Energy Economics 67 (2017) 60–71

Contents lists available at ScienceDirect

Energy Economics

journal homepage: www.elsevier.com/locate/eneco

Forecasting quantiles of day-ahead electricity load

Z. Li[a,*], A.S. Hurn[b], A.E. Clements[b]

[a] School of Civil Engineering and Built Environment, Queensland University of Technology, Brisbane, QLD 4000, Australia
[b] School of Economics and Finance, Queensland University of Technology, Brisbane, QLD 4000, Australia

ARTICLE INFO

Article history:
Received 20 February 2017
Received in revised form 21 July 2017
Accepted 3 August 2017
Available online 17 August 2017

JEL classification:
Q41
Q48

Keywords:
Load forecasting
Quantile forecasts
Bayesian quantile regression

ABSTRACT

Accurate load forecasting plays a crucial role in the decision making process of many market participants, but probably is most important for the dispatch planning of an electricity market operator. Despite the competitive forecast accuracy achieved by existing point forecast models, point forecasts can only provide limited information relating to the expected level of future load. To account for the uncertainty of future load, and provide a more complete picture of the future load conditions for dispatch planning purposes, quantile forecasts can be useful. This paper proposes a computationally efficient approach to forecasting the quantiles of electricity load, which is then applied to forecasting in the National Electricity Market of Australia. The proposed model performs competitively in comparison with one industry standard and two recently proposed quantile forecasting methods. One of the main advantages of the proposed approach is the ease with the number of covariates can be expanded. This is a particularly important feature in the context of load forecasting where large numbers of important drivers are usually necessary to provide accurate load forecasts.

1. Introduction

Accurate load forecasting is a crucial contributor to the decision processes of the electricity market regulator and other electricity market participants primarily because load prediction provides an indication of the future state of the demand side of the market. The literature on load forecasting has concentrated mainly on producing accurate point forecasts (Harvey and Koopman, 1993; Hippert et al., 2001; Espinoza et al., 2005; Cancelo et al., 2008; Amaral et al., 2008; Taylor, 2012; Clements et al., 2016). While this type of forecast does provide valuable information, it only gives an incomplete snapshot of the future state of the market. For example, from the perspective of a market operator, the main goal of dispatch planning is to meet future demand at minimum cost, while at the same time maintaining adequate generation capacity to cover (without disruption) all possible supply-side conditions. In this context, reliance on point forecasts is an extremely risky strategy, especially during peak periods when load is highly volatile. To prevent possible shortages and associated disruptions to electricity dispatch, at the very least it would be natural to consider a number of different possible future demand scenarios. In other words, a well-designed dispatch plan requires the regulator to balance the operational cost with the security of supply. To achieve this balance, it is important to not only have knowledge of the expected level of future demand as provided by a point forecast, but also to be aware of the uncertainty or variability of future demand.

A separate strand of the load forecasting literature, attempts to forecast the *uncertainty* of future load. The most straightforward method to achieve this objective is to use simulation methods to construct interval forecasts of load (Fan and Hyndman, 2012). By contrast, Engle et al. (1992), Hyndman and Fan (2010) and Sigauke and Chikobvu (2011) forecast the maximum load in a day or a week. Forecasting maximum load provides additional information about the upper tail behaviour of the load distribution. However, this approach is still technically a point forecast and is produced using traditional mean regression models. Recognising the fact that forecasting load uncertainty is still relatively underdeveloped, this paper seeks to broaden the menu of options available to the applied researcher or policymaker. To this end, the idea of forecasting quantiles of future load directly using quantile regression is explored in detail.

Although the application of quantile regression to forecasting quantiles of various economic variables is not new, it has rarely been applied in the context of load forecasting. Instead, quantile forecasts of load are normally produced by means of simulation or bootstrapping (McSharry et al., 2005; Fan and Hyndman, 2012). One reason for the relative paucity of studies using quantile regression is that it can be computationally expensive, even in models with a small number of parameters. This problem is particularly acute in the context of load forecasting where complex periodicity together

* Corresponding author.
E-mail address: zili.li@connect.qut.edu.au (Z. Li).

http://dx.doi.org/10.1016/j.eneco.2017.08.002
0140-9883/

with temperature effects and differential load patterns on holidays and weekends means that a large number of covariates need to be included in order to achieve satisfactory forecasting accuracy.

In order to address the computational challenge posed by the large number of covariates in quantile forecasting models, several remedies have been suggested. López Cabrera and Schulz (2014) divide the quantile regression problem into multiple steps and treat the load quantiles as a combination of many principal components with the weights for these components being modelled using autoregressive models. Liu et al. (2017) adopt a forecast combination technique in which quantile regression is performed on point forecasts from a number of models, instead of directly on a large number of covariates. However, since these point forecasts differ only because of differences in the specification of the models, the quantile forecasts only account for the uncertainty in model specification and cannot capture the uncertainty of future load itself conditional on all observed covariates. One of the most promising avenues of recent research uses a technique borrowed from the machine learning literature. Taieb et al. (2016) apply a quantile gradient boosting algorithm that allows the number of covariates used in the load forecast to grow in an optimal manner.

This paper suggests a different solution to the parameter estimation problem in quantile regressions with a high dimensional covariate vector. Specifically, a Bayesian estimation approach is proposed which a large number of covariates can be accommodated quite easily. Two models, which differ only in the assumed form of the underlying distribution of the disturbance terms, are specified and compared. The first model assumes that the disturbances follow an Asymmetric Laplace distribution (ALD) which is conveniently represented as a mixture of a normal and an exponential distribution. This assumed parametric form for the disturbances makes the model relatively easy to estimate (Kozumi and Kobayashi, 2011). The second model treats the distribution of the disturbances in a nonparametric way using a mixture of kernel functions, controlled by a Dirichlet process (Kottas and Krnjajić, 2009). Forecasts for these two models are then compared to a number of existing approaches and also with the quantile forecasts of load reported by the Australia Energy Market Operator (AEMO).

The results of the forecasting comparison may be summarised succinctly as follows. The ALD model is found to perform better than the non-parametric model and this approach is then compared to the combination of point forecasts method suggested by Liu et al. (2017) and the gradient boosting method of Taieb et al. (2016). The ALD based model easily outperforms the combination of forecasts approach but is only marginally superior to the gradient boosting method. Given the large number of arbitrary tuning parameters required to be chosen by the researcher in order to operationalise the gradient boosting algorithm and the importance of these choices to the performance of the algorithm, the Bayesian quantile regression model based on the ALD should be considered as an important addition to the load forecasting arsenal.

Finally, a brief word on the form of the loss functions used in the empirical work of the paper is in order. The loss metrics adopted here, namely, a pinball loss function and a quantile coverage rate, are assumed to be time independent. In other words, there are no differential penalties on forecast errors during peak and off-peak periods. Any time-variation in the loss function would reflect a prior position on the forecaster's attitude toward risk. In this sense, the loss functions used here take an agnostic view on the decision maker's attitude to risk.

2. The modelling framework

To produce quantile load forecasts, quantile regression techniques such as that developed by Koenker and Bassett Jr. (1978), may be applied to the same model specification that is commonly used for making forecasts of the conditional mean. One of the most widely used model specifications is the autoregressive moving average (ARMA) form. Let y_{it} denote the load in half hour i on day t, a basic ARMA model specification for load would then take the form

$$y_{i\ t} = \theta_{0\ i} + \theta_{1\ i} y_{i\ t-1} + \theta_{2\ i} y_{i\ t-7} + \theta_{3\ i} y_{(-t)} + \theta_{4\ i} \varepsilon_{i\ t-1} + \theta_{5\ i} \varepsilon_{i\ t-7} + \varepsilon_{i\ t}, \tag{1}$$

in which $\varepsilon_{i\,t}$ is a disturbance term and the notation $y_{(-t)}$ refers to the latest observed load available at the time of making the forecast. The presence of estimated lagged residuals has proved to be particularly useful in forecasting load (Sigauke and Chikobvu, 2011; Kim, 2013; Clements et al., 2016), since they allow the forecast to be adjusted based on prediction errors from the previous intervals. In the quantile regression context, the only difference is that the residual terms are obtained conditional on quantile estimates of all the parameters. The inclusion of the lagged residual terms can be considered as a special case of the dynamic quantile specification of Engle and Manganelli (2004), in which the lagged quantiles enter the model specification.

In addition to the lag structure in Eq. (1), a good load forecasting model should also account for other factors that may have a significant impact on the variation of load. Paramount among these are the seasonality of load, temperature and special day effects. Incorporating these factors in Eq. (1), the model specification used in this study for the load in half hour i on day t is

$$\begin{aligned} \mathrm{y}_{i\ t} \ = \ & \theta_{0\ i} + \sum_{p=1}^{7} \theta_{1\ i\ p} \mathbb{W}_{t\ p} \mathrm{y}_{i\ t-1} + \theta_{2\ i\ 0} y_{i\ t-7} + \theta_{3\ i} y_{(-t)} + \theta_{4\ i} \varepsilon_{i\ t-1} \\ & + \theta_{5\ i} \varepsilon_{i\ t-7} + \varepsilon_{i\ t} \\ & + \sum_{r=1}^{4} \left[\theta_{2\ i\ r\ 1} \sin\left(2r\pi\left(\frac{t}{365}\right)\right) + \theta_{2\ i\ r\ 2} \cos\left(2r\pi\left(\frac{t}{365}\right)\right)\right] \mathrm{y}_{i\ t-7} \\ & + \sum_{k=1}^{2} \left(\theta_{6\ i\ k\ 1} \mathbb{H}_{i\ t\ k} + \theta_{6\ i\ k\ 2} \mathbb{C}_{i\ t\ k} + \theta_{6\ i\ k\ 3} \mathbb{H}_{i\ t-1\ k} + \theta_{6\ i\ k\ 4} \mathbb{C}_{i\ t-1\ k}\right) \\ & + \sum_{j=1}^{6} \left(\theta_{7\ i\ j\ 1} \mathbb{S}_{t\ j} + \theta_{7\ i\ j\ 2} \mathbb{S}_{t-1\ j}\right). \end{aligned} \tag{2}$$

In this specification, the effect of lagged load is critically dependent on the day of the week. Accordingly, the parameter $\theta_{1\,i}$ on lagged load is allowed to dependent on day-of-the-week dummy variables $\mathbb{W}_{t\,p}$ with $p = 1, \ldots, 7$. The variable $y_{(-t)}$ is context specific and is determined by the time at which the forecast is made. Specifically it relates to the last observed load that may properly be included in the forecasting equation. For example, if the forecast is to be made from 04:00 hours (as is the practice of the Australian market operator) then the observed load in the half hour immediately prior to this cut-off is available for use in forecasting for the entire subsequent 24 hour period.

To accommodate the annual pattern of load induced by seasonal weather changes, the parameter $\theta_{2\,i}$ is specified in terms of Fourier polynomials as

$$\theta_{2\ i} = \theta_{2\ i\ 0} + \sum_{r=1}^{4} \left[\theta_{2\ i\ r\ 1} \sin\left(2r\pi\left(\frac{t}{365}\right)\right) + \theta_{2\ i\ r\ 2} \cos\left(2r\pi\left(\frac{t}{365}\right)\right)\right].$$

Four terms in the Fourier expansion are found to be adequate to model the annual cycle.