

Constructing Bidding Curves for a Price-Taking Retailer in the Norwegian Electricity Market

Stein-Erik Fleten and Erling Pettersen

Abstract—We propose a stochastic linear programming model for constructing piecewise-linear bidding curves to be submitted to Nord Pool, which is the Nordic power exchange. We consider the case of a price-taking power marketer who supplies electricity to price-sensitive end users. The objective is to minimize the expected cost of purchasing power from the day-ahead energy market and the short-term balancing market. The model is illustrated using a case study with data from Norway.

Index Terms—Demand-side bidding, electricity markets, stochastic optimization.

I. INTRODUCTION

THE NORDIC electricity market, comprising Denmark, Finland, Norway, and Sweden, has several hundred retailers, competing to serve a population of about 24 million. This amounts to 385 TWh of electricity per year (2000) at an average price of 20.4 Euro/MWh.¹ The Scandinavian market for retail electricity is very competitive. Retailers get their margins from buying wholesale and selling to end users. Retail electricity is a highly standardized and homogenous good, and both retail and wholesale prices are transparent, with different Web sites keeping track of around 100 retail companies' prices. The difference between the cheapest and most expensive is low, considering the uncertainty in input and competitors' prices and menu costs associated with changing prices. Keeping the costs low is essential for long-term survival as a retailer, and our model should be able to contribute to that goal, by aiming at the lowest level of balancing costs over time.

Although some consolidation is likely to occur, the industry is expected to stay reasonably competitive for many years. Another expected development in this industry is the degree to which two-way communication technology will be adopted. This involves hourly metering of each subscriber and some form of letting the end user know the current price of power. Over time, given large price differences over the course of a day and occurrences of price spikes, such a development will induce households to become more price responsive than today.

Manuscript received March 31, 2004; revised August 23, 2004. E. Pettersen was supported by grants from the Norwegian Electricity Industry Association and the Research Council of Norway. Paper no. TPWRS-00174-2004.

S.-E. Fleten is with the Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway (e-mail: stein-erik.fleten@iot.ntnu.no).

E. Pettersen was with the Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway. He is now with Troms Kraft, NO-9291 Tromsø, Norway (e-mail: erling.pettersen@troms-kraft.no).

Digital Object Identifier 10.1109/TPWRS.2005.846082

¹Average wholesale spot price in Trondheim region, 1997–2003.

In anticipation of this, we propose an optimization model for a price-taking retailer with end users who will reduce load if prices are high and increase load if prices are low.

Bidding in electricity markets has been studied by [1], [2], and others, for players with market power. Bidding for price-taking producers has been studied by, e.g., [3] and [4]. We study bidding for a retailer, and in contrast to previous studies, our approach provides piecewise linear bids (as opposed to, e.g., piecewise constant bids). A game-theoretic assessment of how a large buyer with market power should adapt to an electricity pool market with a day-ahead structure is found in [5]. General continuous bidding curves are used to characterize how such a market will behave, in contrast to our approach, which develops a decision support model for a relatively small retailer.

Today, retailers usually submit price-insensitive bids to the spot market at Nord Pool. This is due to the fact that consumers are not exposed to short-term price fluctuations, with monthly (or less often) metering being the norm. The bids are more or less close to expected demand, which is estimated by statistical techniques. Demand estimation takes into account factors such as past demand and forecasted temperature, which is important since more than 50% of electricity use in Norwegian households is for heating. See [6] for a load forecasting model.

Some managers are speculating that the balancing price will become either higher or lower than the spot price. If a manager is able to guess correctly on, e.g., the balancing price becoming larger than the spot price, he will benefit from submitting a curve that asks for a quantity larger than the expected load. However, it is the policy of the system operator, Statnett, to require all participants to submit bids in accordance with the load that each participant expects. If, e.g., a retailer is discovered deviating from this rule, he will receive a formal request to change bidding practices. This has happened in the past. It is not known exactly how and when the retailer continuing to break Statnett's policy will be penalized.²

Although this research is motivated by the needs of a Norwegian electricity retailer, we believe that the bidding model also can be applied by power marketers elsewhere, given a pool accepting piecewise linear bids and where there is a short-term balancing market whose prices affect the retailer costs whenever he experiences an imbalance between the load planned day-ahead and the realized load.

In Section II, we explain the market rules that are relevant for demand-side bidding in the Norwegian market. In Section III,

²The penalty given by the system operator should not be severe in our opinion, because profitable-biased bidding actually helps fulfill the intention of the policy, namely, to ensure that as much physical power as possible is exchanged on the day-ahead market as compared with on the balancing market.

we specify a stochastic linear program to be solved by the retailer to obtain a bidding curve that gives an optimal combination of expected cost and risk, according to the retailer's preferences. Section IV describes how the scenarios and other data can be set up for a specific case, in order to optimize a retailer's bid for a given hour. We present and discuss some solutions to our model in Section V, and in Section VI, a brief conclusion is given.

II. BIDDING PROCESS

For a Norwegian retailer, there are two important markets for physical exchange of electricity: the day-ahead market at Nord Pool, "Elsport," and the regulating market organized by Statnett, the independent system operator. In the day-ahead market, producers and retailers submit price-quantity bids for buying and selling electricity every day before noon. The bids are specified for each of the next 12–36 h, and Nord Pool subsequently collects all bids and calculates market clearing prices for each hour and each price area. Each country in the region constitutes a price area, with Norway divided further into two or three areas and Denmark divided into two areas. A retailer only submits bids for areas where he has customers.

The regulating market is used by the system operator to ensure real-time balance between supply and demand. Only producers with an ability to ramp up or down significantly on 15-min notice are allowed to participate.³ Whenever there is a load greater than was committed in the day-ahead market, there is a need for up regulation and vice versa. Statnett has collected bids for such up and down regulations for each participant and chooses to use the cheapest feasible source for such ramping.

When submitting an up-regulation bid, the producer obliges himself to increase production by certain amounts at certain prices on 15-minutes' notice. Up-regulating power must be offered at a higher price than the spot price. This implies that in an hour of up regulation, the price in the regulating market will be higher than the spot price. Submitting down-regulation bids means that the producer agrees to decrease production by certain amounts at certain prices on 15-minutes' notice. Down-regulating power must be offered at a lower price than the spot price, which implies that in an hour of down regulation, the price in the regulating market will be lower than the spot price.

If the actual consumption within a price area is equal to or very close to the volume dispatched in the spot market for that region, there is no need for regulation. The price in the regulating market then equals the spot price.

Now, assume that the market is being up regulated. We have explained that during an hour of up regulation, the regulating market price is higher than the spot price. Since there are many retailers operating within a price area, however, there is a chance that not all retailers are regulated in the same direction.

A retailer whose customers have used less electricity than the volume that particular retailer was dispatched at in the spot market will be down-regulated. This retailer will have to sell some power back to the regulating market at the regulating price,

which is higher than the spot price. In other words, that retailer has, in this case, actually made a profit by ordering more electricity in the spot market than his customers actually needed and then selling it back to the market at a higher price.

Other retailers in the same area may be up regulated, however, if they have ordered less electricity in the spot market than their customers actually consume. They will have to purchase the excess demand in the regulating market at a price that is higher than the spot price. These retailers would have been better off if they had been dispatched at a higher volume in the spot market.

A similar argument may be used to explain how a retailer that is up regulated individually would find it beneficial if the market is down regulated. Also, a down-regulated market would not benefit retailers that are down regulated individually. In other words, it is advantageous for a retailer to be regulated in the opposite direction of the market.

The Norwegian approach to handling imbalances between the spot market volumes and the actual consumption is different from the approach used in Denmark, Finland, and Sweden.

III. MODEL SPECIFICATION

We assume that the bidding actions of this retailer do not influence prices in the electricity markets, i.e., the retailer is a price taker operating in a competitive market. For Norway, this assumption is reasonable for the vast majority of retailers, who are relatively small, for example, compared to the largest producers and the largest consumers, the power-intensive industry.

We define β and π to be the regulating market price and the day-ahead ("spot") price, respectively. Based on the brief explanation of the regulating market design given in the introduction, we see that during up regulation $\beta - \pi > 0$, during down regulation $\beta - \pi < 0$, and when there is no regulation, $\beta - \pi = 0$.

The retailer's cost of purchasing electricity may be written as

$$C = C_d + C_r \quad (1)$$

where C_d is the amount paid in the day-ahead market, and C_r is the amount paid in the regulating market. They can be expressed as

$$C_d = \pi y \quad (2a)$$

$$C_r = (\xi - y) \beta \quad (2b)$$

where ξ is the load, and y is the volume knocked down in the day-ahead market. Observe that the regulating cost C_r is negative if $(\xi - y) < 0$. One may consider including the benefit from selling electricity to end users here. However, the benefit is the result of two exogenous variables: end-user price and load. Bidding does not influence the benefit of selling electricity; therefore, this benefit is not included. Also observe that the day-ahead volume y is the only variable that is controlled by the retailer; the parameters ξ , β , and π are exogenous.

Now, we combine the steps in (2) and get the following cost

$$C = \pi y + (\xi - y) \beta = \xi \beta - y \delta \quad (3)$$

where $\delta = \beta - \pi$ is the difference between the regulating price and the spot price.

³Retailers that are able to cut off some of their customers on 15-min notice may also participate in the regulating market. Demand-side bidding in the regulating market is still rather immature, however, so we leave that out for now.

The first term $\xi\beta$ does not include any decision variables. Hence, the retailer should seek to

$$\min E \{-y\delta\} \quad (4)$$

which is equivalent to minimizing the expected regulating market loss $E\{(\xi - y)\delta\}$. This relation denotes the money lost by purchasing the excess demand in the regulating market instead of in the spot market. Note that regulating market loss may well be negative, implying that the retailer may make extra profits in the regulating market.

Since the expected cost turns out to be negative in the case study, we have chosen to turn this around, so that the retailer's objective is to maximize expected regulating market profit or

$$\max E \{y\delta\}. \quad (5)$$

The expression in (5) is the main part of the objective function in our optimization model.

A. Bidpoints

To obtain a piecewise-linear strictly decreasing curve for the day-ahead auction, which is consistent with the bidding rules on the Nord Pool day-ahead market, the retailer submits n price-volume pairs $(P_0, x_0), (P_1, x_1), \dots, (P_n, x_n)$, where $P_0 \leq P_1 \leq \dots \leq P_n$ and $x_0 \geq x_1 \geq \dots \geq x_n$ to the pool. A linear interpolation between the pairs (P_i, x_i) and (P_{i+1}, x_{i+1}) , $i = 1, \dots, n-1$ gives the resulting $n-1$ line segments that decide at what price π and volume y the retailer is dispatched. Fig. 1 shows an example of such a bidding curve with three line segments. If the day-ahead price is π^* , the retailer is dispatched at the volume y^* .

The bidding curve is described by the following relation

$$\pi = \begin{cases} P_1 + \frac{P_0 - P_1}{x_0 - x_1} (y - x_1), & \text{if } x_1 < y \leq x_0 \\ P_2 + \frac{P_1 - P_2}{x_1 - x_2} (y - x_2), & \text{if } x_2 < y \leq x_1 \\ \vdots & \vdots \\ P_n + \frac{P_{n-1} - P_n}{x_{n-1} - x_n} (y - x_n), & \text{if } x_n \leq y \leq x_{n-1}. \end{cases} \quad (6)$$

Assuming that the retailer is a price taker, the relevant decision problem for him is to select volumes at which to be dispatched at the different prices. As is apparent from the relation in (6), the problem of selecting values of both P_i and x_i , $i = 1, \dots, n$ is nonlinear.

As an approximation to ease the solving of the problem, we make the model linear by fixing the price points P_0, \dots, P_n in

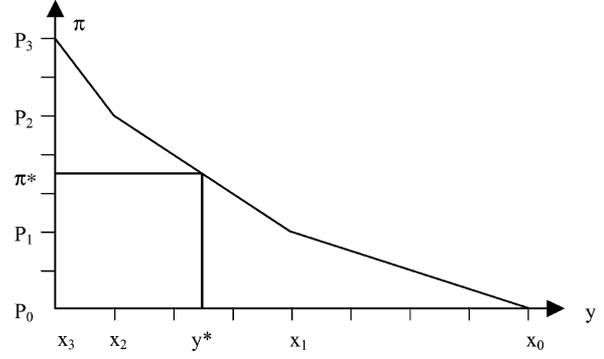


Fig. 1. Possible bidding curve with three line segments. The retailer submits the points (x_i, P_i) , $i = 0, \dots, 3$ to the pool, and the bidding curve emerges from a linear interpolation between those points.

advance. One plausible way of deciding predetermined values of the price points is to require that

$$\Pr \{P_i \leq \pi^* \leq P_{i+1} \cup x_{i+1} \leq y^* \leq x_i\} = \frac{1}{n-1} \quad \forall i \in \mathcal{I}. \quad (7)$$

Hence, the probability of being dispatched on any of the $n-1$ line segments is the same. This may be done by sorting the scenarios generated by price, π_s , $s \in \mathcal{S}$, where \mathcal{S} is the set of scenarios, and fix the values of P_i such that the number of scenarios between any price-volume pair is the same.

Since our problem is to select the volumes x_i corresponding to P_i , we rewrite (6) to make it linear in x , which is shown in (8) at the bottom of the page.

In (8), we have included the subscripts s to point out that y and π are scenario dependent. As indicated, we will generate scenarios for the spot price π .

For each scenario, the day-ahead price π_s will lie between two certain price points P_i and P_{i+1} . All the other $n-2$ price points are irrelevant for computing the knocked-down volume y_s for that scenario. Let $i(s)$ denote the largest i for which $P_i < \pi_s$: $i(s) = \sup_{i \in \mathcal{I}} \{i | P_i < \pi_s\}$. Hence, the point of dispatch (y_s, π_s) will lie on a line segment on the bidding curve that is succinctly described by the linear interpolation between the price-volume pairs $(P_{i(s)}, x_{i(s)})$ and $(P_{i(s)+1}, x_{i(s)+1})$. Relation (8) can now be written as

$$y_s = \left(1 - \frac{\pi_s}{P_{i(s)+1} - P_{i(s)}} + \frac{P_{i(s)}}{P_{i(s)+1} - P_{i(s)}}\right) x_{i(s)} + \left(\frac{\pi_s}{P_{i(s)+1} - P_{i(s)}} - \frac{P_{i(s)}}{P_{i(s)+1} - P_{i(s)}}\right) x_{i(s)+1}. \quad (9)$$

$$y_s = \begin{cases} \left(1 - \frac{\pi_s}{P_1 - P_0} + \frac{P_0}{P_1 - P_0}\right) x_0 + \left(\frac{\pi_s}{P_1 - P_0} - \frac{P_0}{P_1 - P_0}\right) x_1, & \text{if } P_0 \leq \pi_s < P_1 \\ \left(1 - \frac{\pi_s}{P_2 - P_1} + \frac{P_1}{P_2 - P_1}\right) x_1 + \left(\frac{\pi_s}{P_2 - P_1} - \frac{P_1}{P_2 - P_1}\right) x_2, & \text{if } P_1 \leq \pi_s < P_2 \\ \vdots & \vdots \\ \left(1 - \frac{\pi_s}{P_n - P_{n-1}} + \frac{P_{n-1}}{P_n - P_{n-1}}\right) x_{n-1} + \left(\frac{\pi_s}{P_n - P_{n-1}} - \frac{P_{n-1}}{P_n - P_{n-1}}\right) x_n, & \text{if } P_{n-1} \leq \pi_s \leq P_n. \end{cases} \quad (8)$$

The corresponding volume y_s dispatched in each scenario $s \in \mathcal{S}$ is then derived by relation (8).

B. Load

As explained in the introduction, the goal of this paper is to present a model for a retailer whose customers' load ξ is price flexible. A commonly used demand curve in the economics literature is a constant elasticity curve. Assume that for some price π_0 , the customers have an expected load $E[\xi] = \xi_0$, given the expected weather conditions. Then, the expected load at price π is derived by

$$E[\xi] = \xi_0 \left(\frac{\pi}{\pi_0} \right)^\eta \quad (10)$$

where η is the price elasticity.

Since the bidding curve for a certain hour is submitted to the pool between 12 to 36 hours ahead of the physical dispatch, the actual load is random. For example, the weather conditions may change, and even if they do not, the customers' behavior will be stochastic.⁴ We derive the load in each scenario by adding a prediction error term ε_s for which we generate scenarios. The load ξ_s is then derived by the relation

$$\xi_s = \xi_0 \left(\frac{\pi_s}{\pi_0} \right)^\eta + \varepsilon_s. \quad (11)$$

We denote $\xi_s - y_s$ the volume deviation in scenario s . Note that this is different from prediction error ε_s , which is actual load ξ_s minus expected load $E[\xi]$.

C. Risk

The retailer will be subject to two different risk factors that may influence his behavior in the day-ahead market.

The first risk factor is the profit risk, which is simply the risk of making less money than desired on the end-user sales. If both the volume deviation and the difference between the spot price and the regulating market price δ are large and have the same sign, the retailer may lose a considerable amount of money during one single hour, potentially causing financial problems. However, the balance settlements are made a few weeks after the physical dispatch, and then, the regulating costs are charged for one full week. Hence, the actual invoice (or credit memo) from the system operator will include the regulating costs for 168 hours. If we assume that our model gives results that are optimal in the long run, we find it fair to assume that this risk is rather irrelevant. In the model presented here, we have assumed the retailer to be risk neutral. Later in this section, however, we suggest how to include this risk if so desired.

The second risk factor is the volume deviation risk. Since all derivatives are quoted with respect to the day-ahead price, the system operator wants the day-ahead market to reflect the physical conditions. In short, they do not like demand-side speculation in the regulating market. Therefore, they are inclined to take measures toward retailers that are suspected of bidding too high or too low volumes in the day-ahead market on purpose.

⁴The firm that supplies the most commonly used model for predicting load in Norway asserts that the model has an average error of 2% when the weather conditions are known.

According to our data, $E[\delta] < 0$, which implies that the market is down regulated on expectation. If it were not for the volume deviation risk, a risk-neutral retailer would, therefore, order nothing in the day-ahead market and purchase the entire demand in the regulating market at a price that is likely to be lower than the day-ahead price. Hence, without possible volume deviation punishment, the day-ahead market would be rendered obsolete if the retailers are risk neutral.

An alternative way of modeling risk is by using shortfall costs, as shown, for example, in [7]. This may be done by defining variables $w_{1s}^+, w_{2s}^+, \dots, w_{ms}^+$ and $w_{1s}^-, w_{2s}^-, \dots, w_{ms}^-$. When the retailer is up regulated, $\sum_{m \in \mathcal{M}} w_{ms}^+ = \xi_s - y_s$ is positive, and when he is down regulated, $\sum_{m \in \mathcal{M}} w_{ms}^- = y_s - \xi_s$ is positive. Now, we let $T_m^+ > 0$ and $T_m^- > 0$ denote the marginal cost of piece m on the volume deviation risk function for positive and negative deviations, respectively. Then, we may penalize volume deviations in the objective function by adding the term

$$-V \sum_{m \in \mathcal{M}} (T_m^+ w_{ms}^+ + T_m^- w_{ms}^-) \quad (12)$$

where V quantifies the retailer's aversion to the volume deviation risk, relative to other objective terms.

This term will penalize volume deviations more the higher they get, and if we define $T_1^\pm < T_2^\pm < \dots < T_m^\pm$, the marginal penalty will increase with increasing deviations. Also, this would require the following constraints (W_m denotes the size of the m th piece)

$$\sum_{m \in \mathcal{M}} w_{ms}^+ + y_s \geq \xi_s, \quad \forall s \quad (13a)$$

$$\sum_{m \in \mathcal{M}} w_{ms}^- - y_s \geq -\xi_s, \quad \forall s \quad (13b)$$

$$0 \leq w_{ms}^\pm \leq W_m, \quad \forall m, s. \quad (13c)$$

By using shortfall costs, the profit risk could also be plausibly modeled. We define the variables $u_{1s}, u_{2s}, \dots, u_{ks}$, with the constraint

$$\sum_{k \in \mathcal{K}} u_{ks} \geq B - y_s \delta_s + \xi_s (\beta_s - U) \quad (14a)$$

$$0 \leq u_{ks} \leq Q_k \quad \forall k, s \quad (14b)$$

where B defines a target for the regulating market profit in each hour, U is the end-user price (assumed fixed here), $y_s \delta_s - \xi_s (\beta_s - U)$ is the regulating market profit, and Q_k is the size of the k th piece of the profit risk function. If the regulating profit exceeds the budget, the constraint is not active. Next, we let C_k denote the marginal cost of piece k on the profit risk function. Then, the term is added to the objective function

$$-R \sum_{k \in \mathcal{K}} C_k u_{ks} \quad (15)$$

where R denotes the retailer's aversion to profit risk.

As mentioned, losses incurred during one hour are unlikely to be a big issue to the retailer. A loss in one hour may be leveled by a gain the next hour, and therefore, it may seem a bit irrelevant to impose an hour-by-hour budget. Nevertheless, big losses may be

unpleasant, even in just a single hour. Also, if the company has suffered heavy regulating market losses recently, it may want to take back lost ground. The retailer would perhaps want to change B on a regular basis, e.g., each day, based on previous performance.

D. Full Model

Now, using the volume deviation approach to risk, the full model looks as follows.

$$\max \sum_{s \in \mathcal{S}} \rho_s \left\{ y_s \delta_s - V \sum_{m \in \mathcal{M}} (T_m^+ w_{ms}^+ + T_m^- w_{ms}^-) \right\} \quad (16)$$

subject to

$$y_s = \left(1 - \frac{\pi_s}{P_{i(s)+1} - P_{i(s)}} + \frac{P_{i(s)}}{P_{i(s)+1} - P_{i(s)}} \right) \cdot x_{i(s)} + \left(\frac{\pi_s}{P_{i(s)+1} - P_{i(s)}} - \frac{P_{i(s)}}{P_{i(s)+1} - P_{i(s)}} \right) \cdot x_{i(s)+1} \quad \forall s \quad (17a)$$

$$\sum_{m \in \mathcal{M}} w_{ms}^+ + y_s \geq \xi_s, \quad \forall s \quad (17b)$$

$$\sum_{m \in \mathcal{M}} w_{ms}^- - y_s \geq -\xi_s, \quad \forall s \quad (17c)$$

$$0 \leq w_{ms}^\pm \leq W_m, \quad \forall m, s \quad (17d)$$

$$x_i \geq x_{i+1}, \quad \forall i \in \mathcal{I} \quad (17e)$$

$$x_i \geq 0, \quad \forall i \in \mathcal{I}. \quad (17f)$$

The model uses the following indices/sets and parameters:

- s, \mathcal{S} count the scenarios;
- i, \mathcal{I} count the price-volume pairs;
- m , count pieces of the volume deviation risk function;
- \mathcal{M}
- π_s spot price in scenario s [Euros per megawatt-hour];
- ξ_s load from the retailer's customers [in megawatts];
- δ_s difference between the regulating market price and the spot price $\delta_s = \beta_s - \pi_s$ [in Euros per megawatt-hour];
- ρ_s probability of scenario s ;
- P_i price points on the price grid [in Euros per megawatt-hour];
- V weight on volume deviation risk in the objective function;
- T_m^\pm marginal cost of segment m on the piecewise linear volume deviation cost function, for positive and negative deviations [in Euros per megawatt-hour];
- W_m size of segment m on the piecewise linear volume deviation cost function [in megawatts].

Variables:

- x_i bidding volume corresponding to P_i [in megawatts];
- y_s volume dispatched in the day-ahead market [in megawatts];
- w_{ms}^\pm volume deviation, i.e., the degree to which dispatched volume y_s is different from actual load ξ_s [in megawatts].

This is a linear programming problem. The variables measuring dispatched volume y_s can be substituted using (17a).

TABLE I
SOME DESCRIPTIVE STATISTICS OF DATA REGARDING ELECTRICITY PRICES π
AND δ AND LOAD PREDICTION ERRORS ε

	π	δ	ε
Number of entries	25500	25500	4376
Highest entry	480.1	149.2	58.52
Lowest entry	2.19	-467.6	-70.86
Mean	22.3	-0.52	-1.34
Std. dev.	13.3	7.2	16.33

After this substitution, the number of variables is $|\mathcal{I}| + 2|\mathcal{M}||\mathcal{S}|$, and the number of constraints is $|\mathcal{I}| + 2|\mathcal{S}|$, not counting simple bounds.

IV. SCENARIO GENERATION

We have generated scenarios for spot price π , for load prediction error ε , and for the difference between the regulating market price and the spot price δ .

For π and δ , we have used hourly spot prices and regulating prices for the Trondheim region in the period 10 March 1997–16 December 2003. From these data, we have removed the prices in weekends and holidays, in addition to the prices between 24 hr and 07:00 hr. We removed the mentioned data, because we would like to consider bidding for a normal day-time hour. We were then left with 25 500 entries for π and δ .

Note that we are considering a single hour, e.g., in a shoulder period of a given day. The resulting bid curve will be optimal only for that hour. Typically, a retailer already has forecasting models that can be used to estimate load and prices, given the hour of the day, the day of the week, and the season and taking into account weather and possibly other factors. The output from such in-house forecasting models could be used in combination with historical data, as explained below.

For load prediction error ε , we have used measured and estimated load for the concession area of Trondheim Energiverk Nett AS⁵ (TEV) for every hour during the period 1 October 2002–31 March 2003. We subtracted estimated load from measured load to obtain ε , so that $\varepsilon > 0$ means that the retailer has been up regulated, while the retailer has been down regulated if $\varepsilon < 0$. As this is a rather short time series compared to those for π and δ , we have chosen not to remove any data. This leaves us with 4367 entries for ε . If we assume that our model is for a retailer serving the entire TEV concession area, this retailer was regulated in the opposite direction of the market in 1518 of the hours, while he was regulated in the same direction in 1983 of the hours. In 866, there was no regulation ($\delta = 0$).

Some descriptive statistics of the data set are shown in Table I. Prices are listed in Euros per megawatt-hour, while load prediction errors are listed in megawatt-hours.

According to [9], the regulating price strongly depends on the level of the spot price. Also, with such a huge data set, we find it natural to divide the data into subsets so that we may get a more robust treatment of data that are in the outer edges of our sample. Therefore, we have divided the price data π and δ into three sets with respect to π . The first set is made up by the 15% lowest

⁵Trondheim Energiverk Nett AS is the network operator serving Trondheim and Klæbu municipals, which is an area with 157 887 inhabitants as of 1 January 2003. Source: [8].

TABLE II
NUMBER OF ENTRIES IN EACH SUBSET AND PROBABILITIES

Subset	Number of entries	Pr $\{h \mathcal{T}_l\}$	Pr $\{\mathcal{T}_{l,h}\}$
$\mathcal{T}_{1,a}$	1729	0.452	0.0678
$\mathcal{T}_{1,b}$	1185	0.310	0.0464
$\mathcal{T}_{1,c}$	911	0.238	0.0357
$\mathcal{T}_{2,a}$	7234	0.405	0.284
$\mathcal{T}_{2,b}$	6885	0.388	0.270
$\mathcal{T}_{2,c}$	3730	0.209	0.146
$\mathcal{T}_{3,a}$	1347	0.352	0.0528
$\mathcal{T}_{3,b}$	1666	0.435	0.0653
$\mathcal{T}_{3,c}$	813	0.212	0.0319

values of π together with their corresponding values of δ . The second set is made up by the middle 70% values of π together with the corresponding values of δ , while the remaining highest 15% of the values of π with their corresponding values of δ form the third set. We denote these sets \mathcal{T}_1 , \mathcal{T}_2 , and \mathcal{T}_3 , respectively.

To generate scenarios, we need to estimate the first four moments of the variables' distributions. However, in about 21% of the 25 500 hours for which we have price data, there was no regulation, implying $\delta = 0$. This would significantly bias the estimates of the moments for δ , and therefore, we need to treat hours with no regulation separately.

When there is regulation, that is, $\delta \neq 0$, it is the absolute value of δ that counts toward the retailer's regulating cost. Hence, estimating moments for all $\delta \neq 0$ together would give a mean that is too close to zero. One way of dealing with this could be to estimate moments based on $|\delta|$. However, both [9] and our own analysis indicate that the characteristics of $|\delta|$ are influenced by the direction of regulation.

Due to the above-mentioned reasons, we divided each of the sets \mathcal{T}_1 , \mathcal{T}_2 , and \mathcal{T}_3 into three subsets according to the direction of the regulating price. That is, we divided each set into one subset for $\delta < 0$ (down regulation), one subset for $\delta > 0$ (up regulation), and one set for $\delta = 0$ (no regulation). We denote these subsets $\mathcal{T}_{l,h}$, for $h = a, b, c$, respectively, with $l = 1, 2, 3$. Hence, the subset $\mathcal{T}_{1,a}$ includes the price pairs for down regulation among the 15% lowest spot prices.

As mentioned, our time series for ε is not large enough to get one ε for each pair of π and δ . Therefore, we have estimated only one set of moments for ε , which has been added to the nine series with pairs of π and δ . Since we are dealing with a price-taking retailer, we may assume that the load prediction error of this particular retailer has no influence on the prices. Hence, the fact that our volume data are a bit fewer in number than the price data should not disturb our analysis significantly, even though it would have been preferable to have equal time series of data for all three variables.

In Table II, we have listed the number of entries in each subset. The table also shows some probabilities. The column Pr $\{h|\mathcal{T}_l\}$ lists the probability of the market being regulated in direction h , $h = a, b, c$ for set \mathcal{T}_l , $l = 1, 2, 3$. We see that the probability of down regulation is significantly higher for the lowest 15% of spot prices, that is, the set \mathcal{T}_1 . The probability of up regulation is significantly higher for the highest spot prices than for the lowest spot prices. The probability of no regulation seems quite stable.

TABLE III
DESCRIPTIVE STATISTICS FOR δ , THE DIFFERENCE BETWEEN REGULATING MARKET PRICE AND SPOT PRICE

Subset	Mean	Std. dev.	Skewness	Kurtosis
$\mathcal{T}_{1,a}$	-3.60	1.63	-0.400	2.99
$\mathcal{T}_{1,b}$	2.90	2.26	1.77	7.51
$\mathcal{T}_{1,c}$	0	0	0	0
$\mathcal{T}_{2,a}$	-3.88	2.11	-2.05	12.9
$\mathcal{T}_{2,b}$	3.23	5.15	17.3	429
$\mathcal{T}_{2,c}$	0	0	0	0
$\mathcal{T}_{3,a}$	-9.56	23.4	-13.3	216
$\mathcal{T}_{3,b}$	4.78	8.50	6.50	56.1
$\mathcal{T}_{3,c}$	0	0	0	0

TABLE IV
DESCRIPTIVE STATISTICS FOR SPOT PRICE π

Subset	Mean	St. dev.	Skewness	Kurtosis
$\mathcal{T}_{1,a}$	10.5	2.64	-0.700	2.49
$\mathcal{T}_{1,b}$	10.4	2.75	-0.634	2.51
$\mathcal{T}_{1,c}$	10.3	2.91	-0.770	2.68
$\mathcal{T}_{2,a}$	19.8	4.46	0.955	3.24
$\mathcal{T}_{2,b}$	20.1	4.70	0.940	3.12
$\mathcal{T}_{2,c}$	20.1	4.70	0.798	2.89
$\mathcal{T}_{3,a}$	48.8	25.0	8.22	111
$\mathcal{T}_{3,b}$	45.4	16.1	3.22	18.8
$\mathcal{T}_{3,c}$	47.0	19.4	3.38	17.8

TABLE V
DESCRIPTIVE STATISTICS FOR LOAD PREDICTION ERROR ε

	Mean	St. dev.	Skewness	Kurtosis
All subsets	-1.34	16.3	$1.87 * 10^{-4}$	$4.12 * 10^{-3}$

The column Pr $\{\mathcal{T}_{l,h}\}$ lists the probability of a randomly selected pair of price data π and δ , belonging to the subset $\mathcal{T}_{l,h}$, $l = 1, 2, 3$ and $h = a, b, c$. This statistic is calculated as Pr $\{\mathcal{T}_{l,h}\} = \text{Pr}\{h|\mathcal{T}_l\} \text{Pr}\{\mathcal{T}_l\}$, $l = 1, 2, 3$ and $h = a, b, c$, where Pr $\{\mathcal{T}_1\} = \text{Pr}\{\mathcal{T}_3\} = .15$ and Pr $\{\mathcal{T}_2\} = .7$.

Table III shows descriptive statistics for the variable δ . We observe that the mean δ at down regulation is much lower when the spot price is high ($\mathcal{T}_{3,a}$) than otherwise. Also, the absolute value of the mean of δ is a bit higher for down regulation than for up regulation in all sets. The means are listed in Euros per megawatthour. Furthermore, we observe very high kurtosis for $\mathcal{T}_{2,b}$ and $\mathcal{T}_{3,a}$, which is due to some extreme events in these subsets. For the lowest spot prices, δ seems to be quite stable, with all four moments relatively low. Obviously, all moments equal zero when there is no regulation.

Table IV shows descriptive statistics for π . We observe that the variability of π seems to show some more variation in the set of high spot prices \mathcal{T}_3 . The high positive skewness and the high kurtosis in this set indicates that some of the spot prices in this set are significantly higher than the mean. The mean spot prices are listed in Euros per megawatthour.

Table V shows descriptive statistics for ε . The mean prediction error is measured in megawatthours.

To generate scenarios, we also use the correlations between the variables. In subset $\mathcal{T}_{2,a}$, the correlation between π and δ was estimated to be $-.435$. Apart from that, the correlations were all rather low. We have chosen not to display the correlation matrices in the text, although they are allowed for in the analysis.

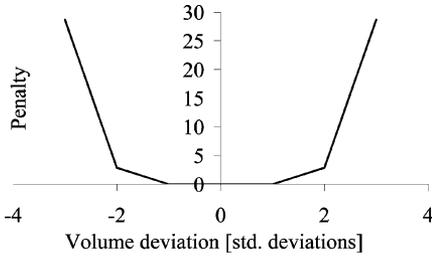


Fig. 2. Volume deviation penalty function.

Based on the descriptive statistics presented in Tables III–V, we generated scenarios for each subset by using the method described in [10]. The advantage of this method is that it is quite simple since one only needs the first four moments of each variable, plus the correlations. This requires, however, that the distributions and the relationships between them be succinctly described by those statistics. This information can be extracted from, e.g., in-house forecasting models and/or historical data (as we have done here). For most sets, we generated 100 scenarios. For the sets $\mathcal{T}_{2,b}$ and $\mathcal{T}_{3,a}$, however, the algorithm used did not converge for 100 scenarios, possibly due to their relatively extreme skewnesses and kurtosis. For those, we generated 600 and 300 scenarios, respectively.

V. SOLUTION

The specification of the volume deviation penalty function is necessarily *ad hoc*. The user has to guess how to set this function, whose form will have a major impact on retailer bidding decisions. We chose to work with the piecewise linear penalty function, as shown in Fig. 2.

If the realized load is larger than one standard deviation from its expected value, one starts penalizing using a marginal penalty of 2.86, which is equal to the expected spot price. This holds for deviations up to two standard deviations. From two standard deviations and up, the marginal penalty is increased by a factor of ten.

We modeled and solved the linear program using XPRESS [11]. A typical problem was compiled and solved in about 16 seconds on a 1-GHz Pentium PC with 524 MB RAM. No special effort was made to make the code efficient with respect to solution time, as only a fraction of the time mentioned (ca. 0.5 s) is used in the dual simplex algorithm. This problem is a simple recourse problem for which there exists special purpose algorithms (see, [12]). Such algorithms may become useful if the model is expanded.

An example of an optimal bidding curve is shown in Fig. 3. This curve has 7 line pieces. Also shown is the expected load curve and the load-price scenarios.

As mentioned, since $E[\delta] < 0$ in our data, the optimal curve asks for less power than the end users are expected to consume.

Using the data described in the previous section, we selected $n = 64$ price points (63 line pieces) according to (7). We do not know what the price elasticity of demand will be in a future with two-way communication. The retailer has to estimate this for the given hour for which a bidding curve is to be constructed. Thus, we have constructed bidding curves with a range of elasticities

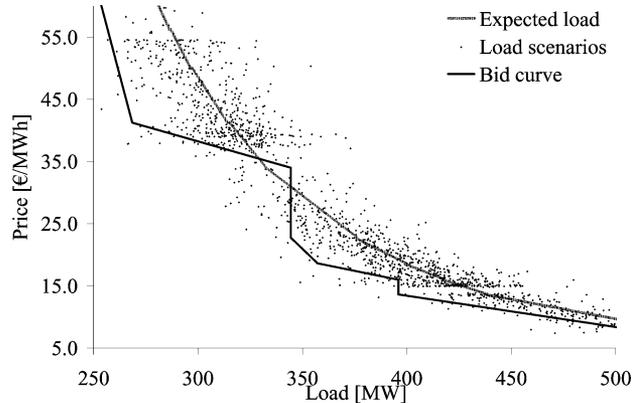


Fig. 3. Example of optimal bid curve, expected load curve, and load scenarios.

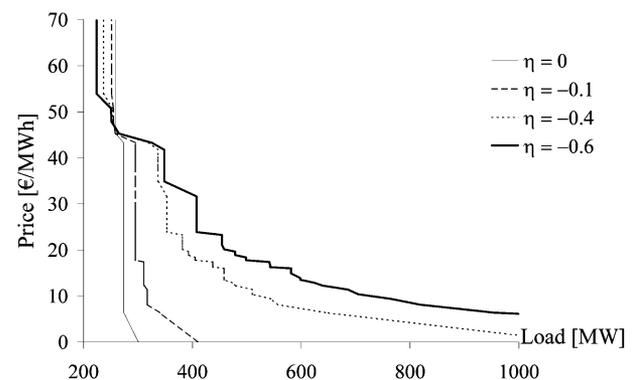


Fig. 4. Bidding curves with varying price elasticity of demand.

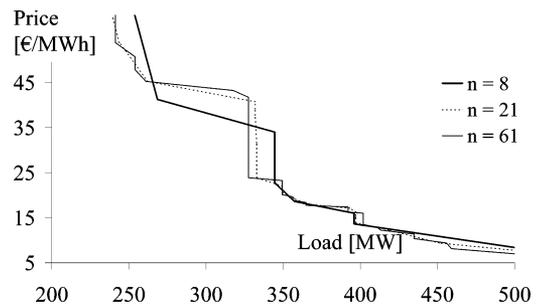


Fig. 5. Bidding curves with a varying number of fixed price points.

from $\eta = -0.6$ to $\eta = 0$. See [6] for a discussion. Fig. 4 shows the optimal bid curves for varying elasticity.

Using a price elasticity of demand of $\eta = -0.3$, we show the effect of varying the number of fixed price points (see Fig. 5).

If the number of price points is decreased, the result is a more crude bidding curve. Notice also that the bidding curves are cruder at prices that are less likely, e.g., above 50 Euro/MWh. This is reasonable and is due to (7).

Fig. 6 shows the cumulative profit distribution and illustrates the risk associated with bidding in this particular case. The expected profit is 26.5 Euros, which is a small amount. The risk is also relatively small, with a one-day 95% VaR of 164 Euros.

Next, we show the effect of varying the weight parameter V in the objective function and also varying the number of fixed price points (see Fig. 7).

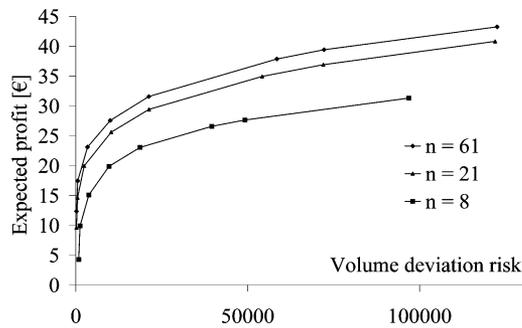


Fig. 6. Cumulative distribution function for profit.

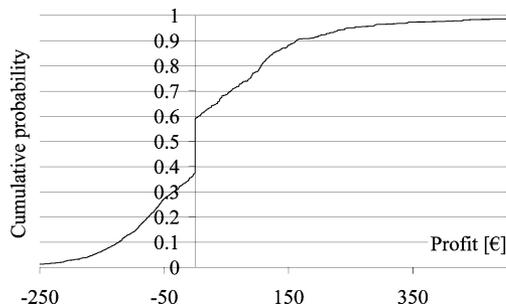


Fig. 7. Tradeoff between expected profit and volume deviation risk.

Volume deviation risk is measured as the expected volume deviation penalty $\sum_{s \in \mathcal{S}} \rho_s \sum_{m \in \mathcal{M}} (T_m^+ w_{ms}^+ + T_m^- w_{ms}^-)$. We see that there is a tradeoff between expected profit and the expected penalty of the difference between volume bid and realized load. If the decision maker wants high expected profit, he has to submit a bid that deviates from the expected load. The larger the volume deviation, the higher the expected profit.

From the figure, we can also see that the retailer is better off by using many price points, and the marginal benefit of increasing the number of price points is a decreasing function.

VI. CONCLUSIONS AND FUTURE RESEARCH

In summary, we have suggested a model supporting the construction of a bidding curve that fits the rules of the Nord Pool day-ahead market. The intended user is a retailer having end users with price-sensitive demand. The case study indicates that the model is able to construct plausible bidding curves, and there is a tradeoff between the expected cost of purchasing electricity and the risk of getting penalized from the system operator because bids deviate from the expected load.

Model validation is left for future work. There is a growing number of tests one can make in order to learn about the quality of the model, e.g., regarding the generation of scenarios, as explained by [13].

A retailer can view each of the next 12–36 h as candidate problems for which he can use our model individually. This is in contrast to a producer, who may have costs associated with changing the output schedule from one hour to the next. The research challenge in this situation is to appropriately model

the joint time-series properties of spot prices, regulating prices, and customer loads. Our paper does not contribute regarding the forecasting and time-series modeling of loads and prices. A research effort trying to properly address the multiperiod aspect of this problem should probably wait until data regarding, for example, dynamic price responsiveness can be obtained.

REFERENCES

- [1] E. J. Anderson and A. B. Philpott, "Optimal offer construction in electricity markets," *Math. Oper. Res.*, vol. 27, no. 1, pp. 82–100, 2002.
- [2] R. Green and D. M. Newbery, "Competition in the British electricity spot market," *J. Polit. Econ.*, vol. 100, no. 5, pp. 929–953, 1992.
- [3] A. K. David, "Competitive bidding in electricity supply," *Proc. Inst. Elect. Eng., pt. C, Gen., Transm., Distrib.*, vol. 140, no. 5, pp. 421–426, 1993.
- [4] A. J. Conejo, F. J. Nogales, and J. M. Arroyo, "Price-taker bidding strategy under price uncertainty," *IEEE Trans. Power Syst.*, vol. 17, no. 4, pp. 1081–1088, Nov. 2002.
- [5] A. B. Philpott and E. Pettersen, "Optimizing Demand-Side Bids in Day-Ahead Electricity Markets," *IEEE Trans. Power Syst.*, submitted for publication.
- [6] D. W. Bunn, "Forecasting loads and prices in competitive power markets," *IEEE*, vol. 88, no. 2, pp. 163–169, Feb. 2000.
- [7] S.-E. Fleten, S. W. Wallace, and W. T. Ziemba, "Hedging electricity portfolios using stochastic programming," in *Decision Making under Uncertainty: Energy and Power*, C. Greengard and A. Ruszczyński, Eds. New York: Springer-Verlag, 2002, vol. 128, IMA Volumes on Mathematics and Its Applications, pp. 71–94.
- [8] Statistics Norway, 19 Population, by Sex and Age. Municipality. 1 January 2003, Statistics Norway (2003). [Online]. Available: http://www.ssb.no/english/subjects/02/01/10/folkemengde_en/tab-2003-03-17-19-en.html
- [9] K. Skytte, "The regulating power market on the Nordic power exchange Nord Pool: An econometric analysis," *Energy Econ.*, vol. 21, pp. 295–308, 1999.
- [10] K. Høyland, M. Kaut, and S. W. Wallace, "A heuristic for moment-matching scenario generation," *Computat. Optimiz. Applicat.*, vol. 24, no. 2–3, pp. 169–185, 2003.
- [11] *Xpress-Mosel User Guide Release 1.2*, Dash Associates, 1998.
- [12] P. Kall and S. Wallace, *Stochastic Programming*. Chichester, U.K.: Wiley, 1994.
- [13] M. Kaut and S. W. Wallace, "Evaluation of scenario-generation methods for stochastic programming," SPEPS, Working Paper 14, 2003.



Stein-Erik Fleten received the Ph.D. degree from the Norwegian University of Science and Technology (NTNU), Trondheim, Norway, in 2000.

He is an Associate Professor at NTNU. Past and current research topics include real options, market power in electricity markets, hydro scheduling under uncertainty, and hedging and risk management in electricity utilities. His main research fields are stochastic programming and finance and energy economics, with applications for electricity companies.



Erling Pettersen received the "Sivilingeniør" degree (which is similar to the M.Sc. degree) in 1999 and the Ph.D. degree in 2004 from the Norwegian University of Science and Technology, Trondheim, Norway.

His main research interest is electricity market modeling, and he has been particularly concerned with end-user flexibility in electricity markets. He is currently an analyst with Troms Kraft, Tromsø, Norway.