



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Energy Economics 25 (2003) 409–424

Energy
Economics

www.elsevier.com/locate/eneco

Constructing forward price curves in electricity markets

Stein-Erik Fleten^a, Jacob Lemming^{b,*}

^a*Department of Industrial Economics and Technology Management,*

Norwegian University of Science and Technology, Trondheim NO-7491, Norway

^b*Systems Analysis Department, Risø National Laboratory, Roskilde DK-4000, Denmark*

Abstract

We present and analyze a method for constructing approximated high-resolution forward price curves in electricity markets. Because a limited number of forward or futures contracts are traded in the market, only a limited picture of the theoretical continuous forward price curve is available to the analyst. Our method combines the information contained in observed bid and ask prices with information from the forecasts generated by bottom-up models. As an example, we use information concerning the shape of the seasonal variation from a bottom-up model to improve the forward price curve quoted on the Nordic power exchange. © 2003 Elsevier Science B.V. All rights reserved.

JEL classifications: G13; C61; C51; Q49

Keywords: Electricity prices; Forward price curves; Electricity markets; Price modelling

1. Introduction

Commodity forward and futures prices serve an important role as information carriers for operational and investment decisions. However, the term structure that can be observed at any given time will be based on a limited number of products regardless of liquidity in the market. Even in highly liquid bond markets, one will often find that the product one needs to price is not traded in the market. There is often a need to estimate the prices for more maturity dates than are observed in the market.

*Corresponding author. Tel.: +45-46775142; fax: +45-46775199.

E-mail address: jacob.lemming@risoe.dk (J. Lemming).

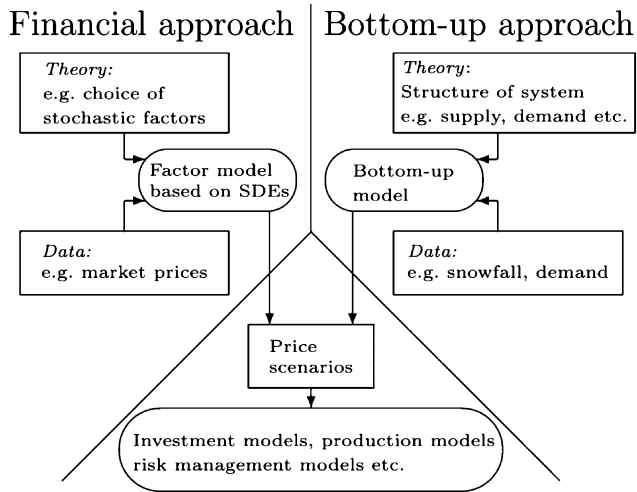


Fig. 1. Two approaches for transforming data and theory into price scenarios, which can be used as input to decision models.

Depending on the required level of accuracy an analyst can choose to use an interpolation between prices of traded products or perform a regression using some smooth function. Such an approach, suggested, e.g. by Adams and van Deventer (1994), will generally suffice if the market is mature and if liquidity of the quoted products is high. In electricity markets, neither of these criteria apply, as markets are generally young and many products struggle with low liquidity. Recognizing the insufficiency of market data, it therefore seems natural to look for additional information. In electricity and other commodity markets, comprehensive bottom-up models often exist, and when used as a supplement rather than a competing methodology, they can often provide additional valuable information.

To compensate for the deficiencies that arise from separate use of either market data or bottom-up models, we suggest a Bayesian inspired approach. The idea is to use market data as an a priori set of information and then form an a posteriori information set, by combining the market prices with forecasts from a bottom-up model. The proposed model constructs daily forward prices forming a smooth forward curve. The model is a quadratic program that uses bid and ask prices to constrain the forward prices from below and above. The objective function ensures both smoothness of the curve and that the curve follows the seasonality of the price forecast of a bottom-up model. Though we restrict ourselves to a model of the Scandinavian futures/forward market and the use of a single bottom-up model, the idea can be seen as a general approach where several bottom-up models might be used to add information to the initial set of market data. Fig. 1 provides an illustration.

Testing the approach against alternatives such as using a truncated Fourier series or maximum smoothness, we find that our model is robust and that its ability to give forward prices within traded maturity is promising.

The article is structured as follows: Section 2 discusses the electricity forward and futures prices; Section 3 presents the model; Section 4 reports on tests of the quality of the generated curves; and Section 5 concludes.

2. Electricity forward markets

An important characteristic of the electricity market is that the power once generated, cannot be stored economically. An exception is mountain reservoirs connected to hydroelectric plants and also manmade hydro pumped-storage facilities that usually are employed for shorter term storage. Producers owning storage facilities benefit from the varying electricity prices and use their production flexibility to produce little (or consume electricity when using pumps) when prices are low and produce near capacity when prices are high. Pumping entails energy losses in the order of 30% however, and aggregate pumping and storage capacity in the energy systems is usually low. Due to the limited storage capacity and flexibility, even hydro-dominated power systems have periodic variations in prices.

Limited storage means that forwards and futures cannot be priced using the standard arbitrage arguments involving cost-of-carry relationships. Forward and futures prices are the result of supply and demand for hedging and speculation. Producers hedge by selling (going short) and power marketers and power-intensive industry hedge by buying (going long). Speculators, which often include producers and power marketers/consumers, enter both sides of the market depending on their expectations and risk-taking ability. These expectations about future spot prices are often formed by price forecasts from bottom-up models.

In the former regulated regime, bottom-up models served multiple purposes including prediction of the marginal cost of electricity production. The merits of such models typically include a detailed technical description of generation, transmission and distribution systems as well as an extensive set of data on hydrological conditions, fuel prices and consumer behaviour. The main drawback of these models is that they cannot estimate or capture the risk premium or market price of risk determined by the market forces.

A class of commodity price models designed to capture the market price of risk is the financial asset pricing models, originating from Black (1976). In this approach, a set of coupled stochastic differential equations is used to describe the time dynamics of prices. As described in Clewlow and Strickland (1999) most literature on commodity pricing fall into one of two categories that originates from the theory of bond pricing. The first category can be seen as an extension of the SDE based short rate models from bond pricing. The idea is to formulate a set of SDEs for the spot price and additional factors that affect the spot price. Forward prices can then subsequently be derived from this representation. Recent contributions generally include more than one stochastic factor e.g. stochastic convenience yield and stochastic interest rates (Schwartz, 1997). Lucia and Schwartz (2002) study both

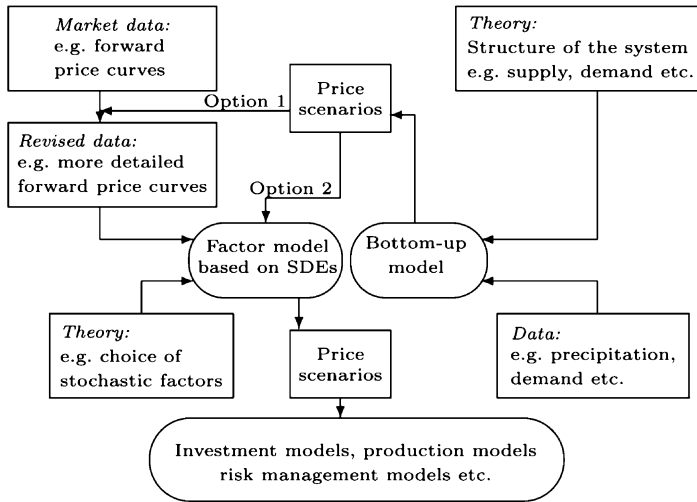


Fig. 2. Relative approach, combining the bottom-up and financial approaches. Both empirical information and theory is used to create price scenarios that are input to decision support models.

single- and two-factor models with stochastic volatility and various forms of deterministic seasonal components, using data from the Nordic market.

The second category of asset pricing models is based on the Heath-Jarrow-Morton (HJM) framework (Heath et al., 1992). In this framework, the time dynamics of the entire forward price curve is described using a multi-dimensional Brownian motion or state variable. The basic idea is to simultaneously describe the stochastic time dynamics of the entire forward price curve using a volatility function and the initial forward price curve observed in the market. Bjerk Sund et al. (2000) use a sum of three uncorrelated Brownian motions to capture the time dynamics in the shape of the electricity forward price curve. To find forward prices for dates that fall between traded maturity dates, the maximum smoothness criterion is applied.

Given the credibility that such factor models have obtained, it seems reasonable to argue that in deregulated markets it is more natural to price electricity using the financial approach than to repair bottom-up models that were created to work well in a different environment. The two financial approaches have already been applied to various electricity markets. The case studies are, however, generally based on a rather scarce data material due to the youth of most markets.

Whatever structural choices the analyst makes with respect to stochastic factors, the quality of the model output will depend crucially on the input data used to estimate the parameters of the model. Rather than relying on a specific financial model, we prefer to use a non-parametric approach taking the market forward prices as given. We suggest using bottom-up data to create an improved forward price curve as illustrated in Fig. 2. An alternative can be to include bottom-up data in the parameters estimation process and thereby basically introduce information from the

bottom-up model into the construction of the factor model (shown as option 2 in Fig. 2).

In the Nordic market, seasonality is an essential characteristic of electricity prices due to forces on both the demand and the supply side. On the supply side, the region relies heavily on Norwegian and Swedish hydropower plants that receive high levels of inflow in spring and summertime when snow in the mountain melts. Due to capacity constraints, these plants must produce at high levels during the summertime in order to avoid costly spill resulting from overflow in the reservoirs. This naturally creates a downward pressure on prices, which is exacerbated by a low demand. Unlike other regions, the cold climate in Scandinavia means that there is little need for air conditioning. Conversely, in the winter period there is an upward price pressure from the demand side as a result of high electricity demand for heating purposes especially in Norway, Finland and Sweden.

Futures and forwards are traded over the counter and in a financial market (Eltermin) at the Nordic power exchange Nord Pool.¹ Unlike other energy commodities such as oil or gas, electricity is a flow commodity implying that delivery of a specified constant (or deterministically time-varying) power level takes place over a period of time (delivery period) rather than at a specific point in time. A forward/futures electricity contract can, therefore, be viewed as a portfolio of basic forward/futures contracts each with different time to maturities, one for each point in time during the delivery period. By point in time is meant a specific hour of a day, as the underlying asset for the contracts is the hour-by-hour spot price in the Nordic market. This spot price is calculated and published every day by Nord Pool. We assume throughout this article that interest rates are deterministic and constant in time, in which case forward and futures contract prices are equal and can be treated as similar products for modelling purposes.

Fig. 3 illustrates the structure of both a forward and a futures contract in the shape provided by Nord Pool. Notice that at T_1 (the start of the delivery period) the spot price and futures price does not necessarily coincide as theory would require if the commodity was to be delivered at that specific point in time. Because the electricity futures concerns power delivered over a period of time $[T_1, T_2]$ the relevant statistic for the contract is not the spot price at T_1 , but the average spot price during the interval $[T_1, T_2]$. Ex post, the difference $F(T_1, T_2) - \sum_{t=T_1}^{T_2} S_t / (T_2 - T_1)$ between the closing futures price at T_1 , $F(T_1, T_2)$ and the average of the spot price S_t in $[T_1, T_2]$, is paid to the buyer and is charged the seller (if the difference is negative the seller gains and buyer loses). A futures contract thus offers a perfect hedge against the risk in the $[T_1, T_2]$ -period average spot (system) price for a constant MW level position held throughout the delivery period.

The forwards and futures products listed by Nord Pool differ not only in terms of time to maturity but also in terms of the length of the delivery period. This

¹ A basic forward contract is an agreement between a buyer and a seller on the future delivery of a product at an agreed price. A basic futures contract is a similar agreement but it is generally a more standardized product and has daily financial settlement until time of maturity through the use of margin accounts.

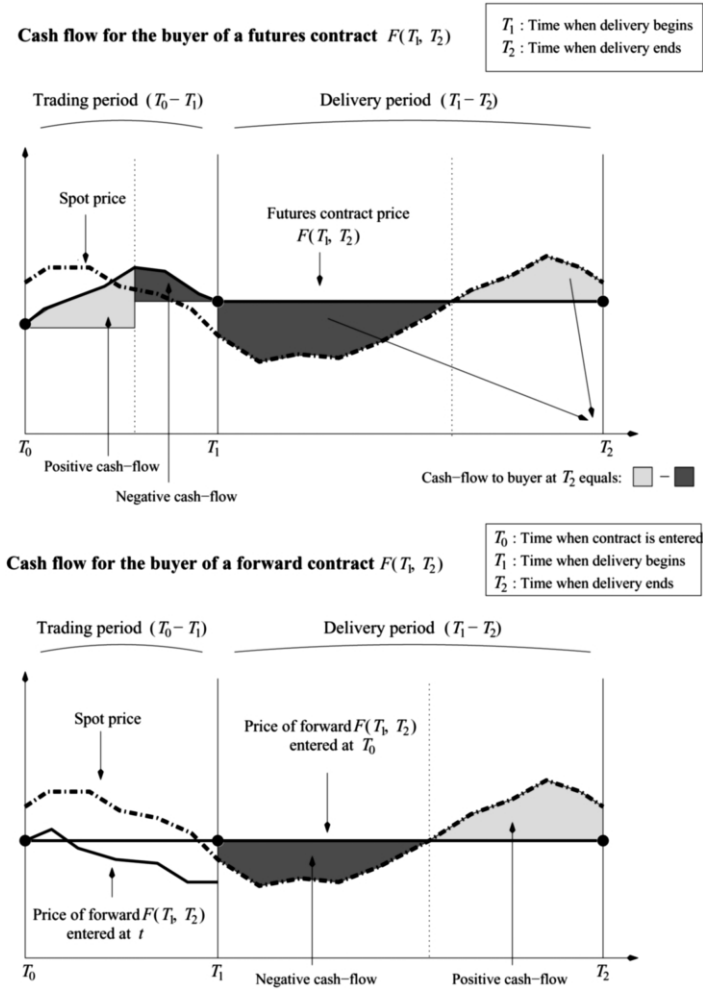


Fig. 3. The structure of future contracts (top figure) and forward contracts (bottom figure) supplied by Nord Pool.

means that at any point in time the decision maker has only a partial picture of the forward price curve available for analysis. Fig. 4 shows a series of typical contract prices where the horizontal regions indicate delivery period intervals. There is no trading in contracts with start of delivery *within* these intervals.

Futures contracts having delivery periods of 4 weeks are traded for maturities ranging from 4 weeks to 12 months into the future. As maturity draws nearer, the nearest contract is split into 4 weekly contracts, which are traded until the working day before the start of delivery. For forward contracts, the delivery periods are seasons and years. Seasons refer to early winter (1 January to 30 April), summer

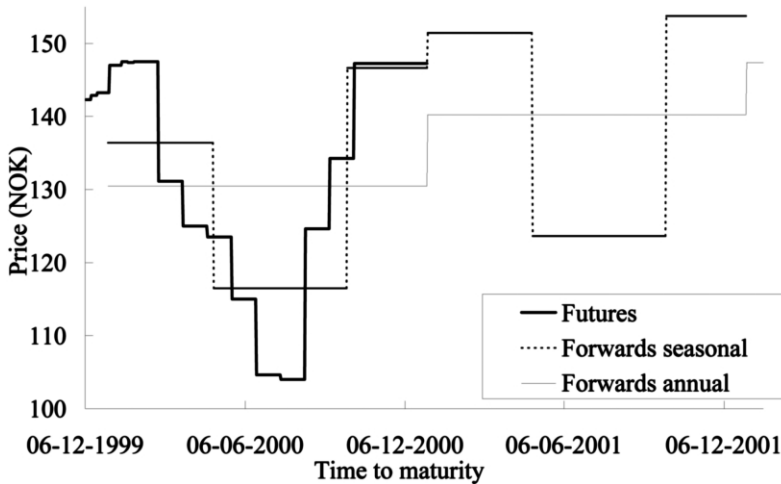


Fig. 4. The term structure of futures and forward prices at Nord Pool on 6-12-1999.

(1 May to 30 September) and late winter (1 October to 31 December). In contrast to futures, the forward contracts are not split into smaller blocks but are traded until the beginning of the delivery period. The fact that contracts are traded in relatively large chunks means that the problem of finding prices for specific maturity times or in general constructing a high-resolution term-structure curve is significant.

The volume has grown extensively at Nord Pool in the last couple of years.² However, many contracts have had small trading volumes. For example, the contract for delivery of power in the period 04-01-1999 to 25-04-1999 had zero turnovers at Nord Pool in 58 out of the 71 days when it was possible to trade the contract. Still, there is an active over-the-counter market where roughly three quarters of all contracts are traded; Nord Pool has a 25% market share for financial products.

3. Model

The described shortcoming of available market data indicates a possible gain from combining these data with forecasts from bottom-up models. In the Scandinavian market, seasonal variation is an important characteristic. However, throughout the year products with maturities exceeding 6–12 months are only represented by 1–3 products per year, as seen in Fig. 4. Clearly, the available market data provide relatively good information about the short end of the term structure, but a much less detailed picture of the long end and the important seasonal component.

A commonly used bottom-up model in the Nordic region is the MPS model (Botnen et al., 1992). This model consists of a number of interconnected subsystems

² 'The total volume of power contracts traded on Nord Pools financial market in January 2001 was 101.8 TWh. This is nearly quadruple the volume traded in January 2000. The monthly volume exceeded the entire annual volume traded on the Exchange in 1998' (Nor, 2001).

(regions). Each region consists of a single-node electricity market where production, consumption and exchange (transmission) with adjacent areas are modeled. A solution of the model results in a set of equilibrium prices and production quantities, for each week over the time horizon considered (usually 3 years), and for each historical inflow year (usually 70). Important input to the model, when determining future spot prices, include demand and its dependence on temperature, fuel costs, exchange with other countries (import/export prices and quantities) and initial reservoir and snow accumulation levels. Since the MPS model explicitly takes into account the dynamic stochastic behavior of demand and of detailed hydropower production, we believe that one of the model's strengths is the ability to capture seasonal variations in prices (Haugstad and Rismark, 1998), at least *within* summer and winter seasons, and in the transition between these.

To illustrate how an improved term structure curve can be constructed by including information from a bottom-up model we propose a small optimization model based on the following criteria: First, since we view market data as providing fundamental information we impose a strong constraint on the relationship between optimized prices and the bid/ask prices observed in the market. However, to properly express this constraint we need the theoretical relationship between forward/futures contract that entails delivery over a period $[T_1, T_2]$ and contracts with delivery in the subperiods of $[T_1, T_2]$, where each subperiod is a day or a week. Denote f_t the price of the forward contract with delivery in day or week t , and let $F(T_1, T_2)$ be the price of the forward contract with delivery in the interval $[T_1, T_2]$. This latter contract is a portfolio of basic forward contracts f_t and by arbitrage arguments its price must be:

$$F(T_1, T_2) = \sum_{t=T_1}^{T_2} \frac{1}{\sum_{t=T_1}^{T_2} e^{-rt}} \sum_{t=T_1}^{T_2} e^{-rt} f_t \quad (1)$$

Therefore, the forward price $F(T_1, T_2)$ can be seen as a weighted average of a series of forward prices f_t over the interval $[T_1, T_2]$. Since we do not observe exact prices but rather a bid/ask spread in the market we replace the equality condition with the following:

$$F(T_1, T_2)_{\text{bid}} \leq \frac{1}{\sum_{t=T_1}^{T_2} e^{-rt}} \sum_{t=T_1}^{T_2} e^{-rt} f_t \leq F(T_1, T_2)_{\text{ask}} \quad (2)$$

The proposed optimization model uses a daily or weekly resolution and we require as the set of strong constraints that the optimized prices do not violate Eq. (2) for any of the products observed in the market. Letting P equal the number of observed futures/forwards and N the number of days or weeks in the period spanning the entire term structure, we have a model with a set of $2P$ constraints and N variables for any given point in time. Many of the contracts are overlapping, and our model will not find a feasible solution if there are arbitrage opportunities among such overlapping contracts.

The second criterion is due to the degrees of freedom for optimization, which are induced by the bid/ask spread and the large blocks in the long end of the forward price curve. The aim is to capture information about the shape of the seasonal variation from the bottom-up model forecasts, and intuitively this leads us to look for a function or curve that has the same shape of seasonal characteristics as the bottom-up forecast. A simple objective function that will serve this purpose is a constrained least square (LSQ) curve fitting. Using this procedure, the model simply minimizes the squared differences between the decision variables and the bottom-up forecast values subject to the bid/ask constraints described above. Squared differences are clearly preferable over absolute differences because we seek to fit the term structure to the shape of the bottom-up forecast but generally at a different level.³

We expect forward contract with neighboring delivery points to exhibit relatively small price differences. In other words, we expect trading to create a smooth forward curve, therefore, we add a smoothing term to the objective function. In the absence of such a term there is no mechanism in the model to prevent large jumps in the forward curve, e.g. on the transition from the end of the summer contract delivery and the beginning of the late winter contract delivery. This combination of smoothing and LSQ constitutes the basic version of our model stated below.

The MPS model usually predicts less seasonality than can be observed in the futures and forward market, i.e. the difference between summer and winter prices is less in the MPS forecasts than what is observed in the derivative market. However, this does not mean that our approach will lead to a too small difference between summer and winter prices. The constraints Eqs. (1) and (2) ensure that the difference between average summer prices and average winter prices is the same as is observed in the contract market.

This illustrates an important feature of the model, namely that the absolute level of the bottom-up forecasted prices or prior periodic function used, do not matter for the generated forward curve. Again, this is due to the fact that we constrain the forward prices implied by the generated curve to be within the relatively narrow bid and ask spread of the observed (given) market prices. This means, for example, that the bottom-up forecast does not have to be an unbiased forecast of future spot prices, nor does it have to be adjusted for risk.

MODEL: Term Structure Generation (TSG)

$$\text{Minimize } \sum_{t=1}^T (f_t - B_t)^2 + \lambda \sum_{t=2}^{T-1} (f_{t-1} - 2f_t + f_{t+1})^2 \quad (3)$$

³ Squaring the differences will increase the penalty for large differences in relative terms and it will therefore tend to equalize differences at all points. Clearly, if the price level dictated by the strong bid/ask constraints lie far from those dictated by the MPS model, LSQ fitting will generally cause a parallel shift in the curve, whereas absolute differences will be more inclined towards also changing the shape of the curve.

Subject to

$$F(T_{1i}, T_{2i})_{\text{bid}} \leq \frac{1}{\sum_{t=T_{1i}}^{T_{2i}} e^{-rt}} \sum_{t=T_{1i}}^{T_{2i}} e^{-rt} f_t; \quad \forall i \in \mathcal{I} \quad (4)$$

$$F(T_{1i}, T_{2i})_{\text{bid}} \geq \frac{1}{\sum_{t=T_{1i}}^{T_{2i}} e^{-rt}} \sum_{t=T_{1i}}^{T_{2i}} e^{-rt} f_t; \quad \forall i \in \mathcal{I} \quad (5)$$

The following notation is used for parameters, variables and indices:

Indices and sets:

T : The final day or week over the entire period.

\mathcal{I} : The set of products with observed market prices.

T_{1i} : Beginning of delivery period for product i in \mathcal{I} .

T_{2i} : End of delivery period.

Parameters:

B_t : Price in bottom-up forecast for day/week t .

$F(T_{1i}, T_{2i})_{\text{bid}}$: Bid price for product i .

$F(T_{1i}, T_{2i})_{\text{ask}}$: Ask price for product i .

λ : Weight that scales smoothing term relative to LSQ term.

Variables:

f_t : Basic forward price with maturity t .

The end point of the curve ($t=T$) should exhibit seasonality. However, if the weight on smoothing is high compared to the weight on LSQ, then the end point will tend to be shaped as a straight line. To avoid this problem, the generated points near the end of the curve is encouraged in the objective function via an additional term penalizing deviation between the shapes of the end of the last two years constrained to equal the first and second derivative at the previous year (i.e. at $T-365$ for daily resolution). Thus, we assume that the shape of the curve at the end point is well represented by the shape 1 year before.

4. Experimental results

Figs. 5 and 6 illustrate forward curves generated for an example where observed bid/ask prices on 01-09-1998 was chosen as apriori data and an average scenario from a simulation with the MPS model was used as the bottom-up data. From Fig. 5 it is clear that the forward price curves generated with LSQ, i.e. $\lambda=0$, have inherited the discontinuities in the market data.

Fig. 6 shows the same simulation except that the weight of the smoothing parameter has been set so that the two terms in the objective function are of the same order of magnitude. It is clear from the figure that the smoothing parameter

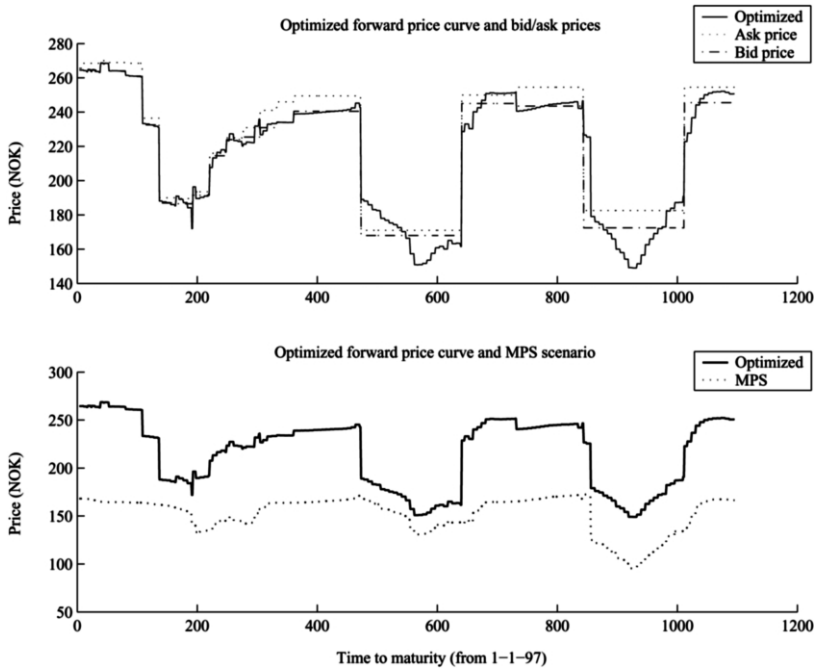


Fig. 5. Forward price curves generated using only least squares in the objective.

effectively removes the undesirable jumps and provides a more realistic picture of the forward price curve.

Testing the quality of the generated forward price curves is not a straightforward task. As the curves themselves represent approximations to realized term structures there are no empirical data to test them against, and the quality should, therefore, ideally be assessed by testing it in the context in which the curves are to be applied. One such application could be as input to a factor model of the type described in Section 2. In order to test the quality of the generated forward price curves against other forms of input data⁴ one could choose a specific factor model and compare output in the form of price forecasts (in-sample and out-of sample analysis of predictive power, see, e.g. Schwartz (1997)) and goodness of fit tests. However, construction of a factor model would exceed the scope of this paper.

Since the TSG model is basically a hybrid where information from bottom-up forecasts supplement market data, it is interesting to examine whether or not the information from the bottom-up model does in fact capture the shape of the seasonal variation. One-way to measure this is to test the generated curves against the price that would have been set by the market, if all the points on the curve where traded

⁴ Other candidates can be market prices, simple periodic functions fitted to market prices and bottom-up forecasts.

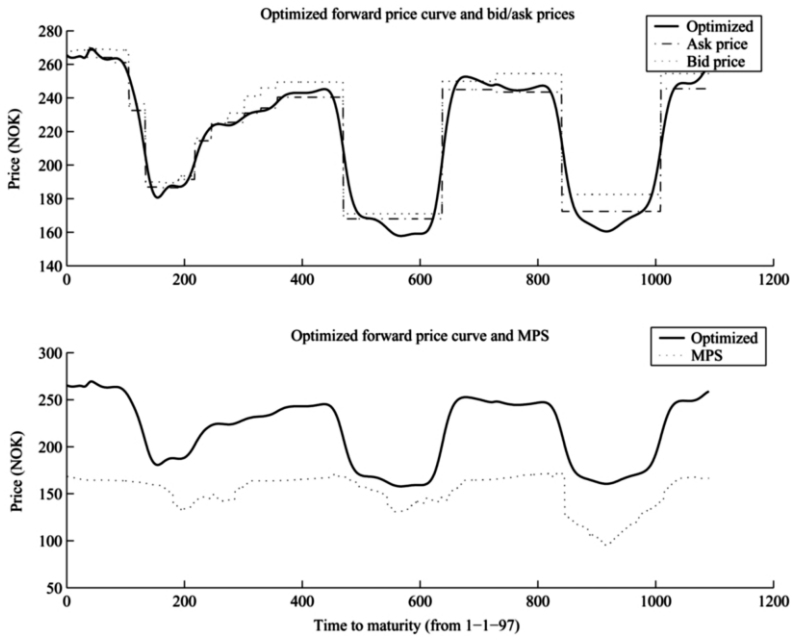


Fig. 6. Forward price curves generated using a combination of least squares and smoothing as the objective.

products.⁵ We do not have market prices beyond those already used to generate the curves, but at dates where new products are listed, i.e. when a season contract is split into 4-week blocks or when a 4-week block is split into weekly contracts we gain new information overnight or over the weekend. This information can be used by comparing the prices that the model generated at the last day before the split, with the actual market prices at the first day after the split. Therefore, to make such a comparison we must in principle be able to extrapolate the forward price curve from one day to another.

We give a short example to clarify. Consider a forward price curve generated based on all products traded at the market on the date 18-04-1997. On the subsequent trading day 21-04-1997, the seasonal product S01-98 ceases to exist and is replaced by four new 4-week products B01-98 to B04-98. The price of these four new products provides information about the seasonal shape within the delivery period previously represented only by the seasonal product S01-98. We now wish to examine how the generated forward price curve based on the data available at 18-04-1997, would have priced the products B01-98 to B04-98. We know the price

⁵ Clearly, one could imagine using the presented framework to beat the market by including new information from bottom-up models thought to be superior to the market. However, here we have restricted ourselves to an analysis of the degree of agreement between market data and the generated curves.

changes for all the products that existed on both dates and possible strategies, therefore, include: (a) to extrapolate using an average of price changes taken over all traded products, (b) to find a functional relation across maturities and then extrapolate the desired interval or (c) to extrapolate using an average of the price changes in products with maturities close to the new products. In the following example we tried all three strategies, but found little or no change compared to the reported results based on no extrapolation.

For comparison with the TSG model, we fitted to market data both a curve generated using only the smoothness part of the objective function and a truncated Fourier series on the form:

$$g(t) = \alpha + \gamma t + \sum_{j=1}^J (\beta_{1,j} \cos(\omega t) + \beta_{2,j} \sin(\omega t)) \quad (6)$$

To approximate the true seasonality and simultaneously avoid overfitting, we set $J=2$. The maximum smoothness approach can be viewed as a best alternative if bottom-up models are considered to hold no information. The truncated Fourier series was included because it is often used in factor models to approximate the seasonal variation.

The forward and futures data used is from Nord Pool, with observations for each trading day in the period 25-09-1995 to 05-04-2002. Four MPS forecasts were used as input. Each column in Table 1 shows the date at which a seasonal product is split into 4-week products (column 1), the number of 4-week products created at that date (column 2) and the average pricing error with respect to the generated prices and the new market prices of the 4-week products (columns 3–5). The error definition is adopted from Bliss (1997) and implies that there is no error if the model prices are within the observed bid–ask spread:

$$\varepsilon = \begin{cases} \frac{F_{\text{ask}} - P}{F_{\text{ask}}} & \text{if } P > F_{\text{ask}} \\ 0 & \text{if } F_{\text{bid}} \leq P \leq F_{\text{ask}} \\ \frac{F_{\text{bid}} - P}{F_{\text{bid}}} & \text{if } P < F_{\text{bid}} \end{cases} \quad (7)$$

where P is the fitted contract price. Each error is calculated as the square root of the average square error, where the average is taken over all new contracts introduced at the date indicated in column 1 (and for all dates in the bottom row). The average errors found using the TSG and maximum smoothness are not significantly different, but the TSG approach seems to perform best. Fig. 7 shows each of the three approaches applied to the date 08-10-1999 where a seasonal product is split into six 4-week products. The curve found by the TSG model captures the bottom peak of 4-week market data more accurately than the smoothness curve. Though this indicates that useful information has been extracted from the MPS forecast, one can also notice that some inaccurate information causes the TSG prices to be lower than

Table 1

Comparing the three approaches with market prices of newly listed block products, i.e. having a four week delivery period. Each percentage error shown is an average for the 3 to 6 new products. The total averages are calculated as the average of errors for each of the 87 monthly products

Date	New prod. (%)	TSG model (%)	Max. smooth (%)	Trunc Fourier (%)
02-01-1996	3	2.62	2.88	1.96
22-04-1996	4	2.43	2.45	6.32
07-10-1996	6	12.10	14.07	14.65
30-12-1996	3	2.67	3.08	5.40
21-04-1997	4	0.78	0.90	2.35
06-10-1997	6	3.60	2.91	4.49
30-12-1997	3	0.27	1.86	1.67
20-04-1998	4	0.77	1.29	2.41
05-10-1998	6	3.04	3.62	5.32
04-01-1999	3	0.65	0.66	0.18
26-04-1999	4	0.00	1.73	0.64
11-10-1999	6	3.05	3.55	4.38
03-01-2000	3	1.17	1.30	0.69
26-04-2000	4	1.76	0.68	3.06
09-10-2000	6	3.79	4.03	7.71
02-01-2001	3	0.50	0.38	1.52
23-04-2001	4	3.30	2.85	5.79
08-10-2001	6	2.62	3.41	5.74
02-01-2002	3	2.98	2.99	1.67
Total	87	2.70	2.87	4.01

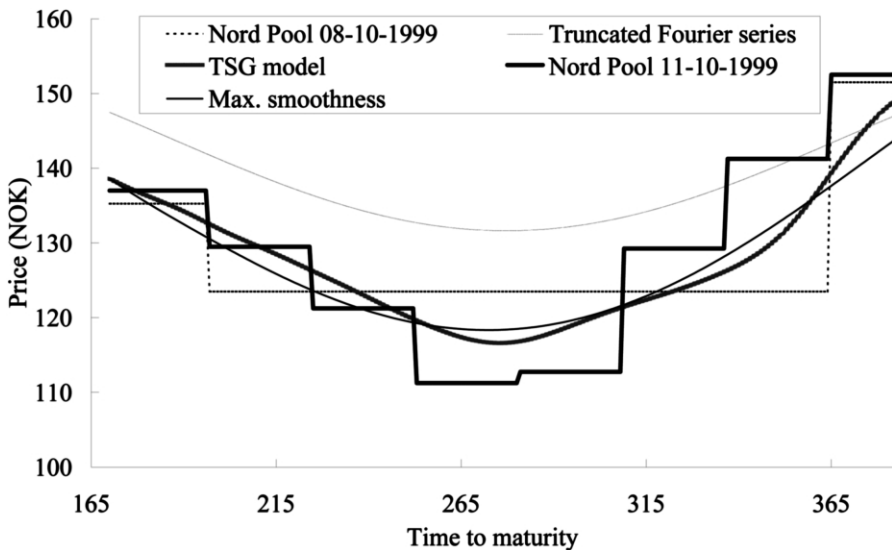


Fig. 7. Pricing the monthly products created on 11-10-1999.

Table 2

Summary of pricing errors in NOK/MWh. RMSE is the square root of the average square error (each trading day is given equal weight)

RMSE (NOK/MWh)	TSG model $\lambda = 500$	TSG model $\lambda = 1000$	Trunc. Fourier	Max. smoothness
New weeks	5.068	5.058	4.774	5.066
New blocks	7.053	7.410	10.01	8.260
Overlapping forwards	4.766	4.760	5.449	4.819

4-week market prices in the region $t=330$ to $t=350$. In this specific case the average error of the TSG model is, therefore, approximately equal to that of the maximum smoothness model.

A summary of the root mean square pricing errors of the TSG model with two different smoothness weights, the truncated Fourier series and the maximum smoothness approach are listed in Table 2. Pricing errors for ‘new weeks’ pertain to splitting of 4-week products into products with delivery period length 1 week. ‘New blocks’ pertain to pricing of newly listed 4-week products. ‘Overlapping forwards’ pertain to pricing of contracts with long delivery periods that overlap with other products—in this case these long delivery period products are excluded from the model. Thus, all errors are out of sample. We see that the Fourier series approach works well in the very short end of the forward curve, the first 4–7 weeks. The TSG model performs no better than the maximum smoothness approach for this part of the forward curve. For maturities 26–52 weeks into the future (new blocks) the differences are greater, the TSG approach performs best. For the overlapped seasonal and annual forward contracts the TSG and maximum smoothness perform somewhat better than the Fourier approach.

5. Conclusions

In this paper we have suggested a new approach for generating forward price curves based on a combination of market data and forecasts from bottom-up models. We have illustrated how to apply the suggested framework using the Scandinavian electricity market as a case study.

The ability to obtain an acceptable approximation in the form of a high-resolution forward price curve is relevant not only for the pricing of simple electricity derivatives such as futures and forwards but also to extract information about volatility and seasonal variation. Finally, our approach enables the decision maker to include additional information not contained in the market or perhaps simply not revealed by the market.

We examine the model’s quality by its ability to price non-traded maturities. By assuming a simple functional relationship along the time to maturity dimension of the forward price curve we showed that at maturities of approximately 26–52 weeks one captures the view of the market far better by using our model than by fitting a

truncated function to the curve. Our model also performs slightly better than a maximum smoothness model.

Acknowledgments

Lemming is supported by the Nordic Energy Research program. Fleten acknowledges partial support from ECs 5th framework programme project IST-1999-12088. Nord Pool is the source of market data, and Nordenfjellske Energi and SINTEF Energy Research provided MPS model forecasts.

References

- Adams, K.J., van Deventer, D.R., 1994. Fitting yield curves and forward rates with maximum smoothness. *J. Fixed Income* 4, 52–62.
- Bjerksund, P., Rasmussen, H., Stensland, G. (2000). Valuation and risk management in the Norwegian electricity market, Discussion paper 20/2000, Norwegian School of Economics and Business Administration.
- Black, F., 1976. The pricing of commodity contracts. *J. Financ. Econ.* 3, 167–179.
- Bliss, R.R., 1997. Testing term structure estimation methods. *Adv. Futures Options Res.* 9, 197–231.
- Botnen, O.J., Johannesen, A., Haugstad, A., Kroken, S., Froystein, O. (1992). Modelling of hydropower scheduling in a national/international context. In: Broch E., Lysne D. (Eds.), *Proceedings of the Second International Conference on Hydropower (Hydropower '92)*, Balkema, Rotterdam, pp. 575–584.
- Clewlow, L., Strickland, C. (1999). Valuing energy options in a one factor model fitted to forward prices. Working paper, School of Finance and Economics, Technical University of Sydney.
- Haugstad, A., Rismark, O. (1998). Price forecasting in an open electricity market based on system simulation, In: *Proceedings from EPSOM '98*, Zurich, Switzerland.
- Heath, D., Jarrow, R., Morton, A., 1992. Bond pricing and the term structure of interest rates. *Econometrica* 60, 77–106.
- Lucia, J.J., Schwartz, E.S., 2002. Electricity prices and power derivatives: evidence from the Nordic power exchange. *Rev. Derivatives Res.* 5 (1), 5–50.
- Nord Pool (2001), Press Release 04/2001: greater liquidity on the Nordic Power Exchange. <http://www.nordpool.com/information/index.html>.
- Schwartz, E.S., 1997. The stochastic behavior of commodity prices: implications for valuation and hedging. *J. Finance* 52 (3), 923–973.