

Triple Seasonal Methods for Short-Term Electricity Demand Forecasting

James W. Taylor

Saïd Business School

University of Oxford

European Journal of Operational Research, 204, 139-152, 2010.

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

Triple Seasonal Methods for Short-Term Electricity Demand Forecasting

Abstract

Online short-term load forecasting is needed for the real-time scheduling of electricity generation. Univariate methods have been developed that model the intraweek and intraday seasonal cycles in intraday load data. Three such methods, shown to be competitive in recent empirical studies, are double seasonal ARMA, an adaptation of Holt-Winters exponential smoothing for double seasonality, and another, recently proposed, exponential smoothing method. In multiple years of load data, in addition to intraday and intraweek cycles, an intrayear seasonal cycle is also apparent. We extend the three double seasonal methods in order to accommodate the intrayear seasonal cycle. Using six years of British and French data, we show that for prediction up to a day-ahead the triple seasonal methods outperform the double seasonal methods, and also a univariate neural network approach. Further improvement in accuracy is produced by using a combination of the forecasts from two of the triple seasonal methods.

Keywords: Forecasting; Electricity demand; Seasonality; Exponential smoothing; ARMA.

1. Introduction

The efficient management of power systems requires accurate forecasts of electricity demand (load) for a range of lead times. To enable real-time scheduling, system operators produce forecasts for short lead times using online automated methods. In deregulated energy markets, such predictions are also of importance to market participants to support energy transactions (Bunn, 2000).

Electricity demand is often modelled in terms of weather variables (e.g. Hor et al., 2005; Cancelo et al., 2008). However, univariate methods are frequently considered to be sufficient for short lead times because the weather variables tend to change in a smooth fashion over short time frames, and this will be captured in the demand series itself. Furthermore, weather-based online systems require default procedures in order to ensure robustness (Bunn, 1982). In a recent empirical study (Taylor, 2008), a univariate method was shown to outperform a multivariate method up to about four hours ahead, and a combination of forecasts from the two methods was found to be the most accurate approach up to a day ahead. This shows that univariate methods have a valuable role to play in short-term load forecasting. Of course, such methods are the only option when forecasting load in locations where weather forecasts are either unavailable or too costly (Soares and Medeiros, 2008). In this paper, we focus on the use of univariate methods for prediction up to one day ahead.

A variety of univariate methods have been used for short-term load forecasting. Seasonal ARIMA models of various types have been considered by, among others, Hagan and Behr (1987), Moghram and Rahman (1989), Mbamalu and El-Hawary (1993), Taylor and McSharry (2007) and Soares and Medeiros (2008). The approach of Park et al. (1991) mixes autoregressive modelling with exponential smoothing and exponentially weighted least squares. The latter is also known as general exponential smoothing, and this was the focus of the early study of Christiaanse (1971). Given its widespread use in many business forecasting applications (see Hyndman et al., 2008), it is surprising that exponential

smoothing had not received more attention from load forecasting researchers prior to the study by Taylor (2003). Although the nonlinear and nonparametric features of neural networks are attractive for modelling load in terms of weather variables (see Hippert et al., 2001), their appeal for univariate modelling is less obvious and rests on the possible existence of nonlinearities in the structure of the load time series. Pai and Hong (2005) implement an ARIMA model and a neural network as benchmarks in their study of support vector machines, which also enable nonlinear regression modelling. A univariate method based on principal component analysis is considered by Taylor and McSharry (2007). The approach aims to simplify the forecasting task by extracting the important underlying seasonal components in the time series.

A feature of time series of intraday electricity demand is the presence of intraday and intraweek seasonal cycles. In two recent empirical studies, several univariate statistical methods were considered that aim to accommodate both of these seasonal cycles (Taylor and McSharry, 2007, Taylor, 2008). Three of the more successful methods were a double seasonal ARMA model, an adaptation of Holt-Winters exponential smoothing for double seasonality, and an exponential smoothing method recently proposed by Gould et al. (2008). Those empirical studies used intraday time series each of length 30 weeks. With multiple years of observations, in addition to intraday and intraweek cycles, a further seasonal cycle is evident, namely the annual cycle. In this paper, our primary contribution is the extension of the three double seasonal methods to include the intrayear seasonal cycle. Although these extensions for triple seasonality are relatively simple, they have not previously been considered in the literature.

Of course, there is no guarantee that allowing for the intrayear cycle in a statistical model will lead to improved accuracy for prediction up to just one day ahead. Therefore, an important contribution of this paper is the empirical study in which we use two time series of half-hourly demand to compare post-sample forecasting results for single, double and triple

seasonal formulations of the methods. The empirical results provide insight for the number of seasonal cycles to include in a method, for the relative importance of the different types of seasonal cycles, and for the choice between the two exponential smoothing methods and ARMA modelling. The study also evaluates the benefit in incorporating a residual autocorrelation term in the exponential smoothing method of Gould et al.

Research in this area would seem to be timely, given the increasing availability of intraday data in a variety of other applications, such as traffic management and call centre staff scheduling (see, for example, Lam et al., 2006; Taylor, 2008). Another potential application area for the methods in this paper is the modelling of intraday electricity prices, which is of primary importance for trading electricity. The methods would have to be carefully adapted for prices because they typically possess rather complex structure with substantial volatility and mean-reverting spikes.

In Section 2, we describe the two load series. Section 3 presents ARMA model formulations for data with single, double and triple seasonality. The focus of Section 4 is exponential smoothing formulations for single, double and triple seasonality based on the Holt-Winters method. In Section 5, the method of Gould et al. (2008) is presented, along with its extension for triple seasonality. In Section 6, we describe the neural network that we use as a sophisticated benchmark in our empirical comparison of methods. Section 7 presents our empirical study, which evaluates forecast accuracy for lead times from one half-hour ahead up to one day ahead. In Section 8, we summarize and provide concluding comments.

2. The British and French load series

In this study, we consider two load series; one is for Great Britain and the other is for France. Each consists of six full years of half-hourly observations for electricity demand from 2001 to 2006, inclusive. Our empirical analysis used the first five years of data to estimate forecasting method parameters and the remaining year to evaluate post-sample

forecast accuracy. This gave 87,648 half-hourly observations for estimation, and 17,520 for evaluation. The two series are shown in Figure 1. Both show a strong intrayear seasonal cycle with demand much lower during the warmer months of each year. A closer look at these warmer months in the French series reveals a period of four weeks of particularly low demand, which corresponds approximately to August. Although this is not apparent from a visual inspection of the British series in Figure 1, the British transmission company does try to model the effect of the three summer weeks when a large amount of industry closes.

----- Figures 1 and 2 -----

Figure 2 shows a winter fortnight and a summer fortnight from both series. Within each fortnight, the weekdays show similar patterns of demand, but the Saturdays and Sundays are somewhat different. These features are typical of series of intraday electricity demand. Interestingly, for both the British and French data, the seasonality is considerably different for the winter and summer. This indicates that, when modelling the seasonality, the pattern should be allowed to change across the year.

In each series, we identified a number of unusual days, termed ‘special’ days, on which demand differed considerably from the regular seasonal pattern. These were either public holidays or days adjacent or close to public holidays. In short-term online prediction systems, these days are typically replaced by forecasts prepared offline. As our study is concerned with online prediction, prior to fitting and evaluating the forecasting methods, we chose to smooth out the special days leaving the natural periodicities of the data intact. In our study, this smoothing was performed by simple averaging procedures, which usually involved replacing demand on each special day period by the mean of the demand in the corresponding periods of the two adjacent weeks. In our post-sample evaluation of the forecasting methods, we did not include the special days occurring in the evaluation period.

Our notation for the lengths of the intraday, intraweek and intrayear cycles is s_1 , s_2 and s_3 , respectively. As our data in this paper is half-hourly, we set $s_1=48$ and $s_2=336$. In our

modelling of the intrayear cycle, for most periods we used $s_3=52 \times 336$. However, there were 53 weeks between the periods in March 2001 and March 2002 when the clocks went forward by one hour. In view of this, we set $s_3=53 \times 336$ for all the periods of the week before and the week after the clock-change period in March 2002. For a similar reason, we set $s_3=53 \times 336$ for all the periods of the week before and the week after the clock-change period in October 2004. For all other periods in our sample, we considered the intrayear cycle to consist of 52 weeks. Note that, in the French series, the time between the onset of the particularly low demand weeks in August was 52 weeks throughout the six years of the series.

3. Seasonal ARMA

3.1. Single seasonal ARMA for the intraweek cycle

Although the literature contains models that aim to capture both the intraday and intraweek seasonal cycles in intraday load data, it could be argued that, if the intraweek cycle is modelled well, there is no need for additional terms to capture the intraday cycle. In view of this, one could opt to use the standard ARMA model for a single seasonal cycle with that cycle specified as the intraweek cycle as in expression (1).

$$\phi_p(L)\Omega_{P_2}(L^{s_2})(y_t - a - bt) = \theta_q(L)\Psi_{Q_2}(L^{s_2})\varepsilon_t \quad (1)$$

y_t is the load in period t ; a is a constant term; b is the coefficient of a linear deterministic trend term; ε_t is a white noise error term; L is the lag operator; and ϕ_p , Ω_{P_2} , θ_q and Ψ_{Q_2} are polynomial functions of orders p , P_2 , q and Q_2 , respectively. We considered lag polynomials up to order three.

The plot of the French series in Figure 1 suggests that the series may be trending. We performed Dickey-Fuller tests for stationarity for each of the two load series. For the French series, after controlling for a deterministic linear trend, there was no significant evidence of a stochastic trend. This prompted us to include the deterministic trend term shown in

expression (1). We experimented with differencing the data before fitting each of the ARMA models considered in this paper, but this led to noticeably poorer forecasting results than the use of the deterministic trend. Another transformation that we considered was logarithmic, however, this did not lead to improved forecast accuracy. For the British data, we found no significant evidence of stochastic or deterministic non-stationarity, so we removed the trend term from the model by setting $b=0$.

We estimated all of the ARMA models considered in this paper using the standard procedure of maximum likelihood. We based model selection on the Schwarz Bayesian Criterion, with the requirement that all parameters were significant (at the 5% level).

3.2. Double seasonal ARMA for the intraday and intraweek cycle

Box, Jenkins and Reinsel (1994, p. 333) write that the standard ARMA model for single seasonality can be extended to multiple seasonal cycles. An ARMA model capturing both the intraday and intraweek seasonal cycles can be expressed as

$$\phi_p(L) \Phi_{P_1}(L^{s_1}) \Omega_{P_2}(L^{s_2})(y_t - a - bt) = \theta_q(L) \Theta_{Q_1}(L^{s_1}) \Psi_{Q_2}(L^{s_2}) \varepsilon_t \quad (2)$$

In comparison with the ARMA model for single seasonality of expression (1), there are two new terms in the double seasonal formulation. These terms are Φ_{P_1} and Θ_{Q_1} , which are polynomial functions of orders P_1 and Q_1 , respectively. These additional lag polynomials enable an ARMA modelling of the intraday cycle. As in the ARMA model for single seasonality, we considered lag polynomials up to order three.

3.3. Double seasonal ARMA for the intraweek and intrayear cycle

In view of the clear intrayear seasonal cycle present in both of the intraday load series of Figure 1, it seems natural to try to include this seasonal cycle in an ARMA model. This is the approach taken by Cancelo et al. (2008) as part of their modelling of hourly Spanish

electricity consumption. They use a separate model for each hour of the day, which eliminates the need to model the intraday seasonal cycle. Their study is concerned with lead times of one day to one week ahead. For the shorter lead times in our paper, the use of separate models for each hour of the day has much less appeal as the models would not capture the level of the load at, or just prior to, the forecast origin. For their daily data, Cancelo et al. apply an ARMA model for double seasonality that incorporates the intraweek and intrayear seasonal cycles. In this paper, we model the intraweek and intrayear seasonal cycles in our half-hourly data by using the double seasonal ARMA model of expression (2) with s_1 replaced by s_3 . As with the ARMA models of the previous two sections, we considered lag polynomials up to order three.

3.4. Triple seasonal ARMA

A natural extension of the two double seasonal ARMA models of the previous two sections is a model that includes all three seasonal cycles. The result is the ARMA model for triple seasonality presented in the following expression:

$$\phi_p(L) \Phi_{P_1}(L^{s_1}) \Omega_{P_2}(L^{s_2}) \Gamma_{P_3}(L^{s_3})(y_t - a - bt) = \theta_q(L) \Theta_{Q_1}(L^{s_1}) \Psi_{Q_2}(L^{s_2}) \Lambda_{Q_3}(L^{s_3}) \varepsilon_t$$

This formulation is similar to the double seasonal model of expression (2), with the difference being the inclusion of two new terms in the triple seasonal model. These are Γ_{P_3} and Λ_{Q_3} , which are polynomial functions of orders P_3 and Q_3 , respectively. These lag polynomials enable ARMA modelling of the intrayear cycle. As in the ARMA models for single and double seasonality, we considered lag polynomials up to order three.

For the British series, the model with the lowest SBC and with all parameters significant was the following:

$$\begin{aligned} & (1 - 0.10L - 0.40L^2 - 0.43L^3) (1 - 0.41L^{s_1} - 0.10L^{2s_1} - 0.26L^{3s_1}) \\ & \times (1 - 0.63L^{s_2} + 0.08L^{2s_2} - 0.02L^{3s_2}) (1 - 0.36L^{s_3} - 0.21L^{2s_3} - 0.14L^{3s_3}) (y_t - 38,893) \\ & = (1 + 1.13L + 0.73L^2 + 0.15L^3) (1 - 0.17L^{s_1} - 0.20L^{3s_1}) (1 - 0.53L^{s_2} + 0.10L^{2s_2}) (1 - 0.08L^{s_3}) \varepsilon_t \end{aligned}$$

For the French series, the selected model was the following:

$$\begin{aligned} & (1 - 0.39L - 0.62L^2 + 0.02L^3) \left(1 - 0.66L^{s_1} - 0.12L^{2s_1} - 0.09L^{3s_1} \right) \\ & \times \left(1 - 0.62L^{s_2} - 0.05L^{3s_2} \right) \left(1 - 0.26L^{s_3} - 0.17L^{2s_3} - 0.21L^{3s_3} \right) (y_t - 49,424 - 0.06t) \\ & = \left(1 + 0.68L + 0.14L^2 + 0.07L^3 \right) \left(1 - 0.34L^{s_1} - 0.12L^{2s_1} - 0.09L^{3s_1} \right) \left(1 - 0.55L^{s_2} - 0.01L^{3s_2} \right) \varepsilon_t \end{aligned}$$

It is interesting to note that, for both load series, model parameters were statistically significant in at least one of the lag polynomials for each of the three seasonal cycles. This provides in-sample evidence of the importance of modelling all three cycles. It is worth noting that both of these ARMA models for triple seasonality contain 21 parameters, and that this is far more than the number required for the exponential smoothing models presented in the next section.

4. HWT exponential smoothing

4.1. Single seasonal HWT exponential smoothing for the intraweek cycle

As with the ARMA modelling in Section 3, we first present an exponential smoothing formulation that accounts for just the intraweek seasonal cycle. The method of expressions (3) to (5) is based on the Holt-Winters exponential smoothing method.

$$\hat{y}_t(k) = l_t + w_{t-s_2+k} + \phi^k (y_t - (l_{t-1} + w_{t-s_2})) \quad (3)$$

$$l_t = \lambda (y_t - w_{t-s_2}) + (1 - \lambda) l_{t-1} \quad (4)$$

$$w_t = \omega (y_t - l_{t-1}) + (1 - \omega) w_{t-s_2} \quad (5)$$

l_t is the smoothed level; w_t is the seasonal index for the intraweek seasonal cycle; λ and ω are smoothing parameters; and $\hat{y}_t(k)$ is the k step-ahead forecast made from forecast origin t . As it is written, the forecast function in expression (3) is valid for $k \leq s_2$, but it is straightforward to rewrite the expression for longer lead times. The term involving the parameter ϕ , in the forecast function of expression (3), is an adjustment for first-order error autocorrelation. As shown by Taylor (2003), this adjustment substantially improves the performance of the Holt-Winters method for intraday electricity data, and so is integral to the method. We distinguish

the method of expressions (3) to (5) from the standard Holt-Winters method by referring to it as the HWT method for single seasonality. A trend term is not included in any of the exponential smoothing formulations in this paper because, in our empirical work, the inclusion of a trend term resulted in no change in forecast accuracy. We present all exponential smoothing formulations with additive seasonality. Multiplicative seasonality formulations led to similar results to the analogous additive versions.

The initial smoothed values for the level and seasonal components are estimated by averaging the early observations. For all exponential smoothing methods considered in this paper, we constrained the parameters to lie between zero and one, and estimated them by minimizing the sum of squared in-sample forecast errors (SSE). We found that the success of various optimization algorithms depended on the choice of initial values for the parameters. To address this, we followed an optimization procedure similar to that described in Engle and Manganelli (2004). For each method, we first generated 10^4 vectors of parameters from a uniform random number generator between 0 and 1. For each of the vectors, we then evaluated the SSE. The 10 vectors that produced the lowest SSE values were used, in turn, as the initial vector in a quasi-Newton algorithm. Of the 10 resulting vectors, the one producing the lowest SSE value was chosen as the final parameter vector. For the single seasonality formulation of expressions (3) to (5), the optimized parameters are shown in the rows labelled ‘Single - intraweek’ of Table 1. As in the study by Taylor (2003), the values of ϕ are high and the values of λ are low, indicating that the adjustment for first-order autocorrelation has, to a large degree, made redundant the smoothing equation for the level.

----- Table 1 -----

4.2. Double seasonal HWT exponential smoothing for the intraday and intraweek cycle

Taylor (2003) adapts the Holt-Winters method in order to accommodate the intraday and intraweek seasonal cycles in intraday load data. The method, which is presented in

expressions (6) to (9), is an extension of the HWT method described in the previous section.

$$\hat{y}_t(k) = l_t + d_{t-s_1+k} + w_{t-s_2+k} + \phi^k \left(y_t - (l_{t-1} + d_{t-s_1} + w_{t-s_2}) \right) \quad (6)$$

$$l_t = \lambda (y_t - d_{t-s_1} - w_{t-s_2}) + (1 - \lambda) l_{t-1} \quad (7)$$

$$d_t = \delta (y_t - l_{t-1} - w_{t-s_2}) + (1 - \delta) d_{t-s_1} \quad (8)$$

$$w_t = \omega (y_t - l_{t-1} - d_{t-s_1}) + (1 - \omega) w_{t-s_2} \quad (9)$$

d_t is a new term representing the seasonal index for the intraday seasonal cycle; w_t is now defined as the seasonal index for the intraweek seasonal cycle that remains after the intraday cycle is removed; δ is a new smoothing parameter; and the forecast function of expression (6) is valid for $k \leq s_1$. As with the single seasonal HWT method, we estimated the initial smoothed values for the level and seasonal components by averaging the early observations.

The estimated values of the parameters of the method are presented in Table 1, in the rows labelled ‘Double – intraday and intraweek’. These values are similar to those obtained by Taylor and McSharry (2007) for much shorter series. In Table 1, for both series, the values of λ and ϕ are reasonably similar to those of the single seasonal method. However, the value of ω is noticeably lower in the double seasonal method. Our explanation for this is that the inclusion of the intraday cycle provides daily updating of the seasonality in the double seasonal method, with the result that the remaining intraweek seasonality needs less frequent updating than in the single seasonal formulation.

4.3. Double seasonal HWT exponential smoothing for the intraweek and intrayear cycle

In Sections 3.2 and 3.3, we described how the double seasonal ARMA formulation can be used with the seasonal cycles defined as the intraday and intraweek cycles or, alternatively, with the cycles defined as the intraweek and intrayear cycles. We make a similar suggestion here for the double seasonal HWT exponential smoothing method. To include the intraweek and annual cycles in a double seasonal formulation, we propose the use of the method of expressions (6)-(9) with s_1 replaced by s_3 . We experimented with several

different initialisation procedures for the smoothed annual cycle, but we found that none offered noticeable improvement over simply setting all elements of the annual cycle to zero, and so we adopted this approach in this paper. We estimated the initial smoothed values for the level and the intraweek seasonal components by averaging the early observations. The estimated values of the method's parameters are shown in the rows labelled 'Double – intraweek and intrayear' of Table 1. For each series, the value of α is low indicating slow evolution of the intrayear seasonal cycle.

4.4. Triple seasonal HWT exponential smoothing

Expressions (10) to (14) present a new HWT exponential smoothing method designed to capture the intraday, intraweek and intrayear seasonal cycles. This triple seasonal method is a simple adaptation of the double seasonal method of expressions (6) to (9), and involves the inclusion of a third seasonal index and an additional smoothing equation for this index.

$$\hat{y}_t(k) = l_t + d_{t-s_1+k} + w_{t-s_2+k} + a_{t-s_3+k} + \phi^k \left(y_t - (l_{t-1} + d_{t-s_1} + w_{t-s_2} + a_{t-s_3}) \right) \quad (10)$$

$$l_t = \lambda (y_t - d_{t-s_1} - w_{t-s_2} - a_{t-s_3}) + (1 - \lambda) l_{t-1} \quad (11)$$

$$d_t = \delta (y_t - l_{t-1} - w_{t-s_2} - a_{t-s_3}) + (1 - \delta) d_{t-s_1} \quad (12)$$

$$w_t = \omega (y_t - l_{t-1} - d_{t-s_1} - a_{t-s_3}) + (1 - \omega) w_{t-s_2} \quad (13)$$

$$a_t = \alpha (y_t - l_{t-1} - d_{t-s_1} - w_{t-s_2}) + (1 - \alpha) a_{t-s_3} \quad (14)$$

In this formulation, l_t , d_t and w_t have the same definitions as in the double seasonal method of Section 4.2. a_t is defined as the seasonal index for the annual cycle that remains after the intraday and intraweek seasonal cycles have been removed; and α is a smoothing parameter for the new index, a_t . We initialised the annual cycle by setting all elements of it to zero, and we initialised the other seasonal cycles and the level by averaging early observations in the series. The parameters of the method are presented in the rows labelled 'Triple - intraday, intraweek and intrayear' of Table 1. The table shows that, for both series, the parameter values are reasonably similar to those of the double seasonal methods.

5. IC exponential smoothing

5.1. Double seasonal IC exponential smoothing for the intraday and intraweek cycles

Gould et al. (2008) argue that an unappealing feature of the HWT formulation for double seasonality of Section 4.2 is that it assumes that the intraday cycle is the same for each of the seven days of the week. To address this, they allow the intraday cycle for the different days to be represented by different seasonal components. By contrast with double seasonal HWT, there is no representation in the formulation for the intraweek seasonal cycle. Due to its focus on intraday cycles, we term this method ‘intraday cycle (IC) exponential smoothing’.

To ensure a more parsimonious formulation, Gould et al. propose that a common intraday cycle is used for days of the week that exhibit similar patterns of demand. Before presenting the model formulation, let us first consider which days of the week have a similar intraday cycle. To do this, we use the same plot considered for this purpose by Gould et al. The plot for the British data is presented in Figure 3. For each day of the week, the figure shows the average demand for each period of the day, calculated using only the in-sample observations. The figure confirms that Saturday and Sunday possess their own distinct cycles. While Tuesday, Wednesday and Thursday show very similar profiles, the demand on Monday mornings and on Friday afternoons and evenings suggests that these days should each be treated as having distinct intraday cycles. Therefore, we concluded that the British data consists of five distinct cycles: one for each of Friday, Saturday, Sunday and Monday, and a common intraday cycle for the other three days of the week.

----- Figures 3 and 4 -----

Figure 4 presents the plot of daily profiles for the French data. In essence, we can make similar conclusions for the French data to those that we drew for the British data. However, in contrast to the British data, French demand on Fridays is not as distinct from the other weekdays as it is in the British data. It is, therefore, not obvious whether or not to treat

Fridays as having the same intraday cycle as Tuesdays, Wednesdays and Thursdays. The lack of clarity as to how to choose between different specifications of the model is an obvious disadvantage of the method. For the French data, we opted to implement a version of the model with the same five distinct intraday cycles used for the British data, and also a version with the following four distinct cycles: one for each of Saturday, Sunday and Monday, and a common intraday cycle for the other four days of the week. In our empirical analysis, we found that the post-sample forecasting results for the model with five distinct cycles was considerably better than the version with just four. For simplicity, in the remainder of Section 5, we present formulations for the model with five distinct cycles.

In the model formulation, for any period t , the estimated values of the five cycles are given as c_{it} , for $i=1$ to 5. Five corresponding dummy variables are defined as follows:

$$x_{it} = \begin{cases} 1 & \text{if time period } t \text{ occurs in a day of type } i \\ 0 & \text{otherwise} \end{cases}$$

Gould et al. present their approach in the form of a state space model, and we follow this convention in our presentation of the model in the following expressions:

$$\begin{aligned} y_t &= l_{t-1} + \sum_{i=1}^5 x_{it} c_{i,t-s_1} + \varepsilon_t \\ l_t &= l_{t-1} + \lambda \varepsilon_t \\ c_{it} &= c_{i,t-s_1} + \left(\sum_{j=1}^5 \gamma_{ij} x_{jt} \right) \varepsilon_t \quad (i = 1, 2, \dots, 5) \end{aligned}$$

where l_t is the smoothed level, and λ and the γ_{ij} are the smoothing parameters. The γ_{ij} can be viewed as a 5×5 matrix of parameters that enables the five types of intraday cycle to be updated at different rates, which is a nice feature of the model. It also enables intraday cycle of type i to be updated even when the current period is not in a day of type i . Gould et al. propose several alternative restrictions for the matrix of γ_{ij} parameters. We included in our empirical study two forms of the method; one involved estimation of the matrix of γ_{ij} parameters with just the basic constraint that the parameters lie between zero and one, and

the other involved the additional restrictions of common diagonal elements and common off-diagonal elements. In our discussion of the forecasting results in Section 7, we refer to these two forms of the model as unrestricted and restricted, respectively.

As with the double seasonal HWT method, we estimated the initial smoothed values for the level and seasonal components by averaging the early observations. The parameters were estimated using the same procedure employed for the HWT methods, with the one difference being that we started the procedure with 10^5 randomly generated vectors of parameters due to the larger dimension of the parameter vector for the IC exponential smoothing method.

We found that the results were considerably improved with the inclusion of the adjustment for first-order residual autocorrelation that was used in the double seasonal HWT method. In Section 7, we report only the results for this improved form of the method. For clarity, in the following expressions, we present the formulation with this adjustment.

$$y_t = l_{t-1} + \sum_{i=1}^5 x_{it} c_{i,t-s_i} + \phi e_{t-1} + \varepsilon_t \quad (15)$$

$$e_t = y_t - \left(l_{t-1} + \sum_{i=1}^5 x_{it} c_{i,t-s_i} \right) \quad (16)$$

$$l_t = l_{t-1} + \lambda e_t \quad (17)$$

$$c_{it} = c_{i,t-s_i} + \left(\sum_{j=1}^5 \gamma_{ij} x_{jt} \right) e_t \quad (i = 1, 2, \dots, 5) \quad (18)$$

Gould et al. note that the restricted form of the model is identical to the double seasonal HWT method of expressions (6) to (9), provided seven distinct intraday cycles are used and the seasonal component smoothing parameters can be written in terms of the HWT parameters as $\gamma_{i1} = \delta + \omega$ and $\gamma_{ij} = \delta$. The optimized parameters for the restricted form of the IC exponential smoothing method, with five distinct intraday cycles, are shown in the rows labelled ‘Double - intraday and intraweek’ of Table 2. Let us briefly compare the parameters for this model with those in Table 1 for the double seasonal HWT method for intraday and intraweek seasonal cycles. For both series, in the two tables, the values of λ and ϕ are similar,

and the smoothing parameters for the seasonal terms can be written as $\gamma_i \approx \delta$ and $\gamma_j \approx \delta + \omega$. It would, therefore, seem that for both series the restricted IC model with five distinct cycles is very similar to the double seasonal HWT method for intraday and intraweek seasonal cycles, and that we should not be surprised to achieve rather similar results with the two methods.

----- Table 2 -----

In Sections 3 and 4, for both the ARMA and HWT exponential smoothing approaches, we presented two forms of double seasonality, one for the intraday and intraweek cycles and a second form for the intraweek and intrayear cycles. Due to the focus of the IC exponential smoothing method on intraday cycles, it makes no sense to try to consider an IC method formulation that does not try to model the intraday cycle. Therefore, for the IC exponential smoothing method, the only double seasonal formulation is the one presented in expressions (15) to (18).

5.2. Triple seasonal IC exponential smoothing

In the same way that we have extended Taylor's (2003) exponential smoothing method for double seasonality, it is straightforward to extend the exponential smoothing model proposed by Gould et al. (2008). The triple seasonal formulation is given as:

$$\begin{aligned}
 y_t &= l_{t-1} + \sum_{i=1}^5 x_{it} c_{i,t-s_i} + a_{t-s_3} + \phi e_{t-1} + \varepsilon_t \\
 e_t &= y_t - \left(l_{t-1} + \sum_{i=1}^5 x_{it} c_{i,t-s_i} + a_{t-s_3} \right) \\
 l_t &= l_{t-1} + \lambda e_t \\
 c_{it} &= c_{i,t-s_i} + \left(\sum_{j=1}^5 \gamma_{ij} x_{jt} \right) e_t \quad (i = 1, 2, \dots, 5) \\
 a_t &= a_{t-s_3} + \alpha \varepsilon_t
 \end{aligned}$$

In this formulation, most of the terms have the same interpretations as in the double seasonal method of expressions (15) to (18). a_t is a new seasonal index for the annual cycle that remains after the level and intraday seasonal cycles have been removed, and α is a

smoothing parameter for the new index, a_t . The optimized parameters for the restricted form of the triple seasonal IC exponential smoothing method, with five distinct intraday cycles, are shown in the rows labelled ‘Triple – intraday, intraweek and intrayear’ of Table 2. Comparing the parameters for this model with those in Table 1 for the triple seasonal HWT method, we note that the values of λ and ϕ are similar, while the seasonal component smoothing parameters do not closely obey the conditions for equivalence ($\gamma_{i1}=\delta+\omega$ and $\gamma_{ij}=\delta$) discussed at the end of Section 5.1.

6. Artificial Neural Network

As a sophisticated benchmark method, we included a neural network in our empirical study. Darbellay and Slama (2000) and Taylor et al. (2006) build univariate neural network models to model the intraday and intraweek cycles in series of intraday load data. As inputs, they use load at various lags. In this paper, we implement similar neural network models, although ours differ in two significant ways. Firstly, we extend the networks to include the intrayear cycle. Secondly, we use a separate neural network model for each lead time. By contrast, Darbellay and Slama and Taylor et al. use a single model, and generate multi-step-ahead predictions by using a forecast for an earlier lead time as the input when the input corresponds to lagged load occurring after the forecast origin. We followed the advice of a referee by avoiding the use of a single model for all lead times. Using cross-validation with a hold-out sample of one year of the estimation sample, we found that forecast accuracy was improved by applying the differencing operator $(1-L^{s_1})(1-L^{s_2})$ to our series of electricity load prior to neural network modelling.

We used a single hidden layer feedforward network consisting of a set of k inputs, x_{it} , connected to each of m units in a single hidden layer, which, in turn, are connected to a single output variable, y_t . The output was specified as the electricity load variable, differenced as described above, and normalised by subtracting the estimation sample mean and dividing by the

standard deviation. As inputs, we used normalised lagged values of the differenced load variable, with lags chosen to be consistent with the ARIMA modelling. In the neural network for h step-ahead prediction, as inputs, we used the load at the forecast origin and the load at the following lags: 1, 2, s_1-h , $2s_1-h$, $3s_1-h$, s_2-h , $2s_2-h$, $3s_2-h$, s_3-h , $2s_3-h$ and $3s_3-h$. The resultant model is written as

$$f(\mathbf{x}_t, \mathbf{v}, \mathbf{w}) = g_2 \left(\sum_{j=0}^m v_j g_1 \left(\sum_{i=0}^k w_{ji} x_{it} \right) \right)$$

where the functions $g_1(\cdot)$ and $g_2(\cdot)$ are chosen as sigmoidal and linear respectively, and w_{ji} and v_j are the weights. The network weights were estimated using least squares with regularisation parameters, λ_1 and λ_2 , included to penalise complexity and thus avoid overfitting (see Bishop, 1997, §9.2). With n in-sample observations, the minimisation was of the following form:

$$\min_{\mathbf{v}, \mathbf{w}} \left(\frac{1}{n} \sum_{t=1}^n (y_t - f(\mathbf{x}_t, \mathbf{v}, \mathbf{w}))^2 + \lambda_1 \sum_{j=0}^m \sum_{i=0}^k w_{ji}^2 + \lambda_2 \sum_{j=0}^m v_j^2 \right)$$

We performed the minimisation using the popular backpropagation algorithm with learning rate η and momentum parameter μ (see Bishop, 1997, §7.5). To established suitable values for λ_1 , λ_2 , η , μ , m and the choice of lags to use as inputs (see Bishop, 1997, §9.8), we used cross-validation with a hold-out sample of one year of the estimation sample. We did this using the model for one step-ahead prediction, and for simplicity, we used the same set of values and choice of lags for the models for all other lead times. For the British load series, this led to the inclusion of all inputs, and $\lambda_1=0.001$, $\lambda_2=0.001$, $\eta=0.1$, $\mu=0.9$ and $m=24$. For the French data, we also found that all inputs should be included, and we obtained the same parameter values, with the exception of $m=12$. As with the other models fitted in this paper, parameters were estimated using the estimation sample of data for each series, and the parameters were not updated after this.

7. Empirical forecasting study

We evaluated post-sample forecasting accuracy for the ARMA and exponential smoothing methods. In this section, we present the mean absolute percentage error (MAPE) results for half-hourly lead times up to one day ahead calculated for the one-year post-sample period of each of the two series of load data. For most of the methods considered, their relative performances were similar for the two load series, and so, in general, for each method, we present the MAPE results averaged over the two series. We also evaluated the root mean squared percentage error, mean absolute error and root mean squared error, but we do not report these results here because the rankings of the methods for these measures were very similar to those for the MAPE. In Section 7.1, we evaluate the worth in extending the various methods from single to double seasonal formulations, and from double to triple seasonal formulations. In Section 7.2, we compare the results for the ARMA and exponential smoothing methods.

7.1. Evaluating the worth of extending double to triple seasonal formulations

In Figure 5, we present the average MAPE results for the ARMA models at each lead time. The plots show that there is a notable improvement in forecast accuracy when using the double seasonal rather than the single seasonal formulation, and that there is a further noticeable improvement achieved by using the triple seasonal model. A similar story can be seen from Figure 6 for the different versions of the HWT method. Figure 5 shows that, for the ARMA approach, there is little difference between the two double seasonal versions of the method. By contrast, for the HWT method, the double seasonal formulation for intraday and intraweek seasonal cycles is superior to the double seasonal formulation for the intraweek and intrayear cycles.

----- Figures 5 to 7 -----

In their study, Gould et al. omitted the residual autocorrelation adjustment term in the HWT method for double seasonality. They also did not consider the inclusion of such a term in their own IC exponential smoothing method. We implemented the methods with and without the adjustment for residual autocorrelation. Typical results are shown in Figure 7. Note that the results are very similar for both methods when the adjustment is included. The results show that the adjustment term is needed in the HWT method, and that it also leads to substantial improvement in the IC method of Gould et al.

In Section 5.1, we described how, for the British data, five distinct intraday cycles seemed appropriate for the IC exponential smoothing method. For the French data, it was unclear from the plot of intraday profiles in Figure 4 whether four or five distinct cycles were needed in the IC exponential smoothing method. The results of Figure 8 indicate that, for the French series, using five distinct cycles is notably better than using four. With regard to the issue of whether to use the restricted or unrestricted version of the IC method, we found that the unrestricted form was a little better for the British data, while the restricted version was slightly better for the French series. Figure 9 shows that extending the unrestricted and restricted forms of the model to include the annual seasonal cycle leads to increased accuracy.

----- Figures 8 to 10 -----

7.2. Comparing seasonal ARMA and HWT exponential smoothing

We found strong similarity between the results of the best of the double seasonal IC exponential smoothing methods and the analogous double seasonal HWT method, and also strong similarity between their respective triple seasonal versions. In view of this, in this section, we focus on just the HWT and ARMA formulations. Figure 10 shows that double seasonal ARMA outperformed double seasonal HWT exponential smoothing when the intraweek and intrayear cycles were used. With only the intraday and intraweek cycles

incorporated, double seasonal HWT exponential smoothing was a little more accurate than double seasonal ARMA, which is consistent with the findings in Taylor and McSharry (2007). In that study, time series of length just 30 weeks were used, while the results reported in Figure 10 provide insight for the case where multiple years of data are used. Turning to the results for the two triple seasonal methods, we see that the methods performed similarly.

Figure 10 also presents the results of our sophisticated benchmark, the neural network. The method can be seen to produce reasonable results, particularly beyond 12 hours ahead, where it competes well with the double seasonal methods that model the intraweek and intrayear seasonal cycles. However, the method is beaten at all lead times by the two double seasonal methods that model the intraweek and intrayear cycles, and it is comfortably outperformed at all lead times by the triple seasonal methods. Our view is that neural networks are more suited to modelling the nonlinear relationship between load and weather variables, and that univariate load modelling does not require the rich nonlinear modelling capabilities of neural networks. However, although we feel confident that our neural network implementation was reasonable, we accept that the neural network literature contains many different approaches and designs, and that one of these may deliver better accuracy for data such as ours. For example, it could be that accuracy is improved by using a single neural network with 48 outputs, where each corresponds to a different lead time. If we then included, in this new model, all inputs used for our 48 separate models, this would amount to a model with more than 400 inputs. Although this is rather large, by contrast with other modelling approaches, it is within the bounds of technical feasibility for neural networks. Consideration of large numbers of inputs might prompt experimentation with the use of a larger set of inputs for each of the 48 separate models employed in our approach. These are interesting and open research questions. In view of these issues, we avoid drawing firm conclusions from our study regarding the potential for other neural network specifications.

There is a sizeable literature on combining the forecasts produced by different methods. Combining has particular appeal in cases when the methods are based on different information. For example, combining the load forecasts from a weather-based method with those from a univariate method can lead to improved forecast accuracy (see Taylor, 2008). Although ARMA and HWT exponential smoothing for triple seasonality are both univariate, they use the historical data in different ways, so we felt that there may be value in combining. For simplicity, we considered a simple average of the forecasts from the two methods. Figure 10 shows that this led to better MAPE values than those of the two individual methods.

8. Summary and concluding comments

The existing literature has considered double seasonal ARMA and exponential smoothing methods designed in order to capture the intraday and intraweek seasonal cycles present in intraday load data. The strong annual fluctuation, present in multiple years of such data, motivates the modelling of the intrayear cycle. One simple approach to this is to use the double seasonal models with the two seasonalities specified as the intraweek and intrayear cycles. However, a more satisfactory approach is to extend the double seasonal methods so that they capture all three seasonal cycles: intraday, intraweek and intrayear. To ignore the annual cycle is likely to result in some degree of misspecification of the other components of the method. Furthermore, explicitly modelling features of the series that repeat on an annual basis enables these features to be anticipated, leading to improved forecast accuracy at all lead times. After smoothing public holidays from our data, such features included the industrial holiday periods and the periods around clock-change.

Although there may be intuitive appeal in modelling the annual seasonal cycles, its benefit for short-term prediction is ultimately a matter for empirical testing. Our empirical study involved two series consisting of six years of half-hourly electricity demand. Consistent with previous studies involving much shorter series, our results show that the

single seasonal versions of HWT exponential smoothing and ARMA were outperformed by the double seasonal formulations. Of particular interest was the finding that the triple seasonal versions of all the methods were more accurate than the double seasonal methods. For both series, there was little difference between the results of the ARMA and exponential smoothing methods for triple seasonality, and we found that a simple average combination led to the greatest accuracy. Comparing exponential smoothing methods, we found that the best performing versions of the IC method were no better than the HWT method.

Although forecast accuracy is of great importance, it is not the only criterion to consider when selecting a forecasting method. In comparison with the HWT method, ARMA models have the disadvantage of requiring substantial specification and a more demanding optimisation due to the far larger number of parameters. The issue of specification is also a potential obstacle to the implementation of the IC exponential smoothing method. For this method, there is currently no clear way to decide upon the number of distinct cycles and the choice between unrestricted and restricted versions of the method.

Acknowledgements

We would like to thank Matthew Roberts of National Grid for supplying the British data, and RTE for making the French data available on their website. We are also grateful for the useful comments of two anonymous referees.

References

- Bishop CM. Neural Networks for Pattern Recognition, Oxford University Press: Oxford; 1997.
- Box GEP, Jenkins GM, Reinsel GC. Time series analysis: Forecasting and control, third edition. Prentice Hall: New Jersey; 1994.

Bunn DW. Short-term forecasting: A review of procedures in the electricity supply industry, *Journal of the Operational Research Society* 1982; 33; 533-545.

Bunn DW. Forecasting loads and prices in competitive power markets, *Proceedings of the IEEE* 2000; 88; 163-169.

Cancelo JR, Espasa A, Grafe R. Forecasting from one day to one week ahead for the Spanish system operator, *International Journal of Forecasting* 2008; 24; 588-602.

Christiaanse WR. Short-term load forecasting using general exponential smoothing, *IEEE Transactions on Power Apparatus and Systems* 1971; PAS-90; 900-910.

Darbellay GA, Slama M. Forecasting the short-term demand for electricity – Do neural networks stand a better chance?, *International Journal of Forecasting* 2000; 16; 71-83.

Engle RF, Manganelli S. CAViaR: Conditional autoregressive value at risk by regression quantiles, *Journal of Business and Economic Statistics* 2004; 22; 367-381.

Gould PG, Koehler AB, Ord JK, Snyder RD, Hyndman RJ, Vahid-Araghi F. Forecasting time-series with multiple seasonal patterns, *European Journal of Operational Research* 2008; 191; 207-222.

Hagan MT, Behr SM. The time series approach to short-term load forecasting, *IEEE Transactions on Power Systems* 1987; PWRS-2; 785-791.

Hippert HS, Pedreira CE, Souza RC. Neural networks for short-term load forecasting: a review and evaluation, *IEEE Transactions on Power Systems* 2001; 16; 44-55.

Hor C-L, Watson SJ, Majithia S. Analyzing the impact of weather variables on monthly electricity demand, *IEEE Transactions on Power Systems* 2005; 20, 2078-2085.

Hyndman RJ, Koehler AB, Ord JK, Snyder RD. *Forecasting with exponential smoothing: The state space approach*. Springer-Verlag: Berlin Heidelberg; 2008.

Lam WHK, Tang YF, Chan KS, Tam M-L. Short-term hourly traffic Forecasts using Hong Kong Annual Traffic Census, *Transportation* 2006; 33; 291-310.

Mbamalu GAN, El-Hawary ME. Load forecasting via suboptimal seasonal autoregressive models and iteratively reweighted least squares estimation, *IEEE Transactions on Power Systems* 1993; 8; 343-348.

Moghran I, Rahman S. Analysis and evaluation of five short-term load forecasting techniques, *IEEE Transactions on Power Systems* 1989; 4; 1484-1491.

Pai P-F, Hong W-C. Support vector machines with simulated annealing algorithms in electricity load forecasting, *Energy Conversion and Management* 2005; 46; 2669-2688.

Park JH, Park YM, Lee KY. Composite modeling for adaptive short-term load forecasting, *IEEE Transactions on Power Systems* 1991; 6; 450-7.

Soares LJ, Medeiros MC. Modeling and forecasting short-term electricity load: A comparison of methods with an application to Brazilian data, *International Journal of Forecasting* 2008; 24; 630-644.

Taylor JW. Short-term electricity demand forecasting using double seasonal exponential smoothing, *Journal of the Operational Research Society* 2003; 54; 799-805.

Taylor JW. An evaluation of methods for very short-term load forecasting using minute-by-minute British data, *International Journal of Forecasting* 2008; 24; 645-658.

Taylor JW. A comparison of time series methods for forecasting intraday arrivals at a call center, *Management Science* 2008; 54; 253-265.

Taylor JW de Menezes LM, McSharry PE. A comparison of univariate methods for forecasting electricity demand up to a day ahead, *International Journal of Forecasting* 2006; 22; 1-16.

Taylor JW McSharry PE. Short-term load forecasting methods: An evaluation based on European data, *IEEE Transactions on Power Systems* 2007; 22; 2213-2219.

Figure 1 Electricity load in France and Great Britain for 2001 to 2006.

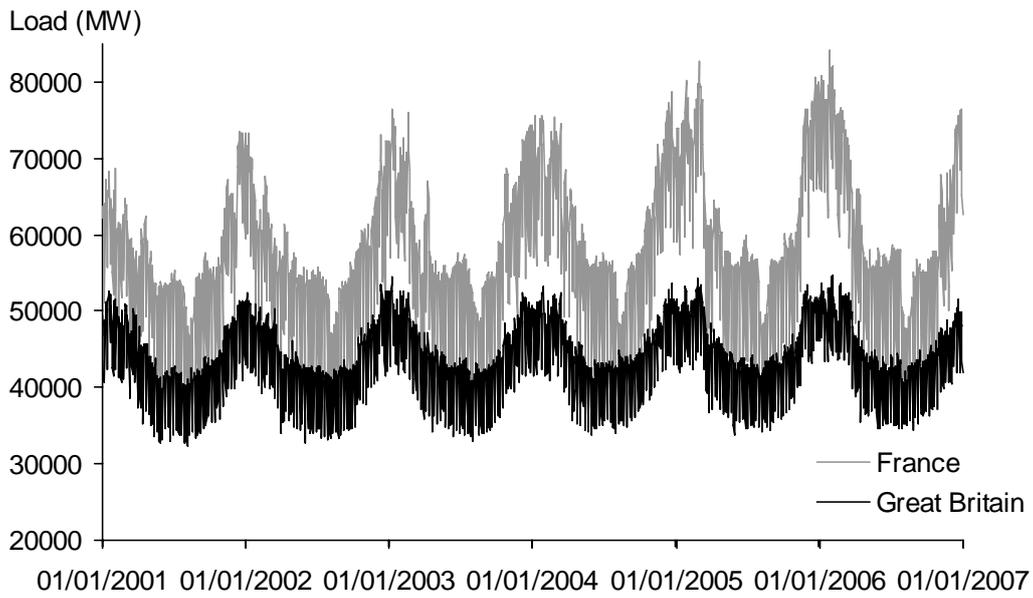


Figure 2 Half-hourly load in France (primary y-axis) and Great Britain (secondary y-axis) for a winter fortnight (8 to 21 January 2006) and a summer fortnight (16 to 29 July 2006).

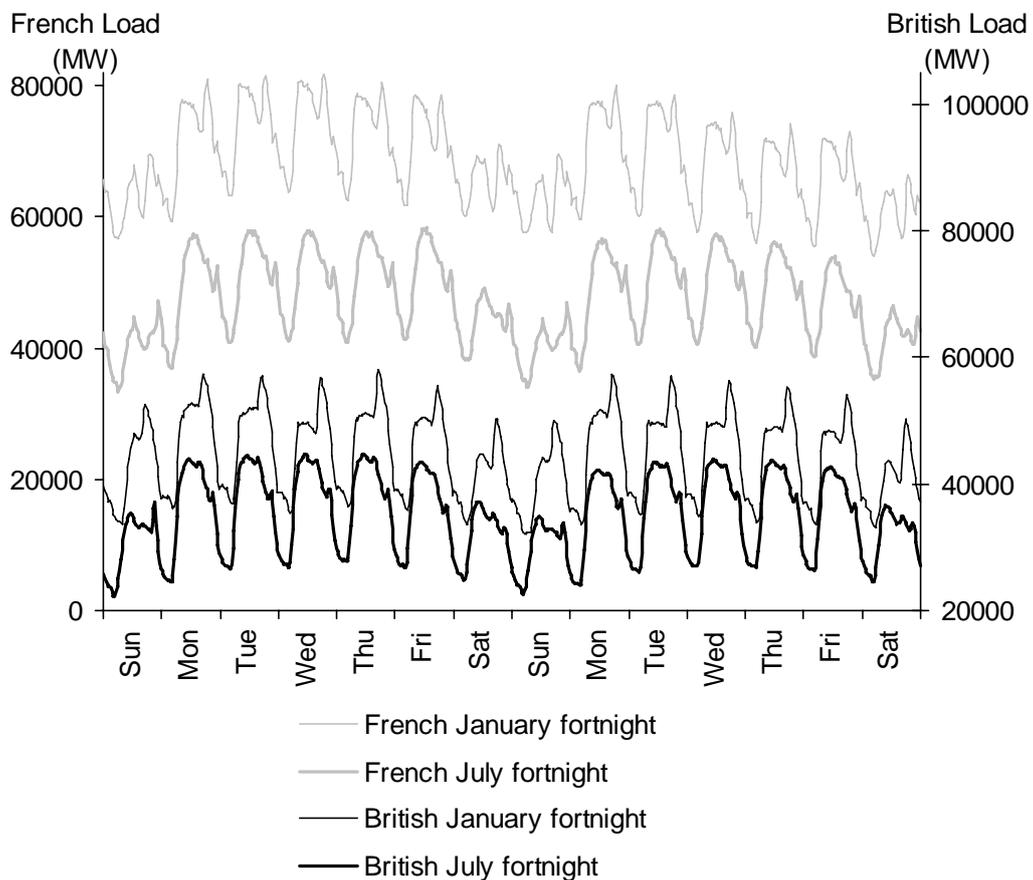


Figure 3 Average intraday cycle for each day of the week constructed using only in-sample observations for Great Britain.

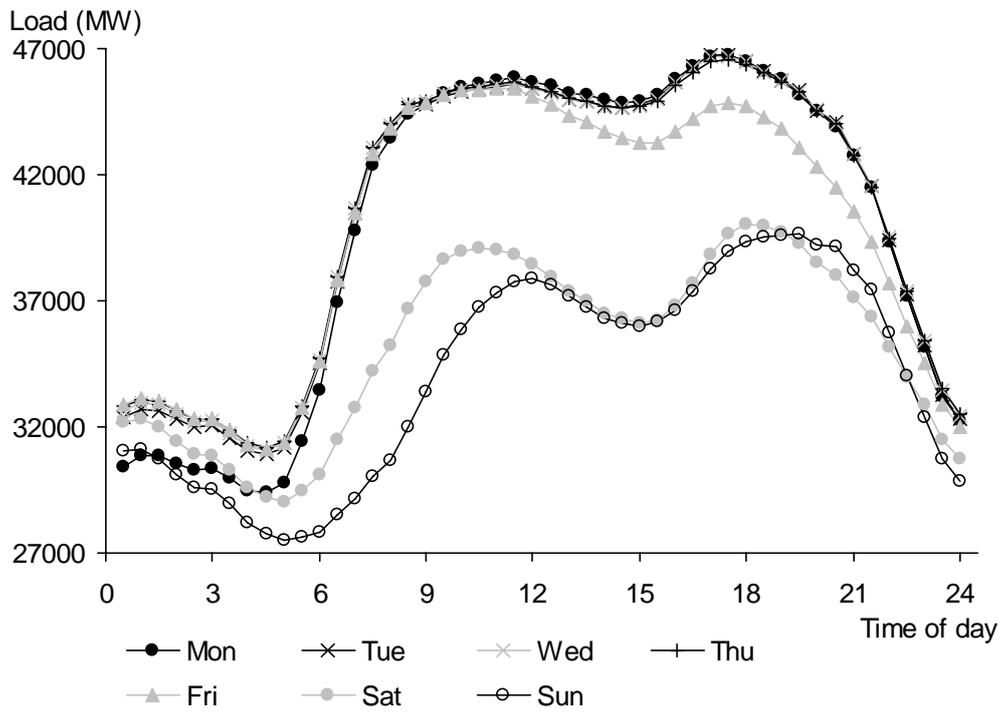


Figure 4 Average intraday cycle for each day of the week constructed using only in-sample observations for France.

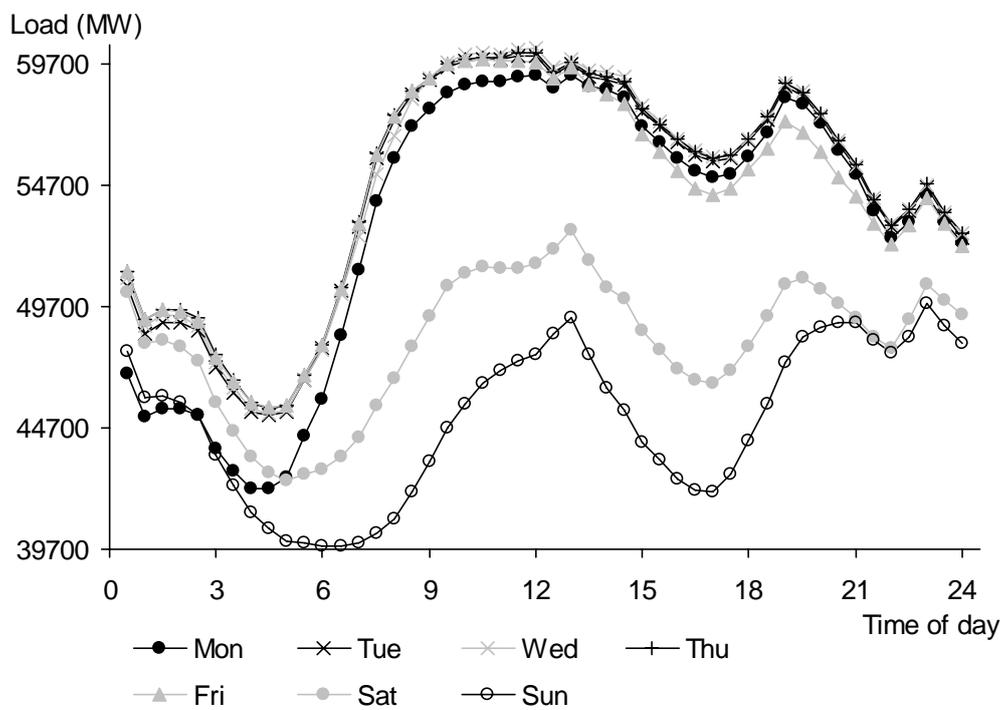


Figure 5 For the British and French load series, average MAPE for seasonal ARMA.

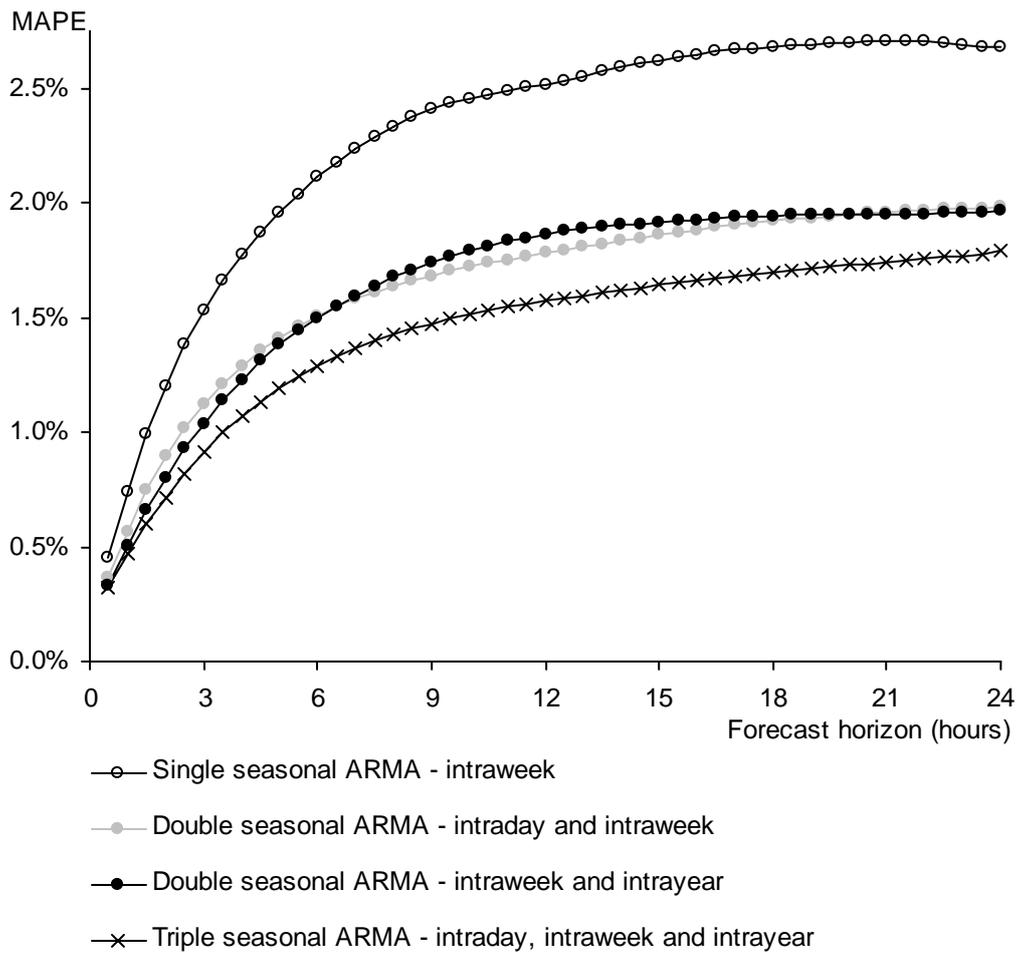


Figure 6 For the British and French load series, average MAPE for HWT exponential smoothing.

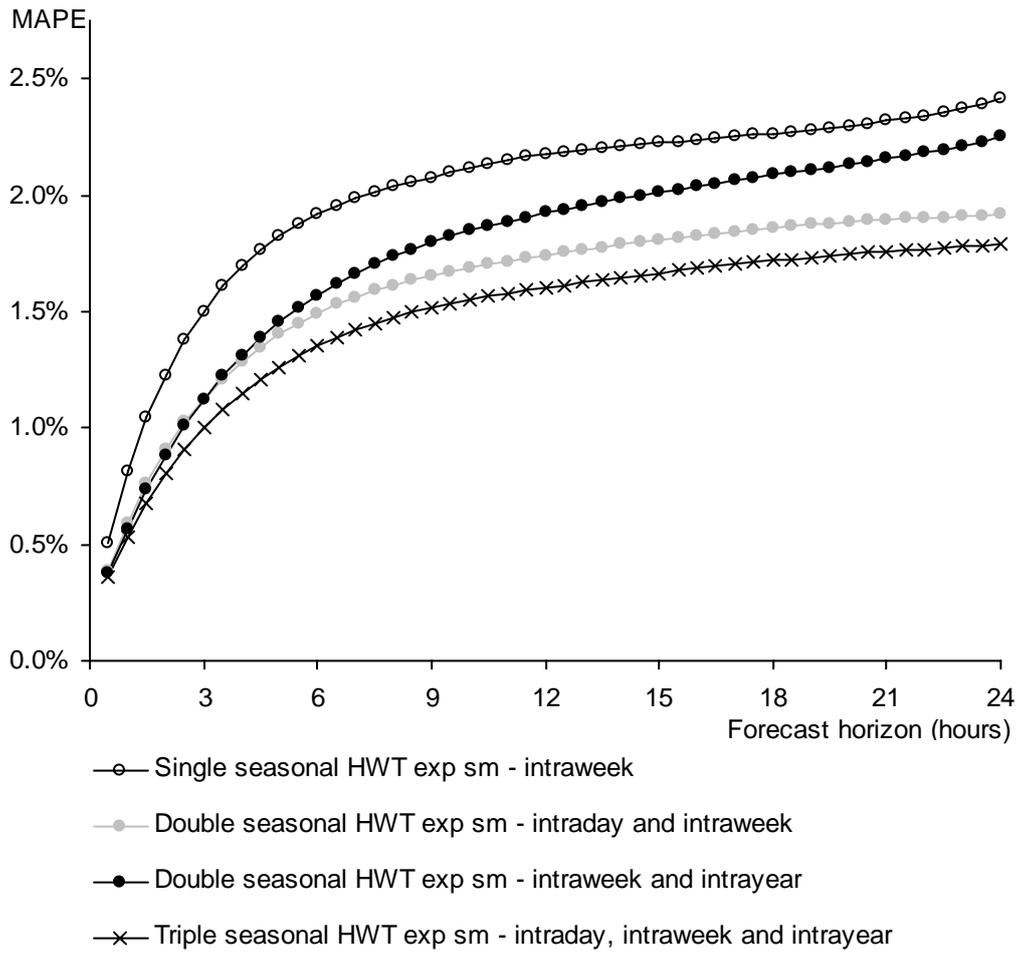


Figure 7 For the British and French load series, average MAPE for double seasonal HWT and IC exponential smoothing methods with and without the residual autocorrelation adjustment term. Both methods are formulated for intraday and intraweek seasonal cycles. The IC method uses 5 cycles and is the restricted form.

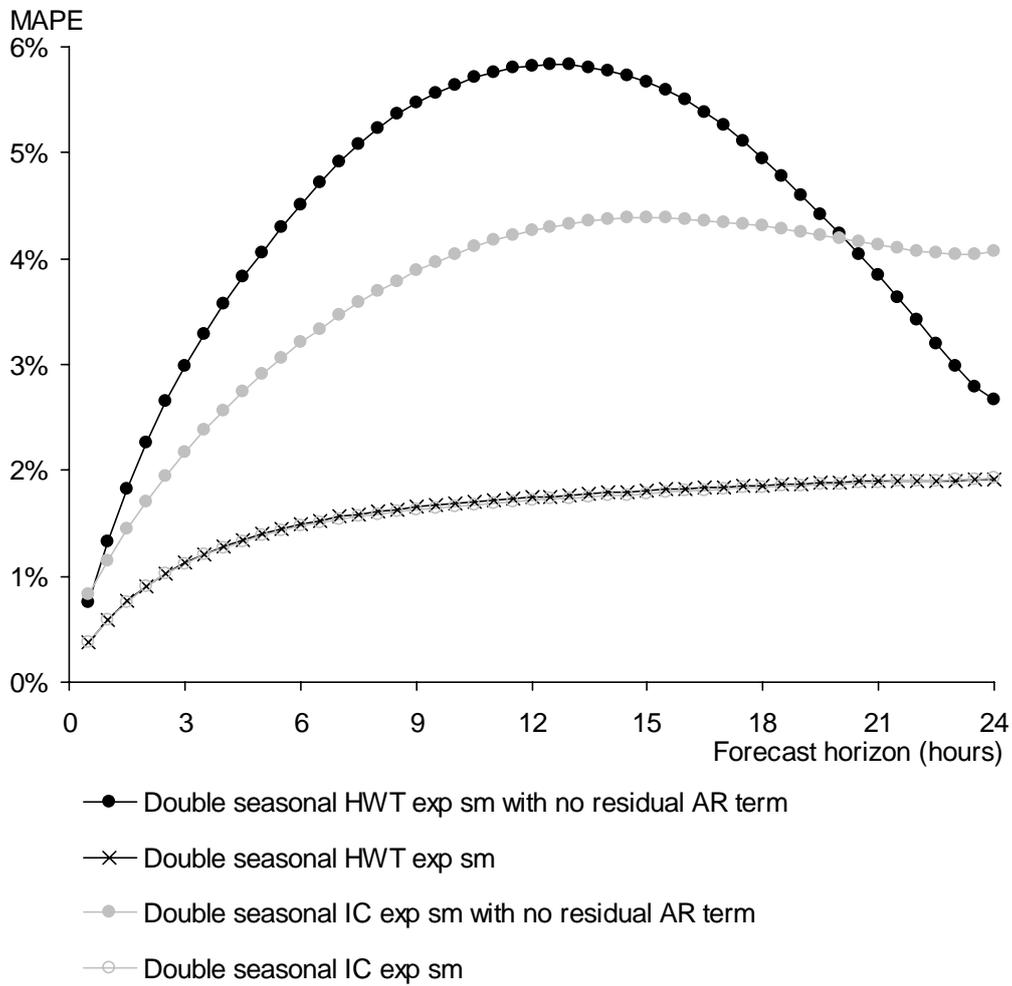


Figure 8 For the French load series, MAPE for double seasonal IC exponential smoothing.

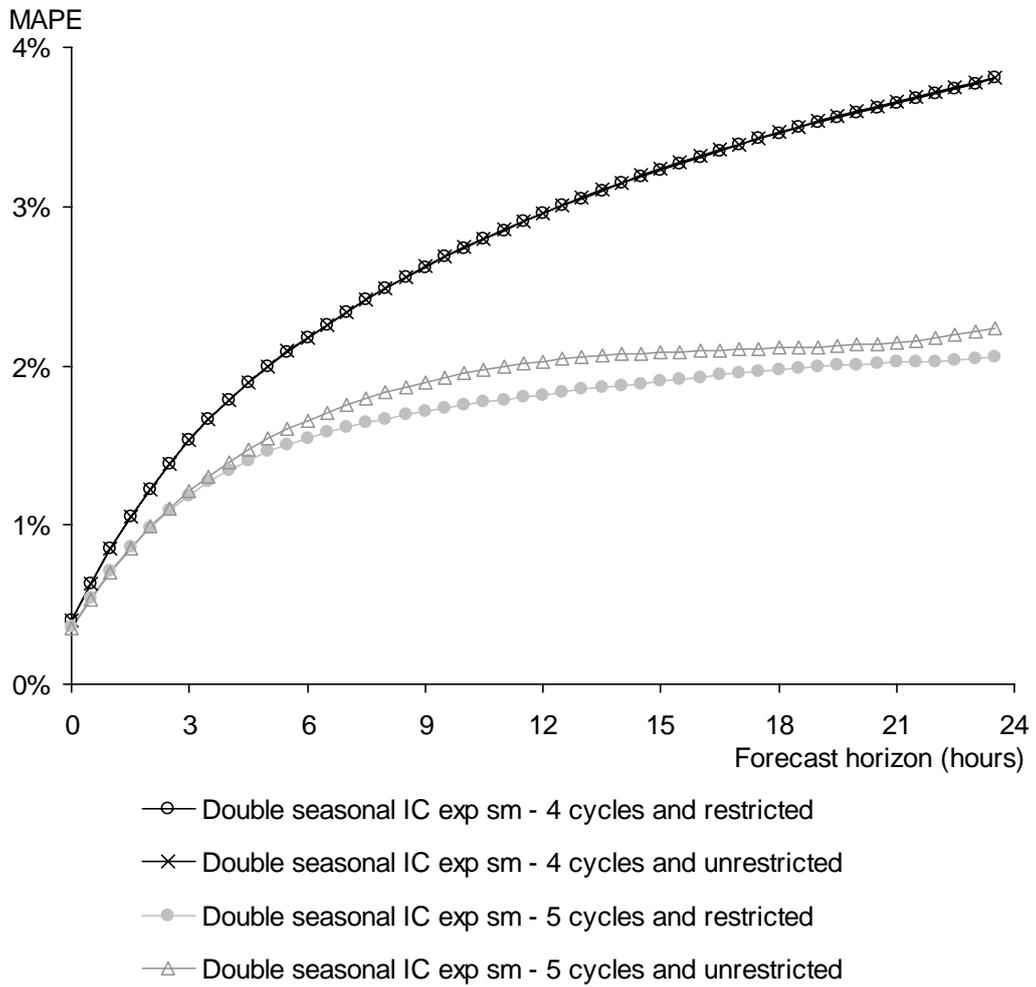


Figure 9 For the British and French load series, average MAPE for IC exponential smoothing with five distinct intraday cycles.

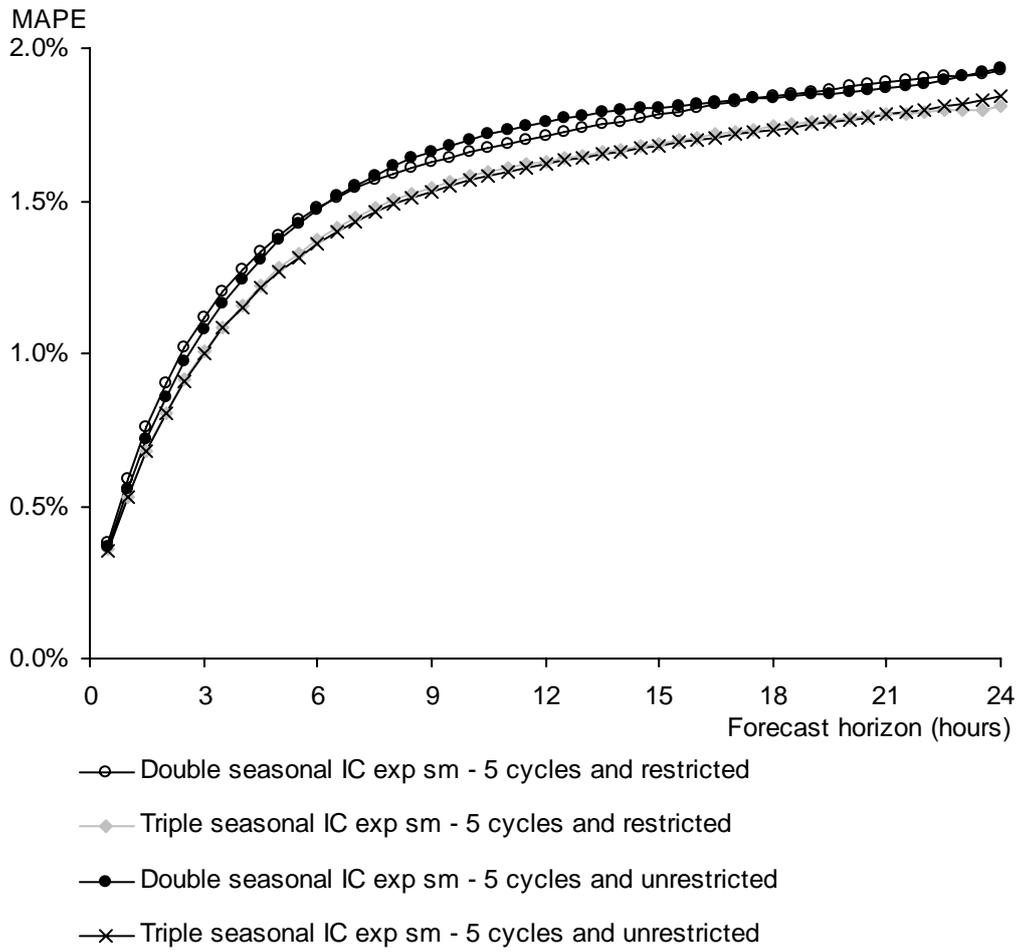


Figure 10 For the British and French load series, average MAPE for seasonal ARMA, HWT exponential smoothing and the neural network.

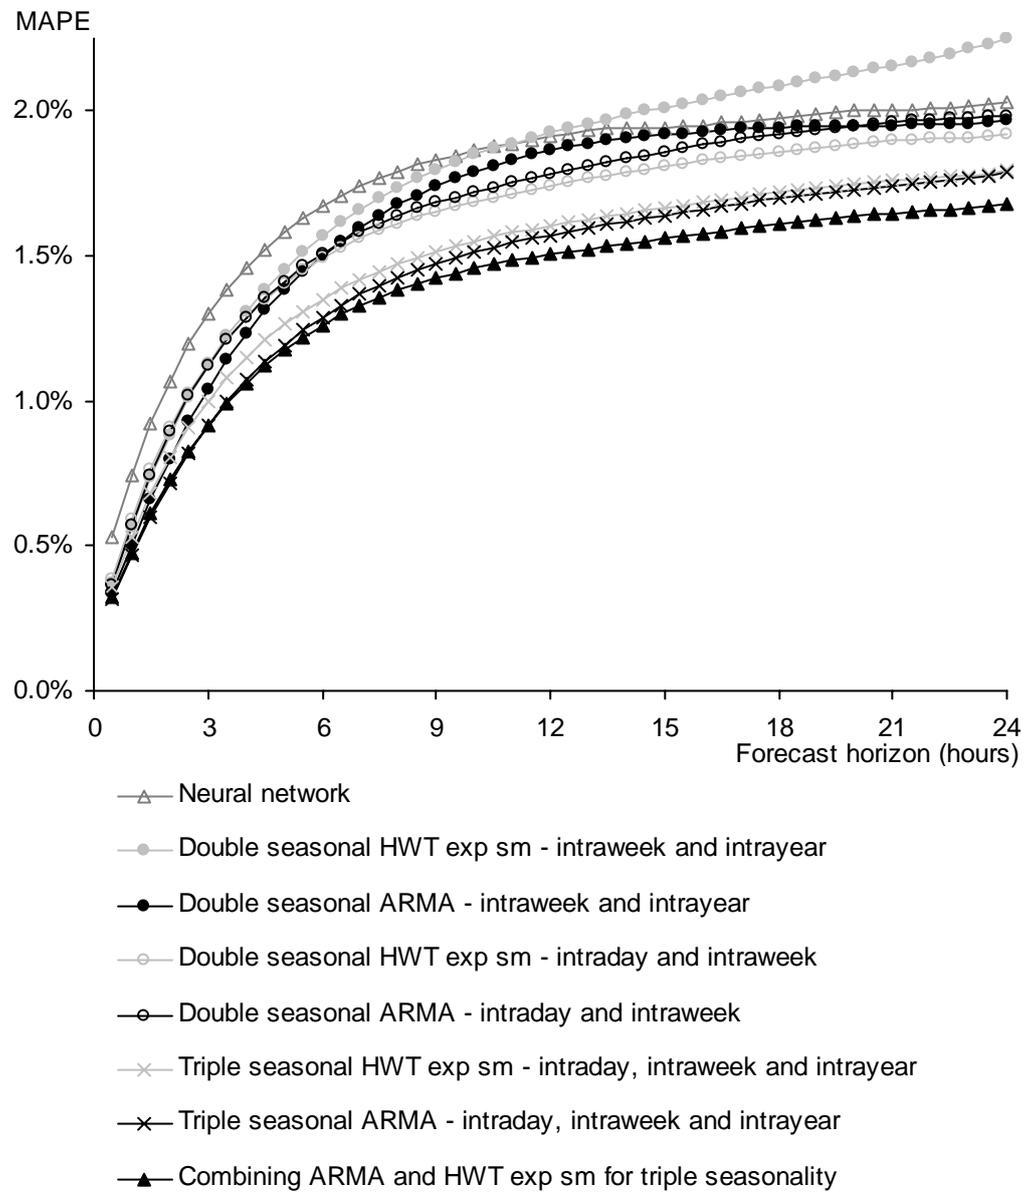


Table 1 For the British and French load series, parameters of the HWT exponential smoothing methods.

| | λ | δ | ω | α | ϕ |
|--|-----------|----------|----------|----------|--------|
| British | | | | | |
| Single – intraweek | 0.028 | | 0.912 | | 0.939 |
| Double – intraday and intraweek | 0.024 | 0.306 | 0.391 | | 0.943 |
| Double – intraweek and intrayear | 0.014 | | 0.677 | 0.010 | 0.980 |
| Triple – intraday, intraweek and intrayear | 0.021 | 0.406 | 0.327 | 0.016 | 0.944 |
| French | | | | | |
| Single – intraweek | 0.030 | | 0.763 | | 0.955 |
| Double – intraday and intraweek | 0.002 | 0.380 | 0.030 | | 0.987 |
| Double – intraweek and intrayear | 0.018 | | 0.797 | 0.007 | 0.957 |
| Triple – intraday, intraweek and intrayear | 0.001 | 0.151 | 0.368 | 0.024 | 0.989 |

Table 2 For the British and French load series, parameters of the restricted form of the IC exponential smoothing method with five distinct intraday cycles.

| | λ | γ_i | $\gamma_j (i \neq j)$ | α | ϕ |
|--|-----------|------------|-----------------------|----------|--------|
| British | | | | | |
| Double – intraday and intraweek | 0.020 | 0.651 | 0.290 | | 0.944 |
| Triple – intraday, intraweek and intrayear | 0.004 | 0.356 | 0.337 | 0.396 | 0.970 |
| French | | | | | |
| Double – intraday and intraweek | 0.009 | 0.390 | 0.354 | | 0.980 |
| Triple – intraday, intraweek and intrayear | 0.002 | 0.380 | 0.355 | 0.156 | 0.988 |