

Option Formulas for Mean-Reverting Power Prices with Spikes

Cyriel de Jong, Ronald Huisman

ERIM REPORT SERIES <i>RESEARCH IN MANAGEMENT</i>	
ERIM Report Series reference number	ERS-2002-96-F&A
Publication status / version	October 2002
Number of pages	32
Email address corresponding author	r.huisman@fbk.eur.nl
Address	Erasmus Research Institute of Management (ERIM) Rotterdam School of Management / Faculteit Bedrijfskunde Erasmus Universiteit Rotterdam PoBox 1738 3000 DR Rotterdam, The Netherlands Phone: # 31-(0) 10-408 1182 Fax: # 31-(0) 10-408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES *RESEARCH IN MANAGEMENT*

BIBLIOGRAPHIC DATA AND CLASSIFICATIONS		
Abstract	Electricity prices are known to be very volatile and subject to frequent jumps due to system breakdown, demand shocks, and inelastic supply. Appropriate pricing, portfolio, and risk management models should incorporate these spikes. We develop a framework to price European-style options that are consistent with the possibility of market spikes. The pricing framework is based on a regime jump model that disentangles mean-reversion from the spikes. In the model the spikes are truly time-specific events and therefore independent from the mean-reverting price process. This closely resembles the characteristics of electricity prices, as we show with Dutch APX spot price data in the period January 2001 till June 2002. Thanks to the independence of the two price processes in the model, we break derivative prices down in a mean-reverting value and a spike value. We use this result to show how the model can be made consistent with forward prices in the market and present closed-form formulas for European-style options.	
Library of Congress Classification (LCC)	5001-6182	Business
	5601-5689	Accountancy, Bookkeeping
	4001-4280.7	Finance Management, Business Finance, Corporation Finance
	HG 6024+	Options
Journal of Economic Literature (JEL)	M	Business Administration and Business Economics
	M 41	Accounting
	G 3	Corporate Finance and Governance
	G 19	General Financial Markets: Other
European Business Schools Library Group (EBSLG)	85 A	Business General
	225 A	Accounting General
	220 A	Financial Management
	220 R	Options market
Gemeenschappelijke Onderwerpsontsluiting (GOO)		
Classification GOO	85.00	Bedrijfskunde, Organiseatiekunde: algemeen
	85.25	Accounting
	85.30	Financieel management, financiering
	85.30	Financieel management, financiering
Keywords GOO	Bedrijfskunde / Bedrijfseconomie	
	Accountancy, financieel management, bedrijfsfinanciering, besliskunde	
	Optiehandel, energieprijzen, econometrische modellen	
Free keywords	Option Pricing, Electricity Price Modelling, Power Spikes, Energy Markets, Mean Reversion	

Option Formulas for Mean-Reverting Power Prices with Spikes

Cyriel de Jong
Ronald Huisman

Energy Global
Rotterdam School of Management at Erasmus University

September 2002

Electricity prices are known to be very volatile and subject to frequent jumps due to system breakdown, demand shocks, and inelastic supply. Appropriate pricing, portfolio, and risk management models should incorporate these spikes. We develop a framework to price European-style options that are consistent with the possibility of market spikes. The pricing framework is based on a regime jump model that disentangles mean-reversion from the spikes. In the model the spikes are truly time-specific events and therefore independent from the mean-reverting price process. This closely resembles the characteristics of electricity prices, as we show with Dutch APX spot price data in the period January 2001 till June 2002. Thanks to the independence of the two price processes in the model, we break derivative prices down in a mean-reverting value and a spike value. We use this result to show how the model can be made consistent with forward prices in the market and present closed-form formulas for European-style options.

All errors pertain to the authors. Cyriel de Jong is the contacting author. He can be contacted at Energy Global, Wilhelminaplein 5, 3072 DE, Rotterdam, The Netherlands, tel. +31 10 485 6033, fax +31 10 485 1502, cyrieldejong@energyglobal.nl. The authors would like to thank Ronald Mahieu for his comments on a previous version of the paper.

1. Introduction

Since the early 90's electricity markets are being reformed worldwide from a highly government controlled and vertically integrated environment into competitive markets. Before the deregulation, government authorities fixed prices based on (marginal) production costs in a very predictable manner. Now that many wholesale markets are deregulated, market participants have to get used to an environment with very volatile prices and high uncertainty. Participants face the added complexity that volatility far exceeds the volatility in markets that are considered relatively risky, such as those for stocks, bonds, and other commodities. At the same time, the number of available instruments to control risks has grown radically. Markets gradually extend trading in day-ahead physical deliveries (spot or pool), and forward contracts with physical deliveries, to more advanced physical and financial products, such as swaps, futures, options, caps, floors and spark spreads. Most derivatives trade in over-the-counter markets, but increasingly on exchanges as well. Examples of such trading venues are the NYMEX, Nord Pool / Eltermin / Eloption, European Energy Exchange (Germany), Chicago Board of Trade, Minneapolis Grain Exchange, Sydney Futures Exchange, and New Zealand Futures Exchange.

For the valuation of electricity derivatives we cannot simply rely on models for financial and other commodity contracts. Electricity is a pure flow commodity with limited storability and transportability that strongly affect the behaviour of electricity spot and derivatives prices. Since we face serious obstacles to transport electricity from one place to the other, and even larger obstacles to carry it from one time period to the other, arbitrage is hardly feasible. This lack of flexibility causes spot prices to depend largely on local and temporal supply and demand conditions. If demand and supply would respond promptly to price movements, prices would not deviate much from other commodity prices. The elasticity of both supply and demand however is relatively limited (see for example Pirrong and Jermakyan, 2001). Only a few large industrial customers have the flexibility to vary their power demand in response to market conditions, whereas most power plants can gear up generation only with a significant time lag. This time lag causes occasional extreme prices, which do however revert within hours, days, or weeks from temporarily extreme levels to a more stable level. All this results in the well-documented characteristics of electricity spot prices, including spikes, mean-reversion, large seasonal variations and extremely high volatility.

These peculiar characteristics of electricity prices have induced researchers to develop special models for electricity prices. Such models are the basis for risk management applications, for the pricing of physical and financial contracts, and for the valuation of real assets. The models needed are relatively complex in order to capture seasonality, mean-reversion, (time-varying) high volatility and spikes at the same time. With the increasing number of tradable contracts, the main challenge for researchers is the development of models to price those contracts. Although interest is growing, the number of papers that address the specific valuation problems is relatively limited, and only a few have been published to date. We recognise two different sets of electricity contract valuation approaches. The first approach is most popular among academics and consists of modelling simultaneously spot and forward contracts. Examples are Schwartz (1997), Hilliard and Reis (1998), Pilipovic (1998), Pirrong and Jermakyan (1999, 2001), Deng (2000), Karsen and Husby (2001), and Lucia and Schwartz (2002). Those papers face the difficulty that standard arbitrage principles cannot be applied to map spot prices to forwards and futures. Therefore, the proposed solutions are naturally derived from the bond or storable-commodity pricing literature. Risk-neutral processes are obtained either through the specification of risk premia or convenience yields. A theoretical drawback of models based on convenience yields is that electricity is not storable, and therefore the interpretation of convenience yields is questionable. Moreover, fitting the theoretical forward curve to market data is a serious problem for both methods, since data is limited and several institutional factors influence power forward price dynamics.

The second approach describes how to price options on spot, forwards or futures, and takes the forward curve as given. It is the approach we take. This is in contrast with the models above that derive the forward curve endogenously. Our approach makes option valuation consistent with market prices, and we believe this is necessary to bridge the gap between academic theory and the derivative models that are predominantly being used in practice (Black & Scholes (1973) and Black (1976)). For option valuation we simply need models that adequately describe the dynamics of spot or forward prices, align them with the market forward curve, and then use arbitrage or “fair-pricing” principles to derive option prices¹. If we need to price an option on a forward for example, we just want to have a good description of the evolution of forward prices. Similarly, if we need to price an option on the spot, we simply need a good spot price model. In other words, we want our assessment of an option value to deviate from market option prices, not because we have a different opinion on the spot and

¹ It should be noted that forward trading is often liquid for only a few maturity series, but combinations of market prices with bottom-up models (see e.g. Fleten and Lemming (2001)) may be used to complete the forward curve.

forward *level*, but a different opinion on the spot and forward *dynamics*. While being simpler at first sight, our approach has at the same time more degrees of freedom to describe the underlying time series accurately. Other examples of this approach are Miltersen and Schwartz (1998), Clewlow and Strickland (1999), Bjerksund, Rasmussen and Stensland (2000) and Koekebakker and Ollmar (2001).

To understand what type of model we exactly need for option pricing, it is important to note that most electricity forwards and futures are defined slightly differently from forwards and futures on other markets. Because electricity is a pure flow variable, the pay-off on longer-term forwards and futures in electricity market is based on the average spot price during the settlement period (e.g. a month, quarter or year), instead of a specific spot closing price, as is common in most other markets. Since this averaging strongly affects the price behaviour, we need to differentiate between valuation models for options on the longer-maturity (average-price) forwards (and futures) and options on the spot² and shorter-maturity forwards or futures. The situation on the Nord Pool market might clarify the distinction. On its Eltermin market we encounter on the one hand futures on individual days up to one week ahead, and futures on individual weeks up to 4 weeks ahead. Due to the short (1 to 7 days) averaging period, we should definitely take spikes and mean reversion into account for those futures. On the other hand we can trade monthly (4-weekly), quarterly and yearly futures that are traded quite some time ahead. Returns on the latter futures exhibit only dampened fluctuations hardly related to the spot price, so spikes or mean reversion can safely be ignored³. The futures price returns fulfil the conditions for normality relatively well, and hedging related options with forwards or futures is feasible to some degree. Consequently, standard arbitrage-based pricing-formulas may yield reliable results for options on longer-maturity forwards and futures.

Spot returns are clearly not lognormally distributed (see e.g. Lucia and Schwartz, 2002) and the standard option pricing formulas may yield totally incorrect outcomes. That's why we need a different type of pricing approach, especially for in-the-money and out-of-the-money options. This paper presents such an approach. It is similar in spirit to Clewlow and Strickland (1999), but extends its mean-reverting framework with the necessary spikes, modelled as a separate and independent regime⁴. At the same time it keeps the attractive feature of closed-

² With spot electricity prices, we mean the prices for electricity that are set one day in advance in spot markets or pool systems. We can have prices for time periods ranging from individual quarters of an hour, hours, blocks of hours (e.g. peak and off-peak) to daily averages. The latter are the most well known 'spot' prices.

³ One finding of Hilliard and Reis (1998) is that jumps and spikes in the spot price do not affect forwards, futures and futures options.

⁴ Deng (2000) also proposes to model spikes in a regime-switching model. He derives formulas for pricing futures, forwards and standard options, but requires simulations to evaluate the outcomes.

form formulas, where other spot price models (Bhanot (2000), Knittel and Roberts (2001), Huisman and Mahieu (2001), and Escribano, Pena and Villaplana (2002)) would require extensive simulations. The two regimes (one for the ‘normal’ process, one for the spikes) capture the systematic alternations between stable and unstable states of demand and supply. The price process takes the main dynamics of electricity prices into account, such as seasonality, mean-reversion and, most importantly, spikes. Furthermore, a major contribution of the model is that it allows for multiple consecutive spikes, which is important for risk management and derivative pricing purposes. Based on this spot price model we present closed-form formulas to price European-style options on spot electricity prices. We furthermore show how the underlying spot price model can be aligned with the observed forward curve in the market, which makes the option valuation consistent with market data.

The paper is built up as follows. First, we present an intuitive spot price model that incorporates the most prominent features of electricity spot prices: mean-reversion and spikes. Next, in section 3 we estimate the model parameters for Dutch APX baseload, peak and off-peak spot prices. In section 3 we also show how the model can be aligned with market forward curves in a practical manner that avoids the separate modelling of seasonalities and risk premia. Finally, section 4 presents closed-form formulas for a set of European-style options and describes how other types of options can be priced. The closed-form formulas are applicable to value European-style options on contracts that are characterized by a mean-reverting ‘spiky’ price process, such as individual hours, block hours (e.g. peak, off-peak, super-peak), daily averages (baseload), and weekly averages. We end with some concluding remarks.

2. Spot Price Model

In this section we specify the models (in discrete time) that we use to examine the dynamics of electricity spot prices. We sequentially move from a basic mean-reverting model and the regime-switching model of Huisman and Mahieu (2001) to a regime-switching model that incorporates independent spikes⁵. We focus so much on spikes, because they are the main source of risk in electricity markets and the main value-driver of options. Currently, modelling spikes correctly is a great challenge to researchers and practitioners in electricity markets. If we don't model spikes at all, or incorrectly, any risk management calculation or derivative valuation will be obsolete. De Jong and Huisman (2002) discuss how well our model describes electricity prices in comparison to other specifications. In section 3 of this paper we show what estimates we obtain on the Dutch (APX) spot market and use those estimates in section 4 to analyse option prices.

The first step in modelling electricity prices, P_t , is the separation of the predictable component from the stochastic component (Lucia and Schwartz, 2002).

$$\ln P_t = f(t) + \ln S_t \quad (1)$$

The first component, $f(t)$ accounts for predictable regularities, such as any genuine periodic behaviour and any trend, and is a deterministic function of time. Seasonalities can be modelled with for example sinusoidal functions or with dummy's for different seasons, to which we come back later in section 3. The stochastic second component, S_t , is the more interesting and we continue with its specification below. In the remaining we refer to the stochastic part S_t as the “spot price”, but remember that in fact it is the spot price from which predictable trends are removed.

2.1 Mean-Reverting Model

Suppose we have a time-series of spot-power prices S_t , $t = 1, \dots, T$. A standard mean-reverting process has the following specification:

⁵ De Jong and Huisman (2002) analyse regime models more extensively, with different datasets and different specifications, and compare them to stochastic jump and GARCH-models. They show that this regime-specification is a robust approach to model spikes. In this paper we only use the APX-data as an example and focus on the implications for option pricing, rather than the optimal specification or fit.

$$\ln S_t = \ln S_{t-1} + \alpha(\mu - \ln S_{t-1}) + \varepsilon_t \quad (2)$$

where $\varepsilon_t \stackrel{i.i.d.}{\sim} N(0, \sigma^2)$

The parameter μ is the long-run stationary level for the natural logarithm of spot prices. It determines to what value spot prices converge. The parameter α measures the speed of convergence from the current to the long-run level and can be related to the concept of half-life, which is well known in physics: the time it takes to move on average halfway from the current level to the long-term level. The parameters of the mean-reversion specification can be estimated by maximum likelihood, using the fact that the error terms are normally distributed. The loglikelihood function equals:

$$LL_M(S_t, \theta_M) = \sum_t \left\{ -\frac{1}{2} e_t^2 - \ln \sigma - \frac{1}{2} \ln 2\pi \right\} \quad (3)$$

where $e_t = \ln S_t - \ln S_{t-1} - \alpha(\mu - \ln S_{t-1})$
 $\theta_M = \{\alpha, \mu, \sigma\}$

2.2 Mean-Reverting Model with Spikes

The mean-reverting specification above is relatively successful in modelling commodities such as oil and gas⁶. It is also frequently applied to electricity markets, but there it performs considerably worse, because of the so-called 'spikes' that are inherent to electricity prices. This 'spiky' behaviour of electricity prices has mainly to do with the non-storability of the commodity and the relative inelasticity of demand and supply, as we discussed in the introduction. Due to the spikes the mean-reverting model performs poorly for electricity. Parameter calibration generally leads to unrealistically high volatility, sometimes incorrect mean reversion parameters and too high levels to which the spot prices would converge. Modelling spikes in a satisfactory framework has turned out to be a major challenge for researchers and practitioners in electricity markets. The most common approach is the addition of a jump diffusion process to the mean-reverting process. Jump models allow for sudden extreme returns and are quite successful in stock markets, but do not incorporate an important characteristic of electricity prices: spikes are relatively short-lived. With a jump diffusion

⁶ See for example Pindyck (1999)

process we allow for large price movements, but we do not deal well with the fact that after a spike prices simply bounce back to normal levels (see Huisman and Mahieu, 2001). In a mean-reverting jump diffusion process this can only be achieved by an unrealistically high mean reversion parameter that forces prices back to normal levels after a spike.

Regime-switch models have the potential to solve some of these deficiencies, since they allow for distinct time-series behaviour in different periods of time. The basic regime model has the following specification (Hamilton, 1989):

$$\ln S_t = \mu_{\lambda_t} + \varepsilon_t \quad \text{where} \quad \varepsilon_t \sim N(0, \sigma_{\lambda_t}^2) \quad (4)$$

Here λ_t is a latent variable representing the regime of the process in time period t . The process can thus be in one of the regimes at each time t . Huisman and Mahieu (2001) propose a regime-shift model with three regimes: there is a mean-reverting regime with moderate mean-reversion and volatility, an initial jump regime that models the process when prices suddenly increase or decrease, and a subsequent jump regime, that describes how prices are forced back to the stable regime. The two jump regimes both have a more extreme expected return and volatility than the mean-reverting regime. The subsequent jump regime has a zero probability of occurrence if prices in the previous time period are in the mean-reverting regime, but a probability of one if they are in the initial jump-regime.

The main drawback of the Huisman and Mahieu (2001) model is that it does not practically allow for multiple consecutive jumps⁷, which are a frequent phenomenon in electricity markets (see figure 2) and crucial for risk management purposes and derivative valuation. Therefore, we introduce a model with only two regimes: a stable mean-reverting regime, and a spike regime. It might seem surprising that the omission of one regime gives the model the flexibility to capture consecutive jumps. However, we don't need a third regime to pull prices back to stable levels, because we assume that prices in the two regimes are independent from each other. Put differently, if there is a generator outage for example, prices may be high for some time period, but once the generator is repaired, prices continue as normal. We believe this regime specification fits well with the structure of electricity markets. As a side effect of the independence of the two processes, we can combine closed-form formulas of mean-reverting time series, with those of lognormally distributed spikes to simplify derivative valuations and to circumvent time-consuming Monte Carlo simulations.

⁷ Multiple consecutive jumps may theoretically be incorporated, but would require the estimation of a large number of switching probabilities.

Our first regime, which we call the mean-reverting regime, equals that in equation (1). The spikes in the second regime are modelled with a simple lognormal distribution whose standard deviation and mean⁸ are higher than those of the mean-reverting process. We have the following specifications:

$$\text{Mean-reverting regime:} \quad \ln S_{M,t} = \ln S_{M,t-1} + \alpha(\mu - \ln S_{M,t-1}) + \varepsilon_{M,t} \quad (5)$$

$$\varepsilon_{M,t} \sim N(0, \sigma_M)$$

$$\text{Spike regime:} \quad \ln S_{S,t} = \mu_S + \varepsilon_{S,t} \quad (6)$$

$$\varepsilon_{S,t} \sim N(0, \sigma_S)$$

At any point in time the spot price S_t is either $S_{M,t}$ or $S_{S,t}$. However, the model specification does not end with the two regimes, since we haven't defined yet how we go from one regime to the other and thus their relative frequencies. For this we use a Markov transition matrix, which contains the probabilities of switching from one regime to the other. Since we have two regimes, our Markov transition matrix MT is a 2x2 matrix. The element in column j and row i contains the probability $\pi_{i,j}$ of going from regime i in period t to regime j in period $t+1$, which are the mean-reverting regime (1) and the spike regime (2) in our case.

$$MT = \begin{bmatrix} 1 - \pi_{MS} & \pi_{SM} \\ \pi_{MS} & 1 - \pi_{SM} \end{bmatrix} \quad (7)$$

If we are in the mean-reverting regime today, we know that with probability π_{MS} the next day is a spike, and with probability $1 - \pi_{MS}$ the mean-reverting regime continues. Similarly, we know that a spike is followed by another spike with probability $1 - \pi_{SM}$, and otherwise the mean-reverting regime resumes.

The parameters of the two regimes can easily be calibrated by maximum likelihood when we condition on the regimes. Based on the normal distribution of the error terms, the loglikelihoods have the following form:

⁸ So-called off-peak hours (when demand is low) are characterized by negative spikes, due to the abundance of supply relative to demand during those periods of the day. In those cases the mean of the spike becomes negative.

$$LL_{M,t} = -\frac{(\ln S_{M,t} - \ln S_{M,t-1} - \alpha(\mu - \ln S_{M,t-1}))^2}{2\sigma_M^2} - \ln \sigma_M - \frac{1}{2} \ln 2\pi \quad (8)$$

$$LL_{S,t} = -\frac{(\ln S_{S,t} - \mu_S)^2}{2\sigma_S^2} - \ln \sigma_S - \frac{1}{2} \ln 2\pi \quad (9)$$

In practice parameter estimation is not that easy: we don't know what regime reigns on each day. The type of regime is only latent and thus not directly observable. The Kalman filtering methodology circumvents the unobservability problem and uses the prior and posterior beliefs to weigh each likelihood properly in the objective function. In fact, on each day, the likelihood of each regime is weighed by its prior probability. A prior is then updated into a posterior probability after we observe today's price. Finally, the posterior beliefs are multiplied by the Markov transition matrix to determine the posterior beliefs of next day's regimes. See for example Harvey (1989) for a description of this methodology that is the common method to estimate latent variables.

The specification for spot prices does have an extra complexity: apart from containing latent regimes the model contains latent prices for the following reason. The mean-reverting price at time t depends on the mean-reverting price at time $t-1$. When the spike at $t-1$ is a spike, we assume that the price in the mean-reverting regime continues as a latent process, and is thus unobservable. Remember that spikes are independent from the other prices, so previous spikes do not influence the magnitude of later spikes. However, previous prices in the mean-reverting regime $S_{M,t}$ do have an impact on later mean-reverting prices. This poses a difficulty in the parameter estimation, which can be solved as follows. Suppose we are at time t . Then we use the model's posterior probabilities to calculate for each of the k last prevailing spot prices (S_{t-1}, \dots, S_{t-k}) the probability that it was the last mean-reverting regime price. Next, we calculate the likelihood in equation (8) assuming alternatively that the price in period $t-1, \dots, t-k$ was the last mean-reverting price. If we look i periods back, in the likelihood equation we use $E[\ln(S_{M,t}) | S_{M,t-i}]$ and the appropriate (higher) variance $\text{Var}[\ln(S_{M,t}) | S_{M,t-i}]$, to capture the effect that prices are stochastically mean-reverting. These expected values and variance can be determined recursively as follows:

$$E[\ln S_{M,t} | S_{M,t-1}] = \alpha\mu + (1 - \alpha) \cdot E[\ln S_{M,t-1} | S_{M,t-2}] \quad (10)$$

$$\text{Var}[\ln S_{M,t} | S_{M,t-1}] = (1 + (1 - \alpha)^2) \cdot \text{Var}[\ln S_{M,t-1} | S_{M,t-2}] \quad (11)$$

This completes the specification of the mean-reverting regime model with independent spikes. In the next section we evaluate the parameters estimated from Dutch APX prices⁹.

⁹ See www.apx.nl

3. Model Estimation Results

This section reports the estimation results of the models presented in the previous section. We discuss some data issues and evaluate the parameter estimates to see whether the regime-switching model picks up mean-reversion and spikes sufficiently well. We use those parameter estimates to price options on spot prices in the next section.

3.1 Data

We use data from the Dutch spot market APX. The Dutch power market was liberalized for large consumers on January 1st 2001. From that day onwards the wholesale day-ahead prices on the Amsterdam Power Exchange (APX) reflect the forces of demand and supply. Volume on the electronic exchange has increased steadily, notably since January 2002, when medium-sized consumers became free to choose their energy-supplier as well. By June 2002 the APX-volume represented approximately 14% of total electricity consumption. We use data from January 2nd 2001 till June 30th 2002 of baseload, peak and off-peak day-ahead prices, totalling 545 observations for each index. The exchange defines baseload prices as the equally weighted average of the 24 individual hourly prices. The peak prices are the average of hour 8 till 23 (7:00 – 23:00); the off-peak prices are the average of the 8 remaining hours 1-7 and 24 (0:00 – 7:00, 23:00 – 24:00).

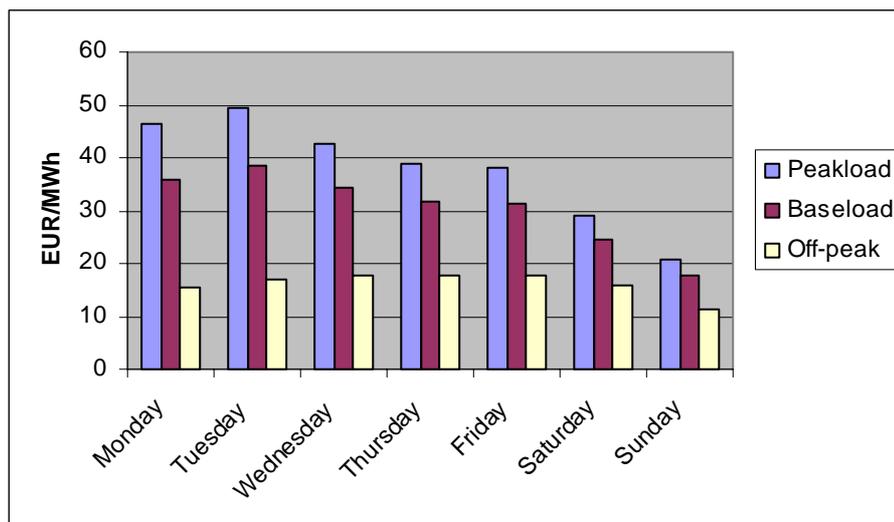


Figure 1. Average APX prices on individual weekdays in the period 2 January 2001 – 14 June 2002 for baseload, peak and off-peak hours.

The division in baseload, peak and off-peak prices reflects part of the seasonality during a day, but prices also exhibit considerable seasonality during a week (see figure 1). In general, prices (and electricity consumption) are lower during the weekend, especially on Sundays. At first sight there seems to be a downward trend from Monday to Friday, but this is probably due to some outliers¹⁰. If we deseasonalise the natural logarithm of spot prices by reducing them with their weekday average (and supplementing them with the overall average) total variance is reduced on average with 20%: weekday influences explain about one fifth of total variance. In terms of our mathematical formulation (equation (1)) $f(t)$ includes a dummy for Saturdays, and a dummy for Sundays. In the Dutch spot prices we found only very weak evidence of seasonality over the year, so no specification is included for it. The seasonal component $f(t)$ is estimated jointly with the stochastic model parameters.

Prices in the Dutch market have witnessed already some serious spikes (see figure 2). For example, the summer of 2001, which was expected to be a quiet period, contained some unexpected price movements. Those spikes even triggered an official investigation, but no irregularities were reported: it was reported to be a simple coincidence of generator outages. Another series of high prices were observed close to the end of the year 2001. Since January 2002 prices were relatively stable till the second week of June, but June ended with some high prices.

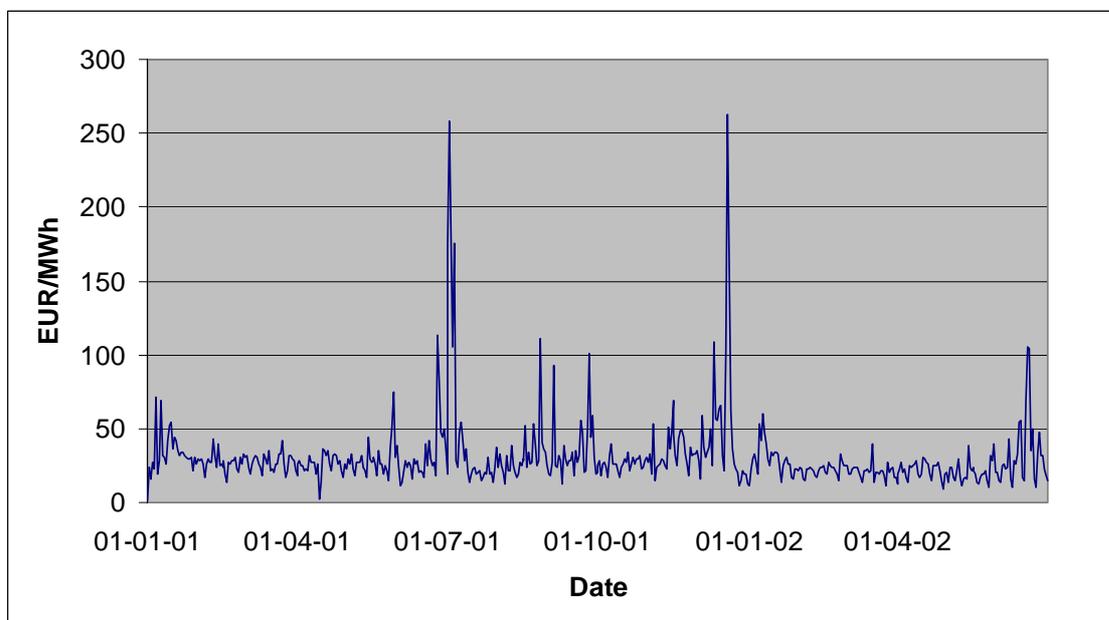


Figure 2. Baseload APX prices in the period 2 January 2001 – 30 June 2002.

¹⁰ Median prices for each weekday very stable from Monday to Friday, so the differences in average prices are mainly a result of a few outliers (or spikes).

APX-prices are no exception to the phenomenon that prices in electricity markets are different from those in most other financial markets. This is clearly reflected in the summary statistics for the baseload, peak and off-peak returns (see table 1). All three series are characterized by a relatively high daily standard deviation of between 35 and 42%. For comparison: most individual stocks have daily standard deviations of 1-2% and only exceed 5% for the most risky stocks. The daily fluctuations in APX-prices can be enormous, reaching levels of over 200%. Especially the off-peak hours, with sometimes very low price levels, exhibit tremendous outliers. The extraordinary high kurtosis level of off-peak returns indicates that the fourth moment is probably not even defined. Returns are left-skewed, which may be surprising at first sight, since *prices* are clearly right-skewed due to the spikes. This indicates that prices do not only spike upwards very fast, but come down even faster, as the minimum and maximum returns indicate as well.

Table 1: Summary Statistics APX Returns

	Baseload	Peak	Off-peak
Average	0.0%	0.0%	0.1%
Standard deviation	35.9%	39.5%	41.2%
Minimum	-204.4%	-186.2%	-452.3%
Maximum	161.3%	179.0%	396.3%
Skewness	-0.17	-0.07	-0.78
Kurtosis	4.60	3.80	45.19

This table presents summary statistics for the daily logreturns of APX day-ahead baseload, peak and off-peak indices in the period 1 January 2001 till 30 June 2002. Weekday influences were first removed from the price series before calculating the statistics.

3.2 Results

We use the sample with one and a half year of daily observations and three time-series to estimate model parameters. Even though the loglikelihoods of the regime-switching models are relatively complex, convergence was never a problem and independent of starting values. Results (table 2) indicate that the regime models improve the fit considerably compared to the mean-reverting specification. Apparently, spikes that destroy the loglikelihood in the mean-

reverting model are ‘transferred’ to the different spike and jump regimes where they do a much better job.

Our model picks up on average 50% more spikes than the Huisman-Mahieu¹¹ model. This is probably explained by the fact that their model requires that an up-jump is immediately followed by a down-jump, and is thus more restrictive on jumps. The mean-reverting components of the two regime models are very similar in terms of mean-reversion speed and mean-reversion level. Moreover, since the regime models remove spikes from the stable process, their volatilities are considerably lower compared to the pure mean-reverting model. The mean-reverting volatility in our model is lowest, because it transfers most erratic prices to the spike regime. The regime models indicate that the long-run average target levels for the baseload and peak spot prices are 4 and 7 €/MWh lower compared to the pure mean-reverting model, which is an economically considerable difference.

Our regime specification picks up spikes well: expected spikes are positive for baseload and peak, negative for off-peak¹², and have a much higher volatility than the stable mean-reverting process. When prices were mean-reverting on the previous day, a spike can be expected with a probability of 10-13%. Overall, between 20 and 30% of the prices are spikes, since spikes last on average 2-3 days. For the pricing of far-out-of-the-money options, it might be considered that our regime model picks up small spikes too easily. This is a common problem as well for jump models, and partly related to parameter calibration with maximum likelihood. It is possible to mitigate this problem with simple weights in the likelihood function.¹³

The spikes in both regime models have an expected magnitude ($E[S_i|\text{spike}]$) of 50-65 €/MWh, and have large possible swings. There is no doubt that the spikes deviate largely from the stable price levels and the data show it is crucial to separate them properly. For example, in the mean-reverting model a baseload price of over 100 €/MWh is nearly impossible (around 0.02% probability). In our regime-switch model its probability of 2.2% is much closer to the observed frequency of 2.6%. Moreover, for risk management purposes and derivative pricing, it is important that a model not only allows for spikes, but also for multiple consecutive spikes.

¹¹ In the Huisman-Mahieu model we count the frequency of spikes as the sum of the up-and down-spikes.

¹² See footnote 8.

¹³ We tested that a very small adaptation to the loglikelihood function, that disfavours spikes, can reduce the frequency of spikes by over 50%, while at the same time keeping the loglikelihood within a few basis points from its maximum. We disfavoured spikes, by increasing the volatility in the $\ln(\sigma_t)$ term by 25% in the loglikelihood specification (8).

Table 2: Parameter estimates

		baseload	peakload	offpeak
Mean-Reverting Model	α	0.384	0.398	0.520
	μ	3.407	3.600	2.792
	σ	0.323	0.355	0.304
	$E[S_t]$	31.794	38.963	17.084
	Sunday	-0.590	-0.630	-0.410
	Saturday	-0.239	-0.280	-0.066
	loglikelihood	-0.289	-0.384	-0.230
Huisman and Mahieu (2001) Regime Model	α	0.395	0.392	0.281
	μ_M	3.329	3.491	2.868
	σ_M	0.196	0.203	0.148
	$E[S_t \text{no spike}]$	33.938	40.176	20.414
	Sunday	-0.515	-0.536	-0.297
	Saturday	-0.237	-0.272	-0.102
	μ_S	0.443	0.489	-0.530
	σ_S	0.587	0.599	0.692
	$E[S_t \text{spike}]$	51.616	64.009	13.170
	π_{MS}	0.088	0.109	0.082
	loglikelihood	-0.111	-0.204	0.093
Mean-Reverting Regime Model with Spikes	α	0.252	0.239	0.156
	μ_M	3.304	3.452	2.856
	σ_M	0.145	0.130	0.119
	$E[S_t \text{no spike}]$	27.50	31.83	17.516
	Sunday	-0.472	-0.481	-0.293
	Saturday	-0.225	-0.263	-0.105
	μ_S	3.678	3.870	2.356
	σ_S	0.685	0.672	0.541
	$E[S_t \text{spike}]$	50.04	60.12	12.208
	π_{MS}	0.107	0.127	0.105
	π_{SM}	0.353	0.290	0.520
loglikelihood	-0.101	-0.168	0.187	

This table presents the parameter estimates and loglikelihood of three different time series models for APX day-ahead baseload, peak and off-peak hours. Estimates were obtained by maximum likelihood using data from 2 January 2001 till 30 June 2002. See the text for an explanation of all parameters and symbols.

4. Option Valuation

In this section we discuss the valuation of European-style options on spot prices using the proposed spot price model with a mean-reverting and spike regime. We present closed-form formulas for standard calls and puts. Pricing of caps, floors and swaptions is then straightforward as well. Closed-form formulas are important for various reasons. First, closed-form formulas may be more insightful than simulation-based calculations. For example, we will obtain two option value parts: one related to the mean-reverting process, one to the spikes. Second, electricity traders often need to get quick answers in their day-to-day activities on the relative pricing of different options in the market. For them speed is often so important that it is necessary to use closed-form formulas instead of simulation-based methods. Finally, closed-form formulas are very useful if options are being valued in a risk management application. Risk management statistics, such as Value-at-Risk can be computed much more quickly if we do not need simulations for individual products in the portfolio.

Our results apply to European-style options on the spot price, which excludes a range of options that are traded in the marketplace. In electricity markets we also observe for example options on average-price forwards¹⁴, and early exercise is sometimes allowed (American-style options). At the end of this section we describe briefly how these kinds of options can be valued, especially American-style options, within or without our regime framework.

4.1 The forward curve

Before we can price options, we have to describe how we align the spot price model with the market forward curve. As shown in Clewlow and Strickland (1999), the mean-reverting specification can be made consistent with the current forward curve $F_{0,\tau}$ (and thus at the same time risk-neutral) if we make μ a function of maturity τ as follows, and take care of the predictable component $f(\tau)$ of the spot price model (equation (1)) :

$$\mu_\tau = \frac{\partial \ln F_{0,\tau}}{\partial \tau} + \ln F_{0,\tau} + \frac{1}{4} \sigma^2 (1 - \exp(-2\alpha\tau)) + \frac{f(\tau)}{\alpha} \quad (12)$$

¹⁴ In most markets a forward contract for the period of July-02 for example, entails the delivery of a constant electricity flow during the whole month of July at a fixed price. The value of the forward does therefore not only depend on the price on some particular day in July, but on the *average* price during the whole of July.

In fact, we adjust the 'long-term' level over time to align the model with market expectations. The above adjustment avoids the tedious modelling and estimation of all seasonal influences and risk premia. For example, in our spot price data we could not find significant seasonal variations over the year, but the current forward curve suggests that prices in the winter will be higher than in the summer. With equation (12) seasonal influences that are reflected by the forward curve, can easily be implemented. This is extremely useful: we don't want our derivatives to deviate from market prices because we have a different view on forward prices, but because we have a different view on the spot price process, such as its volatility and level of mean reversion. For this reason we will apply a similar adjustment to the mean-reverting regime in our model as follows.

Suppose we are at time $t=0$ and the current market forward curve contains prices $F_{0,\tau}$ for maturities $\tau = 1, \dots, T$. First, we split this forward curve in a part that reflects the stable mean-reverting regime $FM_{0,\tau}$ and a part that reflects the expected spike, $FS_{0,\tau}$:

$$FS_{0,\tau} = \pi_{S,\tau} \cdot \exp\left(\mu_S + \frac{1}{2}\sigma_S^2 + f(\tau) + r\tau\right) \quad (13)$$

where $\pi_{S,\tau}$ is the probability that S_τ is a spike, and the second part of the equation is the expected magnitude of a spike valued forward with risk-free interest rate r . The probability of a future spike depends on the regime today and the regime switching probabilities. It can be calculated by τ times premultiplying today's posterior regime probabilities with the transition matrix MT , defined in equation (7). Next, the mean-reverting part of the forward curve is obtained if we deduct the spike part from the market forward curve:

$$FM_{0,\tau} = F_{0,\tau} - FS_{0,\tau} \quad (14)$$

Finally, the first order derivative of $FM_{0,\tau}$, with respect to maturity τ is calculated and entered into equation (12) to obtain a time-varying long-run average level of the mean-reverting regime that makes our spot price model consistent with observed market forward prices.

As an example we apply the above-described procedure to the forward curve in the Dutch market on 1 July 2002 (see table 3). The forward values, except for the spot day-ahead, depend on the average spot prices in the respective periods. This is slightly different from the definition in our model of a forward price maturing on one single day. As a solution we assume the above prices are the expected spot prices in the middle of the period and calculate the

number of days τ till those dates (see table 3). Based on the moderate preceding weekend prices, we assume a prior spike probability of 0%¹⁵. Now we construct the ‘spike-free’ forward curve $FM_{0,\tau}$ as follows. First we calculate the probability of a spike for each maturity τ based on the initial spike probability and the Markov transition matrix MT. Second, we calculate the expected magnitude of a spike and multiply this by the probability of a spike (equation 13) to obtain the ‘spike-component’ of the forward curve. The ‘spike-free’ component of the forward curve $FM_{0,\tau}$ is now obtained by subtracting the spike-component of the forward curve from the market forward curve (equation 14). Finally, we determine the mean-reverting level μ_τ for the spot price model in each period using equation (12), where we fill in $F_{0,\tau} = FM_{0,\tau} / (1 - \pi_{S,\tau})$. This completes the procedure that aligns the spot price model to the market forward curve.

¹⁵ Advancing our model till 1 July 2002 yielded a similar probability.

Table 3: Forward Curve Construction

Contract	τ	Market Price	Pr[Spike]	E[Spike]	FS	FM	μ
Panel A: Baseload contracts							
Day ahead	1	33.72	10.70%	50.03	5.35	28.37	3.07
Week ahead	7	32.50	23.09%	50.03	11.55	20.95	2.93
July	15	30.25	23.26%	50.03	11.64	18.61	2.81
August	46	28.43	23.26%	50.03	11.64	16.79	2.71
Q4-02	138	30.13	23.26%	50.03	11.64	18.49	2.81
Q1-03	230	31.55	23.26%	50.03	11.64	19.91	2.88
Q2-03	319	34.25	23.26%	50.03	11.64	22.61	3.01
Panel B: Peakload contracts							
Day ahead	1	42.69	12.70%	60.09	7.63	35.06	3.25
Week ahead	7	41.75	30.05%	60.09	18.06	23.69	3.09
July	15	40.75	30.46%	60.09	18.30	22.45	3.05
August	46	41.25	30.46%	60.09	18.30	22.95	3.07
Q4-02	138	43.00	30.46%	60.09	18.30	24.70	3.14
Q1-03	230	46.25	30.46%	60.09	18.30	27.95	3.27
Q2-03	319	52.13	30.46%	60.09	18.30	33.83	3.46

Table 3 describes the process of moving from market forward prices to a spot price model that is in line with the market. The first column contains the contract type, the second column the time-to-maturity in days. The third column lists the Dutch forward prices in €/MWh on 1 July 2002 for baseload and peak electricity respectively. Source: Platts Dutch Power Assessments delivered through Moneyline/Telerate. The other columns combine the parameter estimates with the market prices to show the probability of a spike (Pr[Spike]), the expected magnitude of a spike (E[Spike]), the ‘spike-component’ of the forward curve (FS), the ‘Mean-reverting component’ of the forward curve (FM) and the time-varying mean-reverting level in the spot price model (μ).

4.2 European Options on the Spot

In this subsection we show how to value European-style options on the spot and the related valuation of caps and floors (maximum and minimum price guarantees). It should be noted that the value of an option on a forward that matures at one point in time is the same as on the spot, since the value of such forwards equals the spot price at maturity. So the procedure described here is applicable as well to those types of forwards. The idea behind the valuation is to split up the option price in a mean-reverting component and a spike component, the same way as we split up the forward curve. For the valuation of the mean-reverting component, we apply Black’s (1976) formula the same way as Clewlow and Strickland (1999) and Lucia and Schwartz (2002). For the valuation of the spike-component we apply Black’s (1976) result again, but with different inputs. The weight that each component receives, is determined by the

probability of a spike. The reason that the two option components can be added up, is that the spikes are independent from the mean-reverting prices. For example, a European-style call option with maturity τ and strike X has a ‘fair’ value of (ignoring time-value):

$$\begin{aligned}
Call &= E[\max(0, S_\tau - X)] \\
&= \pi_{M,\tau} \cdot E[\max(0, S_\tau - X) | \lambda = M] + \pi_{S,\tau} \cdot E[\max(0, S_\tau - X) | \lambda = S] \\
&= \pi_{M,\tau} \cdot Call_M + \pi_{S,\tau} \cdot Call_S
\end{aligned} \tag{15}$$

where $\pi_{M,\tau}$ and $\pi_{S,\tau}$ are the probability of a mean-reverting regime and spike regime at maturity τ respectively, and λ indicates the type of regime. It is important to understand that τ periods from now, the spot price is not a weighted *combination* of mean-reverting price and spike price, but either a mean-reverting price *or* a spike price.

We continue with a brief review of option pricing when prices are mean-reverting, then we explain the addition of spikes. In electricity markets longer term forwards have a significantly lower volatility than the spot price or short term forwards. This is because the expectation about spot prices varies considerably more than the expectation about price levels in the far future. This phenomenon creates a term structure in volatility: each maturity has its own volatility, which declines progressively. Mathematically, this term structure is a result of mean-reversion in prices, which exemplifies the importance of reliable mean-reversion estimation. Clewlow and Strickland (1999) and Lucia and Schwartz (2002) show how volatility responds to different maturities in a mean-reverting framework. They show that the volatility of a forward with maturity τ equals:

$$vol_\tau = \frac{\sigma_M}{2\alpha\tau} (1 - \exp(-2\alpha\tau)) \tag{16}$$

Straightforward application of the Black (1976) result yields the value of a European call with maturity τ and strike price K :

$$d = \frac{\ln\left(\frac{F_{0,\tau}}{K}\right) + \frac{1}{2} vol_\tau^2 \cdot \tau}{vol_\tau \cdot \sqrt{\tau}} \tag{17a}$$

$$Call(\tau, K, F_{0,\tau}, vol_\tau) = \exp(-r\tau) \cdot \{F_{0,\tau} N(d) - KN(d - vol_\tau \sqrt{\tau})\} \tag{17b}$$

The value of a European put on the spot can be obtained by put-call parity.

We have shown how the mean-reverting model can be extended with independent lognormally distributed spikes. We now show how this affects option valuation. The idea is that the option value must be broken down in a mean-reverting and a spike component. The first component can be calculated with equation (17), where we replace the forward price by the mean-reverting forward price, which is equation (14) divided by the probability of a mean-reverting price.

$$Call_M = Call\left(\tau, K, \frac{FM_{0,\tau}}{\pi_{M,\tau}}, vol_\tau\right) \quad (18)$$

The second component can be calculated with Black's (1976) formula as well. The call option pay-off on this spike equals the maximum of 0 and the difference between spike and strike. Because the spike is lognormally distributed with volatility σ , we can use Black's formula with a forward value equal to the expected spike, a volatility that is independent of time, and a discount period of τ .

$$E[Spikes] = \exp\left(\mu + f(\tau) + \frac{1}{2}\sigma_s^2\right) \quad (19a)$$

$$d' = \frac{\ln\left(\frac{E[Spikes]}{K}\right) + \frac{1}{2}\sigma_s^2}{\sigma_s} = \frac{\mu + f(\tau) - \ln(K) + \sigma_s^2}{\sigma_s} \quad (19b)$$

$$Call_s = \exp(-r\tau) \cdot \{E[Spikes]N(d') - KN(d' - \sigma_s)\} \quad (19c)$$

Since the prices in the two regimes are independent from each other, the actual option value is the sum of the two components, weighted by their respective probabilities (equation 15).

The option values that result from our regime-switch spot price model largely deviate from options in a mean-reverting framework. In table 4 we make a comparison. Option values were calculated of call options that mature in 1, 7, 15 and 46 days, with maturities of 20, 30, 40 and 50 €/MWh. We take the parameter estimates and forward values of 1 July 2002 (table 2 and 3) and assume an interest rate of 0% for simplicity.

Table 4: Call Option Values

Maturity (days)	Mean-Reverting Model				Regime-Switch Model				
	Strike (€/MWh)	20.00	30.00	40.00	50.00	20.00	30.00	40.00	50.00
Panel A: Baseload options									
1	13.72	4.17	0.24	0.00	13.81	4.26	1.89	1.44	
7	12.50	3.36	0.19	0.00	12.70	5.55	4.09	3.10	
15	10.25	1.43	0.01	0.00	10.45	5.45	4.11	3.12	
46	8.43	0.19	0.00	0.00	8.64	5.45	4.11	3.12	
Panel B: Peakload options									
1	22.69	12.70	3.99	0.48	22.75	12.96	4.11	2.55	
7	21.75	11.79	3.55	0.46	21.89	12.46	7.66	6.02	
15	20.75	10.75	2.25	0.08	20.89	11.51	7.75	6.10	
46	21.25	11.25	1.81	0.00	21.39	11.91	7.74	6.10	

This table shows values of European-style options on power spot prices according to the mean-reverting model (panel A) and the regime-switch model (panel B). Parameter estimates are taken from table 2. Forward prices are taken from table 3. Option values are calculated for four different maturities and strikes of 20, 30, 40 and 50 €/MWh.

In the mean-reverting model volatilities decline exponentially to zero with increasing maturity. The theoretical volatility of the baseload contract for example declines progressively from over 13% for the day-ahead contract to less than 1% for the August contract (less than 2 months ahead). The effect of the declining volatility on option values is strong: the mean-reverting model suggests there is hardly any option value (above the forward value) in options on spot prices that are more than 5 or 10 days ahead, since prices always revert back to a long-run level. The regime-switch model however takes into account that the trend might be towards such a long-run level but still prices might deviate from it on individual days due to spikes. The spikes imply that also longer-term options have substantial value.

Similarly, the mean-reverting model suggests that deep out-of-the-money options have hardly any value. The difference between the two models is best understood if we consider options with a strike price of 50 €/MWh. The mean-reverting model indicates those options are close to worthless, although we know that it is certainly not impossible that spot prices will reach levels above 50 €/MWh on individual days. As an illustration, in our sample baseload prices exceeded this level in more than 8% of the days and peakload prices in more than 13% of the days. So even options with high exercise prices have substantial value, which is entirely ignored by the mean-reverting model. For example, the costs of a maximum price (cap) would

severely be underestimated with the mean-reverting model. Caps are equal to a series of call options and frequently embedded in retail electricity contracts, where they form a bridge between fixed and floating price contracts. Let's consider a contract where the end-user pays the daily baseload APX-price on each day in July, but with a cap of 50 €/MWh. If we take the possibility of spikes into account, such a cap would cost approximately 3.12 €/MWh (based on an average maturity of 15 days), whereas a supplier would give it away for free if the wrong model were being used.

5. Concluding Remarks

In this paper we have presented a model to value options on electricity spot prices. It takes into account the two main features of electricity prices: strong mean-reversion and occasional ‘spikes’. Closed-form formulas for European-style options were obtained by disentangling the mean-reverting spot prices from the spikes, such that option values can be broken down in two components that were valued with Black’s (1976) formula for options on forwards and futures. We showed that it is crucial to include spikes in any option price formula, since they represent substantial value, especially for deep out-of-the-money options.

Our results apply to European-style options on the spot price, which excludes a range of other tradable options. In electricity markets we also observe for example options on individual hours, options on forwards, and early exercise is sometimes allowed (American-style options). We believe a regime model could work well for individual hours, since these are characterised by periods of spikes and stable periods of mean-reverting prices as well. A difficulty here is that individual hours are even more volatile and spiky than daily averages, exhibit strong seasonality, and that prices on the same day strongly interact.

Average price forwards on the other hand are hardly affected by the presence of spikes, as long as the averaging period is long enough. The pricing of options on such forwards can therefore best be done by modelling the forwards directly, instead of aggregating spot prices into average price forwards. Moreover, the main uncertainty for Asian-style options emanates from variations in the long-run average price. Our model is a one-factor model, which means that the long-run average price is constant. This is not so relevant for the valuation of options on the spot, but it would be incorrect to use the same one-factor model for long-term Asian-style options. So, for those we need to model the forwards directly.

Early exercise is a difficulty in the valuation of all American-style options, whether they are traded in financial markets or commodity markets, but can add substantial value (see e.g. Longstaff, Santa-Clara and Schwartz, 2001). The common solution is normal backwardation in binomial or trinomial trees to determine optimal exercise strategies. Clewlow and Strickland (1999) present procedures for the construction of a trinomial tree for mean-reverting prices. In our regime-switch model with their procedures it is possible to calculate the early exercise premium for the mean-reverting regime part, and add this to the value of our European-style option including spikes. As long as the possibility of a spike does not affect the early exercise premium, this yields the correct value of an American-style option whose

underlying is a mean-reverting spot price with spikes. If on the other hand we believe the spikes have a non-negligible impact on the early exercise premium, we could use a simulation approach as described in Longstaff and Schwartz (2001), but that is computationally much more intensive.

Our application of Black's formula implies that the risks of the mean-reverting prices and the spikes can be hedged, which is not completely realistic. With increasing liquidity in most electricity forward markets, it might be possible to hedge some part of the uncertainty in option prices, but the largest uncertainty of options on the spot result from the spikes, and there are no financial¹⁶ strategies to hedge spikes properly. Therefore, the derived option values can best be regarded as fair prices if uncertainty is ignored. However, market participants might be willing to price options somewhat higher, because spikes make especially the sale of options risky. The model makes such an adjustment relatively easy, since it yields an explicit value for the spike component of the option value. This component can be multiplied by some factor to include a risk premium.

The separation of the spikes from the mean-reverting prices ensures that only a limited number of parameters needs to be estimated. This is important in electricity markets where we have only a relatively short history of reliable prices, and markets are in constant change. As markets become more mature, it may be worthwhile to include more electricity price characteristics, such as time-varying volatility and time-varying spike intensities. The disadvantage of course would be the loss of closed-form and intuitive formulas, which we believe is one of the main strengths of the model.

¹⁶ A way to hedge spikes physically is by keeping some reserve capacity to use when prices are unexpectedly high.

6. References

- Bhanot, K., 2000, "Behavior of power prices: implications for the valuation and hedging of financial contracts", *Journal of Risk*, 2:3, p. 43-62
- Bjerksund, P., Rasmussen, H. and G. Stensland, 2000, "Valuation and risk management in the Norwegian electricity market", discussion paper 20/200, Norwegian School of Economics and Business Administration
- Black, F., 1976, "The pricing of commodity contracts", *Journal of Financial Economics*, 3, p. 167-179
- Black, F., and M. Scholes, 1973, "The pricing of options and corporate liabilities", *Journal of Political Economy*, 81, p. 637-659
- Clelow, L. and C. Strickland, "Valuing energy options in a one factor model fitted to forward prices", working paper, University of Sydney
- De Jong, C. and R. Huisman, 2002, "Electricity price spikes in independent regimes", working paper, Erasmus University
- Deng, S., 2000, "Stochastic models of energy commodity prices and their applications: mean reversion with jumps and spikes", working paper, University of California Energy Institute
- Fleten, S.E. and J. Lemming, "Constructing forward prices in electricity markets", working paper, Riso National Laboratory Denmark
- Hamilton, J.D., 1989, "A new approach to the economic analysis of non-stationary time series and the business cycle", *Econometrica*, 57, p. 357-384
- Harvey, A.C., 1989, "Forecasting, structural time series and the Kalman filter", Cambridge University Press, Cambridge
- Hilliard, J.E. and J. Reis, 1998, "Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot", *Journal of Financial and Quantitative Analysis*, 33:1, p. 61-86
- Huisman, R. and R. Mahieu, 2001, "Regime jumps in power prices", *Energy & Power Risk Management*, September
- Johnson and Barz, 1999, "Selecting stochastic processes for modelling electricity prices", *Energy Modelling and the Management of Uncertainty*, Risk Publications
- Kaminski, V., 1997, "The challenge of pricing and risk managing electricity derivatives", *The US Power Market*, p. 149-171, Risk Publications

Karesen, K.F. and E. Husby, 2000, "A joint state-space model for electricity spot and futures prices"

Knittel, C.R. and M. Roberts, 2001, "An empirical examination of deregulated electricity prices", *working paper*, University of California Energy Institute

Koekebakker, S. and F. Ollmar, 2001, "Forward curve dynamics in the Nordic electricity market", *working paper*, Agder University College, Denmark

Longstaff, F.A. and E.S. Schwartz, 2001, "Valuing American options by simulation: a simple least-square approach", *Review of Financial Studies*, 14:1, p. 113-147

Longstaff, F.A., P. Santa-Clara and E.S. Schwartz, 2001, "Throwing away a billion dollars: the cost of sub-optimal exercise strategies in the swaption market", *Journal of Financial Economics*, 62:1, p. 39-66

Lucia, J. and E.S. Schwartz, 2002, "Electricity prices and power derivatives: evidence from the Nordic Power Exchange", *Review of Derivatives Research* 5, p. 5-50

Pilipovic, D., 1998, "Energy risk: valuing and managing energy derivatives", McGraw Hill, New York

Pindyck, R.S., 1999, "The long-run evolution of energy prices", *The Energy Journal*, 20:2

Pirrong, C. and M. Jermakyan, 1999, "Valuing Power and Weather Derivatives on a Mesh Using Finite Difference Methods", *Energy Modelling and the Management of Uncertainty*, Risk Books

Pirrong, C. and M. Jermakyan, 2001, *working paper*

Schwartz, E.S., 1997, "The stochastic behavior of commodity prices: implications for valuation and hedging", *Journal of Finance* 52:3, p. 923-973

Miltersen, K. and E.S. Schwartz, 1998, "Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates", *Journal of Financial and Quantitative Analysis*, 33:1, p. 33-59

Publications in the Report Series Research* in Management

ERIM Research Program: "Finance and Accounting"

2002

A Stochastic Dominance Approach to Spanning

Thierry Post
ERS-2002-01-F&A

Testing for Third-Order Stochastic Dominance with Diversification Possibilities

Thierry Post
ERS-2002-02-F&A

Towards a Transaction Cost Theory of Management Control

Roland F. Speklé
ERS-2002-06-F&A

Modeling the Conditional Covariance between Stock and Bond Returns: A Multivariate GARCH Approach

Peter De Goeij & Wessel Marquering
ERS-2002-11-F&A

An Empirical Comparison of Default Swap Pricing Models

Patrick Houweling & Ton Vorst
ERS-2002-23-F&A

Relative Distress and Return Distribution Characteristics of Japanese stocks, a Fuzzy-Probabilistic Approach

Willem-Max van den Bergh, Onno Steenbeek & Jan van den Berg
ERS-2002-29-F&A

Does Risk Seeking Drive Asset Prices? A Stochastic Dominance Analysis of Aggregate Investor Preferences

Thierry Post & Haim Levy
ERS-2002-50-F&A

Reinventing The Hierarchy, The Case Of The Shell Chemicals Carve-Out

Michel A. van den Bogaard, Roland F. Speklé
ERS-2002-52-F&A

A Framework For Managing A Portfolio Of Socially Responsible Investments

Winfried Hallerbach, Haikun Ning, Aloy Soppe, Jaap Spronk
ERS-2002-54-F&A

The Relevance of MCDM for Financial Decisions

Winfried Hallerbach, Jaap Spronk
ERS-2002-69-F&A

A broadband vision of the development of the DAX over time

Winfried Hallerbach, Christoph Hundack, Igor Pouchkarev, Jaap Spronk
ERS-2002-87-F&A

* A complete overview of the ERIM Report Series Research in Management:
<http://www.ers.erim.eur.nl>

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship

Do Countries or Industries Explain Momentum in Europe?
Theo Nijman, Laurens Swinkels, Marno Verbeek
ERS-2002-91-F&A

Measuring Credit Spread Risk: Incorporating the Tails
Rachel Campbell, Ronald Huisman
ERS-2002-95-F&A

Option Formulas for Mean-Reverting Power Prices with Spikes
Cyriel de Jong, Ronald Huisman
ERS-2002-96-F&A