A New Approach for Modelling Electricity Spot Prices Based on Supply and Demand Spreads

Michael Eichler
Maastricht University
Maastricht, The Netherlands
m.eichler@maastrichtuniversity.nl

Johan Sollie
Norwegian University of Science and Technology
Trondheim, Norway
JMSOL@statoil.com

Dennis Tuerk
Maastricht University
Maastricht, The Netherlands
d.tuerk@maastrichtuniversity.nl

Extended abstract

Electricity spot prices are the most important reference measure for all parties of the electricity industry. One important stylized features of electricity spot prices - besides mean reversion and strong seasonality - is the infrequent occurrence of spikes. These are periods of extreme prices that are typically short-lived and during which the spot price exceeds its normal level many times over. As pointed out by, e.g., Knittel and Roberts (2005), spikes occur due to price inelastic market participants when either the demand varies (often as a result of weather conditions) or the supply decreases, e.g., due to outages of generators or transmission lines. The effects can be amplified by generator bidding behavior, to be described later, as the central dispatch process does sometimes require to call generators of high marginal costs into production to satisfy demand. A further phenomena which is to be observed ever more frequently during the last years are negative price occurrences. They are accepted by power suppliers whenever the opportunity costs of a shutdown period would be higher.

In order to capture the various characteristics of electricity prices several approaches of price modeling have been investigated. The most commonly modeling approach are regime-switching models (e.g. Robinson 2000, Huisman and Mahieu 2003, Huisman and De Jong 2003, Rambharat et al. 2005), which distinguish between a base regime for the mean-reverting spot price and a spike regime modeling the extreme prices. The disadvantage of these models is that spikes occur mostly random and the risk of their occurrence is constant over time. In Eichler et al. (2012), it has been shown that forecasting the occurrence of spikes without additional information on the market and possible exogeneous variables is difficult if not impossible.

More promising seems therefore modeling approaches based on fundamentals such as temperatures, loads or reserve margins. For instance, Barlow (2002) uses assumptions about the functional form of the supply and demand curve in order to derive a non-linear Ornstein-Uhlenbeck process as a model for electricity spot prices.

In this paper, we propose a new approach that exploits information available in the supply and demand curves. The underlying idea is that the form of the supply and demand curves or, more precisely, the spread between supply and demand reflects the risk of extreme price fluctuations. In particular, if the spread is constant close to the current price, small fluctuations in supply or demand can lead to large changes in the price. We investigate the approach for the German day ahead market and derive distributions for one-day forecasts of the electricity spot price. For this market the supply and demand curves are available on an hourly basis (see Figure).

As just a part of the system wide load is traded on the German electricity market we admittedly cannot make use of the relation between price and volume (see Boogert and Dupont 2007). Nonetheless, despite the fact that the reported supply and demand curves do not correspond to the curves of the complete market, the quantities do give the clearing price at the exchange. The resulting prices reflect the prices for the whole market because of the no-arbitrage condition. Further a shift in demand and/or supply will affect the price level. By how much the price changes depends on the shape of the supply and demand curves.

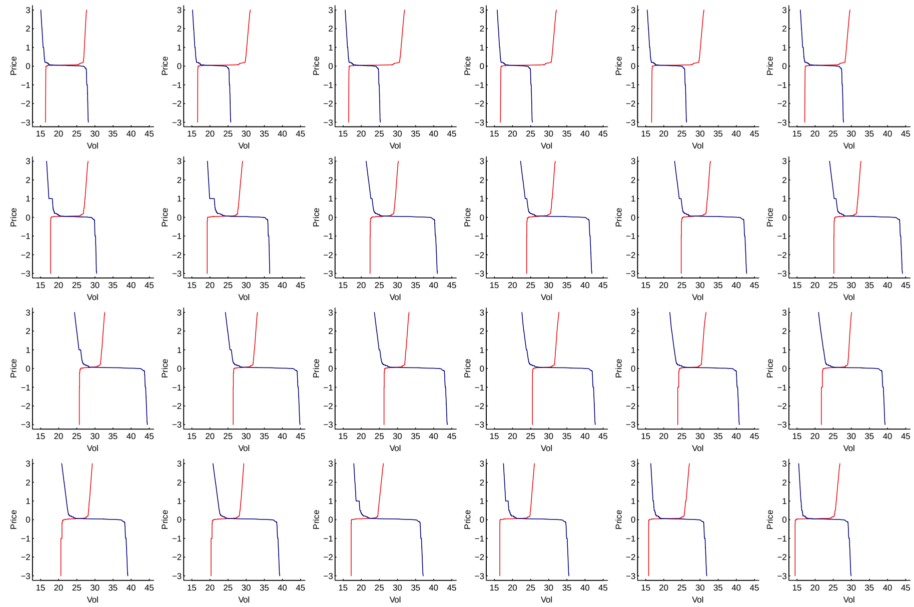

Figure 1: Supply and demand curves for the 24 hours at 1 September 2011 for the Germany/Austrian market.

If the horizontal distance between supply and demand curves is small near the clearing price and if supply and demand curves are further steep, a small shift in one of the curves will induce a large price shift yielding a negative or positive price spike. For this reason, we use the distance between supply and demand curves as a measure of price uncertainty:

$$C_{p,t,h} = \frac{S_{p,t,h} - D_{p,t,h} - min(S_{t,h} - D_{t,h})}{max(S_{t,h} - D_{t,h})}, \qquad (1)$$

with $S_{p,t,h}$ giving a point on the supply curve that corresponds to price p on day t and hour h. the same indices apply to the demand curve yielding $D_{p,t,h}$. Thus $C_{p,t,h}$ is a function of the difference between $S_{p,t,h}$ and $D_{p,t,h}$, which is adapted to start at zero through subtracting the minimal difference between $S_{t,h}$ and $D_{t,h}$ and standardized by the corresponding maximal difference between $S_{t,h}$ and $D_{t,h}$. $C_{t,h}$ thus gives a function which for each price takes a value on the domain $[0, 1]$. Since the shape of supply and demand curves changes due to changing prices of fuels, changes in the production mix in the system, demand, etc., market price uncertainty can be measured through C. Using $C_{t,h}$ to map price $p_{t,h}$ into the domain $[0, 1]$ yields a transformed and much more stable price process, $x_{t,h}$. The resulting surface of the rescaled differences for a given hour over the period from 1st of January 2005 onwards for 600 observations is given in Figure 2. It can be seen, that the steepness of the curves gives an extra information concerning the price uncertainty. When the steepness of a curve is high, the corresponding price uncertainty is low and vice versa. It can further be seen that the mapped time-series, given by the green line, resembles a more stable process than the original price process, which is given by the blue line.

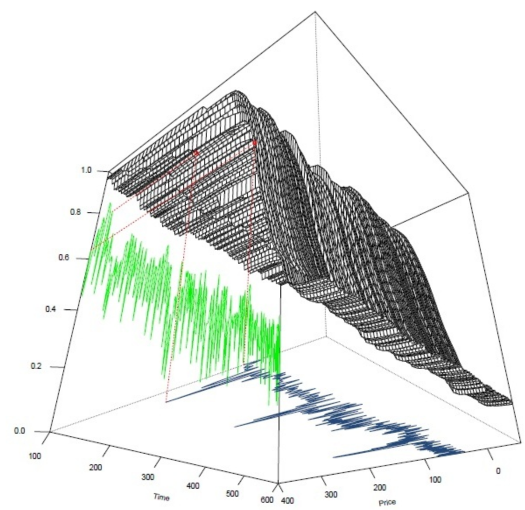

Figure 2: surface of the rescaled differences for a given hour over the period from 1st of January 2005 to 31st of October 2011.

Next we have to fit a model to the transformed spot price in order to be able to forecast the next days value, and simulate from the model. The adjustment of the curves means that the transformed spot price time series only will take values on the interval [0,1]. Modeling this transformed series is not of much use unless the values can be mapped back into the domain of spot prices. Thus, apart from the transformed time-series, we need to simultaneously model next day's distance between supply and demand curves, which we will denote as sd-curve from now on. All the details follow, but first assume that tomorrow's curve easily can be forecasted, say using the observed curve today. The transformed spot price time series, also shown in Figure 2, has a time varying mean, an autoregressive component, a seasonal term, and only takes values on the interval [0,1]. If we have a model for this time series, and can simulate from this model, we can generate a distribution for the spot price of electricity. If the mean of the simulated values corresponds to an area where the distance between the supply and demand curves is steep, it will result in a narrow distribution for the spot price. However, if the mean of the simulated sample is in the higher end, that is close to one, the steepness of the curve is less. This results in a skewed distribution for the spot price, allowing for violent positive price spikes, when the median value of the generated spot price distribution is high. Thus positive price spikes can occur at moderate spot prices, but they will be less probable than when prices are high. A similar argument can be made for low prices.

In order to first model $x_{t,h}$ we need to keep in mind that the series has seasonal patterns, autoregressive components and time varying means. We model the time varying mean as a random walk of second order, and fit the following model to each the 24 transformed spot price time series:

$$x_{t,h} = \mu_t + \gamma_{1,t} + \alpha x_{t-1,h} + \epsilon_t \tag{2}$$

$$\mu_{t+1} = \mu_t + \beta_t \tag{3}$$

$$\beta_{t+1} = \beta_t + \eta_t \tag{4}$$

$$\gamma_{1,t+1} = -\sum_{j=1}^{6} \gamma_{j,t}, \tag{5}$$

where ϵ_t and η_t are independent normally distributed with zero means and variances σ_ϵ^2 and σ_η^2, respectively. The seasonal component is given by $\gamma_{1,t+1}$ and accounts for the weekday-effect. The model is fitted to the series using the Kalman filter. Using a random walk of second order introduces two characteristics of the model; the mean value is time varying, and a forecast obtained from the model will point in the direction that the series have been moving over the last observations. Forecasting over longer time periods with such a model is not suitable, and we will restrict our analysis to one day-ahead forecasts. We find that including one lag is sufficient to model the serial correlation, and the seasonal term removes any seasonality in the standardized one step ahead forecast errors.

In order to simulate from the model we will sample from past residuals while truncating the simulated values to be in the interval $[0, 1]$. Despite the fact, that we model each hour of the day separately, we know, that the prices of a certain day are strongly correlated. Therefore we will always sample from residuals of the same day when simulating different prices over a day.

In order to get a model to forecast tomorrow's supply and demand curves (or specifically, the distance between tomorrow's supply and demand curves) we carefully examine their structure over time. The sd-curves appear to exhibit time varying mean and intraweekly seasonality. Thus we first obtain the expected average value, for each weekday, over our full sample. Next we deseasonalize the curves by subtracting the obtained average values. When sampling from the curves we then use the corresponding mean value, and sample from the last 7 deviations from the expected curve. This can be written as:

$$C_{t+1,h} = \bar{C}_{t+1,h} + (C_{t+1-j,h} - \bar{C}_{t+1-j,h}), \tag{6}$$

with point $\bar{C}_{t+1,h}$ equaling the sum of the corresponding mean over all past points which where observed on the same day and hour as $C_{t+1,h}$. Further we add the difference of a random past value $C_{t+1-j,h}$ and the corresponding mean $\bar{C}_{t+1-j,h}$ with the two values chosen through the random number j which we decided to be between 1 and 7. Using this simple approach, we are able to include the time varying mean and seasonality, and at the same time incorporate the uncertainty in the curve forecast. In order to again take into account that different curves of the same day might be connected we do use the same index j for all 24 curves of a given day.

After having presented the model we now turn to an empirical application in order to show the good quality of the approach. In this context on should keep in mind that a good model for the distribution of the next days spot price of electricity should be robust to changes in the market, and should be able to generate distributions that reflect the true uncertainty of the price process.

To illustrate the forecast distributions, we show in Figure 3 the one-day ahead forecast of the spot price for the period from 8 October 2009 to 26 April 2010. The distribution of the forecast is indicated by the median (solid red) and the 5% and 95% quantiles (dashed red). The four series represent different hours of the day.

To assess the quality of the forecasting distribution, Figure 4 shows the percentiles generated by the model and the theoretical percentiles. The percentiles are calculated in the following way: A transformed spot price series is calculated for every hour, and the model (2) is fitted to the series. We generate one step ahead forecast using the Kalman filter, and draw from historical residuals (taking the correlation structure over

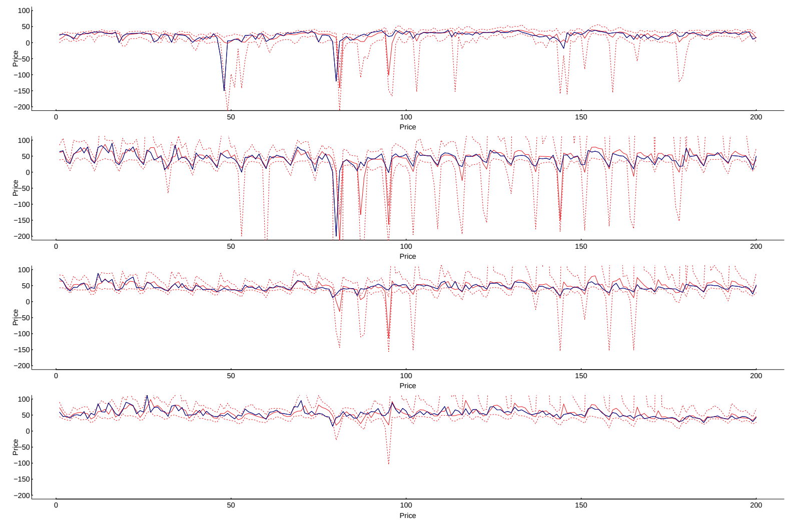

Figure 3: Forecasted (red) and true (black) electricity spot price for the period from 8.10.2009 to 26.4.2010. The solid red line gives the median of the forecasting distribution while the dashed red lines indicate the 5% and 95% quantiles. The four figures from the top show forecasts and true prices for the hours 3-4, 7-8, 13-14, and 18-19.

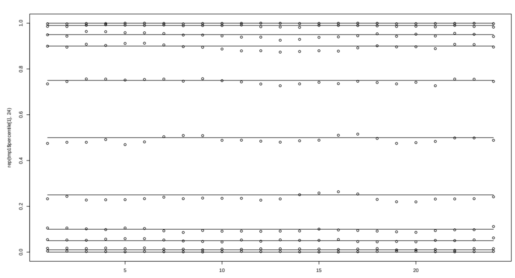

Figure 4: Theoretical percentiles and percentiles obtained from the model.

a day into account as described earlier) to obtain the distribution of the transformed spot price. To translate the transformed spot price into the real spot price, we sample sd-curves as described in (6), again taking the correlation of these curves into account. The steps are repeated in order to obtain 10 000 realizations of each of the 24 hourly spot prices. The day-ahead price is calculated as the mean of the 24 hourly prices for each realization, resulting in 10 000 samples of the day-ahead price. The process is repeated for each day in the sample, and the realized spot price is compared to the generated distributions to se how the model fits the data. The percentiles used in the analysis are $\{0, 1, 5, 10, 25, 50, 75, 90, 95, 99, 1\}$.

We are thus able to show that using information from supply and demand curves will allow us to construct a model that is flexible enough to adapt to changing market conditions. Having a model that is able to do so is of great importance to market participants, and represents a significant contribution to the existing literature on electricity price models.

References

Barlow, M. T. (2002). A diffusion model for electricity prices. *Mathematical Finance* **12**, 287–298.

Boogert, A. and Dupont, D. (2007). When supply meets demand: The case of hourly spot electricity prices. *Birbeck Working Papers in Economics and Finance BWPWF-0707*, Birbeck, University of London.

Eichler, M., Grothe, O., Manner, H. and Tuerk, D. (2012). Modeling spike occurrences in electricity spot prices for forecasting. *Working paper 12/029*, METEOR Research Memorandum. Maastricht University.

Huisman, R. and De Jong, C. (2003). Option pricing for power prices with spikes. *Energy Power Risk Management* **7**, 12–16.

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

Knittel, C. R. and Roberts, M. R. (2005). An empirical examination of restructured electricity prices. *Energy Economics* **27**, 792–817.

Rambharat, B. R., Brockwell, A. E. and Seppi, D. J. (2005). A threshold autoregressive model for wholesale electricity prices. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **54**, 287–299.

Robinson, T. A. (2000). Electricity pool prices: a case study in nonlinear time-series modelling. *Applied Economics* **32**, 527–532.