A Power Market Forward Curve with Hydrology Dependence
An Approach based on Artificial Neural Networks

RIKARD GREEN

KNUT WICKSELL WORKING PAPER 2015:1

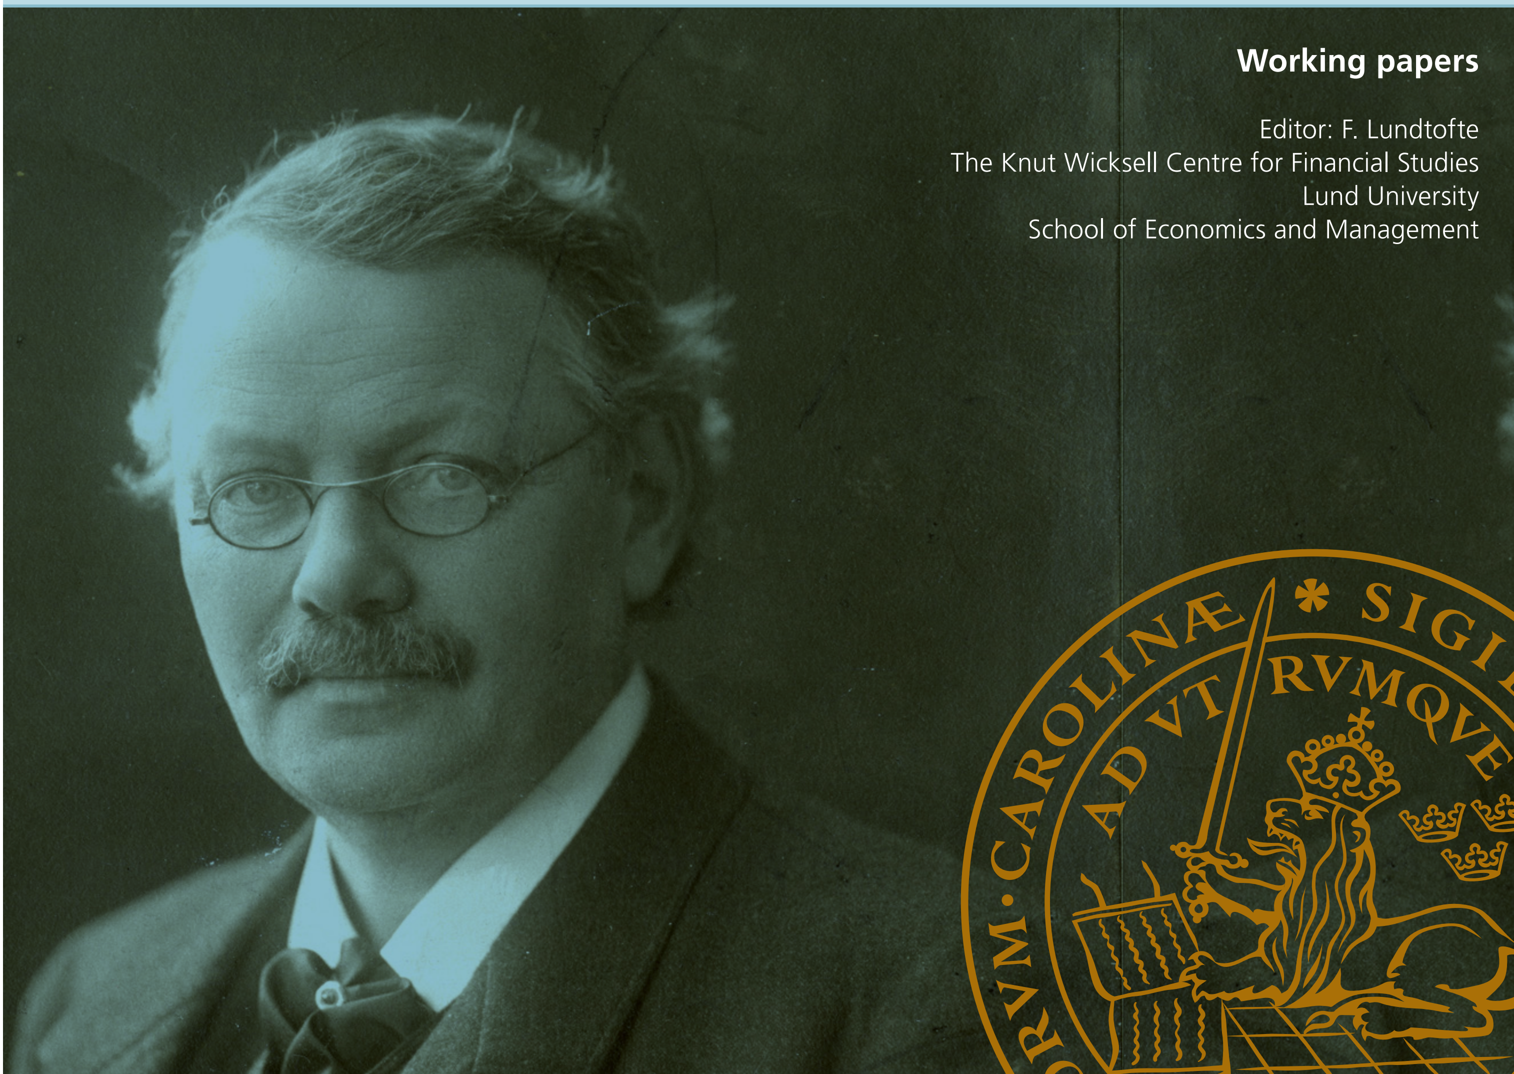

Working papers

Editor: F. Lundtofte
The Knut Wicksell Centre for Financial Studies
Lund University
School of Economics and Management

A Power Market Forward Curve with Hydrology Dependence

An Approach based on Artificial Neural Networks

Rikard Green*

April 11, 2014

Abstract

This paper develops an hourly forward curve for power markets where the intra-day and intra-week shapes (profiles) depend on the level of the hydrological balance. The shaping model is based on a feed-forward Artificial Neural Network (ANN), which is trained on a historical data set of hourly electricity spot prices from the Nord Pool market and weekly measurements of the Nordic hydrological balance. The yearly seasonal cycle is estimated with historical electricity forward prices from the Nasdaq OMX Commodities exchange. We calibrate the shaping model to prevailing electricity forward prices and proceed to demonstrate its most important properties. By using comparative static analysis we particulary focus on the hydro dependence of the shapes. We conclude the paper with a real world valuation task. By combining our proposed forward curve with a simple Ornstein-Uhlenbeck process we price a strip of hourly call options on the electricity spot price under different hydrological scenarios.

JEL classification: C45, C51, G13, Q41,
Keywords: Power markets, forward curve, seasonality, artificial neural networks

*Knut Wicksell Centre for Financial Studies and Department of Economics, School of Economics and Management, Lund University, P.O. Box 7082, S-220 07 Lund, Sweden. E-mail: rikard.green@nek.lu.se

Financial support from the Crafoord Foundation is gratefully acknowledged.

1 Introduction

Electricity is a *flow commodity* which means that it has to be delivered over a period of time. Consequently, exchange traded forward contracts are based on averages of the underlying electricity spot price for different delivery periods (weekly, monthly, quarterly and yearly). The traded term structure is piecewise constant over the quoted delivery periods and it reflects the market participants' believes on the future risk-adjusted average spot price. In this sense the traded term structure is somewhat indistinct and market participants commonly depend on models to get a more informative representation of the term structure. The relevant class of models is given by continuous term structure forward curves and they are designed so that the term structure reflects seasonal shaping patterns while simultaneously being consistent with traded forward prices. In electricity markets such curves are represented at hourly granularity and they are hence known as *hourly forward curves*. The hourly forward curve is a crucial and fundamental tool in any valuation task and it enables the electricity trader to value any forward position regardless if the delivery period is standard or non-standard. It also serves as an essential input to stochastic price models with applications to risk-management and derivative valuation, where it represents the initial value to express today's forward prices. We argue that the hourly forward curve is an instrument of fundamental importance in the modern energy trading business.

The literature has made several contributions to the field of continuous term structure forward curve models. The general method for constructing a continuous forward curve can be decomposed into two steps, where the first step performs estimation of the seasonal shapes (intra-day, intra-week and intra year) and the second step calibrates the estimated seasonal vector to prevailing forward prices. In this paper we use the terms *seasonal shapes* and *seasonal patterns* interchangeably. Most studies have focused on the calibration scheme (second step) while estimation of the seasonal shapes has been left aside. An early study by Fleten and Lemming (2003) proposed an optimization based calibration method to construct a smooth daily forward curve for electricity markets. The discussion on estimation of seasonal shapes was secondary, and the authors used an exogenous fundamental bottom-up model to account for these effects. Another relevant study was carried out by Benth, Koekebakker and Ollmar (2007) where the authors suggested a method to compute a smooth curve from observed forward prices with delivery periods. The curve was decomposed into a seasonal component and a correction term and the model was calibrated to two simplistic seasonal specifications and a fundamental bottom-up model. Again, the study was primarily directed towards the calibration procedure. A simple and practically viable calibration procedure was suggested in Burger, Graeber and Schindlmayr (2008) who employ a linear scaling method to calibrate the estimated seasonal shapes to traded forward prices. The authors note in passing that estimation of the seasonal shapes should be carried out from historical spot prices using forecasting methods similar to those for load forecasting. We argue that neither of the mentioned studies have presented a serviceable approach to model the seasonal shapes.

However, the important connection between seasonality and fundamental factors was pointed out in both Fleten and Lemming (2003) and Benth et al. (2007).

In this paper we suggest a model to estimate the seasonal shapes that provides the basis for an hourly forward curve. Our approach is based on Artificial Neural Networks (ANN) and it is related to the methodology suggested in Crispin and Jacobsson (2007) where the authors compute a smooth daily forward curve for electricity markets using a seasonal model based on ANNs. The main contribution of this paper is that we expand the approach in Crispin and Jacobsson (2007) to include a connection between the seasonal specification and the hydrological balance, which is a fundamental factor related to the supply side of the power system. This extension is in agreement with the observations in Fleten and Lemming (2003) and Benth et al. (2007) that point out the connection between seasonality and fundamental factors. We define the hydrological balance as a measure of the amount of available and potential production resources used for hydro power production. In a power system with a high proportion of hydro generation, such as the Nordic market, the hydrological balance is connected to the production behavior of hydro power producers, which in turn impacts the intra-day and intra-week seasonal shapes. Our model accounts for this connection and by analyzing the estimated model we verify that the intra-day and intra-week seasonal shapes indeed have an impact from changes in the hydrological balance. The proposed ANN-model is trained on a data set of hourly electricity spot prices from the Nord Pool market. We use weekly data for the Nordic hydrological balance for the same time period. In order to estimate the intra-year seasonal shape we use Nordic forward data from the Nasdaq OMX Commodities Exchange. We analyze the seasonality vector given by the model and we pay special attention to the sensitivity between the seasonal shapes and the hydrological balance. In order to calibrate the estimated shaping vector to current forward prices we use the linear scaling method suggested in Burger et al. (2008).

We conclude the paper with a pricing application. By combing our proposed hourly forward curve with a simple stochastic Ornstein-Uhlenbeck model we calculate the price of a strip of hourly call options on the electricity spot price under different hydrological scenarios. The contract specification is a real world example of an OTC product and we therefore employ a pricing approach similar to what is used by practitioners. In a comparative static analysis we conclude that changes in the hydrological balance has a clear effect on the intra-day and intra-week seasonal shapes, which in turn impacts the price of the OTC product.

This paper is organized as follows. Section 2 overviews the data and presents a series of stylized facts. Section 3 describes the shaping model, the calibration method and the stochastic price model used in the application. In Section 4 we analyze the estimated hourly forward curve. Section 5 is devoted to pricing of an OTC product. Finally, we sum up and state our conclusions in Section 6.

2 The Data

In the current section we present the historical data and perform a statistical analysis to establish a series of *stylized facts* that we require our seasonal shaping model to account for.

2.1 Overview of Data

Electricity spot prices exhibit a complex combination of simultaneous seasonal patterns at different frequencies. In order to robustly estimate these patterns we use a rich historical data set of hourly system spot prices from the Nord Pool market between January 7, 2002 and December 25, 2011. Spot data is a natural source to estimate intra-daily and intra-weekly seasonal shapes but to estimate the yearly seasonal cycle we argue that it is better to use data from the forward market. Spot prices strongly depend on the current state of fundamental factors which might cloud the intra-yearly seasonal cycle. This stands in contrast to the forward market which always displays a clear (expected) yearly seasonal shape in its term structure. We estimate this shape using historical forward price quotes (monthly, quarterly and yearly) from the Nasdaq OMX Commodities Exchange between 2008-2009. Finally, since our model connects the seasonal shapes to the hydrological state we use a historical time series of weekly measurements of the Nordic hydrological balance stretching between 2002-2011.[2] The hydrological balance is defined as the sum of the Nordic snow reservoir and hydro reservoir.

2.2 Stylized Facts

Day types

Hourly spot prices show distinct intra-day seasonal patterns with variations for different day types (Monday-Sunday). In Figure 1 we have used our spot data set to calculate the average normalized intra-day hourly weights for each day type and it is clear that Mondays-Thursdays display quite similar behavior, while Fridays are somewhat different with a less pronounced evening peak.[3] Saturdays and Sundays both exhibit a different pattern with delayed morning and evening peaks. We finally note that Sundays show a more protracted evening hump compared to Saturdays.

INSERT FIGURE 1 AROUND HERE

In addition to the intra-daily effects there is a clear intra-weekly pattern with Fridays and weekends displaying lower levels than weekdays. These effects are displayed in Figure 2

[2]The time series measures the deviation of the hydrological balance from its normal state in units of TWh. This measure is commonly used by market participants.

[3]The hourly weights have been normalized to sum to unity for any given day.

which shows the averaged normalized daily weights for each day type.[4] Clearly, Fridays and weekends show lower weights than weekdays. The intra-daily and intra-weekly effects must crucially be accounted for in a shaping model.

INSERT FIGURE 2 AROUND HERE

Season

It is a well established fact that Nordic electricity prices display annual seasonality distinguished by a slowly varying trend due to seasonal changes in temperatures. This is clearly visible in Figure 3, which shows prices for quarterly contracts (settlement period 2013) traded at Nasdaq OMX Commodities on November 25, 2011.

INSERT FIGURE 3 AROUND HERE

In fact, these seasonal changes also affect the appearance of the intra-day and intra-week shapes over the course of the calendar year. This effect is demonstrated in Figure 4 where we display an example of two normalized intra-day shapes for average Wednesdays in February and June respectively. We note that the shape in February carries distinct morning and evening peaks while the corresponding shape in June has a deep trough for the early morning hours while lacking the evening peak. A general notion is that warmer seasons typically show larger differences between peak and off-peak than colder seasons, which is clear from Figure 4.

INSERT FIGURE 4 AROUND HERE

Hence, intra-day shapes for the same day type might be completely different depending on the season. The same line of argumentation holds for the weekly shapes where we typically see larger spreads between weekends and weekdays in warmer seasons. We require our model to reflect (i) the slow annual seasonality due to changes in temperatures and (ii) the seasonal changes in intra-day and intra-week shapes.

Holidays and bridge days

Public holidays, bridge days and school breaks all have major impact on electricity prices. The limited amount of historical data makes it challenging to estimate the size of the impact and simplistic dummy methods are commonly applied, which is also the case in this paper. To account for this information in an hourly forward curve one must keep track of calendar data for all regions/countries affecting the price. The Nordic power market is multinational market and calendar information from Sweden, Norway, Denmark and Finland need to be incorporated into the model.

[4]The daily weights have been normalized to sum to unity for any given week.

Hydrology

The hydrological balance is an important price driver in the Nordic power market. More than 50 percent of the Nordic generation capacity consists of hydro power which makes price levels (and shapes) highly sensitive to changes in the hydrological balance. In situations with considerable hydrological oversupply hydro producers might run into difficulties with managing the water for off-peak hours, resulting in lower prices. This especially holds true during warmer seasons when the power consumption is low to start with. For peak hours, the consumption is higher and the price formation occurs further up the bid-stack making it easier for hydro producers to control the water, and prices will bounce upwards to adequate levels. Hence, ceteris paribus, a situation with hydrological oversupply likely gives a wider price spread between peak and off-peak hours. The same line of argumentation holds for the opposite case, i.e. in situations with hydrological undersupply we generally observe tighter spreads between peak and off-peak prices. We confirm this argument in Figure 5, which displays average normalized intra-day hourly weights for Wednesdays in April, under two different hydrological situations; wet and dry.

INSERT FIGURE 5 AROUND HERE

The calculations are based on data from the years 2006 and 2008, where the average hydrological balance (deviation from normal state) in April was -19.43 TWh and +13.41 TWh, respectively. The wet situation (depicted with circles) displays a pattern with a large spread between off-peak and peak hours, which is in accordance with our argumentation. The dry situation (depicted with triangles) shows a tighter profile. The proposed shaping model will account for the connection between shapes and the hydrological state.

3 The Model

The method of construction of an hourly forward curve for electricity markets is typically subdivided into two steps where the first step concerns shape estimation and the second step calibrates the estimated shape vector to current forward prices. We are primarily concerned with estimation of the seasonal shapes. This section presents the architecture of the Artificial Neural Network (ANN) which is used to estimate the shaping model. A simple empirical approach for monthly shaping is also outlined. We furthermore describe the calibration procedure from Burger et al. (2008) that we employ to calibrate our estimated shapes to traded forward prices. The section is concluded with an overview of a simple stochastic model for the electricity spot price. We utilize this model in conjunction with the proposed hourly forward curve to price a strip of hourly call options on the electricity spot price. The pricing application is done in section 5.

3.1 The Shaping Model

In order to estimate the shaping vector containing all seasonal effects we employ an approach based on Artificial Neural Networks (ANN). For an excellent overview of neural networks see e.g. Ripley (2008). Our approach is a straightforward generalization of the previous work in Crispin and Jacobsson (2007). In fact, our main contribution is that we expand the approach by the previous authors to include dependence between the seasonal specification and the hydrological balance, which according to the analysis in section 2.2 is an important determinant of the seasonal shapes. In addition to the ANN specification for the hourly and daily patterns, we employ a purely empirical method to estimate the annual seasonality (monthly weights). In the end, we combine the results (hours, days and months) to obtain a single vector containing all shaping information.

The hourly and daily seasonal patterns are modeled separately in two unattached networks. The separation into two different models is advantageous since it eliminates the risk of confusing the different patterns.

3.1.1 The Hourly Network

The hourly model is a 6-60-24 feed-forward network. We let $\{(s_t^h, \mathbf{x}_t^h)|t = 1, ..., T\}$ be the available data for training of the network, where $\mathbf{x}_t^h$ denotes the input data and s_t^h is the target data. The first two columns of the input data matrix are given by a sinus-function and a cosine-function with yearly periodicity, which jointly serve as a clock to inform the model on the time of year. Note, that we need two coordinates to represent a unique point in time of the calendar year. By keeping track of the time of year we enable the model to account for potential seasonal changes in the intra-day pattern. In section 2.2 we confirmed the presence of such effects in our data set. Columns three and four of the input matrix serve as a weekly clock, and similar to the previous case they are specified as a sinus-function and a cosine-function however with weekly periodicity. The weekly clock informs the model on the different day types, and enables identification of different shapes for each day type. The analysis in section 2.2 indeed showed that intra-day patterns vary for different day types. The fifth column of the input matrix is a vector of pre-estimated factors to scale down holidays and bridge days. The factors are unit valued for working days and < 1 for holidays or bridge days. Note, for the given historical period the system price is a joint price for Sweden, Norway, Denmark and Finland, and in order to estimate the weights we have utilized calendar information from all countries. The last column of the input data matrix consists of weekly measurements of the hydrological balance from the Nordic market. To obtain daily data points we have interpolated the original data. The original time series measures the deviation of the hydrological balance from its normal state in units of TWh. According to common practice for network training we have standardized the hydrological data to have mean zero and standard deviation one. The target data vector consists of historical hourly spot prices normalized such that the hourly weights for any given day sum to unity.

The network is restricted (guided) in the sense that the neurons governing the yearly/weekly clocks, the holidays and bridge days, and the hydrology, are respectively disconnected from each other. The illustration in Figure 6 shows the essential architecture of the restricted network.

INSERT FIGURE 6 AROUND HERE

Since we require the output of the network to be interpreted as positive weights summing to unity, we employ the *softmax* activation function. The network is trained using a quasi Newton optimization algorithm available in our software package.[5] In order to avoid over-fitting and improve generalization of the network we make use of regularization techniques. We have implemented regularization by Gaussian priors on each network parameter (acting as a decay factor on the weights). To determine the size of the decay factor we retrained several networks using different decay factors and we performed out-of-sample predictions to find a network with good generalization. Specifically, we divided our data set into a design set (90 percent of the observations) and a validation set (10 percent of the observations). We trained three different networks with initial randomized weights, and for each network we forecasted the validation set and calculated the average forecast error. We then re-divided the data set into a new design set and validation set, where the data points in the validation set are non-overlapping with the points in the previous validation set. With randomized initial weights we again trained three networks and calculated the average forecast error. This procedure was iterated 10 times (until all data points in the original data set had been forecasted). According to this method a decay factor of 0.05 proved to be suitable.

3.1.2 The Daily Network

The weekly model is a 10-60-7 feed-forward network. We let $\{(s_t^d, \mathbf{x}_t^d)|t = 1, ..., T\}$ be the available data, where $\mathbf{x}_t^d$ denote the input data and s_t^d is the target data. The weekly model has a similar structure compared to the hourly model. The first two columns of the input data matrix are yearly sinus- and cosine-functions to measure the time of year. Columns three to nine are pre-estimated factors to account for holidays and bridge days. The final column contains weekly measurements of the hydrological balance standardized with mean zero and standard deviation one. The target data vector consists of historical daily spot prices normalized such that the daily weights for any given week sum to unity. The network is trained using the same approach as for the hourly model, with a decay factor of 0.05.

3.1.3 Empirical Monthly Shaping

Spot prices strongly depend the current fundamental situation and for this reason historical price data might cloud the annual seasonal cycle. For estimation of the annual seasonality

[5] We use the Netlab library which is built on top of the Matlab system.

we instead turn to the forward market, which typically displays a clear (expected) yearly seasonal shape in the term structure. We employ a simple empirical method to estimate the seasonal cycle from historical forward price quotes (yearly, quarterly and monthly products). For any given month contained in a certain quarter, and year, let F_t^M, F_t^Q and F_t^M be the respective forward price quotes at time t. We estimate the series of quarter-to-year ratios, $R_t^{Q/Y}$, and month-to-quarter ratios, $R_t^{M/Q}$, simply be calculating

$$R_t^{Q/Y} = \frac{F_t^Q}{F_t^Y} \quad \text{for all } t = 1, ..., T \tag{1}$$

$$R_t^{M/Q} = \frac{F_t^M}{F_t^Q} \quad \text{for all } t = 1, ..., T \tag{2}$$

In order to obtain stable estimates we average the ratios over the entire historical time period (years 2008-2009). The average ratios are computed as

$$\widehat{R}_t^{Q/Y} = \frac{1}{T}\sum_{t=1}^{T} R_t^{Q/Y} \tag{3}$$

$$\widehat{R}_t^{M/Q} = \frac{1}{T}\sum_{t=1}^{T} R_t^{M/Q} \tag{4}$$

We finally calculate the average monthly weights (month-to-year) from

$$\widehat{R}_t^{M/Y} = \widehat{R}_t^{Q/Y} \times \widehat{R}_t^{M/Q} \tag{5}$$

By repeating this operation for all months we end up with a vector of monthly weights, which after normalization, will sum to unity over the calendar year. The vector reflects the annual seasonal cycle as it was perceived by the market during the historical time period. Figure 7 displays the estimated monthly shape vector.

INSERT FIGURE 7 AROUND HERE

The annual seasonal cycle is clearly visible. In particular, we note the drop between June and July, and the corresponding sharp increase between July and August, which is due to the holiday season.

3.2 The Calibration Procedure

There are several available approaches to calibrate an estimated shaping vector to prevailing forward prices. In general, such methods are formulated as constrained optimization problems where the constraints ensure that the curve reflects traded forward prices. The most well known examples are found in Fleten and Lemming (2003) and Benth et al. (2007). In this paper, however, we employ a considerably simpler calibration approach from Burger et al. (2008) which is based on a linear scaling method. This method turns out to be a straightforward and robust alternative sufficient for our purposes.

The method builds on the assumption that the hourly forward prices are given by

$$F(t, T_i) = \sum_{m=1}^{N_b} \alpha_m e_{mi} S_i \tag{6}$$

where $F(t, T_i)$ denotes the forward curve at time t with delivery at T_i and where S_i is the previously estimated shaping vector containing all shaping information. The coefficients α_m are scaling factors chosen such that the forward curve is arbitrage-free against the traded forward contracts F_m^b. The quantity N_b denotes the number of traded forward contracts and e_{mi} is an indicator function that gives the delivery structure of all traded instruments, which is defined as

$$e_{mi} = \begin{cases} 1 & \text{for} \quad i \in J_m^b, \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

where J_m^b is the set of hours corresponding to the delivery period in the traded forward contract F_m^b. By definition we require our forward curve to re-price the traded forward contracts (absence of arbitrage). This property is ensured by setting

$$\frac{1}{|J_k^b|} \sum_{i \in J_k^b} F(t, T_i) = F_k^b \quad \text{for all } k = 1, ..., N_b \tag{8}$$

Substituting the assumed hourly model from equation (6) into the above condition for absence of arbitrage we end up with the following expression.

$$\frac{1}{|J_k^b|} \sum_{i \in J_k^b} \left[\sum_{m=1}^{N_b} \alpha_m e_{mi} S_i \right] = F_k^b \quad \text{for all } k = 1, ..., N_b \tag{9}$$

Rearranging terms allows us to write these equations as

$$\frac{1}{|J_k^b|} \sum_{m=1}^{N_b} \alpha_m \left(\sum_{i \in J_k^b} e_{mi} S_i \right) = F_k^b \quad \text{for all } k = 1, ..., N_b \tag{10}$$

We now have a system of N_b equations for the same number of variables α_m. Absence of arbitrage opportunities in the forward prices implies the existence of a unique solution to this linear system of equations. However, closing prices for Nordic power at Nasdaq OMX Commodities are typically not entirely arbitrage-free, e.g. the price of a yearly contract might be slightly inconsistent with the quarterly contracts for the same delivery year. Since one cannot trade closing prices this is merely a theoretical arbitrage, which does not give rise to real arbitrage opportunities. But from a model point of view we need to handle this by removing redundant forward prices from the system of equations, which makes it possible to still obtain an exact solution. For any given situation with entire overlap of forward products we always remove most granular product furthest out on the term structure. E.g. in a situation where we have a quarterly product being traded simultaneously with the three corresponding monthly products, we remove the last monthly product.

3.3 A Simple Stochastic Model for the Spot Price

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the objective probability space on which we introduce a one-dimensional Wiener processes $W(t)$ on the time interval $[0, T]$, where T is a constant. By letting $\mathcal{F}(t)$, $t \geq 0$ denote the filtration on our space, we interpret it as the information available to agents at time t. Denoting the de-seasonalized electricity spot price by $S(t)$, we assume that $\ln S(t) = X(t)$ where

$$dX(t) = -\alpha X(t)dt + \sigma W(t) \tag{11}$$

$$X(0) = x_0 \tag{12}$$

where σ is the volatility, $\alpha > 0$ is the speed of mean-reversion and where the level of mean-reversion is equal to zero. This type of model is commonly known as an Ornstein-Uhlenbeck process and was first introduced in finance in the context of interest-rate markets, see Vasicek (1977), but has later been frequently employed in the commodity literature, see e.g. Schwartz (1997). We establish the risk-neutral dynamics by standard Girsanov theory, see e.g. Shreve (2004). The risk-neutral specification will have level of mean reversion $\mu(t) = -\lambda(t)\sigma/\alpha$, where the function $\lambda(t)$ denotes the time-dependant market price of risk, which is used to calibrate the spot model to the hourly forward curve in the subsequent application (section 5).

4 Analysis of the Trained Model

We continue to employ the trained networks to generate hourly and daily shapes for the out-of-sample period January 1, 2013 to December 31, 2013. By combining the results from the hourly and daily models with the estimated monthly seasonal vector, we proceed to construct a single hourly vector containing all shaping information. Let $\widehat{S}^H_{hdm}$ be the estimated hourly weight for a given hour h, in a given day d, for a given month, m. Similar notation goes for the estimated daily and monthly weights respectively; $\widehat{S}^D_{dm}$ and $\widehat{S}^M_m$. We let $\hat{\mathbf{S}}$ denote the combined shaping vector for a single calendar year and compute

$$\hat{\mathbf{S}} = \widehat{S}^H_{hdm} \times \widehat{S}^D_{dm} \times \widehat{S}^M_m \tag{13}$$

It is now straightforward to calibrate the estimated seasonal vector $\hat{\mathbf{S}}$ to prevailing forward prices using the calibration procedure from section 3.2. The model is calibrated to forward closing prices from the Nordic power market traded at Nasdaq OMX Commodities per November 20, 2012. Using monthly, quarterly and yearly base products we calibrate the curve for the out-of-sample period January 1, 2013 to December 31, 2013. We deal with arbitrage inconsistencies between closing prices by utilizing the product selection methodology described in section 3.2. The products and input prices used for calibration are given in Table 1. The objective of the forthcoming analysis is to establish that the modeled forward curve reflects the stylized facts outlined in section 2.2.

INSERT TABLE 1 AROUND HERE

Figure 8 displays the calibrated forward curve for the full calendar year 2013. The model is calibrated under the assumption of a normal hydrological balance. The dotted line displays hourly prices and the solid line shows the corresponding monthly averages. Obviously, it is impossible to distinguish the hourly profiles from the figure, but we note that the slow monthly seasonal cycle is present in the model.

INSERT FIGURE 8 AROUND HERE

We emphasize that the monthly shaping only takes place in the absence of traded products. E.g. in the case where we input monthly products to the calibration scheme they will overwrite the monthly shaping from the model. Only in the case where we lack monthly input products will the model perform the monthly shaping. We note from Figure 8 that the out-of-sample monthly shapes for July-December (where no monthly products are available in the model) show a nice seasonal pattern in agreement with the historically estimated monthly shape vector (showed in Figure 7).

We continue to take a closer look at the modeled daily profiles. Figure 9 shows daily averages of the hourly curve for three consecutive weeks starting on a Monday (April 8, 2013) and ending on a Sunday (April 28, 2013).

INSERT FIGURE 9 AROUND HERE

The daily shaping (within the weeks) are in line with the stylized facts pointed out in section 2.2. Indeed, the intra-weekly out-of-sample pattern from the model is very similar to the average historical daily weights displayed in Figure 2. It is clear that Fridays and weekends display lower price levels than weekdays. Adjustments for holidays and bridge days are accounted for internally in the model and they are generated automatically by the networks. In Figure 10 we give an example of a single week containing the Swedish National holiday on June 6th.

INSERT FIGURE 10 AROUND HERE

We now turn to the hourly (intra-daily) pattern generated by the model. In section 2.2 we established that historical hourly spot prices show distinct intra-day seasonal patterns with variations for different day types (Monday-Sunday). These effects are included in the model. In Figure 11 we display an example of the hourly prices in the curve between November 4 and November 24, 2013.

INSERT FIGURE 11 AROUND HERE

It is clear that the hourly profiles are present and that they vary for different day types. We confirm that Mondays-Thursdays exhibit quite similar behavior, while Fridays show a less pronounced evening peak. Saturdays and Sundays both display a different pattern with delayed morning and evening peaks. In section 2.2 we moreover pointed out that intra-daily profiles change appearance over the course of the calendar year. We noted that warmer seasons typically show larger differences between peak and off-peak compared to colder seasons. We will now confirm this property in the model. In Figure 12 we have plotted three consecutive weeks in January next to three consecutive weeks in June.

INSERT FIGURE 12 AROUND HERE

The two series are at different price levels (June obviously at the lower level) but it is clear from the figure that the difference between peak and off-peak prices is larger in June. We calculated the average price difference between peak and off-peak in the curve to be 7.28 EUR/MWh in June vs 5.68 EUR/MWh in January. A final notion is that the evening peak present in January is very modest in June.

The previous analysis is based on a single calibration of the hourly forward curve under the assumption of a normal hydrological state. Let us now change the assumption on the hydrology and generate two additional shape vectors where the first case is for a wet scenario (26 TWh above the normal level), and the second case for a dry scenario (26 TWh below the normal level). Both shape vectors are generated according to the same approach as before, however, now we explicitly inform the model on the hydrological balance for the future time period covering the term-structure of the curve (January 1, 2013 to December 31, 2013). In a practical setting this would typically be a forecast of the hydrological balance. We calibrate the curve to the same input prices as before (see Table 1). The purpose of this comparative static analysis is to study how the shape of the curve is affected by changes in the hydrological balance keeping all other conditions unchanged. Naturally, in a real world setting a wet (dry) scenario would decrease (increase) the price level of the curve, but in this analysis we disregard that effect and focus solely on the changes in the shapes. In fact, in a real world setting the input forward prices would reflect information on the expected hydrological balance. In Figure 13 we show three hourly forward curves (wet, dry and normal) for the time period between September 9 and September 29, 2013.

INSERT FIGURE 13 AROUND HERE

The curves are comparable since they have been calibrated to the same input prices. Comparing the wet scenario to the normal scenario we note that the curve displays a more stretched intra-daily profile with a larger difference between peak and off-peak prices. This is consistent with the reasoning in section 2.2 where we argued that a considerable hydrological oversupply results in a situation where hydro producers might run into difficulties managing the water for off-peak hours, which results in lower prices (in relation to the peak hours). This is moreover in agreement with the previous example from our historical data set showed in Figure 5. Next, we compare the dry scenario to the normal scenario, and we conclude that the intra-daily shape shows a more compressed profile with a smaller difference between the peak and off-peak prices. Again, this is in line with the reasoning in section 2.2. We finally note that the wet scenario has a larger impact on the profiles in the curve compared to the dry scenario. I.e. the hourly profiles appear to have greater sensitivity to wet scenarios. This can be visually verified in Figure 13.

5 Application: Pricing a Strip of Hourly Call Options

5.1 Background

In 2001 the French electricity producer EDF acquired a controlling interest in the German electricity utility EnBW. For reasons of competition the European Commission ruled that EDF was obliged to auction Virtual Power Plant (VPP) capacity in the French electricity market. The auctions took place on a quarterly basis for several years coming to an end in 2011 when EDF was released from the commitment. A similar situation took place in Denmark in 2004, where the Danish Competition Council approved a merger between two major electricity producers under certain conditions, including the commitment to regularly auction VPP capacity. The Danish auctions are still ongoing.

In both of these situations the auctioned contracts are constructed so that the buyer has the right, but not the obligation to purchase hourly power at a fixed price during a pre-defined settlement period. The contract is therefore equivalent to a series of independent call options on the hourly electricity spot price. Contracts of this type are moreover traded in secondary OTC markets where the main players are major energy trading houses and investment banks.

5.2 Valuation Method

Valuation of hourly electricity derivatives is a complicated matter since it involves stochastic modeling of hourly prices that have very complex dynamics. Similar to the approach in Branger, Reichmann and Wobben (2010) we argue that an hourly forward curve is a natural starting point in the valuation of an hourly VPP. An hourly forward curve supplies us with an hourly structure containing all relevant seasonal patterns while simultaneously being consistent with traded forward prices. Hence, it will successfully capture the intrinsic value of the strip of call options that makes the VPP, but being a deterministic object it will not capture the time value. In order to estimate the time value we need to make use of a dynamic stochastic price process for the electricity spot price. In this paper we employ the simple model specified in section 3.3. Admittedly, this price process is not a realistic alternative since it is normally distributed and has no jump component, but it's good enough for our purposes, which is to study how the value of the VPP reacts to changes in the hydrological state. Specification of a realistic model for hourly electricity prices lies outside the scope of this paper.

Since physical electricity is non-storable (it has to be consumed instantly after production) we cannot apply the traditional derivative valuation methods based on the results in Black and Scholes (1973). From a theoretical point of view non-storability of the underlying asset creates a situation with infinitely many martingale measures, which stands in sharp contrast to the Black and Scholes (1973) economy that has one single unique martingale measure. In order to price derivative assets we therefore need to employ methods originally developed in the context of interest rate markets where the situation is similar. These methods are based on calibration of the spot model to the hourly forward curve via the market price of risk. Motivated by Benth, Ekeland, Hauge and Nielsen (2003) we consider a parameterization of the market price of risk which is very flexible when fitting to the hourly forward curve. Specifically, we use the risk-neutral specification of equation (11) and parameterize the market price of risk $\lambda(t)$ to be piecewise constant for each hour. We then pick the equivalent martingale measure $\mathcal{Q}$ by calibrating the market price of risk such that the following arbitrage-free price relation is fulfilled

$$F(t,T) \quad = \quad E_{\mathcal{Q}}[S(T)|\mathcal{F}_t] \quad \text{for all hours} \quad T = 1, ..., \mathcal{T} \tag{14}$$

where $S(T)$ is the de-seasonalized spot price, $F(t,T)$ is the hourly forward curve and $\mathcal{T}$ denotes the last hour on the forward term structure. Within the class of equivalent martingale measures we have picked the probability measure $\mathcal{Q}$ which is closest to the hourly forward curve (which in turn is consistent with the market). This approach is similar to what is used in interest markets, see e.g. Björk (2004). We furthermore note that the our choice of $\mathcal{Q}$ re-introduces all seasonal patterns through the hourly forward curve. With this setup it is now possible to estimate the price of any derivative using standard Monte Carlo methods. In this paper we simulate the process in equation (11) using the exact simulation algorithm (i.e. without any discretization error) outlined in Glasserman (2004).

Before putting the above pricing scheme to work we need to estimate the speed of mean-reversion (α) and the volatility (σ) in equation (11). In order to do this we must first create a de-seasonalized hourly price series that could be used for estimation. We denote the hourly spot price (including all seasonal patterns) as $S^*(T)$ and assume that

$$S(T) \quad = \quad \frac{S^*(T)}{\Lambda(T)} \tag{15}$$

where $S(T)$ is the de-seasonalized price and $\Lambda(T)$ is the seasonal pattern. We estimate $\Lambda(T)$ by calibrating the full shaping model from section 3.1 to yearly historical spot prices. More precisely, we apply the trained shape model (the hourly and daily networks in combination with the empirical monthly shape) to generate a shape vector for the historical time period January 7, 2002 to December 25, 2011. This is the same time period as used in our data set. Each yearly segment of the shape vector is now calibrated (using the method in section 3.2) to the corresponding historical yearly average prices. In this way, we have created a series of historical hourly spot prices, containing all seasonal patterns and holidays/bridge days, which at the same time reflects the correct historical yearly average prices. We regard this series to be an estimate of the function $\Lambda(T)$. It is now straightforward to estimate the de-seasonalized hourly historical (log-)prices $X(t)$ from

$$\widehat{X}(T) = \ln \widehat{S}(T) = \ln \left(\frac{S^*(T)}{\widehat{\Lambda}(T)} \right) \tag{16}$$

With the estimated process $\widehat{X}(t)$ we may proceed to estimate the parameters in equation (11). Model estimation is carried out using Maximum Likelihood (ML), which is a well known method of estimation for processes of the type in equation (11). We estimate the speed of mean-reversion to be $\widehat{\alpha} = 181.95$ and the volatility to be $\widehat{\sigma} = 5.94$.

5.3 Valuation Results

Once more we make use of the three hourly forward curves from section 4. We stress that the curves have been calibrated to the same input prices (given in Table 1), but each shape vector has been generated under a different hydrological scenario; (i) normal, (ii) wet (26 TWh above normal) and (iii) dry (26 TWh below normal). Using the method of valuation outlined above we shall now price an hourly VPP with a yearly settlement period. Similar to the exercise in section 4 we keep the price level of the curve fixed (in line with input prices) and focus solely on the changes in the shapes, and how they impact the price of the VPP. Assuming we stand at November 20, 2012 (which is the trading date for the input prices of the curve) we price an hourly VPP with settlement period between January 1, 2013 to December 31, 2013. The strike price of the VPP is $K = 37$ EUR/MWh for each hour, which roughly corresponds to the yearly average price. Hence, the strike is set at the average at-the-money level. We assume that the risk-free interest rate is $r = 0$ in all calculations. In Table 2 we give the intrinsic value (IV), the time value (TV) and the total option value for the three scenarios. The intrinsic value is calculated according to

$$IV = \frac{1}{N} \sum_{i=1}^{N} \max(F(0, T_i) - K; 0) \tag{17}$$

where $F(0, T)$ is the hourly forward curve at time $t = 0$ (November 20, 2012), and N is the number of hours. In the normal scenario we report an IV of 2.24 EUR/MWh and a TV of 2.90 EUR/MWh, which gives a total option value of 5.15 EUR/MWh.

INSERT TABLES 2 and 3 AROUND HERE

In Table 3 we show the respective percentage contribution of the IV and TV to the total option value. We see that the IV in the normal scenario accounts for about 43.58% of the total option value, while the remaining contribution, 56.42%, consists of TV. If we instead make the same calculation in the wet scenario we see from Tables 2 and 3 that the IV increases to 2.85 EUR/MWh which corresponds to 52.66% of the total option value. The reason for this increase is that a wet scenario gives a larger price spread between peak and off-peak prices (see analysis in section 2.2), which means that the options (hours) that are in-the-money will move even further in-the-money, and the options (hours) that are out-of-the-money will move even further out-of-the-money. This obviously increases the IV. While the wet scenario brings on an increasing IV we simultaneously observe a decreasing TV. From the Tables 2 and 3 we see that the TV decreases to 2.56 EUR/MWh which corresponds to 47.34% of the total option value. The reason for this decrease lies in the nature of options. According to option pricing theory the TV of a given option is at its maximum when it is at-the-money. Any change in the moneyness (away from the at-the-money level) will hence lead to a decrease in the TV. We also note that the total option value in the wet scenario (5.41 EUR/MWh) is higher compared to the premium in the normal scenario (5.15 EUR/MWh). This means that the net effect of the changes in the IV and the TV (due to the change in the hydrological balance) brings on a total increase in the option premium, despite the fact that all input prices have been kept unchanged. We finally make the same calculations for the dry scenario. According to the Tables 2 and 3 we see that the IV in the dry scenario decreases in comparison to the normal scenario to the level of 2.02 EUR/MWh, which corresponds to 40.04% of the total option value. The decrease is due to the shrinking spread between peak and off-peak prices that occurs in dry situations (see analysis in section 2.2). The decreasing price spread gives the consequence that in-the-money options (hours) become less in-the-money, and out-of-the-money options (hours) move towards the at-the-money level. This gives a lower IV. Opposite to the wet case the dry scenario brings on a decreasing IV while simultaneously giving an increase in the TV. The TV increases (in relation to the normal scenario) to 3.02 EUR/MWh which makes 59.96% of the total option value. We finally note that the total option value in the dry scenario (5.04 EUR/MWh) is somewhat lower than the premium in the normal scenario (5.15 EUR/MWh). Similar to the previous case this is due to a positive net effect from the changes in the IV and the TV.

6 Conclusions

We have suggested a model to estimate the seasonal shapes that provides the basis for a power market hourly forward curve. The approach is based on feed-forward Artificial Neural Networks (ANN) in combination with an empirical monthly shaping method. We generalize the methodology suggested in Crispin and Jacobsson (2007) by including a connection between the seasonal specification and a fundamental factor (the hydrological balance). In the initial part of the paper we analyze hourly system price data from the Nordic electricity market and establish a series of stylized facts that we require the trained shaping model to account for. From the analysis we establish that the model is required to account for realistic hourly (intra-daily) and daily (intra-weekly) profiles (including seasonal variations in profiles), holidays and bridge days, annual seasonality in the price level, and finally we want the profiles to depend on the hydrological balance, which is an important fundamental factor in the Nordic power market with an impact on profiles and price levels.

The ANN shaping model consists of two separate networks, one for the hourly profiles and one for the daily. Both networks have been trained on historical price data for the Nordic system spot price between January 7, 2002 and December 25, 2011. We employ a simple empirical method to estimate the annual seasonal cycle from historical forward prices from the Nordic market (yearly, quarterly and monthly products). Estimates are based on historical data between 2008-2009. Finally, we combine the results from the hourly/daily networks and the empirical monthly method to obtain a single hourly vector containing all shaping information. In order to finalize the hourly forward curve we calibrate the shape vector to Nordic forward prices from the Nasdaq OMX Commodities market using the linear scaling method suggested in Burger et al. (2008).

In a subsequent analysis we conclude that the final model successfully accounts for all stylized facts initially stated. In particular, we focus on how the hourly forward curve depends on changes in the hydrological balance. We perform a comparative static analysis (with fixed input prices) where we feed the curve with different hydrological scenarios (normal/wet/dry) only to conclude that the model reacts in agreement with historical data and the expected behavior of hydro producers. A wet scenario (compared to normal) increases the price spread between peak and off-peak prices, while the dry scenario gives the opposite result.

In the final section we combine the proposed hourly forward curve with a simple stochastic Ornstein-Uhlenbeck model in order to calculate the price of a strip of hourly call options on the electricity spot price under different hydrological scenarios. This type of contract is commonly known as a Virtual Power Plant (VPP) and it makes a real world example of and OTC product that depends on hourly profiles. In a comparative static analysis (with fixed input prices) we conclude that changes in the hydrological balance impacts the intrinsic value and the time value of the VPP. It turns out that the intrinsic value increases in the wet scenario (compared to normal) while the time value decreases. The total option value (which gives the net effect of the changes in the intrinsic value and the time value) also increases although input prices are fixed. In the dry scenario we see the opposite effect with decreasing intrinsic value (compared to normal) and increasing time value. The total option value is somewhat lower compared to the normal scenario.

Figures and Tables

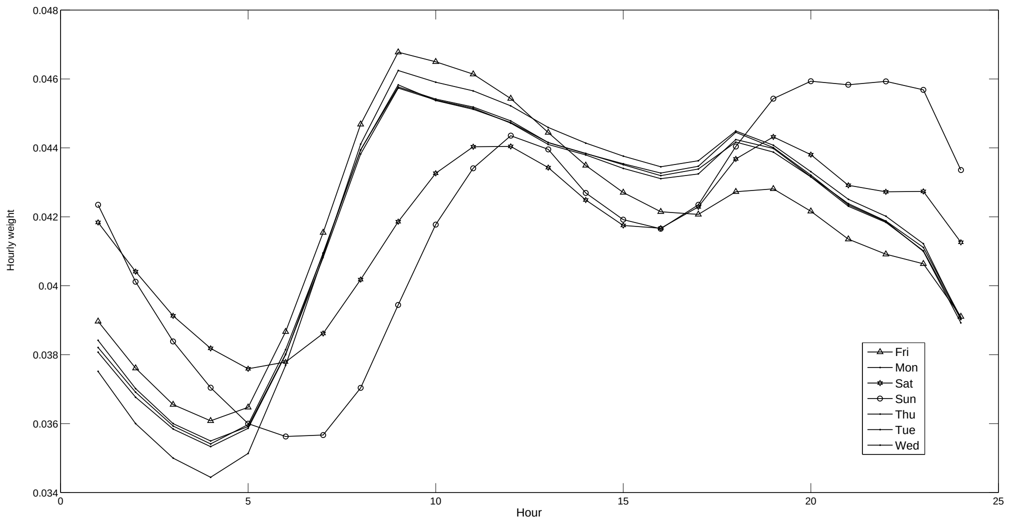

Figure 1: Average hourly weights for all day types (Monday-Sunday) based on historical system price data from the Nordic power market between January 7, 2002 and December 25, 2011. The weights have been normalized to sum to unity for all day types.

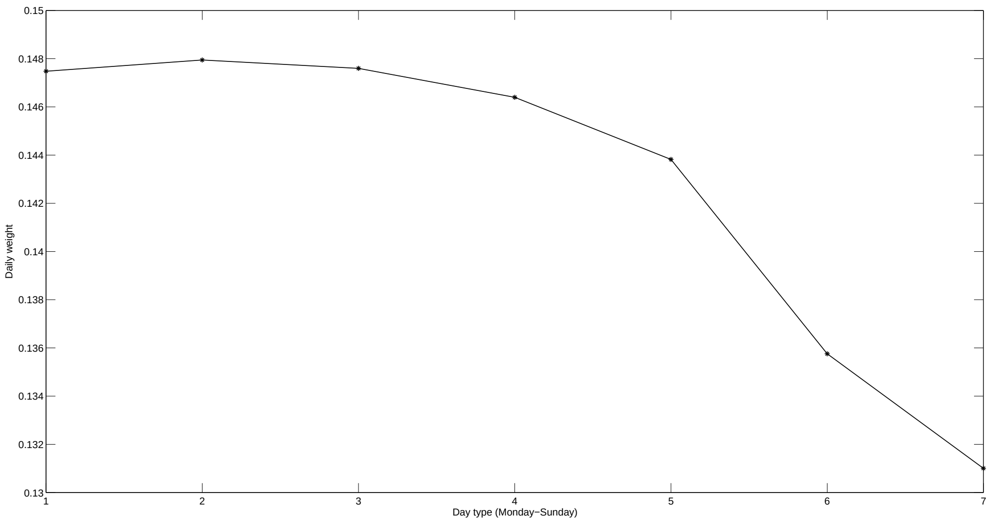

Figure 2: Average daily weights for each day type (Monday-Sunday) based on historical system price data from the Nordic power market between January 7, 2002 and December 25, 2011. The daily weights have been normalized to sum to unity for the week.

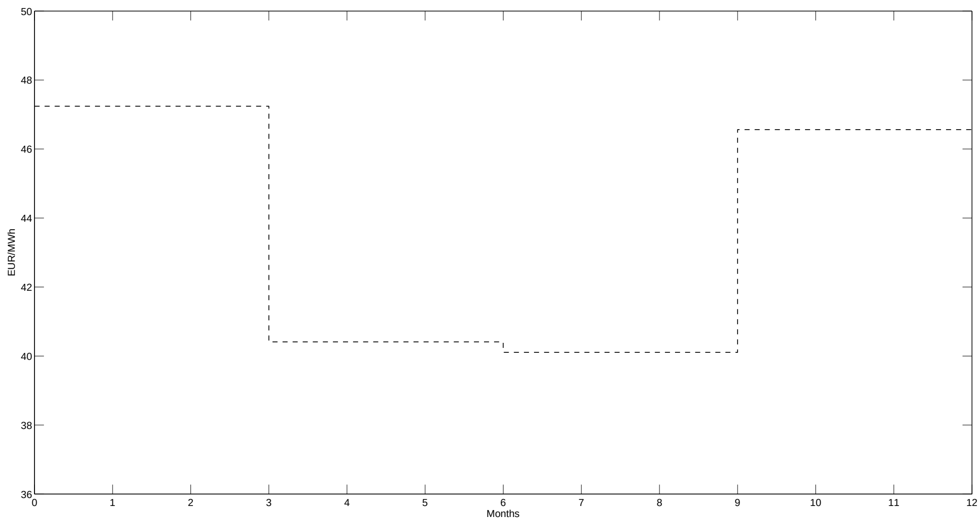

Figure 3: Prices for quarterly Nordic forward contracts (settlement period 2013) traded at the Nasdaq OMX Commodities exchange on November 25, 2011. The annual seasonal cycle is clearly visible.

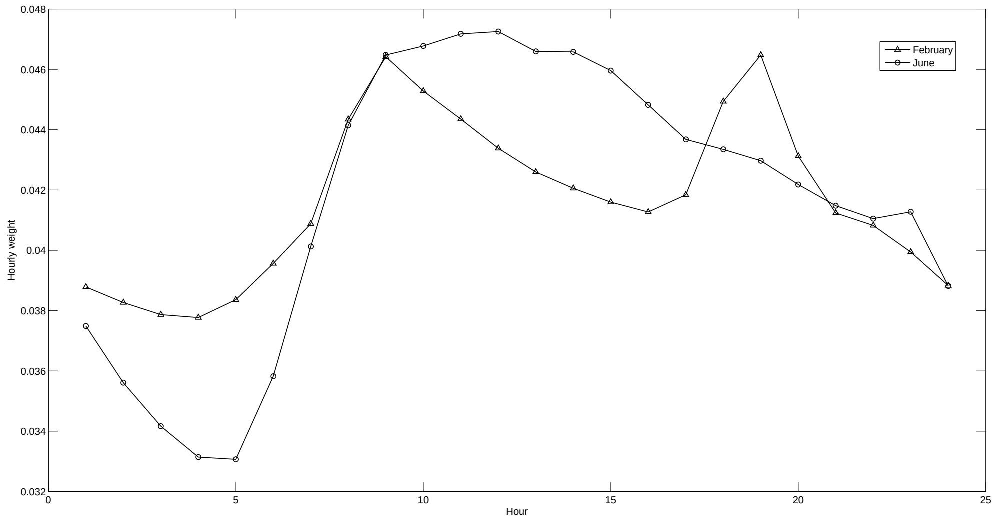

Figure 4: Two normalized intra-day profiles for average Wednesdays in February and June respectively. It is clear that the hourly profiles exhibit distinct differences depending on the season. The calculations are based on data from the Nordic power market between January 7, 2002 and December 25, 2011. The hours have been normalized to sum to unity for each day type.

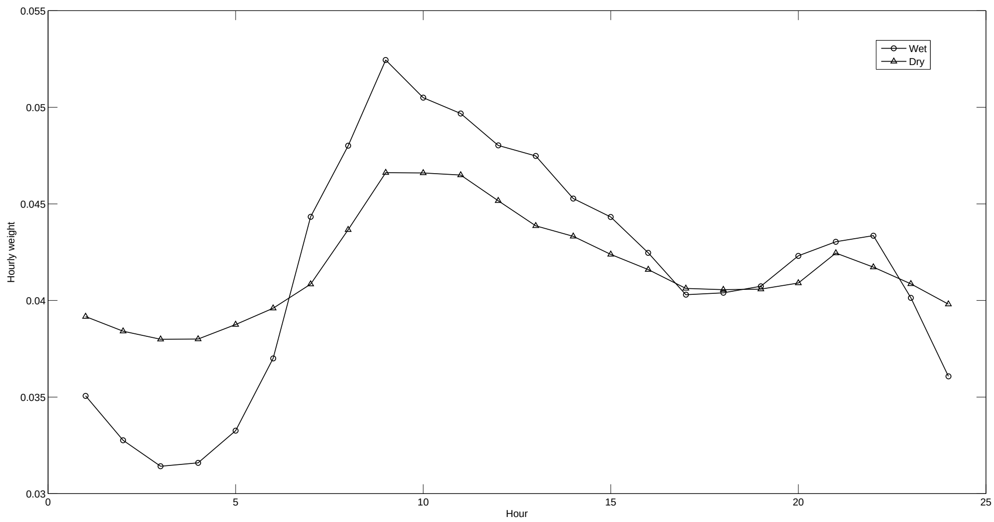

Figure 5: The average normalized intra-day hourly weights for Wednesdays in April for two hydrological situations; wet and dry. Calculations are based on system price data from the Nordic power market for the years 2006 and 2008, where the average hydrological balance (deviation from normal state) in April was -19.43 TWh and +13.41 TWh, respectively. The hours have been normalized to sum to unity.

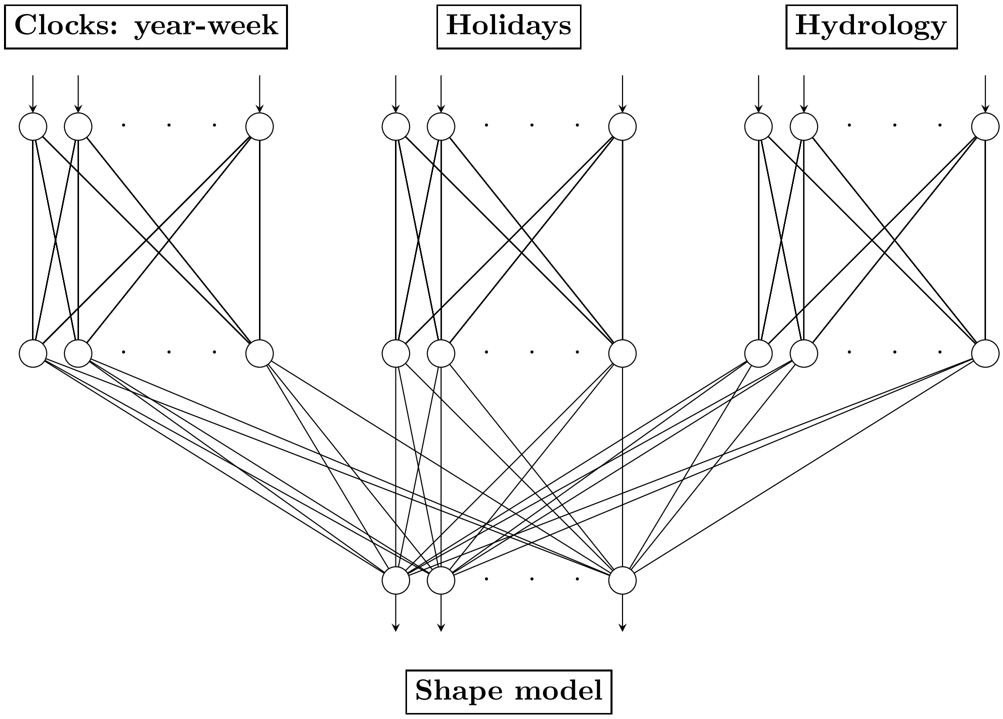

Figure 6: A schematic description of the architecture of the Artificial Neural Networks used for hourly and daily shape estimation. The networks are restricted (guided) since the neurons governing the yearly/weekly clocks, the holidays and bridge days, and the hydrology, are respectively disconnected.

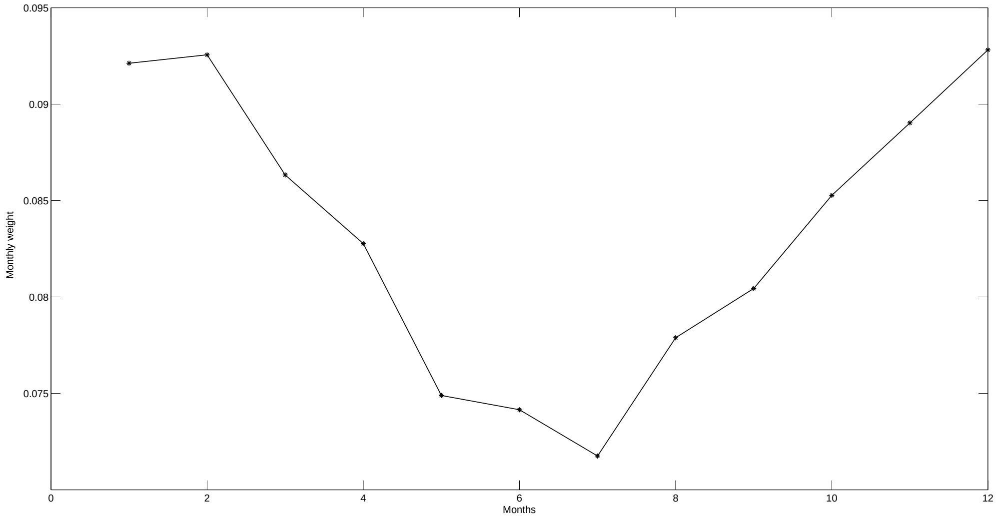

Figure 7: The figure displays the monthly shape vector estimated according to the methodology in section 3.1.3. Estimation is based on historical forward prices from the Nordic power market (monthly, quarterly and yearly contracts) traded at the Nasdaq OMX Commodities exchange between 2008-2009. The monthly weights have been normalized to sum to unity for the year.

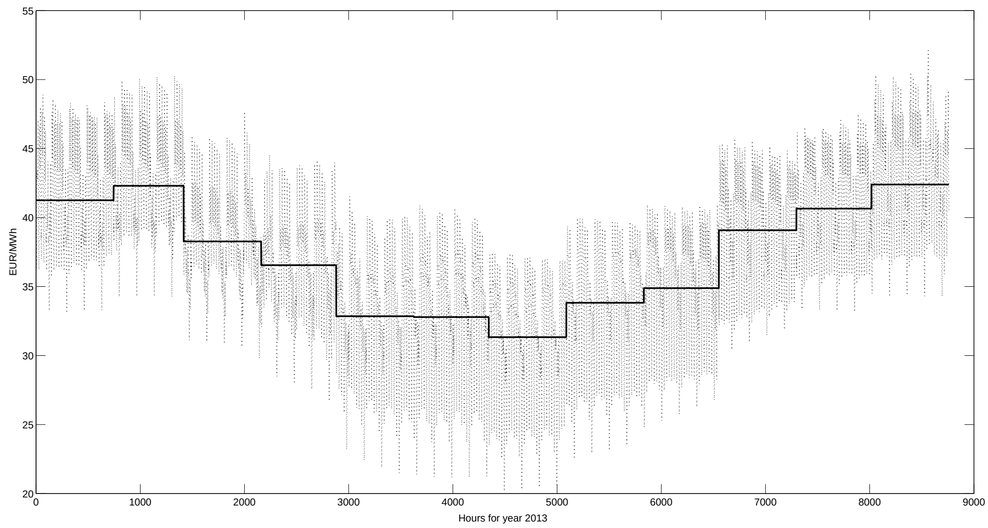

Figure 8: The calibrated hourly forward curve for the full calendar year 2013. Calibration date is November 20, 2012. The plot displays the curve at hourly granularity in combination with its monthly averages. The yearly seasonal cycle is clearly present.

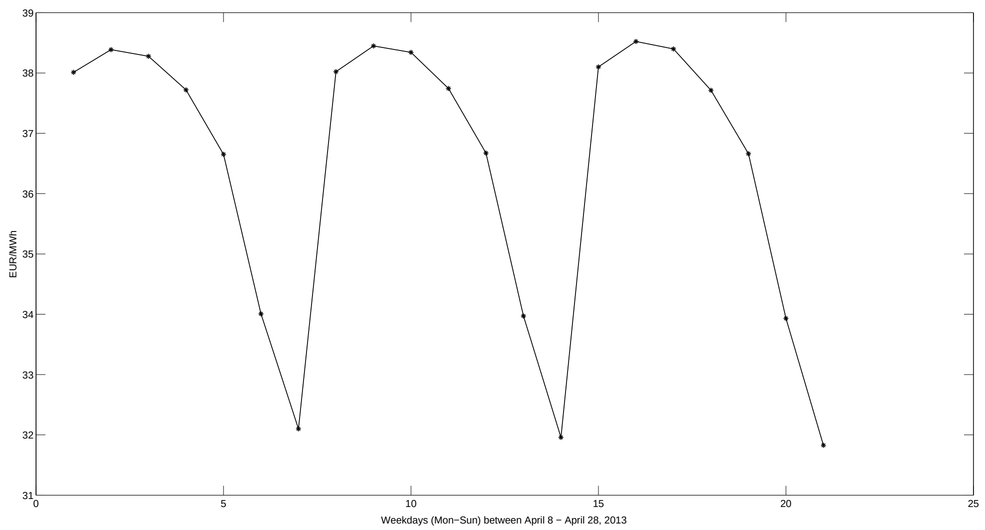

Figure 9: Daily averages of the hourly forward curve to illustrate the daily shaping within a week. Calibration date is November 20, 2012. The plotted time period is between April 8 and April 28, 2013.

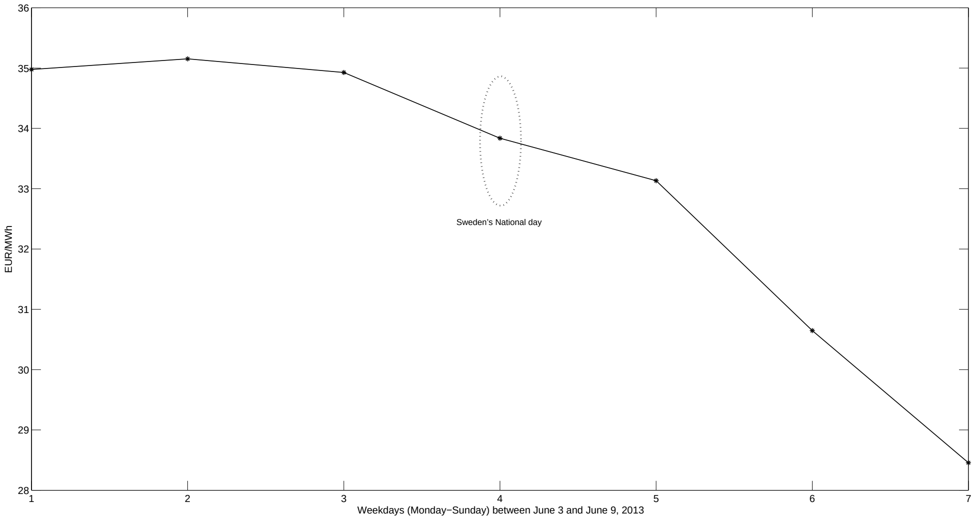

Figure 10: Daily averages of the hourly forward curve for the time period between June 3 and June 9, 2013. Calibration date is November 20, 2012. The price on Sweden's National holiday is scaled downwards according to the pre-estimated weights entered into the networks.

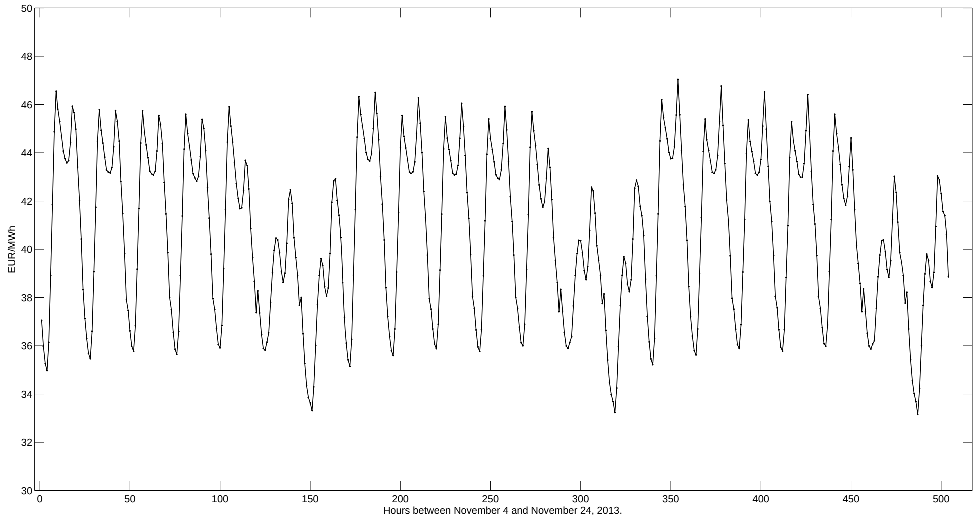

Figure 11: The hourly forward curve for the time period between November 4 and November 24, 2013. Calibration date is November 20, 2012. The hourly profiles are clearly present and they vary for different day types.

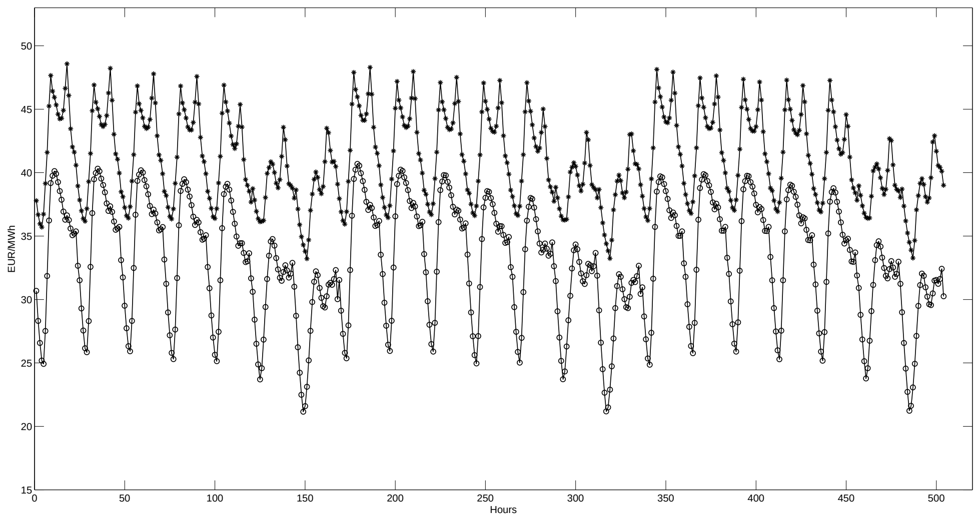

Figure 12: The hourly forward curve for three consecutive weeks in June (June 10 - June 30) vs dito in January (Jan 7 - Jan 27). Calibration date is November 20, 2012. The difference between peak and off-peak prices is evidently larger in June, in combination with a significantly less pronounced evening peak.

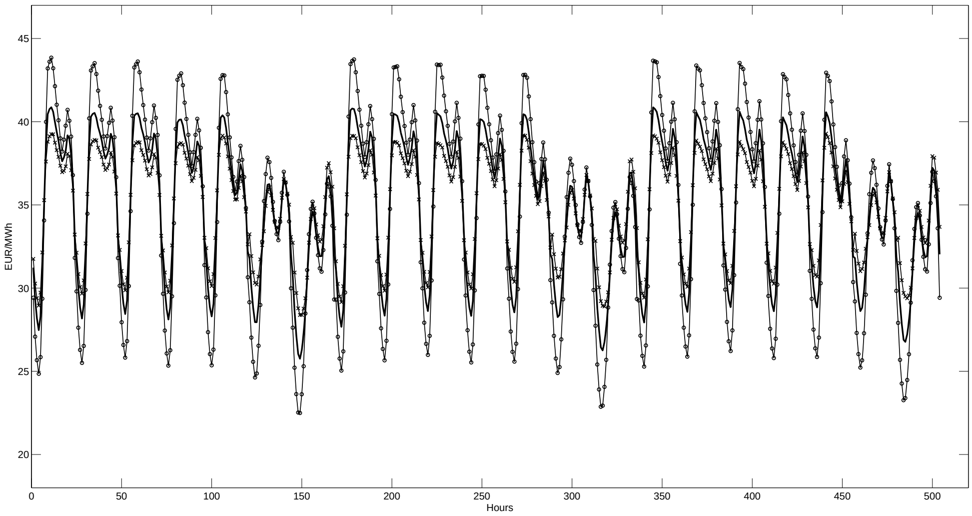

Figure 13: The hourly forward curve for three consecutive weeks in September (September 09 - September 29) generated for three different hydrological scenarios; (i) normal, (ii) wet (26 TWh above normal level) and (iii) dry (26 TWh below normal level). Calibration date is November 20, 2012. The spread between peak and off-peak prices is large (small) in the wet (dry) scenario.

Contract	Input price (EUR/MWh)
$M1$	41.25
$M2$	42.30
$M4$	36.55
$M5$	32.85
$Q1$	40.55
$Q2$	34.05
$Q3$	33.33
$Y1$	37.15

Table 1: Contracts and input prices used for curve calibration. The prefix (M, Q and Y) in the first column respectively denote month, quarter and year, and the suffix (1, 2, 4 and 5) indicate the order of the contract in the term-structure. Price quotes are closing prices for Nordic power forwards traded at Nasdaq OMX Commodities per November 20, 2012.

EUR/MWh	**Normal**	**Wet**	**Dry**
Intrinsic value	2.24	2.85	2.02
Time value	2.90	2.56	3.02
Option value	5.15	5.41	5.04

Table 2: The table shows intrinsic values, time values and total option values (in EUR/MWh) for an hourly Virtual Power Plant (VPP) with settlement period between January 1, 2013 and December 31, 2013. All figures are given for three different hydrological scenarios (normal/wet/dry). Valuation date is November 20, 2012. The strike is set at $K = 37$ EUR/MWh, which is roughly the average at-the-money level for the given settlement period. We assume that the risk-free interest rate is $r = 0$.

Percent of Option value	**Normal**	**Wet**	**Dry**
Intrinsic value	43.58%	52.66%	40.04%
Time value	56.42%	47.34%	59.96%
Option value	100%	100%	100%

Table 3: The table shows the respective percentage contributions of the intrinsic values and the time values to the total option value. The figures are given for the same scenarios and the same contract parameters given in Table 2.

References

Benth, F., Ekeland, L., Hauge, R. and Nielsen, B. (2003), 'A note on arbitrage-free pricing of forward contracts in energy markets', *Applied Mathematical Finance* **10**, 325–336.

Benth, F., Koekebakker, S. and Ollmar, F. (2007), 'Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation', *Journal of Derivatives, Fall* **15**(1), 52–66.

Björk, T. (2004), *Arbitrage Theory in Continuous Time*, Oxford University Press, Oxford.

Black, F. and Scholes, M. (1973), 'The Pricing of Options and Corporate Liabilities', *Journal of Political Economy* **81**, 637–659.

Branger, N., Reichmann, O. and Wobben, M. (2010), 'Pricing electricity derivatives on an hourly basis', *The Journal of Energy Markets* **3**(3).

Burger, M., Graeber, B. and Schindlmayr, G. (2008), *Managing Energy Risk: An Intergrated View on Power and Other Energy Markets*, Wiley Finance, West Sussex.

Crispin, D. and Jacobsson, C. (2007), General Forward Curve Specification: An Optimisation Approach with Price Shape Statistics, Working paper, E.ON Sales and Trading.

Fleten, S. and Lemming, J. (2003), 'Constructing forward price curves in electricity markets', *Energy Economics* **25**, 409–424.

Glasserman, P. (2004), *Monte Carlo Methods in Financial Engineering*, Springer-Verlag, New York.

Ripley, B. (2008), *Pattern Recognition and Neural Networks*, Cambridge University Press, New York.

Schwartz, E. (1997), 'The Stochastic Behaviour of Commodity Prices: Implications for Valuation and Hedging', *Journal of Finance* **LII**(3), 923–973.

Shreve, S. (2004), *Stochastic Calculus for Finance II, Continuous-Time Models*, Springer, New York.

Vasicek, O. (1977), 'An Equilibrium Characterisation of the Term Structure', *Journal of Financial Economics* **5**, 177–188.

A Power Market Forward Curve with Hydrology Dependence

An Approach based on Artificial Neural Networks

RIKARD GREEN

This paper develops an hourly forward curve for power markets where the intra-day and intra-week shapes (profiles) depend on the level of the hydrological balance. The shaping model is based on a feed-forward Artificial Neural Network (ANN), which is trained on a historical data set of hourly electricity spot prices from the Nord Pool market and weekly measurements of the Nordic hydrological balance. The yearly seasonal cycle is estimated with historical electricity forward prices from the Nasdaq OMX Commodities exchange. We calibrate the shaping model to prevailing electricity forward prices and proceed to demonstrate its most important properties. By using comparative static analysis we particulary focus on the hydro dependence of the shapes. We conclude the paper with a real world valuation task. By combining our proposed forward curve with a simple Ornstein-Uhlenbeck process we price a strip of hourly call options on the electricity spot price under different hydrological scenarios.

Keywords: Power markets, forward curve, seasonality, artificial neural networks

JEL: C45, C51, G13, Q41

THE KNUT WICKSELL CENTRE FOR FINANCIAL STUDIES

The Knut Wicksell Centre for Financial Studies conducts cutting-edge research in financial economics and related academic disciplines. Established in 2011, the Centre is a collaboration between Lund University School of Economics and Management and the Research Institute of Industrial Economics (IFN) in Stockholm. The Centre supports research projects, arranges seminars, and organizes conferences. A key goal of the Centre is to foster interaction between academics, practitioners and students to better understand current topics related to financial markets.

LUND UNIVERSITY
School of Economics and Management

LUND UNIVERSITY
SCHOOL OF ECONOMICS AND MANAGEMENT
Working paper 2015:1
The Knut Wicksell Centre for Financial Studies
Printed by Media-Tryck, Lund, Sweden 2015