

# When Supply Meets Demand: The Case of Hourly Spot Electricity Prices

Alexander Boogert, *Student Member, IEEE*, and Dominique Dupont

**Abstract**—We use a supply-demand framework to model the hourly day-ahead price of electricity (the spot price) based on publicly available information. With the model we can forecast the level and probability of a spike in the spot price defined as the spot price above a certain threshold. Several European countries have recently started publishing day-ahead forecasts of the available supply. This paper shows this forecasted capacity is quite successful in predicting spot price movements 24 h ahead.

**Index Terms**—Available capacity, day-ahead electricity market, price forecasting.

## I. INTRODUCTION

**D**AY-AHEAD electricity prices (known as spot prices) serve as an important reference to all members of the electricity industry. These prices are characterized by high volatility and rare but violent spikes. These aspects have motivated numerous research efforts. In this paper we investigate the spot price of electricity in a supply-demand framework.

Spot electricity price models can be used for short- and long-term goals. For short-term tactical planning, a good forecast of the absolute height of the day-ahead prices and the probability of a spike is essential. For long-term strategic planning, an estimate of the volatility of the spot price can be used to assess the value of power plants. In this paper however we focus mainly on the short-term goals.

Our objective is to establish a relation between several fundamental drivers and hourly spot electricity prices. In particular we will investigate the role of available capacity. Using hourly prices instead of daily prices has two main advantages. First, it makes it possible to explain the different patterns of prices over the day. Secondly, it increases the sample size and hence the likelihood of obtaining robust empirical results. Accurate forecasts of demand and supply are of paramount importance to the electricity industry because these two must be balanced at all times to maintain the stability of the power grid.

Forward electricity contracts are traded several years before actual delivery. Contracts are traded both on OTC markets and on organized exchanges. Delivery is typically channelled

Manuscript received June 5, 2007; revised December 7, 2007. The paper was written when D. Dupont was at the University of Twente. The information presented in this paper does not necessarily reflect the views of Essent Energy Trading or RWE Trading. Paper no. TPWRS-00377-2007.

A. Boogert is with the Commodities Finance Centre, Birkbeck College, University of London, London, U.K., and Essent Energy Trading, 5211 AK Hertogenbosch, The Netherlands (e-mail: alexander.boogert@essent.nl).

D. Dupont was with the University of Twente and FELab. He is now with RWE Trading GmbH, 45128 Essen, Germany (e-mail: dominique.dupont@rwe.com).

Digital Object Identifier 10.1109/TPWRS.2008.920731

through day-ahead and real-time markets, whose designs can differ substantially across geographical areas. Examples of design differences include the exact time of settlement, the granularity of the contracts (i.e., the time period over which power is to be delivered), the handling of actual delivery in real time, and the exact information provided to the public. In most cases, the day-ahead market is an auction: an independent auctioneer aggregates buy and sell orders and computes 24 hourly market-clearing prices for the following day. These are the day-ahead prices we discuss in this paper.

Most day-ahead markets operate at an hourly granularity, while the U.K. and Australia operate at an half-hourly granularity. The amount of publicly available information differs across markets. For example, a day-ahead estimate of the available supply is at the moment published for several markets, but it is not yet available everywhere. In this paper we consider the Dutch market where such information is available. The variety of market setups makes it necessary to adapt models to the local conditions to make them useful. However, these local adjustments generally do not affect the cores of the models.

This paper is organized as follows. In Section II we review the literature. In Section III we establish the supply-demand framework for a general electricity market. First we discuss factors which influence spot electricity prices. Then, we introduce an indicator combining supply and demand and use this indicator to forecast hourly spot prices one day ahead. Subsequently, we contrast our non-parametric approach with some parametric ones. In Section IV we introduce the situation in the Dutch market. We specify which data are available and apply the different techniques to forecasting an hourly spot price and the probability of a spike. Section IV ends with a study on the stability of the relationship. In Section V we discuss the implications of our findings for further modeling.

## II. REVIEW OF THE LITERATURE

For a long time, the modeling of electricity spot prices has focused on reduced-form models (e.g., [1], [2]). There are two popular modeling approaches: jump diffusion and regime switching. Both types of models are mathematically tractable and have received considerable attention. However, the estimation of those models tends to be delicate. Another route is provided by fundamental models (e.g., [3]), which carefully describe the characteristics of the supply stack in a market. In the case of a central planner the full supply stack is known and used to serve the load at the lowest cost, while in liberalized markets only the general shape of the daily supply stack is

known. Marginal cost curves obtained with fundamental modeling and estimates of the supply stack need to be transformed into spot prices.

Hybrid models incorporate ideas from both approaches. While reduced-form models use only past prices, hybrid models incorporate additional information, such as load, weather, or power plant availability. Examples of a class of hybrid models based on the assumption that there is an exponential relation between price and load are [4]–[9].

The exponential function captures the empirically observed pattern of strongly increasing prices for high loads and leads in some cases to the closed-form valuation of electricity derivatives. Derivative pricing is based on continuous-time models, while forecasting is mainly based on discrete-time models, e.g., in our approach. Both types of models can be hybrid. We will discuss the connection to continuous-time models in more detail in Section III-F. The approach we develop can be used in both discrete- and continuous-time models.

The main driver in our model is the reserve margin, i.e., the fraction of total supply available to cover demand. In practice, there are several definitions possible of the available capacity, which we will discuss in more detail in Section III-A. Besides spot electricity price modeling, the reserve margin is studied in research on public policy (e.g., [10]) and security of supply (e.g., [11]). In spot electricity modeling, reserve margin is studied in [12]–[16]. Anderson [13] and Burger *et al.* [14] choose a functional form to study the relations between reserve margin on one side and probability of spikes (Anderson) or the level of the spot price (Burger *et al.*) on the other. The index used by Burger *et al.* incorporates the expected relative availability of power plants and load, although the precise form is not given. Mount *et al.* [15] create a regime switching model where the switching probabilities between the regimes and the conditional mean of the power price in each regime vary with time and with the reserve margin. Zareipour *et al.* [16] discuss how different variables and models can improve the forecast of the day-ahead prices. They find a variant of the reserve margin is a useful indicator in the Ontario market.

### III. SUPPLY-DEMAND FRAMEWORK

In this section we discuss factors which influence spot electricity prices. Many factors besides past electricity prices may play a role. We start by investigating available capacity and how to forecast this factor. We then turn to additional price drivers.

#### A. Forecasting Available Capacity

In most electricity markets accurate information is available on the quantity and the price of power traded on the market in the past. By construction, supply equaled demand at those prices. More information about the state of the market can be obtained from the supply and demand curves. In some electricity markets there is a clear relation between the price and the trading volume on the day-ahead market because all supply has to be offered in the day-ahead market (e.g., in the old NETA system in England and Wales or currently in Spain). In these markets the supply and demand curves can be used to characterize the state of the

market. However, in some markets, including the Netherlands, only part of the supply is offered on the day-ahead market and subsequently no apparent relation between price and volume is observed.

The system load is an alternative to the trading volume on the day-ahead market. The system load stands for the demand for power within some area. The relation between system load and market price is known to be stronger, which is confirmed in our sample (see middle panel of Fig. 3). The bidding curve on the day-ahead market gives the demand (and supply) curves as a function of price. This makes it easy to determine the demand elasticity, but it only presents the demand within the day-ahead market. The derivation of the demand elasticity from a single number like the system load is more difficult.

Another reason for the absence of data comes from a delay in data release. For example, in the PJM market the supply and demand curves are released with a delay of 6 months. Besides, Mount *et al.* [15] note that the available capacity could not be fully recovered from the public data on offered capacity and an assumption on the total available capacity was necessary.

In markets where the information from the day-ahead market cannot be used, the elasticity of supply and demand has to be determined by other means. One measure of the supply elasticity is the spare capacity available. This capacity crucially depends on the granularity of the market. In the very short run (e.g., within 15 min) only some flexible units can be turned on and the output of running units can only be marginally increased. On a day-ahead basis more of the capacity will be available.

Demand elasticity is normally not taken into account as consumers are generally price insensitive. However, it may be possible to temporarily reduce the power supply to some selected consumers, e.g., large industrial customers in the metal industry, who accept to run this risk (within contractual limits) in exchange for a discount in the power price. In financial terms, these customers buy an interruptible contract. In general, no public information is available on this type of contracts.<sup>1</sup> These demand elasticity effects are explicitly taken into account in Fezzi and Bunn [17] as a latent variable. In this paper we exclude this effect.

Within the existing literature most articles do not explicitly introduce supply because of a lack of data. Recently, the situation has improved as indicators have been introduced in several European electricity markets. Regulators are currently providing estimates for the available capacity in the Netherlands, U.K. and Germany. In this paper we will focus on the Dutch market, which has the longest history of the three for this type of indicator.

Another grey area in the definition of available capacity is the use of import and export capacity. The question is how to include potential import and export into the total available generation capacity. An important difference is created by the timing of the import/export capacity market relative to that of the day-ahead market. If the capacity market clears before the day-ahead market, the resulting price can serve as a leading indicator of the day-ahead price in the domestic market relative to the foreign market.

<sup>1</sup>We thank the referee for pointing out that Statnett does publish information in Norway.

### B. Additional Price Drivers

Natural price drivers are factors which impact supply or demand or both. Besides there can be feedback effects from prices in either the previous spot price or the most recent real-time prices. To give an indication about the variety of potential price drivers, we refer to [18]. They mention power supply factors (installed capacity, outages, generation resource mix, transmission constraints), demand factors (load duration, weather sensitivity, economic activity, retail price) and market design (retail price caps, revenue share of spot sales, capacity requirements and wholesale price caps) as possible price drivers.

### C. Relation Between Supply, Demand, and Spot Price

One of the goals of this paper is to understand the relation between supply, demand and the spot price. We focus on a simple economically motivated relation which allows us to study the underlying data carefully. Alternatively, one could apply data mining techniques like neural networks and adaptive splines (e.g., [19]). Following Anderson [13] we have decided to consider a demand-supply ratio ( $DSR$ ) of the following kind:

$$DSR := 1 - \frac{\text{demand}}{\text{available capacity}}. \quad (1)$$

This ratio represents the fraction of the supply available to cover demand. This means it is a reserve margin index. Economic intuition suggests this index should be negatively correlated with the spot price. The lower the index, the less capacity is available and the tighter the market becomes: more expensive units come online, and marginal costs increase. This index is closely related to the concept of capacity utilization used in Anderson [13]. Capacity utilization states how much of the available capacity is used to cover the demand.

Instead of a demand-supply ratio, one could opt for a supply-demand ratio (available capacity/demand - 1) or an absolute difference between the supply and demand (available capacity - demand). We prefer a ratio as it is dimensionless. We believe this improves the stability of the relation. We choose the demand-supply ratio because we face mainly measurement errors in supply, which play a lesser role in the demand-supply ratio than in the supply-demand ratio (see [10] for a rigorous support of this claim) or the absolute difference between supply and demand. Moreover, market participants also seem to prefer this ratio.

### D. Forecasting Hourly Spot Prices

One way to forecast the spot price is to consider the average relation between the reserve margin and hourly spot prices. This also allows us to establish a confidence interval around the forecast. From an economic viewpoint we expect this relation to increase for decreasing reserve margins. Moreover, we expect the bandwidth to increase for decreasing reserve margin. In the paper, we use both a piecewise linear fit and a smoothed  $b$ -spline fit.

This bandwidth represents an interval forecast of the price, which appears to receive attention since only recently. For example, Misiorek *et al.* [20, p. 23] note that "interval forecasts have not been investigated to date." The natural extension of

our relation is to consider a two-dimensional version of this approach. Such a step was taken in Lu *et al.* [19] where besides a reserve margin a steepness-of-load indicator is used.

### E. Forecasting the Probability of a Spike

The probability that the prices cross a certain threshold is an important variable for market participants besides the absolute height of the spot price. A spot price above the threshold is called a spike. In this paper we define the threshold as a fixed amount of Euros. An alternative would be to define the threshold in terms of the cost of the marginal unit for the specific hour under consideration. We estimate the probability as the relative number of observations above the threshold in our data sample.

### F. Parametric Approach

In our approach we assume there is a nonlinear relation between the reserve margin and spot prices. Another approach is to parametrize the relation. We can rewrite our reserve margin index as follows:

$$S_t = f\left(1 - \frac{D_t}{C_t}\right) \quad (2)$$

where  $S_t$  is the spot price for hour  $t$ ,  $D_t$  is the demand,  $C_t$  is the available capacity and  $f$  is a nonlinear function.

The variable  $1 - (D_t/C_t)$  takes values between 0 and 1, and  $S_t$  can take very high values. If one assumes a monotonic relation between reserve margin and price, it is reasonable to base  $f$  on the inverse of a cumulative distribution function (cdf) with infinite support and given in closed form, for example, the logistic distribution. This could motivate the function used in Anderson [13] (the function there is given without explicit motivation). Similarly, Barlow [21] makes the power price a function of a latent variable that follows a diffusion process (this variable need not be between 0 and 1). The function is built to contain a singularity, which pushes the price towards infinity in the neighborhood of the singularity. The inverse cdf technique can be seen as refinement of this technique.

Alternatively, one can treat supply and demand as separate stochastic processes, and introduce a functional form for the relation of the form  $S_t = f(D_t, C_t)$ . Using this functional form and an explicit link between day-ahead and forward prices, creates a possibility to study the forward risk premium. Bessembinder and Lemmon [22] formulated a general equilibrium model for the day-ahead forward prices, which they applied to the PJM market in the United States. Villaplana [23] extended the model by considering supply as a random variable and applied the model to the Nordpool market in Scandinavia. In these models the relation is assumed to be of exponential or power form, which simplifies the estimation of the parameters and yields a closed form solution for forward prices.

Simultaneously, it captures the observed feature that prices rise for increasing demand and decreasing capacity. For example, Villaplana [23] estimates

$$S_t = \gamma_1 C_t^{\gamma_2} e^{\gamma_3 D_t} \quad (3)$$

where  $\gamma_1$ ,  $\gamma_2$  and  $\gamma_3$  are constants.

A related version of our approach was independently created by Kanamura and Ohashi [24]. Besides a Box-Cox transformation, they parametrize the relation between price and load by two linear functions and one quadratic function. This comes close to our piecewise linear fit.

#### IV. APPLICATION TO THE DUTCH MARKET

We first discuss the structure of the Dutch market and the availability of data. Secondly, we describe how this data behaves and show how to make a forecast for the spot price and the probability of a spike. Finally, we discuss the stability of our relation.

##### A. Overview of the Dutch Market

The Netherlands was among the first countries in the European Union to liberalize its electricity market. The Dutch Independent System Operator, TenneT, manages the high-voltage grid (380 and 220 kV), which interconnects regional electricity networks and links the Dutch grid to Belgium and Germany. TenneT, a fully state-owned company, ensures access to the domestic high-voltage network and organizes, through its subsidiaries, the day-ahead market for electricity (Amsterdam Power Exchange or APX) and the imbalance market. It also auctions capacity at the five cross-border interconnectors. The maximum import under normal circumstances is 3650 MW, which can be increased to 3850 MW in case of emergency. The scheduled day-ahead import is not exactly realized in real-time. Although the electricity traded on the APX represents about 20% of the Dutch daily consumption, the APX price is considered an important benchmark.

In the Dutch market, the import/export capacity is auctioned before the day-ahead spot electricity while imported electricity has to be offered on the day-ahead market. The Netherlands typically imports power, often up to the full available import capacity. A new development [25] is the introduction of market coupling between the Netherlands, Belgium and France. Under this new system the import/export auction is integrated into the day-ahead auction and the auction is changed from explicit (before the day-ahead electricity market) to implicit (simultaneously with the day-ahead electricity market). We believe our method will remain valid. The new system does not change the load, maximum import or available capacity within the Netherlands, which as discussed below, are the explanatory variables in our approach. However, it might have an impact on the realized spot prices. Whether these have changed significantly lies outside the scope of this paper.

##### B. Available Data in the Netherlands

Before 2004, estimates of the available capacity were made public only if they had dropped below a certain threshold. As this rarely happened, it was difficult to estimate the available capacity and the demand-supply equilibrium in general. Boogert and Dupont [26] show that in that period the water temperature was a good indicator of the spike risk in the electricity price: when water temperatures cross a certain threshold, environmental regulations restrict the production, while demand increases due to cooling needs, which together lead to high prices.

Since 2004, TenneT publishes an estimate of the available capacity in the Dutch grid for the coming 30 days. TenneT gathers statements of the different generators about the availability of their individual plants and combines them in an aggregate index. The TenneT estimate covers the generation in the Netherlands except for wind-based production and smaller units (of less than 10 MW). The TenneT estimate is one way to describe the supply in the Dutch market. Besides this estimate, we think the following variables could help forecast the day-ahead electricity price:

- National load: realized generation including realized net import gives the load which is published by TenneT on a 15-min basis with a delay of two days. The official load data covers only electricity generated by units larger than 10 MW. Moreover, no official forecast is available.
- Realized import or export: history published by TenneT on a 15-min basis. We take import as a positive number since it adds to the available capacity. This information is published with a delay of 30 min.
- Maximum import: the maximum possible import and export is published by TenneT. A day-ahead forecast is available, together with announcements for future maintenance and enlargement in case of emergency. We received historical data from TenneT.
- Wind power: there is no official estimate of the total wind power production in the Netherlands. An internal estimate was provided by Essent Energy Trading.

This paper uses data starting 01/10/2004 and ending 17/06/2006.<sup>2</sup> The starting date coincides with the first publication of the available capacity forecast. As spot prices are published on an hourly scale, we transformed all 15-min data into hourly data by taking the average over the specific hour. Subsequent graphs all show hourly data. In total there are 14 904 hourly data points.

As mentioned in Section III-A, there are several ways to define available capacity. In the Dutch market, we need to make two choices. The first one deals with whether we should include realized imports (or exports) or the day-ahead forecast of maximum imports (or exports). The second one deals with whether we should include wind power capacity.

The potential for wind energy is growing in the Netherlands. Given its size, it could be interesting to include wind power into the total available capacity. However, as the production data are not public, we exclude wind power from the available capacity. Concerning the import/export number, we use the day-ahead forecast of maximum possible import/export. Thus, to create an estimate of the total available capacity  $C_t$ , we sum the TenneT estimate and the maximum import.

##### C. Reserve Margins in the Netherlands

We start by showing the development of the underlying data for the reserve margin over time. Fig. 1 shows the load, forecasted available capacity and day-ahead maximum possible import capacity. For contrast we also include the realized values of the import in the bottom panel. In Fig. 1, load (the upper panel) varies between 7000 and 16 000 MWh while forecasted available capacity (the middle panel) varies between 12 000 and

<sup>2</sup>For convenience we deleted the four days with daylight saving hours in our sample: 31/10/2004, 27/03/2005, 30/10/2005, and 26/03/2006.

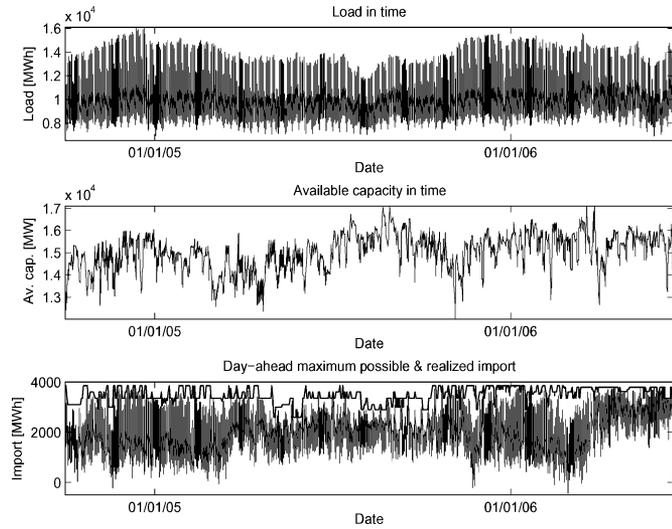


Fig. 1. National load (top), available capacity (middle), maximum import capacity (lower, bold line), and realized net import (lower, thin line), 01.10.2004–17.06.2006.

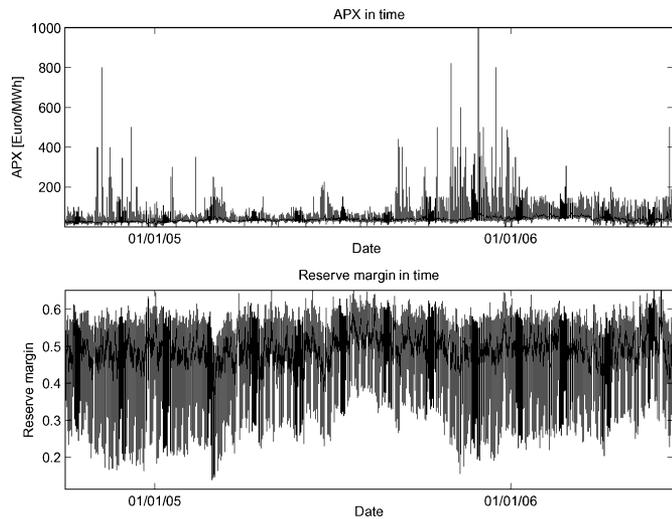


Fig. 2. APX price (top) and reserve margin in the Netherlands (bottom), 01.10.2004–17.06.2006.

17 000 MW. Realized import is significantly more volatile than maximum import (the lower panel). This is due to the uncertainty in the actual demand.

Fig. 2 shows the development of both the APX price and the reserve margin over time. A scatter plot in the upper panel of Fig. 3 reveals a pattern of increasing prices with reserve margin. For completeness, the figure also graphs the APX against the national load (the middle panel) and against the available capacity (the lower panel). High prices are sometimes observed with medium index values. We discuss the stability of our relation in Section IV-F.

D. Forecasting the Spot Price

One way to use the reserve margin to forecast the spot price is to consider the average relation from reserve margin to APX prices. In Fig. 4 we show a piecewise linear fit and a *b*-spline fit. The piecewise linear fit was created by placing the reserve

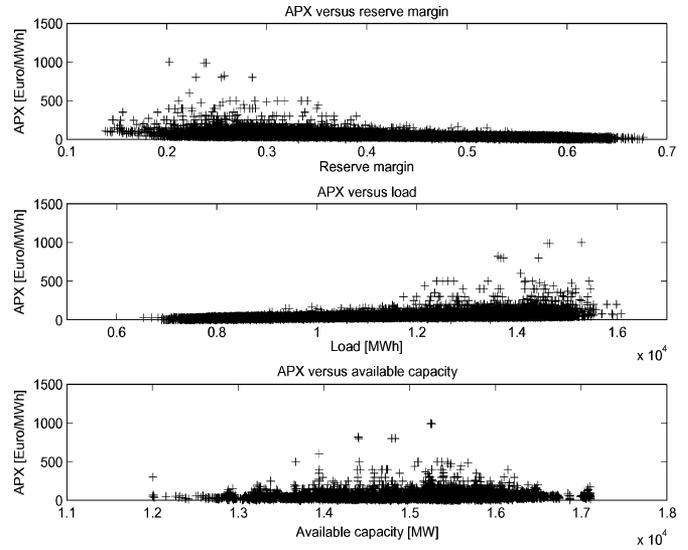


Fig. 3. Scatter plots of APX price versus respectively national load (top), available capacity (middle), and reserve margin (bottom).

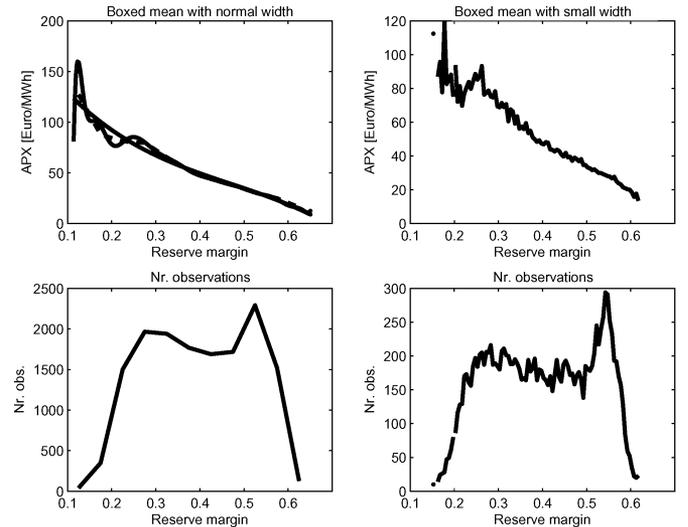


Fig. 4. Top left diagram shows the relation of APX price and reserve margin by means of both smoothed *b*-splines and a piecewise linear function. Top right diagram shows the relation by means a piecewise linear fit with a small width. Lower diagrams show the number of underlying observations in each interval for width 0.05 (left) and width 0.005 (right).

margins into several segments. We create intervals of width 0.05 and take the average of all spot prices within each interval. Reserve margin takes values between 0.10 and 0.70, leading to 12 intervals (0.10–0.15, 0.15–0.20, etc.). We denote an interval by its ending point (so the first interval is 0.15). To show the impact of the width of the interval we included a piecewise linear fit with an interval of width 0.005 in the upper right panel.

From the figure it is clear that odd humps can occur in the piecewise linear fit. A width of 0.05 leads to a monotonic decreasing fit for the existing data set. In contrast, a fit based on a width of 0.005 contains several wiggles. These wiggles contradict the economic intuition that prices should increase as reserve margins decrease.

Producing fitted prices that are monotonously decreasing in reserve margins was achieved by either adjusting the width in

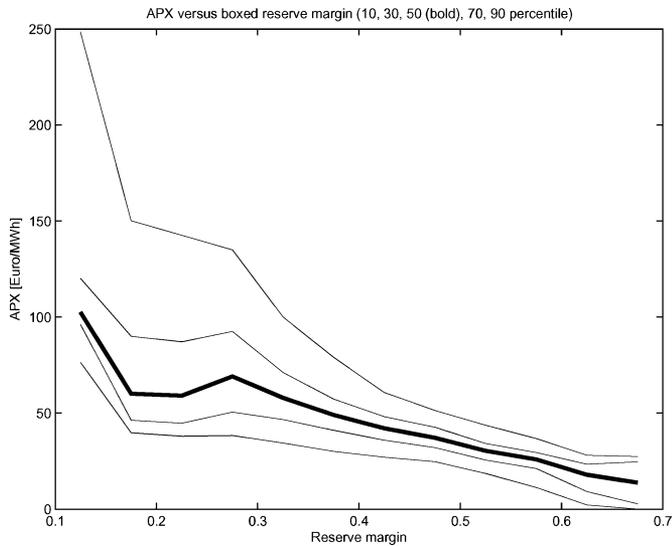


Fig. 5. Five percentiles of the relation between the APX price and the reserve margin (10, 30, 50, 70, and 90 percentile).

the averaging window or by using a non-parametric technique where monotony is built in. In Fig. 4 we show a  $b$ -spline (order 6, 15 evenly spaced knots) fit.<sup>3</sup> We subsequently smoothed the fit, which in our case gave a smooth and monotonic fit.

Two interesting points arise from this figure. In this example the piecewise linear fit and the  $b$ -spline fit are very similar. In addition, we see that the smooth  $b$ -spline fit appears rather linear for the reserve margin where most observations occur: values between 0.20 and 0.60. A similar observation was made by Visudhiphan and Ilic [10] on the NEPOOL market.

Our next question is how spread the real spot prices are around the average relation. In Fig. 5 we show the dispersion around the fit by plotting five different percentiles (10, 30, 50, 70, 90) of the relation between the APX and the reserve margin. From the figure we can conclude the price spread is decreasing with the reserve margin. Again we see a hump around index 0.30.

To gain additional insights, we compute the summary statistics for each of the discussed intervals and display the results in Table I. In this table we see that standard deviation, skewness and kurtosis rise for decreasing reserve margin if we disregard the first and the last interval. This is reasonable considering the limited number of data points in these intervals.

### E. Forecasting the Probability of a Spike

We define a spike as a price above 90 Euros, which is in line with market practice. We will take the threshold as given, and will not include it as one of the parameters to be estimated. A threshold of 90 Euros implies that about 11% of the data qualify as spikes. In Table II we give the percentage of the complete data sample that would qualify as a spike for some other threshold choices.

<sup>3</sup>The  $b$ -spline regression is performed with software provided by J. Ramsay on his website: <ftp://ego.psych.mcgill.ca/pub/ramsay/> which includes  $b$ -spline, smooth  $b$ -spline and smooth, monotonic  $b$ -spline. We used the fourth derivative, and set  $\lambda = 0.02$ .

TABLE I  
SUMMARY STATISTICS OF APX SPIKES FOR DIFFERENT RESERVE MARGIN INTERVALS: NUMBER OF OBSERVATIONS, MEAN, STANDARD DEVIATION, SKEWNESS, AND KURTOSIS

interval	nr obs	mean	std	skew	kurt
0.15	39	127.16	63.49	1.79	5.76
0.20	348	87.80	87.03	5.20	43.08
0.25	1501	81.15	73.06	5.84	56.84
0.30	1966	79.64	50.10	4.17	39.24
0.35	1940	65.42	36.67	4.25	37.94
0.40	1767	52.63	22.83	2.75	19.20
0.45	1689	43.97	16.44	2.28	13.85
0.50	1717	38.02	12.37	2.09	17.73
0.55	2291	30.39	9.98	0.42	5.50
0.60	1519	24.90	9.65	-0.18	3.10
0.65	125	16.10	9.57	-0.22	1.75
0.70	2	13.70	19.35	0.00	1.00

TABLE II  
EXCEEDING PROBABILITIES FOR DIFFERENT THRESHOLD LEVELS. IN THE FULL SAMPLE, THERE ARE 14 904 POINTS

APX threshold [EUR/MWh]	80	90	100	120	150
Exceeding probability	0.1410	0.1084	0.0760	0.0485	0.0214

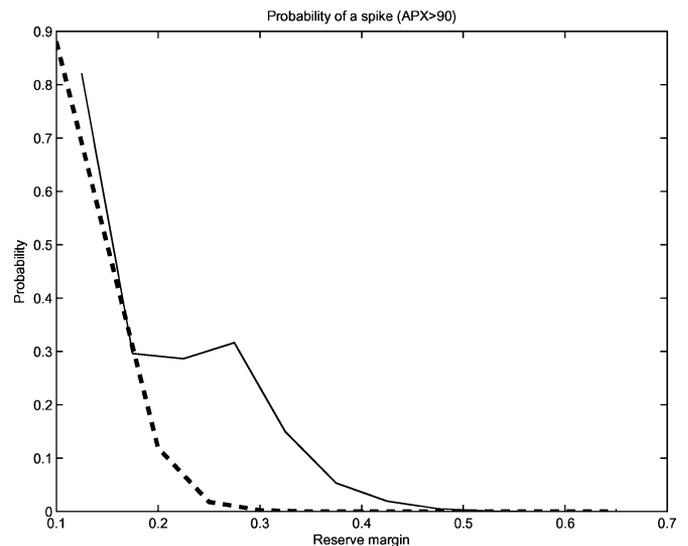


Fig. 6. Probability of a spike versus reserve margin within our data sample (solid line). The dashed line is proposed by Anderson [13] for the PJM market.

In Fig. 6 we show the relation between the probability of a spike and the reserve margin. For the probability of a spike we take the relative number of observations in a specific reserve margin interval above 90 Euros. Varying the threshold yielded similar graphs. The data do not seem to be in line with economic intuition: the probability is sometimes increasing in the reserve margin.

For comparison we include the spike probability function proposed by Anderson [13]. Our data already spike for higher reserve margins and the cut-off point is less clear than in the PJM data. As shown by Birnbaum *et al.* [11] and Mount *et al.* [15] for the PJM market and Ilic and Visudhiphan [10] for the NEPOOL

market, the probability of spike rises fast for reserve margins below 20%.

### F. Stability

The non-monotony of the relation between reserve margin and wholesale electricity price leads us to investigate the stability of the relation. In particular, we consider the time dependence on the daily and the yearly level, one-off events and out-of-sample performance.

1) *Causes of Instability:* The relation between reserve margin and spot prices combines the information from all the hours of the day irrespective whether it is day or night, week or weekend. Since we consider each hour individually, we cannot take the relation between adjacent hours explicitly into account. This information is especially relevant in situations where load is quickly changing or reaches its peak, as start-up costs start to play a role at those moments too. Because the start-up costs for the power plant to cover the maximum load need to be earned back, prices become temporarily higher than the levels we would normally establish. Also peak-load plants have a low number of running hours during which both the capital and operating costs need to be earned back.

Start-up costs also play a role in the weekends. In practice, an operator normally makes a decision whether to run or not to run a unit during the whole weekend. If an operator decides to run, he might be willing to sell under marginal costs (and hence at a loss) in order to avoid start-up costs. Because load can fluctuate a lot during the weekend, and with the load our reserve margin index, we can see different prices than what we expect from the normal price—reserve margin relation.

A similar must-run situation can occur in the winter with power plants that produce both heat and power. In order to cover the heat demand, the power plants become must-run in electricity. With the effect of must-run units known, one can make the hypothesis that hours which are covered with must-run units have a lower price.<sup>4</sup> To test this hypothesis, we need to be able to make the distinction between flexible and inflexible (must-run) units in the forecast of available supply. So far, this type of split is not available in the Dutch market. Therefore, we will test the implication of our hypothesis that, for the same level of reserve margins, wholesale power prices are lower in the weekends.

2) *Dependence on Time of Day:* In this and the next subsection we check whether the relation is similar across different subsets of the data. Trading of electricity normally takes place in fixed subsets of the day. Note that the definitions of such subsets (such as the exact hours to be called peak hours) differ across markets. Table III gives the definitions in the Dutch market and the number of observations in the data sample falling in each subset. A day (or baseload) can be split into a peak and an off-peak period. Peak hours can be split into a weekend-peak, a shoulder and a super-peak period.

In Fig. 7 we see that peak and off-peak prices are in line, though off-peak prices fall below peak prices for a reserve margin between 0.40 and 0.50. This is in line with our hypothesis. In Fig. 8 we consider the average relation for the three subsets of the peak: weekend peak, shoulder and super-peak.

<sup>4</sup>In the Netherlands negative prices can only occur in the real-time market. The minimum possible price at the day-ahead market is 0.01 Euro.

TABLE III  
DEFINITIONS OF DIFFERENT TIME-OF-DAY SUBSETS  
IN THE DUTCH MARKET AND THE NUMBER OF OBSERVATIONS

Products	Hours	Nr. Observations
Baseload	0-23	14904
Off-peak	0-6+23	4968
Peak	7-22	9936
Weekend-peak	7-22 (weekend only)	2800
Shoulder	7+20-22 (week only)	1784
Super-peak	8-19 (week only)	5352

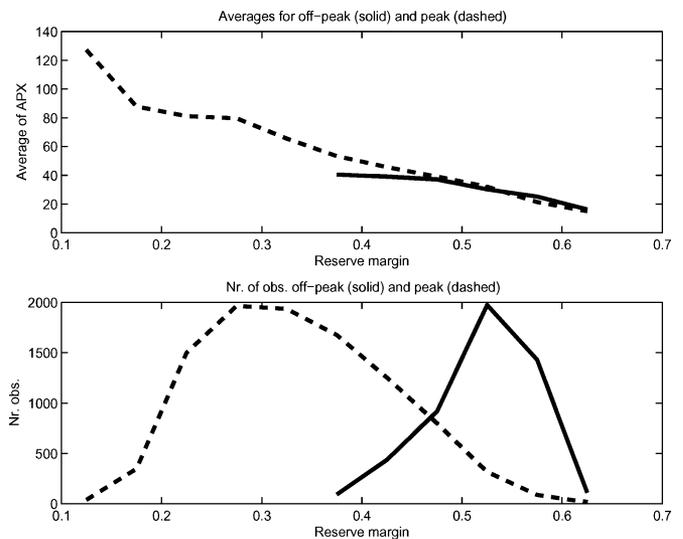


Fig. 7. Average relation during peak and off-peak (top) together with the number of observations (bottom) in each interval. Values based on less than ten observations are omitted.

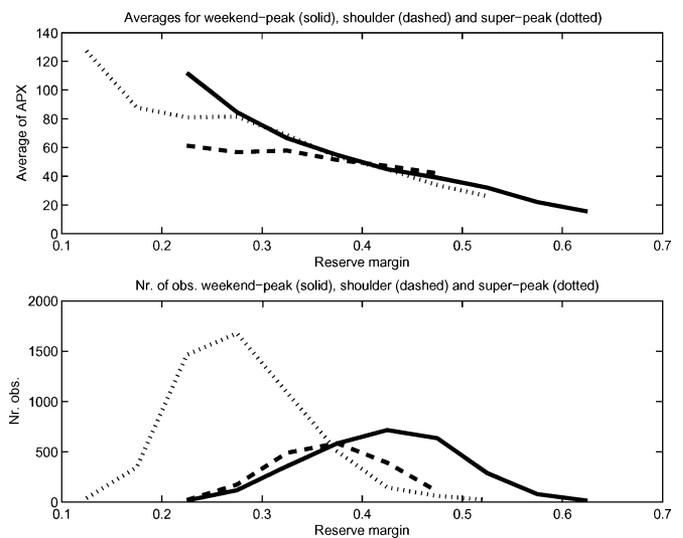


Fig. 8. Average relation during weekend peak (solid), shoulder (dashed), and super-peak (dotted) in the top panel together with the number of underlying observations in each interval in the bottom panel. Values based on less than ten observations are omitted.

In the figure, we see that, keeping reserve margin constant, the power price tends to be lower during shoulder hours than super-peak or weekend-peak hours. In addition, the weekend-peak prices rise above super-peak prices for a reserve margin below

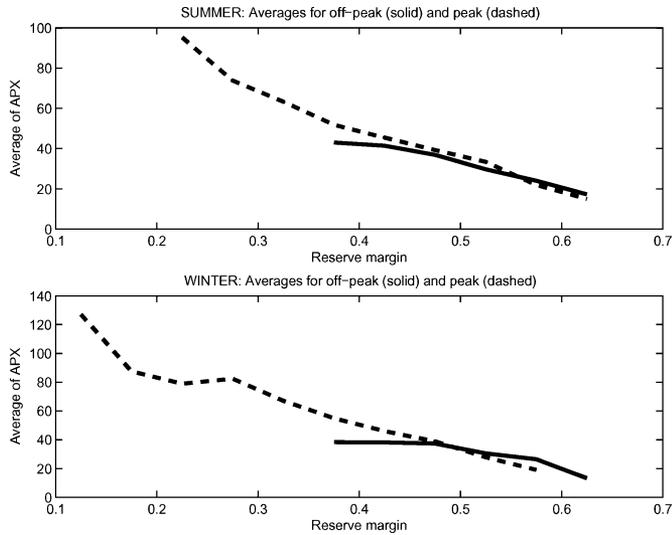


Fig. 9. Average relation during peak (dashed) and off-peak (solid) for summer (top) and winter (bottom).

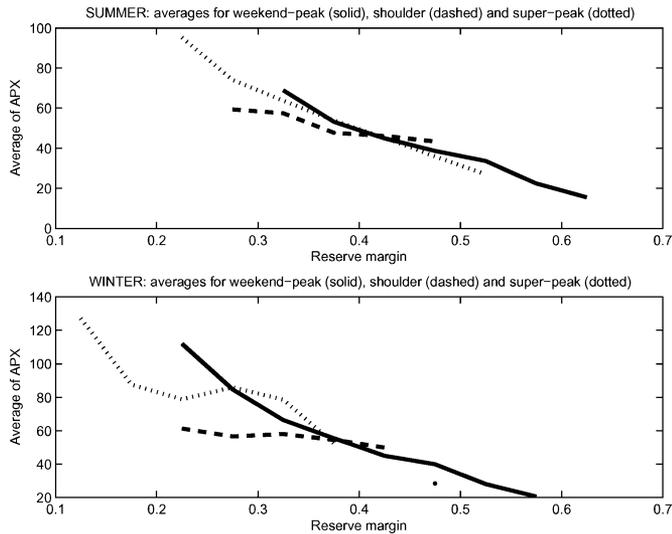


Fig. 10. Average relation during weekend peak (solid), shoulder (dashed), and super-peak (dotted) for Summer (top) and Winter (bottom).

0.30. This appears out of line with our hypothesis. In the next section we investigate seasonal effects.

3) *Dependence on Season:* We divided the data into Summer (April–September) and Winter months (October–March). In our data set of 14 904 data points, we have 8640 data points in the Winter and 6264 in the Summer. In Fig. 9 we see that the difference between peak and off-peak is sustained if we split the data into Summer and Winter. The same conclusion holds for the relation between weekend-peak, shoulder and super-peak hours as can be seen in Fig. 10. This brings us to the conclusion it is better to specify two separate models: one for week and one for weekend days.

4) *Outliers:* In the previous section, we have seen that certain data points were not in line with the average relation. For this reason we consider how these “outlier” data are distributed over the data sample. We decided to look at data points with high APX prices and relatively high values of reserve margin.

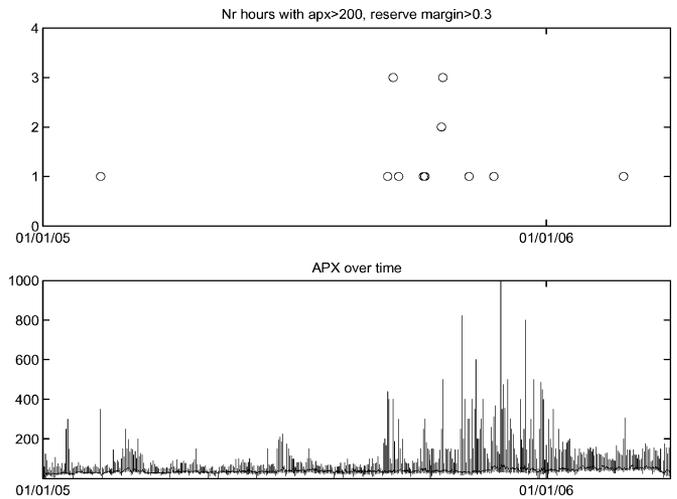


Fig. 11. Top panel shows how many hours on a specific day had both a high APX price ( $>200$ ) and high reserve margin ( $>0.3$ ). The bottom panel shows APX prices, 01.10.2004-17.06.2006.

For example, if we use the (arbitrary) definition of an outlier as  $APX > 200$  and reserve margin  $> 0.3$ , we call 16 data points outliers while in total 128 data points had  $APX > 200$ . In Fig. 11 we graph how many of such data points were clustered in one day over time. We see there are two days with three outliers (that is: on two days there were three hours which are rather out of line with the usual behavior), and that outliers are mainly present around October 2005. An explanation for these outliers could be the low nuclear availability in France during that time, which impacted the different surrounding electricity markets. For comparison, we included again the development of APX prices over time. We conclude that the hump was not due to odd data in the beginning of the data sample. Hence we exclude the possibility that the hump reflects market inefficiencies at the beginning of the sample.

5) *Out-of-Sample:* Until now, we have used the whole sample to draw conclusions about the relation between the reserve margin and the spot prices. In this subsection, we give a first indication of how stable the relation is over different subsamples. In other words: are there different pricing regimes over time? An alternative to our approach could be a regime discovery algorithm proposed by Vucetic *et al.* [27].

For the stability check we divide our sample in three parts: the first 5000, the second 5000 and the remaining 4904 observations. In Fig. 12 we compare the average price and the probability of spike. We see that the average price has increased, and that the relation from the first period understates the average price and probability of spike in the second period. The relation describing the second period is similar to the relation describing the third period. This would show a good out-of-sample behavior. Part of the increase has been due to an increase in marginal costs. This has not been captured by our current definition of a spike.

One way to incorporate the increase of prices would be to include an average error over a certain period. Then, the natural question is how many data points we should use in our data estimation by comparing the errors out-of-sample. We will not address this question in this paper.

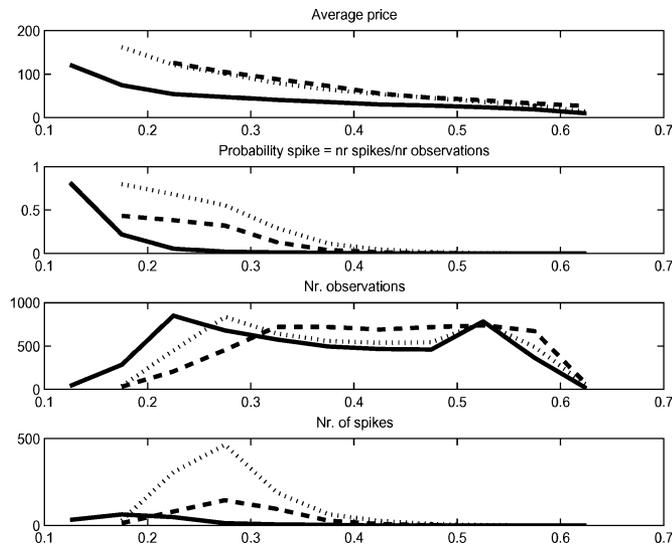


Fig. 12. For three different time periods (period 1 starting 01/10/2004: solid; period 2 starting 29/04/2005: dashed; period 3 starting 24/11/2005: dotted), we show the average price (top), the probability of spike (upper middle), the number of observations (lower middle), and the number of spikes (bottom).

## V. DISCUSSION AND CONCLUSION

In this paper we have shown how to create an estimate for the supply-demand framework and how to build a simple model for it. One of our main findings is that reserve margin should be included into a spot electricity model to enhance performance. Another useful area of application is the development of fundamental models. While most fundamental models can create estimates of future marginal costs, in practice a link from marginal costs to market prices is needed. Our model can provide such a link if marginal costs are driven by the reserve margin.

Our procedure is quite simple and could be applied to other markets where estimates of available capacity are published, e.g., in the U.K. or Germany. The U.K. market is the most similar to the Dutch market. For the German market, it is important to study carefully the role of interconnectors and wind production.

Our piecewise linear fit can contain a double-hump structure which does not comply with standard economic theory. Also our results imply the Dutch market can spike for moderate levels of the reserve margin. A step forward could be made if we could group available capacity by technology, but this information is not available.

The backbone of our relation is an assumption of stability. We have studied the stability of our relation over different times of day and seasons. We found it is better to specify a separate model for different times of day, where it is especially worthwhile to split week and weekend days. Out-of-sample tests gave promising initial results. To improve the relation, one could include an autoregressive part for the error in the spot price prediction for the previous period. This is left for future research.

Our model can be extended in different directions. One of the possible directions is the relation between spot and forward. With a stability assumption it is possible to simulate different underlying drivers and create a simulation of future spot prices. Another direction is the extension to a coupled market. This

type of markets are present in the U.S. and in Nordpool. Market integration of the Dutch, Belgian and French market [25] will provide a new challenge for the presented model.

## ACKNOWLEDGMENT

The authors would like to thank the short-term electricity and risk management department at Essent Energy Trading, A. Musilová, Á. Cartea, R. Brummelhuis, F. Yang, and participants in the Modelling and Measuring Energy Risk 2006 and Commodities 2007 Conference. The authors would also like to thank TenneT for providing historical data on the maximum possible import and export.

## REFERENCES

- [1] A. Cartea and M. G. Figueroa, "Pricing in electricity markets: A mean-reverting jump diffusion model with seasonality," *Appl. Math. Fin.*, vol. 12, no. 4, pp. 313–335, Dec. 2005.
- [2] R. Huisman and R. Mahieu, "Regime jumps in electricity prices," *Energy Econ.*, vol. 25, no. 5, pp. 425–434, Sep. 2003.
- [3] R. Koseckia, "Fuel-Based Power Price Modelling," in *Energy Modelling & The Management of Uncertainty*. London, U.K.: Risk Books, 1999.
- [4] A. Eydeland and H. Geman, "Pricing power derivatives," *Risk*, pp. 71–73, Oct. 1998.
- [5] A. Eydeland and H. Geman, "Fundamentals of electricity derivatives," in *Energy Modelling & the Management of Uncertainty*. London, U.K.: Risk, 1999.
- [6] C. Pirrong and M. Jermakayan, "Valuing power and weather derivatives on a mesh using finite difference methods," in *Energy Modelling & The Management of Uncertainty*. London, U.K.: Risk, 1999.
- [7] C. Pirrong and M. Jermakayan, "The Price of Power: The Valuation of Power and Weather Derivatives," working paper, Olin School of Business, Washington Univ., 2000.
- [8] P. Skantze, A. Gubina, and M. Ilic, "Bid-based Stochastic Model for Electricity prices: The Impact of Fundamental Drivers on the Market Dynamics," MIT Rep. EL00-004, 2000.
- [9] A. Eydeland and K. Wolyniec, *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. New York: Wiley Finance, 2003.
- [10] P. Visudhiphan and M. Ilic, "Dependence of generation market power on the demand/supply ratio: Analysis and modeling," in *Proc. IEEE Power Eng. Soc. Winter Meeting*, 2000, vol. 2, pp. 1115–1122.
- [11] L. Birnbaum, J. M. del Aguila, G. D. Orive, and P. Lekander, "Why electricity markets go haywire," *McKinsey Quart.*, vol. 1, pp. 64–73, Aug. 2002.
- [12] M. Davison, C. L. Anderson, B. Marcus, and K. Anderson, "Development of a hybrid model for electrical power spot prices," *IEEE Trans. Power Syst.*, vol. 17, no. 2, pp. 257–264, May 2002.
- [13] C. L. Anderson, "A hybrid model for electricity spot prices," Ph.D. dissertation, Grad. Depart. Appl. Math., Univ. Western Ontario, London, ON, Canada, 2004.
- [14] M. Burger, B. Klar, A. Muller, and G. Schindlmayer, "A spot market model for pricing derivatives in electricity markets," *J. Quant. Fin.*, vol. 4, pp. 109–122, Feb. 2004.
- [15] T. D. Mount, Y. Ning, and X. Cai, "Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters," *Energy Econ.*, vol. 28, no. 1, pp. 62–80, 2006.
- [16] H. Zareipour, C. A. Canizares, K. Bhattacharya, and J. Thomson, "Application of public-domain market information to forecast Ontario's wholesale electricity prices," *IEEE Trans. Power Syst.*, vol. 21, no. 4, pp. 1707–1717, Nov. 2006.
- [17] C. Fezzi and D. Bunn, "Structural Analysis of High-Frequency Electricity Demand and Supply Interactions," working paper, London Business School, 2006.
- [18] W. R. Hughes and A. Parece, "The economics of price spikes in deregulated power market," *Elect. J.*, pp. 31–44, Jul. 2002.
- [19] X. Lu, Z. Y. Dong, and X. Li, "Electricity market price spike forecast with data mining techniques," *Int. J. Elect. Power Syst. Res.*, vol. 73, pp. 19–29, Jan. 2005.
- [20] A. Misiorek, S. Truck, and R. Weron, "Point and interval forecasting of spot electricity prices: Linear vs. non-linear time series models," *Stud. Nonlin. Dynam. Econometr.*, vol. 10, no. 3, Sep. 2006, article 2.

- [21] M. Barlow, "A diffusion model for electricity prices," *Math. Fin.*, vol. 12, pp. 287–298, Oct. 2002.
- [22] H. Bessembinder and M. L. Lemmon, "Equilibrium pricing and optimal hedging in electricity forward markets," *J. Fin.*, vol. 57, pp. 1347–1382, Jun. 2002.
- [23] P. Villaplana, Valuation of Electricity Forward Contracts: The Role of Demand and Capacity, working paper, Univ. Pompeu Fabra, 2005.
- [24] T. Kanamura and K. Ohashi, "A structural model for electricity prices with spikes: Measurement of spike risk and optimal policies for hydropower plant operation," *Energy Econ.*, vol. 29, no. 5, pp. 1010–1032, 2007.
- [25] Belpex, Nov. 2006, Market Coupling. [Online]. Available: <http://www.belpex.be/index.php?id=4>.
- [26] A. F. Boogert and D. Y. Dupont, "The nature of supply side effects on electricity prices: The impact of water temperature," *Econ. Lett.*, vol. 88, no. 1, pp. 121–125, Jul. 2005.
- [27] S. Vucetic, K. Tomsovic, and Z. Obradovic, "Discovering price-load relationships in California's electricity market," *IEEE Trans. Power Syst.*, vol. 16, no. 2, pp. 280–286, May 2001.

**Alexander Boogert** (S'06) received the M.Sc. degree in 2003 from the University of Twente, Enschede, The Netherlands. He is pursuing the Ph.D. degree in the Commodities Finance Centre at Birkbeck College, the University of London.

He works as a quantitative analyst at Essent Energy Trading, Den Bosch, The Netherlands. In this role, he provides quantitative support to the Risk Management Department while also performing long-term research.

**Dominique Dupont** received the Ph.D. degree in economics in 1995 from the University of Chicago, Chicago, IL.

Before joining the University of Twente, Enschede, The Netherlands, in 2002, he worked at the Federal Reserve Board in Washington DC. His current research centers on modeling derivatives in energy, in particular electricity. In 2007, he joined RWE Trading's Risk Methodology and Policies Department.