

**A Comparison of Univariate Time Series Methods for
Forecasting Intraday Arrivals at a Call Center**

James W. Taylor

Saïd Business School

University of Oxford

Management Science, 2008, Vol. 54, pp. 253-265.

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

**A Comparison of Univariate Time Series Methods for
Forecasting Intraday Arrivals at a Call Center**

A Comparison of Univariate Time Series Methods for Forecasting Intraday Arrivals at a Call Center

Abstract

Predictions of call center arrivals are a key input to staff scheduling models. It is, therefore, surprising that simplistic forecasting methods dominate practice, and that the research literature on forecasting arrivals is so small. In this paper, we evaluate univariate time series methods for forecasting intraday arrivals for lead times from one half-hour ahead to two weeks ahead. We analyze five series of intraday arrivals for call centers operated by a retail bank in the UK. A notable feature of these series is the presence of both an intraweek and an intraday seasonal cycle. The methods considered include seasonal ARIMA modeling; periodic AR modeling; an extension of Holt-Winters exponential smoothing for the case of two seasonal cycles; robust exponential smoothing based on exponentially weighted least absolute deviations regression; and dynamic harmonic regression, which is a form of unobserved component state space modeling. Our results indicate strong potential for the use of seasonal ARIMA modeling and the extension of Holt-Winters for predicting up to about two to three days ahead and that, for longer lead times, a simplistic historical average is difficult to beat. We find a similar ranking of methods for call center data from an Israeli bank.

Key words: call center arrivals; time series forecasting; univariate methods; seasonality

1. Introduction

Call centers have become a convenient and widely used channel through which organizations communicate with their customers. Their importance means that improvements to the management of call center operations can have substantial financial implications. Gans et al. (2003) describe the operational problems associated with call center management, and provide a thorough review of the associated research literature. They write that, for most call centers, capacity costs, and in particular human resource costs, amount to 60-70% of operating expenses. The efficient scheduling and hiring of call center staff rely heavily on accurate forecasts for the number of calls arriving at the center. These forecasts are needed for a range of lead times. Gans et al. explain that an initial schedule might be produced several weeks ahead, but this schedule will be continually updated until the scheduled day itself. On the day, it is common for call centers to track the volume of calls to enable dynamic real-time updating of the agents' deployment. Agents becoming free at short-notice, can be usefully assigned to meetings and training. The latter is particularly relevant, given the typically high turnover of call center agents.

In their assessment of future directions for call center research, Gans et al. highlight the forecasting of arrival rates as one of the most important areas. They write that the practice of time series forecasting at call centers is "still in its infancy". This is supported by the fact that there are so few time series forecasting papers included in Mandelbaum's (2006) extensive call center bibliography. With regard to the broader telecommunications area, Fildes (2002) also highlights the need for more research into time series forecasting methods. In this paper, we evaluate time series methods for forecasting intraday arrivals for lead times from one half-hour ahead to two weeks ahead.

We analyze five series of half-hourly arrivals at call centers operated by a retail bank in the UK. Following the practice of the bank, and with no access to potential explanatory variables, we confined our attention to univariate methods. As the forecasting method must

be used in an automated online system, the focus on univariate methods seems reasonable. Multivariate online systems can be very demanding in terms of inputs and require default procedures to ensure robustness. An interesting feature of the five series is the presence of both an intraweek seasonal cycle and an intraday seasonal cycle. This feature is also present in intraday electricity demand time series, which prompts consideration of methods proposed in the electricity demand literature. One such method is seasonal ARIMA modeling, which has previously been applied to call center arrivals, and another is Taylor's (2003) extension of the Holt-Winters exponential smoothing method for the case of two seasonal cycles. In addition, we consider periodic AR modeling, which allows the parameters in an AR model to vary across the seasonal cycles. We also implement a robust form of exponential smoothing based on exponentially weighted least absolute deviations regression (see Cipra 1992). Finally, we consider dynamic harmonic regression, which was applied by Tych et al. (2002), with some success, to an intraday call center series. The approach involves trigonometric regressors with time-varying parameters formulated as an unobserved component state space model.

In Section 2, we introduce the five UK bank call center time series. In Section 3, we review the literature on time series forecasting of call center arrivals, before describing the five main approaches that we consider. Section 4 evaluates the forecast accuracy of the different approaches. In Section 5, we test the robustness of our findings by comparing the relative accuracy of the methods for the Israeli bank call center arrivals series analyzed by Brown et al. (2005). The final section provides a summary and concluding comments.

2. Description of the UK Bank Time Series

In this paper, our main focus is the analysis of five time series of half-hourly arrivals at the call center of a major retail bank in the UK. The data is for the 36-week period from 3 January 2004 to 10 September 2004, inclusive. The series with by far the largest volumes

corresponds to calls received by the bank's four large call centers dealing with general customer enquiries. These centers, known as Central Telephony centers, together receive more than a quarter of a million calls each week. The other four series correspond to individual call centers, which are known as channels. Each channel specializes in a particular banking product. The average weekly volume of calls for each series, for our 36-week period, is given in Table 1. The table also shows the opening hours for the different centers.

----- Table 1 and Figures 1 to 3 -----

Figures 1 and 2 show the final four weeks of the Central Telephony series and the Channel 3 series, respectively. These four-week periods are very representative of the behavior of the two series throughout our full sample of data. The first point to note is the vast difference in the volumes of arrivals in the two series. Both series exhibit no apparent trend and very clear seasonality. A repeating intraweek cycle can be seen in each series. For the Central Telephony series, the volume of arrivals generally peaks on Mondays and is clearly much lower on Sundays. For the Channel 3 series, arrivals on both Saturdays and Sundays is noticeably lower than weekdays. Although the intraday cycle is less obvious for both series, closer inspection reveals that, at least on weekdays, there is a peak around 11am followed by a second, lower peak around 2pm. Gans et al. (2003) explain that such an intraday pattern is reasonably typical. It is interesting to note that, by comparison with the Central Telephony series, the Channel 3 series seems to possess noticeably greater volatility, when judged relative to the levels of each series. This volatility is not surprising given the low volume of calls at the Channel 3 center. Figure 2 also shows that the Channel 3 series occasionally has unusually large values, suggesting that the conditional distribution is skewed. Our comments here regarding the features of the Channel 3 series also apply to the other three channel series. The one exception to this is that the magnitude of the seasonality in the Channel 4 series changes more substantially over the sample period. This can be seen in Figure 3, which shows the Channel 4 series for the entire 36-week sample of data.

If several years of half-hourly data had been available, the modeling of a possible intrayear seasonal cycle could also have been considered. Gans et al. (2003) discuss how, in other areas of call center management, there is a lack of data, and that this is likely to be due, at least in part, to the common use of relatively simple methods. This is one reason why our collaborating UK bank has, in the past, stored only a limited history of intraday data. Another reason is that the role and organization of the bank's call centers has changed in recent years, with new centers being set up, and old ones being merged.

Prior to fitting and evaluating methods, we smoothed out the 'special days', such as bank holidays, as their inclusion is likely to be unhelpful in our comparison of methods. These days are so very unlike the rest of the year that univariate methods are generally unable to produce reasonable forecasts. In practice, interactive facilities tend to be used for special days, which allow the user to override the system offline. In our study, we replaced all special days by the average of the corresponding period in the two adjacent weeks.

We found that all five series possessed heteroskedasticity, with the variance appearing to be approximately proportional to the volume of arrivals. The existence of a relationship between the level of the series and the variance is consistent with the common assumption that call center arrivals obey a time-inhomogeneous Poisson process (see, for example, Brown et al. 2005). To address this issue further, in a manner consistent with our time series modeling approach, one would need to investigate the relationship between model estimates of the conditional mean and conditional variance. In order to reduce the heteroskedasticity in the Central Telephony series, we applied a logarithmic transformation to this series before fitting each of the forecasting methods in our empirical study. We were unable to apply the transformation to the four channel series because they contained periods with zero arrivals, as can be seen from Figures 2 and 3. Instead, for these series, we applied a square root transformation as this is sometimes used to reduce heteroskedasticity when it is proportional to the level of the series (see, for example, Makridakis et al. 1998, Section 2.7).

Brown et al. (2005) use a similar variance-stabilizing transformation in their analysis. They take the square root of the sum of the arrivals plus 0.25. Their motivation is that the resultant random variable is asymptotically Gaussian if arrivals are Poisson distributed.

We used the first 24 weeks of observations to estimate method parameters, and the remaining 12 weeks for post-sample forecast evaluation. The number of half-hours in the estimation and evaluation samples for each series is given in Table 1. We considered all forecast horizons from one half-hour ahead to two weeks ahead. For each series, we rolled the forecast origin forward through the post-sample evaluation period to produce a collection of forecasts from each method for each horizon. We opted not to re-specify or re-estimate methods, as we rolled the forecast origin forward, because we felt this would be impractical in an online forecasting implementation, for at least some of the methods that we considered.

3. Univariate Time Series Forecasting Methods

3.1. Review of Previous Studies

The literature on time series methods for forecasting call center arrivals is small (Gans et al. 2003). We were able to find only five papers that considered the choice of time series methods for this application. A method considered in all five papers is ARIMA modeling. This is essentially a univariate method that can be adapted for the inclusion of independent variables. For the prediction of daily call volumes at a retailer's call center, Andrews and Cunningham (1995) incorporate advertising effects in an ARIMA model through the use of transfer functions. In a study of daily applications for loans at a financial services call center, Antipov and Meade (2002) implement a univariate ARIMA model, but find that it is outperformed by a multivariate model involving several marketing variables.

Bianchi et al. (1998) used ARIMA modeling for forecasting daily arrivals at a telemarketing call center. They supplemented the approach with intervention analysis to control for outliers, and found that the resultant forecasts outperformed standard Holt-

Winters seasonal exponential smoothing. In our study, we also compare these two approaches, but our versions of these methods differ in that we allow for the two seasonal cycles in intraday data. Of the five time series forecasting papers, the study by Tych et al. (2002) is the only one that deals with intraday data. They consider hourly calls to a credit card company's customer service centers. The paper compares a relatively complex unobserved components state space method with a univariate ARIMA model for double seasonality. Evaluation for lead times up to a week ahead showed the state space method outperforming the ARIMA model.

Seasonal time series are often deseasonalized to enable the use of non-seasonal point forecasting methods (see, for example, Makridakis and Hibon 2000). The resultant forecasts are then reseasonalized prior to being used. In his analysis of daily emergency telephone call arrivals, Mabert (1985) considers several simple methods based on deseasonalized data, and finds that one of them is able to outperform a simple ARIMA model.

In view of the small literature on time series methods for call center arrivals, it is not surprising that simplistic methods dominate forecasting practice in this area. An approach that is often used is the 'top-down' approach (Gans et al. 2003). Indeed, it is the approach currently used by our collaborating UK bank. This method can be viewed as a version of the deseasonalizing approach to seasonal time series forecasting. If applied to half-hourly calls, the top-down approach involves summing over all half-hours in each month to give monthly totals. These totals are then used as a basis for forecasting future monthly totals. This is either done using simple time series methods or, as in the case of our collaborating UK bank, by a judgmental approach. The resulting forecast for a monthly total is then decomposed into forecasts for each week within the month, based on historical records for the proportion of the monthly totals that occur in that week. In a similar way, weekly totals are broken down into daily totals, which in turn are decomposed into forecasts for each half-hour of the day.

Although lacking in sophistication and statistical eloquence, the top-down approach has the appeal of transparency. This is an important issue. Given that their predictions form the basis of potentially emotive decisions regarding recruitment and scheduling of personnel, forecasters are unlikely to adopt a method for which they have little intuition. Of course, it is also worth acknowledging that simpler forecasting methods can sometimes turn out to be more accurate than more sophisticated alternatives.

In the remainder of Section 3, we describe our implementation of five different approaches to forecasting our five series of UK bank call center arrivals. At the end of Section 3, we list all the methods that we included in our empirical study.

3.2. Seasonal ARMA

We fitted ARIMA models of the type considered by Tych et al. (2002). For each of our five UK bank series, we followed the Box-Jenkins methodology to identify the most suitable model based on the estimation sample of 24 weeks. We considered differencing for the Channel 4 series because the magnitude of its seasonality changes to some degree over time, but the resultant model had weaker diagnostics than a model fitted with no differencing. The other four series were clearly stationary, so we concluded that differencing was not necessary for any of the series. We, therefore, change the notation from ARIMA to ARMA.

The multiplicative double seasonal ARMA model (see Box et al. 1994, p. 333) can be written as

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

where y_t represents call center arrivals transformed as described in Section 2; c is a constant term; L is the lag operator; ε_t is a white noise error term; 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, and s_1 and s_2 are the lengths of the intraday and intraweek cycles, respectively. This model can be expressed as

ARMA(p, q) $\times(P_1, Q_1)_{s_1} \times (P_2, Q_2)_{s_2}$. We estimated the model using maximum likelihood with the likelihood based on the standard Gaussian assumption. For the lower volume series, a Gaussian assumption seems unconvincing, but we use it here by default, as it is not obvious what other distribution to use. Although a time-inhomogeneous Poisson distribution might seem appropriate for the number of call center arrivals, we do not use it as a basis for the ARMA likelihood for three reasons. Firstly, we have no evidence that our arrivals possess such a distribution. Secondly, discrete versions of double seasonal ARMA models are not available. Thirdly, transformations have been applied to the arrivals series in order to reduce heteroskedasticity, as described in Section 2.

We considered lag polynomials up to order three because initial experimentation with the data suggested that higher order lags were unnecessary. We based model selection on the Schwarz Bayesian Criterion (SBC), with the requirement that all parameters were significant (at the 5% level). Table 2 presents the orders of the ARMA model selected for each of the five UK bank series. The model for the logged Central Telephony series was of the form shown in expression (1). The parameters of this model are shown in Table 3.

$$\begin{aligned} (1 - \phi_1 L - \phi_2 L^2) (1 - \phi_{32} L^{32} - \phi_{64} L^{64} - \phi_{96} L^{96}) (1 - \phi_{224} L^{224} - \phi_{448} L^{448} - \phi_{672} L^{672}) (y_t - c) \\ = (1 - \theta_2 L^2 - \theta_3 L^3) (1 - \theta_{96} L^{96}) (1 - \theta_{448} L^{448} - \theta_{672} L^{672}) \varepsilon_t \end{aligned} \quad (1)$$

----- Tables 2 and 3 -----

3.3. Periodic AR

Periodic ARMA models allow the parameters in the model to change with the seasons (see Franses and Paap 2004). For economic data, periodic models are often found to be more suitable than standard time-invariant coefficient models (e.g. Osborn et al. 1988). However, although periodic models can improve explanatory power, Ghysels and Osborn (2001, Section 6.6) comment that the evidence on their usefulness for forecasting is less clear.

To assess the potential for periodic ARMA models, we examined whether the autocorrelation at a specified lag exhibited variation across the half-hours of the day or the week. For example, for the Central Telephony series, Figure 4 shows how the autocorrelation at lag $s_2=224$ varies across the 224 half-hours of the week. The confidence intervals in the figure are based on standard error calculated as the inverse of the square root of the number of observations. Although the sample size is not sufficiently large to conclude with confidence, the variation in the autocorrelation in this plot and in the one for lag $s_1=32$, and in similar graphs for the other series, suggested to us that there was some appeal in estimating periodic ARMA models for our data. Interestingly, we found substantially less evidence of seasonal variation in the autocorrelation at lag one.

----- Figure 4 -----

Empirical studies have shown that periodic AR models tend to be sufficient, and that MA terms are unnecessary (Franses and Paap 2004, p. 28). Therefore, we considered only periodic AR models for the call center series. More specifically, we estimated periodic AR models with periodicity in the constant term and the coefficients of AR terms of lags one, s_1 and s_2 . In Section 2, we described how the forecast lead times of interest in this paper are from one half-hour to two weeks ahead. For the Central Telephony series, we found that for lead times beyond two hours ahead, substantial improvement in the accuracy of the approach can be achieved by excluding the lag one AR term from the model. Of the other four series, the inclusion of the lag one term was noticeably beneficial only for the Channel 4 series. In view of this, for simplicity, the results that we present in Sections 4 and 5, for periodic AR modeling, correspond to the use of the method with no lag one AR term. The formulation for this method is presented in the following expressions:

$$y_t = \phi_0(t) + \phi_{s_1}(t) y_{t-s_1} + \phi_{s_2}(t) y_{t-s_2} + \varepsilon_t$$

where $\phi_p(t) = \omega_p + \sum_{i=1}^4 \left(\lambda_{pi} \sin\left(2i\pi \frac{d(t)}{s_1}\right) + \nu_{pi} \cos\left(2i\pi \frac{d(t)}{s_1}\right) + \kappa_{pi} \sin\left(2i\pi \frac{w(t)}{s_2}\right) + \upsilon_{pi} \cos\left(2i\pi \frac{w(t)}{s_2}\right) \right)$

$d(t)$ and $w(t)$ are repeating step functions that number the half-hours within each day and week, respectively. For example, for the Central Telephony series, $d(t)$ counts from 1 to 32 within each day, and $w(t)$ counts from 1 to 224 within each week. ω_p , λ_{pi} , ν_{pi} , κ_{pi} and ν_{pi} are constant parameters. The periodic parameters, $\phi_p(t)$, use a similar flexible fast Fourier form to that employed by Andersen and Bollerslev (1998) in their analysis of the volatility in intraday financial returns. For simplicity, we chose to sum from $i=1$ to 4 for all five series. This choice was made based on inspection of the SBC and significance tests for the constant parameters. These parameters were estimated using OLS regression.

3.4. Exponential Smoothing for Double Seasonality

Exponential smoothing is a simple and pragmatic approach to forecasting, whereby the forecast is constructed from an exponentially weighted average of past observations. The stability and accuracy of the approach has led to its widespread use in applications where a large number of series necessitates an automated procedure, such as inventory control. This motivates its consideration for the automated online forecasting of call center arrivals.

A problem with using the standard Holt-Winters seasonal exponential smoothing method for our call center series is that it can only accommodate a single seasonal cycle. Taylor (2003) extends the formulation in order to accommodate the two seasonal cycles present in half-hourly electricity demand series. As the method involves no model specification, it has the appeal of simplicity and robustness over ARIMA modeling. The formulation for two additive seasonal cycles is given in the following expressions:

$$S_t = \alpha (y_t - D_{t-s_1} - W_{t-s_2}) + (1 - \alpha)(S_{t-1} + T_{t-1}) \quad (2)$$

$$T_t = \gamma (S_t - S_{t-1}) + (1 - \gamma)T_{t-1} \quad (3)$$

$$D_t = \delta (y_t - S_t - W_{t-s_2}) + (1 - \delta)D_{t-s_1} \quad (4)$$

$$W_t = \omega (y_t - S_t - D_{t-s_1}) + (1 - \omega)W_{t-s_2} \quad (5)$$

$$\hat{y}_t(k) = S_t + kT_t + D_{t-s_1+k} + W_{t-s_2+k} + \phi^k \left(y_t - (S_{t-1} + T_{t-1} + D_{t-s_1} + W_{t-s_2}) \right) \quad (6)$$

S_t and T_t , are the smoothed level and trend; D_t and W_t are the seasonal indices for the intraday and intraweek seasonal cycles, respectively; α , γ , δ and ω are the smoothing parameters; and $\hat{y}_t(k)$ is the k step-ahead forecast made from forecast origin t . For simplicity, we have presented, in expression (6), the forecast function for $k \leq s_1$, but it is straightforward to rewrite the expression for longer lead times. The term involving the parameter ϕ , in expression (6), is a simple adjustment for first-order autocorrelation. The parameters are estimated in a single procedure by minimizing the sum of squared one step-ahead in-sample forecast errors. The initial smoothed values for the level, trend and seasonal components are estimated by averaging the early observations.

We fitted the method to the five UK bank series using the estimation sample of 24 weeks. The estimated parameters are shown in Table 4. The values of zero for γ were accompanied by very small values for the smoothed trend, T_t , which seems reasonable since none of the series show any clear trend. For the Central Telephony and Channels 3 and 4 series, the values of the intraweek cycle smoothing parameter, ω , are noticeably higher than the corresponding values of the intraday cycle smoothing parameter, δ . This indicates that the smoothed intraweek cycle is updated to a greater degree each period than the smoothed intraday cycle. By contrast ω is zero for the Channels 1 and 2 series, indicating that the smoothed intraweek cycle is not updated from its initial values.

----- Table 4 -----

3.5. Robust Exponential Smoothing

In Section 2, we commented that the lower volume channel series exhibit sizeable volatility and occasional large values, suggesting that the conditional distribution for the series may be skewed. This motivates consideration of methods that are robust to non-Gaussian distributions and outlying observations. Cipra (1992) proposes the use of

exponentially weighted least absolute deviations (EWLAD) regression as a robust form of exponential smoothing. It is essentially exponential smoothing of the median. For a specified value of the weighting parameter, λ , the EWLAD minimization has the form:

$$\min_{\beta} \sum_t \lambda^{T-t} |y_t - m_t|$$

where T is the sample size and β is a vector of parameters in the model, m_t . This minimization can be formulated as a linear program (LP) for which the dual problem is conveniently solved. As a recursive formula does not exist for updating the estimate, the LP must be solved afresh as each new observation becomes available. With the development of computational power, this should not be a significant obstacle to implementation.

For seasonal data, sinusoidal terms can be included in the EWLAD regression. However, for our call center series, this led to poor results beyond the very early lead times. An explanation for this is that the one parameter is not adequate for smoothing the level and the various seasonal terms. In this paper, we apply the method to deseasonalized data, and use just a constant with no regressors. With just a constant used in the regression, the multiperiod forecasts for the deseasonalised data are equal to the one step-ahead prediction.

Optimization of the weighting parameter λ proceeded by the use of a rolling window of six weeks of observations to produce one step-ahead estimates for each of the remaining observations in the 24-week estimation sample. The choice of six weeks was made arbitrarily, but the procedure was not sensitive to this number. The value of λ deemed to be optimal was the value that minimized the sum of absolute one step-ahead in-sample forecast errors. We considered values for λ between 0.50 and 1, with a step size of 0.005.

3.6. Dynamic Harmonic Regression

Tych et al. (2002) develop their dynamic harmonic regression model (see Young et al. 1999) for application to intraday call center time series. Similarly to the basic structural

model of Harvey (1990), dynamic harmonic regression involves a state space model for the sum of unobserved components representing the trend and seasonal cycles. The basic form of the model is presented in expression (7).

$$y_t = T_t + \sum_i a_{it} \cos(\omega_i t) + \sum_i b_{it} \sin(\omega_i t) + \varepsilon_t \quad (7)$$

where the ω_i are appropriately chosen frequencies, and T_t , a_{it} and b_{it} are stochastic time-varying parameters. The trend component, T_t , is specified to be an integrated random walk process, as in expressions (8) and (9), with dT_t defined as the change in the trend. The parameters a_{it} and b_{it} are modeled as simple random walk processes. The Kalman filter is used to update the time-varying parameters and to generate forecasts.

$$T_t = T_{t-1} + dT_{t-1} \quad (8)$$

$$dT_t = dT_{t-1} + \eta_{dT,t} \quad (9)$$

The unknown constant parameters in the model, termed the hyper-parameters, are the variances of the error terms in expression (7) and in the stochastic models for the time-varying parameters. In state space modeling, the standard approach to the estimation of these hyper-parameters is maximum likelihood. Young et al. (1999) argue that the optimization of the likelihood function can be difficult in practice, and, instead, they propose the novel idea of performing the optimization in the frequency domain. This amounts to a least squares optimization for the difference between the empirical spectrum for the time series and the theoretical pseudo-spectrum for the state space model, which, conveniently, is a linear function of the unknown variance parameters.

The empirical spectrum for the Central Telephony series is shown in Figure 5. The spikes indicate the presence of cycles of lengths: 32/3, 32/2, 32, 224/5, 224/4, 224/3, 224/2 and 224. These cycles clearly correspond to harmonics for the intraday and intraweek seasonal cycles, which are of length $s_1=32$ and $s_2=224$, respectively. Inspection of the

empirical spectra for the other four series led us to use, for all five series, the following set of frequencies in expression (7): $2\pi j/s_1$ and $2\pi j/s_2$ for $j=1$ to 5.

----- Figure 5 -----

Tych et al. describe how the basic dynamic harmonic regression model, presented so far, was unable to adequately model the autocorrelation in their series. This was also the case for our series. To address this problem, Tych et al. fitted to the residuals of the basic model, a separate dynamic harmonic regression model for each day of the week. This involves treating all residuals, for any given day of the week, as a single series, and then estimating a dynamic harmonic regression model for this series. In our implementation of the weekday sub-models, we used the following frequencies in expression (7): $2\pi j/s_1$ for $j=1$ to 4. In our view, the need for weekday sub-models, in a second modeling stage, rather weakens the intuitive appeal of a method that, even in its basic form, is already relatively complex.

Tych et al. also present two further extensions of the method; the use of decimated data and modulated cycles. Decimation is the process by which the frequency of the data is reduced to enable the empirical spectrum to contain information on all the periodic behavior in the time series. In our analysis, we found that decimation had a detrimental effect on accuracy. Modulated cycles essentially enable cycles corresponding to the intraday and intraweek seasonality to be multiplied, rather than simply added as in the model of expression (7). The empirical spectra for our series did not exhibit the necessary features, described by Tych et al., to warrant consideration of modulated cycles.

3.7. Summary of Methods Included in the Empirical Study

We implemented six methods suitable for seasonal data, the first two of which are simplistic benchmark methods.

Method S1 - Seasonal random walk. The forecast for each lead time is the most recently observed value for the same half-hour of the week as the period to be predicted.

Method S2 - Seasonal mean. The forecast for each lead time is the mean of the arrivals for the same half-hour of the week as the period to be predicted. We used only arrivals in a moving window immediately prior to the forecast origin. We set the moving window to be of length equal to the initial estimation sample of 24 weeks.

Method S3 - Seasonal ARMA.

Method S4 - Periodic AR.

Method S5 - Exponential smoothing for double seasonality.

Method S6 - Dynamic harmonic regression.

As we discussed in Section 3.1, time series are often deseasonalized to enable the use of non-seasonal forecasting methods, with the resultant forecasts being reseasonalized prior to being used. One form of this approach is the widely used ‘top down’ approach, which we described in Section 3.1. After deseasonalizing the series using the standard approach for additive seasonality (see Harvey 1990, p. 30), we applied five methods that are applicable to data with no seasonality, the first three of which were simplistic benchmark methods.

Method D1 - Random walk. The forecast for each lead time is the value at the forecast origin.

Method D2 - Mean. The forecast for each lead time is the mean of all periods in a moving window of 24 weeks immediately prior to the forecast origin.

Method D3 - Median. The forecast for each lead time is the median of all periods in a moving window of 24 weeks immediately prior to the forecast origin.

Method D4 - Simple exponential smoothing. We derived the parameter value by the standard procedure of minimizing the sum of squared one step-ahead forecast errors.

Method D5 - Robust exponential smoothing. This is the EWLAD regression method discussed in Section 3.5 with just a constant and no regressors, and with λ optimized.

4. Results for the UK Bank Series

4.1. Results for the Central Telephony Series

We calculated the mean absolute error (MAE) and root mean squared error (RMSE) for the post-sample forecast errors from the 11 methods for each of the 448 forecast lead times (one half-hour ahead to 14 days ahead). We do not discuss here the RMSE results because the relative performances of the methods were similar to those found using the MAE. Table 5 summarizes the MAE results. Note that each column summarizes the MAE for 64 lead times. For example, the column with heading “1-2” days contains the average of the MAE for lead times of one to 64 half-hours ahead. The final column presents the average of the MAE across all lead times. The values in bold highlight the best performing method for each category of lead times.

----- Table 5 -----

The final column in Table 5 indicates that the best method overall was the relatively simple benchmark method, Method S2, which involves averaging past observations occurring on the same day of the week as the forecast period. This method outperformed all others at all lead times beyond six days ahead. Similar results were produced by Method D2, which is based on averaging all in-sample deseasonalized observations. The other methods based on deseasonalized data were not as competitive. The results for robust exponential smoothing are disappointing, but, in defense of this method, the main motivation for its inclusion was the volatile and skewed channel series.

Turning to the more sophisticated methods, we can see that the seasonal ARMA and exponential smoothing method for double seasonality performed very well for the early lead times. This is illustrated more clearly in Figure 6, where we have plotted the MAE results for all 448 forecast lead times for four of the methods. The figure shows that, although exponential smoothing for double seasonality performed poorly for the longer lead times, it was the best performing method up to about three days ahead. Tych et al. (2002) found their

dynamic harmonic regression method to be more accurate than seasonal ARMA when evaluated using the cumulative absolute error for lead times from one hour to one week ahead. Although our results do show that the two methods were similar beyond a week ahead, the dynamic harmonic regression method was clearly weaker for shorter lead times. Table 5 shows that the periodic AR model was slightly poorer than the seasonal ARMA model at all lead times.

----- Table 6 and Figures 6 and 7 -----

We also evaluated the methods using the mean absolute percentage error (MAPE) and root mean squared percentage error (RMSPE). For these measures, the relative performances of the methods were similar to those for the MAE results of Table 5. Table 6 summarizes the MAPE results, and Figure 7 graphically compares the MAPE results for the same four methods plotted in Figure 6 for the MAE. The figure shows that exponential smoothing for double seasonality was the most accurate up to approximately four days ahead.

4.2. Results for the Four Channel Series

As the four channel series each possessed periods with zero arrivals, we were unable to evaluate the forecasts for these series using percentage measures. We, therefore, focused solely on the MAE and RMSE. As with the Central Telephony series, the relative performances of the methods were similar for these two measures. In Table 7, we present the mean of the MAE results for Channels 1, 2 and 3. This averaging of the MAE results seems reasonable, as the results for the three were of a similar order of magnitude. We have had to report the MAE results for Channel 4 separately, in Table 8, because this channel operates for only six days a week.

----- Tables 7 and 8 -----

Understandably, the MAE values in Tables 7 and 8 are substantially smaller than those in Table 5 for the much higher volume Central Telephony series. Nevertheless, it is still

interesting to evaluate the relative rankings of the methods for these rather different series. Table 7 shows that, for Channels 1, 2 and 3, the results are similar for all but the two random walk methods. In Section 4.1, we described how the simplistic benchmark method S2 and D2 were two of the more accurate methods for the Central Telephony series. It is, therefore, interesting to see in Table 8 the relatively poor performance of these methods for the Channel 4 series. The notable feature of this series, that is absent in the other four series, is that the magnitude of the seasonality changes substantially over the sample period. Consequently, as we can see from Table 8, methods that do not adapt over time performed poorly. Although robust exponential smoothing, Method D5, was competitive for Channels 1 to 3, it was outperformed by standard simple exponential smoothing method, Method D4, for Channel 4. The best results for the Channel 4 series correspond to seasonal ARMA and exponential smoothing for double seasonality, with dynamic harmonic regression also performing well.

5. Analysis of the Israeli Bank Series

5.1. Description of the Time Series and Implementation of the Methods

In this section, we investigate the robustness of our findings for the UK bank data by comparing the accuracy of the forecasting methods for the Israeli bank data analyzed by Brown et al. (2005). This data is taken from a small call center 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. Following Brown et al., our analysis focuses on the single time series of total arrivals minus internet assistance calls for the final five months of 1999. More specifically, we used arrivals from 1 August to 25 December, inclusive, with the first 14 weeks used for method estimation, and the remaining seven weeks used for post-sample evaluation. We smoothed out the ‘special days’ in the series prior to method estimation. We analyzed half-hourly arrivals, in order to be consistent with our analysis of the UK bank data.

The more sophisticated methods included in our analysis of the UK bank data were designed specifically for application to data with intraday and intraweek seasonal cycles. The implementation of these methods 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. Although it is an interesting issue how best to adapt the methods for such data, in this paper, our aim is to evaluate the methods in their standard form, and so, for simplicity, we combined the two weekend days into one consisting of the same number of periods as the weekdays. The result was a series with five weekdays plus a single day for the weekend, which consisted of Friday arrivals from 7am to 2pm, zero values from 2pm to 8pm, and Saturday arrivals from 8pm to midnight. Figure 8 presents the final four weeks of the series. Using half-hourly data, the intraday seasonal cycle contains $s_1=34$ periods, and the intraweek cycle contains $s_2=6\times 34=204$ periods.

----- Figure 8 and Tables 9 and 10 -----

As with the UK bank channel series, we applied a square-root transformation to the data to reduce heterokedasticity. We implemented the methods summarized in Section 3.7. The selected seasonal ARMA model is presented in Table 9, and the parameters of the exponential smoothing method for double seasonality are given in Table 10. We used 14 weeks in the moving windows of Methods S2, D2 and D3.

5.2. Results for the Israeli Bank Series

Table 11 presents the MAE results for all 11 methods, and Figure 9 is a plot of the MAE results for the same four methods presented in Figures 6 and 7. In calculating the MAE, errors were recorded only for periods of the week when the call center was open. In essence, the relative performances of the methods are similar to those we found for the UK bank Central Telephony series. Figure 9 shows that the exponential smoothing method for double seasonality was the most accurate method for prediction up to a day ahead, but was

disappointing for longer lead times. The seasonal ARMA method also performed well for prediction up to a day ahead. The results show that beyond the early lead times, as with the UK bank series, the simple benchmark method S2 was hard to beat. By contrast with the results for the UK data, it is interesting to note that the dynamic harmonic regression was noticeably more competitive at all lead times.

----- Figure 9 and Table 11 -----

6. Summary and Concluding Comments

The main focus of this paper has been an empirical comparison of univariate time series forecasting methods for arrivals data recorded at five call centers of a UK bank. Of the five more sophisticated methods considered, seasonal ARMA modeling and dynamic harmonic regression have previously been applied to call center data, but this is not the case for the other three: periodic AR, exponential smoothing for double seasonality and robust exponential smoothing based on EWLAD regression.

The results showed that the exponential smoothing method for double seasonality was the most accurate method for short-term prediction of the high volume Central Telephony series. Beyond a week ahead the best results were achieved with a relatively simple method that is based on averaging past observations occurring on the same day of the week as the forecast period. However, the poor performance of this method for one of our series, the Channel 4 series, highlights the danger of using simplistic historical averages for series for which the magnitude of the seasonality changes substantially over the sample period. A similar comment can be made for series with trend. For such series, our results suggest the use of seasonal ARMA or exponential smoothing for double seasonality. We were interested to see how dynamic harmonic regression would perform, given the promising results of Tych et al. (2002). Our results showed that the method performed poorly for the early lead times, but that it was more competitive beyond about a week ahead. An obstacle to the

implementation of this method is its relative complexity, with its use of state space modeling and the Kalman filter, and parameter estimation in the frequency domain.

We tested the robustness of our findings by comparing the accuracy of the methods for the Israeli bank call center arrivals series analyzed by Brown et al. (2005). Our results for this data were similar to those for the UK bank Central Telephony series, with the main difference being that the seasonal ARMA and exponential smoothing method for double seasonality outperformed the simple benchmark methods only up to a lead time of two days.

In practical terms, it seems sensible to ask whether the forecasting methods that we have considered are sufficiently accurate for call center staffing purposes. There can be little doubt that greater accuracy would deliver substantial cost savings. This is apparent, for example, from Figure 7, which shows the average size of the percentage error to be greater than 8% for prediction beyond a day-ahead. As for whether greater forecast accuracy is achievable, our view is that it is. However, the message from our results is that to use more advanced methods may not be the solution. The relatively simplistic exponential smoothing method for double seasonality performed well for short horizons with the simple benchmarks dominating thereafter. Refinement of the exponential smoothing method would be an interesting line of research. When faced with two or more competing methods, an alternative to method selection is to combine the forecasts from the methods using some form of weighted average. This could certainly be considered for some of the methods in our study. Perhaps a more interesting synthesis of methods is the proposal of Avramidis et al. (2004) to try to incorporate their stochastic models within a time series method. Although we have concentrated on univariate methods, we should acknowledge that if covariates are available, their inclusion may well lead to improved forecast accuracy.

In this paper, we have focused on point forecasting. Gans et al. (2003) emphasize the importance of assessing the uncertainty in call volumes. Our intention is to focus on this issue in future work. For example, prediction intervals and density forecasts can be

calculated for ARIMA models using theoretical error variance formulae. In order to capture heteroskedasticity, these models can be extended to include a seasonal GARCH expression. A density forecast could be produced for the arrival rate, which could then be used within a doubly stochastic Poisson process (see Avramidis et al. 2004; Jongbloed and Koole 2001). Alternatively, if empirical analysis does not support such a process, the density forecast would be needed for the individual arrivals.

Acknowledgements

We are grateful to Doug Lord of the collaborating UK bank for providing the data and background information for the study. We acknowledge Avi Mandelbaum for making available the Israeli bank data. We also acknowledge the help of Robert Fildes, and the very useful comments of two anonymous referees and the Special Issue Editor.

References

- Andersen, T.G., T. Bollerslev. 1998. DM-dollar volatility: Intraday activity patterns, macroeconomic announcements and longer run dependencies. *Journal of Finance* **53** 219-265.
- Andrews, B.H., S.M. Cunningham. 1995. L.L. Bean improves call-centre forecasting. *Interfaces* **25** 1-13.
- Antipov, A., N. Meade. 2002. Forecasting call frequency at a financial services call centre. *Journal of the Operational Research Society* **53** 953-960.
- Avramidis, A.N., A. Deslauriers, P. L'Ecuyer. 2004. Modeling daily arrivals to a telephone call center. *Management Science* **50** 896-908.
- Bianchi, L., J. Jarrett, R. C. Hanumara. 1998. Improving forecasting for telemarketing centers by ARIMA modelling with intervention. *International Journal of Forecasting* **14** 497-504.

- Box, G.E.P., G.M. Jenkins, G.C. Reinsel. 1994. *Time Series Analysis: Forecasting and Control*, third edition, Prentice Hall, NJ.
- Brown, L., N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn, L. Zhao. 2005. Statistical analysis of a telephone call center: A queuing science perspective. *Journal of the American Statistical Association* **100** 36-50.
- Cipra, T., 1992. Robust exponential smoothing, *Journal of Forecasting* **11** 57-69.
- Fildes, R. 2002. Telecommunications demand forecasting - A review. *International Journal of Forecasting* **18** 489-522.
- Franses, P.H., R. Paap. 2004. *Periodic Time Series Models*. Oxford University Press, Oxford, UK.
- Gans, N., G. Koole, A. Mandelbaum. 2003. Telephone call centers: Tutorial, review, and research prospects. *Manufacturing and Service Operations Management* **5** 79-141.
- Ghysels, E., D.R. Osborn. 2001. *The Econometric Analysis of Seasonal Time Series*. Cambridge University Press, Cambridge, UK.
- Harvey, A.C. 1990. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, Cambridge, UK.
- Jongbloed, G., G. Koole. 2001. Managing uncertainty in call centers using Poisson mixtures. *Applied Stochastic Models in Business and Industry* **17** 307-318.
- Mabert, V.A. 1985. Short interval forecasting of emergency phone call (911) work loads. *Journal of Operations Management* **5** 259-271.
- Makridakis S., M. Hibon. 2000. The M3-Competition: Results, conclusions and implications. *International Journal of Forecasting* **16** 451-476.
- Makridakis S., S.C. Wheelwright, R.J. Hyndman. 1998. *Forecasting Methods and Applications*, 3rd edition. John Wiley and Sons, New York.
- Mandelbaum, A. 2006. Call centers: Research bibliography with abstracts. Technical report, Technion, Israel Institute of Technology, downloadable at

<http://iew3.technion.ac.il/serveng2006W/References/references.html>.

Osborn, D.R., S. Heravi, C.R. Birchenhall. 1988. Seasonality and the order of integration for consumption. *Oxford Bulletin of Economics and Statistics* **50**, 361-377.

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

Tych, W., D.J. Pedregal, P.C. Young, J. Davies. 2002. An unobserved component model for multi-rate forecasting of telephone call demand: The design of a forecasting support system. *International Journal of Forecasting* **18** 673-695.

Young, P.C., D.J. Pedregal, W. Tych. 1999. Dynamic harmonic regression. *Journal of Forecasting* **18** 369-394.

Table 1 Summary of the Series of Arrivals at the Call Centers of the UK Bank.

	Average calls per week	Number of days open per week	Opening hours	Half-hours in estimation sample	Half-hours in evaluation sample
Central Telephony	282,146	7	7am-11pm	5,376	2,688
Channel 1	2,817	7	7am-11pm	5,376	2,688
Channel 2	7,804	7	7am-11pm	5,376	2,688
Channel 3	4,003	7	8am-10.30pm	4,872	2,436
Channel 4	1,042	6	9am-7.30pm	3,024	1,512

Table 2 Orders of the Fitted $ARMA(p, q) \times (P_1, Q_1)_{s_1} \times (P_2, Q_2)_{s_2}$ Models For Each of the Five Transformed Series of Arrivals at the Call Centers of the UK Bank.

	p	q	P_1	Q_1	s_1	P_2	Q_2	s_2
Central Telephony	2	[2,3]	3	[3]	32	3	[2,3]	224
Channel 1	3	[3]	0	0	32	[2,3]	3	224
Channel 2	3	[3]	3	0	32	3	[2,3]	224
Channel 3	3	[3]	[1,3]	[1,3]	29	3	[2,3]	203
Channel 4	[1,2]	2	[1,3]	[1,3]	21	3	[2,3]	126

Table 3 Parameters of the Seasonal ARMA Model, Given in Expression (1), for the Log of the UK Bank Central Telephony Series. Standard Errors are in Parentheses.

Lag, i	1	2	3	32	64	96	224	448	672
ϕ_i	0.530 (0.022)	0.409 (0.024)		0.174 (0.022)	0.070 (0.021)	0.331 (0.071)	0.206 (0.023)	0.443 (0.046)	0.336 (0.044)
θ_i		0.256 (0.029)	0.053 (0.024)			0.298 (0.076)		0.332 (0.044)	0.329 (0.040)
c	6.868 (0.989)								

Table 4 Parameters of the Exponential Smoothing for Double Seasonality Method, Given in Expressions (2) to (6), for Each of the Five Transformed Series of Arrivals at the Call Centers of the UK Bank.

	α	γ	δ	ω	ϕ
Central Telephony	0.020	0.000	0.062	0.193	0.779
Channel 1	0.024	0.000	0.035	0.000	0.163
Channel 2	0.070	0.000	0.000	0.000	0.181
Channel 3	0.010	0.000	0.072	0.266	0.321
Channel 4	0.021	0.001	0.058	0.292	0.162

Table 5 MAE for the UK Bank Central Telephony Series.

Forecast Horizon (days)	1-2	3-4	5-6	7-8	9-10	11-12	13-14	All
Methods for seasonal data								
S1 - Seasonal random walk	119	119	118	131	148	145	139	131
S2 - Seasonal mean	108	109	108	108	109	106	100	107
S3 - Seasonal ARMA	94	104	107	111	119	116	110	109
S4 - Periodic AR	99	111	113	116	126	123	115	115
S5 - ES for double seasonality	90	105	118	127	142	148	144	125
S6 - Dynamic harmonic regression	118	117	116	117	121	119	115	118
Methods for deseasonalized data								
D1 - Random walk	132	150	152	149	176	175	165	157
D2 - Mean	108	109	109	109	110	107	101	108
D3 - Median	112	114	114	113	115	112	105	112
D4 - Simple ES	124	142	144	143	169	168	158	150
D5 - Robust ES	123	141	143	141	168	167	156	148

Table 6 MAPE for the UK Bank Central Telephony Series.

Forecast Horizon (days)	1-2	3-4	5-6	7-8	9-10	11-12	13-14	All
Methods for seasonal data								
S1 - Seasonal random walk	11.4	11.3	11.3	12.2	13.1	13.0	12.8	12.1
S2 - Seasonal mean	9.7	9.8	9.8	9.7	9.7	9.6	9.2	9.6
S3 - Seasonal ARMA	8.5	9.5	9.8	10.1	10.7	10.6	10.4	10.0
S4 - Periodic AR	9.3	10.1	10.2	10.4	11.1	11.0	10.6	10.4
S5 - ES for double seasonality	8.0	9.1	10.1	10.7	11.5	12.0	11.9	10.5
S6 - Dynamic harmonic regression	11.0	10.9	10.9	10.9	11.1	11.1	11.1	11.0
Methods for deseasonalized data								
D1 - Random walk	11.1	12.7	13.1	13.1	14.5	14.8	14.5	13.4
D2 - Mean	9.7	9.8	9.9	9.8	9.8	9.7	9.4	9.7
D3 - Median	9.9	9.9	10.0	9.9	9.9	9.8	9.4	9.8
D4 - Simple ES	10.5	12.1	12.5	12.5	14.0	14.3	14.0	12.8
D5 - Robust ES	10.4	12.0	12.4	12.4	13.9	14.2	13.9	12.7

Table 7 Mean MAE for the UK Bank Channel 1, 2 and 3 Series.

Forecast Horizon (days)	1-2	3-4	5-6	7-8	9-10	11-12	13-14	All
Methods for seasonal data								
S1 - Seasonal random walk	5.5	5.5	5.5	5.5	5.6	5.6	5.6	5.5
S2 - Seasonal mean	4.5	4.5	4.4	4.4	4.5	4.5	4.5	4.5
S3 - Seasonal ARMA	4.2	4.2	4.2	4.2	4.2	4.2	4.2	4.2
S4 - Periodic AR	4.4	4.4	4.4	4.4	4.5	4.4	4.4	4.4
S5 - ES for double seasonality	4.2	4.3	4.3	4.4	4.5	4.6	4.6	4.4
S6 - Dynamic harmonic regression	4.4	4.4	4.4	4.4	4.5	4.5	4.5	4.4
Methods for deseasonalized data								
D1 - Random walk	5.9	6.0	6.0	6.0	6.1	6.1	6.0	6.0
D2 - Mean	4.4	4.5	4.4	4.4	4.5	4.5	4.4	4.4
D3 - Median	4.4	4.4	4.4	4.4	4.4	4.4	4.4	4.4
D4 - Simple ES	4.5	4.7	4.7	4.7	4.8	4.8	4.7	4.7
D5 - Robust ES	4.4	4.6	4.6	4.5	4.7	4.7	4.6	4.6

Table 8 MAE for the UK Bank Channel 4 Series. As The Center Operates for Six Days Each Week, 12 Day-Ahead Prediction Corresponds to Two Weeks Ahead.

Forecast Horizon (days)	1-2	3-4	5-6	7-8	9-10	11-12	All
Methods for seasonal data							
S1 - Seasonal random walk	1.9	1.9	1.9	2.0	2.0	1.9	1.9
S2 - Seasonal mean	2.8	2.8	2.9	3.0	3.0	3.1	2.9
S3 - Seasonal ARMA	1.5	1.5	1.5	1.5	1.5	1.5	1.5
S4 - Periodic AR	2.0	2.1	2.2	2.6	2.9	3.0	2.5
S5 - ES for double seasonality	1.5	1.5	1.5	1.5	1.5	1.5	1.5
S6 - Dynamic harmonic regression	1.6	1.6	1.6	1.6	1.6	1.6	1.6
Methods for deseasonalized data							
D1 - Random walk	3.1	3.1	3.0	3.1	3.1	3.0	3.1
D2 - Mean	5.0	5.0	5.1	5.2	5.3	5.3	5.1
D3 - Median	1.9	1.9	1.9	1.9	1.9	2.0	1.9
D4 - Simple ES	2.0	2.0	1.9	2.1	2.0	2.0	2.0
D5 - Robust ES	2.8	2.8	2.7	2.8	2.8	2.7	2.8

Table 9 Parameters of the Seasonal ARMA Model for the Square Root of the Israeli Bank Series. Standard Errors are in Parentheses.

Lag, i	1	2	3	34	68	102	204	408	612
ϕ_i	0.943 (0.017)			0.107 (0.029)	0.086 (0.029)		0.312 (0.028)	0.355 (0.041)	0.239 (0.036)
θ_i	0.567 (0.037)	0.134 (0.034)						0.136 (0.047)	0.184 (0.042)
c	5.805 (1.139)								

Table 10 Parameters of the Exponential Smoothing for Double Seasonality Method, Given in Expressions (2) to (6), for the Square Root of the Israeli Bank Series.

α	γ	δ	ω	ϕ
0.045	0.000	0.023	0.246	0.378

Table 11 MAE for the Israeli Bank Series. As Each Weekend is Treated as One Day, 12 Day-Ahead Prediction Corresponds to Two Weeks Ahead.

Forecast Horizon (days)	1-2	3-4	5-6	7-8	9-10	11-12	All
Methods for seasonal data							
S1 - Seasonal random walk	11.2	11.1	11.1	10.8	10.9	11.0	11.0
S2 - Seasonal mean	9.0	8.9	8.7	8.9	8.7	8.5	8.8
S3 - Seasonal ARMA	8.6	9.1	8.9	9.1	9.0	8.9	8.9
S4 - Periodic AR	9.0	9.6	9.4	9.5	9.4	9.4	9.4
S5 - ES for double seasonality	8.6	10.3	9.7	10.2	10.4	9.9	9.8
S6 - Dynamic harmonic regression	8.8	9.4	9.1	8.6	9.0	8.8	8.9
Methods for deseasonalized data							
D1 - Random walk	10.7	12.2	12.0	11.9	12.0	11.8	11.8
D2 - Mean	9.3	9.3	9.1	9.2	9.0	8.9	9.1
D3 - Median	8.9	9.0	8.8	8.9	8.7	8.7	8.8
D4 - Simple ES	9.1	10.8	10.5	10.5	10.5	10.2	10.3
D5 - Robust ES	8.9	10.5	10.2	10.1	10.2	9.8	10.0

Figure 1 Half-Hourly Arrivals at the UK Bank Central Telephony Call Centers from Saturday 14 August to Friday 10 September 2004. The Centers are Open for 224 Half-Hours Each Week.

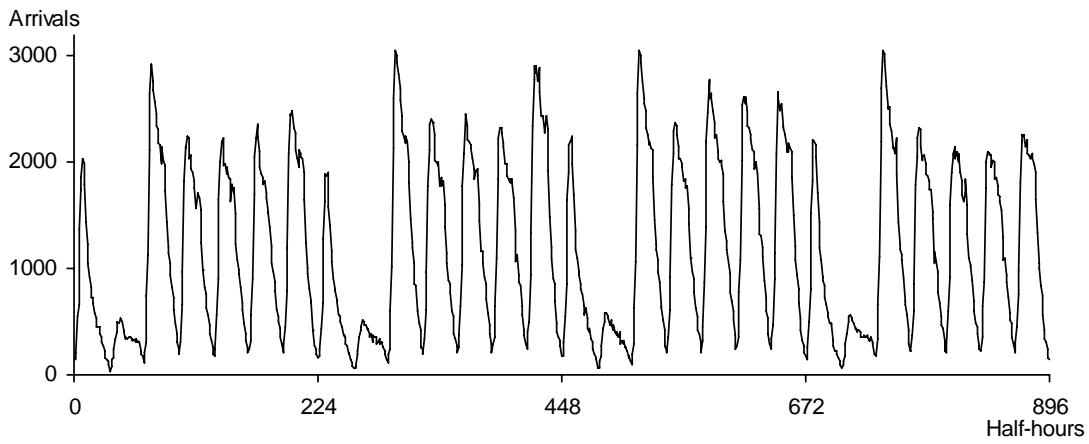


Figure 2 Half-Hourly Arrivals at the UK Bank Channel 3 Call Center from Saturday 14 August to Friday 10 September 2004. The Center is Open for 203 Half-Hours Each Week.

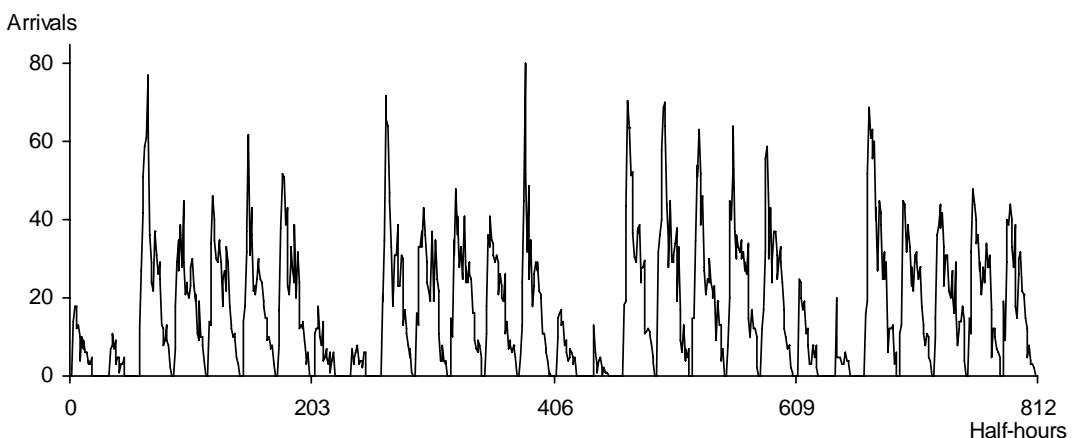


Figure 3 Half-Hourly Arrivals at the UK Bank Channel 4 Call Center from Saturday 3 January to Friday 10 September 2004. The Center is Open for 126 Half-Hours Each Week.

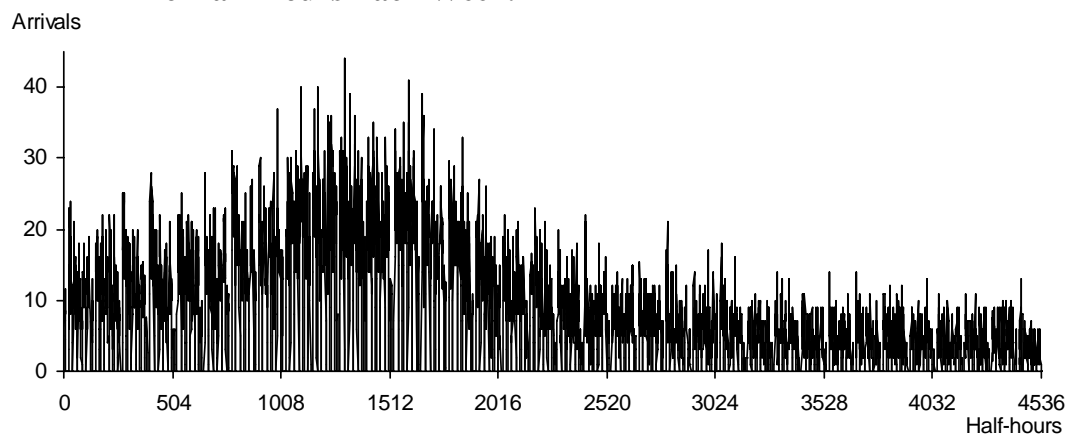


Figure 4 For the Log of the UK Bank Central Telephony Series, Lag 224 Autocorrelation at the 224 Half-Hours of the Week, Calculated using Observations from only the 24-Week Estimation Sample. The First Period on the X-Axis Corresponds to the First Half-Hour on Saturdays. Dotted Lines Indicate 95% Confidence Interval.

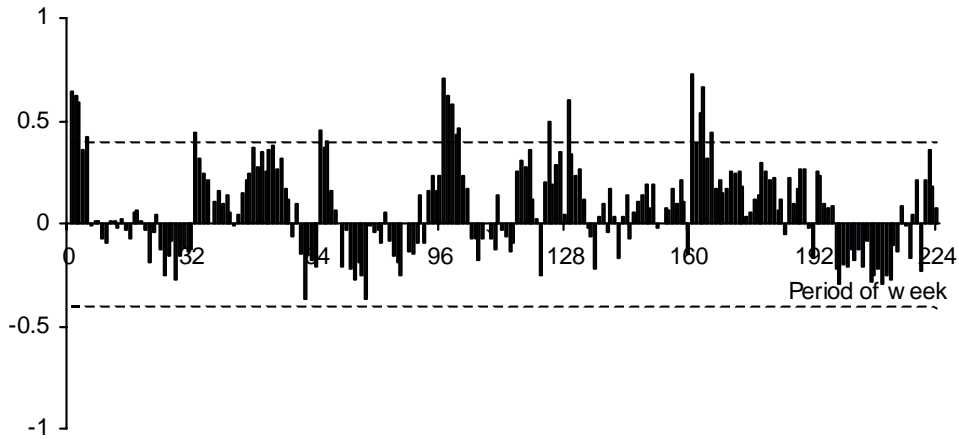


Figure 5 Empirical Spectrum for the Log of the UK Bank Central Telephony Series.

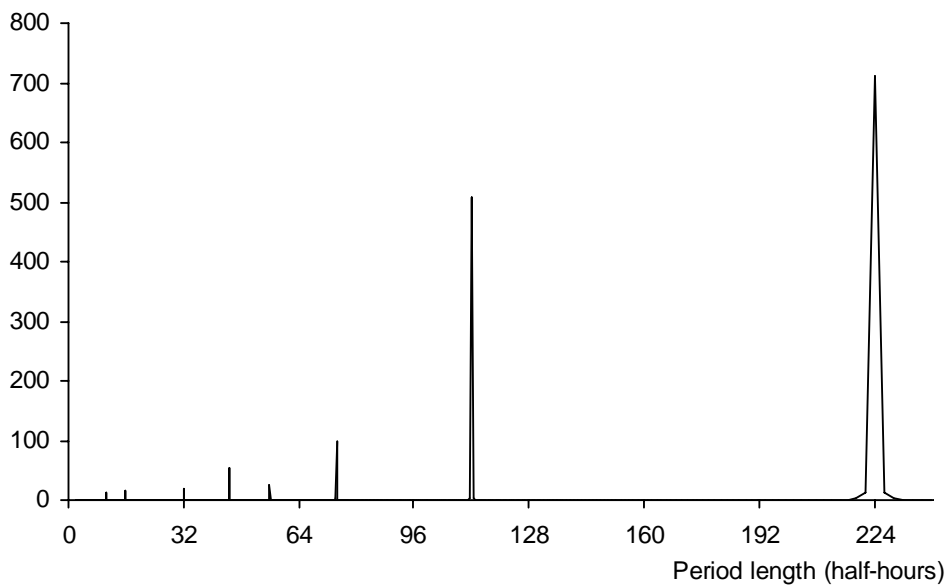


Figure 6 MAE for the UK Bank Central Telephony Series for Lead Times from One Half-Hour to 14 Days.

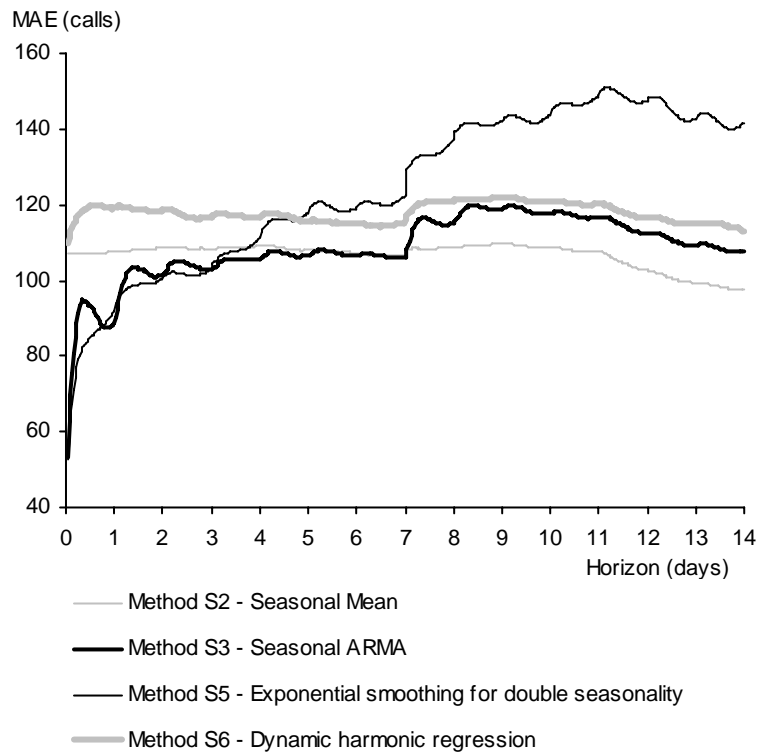


Figure 7 MAPE for the UK Bank Central Telephony Series for Lead Times from One Half-Hour to 14 Days.

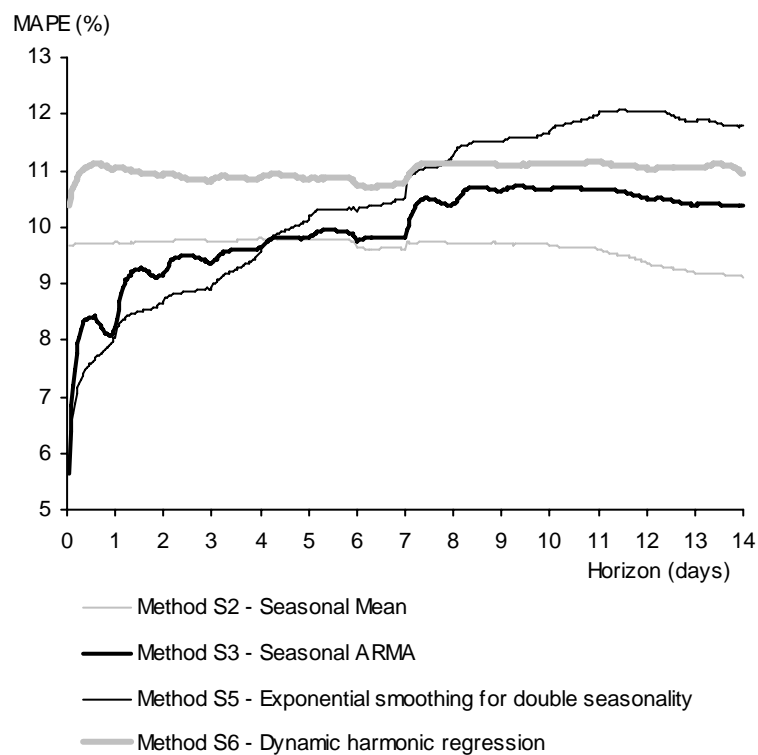


Figure 8 Half-Hourly Weekday Arrivals at the Israeli Bank Call Center from Sunday 28 November to Saturday 25 December 1999. Each Weekend is Treated as One Day, So One Week is Considered to Consist of $6 \times 34 = 204$ Half-Hours.

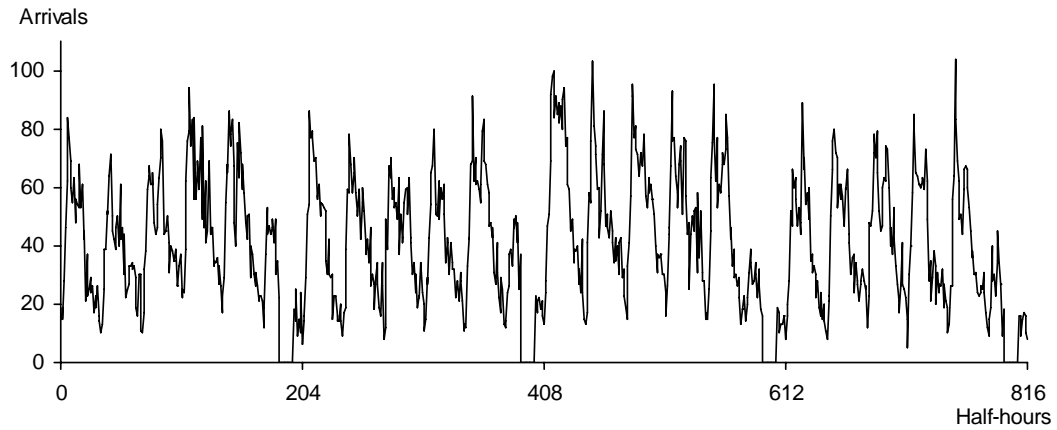


Figure 9 MAE for the Israeli Bank Series for Lead Times from One Half-Hour to 12 Days. As Each Weekend is Treated as One Day, 12 Day-Ahead Prediction Corresponds to Two Weeks-Ahead.

