Communications of the Korean Statistical Society
2011, Vol. 18, No. 2, 237–244
DOI: 10.5351/CKSS.2011.18.2.237

A Comparison of Seasonal Linear Models and Seasonal ARIMA Models for Forecasting Intra-Day Call Arrivals

Myung Suk Kim[1,a]

[a]Sogang Business School, Sogang University

Abstract

In call forecasting literature, both the seasonal autoregressive integrated moving average(ARIMA) type models and seasonal linear models have been popularly suggested as competing models. However, their parallel comparison for the forecasting accuracy was not strictly investigated before. This study evaluates the accuracy of both the seasonal linear models and the seasonal ARIMA-type models when predicting intra-day call arrival rates using both real and simulated data. The seasonal linear models outperform the seasonal ARIMA-type models in both one-day-ahead and one-week-ahead call forecasting in our empirical study.

Keywords: Seasonality, intra-day data, call centre arrival, seasonal linear model, seasonal ARIMA models.

1. Introduction

Various call volume forecasting models have been analyzed by many scholars. As an early approach, the autoregressive integrated moving average(ARIMA) models were applied to the daily call volume prediction by Bianchi *et al.* (1993). Shortly thereafter, seasonal factors were considered to be an important characteristic because multiple seasonal cycles were documented in most call arrivals. A standard Box-Jenkins seasonal ARIMA-type model (an alternative to the pure ARIMA model) has received significant attention ever since. In particular, many new call volume forecasting techniques have been evaluated by the comparison with the seasonal ARIMA models in the literature. For example, Tych *et al.* (2002) compared a dynamic harmonic regression model with a seasonal ARIMA model. Taylor (2008) compared the seasonal ARIMA with dynamic harmonic regression, periodic AR, and exponential smoothing for double seasonality. The seasonal ARIMA model has been popularly suggested as a competing model in the call volume forecasting research as well as many other seasonal time series data prediction studies (Taylor, 2003; Zhang and Qi, 2005).

The seasonal linear models or their modified versions seem to be challenging the seasonal ARIMA-type models to the role of competing models in recent call prediction literature. Weinberg *et al.* (2007) compared their Bayesian forecasting model with the seasonal linear models. Shen and Huang (2008) compared their data-driven methods via singular value decomposition, the Bayesian model (Weinberg *et al.*, 2007), and the seasonal linear models with autoregressive intra-week effect.

Regardless of popular application of both models in call prediction literature, their parallel comparison for the forecasting accuracy was not rigorously examined before. This study is designed to compare the forecasting accuracy of seasonal linear models with seasonal ARIMA-type models using both real and simulated data. Each of these models is evaluated using the real intra-day 5-minute call arrivals in a call center at a large American commercial bank. This data set includes several seasonal

This work was supported by the Sogang University Research Grant of 2008 (200810048.01).
[1] Associate Professor, Sogang Business School, Sogang University, Seoul 121-742, Korea. Email: myungkim@sogang.ac.kr

factors without trend. The seasonal linear models appear to perform better than the seasonal ARIMA-type models in both one-day-ahead and one-week-ahead forecasting in this empirical study. To verify the results in the real data analysis, models are reevaluated using the simulated call arrivals that mimic the actual processes. Similar results are obtained in the simulated call arrivals.

The remainder of this article is arranged in an organized manner. In Section 2, we describe our real call center arrival rates, as well as the simulated data. In Section 3, we explain the seasonal linear models and the seasonal ARIMA-type models, and the model evaluation criteria are also introduced. Section 4 summarizes the simulation and data analysis results. Our concluding remarks are provided in Section 5.

2. Data

2.1. Real data

According to Shen and Huang (2008), the intra-day updating is advantageous over the inter-day one due to the intra-day dependence. In particular, the prediction error can be substantially reduced if appropriate intra-day updating is available. Hence, we apply the intra-day call arrivals. The intra-day call center arrival rates of a retail banking division at a large U.S. commercial bank were examined from April 14, 2003 to October 24, 2003; they were part of the data used in Weinberg *et al.* (2007). The arriving calls were aggregated every 5 minutes for 24 hours a day, and the results showed that the call centers were most active from 7:00 AM to 9:05 PM during weekdays. Following Weinberg *et al.* (2007), we restricted our data set to this time frame. In our study, we observed 169 calls per day over a five-day business week. Since the call center was closed briefly for a holiday, we missed observations on one day. We accounted for the missed call arrivals on the holiday by utilizing the call arrivals for the corresponding weekday in the two adjacent weeks; we averaged these numbers to make the adjustment. This procedure was necessary for the seasonal ARIMA-type modeling due to its recursive structure.

We applied the transformation scheme by Brown *et al.* (2005) to our real data analysis. They suggested transforming the original call volume N_{ij} into $y_{ij} = \sqrt{N_{ij} + 0.25}$ to provide variance stabilization. According to their analysis, y_{ij} is approximately normal as the arrival counts become larger. Note that N_{ij} indicates the number of call arrivals on day i during the intra-day time interval $[t_{j-1}, t_j]$, where $i = 1, 2, \dots, I$ and $j = 1, 2, \dots, 169$. For the ARIMA-type models, N_{ij} will be replaced with N_t, where $t = j + 169 \times (i - 1)$.

2.2. Simulated data

Data sets are generated, which are mimicking the real call arrivals in the previous subsection. The following double seasonal multiplicative model is used for data generation in the simulation analysis

$$y_{ij} = S_j W_{d_i} + e_{ij}, \tag{2.1}$$

where S_j is the intra-day effect with $j = 1, 2, \dots, 169$, and W_{d_i} is the intra-week effect with $d_i = 1, 2, \dots, 5$. Note that S_j is obtained by averaging the real call arrivals for j, and the values for $S_1, S_2, \dots, S_{169}$ will be provided upon request. $W_1 = 1.1129$, $W_2 = 1.0067$, $W_3 = 0.9477$, $W_4 = 0.9480$, $W_5 = 0.9845$ are used in this study. The error term e_{ij} is generated from a normal distribution, $N(0, \sigma^2)$, with three levels of standard deviation: $\sigma = 1, 5, 10$. Due to the error term's normality, we do not have to transform the original data as we did in the real data analysis in the previous subsection. Hence, we can have $y_{ij} = N_{ij}$ in the seasonal linear models, and can use N_t or y_t instead of N_{ij} in the ARIMA-type models.

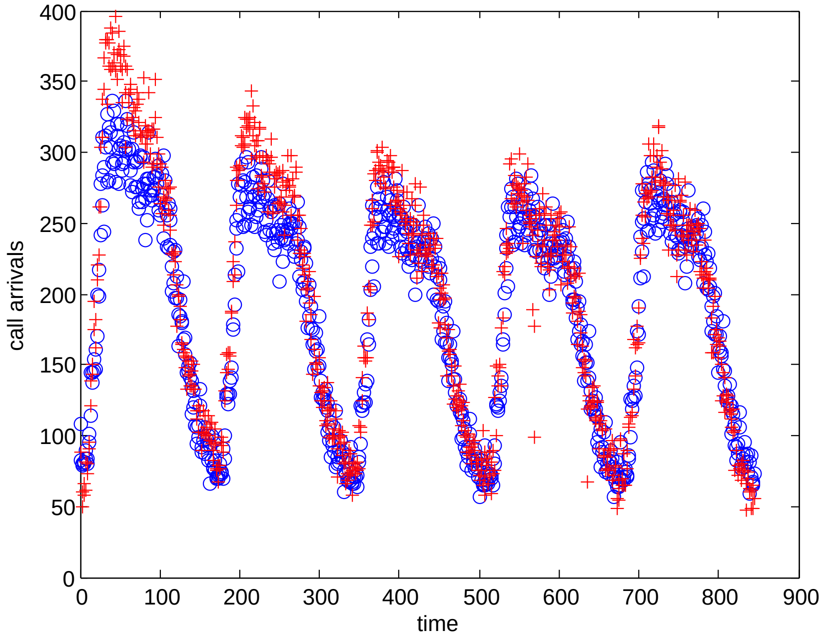

Figure 1: *Samples of call arrivals for the real data and the simulated data with $\sigma = 1$(The symbols "+" and "o" indicate the real and the simulated data, respectively)*

Table 1: Paired t-test results for the difference series between the real data and the simulated data with $\sigma = 1, 5, 10$

	American		
	d1	d2	d3
mean	0.0098	−0.0080	−0.0821
std	21.6000	22.1710	23.7362
test-stat	0.0700	−0.0552	−0.5322

Same number of real data points is generated. Note that every set of 169 data points mimics the 5-minute call arrivals per day. Figure 1 illustrates the comparison of the sample dynamics (845 time points) of the original real data and simulated data with $\sigma = 1$, where the symbols "+" and "o" indicate the real and the simulated data, respectively. The simulated data seem to mimic the pattern of the real data. The average and standard deviation of difference series between the simulated data and real data are summarized in Table 1. Notations d1, d2, d2 represent the difference series between the real data and the simulated data with $\sigma = 1, 5, 10$, respectively. To provide statistical evidences for appropriate imitation of the simulated data, paired t-test results for the difference series are reported under the normal assumption. Any significance from zero was not detected in these tests.

A similar seasonal multiplicative model was used in Zhang and Qi (2005) to generate seasonal time series data. Unlike their model, we leave out the trend term since we could not detect significant trend effects in our real call arrival data. Their model was restricted to only one seasonal factor, whereas the model in this study extends to two seasonal factors: the intra-day and intra-week seasonality. These two seasonal cycles have been employed in many call volume forecasting studies (Tych *et al.*, 2002; Taylor, 2003, 2010; Weinberg *et al.*, 2007; Shen and Huang, 2008).

3. Models and Evaluation

3.1. The seasonal linear model

Weinberg *et al.* (2007) considered the linear additive models on y_{ij} with seasonal variables and their interaction as covariates due to the approximate normal property of y_{ij}. In particular, they refer to the linear additive models as seasonal linear models. The seasonal linear model by Weinberg *et al.* (2007) is as follows:

(Model 1) $\quad y_{ij} = \mu + \alpha_{d_i} + \beta_j + \delta_{d_i j} + \varepsilon_{ij},$

where μ is a constant, and $\varepsilon_{ij} \sim N(0, \sigma^2)$. Note that d_i indicates the weekday for day i. The α_{d_i}, β_j and $\delta_{d_i j}$ indicate the day-of-week effect (intra-week effect), the time of day effect (intra-day effect), and their interaction effect, respectively. This model is a type of regression model with categorical independent variables.

In addition to these two seasonal cycles, we have considered one more seasonal component, the intra-month effect, which is included in Models 2 and 3:

(Model 2) $\quad y_{ij} = \mu + \alpha_{d_i} + \beta_j + \delta_{d_i j} + \rho_{m_i} + \varepsilon_{ij},$

(Model 3) $\quad y_{ij} = \mu + \alpha_{d_i} + \beta_j + \delta_{d_i j} + \lambda_{w_i} + \varepsilon_{ij},$

where ρ_{m_i} and λ_{w_i} indicate the day-of-month effect and the week-of-month effect, respectively. Both models indicate the intra-month effect. In particular, the day-of-month effect was considered in Gans *et al.* (2003).

3.2. The seasonal ARIMA model

The multiplicative double seasonal ARIMA model can be written as

$$\phi_p(B)\Phi_{P_1}\left(B^{S_1}\right)\Phi_{P_2}\left(B^{S_2}\right)(1-B)^d(y_t - c) = \theta_q(B)\Theta_{Q_1}\left(B^{S_1}\right)\Theta_{Q_2}\left(B^{S_2}\right)\varepsilon_t,$$

where c is a constant, ε_t is a white noise error term, and B is the back shift operator with $B^l y_t = y_{t-l}$. $\phi_p(B)$ and $\theta_q(B)$ indicate the autoregressive(AR) and moving average(MA) parts, respectively. They are polynomial functions of orders p and q, respectively. $\Phi_{P_1}(B^{S_1})$ and $\Phi_{P_2}(B^{S_2})$ indicate two seasonal factors related to the AR part; $\Theta_{Q_1}(B^{S_1})$ and $\Theta_{Q_2}(B^{S_2})$ indicate two seasonal factors related to the MA part. P_1, P_2, Q_1 and Q_2 indicate the order of corresponding polynomial functions, and s_1 and s_2 indicate the length of two seasonal cycles. Note that the ARIMA model is reduced to the autoregressive moving average(ARMA) model if $d = 0$ (if there is no differencing). A model fitted with differencing ($d = 1$) would be useful if there is evidence of any significant trend effects. We have considered a model fitted with differencing ($d = 1$) as well a model with no differencing. In this study, the double seasonal ARMA model and the double seasonal ARIMA model are called Model 4 and Model 5. The transformation $y_t = \sqrt{N_t + 0.25}$ is applied to the real data analysis for procuring the normality, whereas $y_t = N_t$ is applied to the simulated case.

We have tried to include the intra-month effect in the seasonal ARIMA-type models, but there were some difficulties. First of all, the length of a month measured by days or weeks varies with the month. Hence, one specific value corresponding to the length of the intra-month effect could not be applied to the model. Furthermore, failures tended to occur in running the computer program due to the considerable memory capacity requirement if more than two seasonal cycles are employed in the model. Therefore, no more than two seasonal factors are utilized in the ARIMA-type models.

3.3. Model evaluation

In our study, the forecasting accuracy is evaluated by computing the prediction of the root mean square error(RMSE) and the average percentage error(APE) for each day i; these are written as

$$\mathrm{RMSE}_i = \left(\frac{1}{J} \sum_{j=1}^{J} \left(N_{ij} - \hat{N}_{ij} \right)^2 \right)^{\frac{1}{2}} \quad \text{and} \quad \mathrm{APE}_i = \frac{100}{J} \sum_{j=1}^{J} \frac{\left| N_{ij} - \hat{N}_{ij} \right|}{N_{ij}},$$

where $\hat{N}_{ij}$ is the predicted call arrival rate, and $J = 169$. For the ARIMA-type models, N_{ij} and $\hat{N}_{ij}$ have been replaced with N_t and $\hat{N}_t$, respectively. These criteria have been applied in Weinberg *et al.* (2007) and Shen and Huang (2008). One-day-ahead forecasting performance is basically examined. In practice the forecasted call arrival rates are utilized for more than one-day-ahead scheduling and the staffing of a call centre; in addition, a one-week-ahead prediction is also considered.

4. Results

4.1. Real data analysis

The prediction accuracy of five models is evaluated in one-day-ahead and one-week-ahead forecasting. Historical data from April 14, 2003 to August 1, 2003 were applied for the model selection. The next-day and the next-week 5-minute call arrival rates are forecasted as the unit of a day using the selected models. Note that the term next week indicates five days ahead, since our data set defines five days as a week. The forecasting period includes August 4, 2003 to October 24, 2003. Hence, the number of prediction days is 60 days and 56 days for one-day-ahead forecasting and one-week-ahead forecasting, respectively. The historical data set or in-sample data set is updated daily by updating the next day-level forecasting evaluation. Hence, the number of in-sample data should be same.

For the seasonal ARIMA-type modeling, we considered double seasonality such as intra-day and intra-week cycles. Therefore, we used $s_1 = 169$ and $s_2 = 845$. Their corresponding periods indicate the lengths of a day and a week. For the model's lag order selection, all maximum likelihood estimates were requested to be significant at level 5%. Insignificant lags were removed from the model. Additionally, the Schwarz Bayesian criterion(SBC) was applied to the model selection. Since the evidence of trend effect was not detected in the unit-root tests, a model with no differencing considered. The selected final model is the double seasonal ARMA model or Model 4, based on the in-sample real data and is written as

$$\left(1-\phi_1 B-\phi_2 B^2\right)\left(1-\phi_{169} B^{169}\right)\left(1-\phi_{845} B^{845}\right)(y_t-c) = (1-\theta_1 B)\left(1-\theta_{169} B^{169}\right)\left(1-\theta_{845} B^{845}\right)\varepsilon_t, \tag{4.1}$$

where the coefficient estimates are reported in Table 2.

Although the evidence of unit root was not detected, we also considered the model with differencing ($d = 1$) for more complete comparison purposes. The final model is the double seasonal ARIMA model or Model 5, which uses the in-sample data expressed as

$$(1-\phi_1 B)\left(1-\phi_{169} B^{169}\right)\left(1-\phi_{845} B^{845}\right)(1-B)(y_t-c) = (1-\theta_1 B)\left(1-\theta_{169} B^{169}\right)\left(1-\theta_{845} B^{845}\right)\varepsilon_t, \tag{4.2}$$

where the coefficient estimation results are presented in Table 2.

The distributions of the RMSE and APE for each model are summarized in Table 3, which indicate the one-day-ahead and one-week-ahead forecasting results, respectively. M1–M5 represent Models 1–5, respectively. Q1, Q2 and Q3 indicate the first quarter, the second quarter, and the third quarter.

Table 2: Parameter estimates in ARIMA-type models for real and simulated data

Model	Lag, i		1	2	5	169	845
Real data Model 4	ϕ_i		1.0192^{**}	-0.0329^{*}		0.9975^{**}	0.9972^{**}
	θ_i		0.7889^{**}			0.9682^{**}	0.9836^{**}
	c	13.4838^{**}					
Real data Model 5	ϕ_i		0.0489^{**}			0.9997^{**}	0.9422^{**}
	θ_i		0.8006^{**}			0.9831^{**}	0.9078^{**}
	c	0.0012					
Simulation	ϕ_i		1.2771^{**}	-0.0776^{**}	-0.2008^{**}	0.9999^{**}	0.9967^{**}
Model 4	θ_i		1.0939^{**}		-0.1438^{**}	0.9939^{**}	0.9118^{**}
($\sigma = 1$)	c	185.9270^{**}					
Simulation	ϕ_i		1.2644^{**}	-0.0655^{**}	-0.2004^{**}	0.9999^{**}	0.9882^{**}
Model 4	θ_i		1.1357^{**}		-0.1658^{**}	0.9886^{**}	0.9468^{**}
($\sigma = 5$)	c	189.3130^{**}					
Simulation	ϕ_i		1.2473^{**}	-0.0487^{**}	-0.2006^{**}	0.9998^{**}	0.9920^{**}
Model 4	θ_i		1.1473^{**}		-0.1713^{**}	0.9871^{**}	0.9716^{**}
($\sigma = 10$)	c	191.5248^{**}					

Note: +: $p < 0.01$; *: $p < 0.005$; **: $P < 0.0001$

Table 3: Forecasting results of one-day-ahead and one-week-ahead for the real data

Model	One-day-ahead					One-week-ahead				
	M1	M2	M3	M4	M5	M1	M2	M3	M4	M5
(a) RMSE										
mean	20.759	21.386	20.540	22.126	25.243	20.441	20.862	20.195	21.183	25.525
min	11.855	12.418	12.444	11.293	12.002	11.855	12.511	13.184	12.040	11.882
Q1	15.897	15.704	15.865	15.672	16.589	15.783	15.634	15.713	15.782	16.745
Q2	18.953	18.472	18.097	20.138	21.805	18.623	17.602	18.196	19.014	20.349
Q3	22.298	24.188	22.472	24.518	28.541	21.824	24.086	21.571	24.422	26.781
max	69.759	61.739	63.699	65.274	113.169	69.759	62.240	62.079	68.436	116.600
(b) APE										
mean	9.393	9.653	9.205	10.174	11.842	9.392	9.563	9.241	10.086	12.314
min	5.666	6.236	6.286	5.379	5.801	5.666	6.194	6.390	5.456	5.637
Q1	7.355	7.290	7.334	7.652	7.951	7.224	7.152	7.270	7.812	8.025
Q2	8.453	8.096	8.020	9.366	9.633	8.128	8.009	8.010	9.142	9.630
Q3	9.886	10.637	9.840	10.721	14.202	9.886	10.705	9.915	10.879	13.083
max	29.943	27.302	27.571	34.616	55.246	29.943	27.370	27.900	35.109	59.020

The bold letters denote the smallest values among the models for each statistic regarding the RMSE and APE. M4 documented the smallest minimum RMSE and APE among the five models. However, the seasonal linear models seem to provide more accurate forecasting results than the seasonal ARMA or ARIMA model by documenting smaller mean or median RMSE and APE in Table 3.

Note that M1 includes the same two cycles that M4 and M5 do. When comparing these three models, M1 still provides the most accurate prediction results, as it documents the smaller mean and median of RMSE and APE than M4 and M5. Among all five models, M3 provides the smallest mean and median of RMSE and APE: The intra-month effect seems to be suspected. Similar results were observed in Table 3. M3 provides the smallest mean RMSE and APE, and M1 provides a smaller mean and median of RMSE and APE than M4 and M5. In summary, the seasonal linear models appear to outperform the seasonal ARMA or ARIMA model. Note that a parallel comparison for one-day-ahead and one-week-ahead forecasting accuracy is not available due to the different prediction sample sizes for two cases.

4.2. Simulated data analysis

To confirm the results in the real data analysis, we conducted a simulation analysis. Using the data

Table 4: Forecasting results of one-day-ahead and one-week-ahead for the simulated data

	One-day-ahead						One-week-ahead					
	(σ = 1)		(σ = 5)		(σ = 10)		(σ = 1)		(σ = 5)		(σ = 10)	
Model	M1	M4	M1	M4	M1	M4	M1	M4	M1	M4	M1	M4
(a) RMSE												
mean	1.035	1.274	5.166	6.808	10.358	12.270	1.036	1.227	5.177	6.661	10.369	12.051
min	0.933	1.011	4.531	4.715	9.043	8.824	0.933	0.961	4.531	4.726	9.043	8.796
Q1	1.003	1.080	4.979	5.288	10.066	10.454	1.003	1.050	4.979	5.227	10.107	10.454
Q2	1.035	1.159	5.121	6.104	10.418	11.389	1.037	1.167	5.127	6.070	10.418	11.274
Q3	1.061	1.384	5.290	7.686	10.705	12.761	1.065	1.258	5.313	6.745	10.705	12.500
Max	1.238	2.042	6.578	12.995	11.966	20.269	1.238	1.720	6.578	11.192	11.966	17.297
(b) APE												
mean	0.530	0.628	2.661	3.341	5.446	6.205	0.529	0.601	2.662	3.246	5.439	6.107
min	0.421	0.496	2.188	2.492	4.287	4.483	0.421	0.467	2.188	2.391	4.287	4.533
Q1	0.502	0.560	2.534	2.768	5.047	5.442	0.502	0.538	2.534	2.715	5.047	5.455
Q2	0.532	0.594	2.646	3.198	5.450	6.103	0.532	0.592	2.646	3.215	5.395	6.095
Q3	0.556	0.664	2.803	3.677	5.767	6.800	0.556	0.657	2.811	3.475	5.767	6.697
Max	0.627	0.890	3.101	5.587	6.948	8.391	0.627	0.748	3.101	4.730	6.948	7.681

generation model in (2.1), we generated three simulation sets of same number of data points as that of the real data. Each simulated data sets are obtained using three different values. Since the data generating procedures do not include the intra-month effect, we did not consider M2 or M3. Moreover, M5 was not considered since there were no trend effects in model (2.1). Therefore, we only compared M1 and M4 in the simulation.

For the seasonal ARMA modeling, the SBC and the 5% significance level were applied to the final model selection. The forms of the final models that were selected for the simulated data were similar to (4.1) and (4.2), but there were some significant higher-order lags unlike the real data analysis case. Maximum likelihood estimation results of M4 for three different σ values are reported in Table 2.

Table 4 summarizes the one-day-ahead and one-week-ahead forecasting results regarding the simulated data for three different σ values. Here, the seasonal linear model M1 is compared with the seasonal ARMA model M4. Our results indicate that M1 tends to outperform M4 for all three σ values. The means and medians of RMSE and APE of M1 were always smaller than those of M4. Note again that the parallel comparison between the one-day-ahead and one-week-ahead forecasting accuracy is not possible due to the different prediction sample sizes for the two cases.

5. Conclusion

The seasonal linear models challenge the seasonal ARIM-type models to the part of competing models in call arrival prediction studies. However, their forecasting accuracy was not rigorously compared before. Therefore, this study compared the seasonal linear models with the seasonal ARMA or ARIMA models in forecasting intra-day call center arrivals. Both the simulated data and the real data were used to evaluate the forecasting models. The seasonal linear models are better than the competing seasonal ARIMA-type models in this empirical analysis.

Regarding the simulation scheme, Curry (2007) pointed out that the ARIMA-type model is not appropriate when the true dynamics are generated from a multiplicative time series model like (2.1) due to the specification problems. As Zhang and Qi (2005) claimed that the multiplicative time series models seem to be able to mimic real seasonal time series data. Their arguments appear to be plausible in our study according to Figure 1 and Table 1.

Our study is limited to an empirical analysis; however, our results provide some possibility that the seasonal linear models can compete with the seasonal ARIMA-type models for versatile seasonal time series data forecasting the underlying series include seasonality without trend. We expect that the seasonal linear models will be more popularly applied to a wide range of future seasonal time series data forecasting articles.

References

Bianchi, L., Jarrett, J. and Hanumara, R. C. (1993). Forecasting incoming calls to telemarketing centers, *Journal of Business Forecasting Methods and Systems*, **12**, 3–12.

Brown, L. D., Gans, N., Mandelbaum, A., Sakov, A., Shen, H., Zeltyn, S. and Zhao, L. H. (2005). Statistical analysis of a telephone call center: A queueing-science perspective, *Journal of the American Statistical Association*, **100**, 36–50.

Curry, B. (2007). Neural networks and seasonality: Some technical considerations, *European Journal of Operational Research*, **179**, 267–274.

Gans, N., Koole, G. and Mandelbaum, A. (2003). Telephone call centers: Tutorial, review, and research prospects, *Manufacturing and Service Operations Management*, **5**, 79–141.

Shen, H. and Huang, J. Z. (2008). Interday forecasting and intra-day updating of call center arrivals, *Manufacturing and Service Operations Management*, **10**, 391–410.

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

Taylor, J. W. (2008). A comparison of univariate time series methods for forecasting intra-day arrivals at a call center, *Management Science*, **54**, 253–265.

Taylor, J. W. (2010). Exponentially weighted methods for forecasting intra-day time series with multiple seasonal cycles, *International Journal of Forecasting*, **26**, 627–646.

Tych, W., Pedregal, D. J., Young, P. C. and Davies, J. (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.

Weinberg, J., Brown, L. D. and Stroud, J. R. (2007). Bayesian forecasting of an inhomogeneous Poisson process with applications to call center data, *Journal of the American Statistical Association*, **102**, 1185–1198.

Zhang, G. P. and Qi, M. (2005). Neural network forecasting for seasonal and trend time series, *European Journal of Operational Research*, **160**, 501–514.

Received October 2010; Accepted January 2011