

When supply meets demand: the case of hourly spot electricity prices

Alexander Boogert* and Dominique Dupont†

October 31, 2006

Abstract

We use a supply-demand framework to model the hourly day-ahead spot price of electricity based on publicly available information. With the model we can forecast the level and the probability of a spike in the spot price defined as the spot price being above a certain threshold. Several European countries have recently started publishing day-ahead forecasts of the available supply. In this paper we show potential uses of such indicators and test their forecasting power in an hourly spot price model. We conclude that a forecast of the available supply can be part of a useful indicator and discuss ways to further improve the forecasts.

1 Introduction

Day-ahead spot electricity prices provide an important reference point to all members of the electricity industry. These prices are characterized by high volatility and rare but violent spikes. These aspects have motivated significant research efforts. In this paper we model the spot electricity price based on the supply-demand equilibrium.

There are several ways spot electricity models can be applied. In short-term tactical planning it is important to forecast the absolute height of the day-ahead spot prices and to forecast the probability of a spike. On a long-term basis the variability of spot prices becomes interesting as well. This variability can be used as an input for the long-term planning of powerplants. In the current article we focus mainly on the short-term horizon.

We focus on the relation between several fundamental drivers and hourly spot electricity prices. Using hourly prices instead of daily prices increases the sample size and hence the likelihood of obtaining robust empirical results. Accurate forecast of demand and supply is of paramount importance to the

*Birkbeck College, University of London, Commodities Finance Centre & Essent Energy Trading, 's-Hertogenbosch, The Netherlands.

†Corresponding author. FELab, University of Twente, Capitoool, P.O. Box 217, 7500 AE Enschede, The Netherlands.

electricity industry because these two must be balanced at any time to maintain the stability of the power grid. Forward electricity contracts are traded several years before actual delivery. Contracts are traded both on the OTC market and on organized exchanges and delivery is normally channelled through a day-ahead market. The market design differs between electricity markets. Examples of design differences include the exact time of settlement, the granularity of the contracts (i.e, the time period for which power is to be delivered), the handling of actual delivery in real-time and the exact information provided to the public. In general day-ahead markets take the form of an auction. First, the independent auctioneer aggregates buy and sell orders from the different market participants for electricity to be delivered the following day for each individual hour. Then it computes 24 market-clearing prices for the next day, which are the day-ahead prices (or spot prices) we discuss in this article.

Some of these market design differences, like the exact time of settlement and granularity of the market, do not have a great impact. Markets settle at different moments in the morning creating small differences in forecasting. Most markets operate at an hourly granularity, while the UK and Australia operate at an half-hourly granularity.

On the other hand, it is worth noting that the relation between day-ahead and real-time markets can be very different dependent on the market studied. Longstaff & Wang (2004) find in the PJM market in the United States that power prices on the day-ahead market are on average higher than on the real-time market and relate this spread to several risk factors. Karakatsani & Bunn (2005) find that in the UK market the difference shows a diurnal pattern. Boogert & Dupont (2005a) find that in the Dutch market differences are rarely positive on average and are always characterized by very large potential losses. Other differences occur in how much information is shown to the public, and of what quality these numbers are. For example, is there a day-ahead estimate of available supply, and do these numbers cover the full market? Is there a day-ahead forecast for the national load?

This creates a situation in which the general structure of the models can span several markets, but where local adjustments are needed to make a useful model in a particular market. In this article we try to discuss both aspects while creating an hourly spot electricity price model for the Dutch market.

The remainder of this article is as follows. In section 2 we give review the literature. In section 3 we establish the supply-demand framework for a general electricity market. First we discuss factors which influence spot electricity prices. Then we introduce a first relation between supply demand and spot prices and show how it can be used for forecasting hourly spot prices and probability of a spike. Subsequently, we contrast our non-parametric approach with some parametric ones. In section 4 we introduce the situation in the Dutch market. We specify which data is available and apply the techniques for forecasting an hourly spot price and the probability of spike. The section ends with a study on the stability of the relation. In section 5 we discuss the implications for further modelling.

2 Review of the literature

The modelling of electricity spot prices has long focused on the reduced form models (e.g. Deng (1999), Cartea & Figueroa (2005), Huisman & Mahieu (2003).) Two popular modelling approaches are jump diffusion and regime switching. Both of these approaches are mathematically tractable and have received considerable attention, especially in order to perform good parameter estimation. Another route is provided by fundamental models (e.g. Kosecki (1999)), which carefully describe the characteristics of the supply stack in a market. In monopolistic markets the full supply stack was known and used to serve the load at the lowest cost. In liberalized markets only the general shape of the daily supply stack is known. The marginal cost curves one obtains with fundamental modelling and estimates of the supply stack need to be transformed into spot prices. How to transform these marginal costs into a market price in a liberalized market setting is not straightforward.

A hybrid model incorporates ideas from the two approaches. Compared to reduced-form models, hybrid models take into account useful additional information besides the price time series like for example weather or availability of power plants. Eydeland and Geman (1998), Eydeland and Geman (1999), Pirrong & Jermakayan (1999), Pirrong & Jermakayan (2000), Skantze, Gubina & Ilic (2000) and Eydeland & Wolyniec (2003) are examples of a class of hybrid models based on the assumption that there is an exponential relation between price and load. This captures the behavior of strongly increasing prices when the load is growing, while it can facilitate closed form solutions for the pricing of electricity derivatives.

The underlying assumption for these models is that there is a clear relation between price and volume on the day-ahead market. This makes them useful in so-called pool markets where all supply has to be offered in the day-ahead market (for example in the old NETA system in England and Wales or currently in Spain.) In other markets a day-ahead forecast for the national load needs to be created. For example, there is no apparent relation between price and volume on the Dutch market (APX) as can be seen in figure 1. The APX represents only about 20 percent of the total national load.

One main ingredient for our model is the reserve margin which covers the fraction of the supply which is still available for covering the demand. Ingredients for our approach have appeared in different forms. Mount et al. (2006) created a regime switching model where the switching probabilities between the regimes and the conditional means for each regime vary with time and with reserve margin. Anderson (2004) prescribes a functional form (unfortunately non-motivated) for the relation between a type of reserve margin and the probability of a spike. One difference is that our transformation is non-parametric. Burger et al. (2004) prescribes a functional form for the relation between an index related to the reserve margin and the spot price together with residual short-term fluctuations and long-term variation of prices. The index incorporates the expected relative availability of power plants and load, though the precise form is not given. The difference with our article is that Burger et al.




Figure 1: APX price versus APX volume

more focus on the long horizon.

Spot electricity price modelling is not the only area where the reserve margin is studied. Actually, the idea seems to stem from the public policy research on market power (e.g. Visudhiphan & Ilic (2000)) and security of supply (e.g. Birnbaum (2002)). The difference with our work is their focus on assessing the market mechanism, and thus on explaining the relation. Our focus is on the description and simulation of prices, and thus on exploiting the relation.

3 The supply-demand framework

In this section we discuss factors which influence spot electricity prices. Besides past spot electricity prices, there is a range of factors which could impact the analysis. As one of our goals is to investigate the forecast of the available capacity, we dedicate the first subsection to this topic. Subsequent we discuss additional price drivers.

3.1 Forecasting available capacity

It is important to distinguish between supply and supply curve. Accurate information is readily available on the amount and the price of power traded on the market in the past. Naturally, by construction, supply equaled demand at those prices. More interesting for forecasting purposes would be information about the supply curve around past equilibria, or, equivalently, information on

the elasticity of supply. The elasticity of supply is determined by spare capacity available. In turn, this capacity crucially depends on the granularity of the market. In the very short run (e.g. in a 15 minutes) some flexible units can be turned on and the output of running units can be increased. On a day-ahead basis one can start up units, though there are technical restrictions on production schedules.

Demand elasticity is normally not taken into account as consumers are generally price insensitive. However, there are groups of large customers (e.g. in the metals industry) for whom it is possible to temporarily shut down. These customers are willing to reduce their demand in exchange for a reduction of the electricity price. In financial terms: the customers buy an interruptible contract. These demand elasticity effects are explicitly taken into account in Fezzi & Bunn (2006). In this article we refrain from this effect, and assume it will enter the model implicitly.

Within the existing literature most articles do not explicitly introduce supply because of an apparent 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, UK and Germany. In this article we will focus on the Dutch market, which has the longest history of these three. In the US, for the PJM market information is published on available capacity with a delay of 6 months (see Mount et al. (2006)).

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. We think this is dependent on the market design and on actual market behavior. An important difference is created by the timing of the import/export capacity market in comparison to the day-ahead markets, and the timing between day-ahead markets in different markets. In case the capacity auction is in advance of the day-ahead market, the resulting price is an indicator for the day-ahead spot price. To be more precise: it is an indication for the upcoming spread between two different day-ahead spot electricity markets.

3.2 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 Hughes & Parece (2002). 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 organization and design (retail price caps, revenue share of spot sales, capacity requirements and wholesale price caps) as possible price drivers.

In this paper, we focus on finding a simple relation between price and available capacity. Alternatively, one could apply data mining techniques like neural

networks and adaptive splines (see e.g. Lu et al. (2005) for reference).

3.3 A relation between supply, demand and spot price

One of the goals of this article is to understand the relation between supply, demand and the spot price. Inspired by Anderson (2004) we have decided to consider a relative demand-supply ratio (*RDS*) of the following kind:

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

This index is a reserve margin index: it covers the fraction of the supply which is available for covering demand. It is realistic to believe this index will have a negative correlation to the spot prices. We note that the lower the index, the less capacity is available and the tighter the market. This will then imply that more expensive units are coming online, and marginal costs increase. In addition, we expect the bandwidth around the average to increase for lower indices. Note that this index is closely related to the concept of capacity utilization as used by Anderson (2004). Capacity utilization states how much of the available capacity is used to cover the demand (that is: capacity utilization equals 1 - reserve margin.)

There are two natural candidates which could provide an alternative to our index. Instead of a relative demand-supply ratio, we could opt for a relative supply-demand ratio (*RSD*) or a difference between the absolute supply and demand (*ADS*). As mentioned by Visudhipan & Ilic (2000), the supply-demand ratio is more sensitive to variation in supply than the demand-supply ratio, while the demand-supply ratio is more sensitive to variation in demand than the supply-demand ratio.

$$\text{RSD} := \frac{\text{available capacity}}{\text{demand}} - 1 \quad (2)$$

$$\text{ADS} := \text{available capacity} - \text{demand} \quad (3)$$

3.4 Forecasting hourly spot prices

One way to forecast the spot price is to consider the average transformation from reserve margin to hourly spot prices. Simultaneously this allows us to produce a confidence interval around the average. From an economic viewpoint we expect this transformation to increase for a decreasing reserve margin. As well, we expect the bandwidth to increase for decreasing reserve margin. In our data part, we have used both a smoothed b-spline fit and a piecewise linear fit.

The natural extension is to consider a two-dimensional version of this approach. Such a step was taken in Lu et al. (2005) where besides a reserve margin a steepness-of-load indicator was used.

3.5 Forecasting the probability of a spike

Besides the absolute height of the prices, an important variable for market participants is the probability that the prices will end up above a certain threshold. We will call a price above the threshold a spike. In this article 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 in the market at the time.

We estimate the probability as the relative number of observations above the threshold in our data sample. We vary the threshold to avoid our results are dependent on the specific threshold. At this moment it is interesting to relate our work to Anderson (2004). She put a parametric transformation from reserve margin to spike probability at a central point in the model and kept the transformation fixed. This motivated us to study the stability of the transformation in the data part. In the following paragraph we discuss more parametric models to contrast such an approach to our non-parametric approach.

3.6 Parametric approach

In our approach we assume there is a non-linear relation between the reserve margin and spot prices. Another approach is to parametrize the relation. In a functional form, we can rewrite our reserve margin index as follows:

$$S_t = f\left(1 - \frac{D_t}{C_t}\right) \quad (4)$$

where S_t is an hourly spot price, D_t is the demand, C_t is the available capacity and f is a non-linear function.

The variable $1 - \frac{D_t}{C_t}$ 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 appears (the function is given without explicit motivation) to be the motivation behind Anderson (2004). Similarly, Barlow (2002) uses a displaced diffusion model, where the power price is 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 toward $+\infty$ in the neighborhood of the singularity. The inverse cdf technique can be seen as refinement of this technique. Such approaches are elegant, but lack flexibility. We focus on non-parametric methods instead.

Another route is obtained if we treat supply and demand separately, and introduce a functional form for the relation. This functional form allows an explicit link between spot and forward prices, and creates a possibility to study the forward risk premium. Bessembinder and Lemmon (2002) formulated a general equilibrium model for the day-ahead forward prices, which they applied to the PJM market. In their model speculators cannot participate and supply is not a random variable. Saravia (2004) extended the model and studied the effect of speculators on the relation between the day-ahead and real-time market

in the New York electricity market. Villaplana (2005) extended the model by considering supply as a random variable and applied the model to the Nordpool market.

In these models the relation is assumed to be of exponential or power form. This creates the possibility to estimate the parameters by a linear equation, and allow for closed form solutions for forward prices. Simultaneously, it captures the empirical phenomenon that prices rise for increasing demand and decreasing capacity. To give an example, Villaplana (2005) estimates

$$S_t = \gamma_1 C_t^{\gamma_2} e^{\gamma_3 D_t} \quad (5)$$

4 Application to the Dutch market

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

4.1 Overview of the Dutch market

The Netherlands was among the first countries in the European Union to liberalize its electricity market. The Dutch ISO, 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 wholly 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 in normal circumstances is 3650 MW, which can be enlarged to 3850 MW in case of emergencies. The scheduled day-ahead import is not exactly realized in real-time. Although the electricity traded on the APX represents about 20 percent of the Dutch daily consumption, the APX is considered an important benchmark.

In the Dutch market import/export capacity is auctioned before the day-ahead spot electricity and imported electricity has to be offered on the day-ahead market. It is worth noting, that there is a need for import to the Dutch market and that the available import capacity is used frequently. This makes it possible to treat the import/export mainly as import.

A new development (starting 21st November 2006) is the introduction of market coupling between the Netherlands, Belgium and France (Belpex (2006)). Under this new system the import/export auction will be integrated into the day-ahead auction. This means that the current explicit auctions, will be transformed into implicit auctions. Similar types of markets are present in the US and in Nordpool.

4.2 Available data in the Netherlands

Before 2004 no public estimates were made of the available supply in normal circumstances. Only when available capacity dropped below a lower threshold the public was informed about the state of the system. As this happened rarely, it was difficult to estimate the available capacity and the demand-supply equilibrium in general. In a previous article, Boogert & Dupont (2005b), we have shown a good indicator for the spot prices in the period 2000-2003 was the water temperature: hot water reduces potential capacity and hot water occurs when temperatures are high leading to high demand.

Since 2004, TenneT publishes an estimate for the available capacity for the coming 30 days in the Dutch grid. TenneT gathers statements of the different generators about the availability of their individual plants, and summarizes them on an aggregate level.

The problems with the TenneT estimate are twofold. On the one hand there is no reliability check on the provided data. The generators are not checked whether the indicated available capacity is indeed available and there exists no penalty in case of bad performance. On the other hand the estimated available capacity does not cover the total available capacity. The estimate contains neither potential wind production nor generation in smaller units (< 10 MW). We decided to use this estimate as it gives a unique opportunity to estimate a supply-demand framework.

The TenneT estimate is one way to describe the supply-demand equilibrium in the Netherlands. On top of that estimate (which we will denote by *TAC*) we think the following types of data could be related to the supply-demand framework:

- IMEX Realized import or export: history published by TenneT on 15 minute basis. We take import as a positive number since it adds to the available capacity. This information is published with a delay of 30 minutes.
- INT Maximum import: the maximum possible import and export is published by TenneT. A day-ahead prediction is available, together with announcements for future maintenance and enlargement in case of emergency. We received a historical database from TenneT.
- NL National load: realized generation including realized net import gives the load which is published by TenneT on 15 minute basis with a delay of two days. There is no official prediction available, though Essent Energy Trading can provide an internal one. Note the official load data covers only electricity generated by units larger than 10 MW.
- IMB Real time imbalance prices and volumes: TenneT publishes the real-time imbalance prices and volumes on their website.
- RP Regulating power: TenneT publishes the prices for up and downward regulation for the coming day. This could provide an additional tweak to the available capacity.

WP Wind power: there is no official estimate for the total wind production in the Netherlands. An internal estimate was provided by Essent Energy Trading, but unfortunately the starting date of this series is 01/01/2005.

In this report we are working with data starting 01/10/2004 and ending 17/06/2006.¹ The starting date coincides with the first publishing date of the forecast for the available capacity. As spot prices are published on an hourly scale, we transformed all 15 minute data into hourly data by taking the average over that specific hour. Subsequent graphs all show hourly data. In total there are 14904 hourly data points.

As we mentioned in section 3.1, the definition of total available capacity is not always clear. We note that in the Dutch market the first choice concerns inclusion of realized or day-ahead maximum import or export. The second choice to be made is whether we include wind power or not. This gives us four different specifications of available capacity.

The potential for wind energy is growing in the Netherlands. In 2004 the total installed capacity was 1073 MW, which grew to 1224 MW in 2005 (CBS, 2005). Given its size, it could be interesting to include wind power into the available capacity. However, as the data is not public, we have for the current version decided to exclude wind power from the available capacity. Concerning the import/export number, we have chosen to use the day-ahead forecast of maximum possible import/export. Real-time flows are not available day-ahead and flows appear more a resultant of our model. Thus we work with the following estimate for the total available capacity C_t :

$$C_t = TAC_t + INT_t \quad (6)$$

Comparing this to the situation studied by Mount et al. (2006), we see the ISO in the PJM market had more information than the average market participant for which the models perform less good. 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.

4.3 Reserve margins in the Netherlands

Let us start with showing the development over time of the underlying data for the reserve margin. In figure 2 we show the load, forecasted available capacity and day-ahead maximum possible import capacity. To contrast we also included the realized values of the import into the bottom panel. In the figure we see load changes between 700 and 1600 MWh, while the forecast of the available capacity moves between 1200 and 1700 MW. The realized import varies significantly, while the maximum possible import capacity.

In figure 3 we show the development over time of both the APX price and the reserve margin. In that figure the relation does not seem to be so strong, but a

¹For 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.




Figure 2: Load, available capacity and day-ahead maximum possible import capacity. In the bottom panel we include realized net import during our data sample

scatter plot in figure 4 reveals a pattern of increasing prices with reserve margin. In the same figure we plotted scatter plots between APX and load respectively available capacity. We see the relation between APX and load appears stronger than the relation between APX and available capacity. It is good to note that there are some high prices for still medium index values. We will come back to this point in section 4.6.

4.4 Forecasting the spot price

Given the reserve margin, one way to forecast the spot price is to consider the average transformation from reserve margin to APX prices. In figure 5 we show a b-spline fit and a piecewise linear fit. The piecewise linear fit was created by a discretization of the reserve margin. We create intervals of width 0.05, and take the average of all spot prices within each interval. In the reserve margin, values were taken 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) This is a data driven approach which does not comply with economical sense that prices should reduce for increasing reserve margin. For comparison reasons we introduced a smoothed, monotonic fit as well using smoothed b-splines.² Looking at the smoothed curve, we see that the b-spline fit appears rather linear most observations occur: values between 0.20 and 0.60. A similar observation

²The b-spline regression is performed with software provided by Jim Ramsay on his website: <ftp://ego.psych.mcgill.ca/pub/ramsay/>




Figure 3: APX price and reserve margin during our data sample

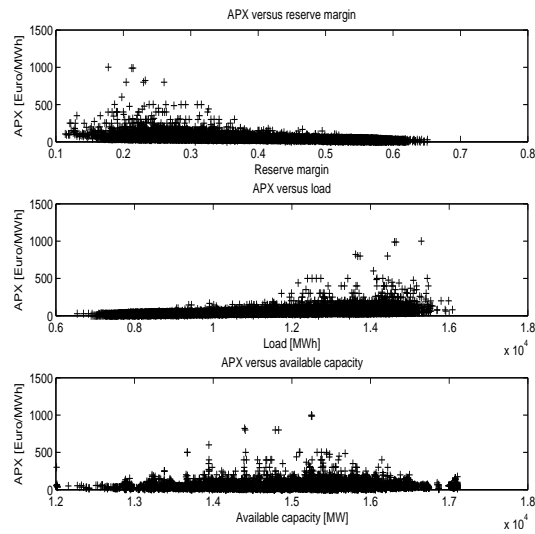


Figure 4: Scatter plots of APX versus respectively load, available capacity and reserve margin

was made by Visudhiphan and Ilic (2000) in the NEPOOL market. Another alternative would be to use outlier detection. If certain points are too unrealistic to fall within our economic theory, we exclude them from the sample. However, we decided to avoid the discussion about what is realistic.




Figure 5: Interpolated graph of the APX-index relation via spline and piecewise linear fit. The lower diagram shows the number of underlying observations in each box.

The next question is how spread the real spot prices are around the average transformation. In figure 6 we show the variability around the fit by the 5 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 local hump around index 0.30.

To get a better feeling for the distributions, we give the summary statistics in table 1. In the table we see standard deviation, skewness and kurtosis are rising for decreasing reserve margin.

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

Table 1: Summary statistics for different intervals: mean, standard deviation, skewness and kurtosis




Figure 6: The top diagram shows the average transformation occurring for the different intervals together with the 10, 30, 70 and 90 percentile.

4.5 Forecasting the probability of a spike

In this article we define a spike as a price above 90 euro. This threshold was originally established by traders as a reasonable benchmark for which they would like to receive a warning. In this article we take the threshold as given, and do not include it as one of the parameters to be estimated. A threshold of 90 euro implies that about 11% of the data sample is qualified as a spike. In table 2 we give the percentage of the data sample that would be qualified as a spike for some other threshold choices.

APX threshold	80	90	100	120	150	200
Exceeding probability	0.1410	0.1084	0.0760	0.0485	0.0214	0.0086

Table 2: Different exceeding probabilities for different threshold levels. In the full sample there are 14904 points.

In figure 7 we show the relation between the probability of a spike and the reserve margin. We take as the probability of a spike, the relative number of observations in a specific reserve margin interval above 90 euro. By varying the threshold we found similar graphs as the ones presented.

Again, we see our data does not follow the economic theory that the probability should increase for decreasing reserve margin. For comparison reasons we included the spike probability function indicated by Anderson (2004). Although




Figure 7: Probability of a spike versus reserve margin. The dashed graph is proposed by Anderson (2004) for the PJM market.

estimated in another market, we see that our data spikes earlier and that the cut off point is not as clear as in the PJM data.

The same result in the PJM market has been established by Birnbaum (2002) and Mount et al. (2006): the probability of spike rises fast for reserve margins below 20%. A similar result was found by Ilic & Visudhiphan (2000) in the NEPOOL market: "Spot prices tend to vary proportionally to the ratio. However, when the ration exceeds 0.8 the proportionality is no longer valid. Therefore 0.80 could be considered to be a good cut-off value for the current NEPOOL market." It is good to note that the results by Mount et al. were for the ISO, who had better data than the average market participant. Results worsened for the average market participant who did not have access to this data. If we would continue on this line of thought, we could put a question mark on the reliability of the current capacity estimate. Another approach would be to question the stability of the relation. This is the approach we will take in the last paragraph of this section.

4.6 Stability

In this paragraph we discuss the stability of the relation we found above. In particular, we consider time dependence on the daily and yearly level, one-off events and out-of-sample performance. Before we start, let us give some reasons why there could be instability in the relation

4.6.1 Reasons for instability

In our approach we do not explicitly consider which behavior the system has in surrounding hours. This makes that we do not explicitly consider the effect of the start-up costs due to steepness of the load or maximum load (so-called peak-shaving units). This can make prices temporarily higher than the levels we would normally establish. A similar start-up cost problem is present in the weekend, where especially Sunday can present a difficulty for our reserve margin index. Turning down on Sunday is normally not an option as the power plant has to participate on Monday again when load goes up. Meanwhile the minimum and maximum load can vary a lot, such that the demand factor moves. Therefore we expect our index to perform worse in those situations.

For the weekend one has to make the decision to run certain units or not. After a decision to run, the unit is must run and prices can drop below normal levels for certain moments. In the winter, power plants who produce both heat and power, can become must-run in electricity to cover the heat demand. In certain cases the prices can even become negative. 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. For this hypothesis to be tested well, we should be able to make the distinction between flexible and inflexible (must-run) units in the forecast of available supply. Until now, this type of split has not been made. A testable implication of our hypothesis is that weekend days have a lower price even if the reserve margin is the same.

4.6.2 Dependence on time of day

In this and the next subsection we check whether the relation is similar among different subsets of the data. Here we consider the usual time-of-day subsets. Note, that the same product names are in use in different electricity markets, but that the exact definitions may vary. Table 3 describes the definitions in the Dutch market.

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

Table 3: Description of different time of day segments and the number of observations. Baseload denotes the full data set.

In figure 8 we see peak and off-peak prices are in line, though off-peak prices fall below peak prices for reserve margin between 0.40 and 0.50. This is in line with our hypothesis.

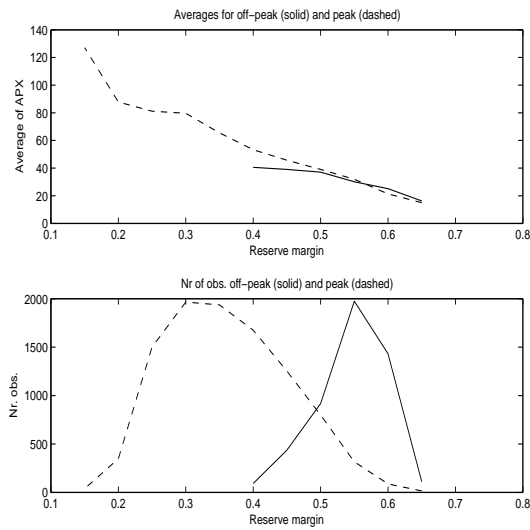


Figure 8: The average transformation occurring for peak and off-peak together with the number of underlying observations in each box. To show significant values, we suppress a time of day in a certain box if it contains less than 10 observations of that kind.

To study the distribution for the peak-prices the average transformation for weekend-peak, shoulder and super-peak are shown in figure 9. In the figure we see the shoulder transformation is sometimes above the super peak transformation (reserve margin 0.45 – 0.55), while the weekend-peak is below both. This confirms our hypothesis about weekend prices being lower than week prices for the same reserve margin. The shoulder hours are not fully explainable with the current figure. Therefore we consider the difference between summer and winter in the following subsection.

4.6.3 Dependence on season

Next, we ask ourselves whether there is a seasonal effect. We divided the data into summer (April-September) and winter months (October-March). In our data set of totally 14904 data points, we have 8640 data points in the winter and 6264 in the summer.

In figure 10 we see the difference between peak and off-peak is sustained if we split the data in summer and winter. The same conclusion holds for the relation between weekend-peak, shoulder and super-peak hours as can be seen in figure 11. This brings us to the conclusion that our hypothesis holds in general, but that the shoulder hours and super peak hours with reserve margin between 0.45 – 0.55 need to be studied in more detail.

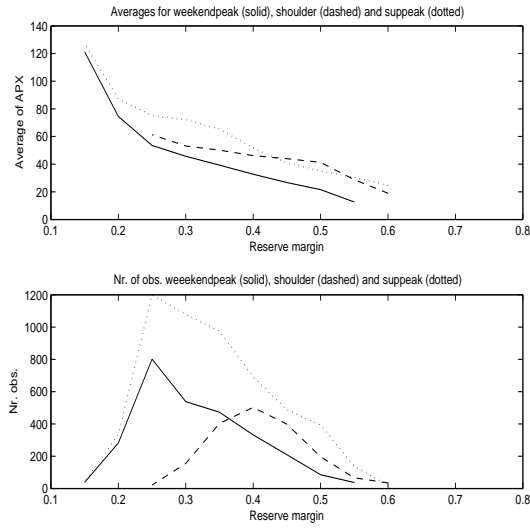


Figure 9: The average transformation occurring for three different time of day indices together with the number of underlying observations in each box. To show significant values, we suppress a time of day in a certain box if it contains less than 10 observations of that kind.




Figure 10: The average transformation occurring for peak and off-peak indices for both summer and winter. To show significant values, we suppress a time of day in a certain box if it contains less than 10 observations of that kind

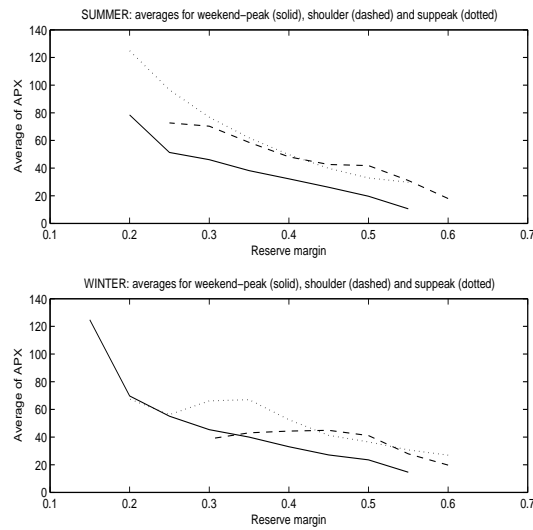


Figure 11: The average transformation occurring for three different time of day indices for both summer and winter. To show significant values, we suppress a time of day in a certain box if it contains less than 10 observations of that kind

4.6.4 Outliers

In the previous section, we have seen that certain data points were not in line with the average transformation. This leads to the question whether something strange has happened in our data sample. To find out whether there has been a bad period in the data, we consider how these “outlier” data are distributed over the data sample.

We decided to look at data points which have high APX values, and relatively high values of reserve margin. For example, if we use a definition of odd value as $APX > 200$, $reserve\ margin > 0.3$, we call 16 data points outliers while in total 128 data points had $APX > 200$. In figure 12 we graph how many of such data points were clustered in one day over time. We see there are two days with 6 outliers, and that they are mainly present around October 2005. For comparison reason, we included again the development of APX prices over time.

4.6.5 Out-of-sample

Up to 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 how stable the relation is over different parts of the sample.

For our stability check we divide our sample in three parts (the first 5000, the second 5000 and the remaining 4904 observations) and compare the average price (figure 13) and the probability of spike (figure 14). We see that the average price




Figure 12: The top panel shows data points with high APX values (> 200) and high reserve margin (> 0.2). The bottom panel shows APX prices

has increased, and that the transformation from the first period understates the average price and probability of spike in the second period. The transformation from the second period is close to the transformation of 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.

The natural next question is how many data points we should use in our data estimation by comparing the errors out-of-sample. We will address this question in the full paper. This means we cannot address the question yet whether we believe it is appropriate to specify a dynamic model like for example discussed by Fezzi & Bunn (2006).

We have not yet considered the impact of different prices on the spot price itself. Here one can think about the current real-time prices or recent spot prices. Also, the prices of the import capacity could prove to be a useful indicator. In case prices are rising, this could be an indicator that something is wrong in the market and prices could potentially spike.

5 Discussion and conclusion

In this work we have shown how to create an estimate for the supply-demand framework and how to build a simple model around it. One of the main findings is that reserve margin matters and should be included into a spot electricity model to enhance performance. Another useful area of application has shown to be the development of a fundamental model. While most fundamental models




Figure 13: Average price for three different time periods

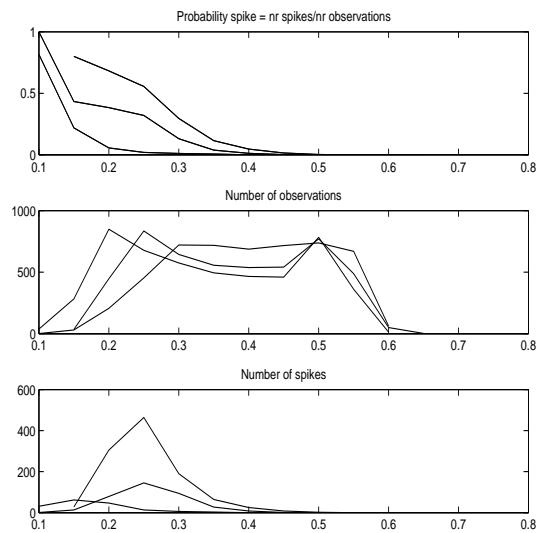


Figure 14: Probability of spike for three different time periods

can create estimates of future marginal costs, it needs a link from marginal costs to market prices. Our model can provide such a link if marginal costs are driven by reserve margin.

Estimates for available capacity are not public data in different electricity markets. Due to the simple nature of our approach, we hope it will be easy to replicate our results in other markets. In Europe, two examples would be the UK and German market who both started to publish estimates. The UK market is most similar to the Dutch market. For the German market, it is important to look carefully at the inter-connectors and wind production.

Our relation initially contained a double hump structure which did not comply with economic theory. As well, our results imply the Dutch market can spike for a still medium values of the reserve margin. This could be a proof of unreliable estimates for the available capacity, or that particular market participants do not believe the true estimate. Such a situation could remain until the regulator enforces that the data must be correct. Another improvement would be to group available capacity by technology. We tested the hypothesis that inflexible power plants have a downward impact on prices. We found the hypothesis holds in general cases, with the exception of the relation between shoulder hours and super peak hours for a reserve margin between 0.45 – 0.55.

The backbone of the our relation is an assumption on stability. We have studied the stability over different time of days and seasons. We found it is better to specify a separate model for different time of days, where especially it is worth to split week and weekend days. Out-of-sample test gave promising initial results, which will be extended.

The model can be extended in different directions. One of the directions is into 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. This road has been followed by Anderson (2004). The study of forward risk premia in a similar perspective has been performed by specifying a functional form for the relation between supply, demand and spot prices.

Another direction is the extension to a coupled market. This type of markets are present in the US and in Nordpool. As indicated by Belpex (2006), the Dutch, Belgian and French market will be integrated too. This will provide a new challenge for the current model.

References

- [1] Anderson, C.L., 2004, A Hybrid Model for Electricity Spot Prices. PhD thesis, The University of Western Ontario, Graduate Department of Applied Mathematics
- [2] Barlow, 2002, A diffusion model for electricity prices. *Mathematical Finance*, vol. 12, p. 287-298.
- [3] Belpex, 2006, announcement on market coupling. Can be downloaded from <http://www.belpex.be/index.php?id=4>

- [4] Bessembinder, H. & M.L. Lemmon, 2002, Equilibrium pricing and optimal hedging in electricity forward markets. *Journal of Finance*, vol. 57, p. 1347-1382
- [5] Birnbaum, L., J.M. del Aguila, G.D. Orive & P. Lekander, 2002, Why Electricity Markets Go Haywire. *McKinsey Quarterly*, vol. 1, p. 64-73
- [6] Boogert, A.F. & D.Y. Dupont, 2005a, On the effectiveness of the anti-gaming policy between the day-ahead and real-time electricity markets in the Netherlands, *Energy Economics*, vol. 27 (5), p. 752-770
- [7] Boogert, A.F. & D.Y. Dupont, 2005b, The nature of supply side effects on electricity prices: the impact of water temperature, *Economics Letters*, vol. 88 (1), p. 121-125
- [8] Burger, M., B. Klar, A. Muller & G. Schindlmayer, 2004, A spot market model for pricing derivatives in electricity markets. *Journal of Quantitative Finance*, vol. 4, p. 109-122
- [9] Cartea, A. & Figueroa, M.G., 2005, Pricing in electricity markets: a mean-reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, vol. 12 (4), p. 313-335
- [10] CBS, 2005, Duurzame energie; jaarcijfers: capaciteit, prod. en vermeden prim. energie.
- [11] Deng, S., 1999, Financial Methods in Competitive Electricity Markets. PhD dissertation, University of California, Berkeley.
- [12] Eydeland, A. & K. Wolyniec, 2003, Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging. Wiley Finance.
- [13] Eydeland, A. & H. Geman, 1998, Pricing power derivatives. *Risk*, p. 71-73
- [14] Eydeland, A. & H. Geman, 1999, Fundamentals of electricity derivatives. Energy Modelling & the Management of Uncertainty. Risk Books
- [15] Fezzi, C. & D. Bunn, 2006, Structural analysis of high-frequency electricity demand and supply interactions. Working paper, London Business School
- [16] Hughes, W.R., & A. Parece, 2002, The Economics of Price Spikes in Deregulated Power Market. *The Electricity Journal*, July, p. 31-44
- [17] Huisman, R. & R. Mahieu, 2003, Regime jumps in electricity prices. *Energy Economics*, vol. 25, p. 425-434
- [18] Karakatsani, N.V. & D. W. Bunn, 2005, Diurnal Reversals of Electricity Forward Premia. Working paper, London Business School
- [19] Kosecki, 1999, Fuel-Based Power Price Modelling. Energy Modelling & The Management of Uncertainty. Risk Books

- [20] Longstaff, F.A. & A.W. Wang, 2004, Electricity forward prices: a high-frequency empirical analysis. *Journal of Finance*, vol. 59 (4), p. 1877-1900
- [21] Lu, X., Z.Y. Dong & X. Li, 2005, Electricity market price spike forecast with data mining techniques. *Electric Power Systems Research*, vol. 73, p. 19-29
- [22] Mount, T.D., Y. Ning & X. Cai, 2006, Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Economics*, vol. 28 (1), p. 62-80
- [23] Pirrong & Jermakyan, 1999, Valuing power and weather derivatives on a mesh using finite difference methods. *Energy Modelling & The Management of Uncertainty*. Risk Books
- [24] Pirrong & Jermakyan, 2000, The Price of Power: the Valuation of Power and Weather Derivatives. Working paper, Olin School of Business, Washington University
- [25] Saravia, C., 2004. Speculative trading and market performance: the effect of arbitrageurs on efficiency and market power in the New York electricity market. Working paper, University of California, Berkeley
- [26] Skantze, P., A. Gubina & M. Ilic, 2000, Bid-based Stochastic Model for Electricity prices: The Impact of Fundamental Drivers on the Market Dynamics. MIT report EL00-004
- [27] Villaplana, P., 2005, Valuation of electricity forward contracts: the role of demand and capacity. Working paper, Universitat Pompeu Fabra
- [28] Visudhiphan, P. & M. Ilic, 2000, Dependence of Generation Market Power on the Demand/Supply Ratio: Analysis and Modeling. *IEEE Transactions on Power Systems*, vol. 2, p. 1115-1122