

Forecasting Intraday Time Series with Multiple Seasonal Cycles
Using Parsimonious Seasonal Exponential Smoothing

James W. Taylor

Saïd Business School, University of Oxford, UK

Ralph D. Snyder

Department of Econometrics and Business Statistics, Monash University, Australia.

Omega, 2012, Vol. 40, pp. 748-757.

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

Forecasting Intraday Time Series with Multiple Seasonal Cycles Using Parsimonious Seasonal Exponential Smoothing

Abstract

This paper concerns the forecasting of seasonal intraday time series that exhibit repeating intraweek and intraday cycles. A recently proposed exponential smoothing method involves smoothing a different intraday cycle for each distinct type of day of the week. Similar days are allocated identical intraday cycles. A limitation is that the method allows only whole days to be treated as identical. We introduce a new exponential smoothing formulation that allows parts of different days of the week to be treated as identical. The result is a method that involves the smoothing and initialisation of fewer terms. We evaluate forecasting up to a day ahead using two empirical studies. For electricity load data, the new method compares well with a range of alternatives. The second study involves a series of arrivals at a call centre that is open for a shorter duration at the weekends than on weekdays. Among the variety of methods considered, the new method is the only one that can model in a satisfactory way this situation, where the number of periods on each day of the week is not the same.

Keywords: Forecasting; Time Series; Exponential smoothing; Electricity load; Call centre arrivals

1. Introduction

This paper is concerned with univariate forecasting methods for time series recorded at an intraday frequency. It is common for such data to exhibit seasonal cycles within the week and also within the day. Examples of the type of series of interest are the two samples of half-hourly electricity load presented in Fig. 1. Both series show a repeating intraweek cycle of length $m_2=336$ periods and also, at least for the weekdays, there is similarity between the intraday patterns of length $m_1=48$ periods. The univariate methods that we consider aim to capture this feature of ‘double seasonality’. Although electricity load is often modelled in terms of weather variables [1], the use of univariate methods is common for lead times less than half a day, and for longer horizons in locations where weather data is not available. Other examples of intraday data exhibiting double seasonality are call centre arrivals, hospital admissions, transportation and internet traffic [2,3,4]. In these applications, explanatory variables are often not available, and so univariate forecasting methods are required.

----- Fig. 1 -----

A variety of univariate methods have been proposed for forecasting intraday data. These include exponentially weighted regression [5]; seasonal ARMA [6]; artificial neural networks [7]; dynamic harmonic regression [2]; vector autoregression [8]; random effects models [9]; and an approach based on singular value decomposition [10]. In many applications with intraday data, the forecasting method must be implemented within an automated system. A class of univariate methods that have been widely used in automated applications due to their robustness, intuitiveness and accuracy is exponential smoothing [11,12]. Indeed, exponential smoothing methods have performed well in empirical studies with intraday data [13,14]. This motivates us to develop, in this paper, the use of exponential smoothing for intraday data.

Two exponential smoothing methods have been presented in the literature for intraday data. Taylor’s [15] adaptation of the Holt-Winters method involves the smoothing of an intraday cycle and intraweek cycle. We term this ‘HWT exponential smoothing’. Although empirical results for the method have been encouraging [13,14], it has two features that are rather unappealing. Firstly, the method uses a common intraday cycle for all days of the week. Secondly, the method can be viewed as being of high dimension because the method’s smoothing of an intraday cycle and an intraweek

cycle requires the initialisation and updating of a large number of terms. Gould *et al.* [4] address these issues by introducing an exponential smoothing method that models the intraweek and intraday cycles through the use of a different intraday cycle for each distinct type of day of the week. For example, if the pattern of demand on Saturdays is different to the rest of the week, it is allocated its own intraday cycle. If all weekdays can be assumed to be the same, a common intraday cycle can be used for these five days. The result is a method with lower dimensionality than the HWT method. Due to its focus on the intraday cycle, the method has been termed ‘intraday cycle (IC) exponential smoothing’.

A limiting feature of IC exponential smoothing is that it allows only whole days to be treated as identical. We would argue that it often makes more sense to assume that just parts of days are identical. For example, it could be that during daylight hours the series differs on each day of the week, but that the pattern during night hours can be treated as identical on all days of the week. In this paper, we propose a new exponential smoothing formulation that has the flexibility to allow parts of different days of the week to be treated as identical. The new method is of lower dimension than the IC method. In view of this, we term the method ‘parsimonious seasonal exponential smoothing’. The HWT and IC methods can be viewed as special cases of this new method. A simple version of the method is briefly discussed by Hyndman *et al.* [11, Section 3.5.5], but our study is the first to consider it formally and to develop the method for modelling intraday data.

In this paper, we present two empirical studies. The first involves two series of electricity load; one for Great Britain and the other for France. Automated electricity demand prediction is required for the control and scheduling of power systems. The liberalisation of electricity markets has meant that load forecasting is also important for market participants to support energy transactions [16,17]. We use the two load series to illustrate the implementation of methods in the earlier sections of the paper before using the series to compare the forecast accuracy of the various methods. The series in Fig. 1 are samples of the two electricity load series. The full series spanned the three-year period from Thursday 1 January 2004 to Sunday 31 December 2006, inclusive. We use the first two years of data for estimation of method parameters, and the remaining year for post-sample evaluation. As the estimation sample contains just two years of data, we make no attempt to model annual seasonality. Both of our empirical studies focus on lead times up to one day-ahead.

The second empirical study focuses on a half-hourly series of the number of calls arriving at a call centre of an Israeli bank. The efficient scheduling of call centre staff relies on accurate forecasts for the number of calls [18]. Although causal models are sometimes used for modelling call arrivals [e.g. 19,20] it is univariate methods that tend to be used for short-term forecasting of intraday call arrivals, which is our focus in this paper. The Israeli bank's call centre is open for a shorter duration at the weekends than on weekdays. This situation, where the number of periods on each day of the week is not the same, cannot be satisfactorily addressed with the HWT and IC methods. By contrast, it can be modelled in a straightforward way using parsimonious seasonal exponential smoothing. Our call centre forecasting study considers lead times up to one day-ahead.

In each of the three time series used in this paper, we identified a number of unusual days, termed 'special' days, on which the pattern of observations differed considerably from the regular seasonal pattern. These were either public holidays or days adjacent or close to public holidays. Online prediction systems are typically replaced for these days by forecasts prepared offline. As the methods considered in this paper are appropriate for automated online prediction, prior to fitting the methods, we smoothed out the special days. This smoothing used the average of the corresponding periods in 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.

In Section 2, we review the HWT and IC exponential smoothing methods. In Section 3, we introduce our new parsimonious seasonal exponential smoothing method. Section 4 uses the load data to compare post-sample forecast accuracy for the exponential smoothing methods, as well as ARMA modelling and an artificial neural network approach. In Section 5, we present our empirical study of intraday call centre arrivals. Section 6 summarises the paper, and provides concluding comments.

2. Previously proposed exponential smoothing methods for intraday data

2.1. HWT exponential smoothing

Taylor's [15] HWT exponential smoothing extends the Holt-Winters method in order to accommodate the intraday and intraweek cycles in intraday data. It is presented in the following expressions:

$$\hat{y}_t(k) = l_t + d_{t-m_1+k_1} + w_{t-m_2+k_2} + \phi^k e_t \quad (1)$$

$$e_t = y_t - (l_{t-1} + d_{t-m_1} + w_{t-m_2}) \quad (2)$$

$$l_t = \alpha(y_t - d_{t-m_1} - w_{t-m_2}) + (1-\alpha)l_{t-1} \quad (3)$$

$$d_t = \delta(y_t - l_t - w_{t-m_2}) + (1-\delta)d_{t-m_1} \quad (4)$$

$$w_t = \omega(y_t - l_t - d_{t-m_1}) + (1-\omega)w_{t-m_2} \quad (5)$$

where l_t is the smoothed level; d_t is the seasonal index for the intraday seasonal cycle; w_t is the seasonal index for the intraweek seasonal cycle that remains after the intraday cycle is removed; m_1 and m_2 are the lengths of the intraday and intraweek cycles, respectively; α , δ and ω are smoothing parameters; and $\hat{y}_t(k)$ is the k step-ahead forecast made from forecast origin t ; and $k_1 = [(k-1) \bmod m_1] + 1$ and $k_2 = [(k-1) \bmod m_2] + 1$. The term involving the parameter ϕ , in the forecast function of expression (1), is an adjustment for first-order residual autocorrelation. This adjustment substantially improves the method's forecast accuracy, and so should be considered integral to the method [15]. A trend term is not included in any of the exponential smoothing formulations in this paper because, in our empirical work, its inclusion 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.

Hyndman *et al.* [11] describe how innovations state space models are an attractive class of statistical models for representing exponential smoothing methods. The following innovations state space model captures the features of HWT exponential smoothing [21]:

$$y_t = l_{t-1} + d_{t-m_1} + w_{t-m_2} + \phi e_{t-1} + \varepsilon_t$$

$$e_t = y_t - (l_{t-1} + d_{t-m_1} + w_{t-m_2})$$

$$l_t = \alpha(y_t - d_{t-m_1} - w_{t-m_2}) + (1-\alpha)l_{t-1}$$

$$d_t = \delta(y_t - l_t - w_{t-m_2}) + (1-\delta)d_{t-m_1}$$

$$w_t = \omega(y_t - l_t - d_{t-m_1}) + (1-\omega)w_{t-m_2}$$

where $\varepsilon_t \sim \text{NID}(0, \sigma^2)$. This notation is used to indicate that the ε_t are independent and normally distributed with zero mean and constant variance σ^2 . An advantage of the HWT method is its relative simplicity. Indeed, its formulation will seem natural and intuitive to the many people familiar with the standard Holt-Winters method. Turning to the method's disadvantages, Gould *et al.* [4] argue that it is

unappealing to use the same intraday cycle state variable, d_t , for each of the seven days of the week. Furthermore, the HWT method can be viewed as being of high dimension, as it involves initialising and updating the level, plus the m_1 periods of the intraday cycle, as well as the m_2 periods of the intraweek cycle. For half-hourly load data, this amounts to $(1+m_1+m_2) = (1+48+336) = 385$ terms.

2.2. IC exponential smoothing

The IC exponential smoothing method of Gould *et al.* [4] allows the intraday cycle for the different days of the week to be represented by different seasonal components. To obtain a more parsimonious formulation than the HWT method, a common intraday cycle is used for days of the week that exhibit similar patterns of demand. In contrast with the HWT method, the formulation does not contain a representation of the intraweek seasonal cycle.

Before presenting the model, let us first consider which days of the week can be considered to have a similar intraday cycle for our electricity load series. To guide us, we follow the approach of Gould *et al.* by using the plots in Figs. 2 and 3, which show, for each day of the week, the average load for each period of the day, calculated using only in-sample data. For the British data, Fig. 2 shows that Saturday and Sunday should each be treated as having its own distinct cycle. In this figure, the patterns for Monday morning and Friday afternoon suggest that there should be a distinct cycle for each of these days. The patterns on the other three weekdays are very similar, which suggests a common cycle for these days. Similar conclusions can be drawn from Fig. 3 for the French series, although the argument for a distinct Friday intraday cycle is less clear. It seems reasonable to implement, for both series, a model with distinct intraday cycles for Friday, Saturday, Sunday and Monday, and a common cycle for the other three days. IC exponential smoothing for the case of five distinct intraday cycles is presented as an innovations state space model in the following expressions:

$$y_t = l_{t-1} + \sum_{i=1}^5 I_{it} d_{i,t-m_i} + \phi e_{t-1} + \varepsilon_t$$

$$e_t = y_t - \left(l_{t-1} + \sum_{i=1}^5 I_{it} d_{i,t-m_i} \right)$$

$$l_t = l_{t-1} + \alpha e_t$$

$$d_{it} = d_{i,t-m_i} + \left(\sum_{j=1}^5 \gamma_{ij} I_{jt} \right) e_t \quad (i=1 \text{ to } 5)$$

$$I_{jt} = \begin{cases} 1 & \text{if time period } t \text{ occurs on a day of type } j \\ 0 & \text{otherwise} \end{cases}$$

where $\varepsilon_t \sim \text{NID}(0, \sigma^2)$; l_t is the smoothed level; d_{it} is the value of the intraday cycle of type i in period t (for $i=1$ to 5); and α and the γ_{ij} are smoothing parameters. Although not included by Gould *et al.*, we find that the addition of the residual adjustment term involving the parameter ϕ leads to noticeable improvement in forecast accuracy. The γ_{ij} can be viewed as a 5×5 matrix of parameters. Through these parameters, the formulation enables the five types of intraday cycle to be updated at different rates. It is interesting to note that these parameters also enable intraday cycle of type i to be updated when the current period is not on a day of type i . Gould *et al.* propose several alternative restrictions for the γ_{ij} matrix; one involves estimation with just the basic constraint that the parameters lie between zero and one, and the other involves the additional restrictions of common diagonal elements and common off-diagonal elements. We refer to these two forms of the model as unrestricted and restricted, respectively. As noted by Gould *et al.*, the restricted form is identical to the HWT method, provided seven distinct intraday cycles are used and the seasonal smoothing parameters are written in terms of the HWT parameters as $\gamma_{ii} = \delta + \omega$ and $\gamma_{ij} = \delta$.

In comparison with the HWT method, the IC method has the appeal of lower dimensionality. As noted at the end of Section 2.1, for our load data, the HWT method involves initialising and updating 385 terms. In contrast, the IC method, with five distinct intraday cycles, involves initialising and updating the level and $5m_1$ seasonal periods, which amounts to $(1+5 \times 48) = 241$ terms.

----- Figs. 2 and 3 -----

3. A new parsimonious exponential smoothing method for intraday data

3.1. Parsimonious seasonal exponential smoothing

A limiting feature of IC exponential smoothing is that it allows only whole days to be treated as identical. We would argue that it often makes more sense to allow parts of days to be treated as the same. For example, in the IC method of Section 2.2, a distinct intraday cycle was used for Fridays.

However, for the British data, Fig. 2 shows that the average Friday pattern of load prior to about 11am seems to be very similar to that on Tuesday, Wednesday and Thursday, and so it seems inefficient to treat these earlier Friday hours as distinct from the rest of the week. We propose a new exponential smoothing formulation that allows parts of different days of the week to be treated as identical, with the remaining parts of these days treated as distinct. In comparison with the IC method, the new formulation offers greater flexibility and parsimony. This prompts us to term the method ‘parsimonious seasonal exponential smoothing’.

In this method, we use the term ‘season’ to refer to a set of periods in the intraweek cycle for which demand is assumed to be identical. The method proceeds by clustering the m_2 periods of the week so that each belongs to one of M distinct seasons, where $M \leq m_2$. (We discuss the clustering for the load data in Section 3.2.) A seasonal state s_{it} is defined for each season i , and exponential smoothing is used to update the M seasonal states. The simplest version of parsimonious seasonal exponential smoothing, which we term ‘Version 1: basic model’, is presented in the form of an innovations state space model in the following expressions:

Parsimonious seasonal exponential smoothing – Version 1: basic model

$$y_t = \sum_{i=1}^M I_{it} s_{i,t-1} + \phi e_{t-1} + \varepsilon_t \quad (6)$$

$$e_t = y_t - \sum_{i=1}^M I_{it} s_{i,t-1} \quad (7)$$

$$s_{it} = s_{i,t-1} + (\alpha + \omega I_{it}) e_t \quad (i = 1, 2, \dots, M) \quad (8)$$

$$I_{it} = \begin{cases} 1 & \text{if period } t \text{ occurs in season } i \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

$\varepsilon_t \sim \text{NID}(0, \sigma^2)$; α and ω are smoothing parameters; and ϕ is the parameter of a residual autoregressive term. Each seasonal state, s_{it} , is updated using expression (8), with the degree of updating depending on whether or not the current period t is in season i . There is no separate smoothing of the level, as in the HWT and IC methods. Instead, the level is captured within the seasonal states, and updated each period through the parameter α . This gives a more concise formulation, and potentially simplifies initialisation of the states. The smoothing of a trend can be incorporated by the addition of a constant term in expression (8).

3.2. Clustering periods of the week into distinct seasons

Drawing on the approach used to specify IC exponential smoothing, we now use Figs. 2 and 3 to dictate the clustering of the periods of the week into seasons for parsimonious seasonal exponential smoothing. For the British plot in Fig. 2, we make the following observations:

- (i) As the average demand on each weekend period is so very different to the other periods of the week, we define 96 distinct seasons for the weekend periods.
- (ii) The average patterns of demand on Tuesdays, Wednesdays and Thursdays are so similar that they should be treated as having a common intraday cycle. For this, we define 48 distinct seasons.
- (iii) After about 8.30am, the average pattern of demand on Mondays is extremely similar to that on Tuesdays, Wednesdays and Thursdays. This implies that we should specify these 31 Monday periods to be identical to the corresponding periods on Tuesdays, Wednesdays and Thursdays. However, for each of the first 17 periods on Mondays, we specify distinct seasons, as they would seem to be noticeably different from the other periods of the week.
- (iv) Up to about 11am, the average pattern of demand on Fridays is very similar to that on Tuesdays, Wednesdays and Thursdays, and so we specify these 22 Friday morning periods to be identical to the corresponding periods on Tuesdays, Wednesdays and Thursdays. However, we specify distinct seasons for the last 26 periods on Fridays.

This assessment of Fig. 2 leads us to define $(96+48+17+26)=187$ distinct seasons for the British demand data. This can be compared with the number of terms that require initialisation and updating in the HWT and IC methods. As we discussed at the end of Sections 2.1 and 2.2, the HWT method involves 385 terms, and our implementation of the IC method has 241 terms.

Let us now consider the French data. A similar set of observations to those made above regarding Fig. 2 can be made for the French average intraday cycles plotted in Fig. 3:

- (i) 96 distinct seasons are defined for the weekend periods.
- (ii) 48 distinct seasons are used to represent the common intraday cycle occurring on Tuesdays, Wednesdays and Thursdays.

(iii) The average pattern of demand up to midday on Mondays is rather different to the other days, and so we define distinct seasons for the first 24 Monday periods. The other 24 periods are defined to be identical to the corresponding periods on Tuesdays, Wednesdays and Thursdays.

(iv) For Fridays, the average demand is somewhat different to the other days beyond about 1pm. Therefore, we specify 22 distinct seasons for these Friday periods. The other 26 periods are defined to be identical to the corresponding periods on Tuesdays, Wednesdays and Thursdays.

The result of this is $(96+48+24+22)=190$ distinct seasons for the parsimonious seasonal exponential smoothing model applied to the French data. We acknowledge that, based on Figs. 2 and 3, other clusters of periods into seasons could be considered, but we have opted here for simplicity. As we have presented it, the selection of clusters relies on the judgement of the user. We also considered statistical approaches to clustering that require only minimal judgemental input. We considered hierarchical and k-mean cluster analysis based on the sum of the Euclidean distances between observations for different periods of the same historical week in the estimation sample. We implemented an unconstrained clustering that allowed any periods of the week to be clustered in the same season, and hence treated as identical. We also experimented with a constrained version that only permitted clusters to be formed from periods of different days if they corresponded to the same half-hourly period of the day. However, none of the statistical clustering approaches led to improved accuracy over the judgemental selections of clusters. In view of this, the empirical results in this paper correspond to a purely judgemental specification of the new method, which is also the approach taken by Gould *et al.* for their IC exponential smoothing method.

3.3. Extending the parsimonious seasonal exponential smoothing model

The HWT and IC exponential smoothing methods are special cases of extended versions of parsimonious seasonal exponential smoothing. In view of the strong empirical results and intuitive appeal of these methods, it seems sensible to consider these extended versions of the parsimonious seasonal exponential smoothing method, as they may lead to improved forecast accuracy over the basic model. To obtain the HWT method, we first set $M=m_2$, so that each period of the week is treated as a distinct season. We then extend the parsimonious seasonal exponential smoothing model to the

formulation of expressions (10)-(14), which we term ‘Version 2: with period of day effect’. In comparison with Version 1 (the basic model), Version 2 includes an additional term in the updating equations of expression (12), and a corresponding additional smoothing parameter δ . The updating equations cause the seasonal states to be updated to differing degrees depending on whether or not the current period t falls in season i or occurs at the same period of the day as season i .

Parsimonious seasonal exponential smoothing – Version 2: with period of day effect

$$y_t = \sum_{i=1}^M I_{it} s_{i,t-1} + \phi e_{t-1} + \varepsilon_t \quad (10)$$

$$e_t = y_t - \sum_{i=1}^M I_{it} s_{i,t-1} \quad (11)$$

$$s_{it} = s_{i,t-1} + (\alpha + \delta I_{it}^{\text{period of day}} + \omega I_{it}) e_t \quad (i = 1, 2, \dots, M) \quad (12)$$

$$I_{it} = \begin{cases} 1 & \text{if period } t \text{ occurs in season } i \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

$$I_{it}^{\text{period of day}} = \begin{cases} 1 & \text{if period } t \text{ occurs at same period of day as season } i \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

The restricted version of the IC method is also a special case of Version 2 of the parsimonious seasonal exponential smoothing model. To obtain the restricted IC method with, say, five distinct intraday cycles, we define $M=(5 \times m_1)$ distinct seasons as the periods of the five intraday cycles. With the same distinct seasons, a further extension of parsimonious seasonal exponential smoothing delivers the unrestricted form of the IC method. This extension is given in expressions (15)-(20). In addition to requiring that each period of the week is assigned to one of the M seasons, this version of the model also requires that each period of the week is classified as belonging to one of the five distinct day types. As in the unrestricted IC method, this parsimonious seasonal exponential smoothing formulation contains a smoothing parameter α for the level and a matrix of parameters γ_{jk} to dictate the updating of the seasonality.

Parsimonious seasonal exponential smoothing – Version 3: with period of day and day type effects

$$y_t = \sum_{i=1}^M I_{it} s_{i,t-1} + \phi e_{t-1} + \varepsilon_t \quad (15)$$

$$e_t = y_t - \sum_{i=1}^M I_{it} s_{i,t-1} \quad (16)$$

$$s_{i,t} = s_{i,t-1} + \left(\alpha + I_{it}^{period\ of\ day} \sum_{j=1}^5 \sum_{k=1}^5 \gamma_{jk} I_{jt}^{period\ day\ type} I_{ki}^{season\ day\ type} \right) e_t \quad (i = 1, 2, \dots, M) \quad (17)$$

$$I_{it}^{period\ of\ day} = \begin{cases} 1 & \text{if period } t \text{ occurs at same period of day as season } i \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

$$I_{jt}^{period\ day\ type} = \begin{cases} 1 & \text{if period } t \text{ occurs on a day of type } j \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

$$I_{ki}^{season\ day\ type} = \begin{cases} 1 & \text{if season } i \text{ occurs on a day of type } k \\ 0 & \text{otherwise} \end{cases} \quad (20)$$

We have introduced Versions 2 and 3 of the parsimonious seasonal exponential smoothing model as being equivalent, with appropriate choices of the M distinct seasons, to the HWT and IC methods. However, our proposal is that these additional versions of the new model could be used with other choices for the M distinct seasons. In Section 4, we consider the three versions of the method with the judgementally selected clusters that in Section 3.2 led to $M=187$ for the British series, and $M=190$ for the French series. To implement Version 3 of the model for the British series, we defined the following five different ‘day types’: Saturday is ‘day type 1’; Sunday is ‘day type 2’; Tuesday, Wednesday, Thursday, Monday periods after 8.30am, and Friday periods before 11am are all ‘day type 3’; Monday morning periods up to 8.30am are ‘day type 4’; and Friday periods from 11am onwards are ‘day type 5’. For the French series, we applied the same ‘day type’ definitions, except that 8.30am was replaced by midday, and 11am was replaced by 1pm.

4. Evaluation of forecasts for the electricity load time series

In this section, we use the two load series to compare the post-sample forecast accuracy of the exponential smoothing methods and several benchmark methods. Estimation was performed once using the first two years of the data, and the forecast origin was then rolled forward through the one-year post-sample evaluation period to produce a collection of forecasts for lead times up to a day

ahead. Before evaluating forecast accuracy, we first describe the estimation of the exponential smoothing methods and introduce the benchmark methods.

4.1. Exponential smoothing methods

We implemented the HWT method; the restricted and unrestricted forms of the IC method; and the three versions of parsimonious seasonal exponential smoothing. For quarterly data, Rasmussen [22] demonstrates that forecast accuracy can be improved if the exponential smoothing initial values are treated as parameters and estimated in the same optimisation routine as the smoothing parameters. Unfortunately, this is impractical for our study due to the number of terms that require initialisation in each of our intraday forecasting methods. In view of this, we estimated initial smoothed values for the level and seasonal components by averaging the observations from the first three weeks of data. We constrained the parameters to lie between zero and one, and estimated them by minimizing the sum of squared in-sample forecast errors (SSE). For each method, the optimisation proceeded by first generating 10^5 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 British series, the optimized parameters for the HWT method were $\alpha=0.003$, $\delta=0.295$, $\omega=0.397$ and $\phi=0.968$, and for the restricted form of the IC method, we derived $\alpha=0.013$, $\gamma_{ij}=0.281$, $\gamma_{ii}=0.640$ and $\phi=0.953$. The low values of α and high values of ϕ for the HWT method are consistent with previous studies with intraday load data [13,15]. Earlier, we noted that the restricted form of the IC method is identical to the HWT method, provided seven distinct intraday cycles are used and the parameters for the two methods obey $\gamma_{ii}=\delta+\omega$ and $\gamma_{ij}=\delta$. Our derived parameters, for both the British and French data, approximately satisfy these two equations. Even though our IC method implementation involves five distinct intraday cycles, it seems likely that the HWT and restricted IC methods will deliver similar results. For Version 1 of the parsimonious seasonal exponential smoothing method (the basic model), the optimised parameter values for the British data were $\alpha=0.035$, $\omega=0.739$ and $\phi=0.936$, and for Version 2 (the version with period of day effect), we

obtained $\alpha=0.018$, $\delta=0.272$, $\omega=0.346$ and $\phi=0.947$. The values for Version 2 are similar to those derived for the HWT method, which is perhaps a little surprising because the two methods have quite different formulations and dimensionality. In implementing the parsimonious seasonal exponential smoothing method, we should emphasise that the judgemental clustering of the seasons, discussed in Sections 3.2 and 3.3, was based on only the in-sample data. To have based the judgemental clustering on the post-sample data would have given the method an unfair advantage over the other methods in our comparison of forecasting accuracy.

4.2. Simplistic benchmark methods

We implemented two simplistic benchmark methods. The first takes as a forecast the most recently observed value for the same half-hour of the week as the period to be predicted. The second produces forecasts as the average of the observations from the most recent four weeks that occur at the same half-hour of the week as the period to be predicted.

4.3. Double seasonal ARMA

Double seasonal ARMA models are often used as sophisticated benchmarks in load forecasting studies [e.g. 6, 7, 23]. The multiplicative double seasonal ARMA model can be written as

$$\phi_p(L) \Phi_{P_1}(L^{m_1}) \Omega_{P_2}(L^{m_2})(y_t - c) = \theta_q(L) \Theta_{Q_1}(L^{m_1}) \Psi_{Q_2}(L^{m_2}) \varepsilon_t$$

where $\varepsilon_t \sim \text{NID}(0, \sigma^2)$; c is a constant term; L is the lag operator; and ϕ_p , Φ_{P_1} , Ω_{P_2} , θ_q , Θ_{Q_1} and Ψ_{Q_2} are polynomial functions of orders p , P_1 , P_2 , q , Q_1 and Q_2 , respectively. The lack of clear evidence from unit root tests initially led us not to apply differencing. However, for similar data to ours, some researchers have applied the differencing operator $(1-L)(1-L^{m_1})(1-L^{m_2})$ [e.g. 6, 7, 23]. In view of this, we implemented two versions of double seasonal ARMA modelling; one with this differencing operator and one with no differencing. We estimated model parameters using maximum likelihood. We considered lag polynomials up to order three, and selected AR and MA terms using the Schwarz Bayesian Criterion, with the requirement that all parameters were significant (at the 5% level).

4.4. Artificial neural networks

Artificial neural networks (ANNs) are predictive models with nonlinear structures based on an imitation of biological neurons [24]. Much research has focused on the use of ANNs to model electricity load in terms of weather variables [25]. In that context, the nonlinear and nonparametric features of ANNs have strong appeal. Their usefulness for univariate modelling is less clear, and relies on the nonlinear structures in the evolution of the load series. Nevertheless, we felt it would be interesting to include an ANN as a second sophisticated benchmark method in our study. Darbellay and Slama [7] and Taylor *et al.* [13] build univariate ANNs to model intraday load data with load at various lags as inputs. They both use a single model, and generate multi-step-ahead predictions by using forecasts for earlier lead times. Our approach differed from this by employing a separate ANN for each lead time. Using a hold-out sample of six months of the estimation sample, we found that forecast accuracy was improved by applying the differencing operator $(1 - L^{m_1})(1 - L^{m_2})$ prior to ANN modelling.

We used a single hidden layer feedforward network with k inputs, x_{it} , connected to each of m units in a single hidden layer, which, in turn, were connected to a single output. The 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 output was specified as the load variable, differenced as described above, and normalised by subtracting the mean and dividing by the standard deviation. In the model for h step-ahead prediction, as potential inputs, we considered this variable at the forecast origin and at the following lags: 1, 2, m_1-h , $2m_1-h$, $3m_1-h$, m_2-h , $2m_2-h$ and $3m_2-h$. This set of candidate input variables was motivated by our consideration of lag polynomials up to order three in the double seasonal ARMA models. The network weights were estimated using least squares. To penalise complexity and avoid overfitting, we added to the objective function the sum of the weights squared multiplied by a regularisation parameter λ [26, §9.2]. The minimisation was performed using backpropagation with learning rate η and momentum parameter μ [26, §7.5].

To derive values for λ , η , μ and m , and to select from the set of potential input variables, we evaluated one step-ahead prediction accuracy using a hold-out sample of six months of the estimation sample. We then used the same values and inputs for the models for all other lead times. For the British series, this led to the inclusion of all inputs except lag 2, and $\lambda=0.0001$, $\eta=0.1$, $\mu=0.9$ and $m=4$. For the French data, we found that all inputs were needed, and we obtained $\lambda=0.0001$, $\eta=0.1$, $\mu=0.9$ and $m=8$.

4.5. Post-sample load forecasting results

In Fig. 4, for each method, we present the mean absolute percentage error (MAPE) results averaged over the two series. Presenting just the average seems reasonable because, for most of the methods considered, their relative performances were similar for the two load series. The relative performances of the two ARMA methods were an exception to this, with differencing being preferable for the British series and no differencing being better for the French data. We also evaluated the root mean squared percentage error (RMSPE), mean absolute error (MAE) and root mean squared error (RMSE), but we do not report these results here because the rankings of the methods for these measures were similar to those for the MAPE. An alternative to the use of statistical errors measures is the evaluation of the economic impact of the forecast errors from different methods, as considered recently for a supply chain application [27]. However, this is beyond the scope of our study.

----- Fig. 4 -----

The results for the restricted IC method were similar to those of the HWT method and Version 2 of the parsimonious seasonal exponential smoothing method (the version with period of day effect), and the unrestricted IC method results were similar to those of Version 3 of the parsimonious seasonal exponential smoothing method (the version with period of day and day type effects). In view of this, for simplicity, we have not included the IC method results in Fig. 4. The results for both of the simplistic benchmark methods were so relatively poor that, with our choice of scaling of the axes, they are not shown in Fig. 4. With regard to the sophisticated benchmark methods, Fig. 4 shows the ARMA approach with differencing slightly outperforming the ARMA

method with no differencing. The ANN is relatively uncompetitive for the early lead times, but outperforms both ARMA methods beyond about 15 hours ahead. The HWT method matches or beats the ARMA and ANN methods at all lead times. Turning to parsimonious seasonal exponential smoothing, we see that Version 1 (the basic model) performed relatively poorly, Version 2 (the version with period of day effect) produced similar results to the HWT method, and Version 3 (the version with period of day and day type effects) produced the best results across all lead times.

5. Empirical study of call centre time series

5.1. The call centre times series

In this section, we consider the Israeli bank data introduced and analysed by Brown *et al.* [28], and also used by Taylor [14]. This data is taken from a relatively small call centre that is open from 7am to midnight on weekdays (Sunday to Thursday in Israel), from 7am to 2pm on Fridays, and from 8pm to midnight on Saturdays. Our analysis focuses on the single time series of total arrivals minus internet assistance calls for the period 1 August to 25 December 1999, inclusive, with the first 14 weeks used for estimation, and the remaining seven weeks used for post-sample evaluation. Estimation was performed once, with no updating of parameters. To be consistent with our analysis of the load data, we analyzed half-hourly arrivals and lead times up to a day ahead. Forecasts for such short lead times are used for dynamic intraday updating of the deployment of call centre agents [29].

In the call centre literature, several authors have applied a variance stabilizing transformation prior to model fitting. Brown *et al.* [28] present one such transformation, which draws on the count data nature of call centre arrivals. They explain that if y is a Poisson random variable, then the transformed variable $\sqrt{y + \frac{1}{4}}$ is asymptotically Gaussian as $\lambda \rightarrow \infty$. We applied this transformation prior to the use of each forecasting method.

For most of the methods considered in this paper, implementation is not straightforward for series in which the number of periods is not the same on each day of the week, which is the case for the Israeli bank series. An exception to this is the parsimonious seasonal exponential smoothing method. Before discussing this further, we describe how, in order to be able to apply the other methods, we converted the Israeli bank series into a new series with the same number of periods in

each day. We achieved this by combining the two weekend days into one single day consisting of the same number of periods as the weekdays. The result is a series with five weekdays plus a sixth day consisting of Friday arrivals from 7am to 2pm, zero values from 2pm to 8pm, and Saturday arrivals from 8pm to midnight. We refer to this as the ‘manipulated’ arrivals series. Fig. 5 presents the final four weeks of the new manipulated series. Each day of this series consists of $m_1=34$ half-hourly periods, and each week contains $m_2=6\times 34=204$ half-hourly periods. In our description of the forecasting methods below, we present separately the methods applied to the manipulated series, and those applied to the original series.

----- Fig. 5 -----

5.2. *Methods applied to the manipulated arrivals series*

For the manipulated arrivals series, we implemented all of the methods used with the load data, except the parsimonious seasonal exponential smoothing method. To specify the IC method, we plotted the average intraday patterns for each day of the week, as shown in Fig. 6. In comparison with the load series, our call centre series consists of relatively few observations and relatively high volatility. Therefore, sampling error is more of a concern for Fig. 6 than for the analogous load patterns in Figs. 2 and 3. In view of this, although there are differences in the average intraday patterns for the Israeli weekdays, Sunday to Thursday, it seems sensible to assume the weekdays possess a common intraday cycle in the IC method. The average intraday patterns on Fridays and Saturdays appear to be lower than those on the other days, and so, for the IC method, we specified a distinct intraday cycle for the day that is made up of the opening hours on Fridays and Saturdays.

----- Fig. 6 -----

5.3. *Methods applied to the original arrivals series*

The original arrivals series consisted of five days with 34 half-hourly periods, one day with 14 periods, and one day with just eight periods. This implies an intraweek cycle of length $m_2=5\times 34+14+8=192$ periods. Due to the unequal number of periods on the days of the week, we were unable to apply to the original series the HWT method with its double seasonality features. However,

by setting $d_t=0$, for all t , in the HWT method formulation, we have a simpler method that is able to model just the intraweek cycle. We applied this ‘single seasonal HWT’ method to the original arrivals series. (Note that this method is not the same as standard Holt-Winters exponential smoothing because of the presence of the residual autocorrelation term in the single seasonal HWT formulation.)

The weakness of single seasonal HWT is that it does not capture the commonality in the intraday cycles for the five weekdays. A method that can model this common feature, as well as the intraweek cycle, is the parsimonious seasonal exponential smoothing method. For this method, we specified 34 distinct seasons to represent the intraday cycle that is common to the five weekdays, 14 distinct seasons for the Friday periods, and a further eight distinct seasons for the Saturday periods. This implies a total of $M=34+14+8=56$ seasons. We found that Version 2 (the version with period of day effect) of the method offered no improvement over Version 1 (the basic model). This suggested to us that Version 3 (the version with period of day and day type effects) would not be useful. Furthermore, we felt the simplicity of the choice of seasons and the relatively short length of the series did not warrant the use of Version 3 of the method, and so we did not implement this version.

5.4. Post-sample call centre forecasting results

We evaluated the various methods using MAE and RMSE. Percentage measures could not be used due to the existence of periods where the series was close to or equal to zero. It is important to note that, regardless of whether a method was applied to the original or manipulated arrivals series, post-sample forecast accuracy was recorded for the same set of periods, consisting of just the opening hours of the call centre. For simplicity, we present and discuss only the MAE results, as the relative rankings of the methods were similar when the RMSE was used. The MAE results are presented in Fig. 7. We do not show the results for the ANN method, as they were not competitive. This is presumably at least partly due to the relatively short length of the time series. We also do not present the results for either the restricted or unrestricted IC methods because the results for these methods were very similar to those of the HWT method.

----- Fig. 7 -----

In Fig. 7, the simplistic benchmark methods are outperformed by all the sophisticated methods, except the ARMA method that involved no differencing. Using differencing can be seen to improve the results of the ARMA method at all lead times. However, the figure shows that this ARMA method is, in turn, beaten at all lead times by the HWT method. Let us now turn to the two methods that were applied to the original arrivals series. The figure shows that single seasonal HWT offers a small improvement on the best of the methods applied to the manipulated arrivals series. For the parsimonious seasonal exponential smoothing method, we show just the results for Version 1 (the basic version) of this method because Version 2 (the version with period of day effect) offered no improvement. The figure shows that the best results at all lead times were achieved by this parsimonious seasonal exponential smoothing method. It is impressive that the method outperforms not only those methods applied to the manipulated arrivals series, but also the single seasonal HWT method, which was applied to the original data.

6. Summary and concluding comments

In comparison with the HWT method, the IC method has the appeal that it involves the initialisation and updating of fewer terms. This dimension reduction is essentially achieved by allowing the same intraday cycle to be used for days that can be assumed to be identical. However, an unappealing feature of the IC method is that it allows only whole days to be treated as identical. It often makes more sense to assume that just parts of days are identical. In this paper, we have introduced a new parsimonious seasonal exponential smoothing model for intraday data that has the flexibility to allow parts of different days of the week to be treated as identical. This flexibility leads to a model of lower dimensionality than the IC method. Post-sample forecasting results for two electricity load series showed that the new method competes well with the HWT and IC methods. Our second empirical study involved a series of call volumes from a call centre that is open for a shorter duration at the weekends than on weekdays. The HWT and IC methods cannot satisfactorily address the situation where the number of periods on each day of the week is not the same. However, it can be modelled in a straightforward way using parsimonious seasonal exponential smoothing. Post-sample

forecasting results supported the use of the new method. With the method presented as an innovations state space model, it is easy to generate prediction intervals using Monte Carlo simulation.

In terms of future work regarding the new method, it may be useful to have some form of statistical measures or procedures to support the clustering of the periods of the week into distinct seasons. Although our experimentation with statistical cluster analysis did not produce promising results, a fresh look at this with alternative data may be more fruitful. It would also be interesting to consider the use of information criteria to select among different specifications of the method. Finally, although we have used the method for data that exhibits repeating intraday and intraweek seasonal cycles, it may also be useful for data that possesses just a single repeating seasonal cycle.

Acknowledgements

We would like to thank Matthew Roberts of National Grid for supplying the British load data, and RTE for making the French load data available on their website. We acknowledge Avi Mandelbaum for making available the Israeli bank data. We are also grateful for the helpful comments of two referees.

References

- [1] Regnier, E. Doing something about the weather. *Omega* 2008;36:22-32.
- [2] Tych W, Pedregal DJ, Young PC, Davies J. An unobserved component model for multi-rate forecasting of telephone call demand: The design of a forecasting support system. *International Journal of Forecasting* 2002;18:673-695.
- [3] 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.
- [4] 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.
- [5] Christiaanse WR. Short-term load forecasting using general exponential smoothing. *IEEE Transactions on Power Apparatus and Systems*. 1971;PAS-90:900-910.

- [6] Hagan MT, Behr SM. The time series approach to short-term load forecasting. *IEEE Transactions on Power Systems* 1987;PWR-2:785-791.
- [7] 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.
- [8] Cottet R, Smith M. Bayesian modeling and forecasting of intraday electricity load. *Journal of the American Statistical Association*. 2003;98:839–849.
- [9] Weinberg J, Brown LD, Stroud JR. Bayesian forecasting of an inhomogeneous Poisson process with applications to call center data. *Journal of the American Statistical Association* 2007;102:1185-1198.
- [10] Shen H, Huang JZ. Interday forecasting and intraday updating of call center arrivals. *Manufacturing and Services Operations Management* 2008;10:391-410.
- [11] Hyndman RJ, Koehler AB, Ord JK, Snyder RD. *Forecasting with Exponential Smoothing: The State Space Approach*. Berlin Heidelberg, Germany: Springer-Verlag; 2008.
- [12] Gardner ES, Jr. Exponential smoothing: the state of the art - Part II. *International Journal of Forecasting* 2006; 22:637-666.
- [13] 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.
- [14] Taylor JW. A comparison of time series methods for forecasting intraday arrivals at a call center. *Management Science* 2008;54:253-265.
- [15] Taylor JW. Short-term electricity demand forecasting using double seasonal exponential smoothing. *Journal of the Operational Research Society* 2003;54:799-805.
- [16] Bunn DW. Forecasting loads and prices in competitive power markets. *Proceedings of the IEEE* 2000;88:163-169.
- [17] Woo C-K, Horowitz I, Olson A, Horii B, Baskette C. Efficient frontiers for electricity procurement by an LDC with multiple purchase options. *Omega* 2006;34:70-80.
- [18] Gans N, Koole G, Mandelbaum A. Telephone call centers: Tutorial, review, and research prospects. *Manufacturing and Service Operations Management* 2003; 5:79-141.

- [19] Antipov A, Meade N. Forecasting call frequency at a financial services call centre. *Journal of the Operational Research Society* 2002; 53:953-960.
- [20] Soyer R, Tarimcilar MM. Modeling and analysis of call center arrival data: A Bayesian approach. *Management Science* 2008; 54:266-278.
- [21] Taylor JW. Density forecasting of intraday call center arrivals using models based on exponential smoothing. Working paper, University of Oxford 2010.
- [22] Rasmussen R. On time series data and optimal parameters. *Omega* 2004;32:111-120.
- [23] Pedregal DJ, Young PC. Development of improved adaptive approaches to electricity demand forecasting. *Journal of the Operational Research Society* 2008;59:1066-1076.
- [24] Demyanyk Y, Hasan I. Financial crises and bank failures: A review of prediction methods. *Omega* 2010;38:315-324.
- [25] 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.
- [26] Bishop CM. *Neural Networks for Pattern Recognition*. Oxford, UK: Oxford University Press; 1997.
- [27] Sanders NR, Graman GA. Quantifying costs of forecast errors: A case study of the warehouse environment. *Omega* 2009;37:116-125.
- [28] Brown LD, Gans N, Mandelbaum A, Sakov A, Shen H, Zeltyn S, Zhao L. Statistical analysis of a telephone call center: A queuing science perspective. *Journal of the American Statistical Association* 2005;100:36-50.
- [29] Mehrotra V, Ozluk O, Saltzman R. Intelligent procedures for intra-day updating of call center agent schedules. *Productions and Operations Management* forthcoming.

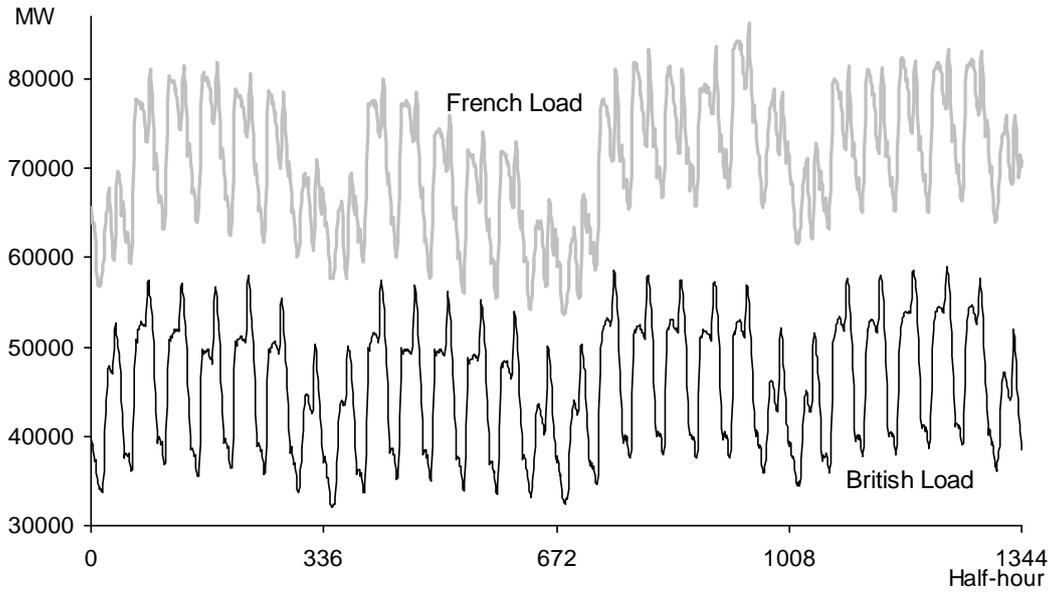


Fig. 1. Half-hourly load in Great Britain and France for the four-week period from Sunday 8 January to Saturday 4 February 2006. Each week contains $7 \times 48 = 336$ periods.

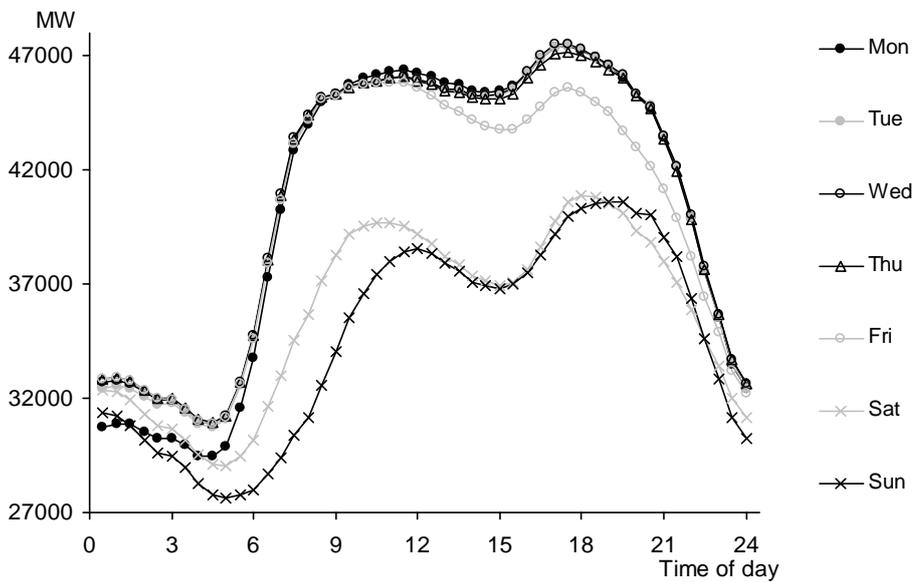


Fig. 2. Average intraday cycle for each day of the week for the British load data. Calculated using the two-year estimation sample.

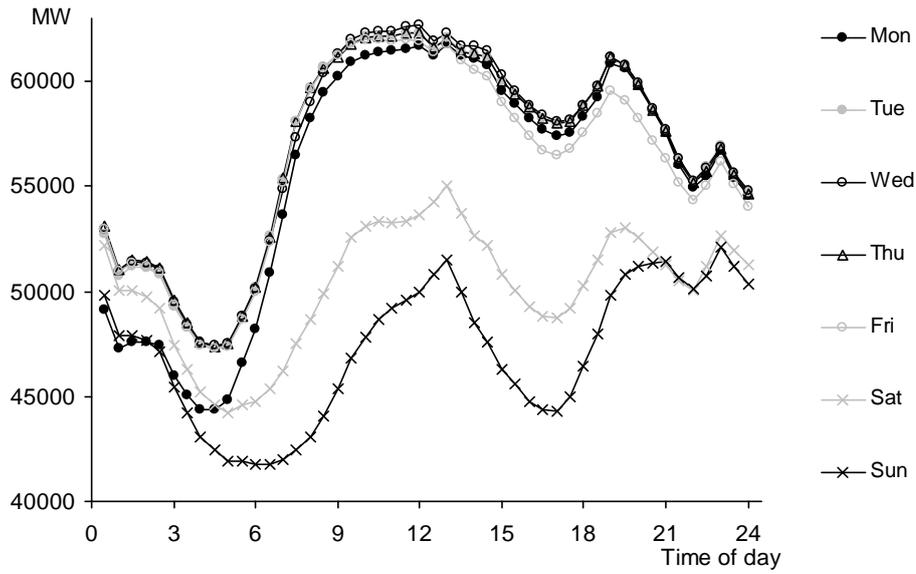


Fig. 3. Average intraday cycle for each day of the week for the French load data. Calculated using the two-year estimation sample.

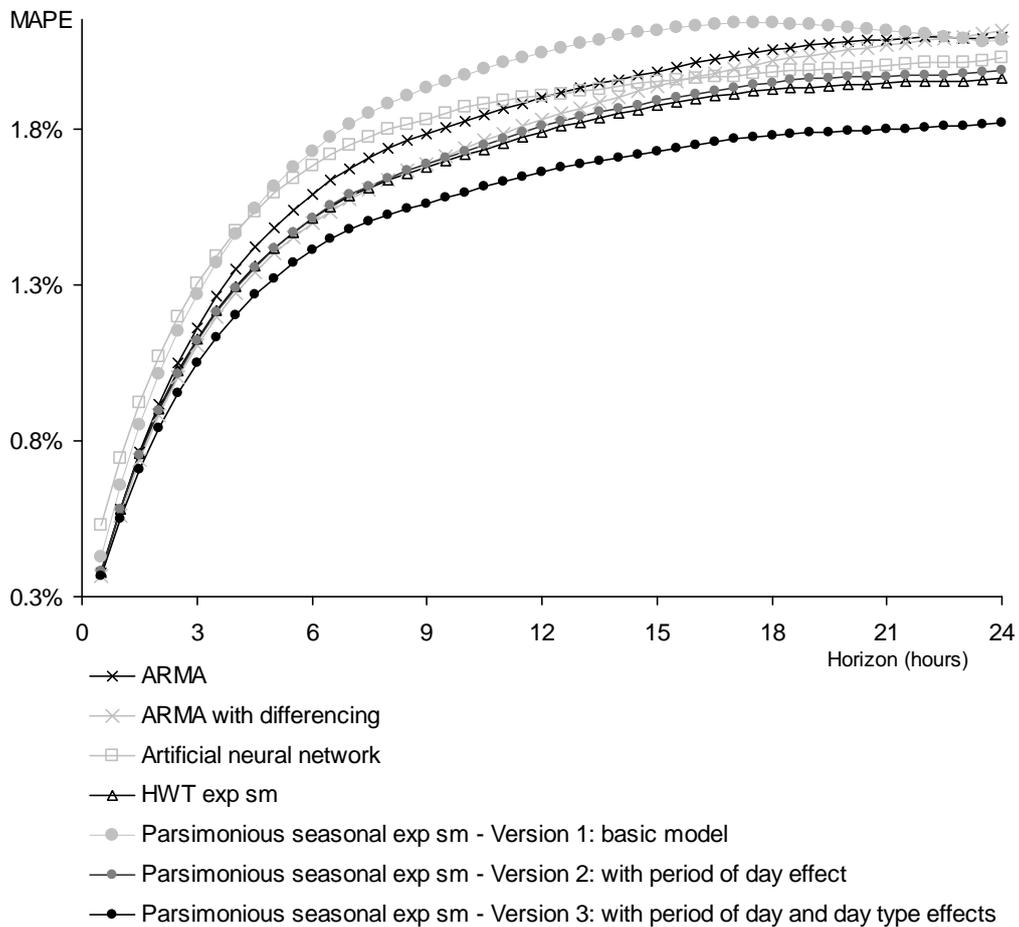


Fig. 4. Average post-sample MAPE for the British and French load series.

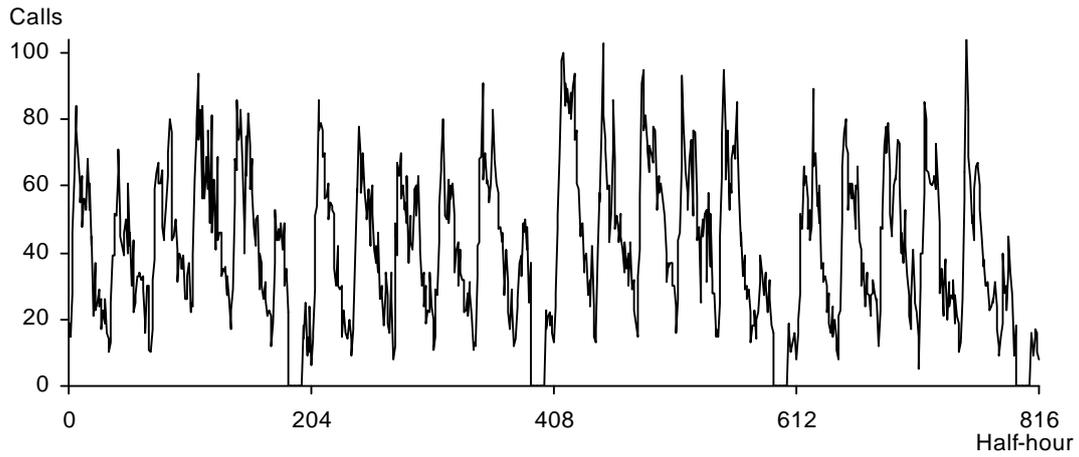


Fig. 5. Half-hourly arrivals at the Israeli bank call centre for the four-week period from Sunday 28 November to Saturday 25 December 1999. ‘Manipulated’ arrivals series is plotted. In this series, each weekend is presented as one day, and each week consists of $6 \times 34 = 204$ half-hourly periods.

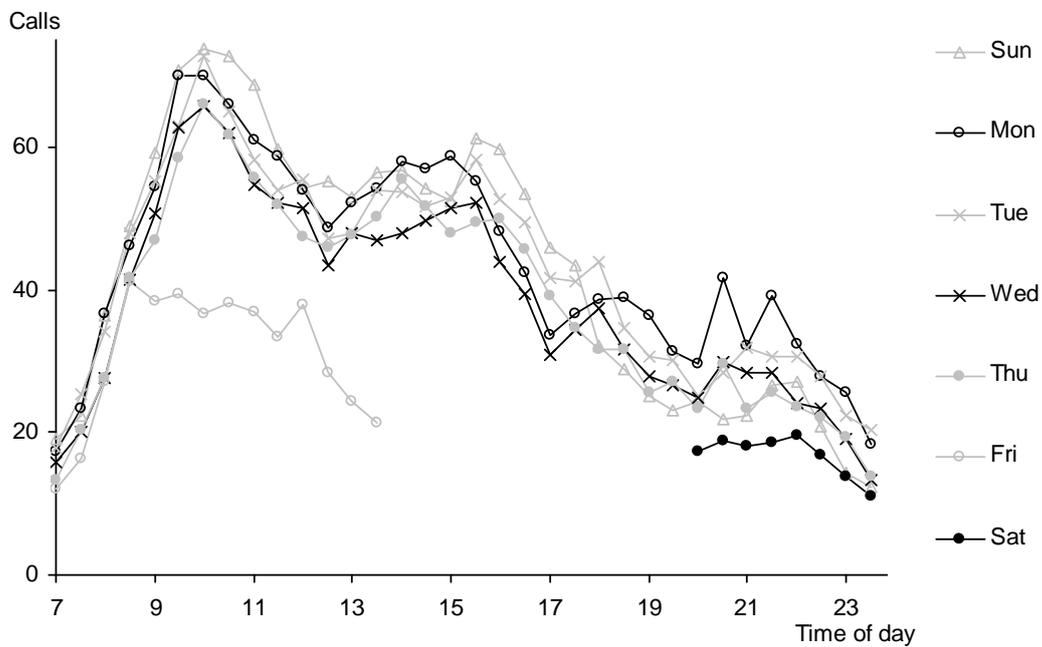


Fig. 6. Average intraday cycle for each day of the week for the Israeli bank call centre data. Calculated using the 14-week estimation sample.

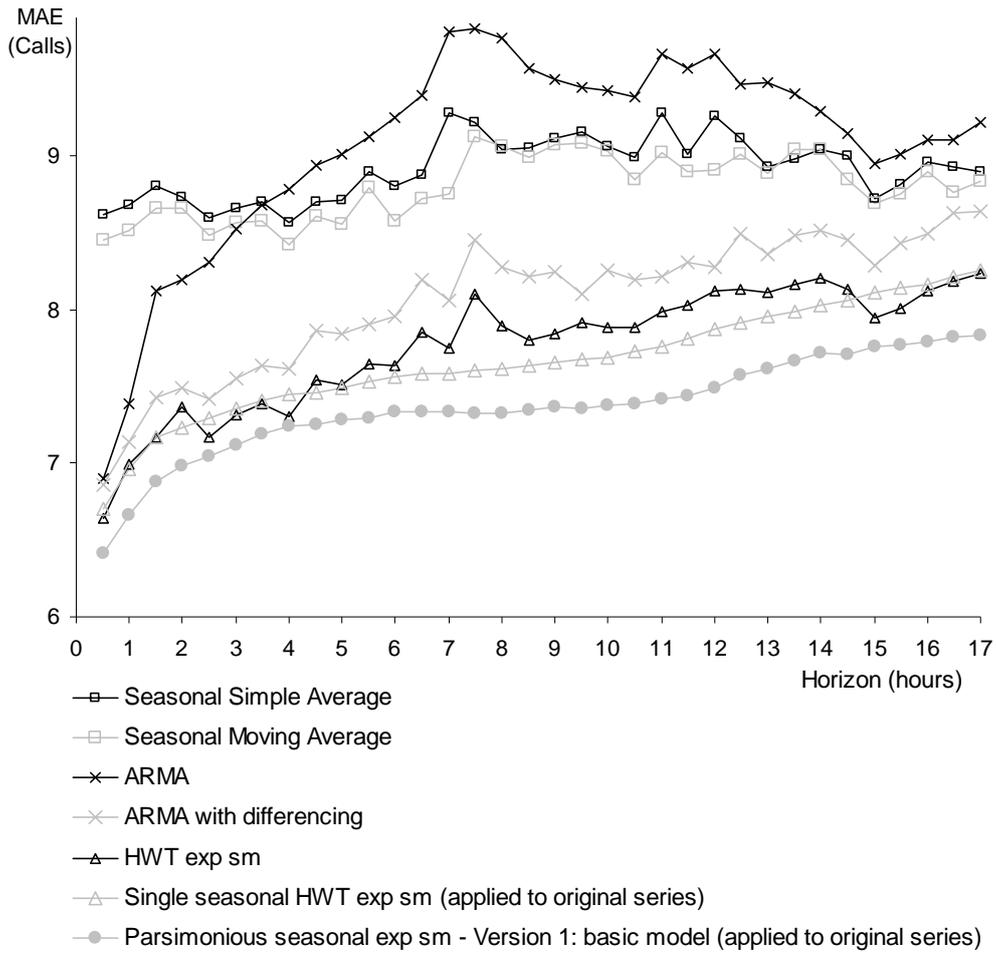


Fig. 7. Post-sample MAE for the Israeli bank call centre series. Two methods were applied to the original arrivals series, and the rest to the ‘manipulated’ arrivals series.