the spot price process S is defined to be the exponential of the sum of three components: a deterministic periodic function f characterising seasonality, an Ornstein-Uhlenbeck (OU) process X and a mean-reverting process with a jump component to incorporate spikes Y :

$$\begin{aligned} S_t &= \exp(f(t) + X_t + Y_t), \\ dX_t &= -\alpha X_t dt + \sigma dW_t, \\ dY_t &= -\beta Y_{t-} dt + J_t dN_t, \end{aligned} \tag{1}$$

where N is a Poisson-process with intensity λ and J is an independent identically distributed (iid) process representing the jump size. We assume W , N and J to be mutually independent processes.

Our model generalises a number of earlier models. For example, the commonly used model of [Lucia and Schwartz, 2002] is obtained by setting $\beta = 0$ and taking $J = 0$. In this model S_t is log-normally distributed giving analytic option price formulae very similar to those in the Black-Scholes model. To allow for a stochastic seasonality a further component can be inserted into the model and, as long as this process has a normal distribution, analytic tractability is maintained.

The main disadvantage of the Lucia and Schwartz models is their inability to mimic price spikes. To overcome this, jumps can be added to the model; for example the case $\beta = 0$ in our model,

$$dX_t = -\alpha X_t dt + \sigma dW_t + J_t dN_t, \quad S_t = \exp(f(t) + X_t). \quad (2)$$

This model is briefly mentioned in [Clewlow and Strickland, 2000, Section 2.8]. Analytic results are given in [Deng, 2000] based on transform analysis described in [Duffie et al., 2000]. Calibration to historical data and the observed forward curve is discussed in [Cartea and Figueroa, 2005] where practical results for the UK electricity market are given. For these models to exhibit typical spikes the mean-reversion rate α must be extremely high, otherwise the jumps do not revert quickly enough.

In [Benth et al., 2005] a set of independent pure mean-reverting jump processes of the form

$$S_t = \sum_i w_i Y_t^{(i)}, \quad dY_t^{(i)} = -\alpha_i Y_t^{(i)} dt + \sigma_i dL_t^{(i)},$$

are introduced, where w_i are some positive weights and the $L^{(i)}$ are independent increasing càdlàg pure jump processes. The spot price process is a linear combination of the pure jump processes and as there is no exponential function involved, positivity of the spot is achieved by allowing positive jumps only. An advantage of this formulation is that semi-analytic formulae for option prices on forwards with a delivery period can be derived. However, a full analysis of this class of models still seems to be in its early stages.

The model (1) we consider in this paper is an extension of (2), in which we allow for two different mean-reversion rates, one for the diffusive part and one for the jump part. The introduction of a mean reverting spike process Y allows us to choose a higher mean-reversion rate β in order for the jump to revert much more quickly and so mimic a price spike. This is crucial for modelling the NordPool market but might not be needed in markets where the speed of mean-reversion α is generally very high, like in the UKPX or EEX.

Returning to our model (1) and solving for X_t and Y_t , we have

$$X_t = X_0 e^{-\alpha t} + \sigma \int_0^t e^{-\alpha(t-s)} dW_s, \quad Y_t = Y_0 e^{-\beta t} + \sum_{i=1}^{N_t} e^{-\beta(t-\tau_i)} J_{\tau_i}, \quad (3)$$

where τ_i indicates the random time of the occurrence of the i -th jump. Thus, given $X_0 = x_0$, $X_t \sim \mathcal{N}\left(x_0 e^{-\alpha t}, \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha t})\right)$. Properties of the spike process Y are not as obvious and will be examined in the following section. At this point we make no assumption on the jump size J but will later give results for exponentially and normally distributed jump sizes.

Note also that, although X and Y are both Markov processes, the price process S is not. We will therefore assume that all three components of the price process, i.e. the time t values $f(t)$, X_t and Y_t are individually observable, implying we expect jumps not to be small. Figure




Figure 1: Simulated sample paths of X , Y and S of the spot model (1). We use the following parameters which are of the same order of magnitude as calibrated values of the Nord Pool spot market, except the seasonality function $f(\cdot)$ which has been chosen arbitrarily: $f(t) = \ln(100) + 0.5 \cos(2\pi t)$, $\alpha = 7$, $\sigma = 1.4$, $\beta = 200$, $J_t \sim \exp(1/\mu_J)$, $\mu_J = 0.4$, $\lambda = 4$.

1 shows a simulated sample path of the processes X , Y and the composed process S .

In Section 2 we derive important stochastic properties of the process, including the moment generating function and various approximations to its probability density function. Pricing of a variety of derivative contracts will be discussed in Sections 3 and 4, using the results obtained in Section 2.

2 Properties of the model for spot prices

2.1 The spike process

The following result is known and given in [Duffie et al., 2000] in a more general framework.

Lemma 2.1 (moment generating function of the spike process, Y_t) Let $\{J_1, J_2, \dots\}$ be a series of iid random variables with the moment generating function $\Phi_J(\theta) := \mathbb{E}[e^{\theta J}]$. Let $\{\tau_1, \tau_2, \dots\}$ be the random jump times of a Poisson process N with intensity λ . Then the process Y with initial condition $Y_0 = 0$ has the moment generating function

$$\Phi_Y(\theta, t) := \mathbb{E}[e^{\theta Y_t}] = \exp\left(\lambda \int_0^t \Phi_J(\theta e^{-\beta s}) - 1 \, ds\right). \quad (4)$$

Furthermore, the first two moments of Y_t are given by

$$\mathbb{E}[Y_t] = \Phi_Y'(0, t) = \frac{\lambda}{\beta} \mathbb{E}[J](1 - e^{-\beta t}), \quad \mathbb{E}[Y_t^2] = \Phi_Y''(0, t) = \mathbb{E}[Y_t]^2 + \frac{\lambda}{2\beta} \mathbb{E}[J^2](1 - e^{-2\beta t}),$$

and in particular we have

$$\text{var}[Y_t] = \frac{\lambda}{2\beta} \mathbb{E}[J^2](1 - e^{-2\beta t}).$$

Remark 2.2 (asymptotics for $\beta \rightarrow \infty$) As remarked above, in practice the timescale $1/\beta$ for mean-reversion of spikes is much shorter than any of: the contract lifetime T ; the diffusive mean-reversion time $1/\alpha$; the volatility timescale $1/\sigma^2$; the mean arrival time of spikes $1/\lambda$. We therefore calculate approximations for the moment generating function of the spike process, and for its distribution, as $\beta \rightarrow \infty$.

To analyse the behaviour of the moment generating function for large β we make the substitution $u = \theta e^{-\beta s}$ in the integrand to obtain

$$\Phi_{Y_t}(\theta) = \exp\left(\frac{\lambda}{\beta} \int_{\theta e^{-\beta t}}^{\theta} \frac{\Phi_J(u) - 1}{u} du\right).$$

For fixed θ, t , as $\beta \rightarrow \infty$ we have

$$\int_0^{\theta e^{-\beta t}} \frac{\Phi_J(u) - 1}{u} du = \theta e^{-\beta t} \mathbb{E}[J] + O(e^{-2\beta t}),$$

because $\Phi_J(u) = 1 + \mathbb{E}[J]u + O(u^2)$, $u \rightarrow 0$, and so

$$\Phi_{Y_t}(\theta) = \exp\left(\frac{\lambda}{\beta} \left(\int_0^{\theta} \frac{\Phi_J(u) - 1}{u} du - \theta e^{-\beta t} \mathbb{E}[J] + O(e^{-2\beta t})\right)\right). \quad (5)$$

Example 2.3 (exponentially distributed jump size) If $J \sim \text{Exp}(1/\mu_J)$ with mean jump size μ_J , then $\Phi_J(\theta) = 1/(1 - \theta\mu_J)$, $\theta\mu_J < 1$. We obtain

$$\Phi_Y(\theta, t) = \left(\frac{1 - \theta\mu_J e^{-\beta t}}{1 - \theta\mu_J}\right)^{\lambda/\beta}, \quad \theta\mu_J < 1.$$

As $t \rightarrow \infty$, we have $\Phi_Y(\theta, t) \rightarrow (1 - \theta\mu_J)^{-\lambda/\beta}$ so the stationary distribution for Y is $\Gamma(\lambda/\beta, 1/\mu_J)$. As $\beta \rightarrow \infty$, we also have $\Phi_Y(\theta, t) = 1 + \theta\mu_J\lambda/\beta + O(\beta^{-2}) \approx (1 - \theta\mu_J)^{-\lambda/\beta}$. Thus Y_t is approximately $\Gamma(\lambda/\beta, 1/\mu_J)$ distributed for large β .

The mean and variance of the spike process Y_t with $Y_0 = 0$ are

$$\mathbb{E}[Y_t] = \frac{\lambda\mu_J}{\beta}(1 - e^{-\beta t}), \quad \text{var}[Y_t] = \frac{\lambda\mu_J^2}{\beta}(1 - e^{-2\beta t}).$$

2.2 The combined process

Having examined the properties of the spike process Y we conclude properties of the sum $X_t + Y_t$ and consequently of the price $S_t = \exp(f(t) + X_t + Y_t)$.

Theorem 2.4 Let the spot process S be defined by (1) and let $Z_t := \ln S_t = f(t) + X_t + Y_t$ with X_0 and Y_0 given. The moment generating function of Z_t is then

$$\mathbb{E} e^{\theta Z_t} = \exp\left(\theta f(t) + \theta X_0 e^{-\alpha t} + \theta^2 \frac{\sigma^2}{4\alpha}(1 - e^{-2\alpha t}) + \theta Y_0 e^{-\beta t} + \lambda \int_0^t \Phi_J(\theta e^{-\beta s}) - 1 ds\right). \quad (6)$$

Proof. The processes X and Y are independent so the expectation of the product is the product of the expectations. The moment generating function of Y_t is given in Lemma 2.1 which yields the result. \square

The expectation value of the spot process at time t , S_t immediately follows by setting $\theta = 1$.

2.3 Approximations

We now derive approximations to the density functions of the spike process at maturity T for large mean-reversion values β of the spike process. Although we can always compute the density by Laplace inversion of the moment generating function, an explicit expression for the density allows for more efficient algorithms and more explicit option pricing formulae. Here we only provide an expression for the density of a ‘truncated’ spike process \tilde{Y}_T as defined below. However, knowledge of the density of \tilde{Y}_T alone will help us to efficiently construct a grid to price swing options. We start by defining the truncated spike process \tilde{Y} , showing that \tilde{Y}_t provides a good approximation to the value Y_t for large values of β , deriving a general formula for the density function of \tilde{Y}_T and finally making it more explicit by considering an exponential jump size distribution.

For very high mean reversion rates β and small jump intensities λ , the dominant contribution to the density of the spike process comes from the last jump. We therefore introduce the truncated spike process

$$\tilde{Y}_t := \begin{cases} J_{N_t} e^{-\beta(t-\tau_{N_t})} & N_t > 0, \\ 0 & N_t = 0. \end{cases} \quad (7)$$

Note that we only consider Y starting from 0; any other starting point can be incorporated by adding the initial value.

Lemma 2.5 \tilde{Y}_t is identically distributed as

$$Z_t := \begin{cases} J_1 e^{-\beta\tau_1} & \tau_1 \leq t, \\ 0 & \tau_1 > t. \end{cases}$$

Proof. We use the reversibility property that if $N = \{N_t; t \in \mathbb{R}_+\}$ is a Poisson process then $\hat{N} = \{-N_{-t}; t \in \mathbb{R}_+\}$ is also a Poisson process. As τ_{N_t} is the jump time of the last jump before t , this translates to the first jump of the reversed process and hence $t - \tau_{N_t}$ and τ_1 are identically distributed, given $N_t > 0$. If $N_t = 0$ then there has been no jump in $[0, t]$ and the same applies for the reversed process and so this is equivalent to $\tau_1 > t$. \square

Lemma 2.6 (moment generating function of the truncated spike process) *The random variable \tilde{Y}_t of the truncated spike process at time t with initial condition $\tilde{Y}_0 = 0$ has the moment generating function*

$$\Phi_{\tilde{Y}_t}(\theta) = 1 + \lambda \int_0^t \left(\Phi_J(\theta e^{-\beta s}) - 1 \right) e^{-\lambda s} ds,$$

and the first two moments are given by

$$\mathbb{E}[\tilde{Y}_t] = \frac{\lambda}{\beta + \lambda} \mathbb{E}[J] \left(1 - e^{-(\beta + \lambda)t} \right), \quad \mathbb{E}[\tilde{Y}_t^2] = \frac{\lambda}{2\beta + \lambda} \mathbb{E}[J^2] \left(1 - e^{-(2\beta + \lambda)t} \right).$$

Proof. By Lemma 2.5 we only need to determine the moment generating function of

$$Z_t := J e^{-\beta\tau} I_{\tau \leq t}, \quad \tau \sim \text{Exp}(\lambda),$$

where I_A denotes the indicator function for the event A . Given the jump time τ we have

$$\mathbb{E}[e^{\theta Z_t} | \tau = s] = \Phi_J(\theta e^{-\beta s} I_{s \leq t}),$$

and so

$$\begin{aligned} \mathbb{E}[e^{\theta Z_t}] &= \mathbb{E}[\mathbb{E}[e^{\theta Z_t} | \tau]] \\ &= \int_0^\infty \Phi_J(\theta e^{-\beta s} I_{s \leq t}) e^{-\lambda s} ds \\ &= \int_0^t \Phi_J(\theta e^{-\beta s}) e^{-\lambda s} ds + e^{-\lambda t}. \end{aligned}$$

The first two moments are given by $\mathbb{E}[\tilde{Y}_t] = \Phi'_{\tilde{Y}_t}(0)$ and $\mathbb{E}[\tilde{Y}_t^2] = \Phi''_{\tilde{Y}_t}(0)$. \square

Remark 2.7 (pointwise convergence of the moment generating functions) The moment generating function of the truncated spike process converges pointwise to the moment generating function of the spike process for either $\lambda \rightarrow 0$ or $\beta \rightarrow \infty$ with t and θ fixed. First consider $\lambda \rightarrow 0$. Fix all other parameters and set $g(s; \beta, \theta) := \Phi_J(\theta e^{-\beta s}) - 1$, then

$$\begin{aligned} \Phi_{Y_t}(\theta) &= \exp\left(\lambda \int_0^t g(s; \beta, \theta) ds\right) = 1 + \lambda \int_0^t g(s; \beta, \theta) ds + O(\lambda^2), \\ \Phi_{\tilde{Y}_t}(\theta) &= 1 + \lambda \int_0^t g(s; \beta, \theta) e^{-\lambda s} ds = 1 + \lambda \int_0^t g(s; \beta, \theta) ds + O(\lambda^2). \end{aligned}$$

To see the convergence for $\beta \rightarrow \infty$ with λ , t and θ fixed, first note that from (5) we have

$$\Phi_{Y_t}(\theta) = 1 + \frac{\lambda}{\beta} \int_0^\theta \frac{\Phi_J(u) - 1}{u} du + O(1/\beta^2).$$

Also from Lemma 2.6,

$$\begin{aligned} \Phi_{\tilde{Y}_t}(\theta) &= 1 + \lambda \int_0^t \left(\Phi_J(\theta e^{-\beta s}) - 1\right) e^{-\lambda s} ds \\ &= 1 + \frac{\lambda}{\beta} \int_{\theta e^{-\beta t}}^\theta \frac{\Phi_J(u) - 1}{u} \left(\frac{u}{\theta}\right)^{\lambda/\beta} du \end{aligned}$$

by setting $\theta e^{-\beta s} = u$. Now as $\beta \rightarrow \infty$, $(u/\theta)^{\lambda/\beta} \rightarrow 1$ except in a small region $u = O(\theta e^{-\beta/\lambda})$, which makes a negligible (exponentially small) contribution to the integral. Likewise we may replace the lower limit of integration by 0 and incur a similarly small error. Hence

$$\begin{aligned} \Phi_{\tilde{Y}_t}(\theta) &= 1 + \frac{\lambda}{\beta} \int_0^\theta \frac{\Phi_J(u) - 1}{u} du + o\left(\frac{1}{\beta}\right) \\ &= \Phi_{Y_t}(\theta) + o\left(\frac{1}{\beta}\right). \end{aligned}$$

Two examples of the approximated and exact moment generating function using our standard parameters can be seen in Figure 2.




Figure 2: Moment generating function of Y_t and \tilde{Y}_t , denoted by mgf and approximation, respectively. In the left we use $J \sim \text{Exp}(1/\mu_J)$ and in the right $J \sim \mathcal{N}(\mu_J, \mu_J^2)$. The same parameters as in Figure 1 are used and we set $t = 1$.

Lemma 2.8 (distribution of the truncated spike process) *Let the jump size distribution have density function f_J . Then the truncated spike process \tilde{Y}_t as defined above has the cdf*

$$F_{\tilde{Y}_t}(x) = e^{-\lambda t} I_{x \geq 0} + \int_{-\infty}^x f_{\tilde{Y}_t}(y) dy, \quad t \geq 0,$$

with

$$f_{\tilde{Y}_t}(x) = \frac{\lambda}{\beta} \frac{1}{|x|^{1-\lambda/\beta}} \left| \int_x^{x e^{\beta t}} f_J(y) |y|^{-\lambda/\beta} dy \right|, \quad x \neq 0. \quad (8)$$

Proof. Based on Lemma 2.5 it suffices to determine the distribution of

$$\tilde{Y}_t = JZ I_{\tau \leq t}, \quad Z := e^{-\beta \tau}, \quad \tau \sim \text{Exp}(\lambda).$$

It follows that Z is the $\frac{\beta}{\lambda}$ th power of an uniformly distributed random variable on $[0, 1]$ and its density is given by

$$f_Z(x) = \frac{\lambda}{\beta} x^{-(1-\frac{\lambda}{\beta})} I_{x \in [0, 1]}.$$

As $\mathbb{P}(\tau > t) = e^{-\lambda t}$ we obtain the cdf of $Z I_{\tau \leq t}$ as

$$F_{Z I_{\tau \leq t}}(x) = e^{-\lambda t} I_{x \geq 0} + \int_{-\infty}^x f_{Z I_{\tau \leq t}}(y) dy,$$

$$f_{Z I_{\tau \leq t}}(x) = \frac{\lambda}{\beta} x^{-(1-\frac{\lambda}{\beta})} I_{x \in [e^{-\beta t}, 1]},$$

and the distribution of the product of two independent random variables J and $Z I_{\tau \leq t}$ is then given by

$$F_{J Z I_{\tau \leq t}}(c) = e^{-\lambda t} I_{c \geq 0} + \int_{-\infty}^c f_{J Z I_{\tau \leq t}}(x) dx,$$

$$f_{J Z I_{\tau \leq t}}(c) = \int_{-\infty}^{\infty} f_{Z I_{\tau \leq t}}(c/x) \frac{f_J(x)}{|x|} dx.$$




Figure 3: Distribution of the spike process (Y_t) at T with a jump size of $J \sim \text{Exp}(1/\mu_J)$. We use approximation (9) and compare it with the exact density as produced by a Monte-Carlo simulation. We use the same parameters as in Figure 1.

With

$$f_{ZI_{\tau \leq t}}(c/x) = \frac{\lambda}{\beta} \frac{1}{c^{1-\frac{\lambda}{\beta}}} I_{x \in [c, c e^{\beta t}]} x^{1-\frac{\lambda}{\beta}}, \quad c > 0,$$

$$f_{ZI_{\tau \leq t}}(c/x) = \frac{\lambda}{\beta} \frac{1}{|c|^{1-\frac{\lambda}{\beta}}} I_{x \in [c e^{\beta t}, c]} |x|^{1-\frac{\lambda}{\beta}}, \quad c < 0,$$

the desired result follows. \square

Example 2.9 (exponential jump size distribution) Let $J \sim \text{Exp}(1/\mu_J)$ be exponentially distributed. The distribution of the truncated spike process \tilde{Y}_t is

$$f_{\tilde{Y}_t}(x) = \frac{\lambda}{\beta \mu_J^{\lambda/\beta}} \frac{\Gamma(1 - \lambda/\beta, x/\mu_J) - \Gamma(1 - \lambda/\beta, x e^{\beta t}/\mu_J)}{x^{1-\lambda/\beta}}, \quad x > 0, \quad (9)$$

where $\Gamma(a, x)$ is the incomplete Gamma function. The approximation is a good fit to the exact density for typical market parameters as can be seen in Figure 3. The only discrepancy occurs at $Y_t = 0$ where the density has a singularity. We use this approximation in Section 4.1 to efficiently generate a grid to price swing options.

3 Option pricing

The electricity market with the model presented is obviously incomplete. Not only are we faced with a non-hedgeable jump risk but also we cannot use the underlying process (S_t) to hedge derivatives due to inefficiencies in storing electricity. Hence, the discounted spot price

process in the risk-neutral measure is not necessarily a martingale. From now on we assume the model is specified in the risk-neutral measure \mathbb{Q} as

$$\begin{aligned} S_t &= \exp(f(t) + X_t + Y_t), \\ dX_t &= -\alpha X_t dt + \sigma dW_t, \\ dY_t &= -\beta Y_{t-} dt + J_t dN_t, \end{aligned} \tag{10}$$

where W is a Brownian motion under \mathbb{Q} and N a Poisson process with intensity λ under \mathbb{Q} . For simplicity of notation we use the same parameters as in (1) but note that they might differ from the parameters under the real world measure. As is reasonable, the risk-neutral model has a jump structure that is similar to that observed under \mathbb{P} .

Lemma 3.1 (seasonal function consistent with the forward curve) *Let $t = 0$ and $F_0^{[T]}$ be the forward at time 0 maturing at time T ; then the risk-neutral seasonality function is given by*

$$f(T) = \ln F_0^{[T]} - X_0 e^{-\alpha T} - Y_0 e^{-\beta T} - \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha T}) - \lambda \int_0^T \Phi_J(e^{-\beta s}) - 1 ds. \tag{11}$$

Proof. The forward price is $F_0^{[T]} = \mathbb{E}^{\mathbb{Q}}[S_T]$ and so the result follows from (6). \square

This result forms an important part of the calibration of the model. As the model is incomplete the calibration procedure depends on the set of liquid derivatives used. Here we assume a continuous forward curve is observable in the market, i.e. values of $F_0^{[T]}$ are given. This is not a realistic assumption but there are ways to generate a continuous curve consistent with discretely observed prices, see Figure 4 and [Kluge, 2006, Section 2.2.2.5] for more details.

For the sake of simplicity we adopt a policy of choosing parameters from the real world measure \mathbb{P} if they are not uniquely determined by the set of observed derivative prices. This is equivalent to saying we choose a risk-neutral measure \mathbb{Q} which changes as few parameters of the model as possible. The volatility parameter σ remains unchanged by any equivalent measure change so we can always determine it from historical data. If we only see forward prices in the market we can also calibrate all other parameters to historical data except the seasonal function. In brief, this can be done by de-seasonalising the data using the real-world seasonality and then making a first estimation of α from which we can start filtering suspected spikes. From the reduced dataset, α can be re-estimated and suspected spikes filtered recursively. Having determined all parameters from historical data we finally calculate the risk-neutral seasonality function f from the observed forward curve based on (11).

3.1 Pricing path-independent options

If the payoff of an option on the spot at maturity T is given by $g(S_T)$ then its arbitrage free price at time t is given by

$$V(x, y, t) = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[g(S_T) | X_t = x, Y_t = y].$$

Although we do not have an expression for the density of S_T we know its moment generating function and so can apply Laplace transform methods to calculate the expectation value.




Figure 4: Interpolation of the forward curve by a seasonal function and spline correction. Three years worth of spot history data has been used to calibrate a seasonality function which is then used as a first approximation of the forward curve. The difference between the seasonality function and the observed forward prices is then corrected by a piecewise quadratic polynomial.

For an overview see [Cont and Tankov, 2004, Section 11.1.3]¹ or [Carr and Madan, 1998] and [Lewis, 2001]. Consider, for example, put or call options. Let $Z_t = \ln S_t$ and let $\Phi_{Z_t}(\theta)$ be its moment generating function, as given in (6). Now define its truncated moment generating function by

$$G_\nu(x, t) := \mathbb{E} \left[e^{\nu Z_t} I_{\{Z_t \leq x\}} \right] = \int_{-\infty}^x e^{\nu y} dF_{Z_t}(y),$$

which can be computed using a generalisation of Lévy's inversion theorem:

$$G_\nu(x) = \frac{\Phi_{Z_t}(\nu)}{2} - \frac{1}{\pi} \int_0^\infty \frac{\Im(\Phi_{Z_t}(\nu + i\theta) e^{-i\theta x})}{\theta} d\theta.$$

The price of a put option is then

$$\mathbb{E}[(K - S_T)^+] = K \mathbb{E}[I_{S_T \leq K}] - \mathbb{E}[S_T I_{S_T \leq K}] = KG_0(\ln K) - G_1(\ln K),$$

and by put-call parity we obtain the price of a call option.

3.2 Pricing options on Forwards

For a forward contract at time t , understood to be today, maturing at T the strike of a zero-cost forward is given by

$$F_t^{[T]} = \mathbb{E}^{\mathbb{Q}}[S_T | \mathcal{F}_t].$$

The most common options on forwards are puts or calls maturing at the same time as the underlying forward, i.e. the payoff is given by $(F_t^{[T]} - K)^+$ which is equivalent to $(S_T - K)^+$. We can price these contracts based on the dynamics of the spot and using methods developed above. However, by analysing the dynamics of the forward curve implied by the spot price model we will gain further insights and be able to relate the price of an option to the Black-76 formula [Black, 1976], which is still widely used in practice.

Recall that the expectation value of S_T is equal to the moment generating function given in (6) at $\theta = 1$. For $F_t^{[T]} = \mathbb{E}^{\mathbb{Q}}[S_T | X_t, Y_t]$ we obtain

$$F_t^{[T]} = \exp \left(f(T) + X_t e^{-\alpha(T-t)} + Y_t e^{-\beta(T-t)} + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds \right). \quad (12)$$

For fixed T , the dynamics of the forward maturing at T is then

$$\frac{dF_t^{[T]}}{F_t^{[T]}} = -\lambda \left(\Phi_J(e^{-\beta(T-t)}) - 1 \right) dt + \sigma e^{-\alpha(T-t)} dW_t + \left(\exp(J_t e^{-\beta(T-t)}) - 1 \right) dN_t. \quad (13)$$

The forward is a martingale under \mathbb{Q} by definition, and so the drift term compensates the jump process. For large time to maturities $T - t$, a jump in the underlying process has only very limited effect on the forward. More precisely, if the relative change in the underlying is $\exp(J_t) - 1$ the forward changes relatively by $\exp(J_t e^{-\beta(T-t)}) - 1$. In addition to the jump component the dynamics follows a deterministic volatility process starting with a low

¹They describe the method in terms of a complex valued characteristic function and Fourier inversion, but by allowing complex values the method can also be written in terms of Laplace transforms.

volatility $\sigma e^{-\alpha T}$ at $t = 0$ and increasing to σ at maturity. Without the jump component there are clear similarities with the Black-76 model.

For pricing purposes we need to find the distribution of $F_T^{[T]}$ in terms of its initial condition $F_t^{[T]}$. We have

$$\begin{aligned}\ln F_T^{[T]} &= f(T) + X_T + Y_T, \\ \ln F_t^{[T]} &= f(T) + X_t e^{-\alpha(T-t)} + Y_t e^{-\beta(T-t)} + \frac{\sigma^2}{4\alpha}(1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 \, ds.\end{aligned}$$

Eliminating the seasonality component $f(T)$, and using the relations

$$X_T - X_t e^{-\alpha(T-t)} = \sigma \int_t^T e^{-\alpha(T-s)} \, dW_s, \quad Y_T - Y_t e^{-\beta(T-t)} = \sum_{i=N_t}^{N_T} J_{\tau_i} e^{-\beta(T-\tau_i)},$$

we finally get

$$\begin{aligned}\ln F_T^{[T]} &= \ln F_t^{[T]} + \sigma \int_t^T e^{-\alpha(T-s)} \, dW_s + \sum_{i=N_t}^{N_T} J_{\tau_i} e^{-\beta(T-\tau_i)} \\ &\quad + \frac{\sigma^2}{4\alpha}(1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 \, ds.\end{aligned}\tag{14}$$

Without the jump component, $F_T^{[T]}$ would be log-normally distributed. In order to relate the pricing of options to the Black-76 formula even in the presence of spike risks, we assume that $F_T^{[T]}$ is log-normally distributed in a first approximation. We basically ignore the heavy tails caused by the spike risk and so expect to underestimate prices of far out-of-the-money calls but should do well with at-the-money calls.

We define the approximation by matching the first two moments but take into account that by definition $F_t^{[T]}$ is a martingale for a fixed maturity T and in order to keep the same property we set

$$\ln F_T^{[T]} \approx \ln F_t^{[T]} + \xi, \quad \xi \sim N\left(-\frac{1}{2}\hat{\sigma}^2(T-t), \hat{\sigma}^2(T-t)\right),$$

and set $\hat{\sigma}^2(T-t) := \text{var}[\ln F_T^{[T]} | \mathcal{F}_t]$, i.e.

$$\begin{aligned}\hat{\sigma}^2(T-t) &= \text{var}\left[\sigma \int_t^T e^{-\alpha(T-s)} \, dW_s + \sum_{i=N_t}^{N_T} J_{\tau_i} e^{-\beta(T-\tau_i)}\right] \\ &= \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha(T-t)}) + \frac{\lambda}{2\beta} \mathbb{E}[J^2](1 - e^{-2\beta(T-t)}).\end{aligned}$$

Remark 3.2 (term structure of implied volatility) Comparing this result with the setting of Black-76 [Black, 1976] where $dF = F\sigma \, dW$ and so $F_T = F_t \exp(\xi)$ with $\xi \sim N(-\frac{1}{2}\sigma^2(T-t), \sigma^2(T-t))$, we conclude that $\hat{\sigma}$ is the implied Black-76 volatility and in a first approximation given by

$$\hat{\sigma}^2 \approx \frac{\frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha(T-t)}) + \frac{\lambda}{2\beta} \mathbb{E}[J^2](1 - e^{-2\beta(T-t)})}{T-t},\tag{15}$$

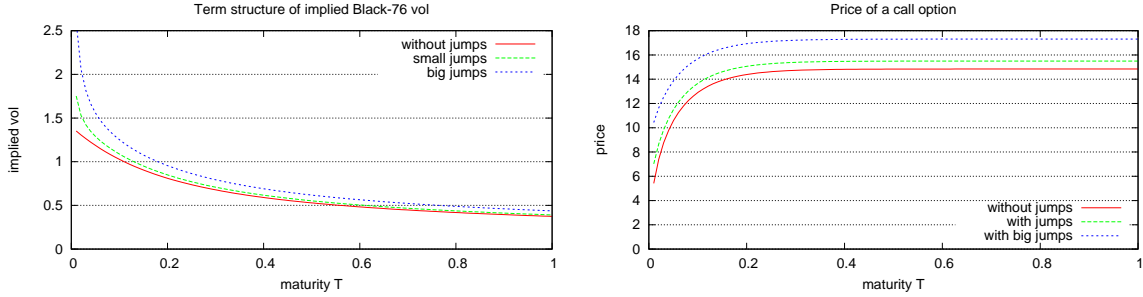


Figure 5: Implied volatilities and prices. The left graph shows implied volatilities with respect to time to maturity where approximation (15) is used. The three lines correspond to no jumps ($\mu_J = 0$), small jumps ($\mu_J = 0.4$) and big jumps ($\mu_J = 0.8$). In the right graph the corresponding prices of an at the money call are plotted. Parameters are $r = \ln(1.05)$, $\alpha = 7$, $\beta = 200$, $\sigma = 1.4$, $\lambda = 4.0$, $F_0^{[T]} = 100$, $K = 100$. Note, for an exponential distributed jump size J we have $\mathbb{E}[J^2] = 2\mu_J^2$.

which is shown in Figure 5. It can be seen that the spike process has a much more significant impact on the implied volatility for short maturities rather than for long term maturities. As far as the price of an at the money call is concerned, the additional jump risk adds an almost constant premium to the price to be paid without any jump risk.

Remark 3.3 (implied volatility across strikes) The approximation does not predict a change of implied volatility across strikes. However, the jump risk introduces a skew as can be seen in Figure 5 where the exact solution based on Section 3.1 has been used to calculate implied volatilities. The bigger the mean jump size and hence the bigger $\mathbb{E}[J^2]$, the more profound is the skew.

3.3 Pricing options on Forwards with a delivery period

As electricity is a flow variable, forwards always specify a delivery period. The results of the previous section can therefore only be seen as an approximation to option prices on forwards with short delivery periods, like one day. Here we only consider options on forwards maturing at the beginning of the delivery period, i.e. the payoff is given by some function of $F_{T_1}^{[T_1, T_2]}$ at time T_1 . An option on such a forward is conceptually similar to an Asian option in the Black-Scholes world. One method of pricing Asian options is to approximate the distribution of the integral by a log-normal distribution and can be done by matching the first two moments, see [Turnbull and Wakeman, 1991] for example. Once the parameters of the approximate log-normal distribution have been determined, pricing options comes down to pricing in the Black-Scholes or Black-76 setting.

The strike price of a zero cost forward with a delivery period is generally given by a weighted average of instantaneous forwards of the form

$$F_t^{[T_1, T_2]} = \int_{T_1}^{T_2} w(T; T_1, T_2) F_t^{[T]} dT,$$

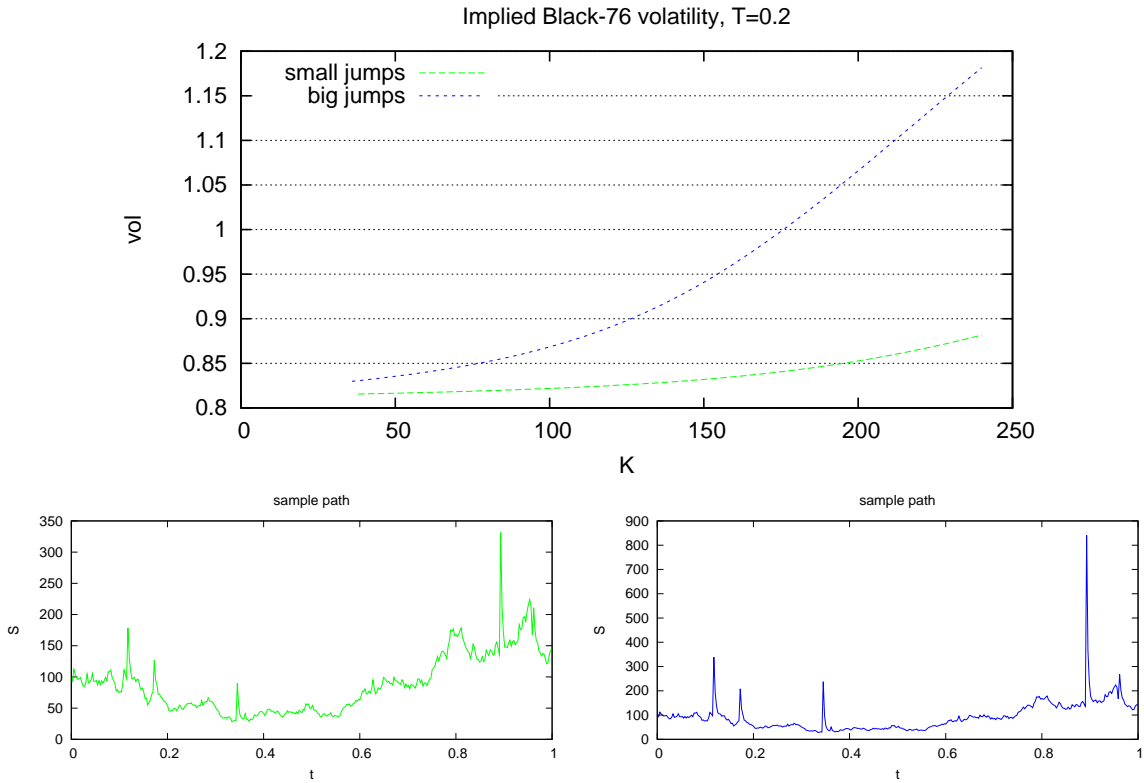


Figure 6: Implied volatilities across strikes and sample paths. The upper graph shows the implied volatility for one maturity $T = 0.2$ based on the exact solution. The approximate solution (15) yields 0.82 and 0.85 for the small and big jumps, respectively. Sample paths of the model with the same parameters are drawn in the lower two graphs, where the left path is generated with a low mean jump size ($\mu_J = 0.4$) and the right with a high mean jump size ($\mu_J = 0.8$). All the other parameters are the same as in Figure 5.

where for a settlement at maturity T_2 the weighting factor is given by $w(T; T_1, T_2) = 1/(T_2 - T_1)$ and for instantaneous settlement the discounting alters the weighting to $w(T; T_1, T_2) = r e^{-rT} / (e^{-rT_1} - e^{-rT_2})$.

The second moment of $F_{T_1}^{[T_1, T_2]}$ is given by

$$\mathbb{E}^{\mathbb{Q}} \left[\left(\int_{T_1}^{T_2} w(T) F_{T_1}^{[T]} \right)^2 dT \middle| \mathcal{F}_t \right] = \int_{T_1}^{T_2} \int_{T_1}^{T_2} w(T) w(T^*) \mathbb{E}^{\mathbb{Q}} \left[F_{T_1}^{[T]} F_{T_1}^{[T^*]} \middle| \mathcal{F}_t \right] dT dT^*,$$

and the expectation of the product of two individual forwards $\mathbb{E}^{\mathbb{Q}} \left[F_{T_1}^{[T]} F_{T_1}^{[T^*]} \middle| \mathcal{F}_t \right]$ can be derived using the solution of the forward (12) as follows:

$$\begin{aligned} \ln F_{T_1}^{[T]} &= \ln F_t^{[T]} + e^{-\alpha(T-T_1)} \sigma \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s + e^{-\beta(T-T_1)} \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} \\ &\quad - \frac{\sigma^2}{4\alpha} (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \\ &\quad + \lambda \int_0^{T-T_1} \Phi_J(e^{-\beta s}) - 1 ds - \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds \\ &= \ln F_t^{[T]} + e^{-\alpha(T-T_1)} \sigma \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s + e^{-\beta(T-T_1)} \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} \\ &\quad - \frac{\sigma^2}{4\alpha} (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) - \lambda \int_0^{T_1-t} \Phi_J(e^{-\beta(T-T_1)} e^{-\beta s}) - 1 ds, \end{aligned}$$

and so

$$\begin{aligned} \ln F_{T_1}^{[T]} + \ln F_{T_1}^{[T^*]} &= \ln F_t^{[T]} + \ln F_t^{[T^*]} \\ &\quad + \left(e^{-\alpha(T-T_1)} + e^{-\alpha(T^*-T_1)} \right) \sigma \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s \\ &\quad + \left(e^{-\beta(T-T_1)} + e^{-\beta(T^*-T_1)} \right) \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} \\ &\quad - \frac{\sigma^2}{4\alpha} (1 + e^{-2\alpha(T^*-T)}) (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \\ &\quad - \ln \Phi_Y(e^{-\beta(T-T_1)}) - \ln \Phi_Y(e^{-\beta(T^*-T_1)}), \end{aligned}$$

which gives

$$\begin{aligned} \mathbb{E}^{\mathbb{Q}} \left[F_{T_1}^{[T]} F_{T_1}^{[T^*]} \middle| \mathcal{F}_t \right] &= \mathbb{E}^{\mathbb{Q}} \left[\exp \left(\ln F_{T_1}^{[T]} + \ln F_{T_1}^{[T^*]} \right) \middle| \mathcal{F}_t \right] \\ &= F_t^{[T]} F_t^{[T^*]} \frac{\Phi_Y(e^{-\beta(T-T_1)} + e^{-\beta(T^*-T_1)})}{\Phi_Y(e^{-\beta(T-T_1)}) \Phi_Y(e^{-\beta(T^*-T_1)})} \\ &\quad \times \exp \left(-\frac{\sigma^2}{4\alpha} (1 + e^{-2\alpha(T^*-T)}) (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \right) \\ &\quad \times \exp \left(\frac{\sigma^2}{4\alpha} (1 + e^{-\alpha(T^*-T)})^2 (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \right). \end{aligned}$$




Figure 7: Distribution of $F_{T_1}^{[T_1, T_2]}$. Parameters are as before, see Figure 5, and $\mu_J = 0.4$, $T_1 = 1$, $T_2 = 1.25$. The red curve is based on a Monte-Carlo simulation with 10^6 sample paths; the green curve is the log-normal approximation matching the first two moments.

How well the moment matching procedure works is shown in Figure 7 where the density of a forward $F_{T_1}^{[T_1, T_2]}$ is compared to the density obtain by the approximation. The shapes of both densities are similar but differences in values are clearly visible. As a result one might not expect call option prices based on the approximate distribution to be very close to the exact prices for all strikes K but still close enough to be useful. For an at-the-money strike call option, prices for varying maturities are shown in Figures 8 and 9. As it turns out, the approximation gives very good results for short delivery periods and is still within a 5% range for delivery periods of one year.

4 Pricing swing options

Swing contracts are a broad class of path dependent options allowing the holder to exercise a certain right multiple times over a specified period but only one right at a time² or per time-interval like a day. Such a right could be the right to receive the payoff of a call option. Other possibilities include the mixture of different payoff functions like calls and puts or calls with different strikes. Another very common feature is to allow the holder to exercise a multiple of a call or put option at once, where the multiple is called volume. This generally involves further restriction on the volume, like upper and lower bounds for each right and for the sum of all trades.

Swing contracts can be seen as an insurance for the holder against excessive rises in electricity prices. Assuming the prices generally revert to a long term mean, even a small number N of exercise opportunities suffices to cover the main risks and hence make the premium of the contract cheaper. Sometimes, swing contracts are bundled with forward contracts. The

²This will also involve a “refraction period” in which no further right can be exercised.




Figure 8: Value of an at-the-money call option on a forward $F_{T_1}^{[T_1, T_2]}$ depending on the delivery period $T_2 - T_1$. Parameters are as before, see Figure 5, and $\mu_J = 0.4$, $T_1 = 1$, $K = 100$. The Monte-Carlo result is based on 100 000 sample paths for each duration. Note, any forward $F_{T_1}^{[T_1, T_2]}$ delivers exactly 1 MWh over the delivery period $[T_1, T_2]$.




Figure 9: Same as in Figure 8 but where the volume of the forward is proportional to the delivery period, i.e. we assume a constant consumption of 1MWh per year, which is about 114 W. Here, the price is simply $T_2 - T_1$ times the price of a standard call option on $F_{T_1}^{[T_1, T_2]}$.

forward contract then supplies the holder with a constant stream of energy to a fixed pre-determined price. If the strike price of the call options of the swing contract is set to the forward price, the swing contract will allow for flexibility in the volume the customer receives for the fixed price. They can either “swing up” or “swing down” the volume of energy and hence the name swing contract. One cannot assume that the holder always exercises the contract in an optimal way to maximise expected profit but they might also exercise according to their own internal energy demands.

It is only very recently that articles on numerical pricing methods for swing options have appeared in the literature. We can identify a few main approaches all based on dynamic programming principles. A Monte-Carlo method and ideas of duality theory are utilised in [Meinshausen and Hambly, 2004] to derive lower and upper bounds for swing option prices. The main advantages of the method being their flexibility as it can be easily adapted to any stochastic model of the underlying and its ability to produce confidence intervals of the price. Monte-Carlo techniques are also used in [Ibanez, 2004] and [Carmona and Touzi, 2004], where the latter uses the theory of the Snell envelope to determine the optimal exercise boundaries and also utilises the Malliavin calculus for the computation of greeks. A constructive solution to the perpetual swing case for exponential Brownian motion is also given in [Carmona and Touzi, 2004]. Unfortunately, these methods only work for the most basic versions of swing contracts where at each time only one unit of an option can be exercised.

More general swing contracts with a variable volume per exercise and an overall constraint can be priced with a tree based method introduced in [Jaillet et al., 2004].

In the above papers a discrete-time model for the underlying is used where one time step corresponds to the time frame in which no more than one right can be exercised, i.e. one day in most of the traded contracts. A special case where the number of exercise opportunities is equal to the number of exercise dates is considered in [Howison and Rasmussen, 2002] and a continuous optimal exercise strategy derived which yields a partial integro-differential equation for the option price.

Our method is based on the tree approach of [Jaillet et al., 2004] with some slight modifications to adapt it to the peculiarities of our model for the underlying electricity price process.

4.1 The grid approach

The tree method of [Jaillet et al., 2004] requires a discrete time model of the underlying. This is due to the fact that their swing contracts allow the holder to exercise at most one option within a specified time interval, say a day, and this is best modelled if the underlying process has the same time discretisation. Assuming (S_t) is some continuous stochastic process for the spot price we obtain a discrete model by observing it on discrete points in time only, i.e.

$$S_{t_0}, S_{t_1}, S_{t_2}, \dots, S_{t_m},$$

with $t_0 = 0$, $t_{i+1} = t_i + \Delta t$, $t_m = T$ and $\Delta t = \frac{1}{365}$, indicating we can exercise on a daily basis.

Let the maturity date T be fixed and the payoff at time t for simplicity³ be given by $(S_t - K)^+$ for some strike price K and we assume only one unit of the underlying can be exercised any time period. Let $V(n, s, t)$ denote the price of such a swing option at time t and spot price s

³We could assume any general payoff function.

which has n out of N exercise rights left. The dynamic programming principle allows us to write

$$V(n, s, t) = \max \left\{ e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s], \right. \\ \left. e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n-1, S_{t+\Delta t}, t + \Delta t) | S_t = s] + (s - K)^+ \right\}, \quad n < N, \quad (16)$$

and $V(n, s, T) = (S - K)^+$, $0 < n \leq N$ and $V(0, s, t) = 0$. The conditional expectations can be written

$$\mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s] = \int_{-\infty}^{\infty} V(n, x, t + \Delta t) f_S(x; s) dx,$$

where $f_S(x; s)$ is the density of $S_{t+\Delta t}$ given $S_t = s$. Discretising the spot variable we approximate

$$\mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s_i] \approx \sum_j V(n, s_j, t + \Delta t) f_S(s_j; s_i) \Delta s_j.$$

This is only one possible approximation; others might be to use higher order integration rules or using only a few grid points in the sum based on the fact that $f_S(x; s) \rightarrow 0$ for $|s - x|$ large. For a trinomial tree one only uses three grid points, i.e.

$$\mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s_i] \approx \sum_{j=-1}^1 V(n, s_{i+j}, t + \Delta t) p_{i,i+j},$$

$p_{i,j}$ being the probability of going from node i to node j . However, such a tree approach is not well suited to our case for two reasons. First, the time step size is determined by the shortest time between two possible exercise dates, which is mainly one day for swing contracts. This limits the accuracy of the algorithm as a refinement of the grid in spot direction will not improve the result. Second, in the presence of jumps, a three-point approximation for the conditional density is insufficient due to the heavy tails in the distribution. As a result, we keep our method general and say

$$\mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s_i] \approx \sum_j V(n, s_j, t + \Delta t) p_{i,j},$$

where $p_{i,j}$ is an approximation to the density $f_S(s_j; s_i) \Delta s_j$ (it can accommodate higher-order integration rules and boundary approximations). With the notation $V_{i,k}^n := V(n, s_i, t_k)$ we can then write the method as

$$V_{i,k}^n = \max \left\{ e^{-r\Delta t} \sum_j V_{j,k+1}^n p_{i,j}, e^{-r\Delta t} \sum_j V_{j,k+1}^{n-1} p_{i,j} + (s_i - K)^+ \right\}, \quad (17)$$

$$V_{i,k}^0 = 0,$$

$$V_{i,m}^n = 0.$$

4.2 Numerical results

We now turn to the model of interest, (10), which exhibits spikes. Assume that the mean-reversion process (X_t) and the spike process (Y_t) are individually observable and so the value

function V of a swing option depends on both variables and the general pricing principle (16) becomes

$$V(n, x, y, t) = \max \begin{cases} e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n, X_{t+\Delta t}, Y_{t+\Delta t}, t + \Delta t) | X_t = x, Y_t = y], \\ e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n-1, X_{t+\Delta t}, Y_{t+\Delta t}, t + \Delta t) | X_t = x, Y_t = y] + (e^{f(t)+x+y} - K)^+. \end{cases}$$

In order to calculate conditional expectations we need to define transition probabilities. Given one starts at node $(X_t, Y_t) = (x_i, y_j)$ the probability to arrive at node $(X_{t+\Delta t}, Y_{t+\Delta t}) = (x_k, y_l)$ is approximately given by

$$p_{i,j,k,l} \approx f_{X_{t+\Delta t}|X_t=x_i}(x_k) f_{Y_{t+\Delta t}|Y_t=y_j}(y_l) \Delta x \Delta y,$$

because X_t and Y_t are independent. The conditional density of the mean-reverting process (X_t) is known as $X_{t+\Delta t}$ given $X_t = x$ is normally distributed with $\mathcal{N}(x e^{-\alpha\Delta t}, \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha\Delta t}))$. As we do not have a closed form expression for the density of the spike process we use approximations developed in Section 2.1. For an exponential jump size distribution $J \sim \text{Exp}(1/\mu_J)$ for example we use approximation (9) for the spike process at time t given zero initial conditions.

The introduction of a second space dimension increases the complexity of the algorithm considerably, essentially by a factor proportional to the square of the number of grid points in the y -direction. To price the swing contract shown in Figure 10 which has 365 exercise dates and up to 100 exercise opportunities, our C++ implementation requires about 10 minutes to complete the calculation on an Intel P4, 3.4GHz, and for a grid of 120×60 points in x and y direction, respectively. The same computation but with no spikes and a grid of 120×1 points only takes about one second.

Based on Figure 10 we make two observations. First, the price per exercise right decreases with the number of exercise rights. This is the correct qualitative behaviour one would expect because n swing options each with one exercise right⁴ only, offer more flexibility than one swing option with n exercise rights.⁵ Second, the premium added due to the spike risk is much more significant for options with small numbers of exercise rights than for a large number. This is also intuitively clear, as an option with say 100 exercise rights will mainly be used against high prices caused by the diffusive part and only occasionally against spiky price explosions.

In Figure 11 we show how sensitive swing option prices are to changes in market parameters. There we consider a swing option with a duration of 60 days and up to 20 exercise opportunities. In each graph we change one parameter by 20% up and down. The most significant change is caused by a change in the volatility parameter σ . Note, the longterm variance of the mean-reverting process (X_t) is $\frac{\sigma^2}{2\alpha}$ and we expect some direct relationship between the long-term variance and the option price. Hence, a change in the mean-reversion parameter α is inversely proportional to the price and quantitatively changes the price less than the volatility σ . The mean-reversion parameter β of the spike process has a similar effect on the option price as α has, but where the influence slightly decreases with the number of options. This is consistent with previous observations of the impact of jumps on option prices as seen in Figure 10. This effect is much more clearly visible for the other jump parameters λ and μ_J

⁴This is actually an American option.

⁵The rights of a swing option can only be exercised one at a time.




Figure 10: Value of a one year swing option per exercise right. Market parameters of the underlying are as before, see Figure 1, i.e. $\alpha = 7$, $\beta = 200$, $\sigma = 1.4$, $\lambda = 4$, $J \sim \text{Exp}(1/\mu_J)$ with $\mu_J = 0.4$, $f(t) = 0$, $r = 0$, and initial conditions $X_0 = 0$ and $Y_0 = 0$. The swing contract delivers over a time period of one year $T \in [0, 1]$ with up to 100 call rights and a strike price of $K = 1$, where a right can be exercised on any day. As a comparison the price of the same swing option is plotted but where the underlying does not exhibit spikes, i.e. $\lambda = 0$.

which have the greatest impact on options with only a few exercise rights. For one exercise right, a 20% change in the jump size μ_J has an even greater effect on the price than a 20% change in volatility σ . A possible explanation is that we deal with the exponential of an exponentially distributed jump size which is very heavy tailed.

References

- [Benth et al., 2005] Benth, F., Kallsen, J., and Meyer-Brandis, T. (2005). A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivative pricing. preprint.
- [Black, 1976] Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3:167–179.
- [Carmona and Touzi, 2004] Carmona, R. and Touzi, N. (2004). Optimal multiple stopping and valuation of swing options. to appear in *Mathematical Finance*.
- [Carr and Madan, 1998] Carr, P. and Madan, D. (1998). Option valuation using the fast Fourier transform. *J. Comput. Fin.*, 2:61–73.
- [Cartea and Figueroa, 2005] Cartea, A. and Figueroa, M. (2005). Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, 12(4):313–335.




Figure 11: Sensitivity of swing option prices with respect to model parameters. A swing option with 60 exercise dates and up to 20 rights is considered, where the red curve is based on the parameters of Figure 10. In each graph one market parameter is shifted up by 20% (green line) and down by 20% (blue line). We always plot the option price divided by the number of exercise rights.

- [Clewlow and Strickland, 2000] Clewlow, L. and Strickland, C. (2000). *Energy Derivatives: Pricing and Risk Management*. Lacima Publications, London, UK.
- [Cont and Tankov, 2004] Cont, R. and Tankov, P. (2004). *Financial Modelling with Jump Processes*. Chapman & Hall/CRC.
- [Deng, 2000] Deng, S. (2000). Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes. Working paper PWP-073.
- [Duffie et al., 2000] Duffie, D., Pan, J., and Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.
- [Howison and Rasmussen, 2002] Howison, S. and Rasmussen, H. (2002). Continuous swing options. working paper.
- [Ibanez, 2004] Ibanez, A. (2004). Valuation by simulation of contingent claims with multiple early exercise opportunities. *Mathematical Finance*, 14(2):223–248.
- [Jaillet et al., 2004] Jaillet, P., Ronn, E., and Tompadis, S. (2004). Valuation of commodity-based swing options. *Management Science*, 50:909–921.
- [Kluge, 2006] Kluge, T. (2006). *Pricing Swing Options and other Electricity Derivatives*. PhD thesis, University of Oxford.
- [Lewis, 2001] Lewis, A. (2001). A simple option formula for general jump-diffusion and other exponential Lévy processes. working paper.
- [Lucia and Schwartz, 2002] Lucia, J. and Schwartz, E. (2002). Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, 5(1):5–50.
- [Meinshausen and Hambly, 2004] Meinshausen, N. and Hambly, B. (2004). Monte carlo methods for the valuation of options with multiple exercise opportunities. *Mathematical Finance*, 14(4):557–583.
- [Turnbull and Wakeman, 1991] Turnbull, S. M. and Wakeman, L. M. (1991). A quick algorithm for pricing european average options. *Journal of Financial and Quantitative Analysis*, 26(3):377–389.