International Journal of Forecasting

Special Issue on Energy Forecasting

Introduction

James W. Taylor

Saïd Business School, University of Oxford

and

Antoni Espasa

Universidad Carlos III Madrid

International Journal of Forecasting, 2008, Vol. 24, pp. 561-565.

Address for Correspondence:

James W. Taylor
Saïd Business School
University of Oxford
Park End Street
Oxford OX1 1HP, UK

Tel: +44 (0)1865 288927
Fax: +44 (0)1865 288805
Email: james.taylor@sbs.ox.ac.uk

Introduction

1. Background

It has been well established for many years that electricity load forecasts are required for the efficient scheduling and control of power systems. In the early days of load forecasting, the methods used involved simple statistical procedures with considerable heuristic or judgmental input. In the 1980's, sophisticated time series models started to be used by some of the system operators, and these approaches, although not always published, were to become pioneering works in this field. At that time, there was little competition in the market and hence little incentive to improve the quality of the procedures used for demand prediction. This situation changed completely in the early 1990's when liberalisation took place in the energy markets of many countries. Industries, once viewed as natural monopolies, now function around competitive markets with many participants. This has prompted an increased interest in predicting energy demand and the development of methods for energy price forecasting.

Energy market participants include system operators, producers, retailers and traders. These parties have diverse objectives and require specific types of forecasts. The result is that in the energy industry there is a real need for different kinds of forecasts which must be based on different information sets and different types of models. To understand the types of models that can be used, it is important to be aware of the behaviour of electricity consumption, which shows huge fluctuations when recorded daily, hourly, half-hourly or minute-by-minute. A promising aspect in forecasting such data is that much of the variation is caused by consumer habits determined by social regulations, conventions and meteorological factors, and consequently the use of quantitative models with explanatory variables can provide accurate forecasts of this highly oscillating consumer behaviour. The causal factors lead to a number of stylized facts regarding energy demand. These include some degree of trend;

several seasonalities of daily, weekly and annual periodicities; complex calendar effects which include weekends, holidays, vacation periods and other special days; and a non-linear dependence with meteorological variables, which could change depending on the day of the week, season of the year or consumer habits. The change in consumer behaviour can also affect the seasonalities and calendar effects.

In general, for prediction of load beyond a few hours ahead, weather variables are used. However, in order to capture the weather effects and the stylised facts listed above, a great variety of modelling approaches are possible. This is apparent from the demand forecasting papers in this issue of the journal. A multi-equation model with time-varying parameters is developed by Virginie Dordonnat, Siem Jan Koopman, Marius Ooms, Alain Dessertaine and Jérôme Collet (2008–this issue); models with different regimes, which are established in accordance with experience, are presented by José Ramón Cancelo, Antoni Espasa and Rosmarie Grafe (2008–this issue); Luiz Felipe Amaral, Reinaldo Castro Souza and Maxwell Stevenson (2008-this issue) consider nonlinear time series models that use exogenous variables in the specification of the nonlinear structure; and artificial neural networks (ANN's) is the subject of the paper by Alexandre Alves de Silva, Vitor Ferreira and Roberto Velasquez (2008-this issue). The papers by James Taylor (2008-this issue) and Lacir Soares and Marcelo Medeiros (2008–this issue) describe how univariate methods can be useful when either the lead time is just a few hours or less, or load forecasts are required for regions where weather predictions are not available. The paper by Marek Brabec, Ondřej Konár, Emil Pelikán and Marek Malý (2008–this issue) broadens the scope of this special issue by looking at the problem of forecasting gas consumption for individual customers.

From the supply side of the energy sector, renewable sources of energy are receiving considerable attention. Scientific evidence indicates that the current level of emissions of greenhouse gases into the atmosphere is leading to climate change. In many countries, ambitious targets are already in place for energy to be generated from renewable sources. This produces new challenges for forecasters. An example is the prediction of wind power, which is the fastest growing source of renewable energy. In contrast to more traditional sources of energy, wind power fluctuates greatly, which has significant cost implications for its integration into the electricity system. These costs will be reduced if wind power forecast accuracy can be improved. The forecasts are required at the wind farm level, and are usually needed online using automated procedures. This is the subject of the papers by Ismael Sánchez (2008–this issue) and René Jursa and Kurt Rohrig (2008–this issue).

Although electricity price series exhibit some of the stylized facts seen in load data, they also possess some very noticeable differences. Series of electricity spot price are typically characterised by relatively high volatility, mean-reverting spikes, and a distribution that exhibits skewness and fat tails. The presence of such features means that it is not straightforward to apply, to electricity prices, models developed for load and for more traditional financial market prices. Nevertheless, some of the ideas from those areas are being adapted for electricity prices, and this can be seen in the price forecasting papers in this issue. Due to the complex distributional properties of spot prices, the three price forecasting papers in this issue, by Anastasios Panagiotelis and Michael Smith (2008–this issue), Kam Fong Chan, Philip Gray and Bart van Campen (2008–this issue), and Rafal Weron and Adam Misiorek (2008–this issue), include not only point prediction, but also volatility, interval or density forecasting. Density forecasting in the fourth paper, by Nektaria Karakatsani and Derek Bunn (2008–this issue), is less straightforward because it involves a rich specification of fundamental variables. The benefit of including exogenous variables is an interesting issue for spot price forecasting studies.

As is clear from the above discussion, research in modelling and forecasting in the energy sector, especially electricity price forecasting, is evolving at a fast pace. This situation has motivated this special issue on energy forecasting. Several of the papers in this special issue were presented at the International Institute of Forecasters Workshop on Energy Forecasting that took place in Rio de Janeiro in January 2007. These, along with independently submitted papers, were evaluated through the journal's normal peer-review process. The resulting papers fall naturally into three categories: energy demand, wind power and electricity price. In the next section, we briefly describe each paper in the order in which they appear in this special issue.

2. The papers in this special issue

2.1. Energy demand

Developments in forecasting methodologies are continuously reflected in new contributions to the literature on electricity demand forecasting. The six electricity demand forecasting papers in this special issue illustrate the current variety of methods available. The seventh paper included in this section concerns gas consumption, which is an area that has received relatively little attention from forecasting researchers.

A common approach to modelling hourly electricity load is to specify a separate weather-based regression model for each hour of the day. The model presented by Virginie Dordonnat, Siem Jan Koopman, Marius Ooms, Alain Dessertaine and Jérôme Collet (2008–this issue) is of this multi-equation type, but it has the key features that the hourly equations are estimated simultaneously, and the parameters in each hourly equation are allowed, not only to be different for each hour, but also to be time-varying. With this dynamic parameter structure, the model can be expressed in a multivariate linear Gaussian state space framework, which leads to the use of the Kalman filter for signal extraction and forecasting. The richness of the model implies a potentially large number of parameters, so the authors propose restrictions on the various covariance matrices in order to limit the parameterisation.

José Ramón Cancelo, Antoni Espasa and Rosmarie Grafe (2008–this issue) describe the load forecasting scheme that they have developed for the Spanish system operator. The paper makes explicit many of the forecasting needs of a national system operator, and provides ways to approach them that could involve the joint use of daily and hourly models. For lead times from one to three days ahead, they use a separate, independently estimated weather-based model for each hour of the day. For longer lead times, a model is built from daily data to deliver a prediction of daily total demand, which is then disaggregated into hourly demand using daily load profiles derived from the hourly models. Notable features of the models are the precise modelling of special days, such as national holidays, and the detailed specification of the nonlinear impact of temperature with different regimes for different periods of the week and year.

The nonlinear impact of temperature on load has very often been captured as a regressor in load forecasting models. In contrast to this, Luiz Felipe Amaral, Reinaldo Castro Souza and Maxwell Stevenson (2008-this issue) work with smooth transition autoregressive models and consider lags of temperature as transition variables. They show that for the Australian half-hourly data that they use, temperature performs better as transition variable than lags of load.

An alternative approach to modelling nonlinearities is provided by ANN's. Indeed, the nonlinear and nonparametric feature of ANN's has led to many researchers and practitioners considering their use for load forecasting. Although the literature has developed greatly since the first ANN load studies of the early 1990's, there is still no clear guidance as to how best to select input variables in an automated way. This issue is the theme of the paper by Alexandre Alves da Silva, Vitor Ferreira and Roberto Velasquez (2008–this issue). They propose two types of automated input selection procedures. The first ignores the ANN structure and estimation algorithm used, and works only on the set of potential inputs. By contrast, the second type, which is far more computationally demanding, aims to assess the usefulness of a set of inputs with respect to a specific ANN.

Although not the main subject of their study, an ANN is included as a benchmark method in the empirical analysis of Lacir Soares and Marcelo Medeiros (2008–this issue). They develop univariate models for forecasting the demand for electricity from one day to one week ahead in part of the Brazilian state of Rio de Janeiro. For such lead times, one might expect weather-based models to be used. However, for the region considered in this paper, reliable weather data and forecasts are simply not available. They use a separate, independently estimated model for each hour of the day. Their proposed models are essentially autoregressive, augmented with deterministic components to capture trends, annual seasonality, and special day effects. The emphasis in the paper is on a thorough statistical testing of their model, rather than on the development of a complex model.

Most short-term electricity demand forecasting studies have used hourly or half-hourly data. However, to enable the real-time scheduling of electricity generation, as well as load-frequency control, online load forecasts are required for lead times shorter than 30 minutes. This is the focus of James Taylor (2008–this issue) who uses minute-by-minute load data to compare the accuracy of predictions from various univariate methods, as well as forecasts derived from a weather-based regression method. In contrast to the univariate models of Soares and Medeiros, the substantially shorter lead times considered by Taylor leads him to fit each univariate method to the complete series of demand.

In the only paper in this special issue that is not concerned with electricity data, Marek Brabec, Ondřej Konár, Emil Pelikán and Marek Malý (2008–this issue) address the modelling and prediction of the daily natural gas consumption of individual large commercial customers. The focus on individual customers contrasts with the other demand forecasting papers in this issue, which address the more standard problem of predicting aggregate demand. An example of the use of individual customer forecasts is to enable better prediction of the total consumption of a time-varying portfolio of customers. In this study, consumption is modelled as a mixture involving a logistic regression model for the probability of non-zero consumption, and a nonlinear regression model for the consumption conditional on it being non-zero. A mixed effects framework is used for both models, which enables the modelling of common features in the consumption patterns of the customers, as well as individual customer effects.

2.2. Wind power

The second part of this special issue consists of two papers on wind power forecasting. Ismael Sánchez (2008–this issue) addresses the situation where wind power predictions are available from more than one source. The relative performances of such forecasts can vary from one wind farm to another, can depend on the forecast horizon, and can vary over time. With the aim of developing methods for an online prediction system, Sánchez proposes linear combining methods that allow the combining coefficients to adapt over time in an automated fashion. As an alternative to combining, he develops a procedure for selecting the best of the available forecasts at each time period, which amounts to switching between methods. As the relative performances of the combining and switching methods can vary over time, Sánchez proposes that a number of these methods be implemented in the online system, which in a final step performs a further adaptive selection from among the methods.

Typically, a statistical forecasting model for a wind farm is based on data for that one location. René Jursa and Kurt Rohrig (2008–this issue) explain that, due to the spatio-temporal development of weather fronts over an extended area, there are dependencies between weather and power data of different locations of wind farms. In view of this, the authors allow wind power for a single wind farm to be modelled in terms of wind power and wind speed data for the wind farms in the surrounding region. As this approach is rather data intensive, and involves nonlinear relationships, the authors employ artificial intelligence (AI) methods, including ANN's. The main theme of the paper by Jursa and Rohrig is how best to automate the specification of the model architecture and choice of input variables, which is particularly important with large datasets. This focus on automated specification of AI methods has similarities with the ANN load forecasting paper by Alves da Silva, Ferreira and Velasquez (2008–this issue).

Although the two wind power papers in this special issue concentrate on point forecasting, the content of the papers presented at the 2008 International Symposium on Forecasting in Nice indicated that interest wind power density forecasting is taking off.

2.3. Electricity price

In view of the relatively short history of deregulated electricity markets, it is understandable that the literature on electricity spot price forecasting is not as developed as that of electricity demand prediction. The third and final part of this special issue consists of four papers that contribute to the price forecasting literature.

Anastasios Panagiotelis and Michael Smith (2008–this issue) consider the modelling of hourly spot price data. Their univariate model involves a vector autoregression (VAR), where the vector is comprised of the 24 hourly log price observations for each day. The use here of a separate equation for each hour of the day seems natural for electricity prices, given the popularity of such multi-equation formulations for modelling hourly electricity demand. Exogenous variables are included to control for deterministic trend and seasonal effects. The number of parameters is reduced substantially by imposing sparse forms on the variance matrix of the model errors and on the VAR coefficient matrix. An important and novel feature of the model is that a multivariate skewed t distribution is used for the model errors. The model is estimated using Markov chain Monte Carlo, which through Bayesian posterior inference delivers density forecasts for the spot price.

Univariate modelling of spot prices is also the theme of the paper by Kam Fong Chan, Philip Gray and Bart van Campen (2008–this issue). They apply ideas recently developed for the analysis of the volatility in more traditional financial markets to the analysis of the variation in a series of half-hourly spot prices. Using quadratic variation theory, financial econometricians have devised a nonparametric method for separating realized financial volatility into a discontinuous jump component and a smooth, continuous non-jump component. This decomposition of the variation is appealing because it allows the distinct dynamics in the two sources of variation to be modelled separately with the hope of improving forecast accuracy. The method relies on the assumption that the process has a zero mean, which for electricity prices is certainly not the case. To overcome this, Chan, Gray and van Campen apply the method to the residuals from a model for the mean of the electricity price series.

Rafal Weron and Adam Misiorek (2008–this issue) provide an empirical comparison of the point and interval forecasts produced by a number of models using two hourly time series of spot prices. They consider a variety of autoregressive models, including forms of spike-preprocessed models and regime switching models. An interesting finding was that models with nonparametric disturbances tended to perform better than alternatives, including Gaussian or heavy-tailed innovations. In contrast to the two price forecasting papers in this issue that we have discussed so far, Weron and Misiorek consider the inclusion of exogenous variables in their models. More specifically, they evaluate the usefulness of including temperature and system load, and find that the latter seems to be of more use than the former.

A thorough consideration of exogenous variables is the theme of the paper by Nektaria Karakatsani and Derek Bunn (2008–this issue). They suggest that a model for spot price that does not invoke market fundamentals and agent behaviour may be inadequate for forecasting. Their view is that, although univariate statistical models may replicate well the characteristics of the historical data, they cannot anticipate the abrupt, fast-reverting spikes, which they reason can only be pre-signalled by predictions of relevant exogenous drivers. They consider the day-ahead forecasting of British spot prices, and propose fundamental price models estimated independently for each half-hour of the day. Their empirical study supports the use of a time-varying parameter regression model that allows for a smoothly evolving impact of the various fundamental variables. The dynamic structure for the parameters has similarities with the electricity demand forecasting study of Dordonnat, Koopman, Ooms, Dessertaine and Collet (2008–this issue)

3. Summary

In terms of methodology, there is great variety amongst the papers of this special issue. The methods considered include univariate, multivariate, regression, ARIMA, VAR, Bayesian, ANN, state-space, regime switching, smooth transition, mixed effects, exponential smoothing, heuristics, semi-parametric, spatio-temporal, time-varying parameters, and various mixtures and combinations of methods. The data used includes minute-by-minute, half-hourly, hourly and daily with lead times ranging mainly from a few minutes to 10 days. In view of the title and readership of the journal, it seems appropriate that the authors of the papers in this special issue, and the data that they analyse, have a very international feel. In our opinion, the papers in this issue provide a very useful contribution to the literature on the forecasting of energy demand, energy prices and power generation from renewable sources. However, there is clearly great potential for further work in each of these areas. In view of this, perhaps the most important contribution of this special issue is that it will hopefully stimulate new research on energy forecasting. Finally, we should say that we are extremely grateful to the diligent work of the referees, whose comments and criticisms greatly assisted the authors and ourselves.