A Two-Factor Model for the Electricity Forward Market

Rüdiger Kiesel (Ulm University)
Gero Schindlmayr (EnBW Trading GmbH)
Reik H. Börger (Ulm University)

August 20, 2007

Abstract

This paper provides a two-factor model for electricity futures, which captures the main features of the market and fits the term structure of volatility. The approach extends the one-factor-model of Clewlow and Strickland to a two-factor model and modifies it to make it applicable to the electricity market. We will especially take care of the existence of delivery periods in the underlying futures. Additionally, the model is calibrated to options on electricity futures and its performance for practical application is discussed.

Keywords: Energy derivatives, Futures, Option, Two-Factor Model, Volatility Term Structure

1 Introduction

Since the deregulation of electricity markets in the late 1990s, power can be traded on spot and futures markets at exchanges such as the Nordpool or the European Energy Exchange (EEX). Power exchanges established the trade of forwards and futures early on and by now large volumes are traded motivated by risk management and speculation purposes.

Spot electricity is not a tradable asset which is due to the fact that it is non-storable. Thus, the spot trading and prices in electricity markets are not defined in the classical sense. Similarly, electricity futures show contract specifications, that are different to many other futures markets. In addition to the characterisation by a fixed delivery price per MWh and a total amount of energy to be delivered, power forward and futures contracts specify a delivery period, which will directly influence the price of the contract.

While much of the research in electricity markets focuses on the spot market (cp. Geman and Eydeland [1999] for an introduction, Ventosa et al. [2005] for a survey of modelling approaches, the monographs Eydeland and Wolyniec [2003], Geman [2005] for detailed overviews, Weron [2006] for a discussion of time-series characteristics and distributional properties of electricity prices), only few results are available for electricity futures and options (on such futures).

For commodity futures modelling approaches can broadly be divided into two categories. The first approach is to set up a spot-market model and derive the futures as expected values under a pricing measure. The best-known representative is the two-factor model by Schwartz and Smith [2000], which uses two Brownian motions to model short-term variations and long-term dynamics of commodity spot prices. The authors also compute futures prices and prices for options on futures. However, electricity futures are not explicitly modelled. In particular, the fact that electricity futures have a delivery over a certain period (instead of delivery at a certain point in time) is not taken into account. Thus, the applicability to pricing options on electricity futures is limited.
Several models more specific to electricity spot markets have extended the approach by Schwartz and Smith [2000] during the last few years. The typical model ingredients are a deterministic seasonality function plus some stochastic factors modelled by Lévy processes. Typical representatives are Geman and Eydeland [1999], who use Brownian motions extended by stochastic volatilities and poisson jumps, Kellerhals [2001] and Culot et al. [2006], who use affine jump-diffusion processes, Cartea and Figueroa [2005], in which a mean-reverting jump-diffusion is suggested, Benth and Saltyte-Benth [2004], who apply Normal Inverse Gaussian processes, which is extended to non-Gaussian Ornstein-Uhlenbeck processes in Benth et al. [2005]. Furthermore, a regime-switching factor process has been used in Huisman and Mahieu [2003]. All of these models are capable of capturing some of the features of the spot price dynamics well and imply certain dynamics for futures prices, but these are usually quite involved and difficult to work with. Especially, these models are not suitable when it comes to option pricing in futures markets, since the evaluation of option pricing formulae is not straightforward. In particular, none of the models has been calibrated to option price data.

The other line of research is to model futures markets directly, without considering spot prices, using Heath-Jarrow-Morton-type of models (HJM). Here Chapter 5 of the monograph by Eydeland and Wolyniec [2003] provides a general summary of the modelling approaches for forward curves, but it does not apply a fully specified model to electricity futures data. They rather point out that the modelling philosophy coming from the interest rate world can be applied to commodities markets in a refined version as well. The paper by Hinz et al. [2005] follows this line of research and provides arguments to justify this intuitive approach.

More direct approaches start with the well-known one-factor model by Clewlow and Strickland [1999] for general commodity futures. In principle, they can capture roughly term-structure features, which are present in many commodities markets, but their emphasis is on the evaluation of options such as caps and floors, the derivation of spot dynamics within the model and building of trees, which are consistent with market prices and allow for efficient pricing routines for derivatives on spot prices. Again, they do not apply their model to the products of electricity markets and do not discuss an efficient way of estimating parameters. A general discussion of HJM-type models in the context of power futures is given in Benth and Koekebakker [2005] (which can be viewed as an extension of Koekebakker and Ollmar [2005]). They devote a large part of their analysis to the relation of spot-, forward and swap-price dynamics and derive no-arbitrage conditions in power future markets and conduct a statistical study comparing a one-factor model with several volatility specifications using data from Nord Pool market. Among other things, they conclude that a strong volatility term-structure is present in the market. The main empirical focus of both papers is to assess the fit of the proposed models to futures prices. While there are further studies using variants of HJM-type model, e. g. Vehviläinen [2002], a successful application to the pricing of options on electricity futures is still lacking.

One of the contributions of this paper is to address the issue of this pricing problem. In order to obtain an option pricing formula, we will follow the second line of research, that is, we will model the futures directly. We extend Clewlow and Strickland [1999] and Benth and Koekebakker [2005], in that we propose an explicit two-factor model and fit it to option price data. Our main objective is to formulate a model and specify a certain volatility function, so that we are able to resemble the volatility term-structure. Thus, we will not use an Heath-Jarrow-Morton-type model, but rather set up a market model in the spirit of LIBOR market models in the interest rate world. We regard that as the second main contribution of the paper, since this approach enables us to model one-month futures as building blocks and derive prices of futures with different delivery period as portfolios of the building blocks. With this model we are not only able to price standardized options in the market but also to provide consistent prices for non-standard options such as options on seasonal contracts. In order to provide pricing formulae for options on futures with a variety of delivery periods, we use an approximation of the portfolio distribution and assess the quality of this approximation. We then provide option pricing formulae for all options in the market and use the market-observed option prices to infer the model parameters. General two-factor model specification can be found in Schwartz and Smith [2000], Benth and Koekebakker [2005] and Lucia and Schwartz [2002] as well, but the parameter estimation uses time series techniques. By inferring risk-neutral parameters we avoid all complications related to the specification of the market price of risk.

We will show, that our model is robust, captures the term-structure of volatility and includes the delivery over a period in futures and option prices. Thus, after the model has been calibrated to plain-vanilla calls it can be used to price exotic options consistently (as soon as such options are traded).

The remainder of the paper is organised as follows. The following section describes the EEX Futures and Option market to which our model is eventually calibrated. Section 3 develops our general modelling approach, while in section 4 a specific model is formulated. In section 5 we calibrate our model to market data and discuss the quality of the fit. Section 6 concludes.

2 The EEX Futures and Options market

An electricity future is the obligation to buy or sell a specified amount of power at a predetermined delivery price during a fixed delivery period. The contract (i.e. the delivery price) is set up such that initially no payment has to be made. While the delivery price is fixed, the price which makes the contract have zero value will change over time. This price is called futures price and is quoted at the exchanges.

The futures are standardized by the following characteristics: Volume, delivery period and settlement.

The volume is fixed to a rate (energy amount per hour) of 1 Megawatt (MW). For a delivery period of e.g. September, this means a total of 1MW x 30days x 24h/day = 720MWh. Quoted is the futures price per 1MWh. Smallest tick size is 0.01EUR per MWh.

The delivery periods are fixed to each of the 12 calendar months, the four quarters of the calendar year or the whole calendar year. When a year-contract comes to delivery, it is split up into the corresponding four quarters. A quarter, which is at delivery, is split up into the corresponding three months and only the month at delivery will be settled either physically or financially.

Additional to these baseload contracts, there are peakload contracts, which deliver during the day from 8am to 8pm Monday to Friday in the delivery period only. These are not considered in this paper.

At any fixed point in time, the next 6 months, 7 quarters and 6 years can be traded, but usually only the next 4 or 5 months, 5 quarters and 2 or 3 years show activity. Figure 1 shows the available forward prices at September 14, 2005. (This represents a typical trading day and will be used throughout the text.) One might observe seasonalities in the maturity variable T (not the time variable t), especially in the quarterly contracts: Futures during winter months show higher prices than comparable contracts during the

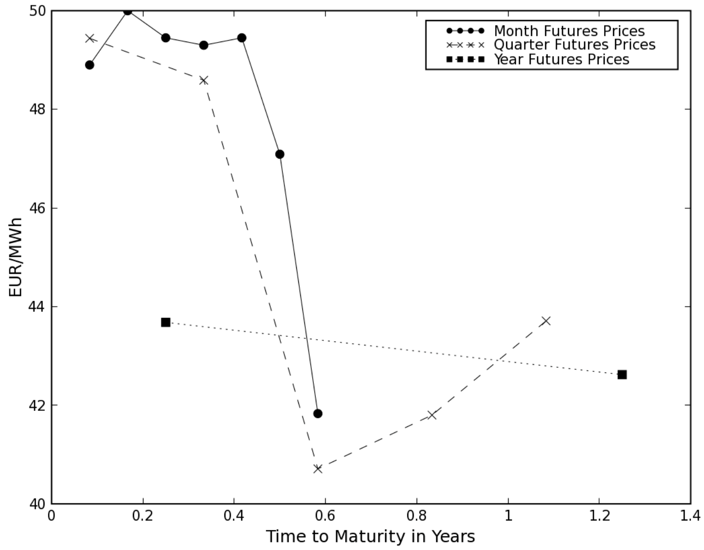

Figure 1: Forward prices of futures with different maturities and delivery periods

summer

All options under consideration are European call options on baseload futures described above, which can be exercised only on the last day of trading, which coincides with the options maturity. The maturity is fixed according to a certain scheme, but usually it is on the 3rd Thursday of the month before delivery. The options are settled by the opening of a position in the corresponding future. The option prices are quoted in EUR per MWh and the smallest tick size is 0.001EUR.

Available are options on the next five month-futures, six quarter-futures and three year-futures.

Especially the length of the delivery period and the time to maturity determine the value and statistical characteristics of the futures and options vitally. One can observe, that contracts with a long delivery period show less volatile prices than those with short delivery. This is called term structure of volatility and is present in most power futures markets. The term structure has to be modelled accurately in order to be able to price options on futures. Figure 2 gives an example of such a term structure for futures traded at the EEX. The figure shows the volatility of futures contracts, which is obtained by inverting the Black-formula that evaluates options on these futures. From an economic point of view it is clear, that futures with long delivery period are less volatile than those with short delivery, since the arrival of news such as temperature, outages, oil price shocks etc. influence usually only certain months of the year and

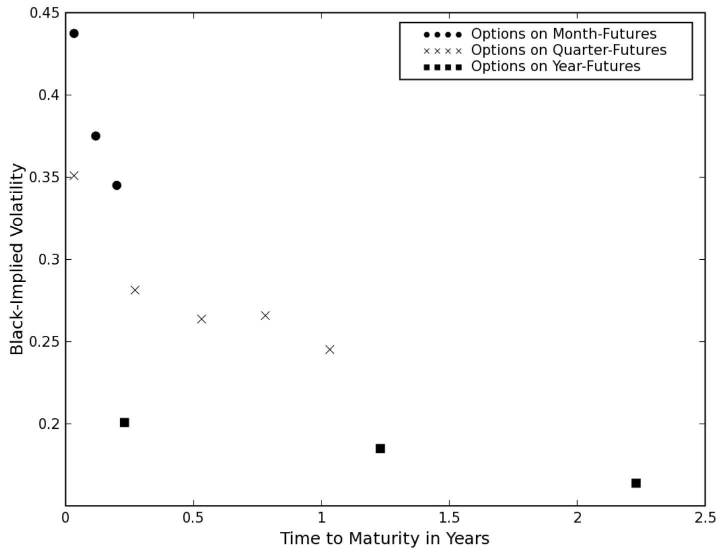

Figure 2: Option-implied volatilities of futures with different delivery periods and delivery starting dates

will average out in the long run with opposite news for other months. Only if all one-months contracts move in the same direction, the corresponding year contract will move as well. Furthermore, the arrival of news will accelerate when a contract comes to delivery, since temperature forecasts, outages and other specifics about the delivery period become more and more precise. Thus, the volatility increases.

We will show, that our modelling approach using one-month contracts as building blocks, will enable us to capture the term-structure of volatility and the influence of the period of delivery.

3 Description of the Model and Option Pricing

3.1 General Model Formulation

In energy markets we observe futures with different delivery periods. In the following, energy futures with 1-month delivery are the building blocks of our model. Note, that futures with other delivery period are derivatives now, i.e. a future with delivery of a year is a portfolio of 12 appropriate month-futures. For a full treatment of no-arbitrage conditions and the relations between electricity futures compare Benth and

Koekebakker [2005].

Let $F(t,T)$ denote the time t forward price of 1MWh electricity to be delivered constantly over a 1-month-period starting at T. Then, assuming a deterministic and flat rate of interest r, the time t value of this futures contract with delivery price D is given by

$$V^{future}(t,T) = e^{-r(T-t)}\left(F(t,T) - D\right).$$

D is the price for 1MWh electricity delivered constantly during the 1-month-period agreed upon at the time of signing the forward contract. Assuming the existence of a risk-neutral measure, discounted value processes have to be martingales under this measure, which in this case is equivalent to forward prices being martingales.

Thus, in the spirit of LIBOR market models, we model the observable forward prices directly under a risk-neutral measure as martingales via the stochastic differential equation

$$dF(t,T) = \sigma(t,T)F(t,T)dW(t),$$

where $\sigma(t,T)$ is an adapted d-dimensional deterministic function and $W(t)$ a d-dimensional Brownian motion. The initial value of this SDE is given by the condition to fit the initial forward curve observed at the market. This takes care of the seasonality in the maturity variable T.

The solution of the SDE is given by

$$F(t,T) = F(0,T)\exp\left(\int_0^t \sigma(s,T)dW(s) - \frac{1}{2}\int_0^t ||\sigma(s,T)||^2 ds\right)$$

where $||\cdot||$ is the standard Euclidean norm on $\mathbb{R}^d$.

3.2 Option Pricing

A European call option on $F(t,T)$ with maturity T_0 and strike K can be easily evaluated by the Black-formula

$$V^{option}(t) = e^{-r(T_0-t)}\left(F(0,T)\mathcal{N}(d_1) - K\mathcal{N}(d_2)\right), \tag{1}$$

where $\mathcal{N}$ denotes the normal distribution and

$$\begin{aligned} d_1 &= \frac{\log\frac{F(0,T)}{K} + \frac{1}{2}\mathrm{Var}(\log F(T_0,T))}{\sqrt{\mathrm{Var}(\log F(T_0,T))}} \\ d_2 &= d_1 - \sqrt{\mathrm{Var}(\log F(T_0,T))} \end{aligned}$$

We now use month-futures to describe futures with longer delivery period. In particular, options on year-futures are the most heavily traded products in the option market. To price such options, year-futures (i. e. futures with delivery period 1 year) are a portfolio of the building blocks, the month-futures. Thus, the pricing of an option on such a portfolio is not straightforward and a closed form formula is not known in general. The issue is closely related to the pricing of swaptions in the context of LIBOR market models and is discussed in Brigo and Mercurio [2001].

The value of a portfolio of month-futures (e. g. a year-future) with delivery starts at $T_i, i = 1, \dots n$ normalized to the delivery of 1 MWh and delivery price D is given by

$$V^{portfolio}(t, T_1, \dots, T_n) = \frac{1}{n} \sum_{i=1}^{n} e^{-r(T_i - t)} \left(F(t, T_i) - D \right).$$

In the context of interest rate swaps, the value of a swap is expressed in terms of a swap rate Y, which is here:

$$V^{portfolio}(t, T_1, \dots, T_n) = \frac{1}{n} \sum_{i=1}^{n} e^{-r(T_i - t)} \left(Y_{T_1, \dots, T_n}(t) - D \right),$$

where

$$Y_{T_1, \dots, T_n}(t) = \frac{\sum_{i=1}^{n} e^{-r(T_i - t)} F(t, T_i)}{\sum_{i=1}^{n} e^{-r(T_i - t)}}.$$

In case the portfolio represents a 1-year-future, the swap rate is the forward price of the 1-year-future, which can be also observed in the market.

Evaluating an option on this 1-year-forward price (i. e. on the swap par rate) poses the problem of computing the expectation in

$$e^{-rT_0} \mathbb{E} \left[(Y(T_0) - K)^+ \right],$$

where the distribution of Y as a sum of lognormals is unknown. We use an approximation as suggested by e. g. Brigo and Mercurio [2001], which assumes Y to be lognormal. Formally, we can approximate Y by a random variable $\hat{Y}$, which is lognormal and coincides with Y in mean and variance. Then,

$$\log \hat{Y} \sim \mathcal{N}(m, s)$$

with s^2 depending on the choice of the volatility functions $\sigma(t, T_i)$.

Using this approximation, it is possible to apply a Black-option formula again to obtain the option value as

$$\begin{aligned} V^{option} &= e^{-rT_0} \mathbb{E} \left[(Y(T_0) - K)^+ \right] \\ &\approx e^{-rT_0} \mathbb{E} \left[\left(\hat{Y}(T_0) - K \right)^+ \right] \\ &= e^{-rT_0} \left(Y(0) \mathcal{N}(d_1) - K \mathcal{N}(d_2) \right) \end{aligned} \tag{2}$$

with

$$
\begin{aligned}
d_1 &= \frac{\log \frac{Y(0)}{K} + \frac{1}{2}s^2}{s} \\
d_2 &= d_1 - s
\end{aligned}
$$

The approximation has been proposed by Lévy [1992] in the context of pricing options on arithmetic averages of currency rates. This density approximation competes mainly with Monte Carlo methods and modifications of the price of the corresponding geometric average option. The advantage of the approximation to Monte Carlo simulation is clearly the difference in speed in which an option evaluation can be carried out, which becomes even more dramatic when turning the focus to calibration. The main drawback of the manipulation of arithmetic average options is that it yields pricing formulae, which do not satisfy the put-call parity in general.

An empirical discussion of the goodness of the approximation in the context of currency exchange rates is also provided by Lévy [1992]. A comparison of second moments leads to errors that are usually much smaller than 1%, especially when the underlying's volatility is below 50%. While skewness is present in the true but not in the approximated distribution, kurtosis is matched very well. Another study emphasizing the applicability of this approximation in interest rates markets can be found in Brigo and Liinev [2005]. A simulation study done by us in the case of electricity futures shows, that the difference between the true distribution of the sum and the approximating distribution is very small and becomes negligible considering other uncertainties in the application such as quality of market quotes.

4 The Special Case of a Two-Factor-Model

As motivated in the introductory part, a special choice of the volatility function is needed to resemble market observations of the term structure of volatility in the futures contracts. In this section we will use a two-factor model given by the SDE

$$
\frac{dF(t,T)}{F(t,T)} = e^{-\kappa(T-t)}\sigma_1\, dW_t^1 + \sigma_2\, dW_t^2, \tag{3}
$$

for a fixed T. The Brownian motions are assumed to be uncorrelated.

This is a special case of the general setup in the previous section with

$$
\sigma(t,T) = \left(e^{-\kappa(T-t)}\sigma_1,\ \sigma_2\right)
$$

and $W(t)$ a 2-dimensional Brownian motion.

For ease of notation assume in the following that today's time $t = 0$.

This choice of volatility is motivated by the shape of the term structure of volatility (cp. Figure 2). The strong decrease will be modelled by the first factor with an exponentially decaying volatility function. Thus, futures maturing later will have a lower volatility than futures maturing soon. Finally, as $T - t$ becomes very large, the volatility assigned to the contract by this factor will be close to zero. As this is not the case in practice, we introduce a second factor, which will keep the volatility away from zero.

Another way of viewing the two factors comes from the economic interpretation: The first factor captures the increased trading activity as knowledge about weather, unexpected outages etc. becomes available. The second factor models a long-term uncertainty, that is common to all products in the market. This uncertainty can be explained by technological advances, political changes, price developments in other commodity markets and many more.

Within this two-factor model, the variance of the logarithm of the future contract at some future time T_0 can be computed easily (see the Appendix):

$$\mathrm{Var}(\log F(T_0, T)) = \frac{\sigma_1^2}{2\kappa}(e^{-2\kappa(T-T_0)} - e^{-2\kappa T}) + \sigma_2^2 T_0 \tag{4}$$

This quantity has to be used to price an option on month-futures with maturity T_0 with the option formula (1).

In the case of options on quarter- or year-futures, it is necessary to compute the quantity s^2 of the lognormal approximation in equation (2). The derivation of s^2 in this two-factor model can be found in Appendix and is given by

$$\begin{aligned} \exp(s^2) &= \frac{\sum_{i,j} e^{-r(T_i+T_j)} F(0,T_i) F(0,T_j) \cdot \exp\left(Cov_{ij}\right)}{\left(\sum e^{-r \cdot T_i} F(0,T_i)\right)^2} \qquad (5) \\ Cov_{ij} &= \mathrm{Cov}(\log F(T_0,T_i), \log F(T_0,T_j)) \\ &= e^{-\kappa(T_i+T_j-2T_0)} \frac{\sigma_1^2}{2\kappa}(1 - e^{-2\kappa T_0}) + \sigma_2^2 T_0 \end{aligned}$$

5 Fitting the model

5.1 Calibration procedure

In order to calibrate the two-factor model to market data, we need to estimate the parameters $\phi = (\sigma_1, \sigma_2, \kappa)$ such that the model fits the market behaviour. Since we have modelled under a risk-neutral measure, we need to find risk-neutral parameters, which can be observed using option-implied parameters.

Given the market price of a futures-option (month-, quarter- or year-futures), we can observe its implied variance $\mathrm{Var}(\log(F(T_0, T))$ for month-future or s^2 for quarter- or year-futures.

Furthermore, we can compute the corresponding model implied quantities, which depend on the choice of the parameter set $\phi = (\sigma_1, \sigma_2, \kappa)$ as described in the previous section. We will estimate the model parameters such that the squared difference of market and model implied quantities is minimal, i. e.

$$\sum_i \left(\mathrm{Var}_{market}(\log Y_{T_{1_i},\ldots T_{n_i}}(T_{0_i})) - \mathrm{Var}^{\phi}_{model}(\log Y_{T_{1_i},\ldots,T_{n_i}}(T_{0_i})) \right)^2 \to \underset{\phi}{\operatorname{argmin}}, \tag{6}$$

where i represents an option with maturity T_{0_i} and delivery covering the months $T_{1_i}, \ldots, T_{n_i}$. Depending on the delivery period, which of course may be longer than one month, the model variance is either the true model implied variance according to equation (4) or the approximated variance according to equation (5). The minimum is taken over all admissible choices of $\phi = (\sigma_1, \sigma_2, \kappa)$, that means $\sigma_{1,2}, \kappa > 0$.

Since our model is not capable of capturing volatility smiles, which can be observed in option prices very often, we will use at-the-money options only.

The minimization can be done with standard programming languages and their implemented optimizers. The objective function (6) is given to the optimizer, which has to compute the model implied variances of all options for different parameters. The worst case (the computation of the variance of a year-contract) involves the evaluation of all covariances Cov_{ij} between the underlying month-futures in equation (5), which is a 12 by 12 matrix, thus computationally not too expensive. As there are usually not more than 15 at-the-money options available (e. g. at the EEX), the optimization can be done within a few minutes.

Additionally, it is possible to use the gradient of the objective function for the optimization. The gradient can be computed explicitly, which makes the numerical evaluation of the gradient in the optimizers unnecessary. Usually, there is a smaller number of function calls necessary to reach the optimal point within a given accuracy using the gradient than using a numerical approximation. But, the explicit calculation again involves matrices up to size 12 by 12. We found, that the time saved by less function calls is eaten up by the increased complexity of the problem. Both methods end up with about the same optimization time, though the gradient method finds minima, which usually give slightly smaller optimal values than methods without gradient.

5.2 Calibration to Option Prices

In the following we will apply the two-factor model introduced in Section 4 to the German market, i. e. we will calibrate it to EEX prices. We repeated the procedure

Product	Delivery Start	Strike	Forward	Market Price	Implied Vola
Month	October 05	48	48.90	2.023	43.80%
Month	November 05	49	50.00	3.064	37.66%
Month	December 05	49	49.45	3.244	34.72%
Quarter	October 05	48	49.44	2.086	35.15%
Quarter	January 06	47	48.59	3.637	28.43%
Quarter	April 06	40	40.71	3.421	26.84%
Quarter	July 06	42	41.80	3.758	27.19%
Quarter	October 06	43	43.71	4.566	25.35%
Year	January 06	44	43.68	1.521	20.19%
Year	January 07	43	42.62	3.228	19.14%
Year	January 08	42	42.70	4.286	17.46%

Table 1: ATM calls and implied Black-volatility

Method	Constraints	σ_1	σ_2	κ	Time
Function calls and numerical gradient	yes	0.37	0.15	1.40	<1min
Least Square Algorithm	no	0.37	0.15	1.41	<1min

Table 2: Parameter estimates with different optimizers

on different days with similar estimates, but the results will be discussed exemplary using a typical day. Properties of the term structure of volatility on this day have been discussed in the introductory part (cp. Figure 2). The data set is shown in Table 1.

The column implied volatility in Table 1 shows the implied volatility by the Black76-formula. After filtering out data points, where the option price is only the inner value of the option, these 11 options are left out of 15 observable in the EEX data set. Now, one can observe a strong decreasing term structure with increasing time to maturity (there are two outliers, which do not confirm the hypothesis) and a decreasing volatility level with increasing delivery period.

The optimizers converge in less than a minute and optimizing with and without gradient delivers the same results up to two decimal places. Even not restricting the parameters does not change the estimates (see Table 2).

The calibration leads to parameter estimates $\sigma_1 = 0.37$, $\sigma_2 = 0.15$, and $\kappa = 1.40$. This implies, that options, which are far away from maturity, will have a volatility of about 15%, which can add up to 40%, when time to maturity decreases. A κ value of 1.40 indicates, that disturbances in the futures market halve in $\frac{1}{\kappa} \cdot \log 2 \approx 0.69$ years.

The model-implied volatility term structure is shown in Figure 3 together with the

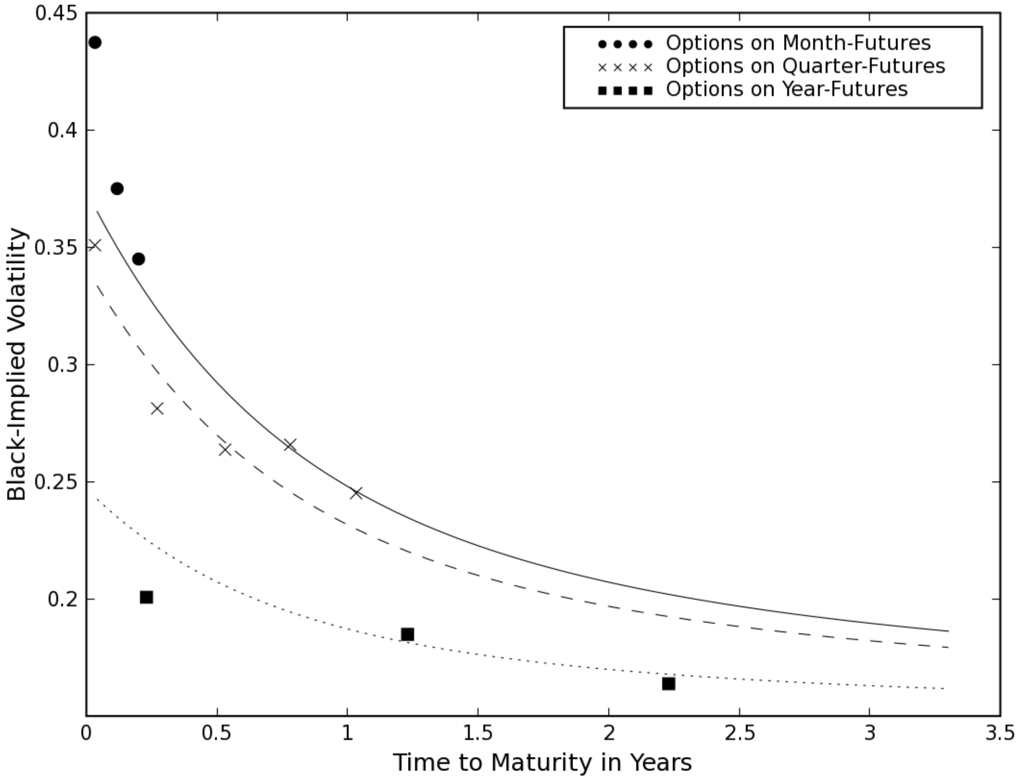

Figure 3: Market option-implied volatilities of futures (symbols) compared to model option-implied volatilities (lines) for different delivery periods and starting dates

Delivery Start	Market Price	Model Price	Market Volatility	Model Volatility
M-October 05	2.023	1.844	43.80%	38.52%
M-November 05	3.064	3.000	37.66%	36.70%
M-December 05	3.244	3.279	34.72%	35.13%
Q-October 05	2.086	2.089	35.15%	35.25%
Q-January 06	3.637	3.865	28.43%	30.82%
Q-April 05	3.421	3.539	26.84%	27.88%
Q-July 06	3.758	3.520	27.19%	25.51%
Q-October 06	4.566	4.315	25.35%	23.83%
Y-January 06	1.521	1.746	20.19%	22.92%
Y-January 07	3.228	3.074	19.14%	18.28%
Y-January 08	4.286	4.131	17.62%	16.93%

Table 3: Comparison between market and model quantities

observed market values. One can see, that, qualitatively, most of the desired properties are described by the model. Futures with a long delivery period show a lower level of volatility compared to those with short delivery. The volatility term structure is decreasing as the time to maturity increases, but it does not go down to zero. Yet, quantitatively, there are some drawbacks. Especially the month-futures show a volatility term structure, that has a much steeper slope than the model implies. Also, the level of volatility is mostly higher than observed in the market, which seems to be the trade-off between fitting all options well at the short end and at the long end. Yet, the model implies reasonable values for all contracts. Absolute values in terms of volatilities and option prices can be taken from Table 3.

5.3 Parameter Stability

In the following we discuss parameter stability. Common procedures are to recalibrate the model frequently on several days and analyze the stability of parameter over time, for different stikes and infer confidence intervals for the parameter. All this requires an active market. This is still not the case for the EEX option market. Almost simultaneously to the introduction of option trading, electricity prices increased sharply. This left the market with deep in-the-money call options, which are not suitable for testing the model introduced above. More recently, the EEX partially overcompensated this effect, leading to many options which are deep out-of-the-money. Only during the last couple of months, it seems as if the market has stabilized and there are quotes for (almost) ATM options and a preliminary analysis is possible. Yet, we want to point out that the longer we go back into recent history, options on month-futures are more and more out-of-the-money.

We estimated parameters daily using the valid forward curves and market prices of options that are closest to at-the-money. We use a history of 52 days, i. e. $2\frac{1}{2}$ months. Figure 4 shows the evolution of parameters over time.

We can clearly identify very stable estimates for the parameter σ_2, which reflects the volatility at the long end of the term-structure. It is mainly determined by options on year-futures, which are available at many strike levels, in particular we observe every day at-the-money option prices. The stability of this estimate is in line with the intuition that long-delivery futures are less sensitive to market changes which should result in low volatility and a low variability of the volatility estimates.

σ_1 and κ describe the volatility level at the short end of the term-structure and the speed of decrease of the term-structure, respectively. The estimate of σ_1 fluctuates more than that of σ_2 but changes are not substantial. The changes can be explained by the sensitivity of option prices to the changes in the underlying. When underlyings move, the option prices move, especially for options with one-month delivery. Since these options are rarely exactly at-the-money, the change in one underlying can change the volatility term-structure drastically, in particular when considering that the volatility smile can be quite steep in electricity markets. This forces the estimate σ_1 to react to the changed condition. In this light, the minor changes over time of σ_1 are acceptable.

In fact, most of the imperfectness of the at-the-money-assumptions are captured by estimates of κ. It is not only time-varying substantially, it seems to have a tendency to decrease during the time under consideration. This can be explained by observing that especially the futures prices at the short end of the curve have moved constantly to the smallest strike level of the corresponding options from below. In other words, options, that have been out-of-the-money at the beginning of the period, are less so at the end. These options have an increased volatility compared to ATM options, which leads to a bigger gap between the short end of the term-structure (mostly options on one-month futures) and the long end of the term-structure (options on year-futures). A large value of κ is required to induce a steep term-structure within in the model. Since the steepness in the market is decreasing over time, κ is decreasing, too.

Considering the correlation of parameter estimates, we find support for the argumentation above (cp. Table 4). The correlation between κ and σ_1 of about 0.64 is substantial, which coincides with the argument that both are driven by the moneyness of the short-end volatilities.

Computing standard deviations from the estimates yields Table 4, which allows to compute asymptotic confidence intervals for each parameter. While standard deviations are rather low for σ_1 and σ_2 as one might expect from Figure 4, it is much higher for κ. Yet, constructing standard confidence intervals for each parameter leads to a decisive rejection of the hypothesis of zero values for any of the parameters at a 99% level of

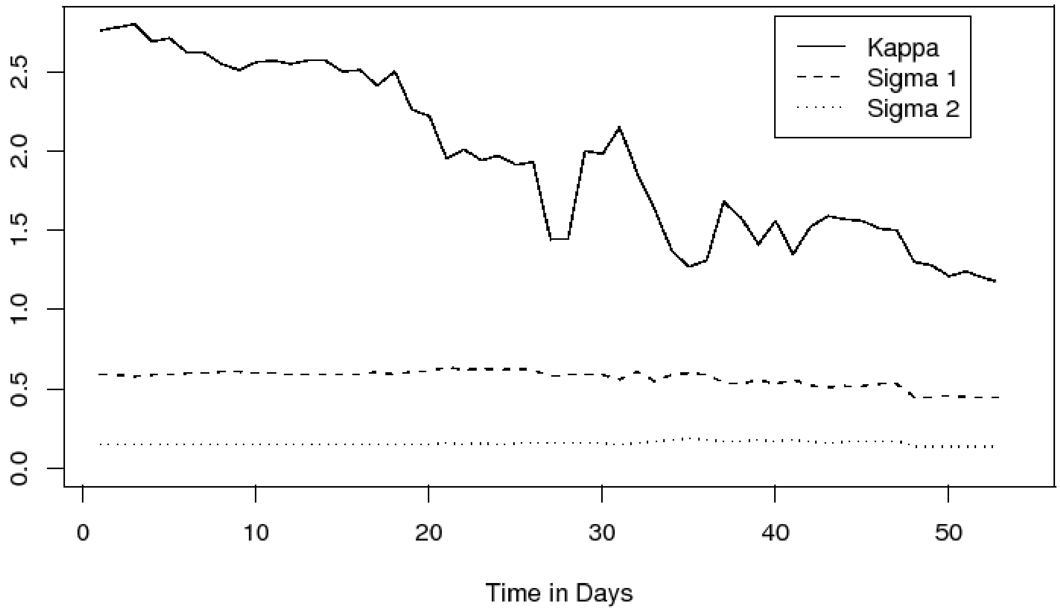

Figure 4: Change in parameter estimates over a 10 week period

confidence.

6 Conclusion

We have presented a two-factor-model for the electricity futures market. It is embedded in a bigger class of market models, which are similar to the very popular market models in the interest rate markets. We have developed pricing formulae for relevant products in the market and shown a procedure to fit the market data.

Correlation	σ_1	σ_2	κ	Mean	Std
σ_1	1.00	0.16	0.63	0.57	0.052
σ_2		1.00	-0.41	0.16	0.012
κ			1.00	1.96	0.536

Table 4: Correlation, mean and standard deviation of parameter estimates

The main results of the work are, firstly, a very good overall fit of the model to option implied volatilities, especially for options on futures with long delivery period and not too close to maturity. Secondly, an excellent incorporation of the length of delivery period into the option pricing. Thirdly, we showed that the model and the calibration procedure work fast and reliable and are ready to use for day-to-day application.

Nevertheless, we have to mention an unsatisfactory fit at the very short end of the volatility term-structure for one-month futures. This might be due to the assumption of lognormal returns. While this assumption is acceptable for some futures contracts in the electricity market for practical applications, it breaks down as the delivery period is short (i. e. one month) and the contract is close to maturity (i. e. last two months of trading), in other words, the more the futures contract approximates the spot price. Possible solutions to this shortcoming might be the inclusion of jumps into the model or another specification of Levy-process driving the background noise. Additionally, this can broaden the data basis for fitting purposes, in that we can price and use other than at-the-money options only.

7 Appendix

All following results will be derived under the assumption of correlated Brownian motions, i. e. $dW_t^{(1)}dW_t^{(2)} = \rho dt$ (in contrast to the independence assumption in the article). The results of the article will be obtained by setting $\rho = 0$.

Derivation of the Variance of a Month-Futures Contract
We will derive equation (4), i. e. $\mathrm{Var}(\log F(T_0, T)$. The SDE describing the futures dynamics in equation (3) can be solved by

$$\begin{aligned} F(t,T) &= F(0,T)\exp\left\{-\frac{1}{2}\int_0^t \tilde{\sigma}^2(s,T)ds + \int_0^t e^{-\kappa(T-s)}\sigma_1 dW_s^{(1)} + \int_0^t \sigma_2 dW_s^{(2)}\right\} \\ \tilde{\sigma}^2(s,t) &= \sigma_1^2 e^{-2\kappa(t-s)} + 2\rho\sigma_1\sigma_2 e^{-\kappa(t-s)} + \sigma_2^2. \end{aligned}$$

Now

$$\mathrm{Var}(\log F(T_0,T)) = \frac{\sigma_1^2}{2\kappa}(e^{-2\kappa(T-T_0)} - e^{-2\kappa T}) + \sigma_2^2 T_0 + 2\frac{\rho\sigma_1\sigma_2}{\kappa}(e^{-\kappa(T-T_0)} - e^{-\kappa T})$$

Derivation of the Variance of Quarter- and Year-Futures Contracts
We will derive equation (5), i. e. s^2 at time T_0. We have

$$\mathbb{E}(Y) = \mathbb{E}(\tilde{Y}), \;\; \mathrm{Var}(Y) = \mathrm{Var}(\tilde{Y}), \;\; \log\tilde{Y} \sim \mathcal{N}(m, s^2).$$

Moments of normal and lognormal distributions are related via

$$\mathbb{E}(Y) = \exp(m + \frac{1}{2}s^2), \;\; \mathrm{Var}(Y) = \exp(2m + 2s^2) - \exp(2m + s^2)$$

Solving this system, we get $\exp(s^2) = \frac{\mathrm{Var}(Y)}{(\mathbb{E}(Y))^2} + 1 = \frac{\mathbb{E}(Y^2)}{\mathbb{E}(Y)^2}$. It can be seen easily that

$$\mathbb{E}(F_{T_0,T_i}) = F_{0,T_i}, \;\; \mathbb{E}(Y_{T_1,\ldots T_n}(T_0)) = \frac{\sum e^{-r(T_i-T_0)}F_{0,T_i}}{\sum e^{-r(T_i-T_0)}}$$

Further

$$\begin{aligned} \mathbb{E}(Y_{T_1,\ldots T_n}(T_0)^2) &= \frac{1}{\left(\sum e^{-r(T_i-T_0)}\right)^2}\cdot \ldots \\ &\quad \sum_{i,j} e^{-r(T_i+T_j-2T_0)}F_{0,T_i}F_{0,T_j}\cdot \exp Cov_{ij} \\ Cov_{ij} &= \mathrm{Cov}(\log F(T_0,T_i), \log F(T_0,T_j)) \end{aligned}$$

The covariance can be computed directly from the explicit solution of the SDE

$$\begin{aligned} \mathrm{Cov}(\log F(T_0,T_i), \log F(T_0,T_j)) &= e^{-\kappa(T_i+T_j-2T_0)}\frac{\sigma_1^2}{2\kappa}(1-e^{-2\kappa T_0}) + \sigma_2^2 T_0 + \ldots \\ &\quad + \frac{\rho\sigma_1\sigma_2}{\kappa}(1-e^{-\kappa T_0})(e^{-\kappa(T_i-T_0)} + e^{-\kappa(T_j-T_0)}) \end{aligned}$$

References

F. E. Benth, J. Kallsen, and T. Meyer-Branids. A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivative pricing. E-print 15, Department of Mathematics, University of Oslo, 2005.

F. E. Benth and S. Koekebakker. Stochastic modeling of financial electricity contracts. Preprint, University of Oslo, March 2005.

F. E. Benth and J. Saltyte-Benth. The normal inverse Gaussian distribution and spot price modelling in energy markets. *International Journal of Theoretical and Applied Finance*, 7(2):177–192, 2004.

D. Brigo and J. Liinev. On the distributional distance between the lognormal LIBOR and swap market models. *Quantitative Finance*, 5(5):433–442, 2005.

D. Brigo and F. Mercurio. *Interest Rate Models – Theory and Practice*. Springer-Verlag, 2001. ISBN 3-540-41772-9.

Á. Cartea and M. G. Figueroa. Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, 12(4), 2005.

L. Clewlow and C. Strickland. Valuing energy options in a one factor model fitted to forward prices. Research Paper Series 10, Quantitative Finance Research Centre, University of Technology, Sydney, 1999.

M. Culot, V. Goffin, S. Lawford, and S. d. Menten. An affine jump diffusion model for electricity. Seminars, Groupement de Recherche en Economie Quantitative d'Aix-Marseille, 2006.

A. Eydeland and K. Wolyniec. *Energy and Power Risk Management - New Developments in Modeling, Pricing and Hedging*. Wiley & Sons, 2003. ISBN 978-0471104001.

H. Geman. *Commodities and commodity derivatives: Modelling and Pricing for Agriculturals, Metals and Energy*. Wiley Finance Series. Wiley & Sons, 2005. ISBN 978-0470012185.

H. Geman and A. Eydeland. Some fundamentals of electricity derivatives. In *Energy Modelling & The Management of Uncertainty*, pages 35–43. Risk Books, 1999. ISBN 978-1899332434.

J. Hinz, L. v. Grafenstein, M. Verschuere, and M. Wilhelm. Pricing electricity risk by interest rate methods. *Quantitative Finance*, 5(1):49–60, 2005.

R. Huisman and R. Mahieu. Regime jumps in electricity prices. *Energy Economics*, 25 (5):425–434, 2003.

B. P. Kellerhals. Pricing electricity forwards under stochastic volatility. Technical Report 209, Universität Tübingen, 2001.

S. Koekebakker and F. Ollmar. Forward curve dynamics in the nordic electricity market. *Managerial Finance*, 31(6):73–94, 2005.

E. Lévy. Pricing European average rate currency options. *Journal of International Money and Finance*, 11(5):474–491, 1992.

J. Lucia and E. Schwartz. Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, 5(1):5–50, 2002.

E. S. Schwartz and J. E. Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911, 2000.

I. Vehviläinen. Basics of electricity derivative pricing in competetive markets. *Applied Mathematical Finance*, 9(1):45–60, 2002.

M. Ventosa, A. Baillo, A. Ramos, and M. Rivier. Electricity market modeling trends. *Energy Policy*, 33:897–913, 2005.

R. Weron. *Modelling and forecasting electricity loads and prices: A statistical approach*. Wiley Finance Series. Wiley & Sons, 2006. ISBN 978-0470057537.