Nigerian Journal of Basic And Applied Sciences (2008): 16 (1): 36-47 ***ISSN 794-5696***

Forcasting of Temperature in Sokoto Metropolis Using Seasonal Modelling

Yakubu M. and Gulumbe S.U.
Department of Mathematics
Usmanu D/fodiyo University, Sokoto

Abstract: In this paper, the results of seasonal modeling and subsequent forecasting of Sokoto monthly average temperature have been obtained using seasonal autoregressive integrated moving average modeling approach. Based on this seasonal modeling analysis, we conclude that the SARIMA(1,0,0)(0,1,1)12 ,SARIMA (3,0,1)(4,1,0)12 and SARIMA (4,0,2)(5,1,1)12 models are adequate for a good description of temperature in Sokoto. We further asses their forecastibility from the out-of-sample forecast statistic, Results show that for the short forecast statistics SARIMA (3,0,1)(4,1,0)12 model minimizes the mean squares error of the forecast, while the middle forecast and long forecast statistics results have shown that SARIMA (4,0,2)(5,1,1)12 model has optimal forecast to the Sokoto temperature, hence this models have the advantage of capturing and describing and forecasting the seasonal dynamics of Sokoto city temperature.
Key Words: Seasonality, SARIMA, Identification, Estimation, Diagnostics test, and forecasting.

INTRODUCTION

Sokoto is a city located in the extreme northwest of Nigeria. The location of Sokoto in Nigeria is at Latitude 13^0 02 N and Longitude 05^0 15 E. Sokoto State is in the dry Sahel, surrounded by sandy Savannah and isolated hills. Sokoto as a whole is very hot area. The warmest months are February to April when daytime temperature is rising. The raining season is from June to October during which showers are a daily occurrence. From late October to February, during the cold season, the climate is dominated by the Hamattan wind blowing Sahara dust over the land. The dust dims the sunlight there by lowering temperatures significantly and also leading to the inconvenience of dust everywhere in houses.

Seasonality modeling is a major interested area in Univariate time series modeling. These models have the advantage that behavioral patterns can be predicted simply by analyzing the past history of a variable, reflecting these patterns. The most important aspect of building such a model is learning about the intrinsic time patterns of a variable or its underlying generating process.

Seasonality is the systematic although not necessarily regular, intra-year movement caused by the changes of weather, the calendar, and timing of decisions, directly or indirectly through the production and consumption decisions made by the agents of the economy (Hylleberg, 1990). These decisions are influenced by endowments, the expectations and preferences of the agents, and the production techniques available in the economy. Such seasonal patterns can be observed for many macroeconomic time series like gross domestic product, unemployment, weather, industrial production or construction. Temperature is indispensable for sustaining life. Even a brief rise and falling of it can cause a serious effect on human and his economic activities. However, the term seasonality is also used in a broader sense to characterize time series that show specific patterns that regularly recur within fixed time intervals (e.g. a year, a month or a week). For example it is very cold and dusty during the hamattan period compared to other time period within the year and this pattern will be the same for each year.

Statistical Software: We used **ASTSA and Gretl** (version 6.0) for the analysis. These softwares are also used to plot the graphs and the autocorrelation function

SARIMA MODELLING

The multiplicative seasonal autoregressive integrated moving average model, SARIMA is denoted by SARIMA v *(P,d,q) (P,D,Q)s* (Box and Jenkins 1976), where P denotes the number of autoregressive terms, q denotes the number of moving average terms and d denotes the number of times a series must be differenced to induce stationarity. P denotes the number of seasonal autoregressive components, denotes the number of seasonal moving average terms and D denotes the number of seasonal differences required to induce stationarity and is the period. The seasonal autoregressive integrated moving average model has the following representation:

$$(1-L)^d (1-L^s)^D \phi(L)\Phi_s(L)X_t = a + \theta(L)\Theta_s(L)e_t \qquad (2.01)$$

where:

a is a constant,

$\{e_t\}$ is a sequence of uncorrelated normally distributed random variables with the same mean (μ) and the same variance (σ^2)

L is the lag operator defined by

$$\theta(L) = 1 + \theta_1 L + \theta_2 L^2 + \dots + \theta_q L^q$$

$$.\phi(L) = 1 - \phi_1 L - \phi_2 L^2 - \dots - \phi_p L^p$$

$$\Phi(L) = 1 - \Phi_{1s} L^{1s} - \Phi_{2s} L^{2s} - \dots - \Phi_{Ps} L^{Ps}$$

$$\Theta(L) = 1 + \Theta_{1s} L^{1s} + \Theta_{2s} L^{2s} + \dots + \Theta_{Qs} L^{Qs} .$$

The selection of the appropriate seasonal ARIMA model for the data is achieved by an iterative procedure based on three steps (Box *et al*, 1994). This is shown in fig 2.01 :

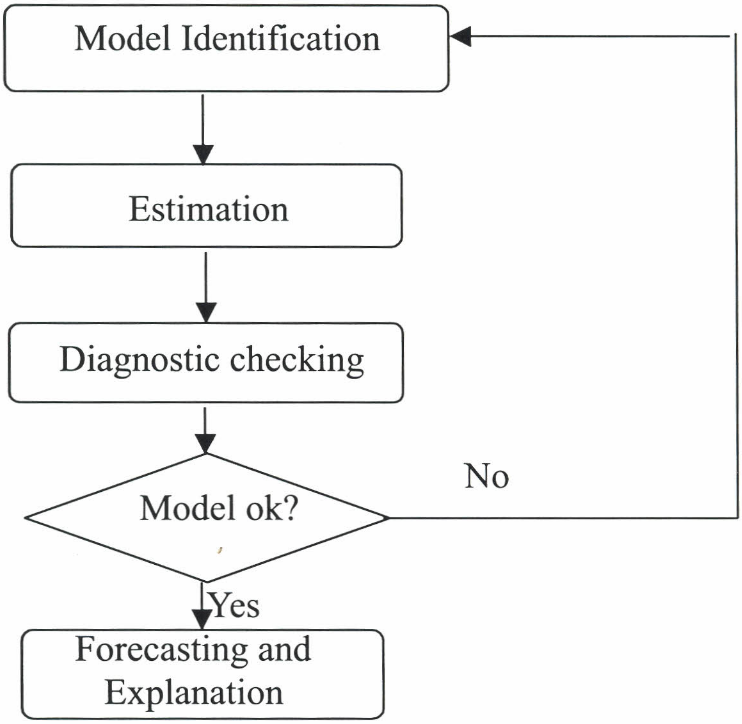

Fig 2.01: Flow chart illustrating modeling strategies

Model Identification

The Model Identification stage enables us to select a subclass of the family of SARIMA models appropriate to represent a time series. This involves stationary transformation, regular differencing, seasonal differencing and the Unit root and Stationarity tests (ADF, KPSS and HYGY).

Stationary transformations: Our task here is to identify if the time series could have been generated by a stationary process. First, we use the time plot of the series to analyze if it is variance stationary. The series departs from this property when the dispersion of the data varies with time. In this case, the stationarity in variance is achieved by applying the appropriate Box- Cox transformation (Box and Jenkins, 1976):

$$X_t^{(\lambda)} = \begin{cases} \frac{X_t^{\lambda}-1}{\lambda} & \lambda\neq0 \\ \ln(X_t) & \lambda=0 \end{cases}$$

And as a result, we get the series . $X_t^{(\lambda)}$

In some cases, especially when variability increases with level, such series can be transformed to stabilize the variance before being modeled with the Univariate Box and Jenkins method. A common

transformation involves taking the natural logarithms of the original series. The second part is the analysis of the stationarity in mean. The instruments are the time plot, the sample correlograms (ACF and PACF) and the tests for unit roots and stationarity. The path of a nonstationary series usually shows an upward or downward slope or jumps in the level whereas a stationary series moves around a unique level along time. The sample autocorrelations of stationary processes are consistent estimates of the corresponding population coefficients, so the sample correlograms of stationary processes go to zero for moderate lags.
When the series shows nonstationary patterns, we should take first differences and analyze if $\Delta X_t^{(\lambda)}$ is stationary or not in a similar way. This process of taking successive differences will continue until a stationary time series is achieved.

Regular differencing : To difference a data series, we define a new variable (W_t) which is the change in Z_t from one time period to the next; that is,

$$W_t = (1-L)Z_t = Z_t - Z_{t-1} \quad , \qquad t = 1,2,...,n \tag{2.02}$$

This working series" W_t is called the first difference of e_t. If the first differences do not have a constant mean, we might try a new W_t, which will be the second differences of Z_t, that is:

$$W_t = (Z_t - Z_{t-1}) - (Z_{t-1} - Z_{t-2}) = Z_t - 2Z_{t-1} + Z_{t-2},$$

Using the lag operator as shorthand (1-L) is the differencing operator since $(1-L)Z_t = Z_t - Z_{t-1}$. Then, in general, $W_t = (1-L)^d Z_t$ is a d-*th* order regular difference. That is, d denotes the number of nonseasonal differences.
Seasonal differencing: For seasonal models, seasonal differencing is often useful. For example,

$$W_t = (1-L^{12})Z_t = Z_t - Z_{t-12} \tag{2.03}$$

Equation (2.03) is a first-order seasonal difference with period 12, as would be used for monthly data with 12 observations per year. Rewriting (2.03) and using successive resubstitution (i.e., using $W_{t-12} = Z_{t-12} - Z_{t-24}$) gives

$$\begin{aligned} Z_t &= Z_{t-12} + W_t \\ &= Z_{t-12} + W_{t-12} + W_t \\ &= Z_{t-36} + W_{t-24} + W_{t-12} + W_t \end{aligned}$$

and so on. This is a kind of "seasonal integration ", in general , $W_t = (1-L^s)^D Z_t$ is a Dth order seasonal difference with period S where D denotes the number of seasonal differences.

Unit Roots and Stationarity: Because the order of integration of a time series is of great important for the analysis, a number of statistical tests have been developed for investigating it. In this case, we have to test the data, to know the level or if there is any need for seasonal and nonseasonal differencing before modeling the data. The following are the unit roots and stationarity tests:

? **Augmented Dickey-Fuller (ADF) Test**

This test was first introduced by Dickey and Fuller (1979) to test for the presence of unit root(s). The regression model for the test is given as:

$$\Delta X_t = \phi X_{t-1} + \sum_{j=1}^{p-1} \alpha_j^* \Delta X_{t-j} + u_t \qquad (2.04)$$

in this model the pair of hypothesis

$$H_0 : \phi = 0 \qquad \text{Versus} \qquad H_1 : \phi < 0$$

H_0 is rejected if the t-statistics is smaller than the relevant p-values (critical value). If ϕ =0 (that is, under H_0) the series X_t has a unit root and is nonstationary, whereas it is regarded as stationary if the null hypothesis is rejected.

? **KPSS Test**

This test (KPSS) has been proposed by Kwiatkowski *et al* (1992) where the hypothesis that the Data generating process (DGP) is stationary is tested against a unit root. The Data generating process is given by

$$X_t = y_t + z_t \qquad (2.05)$$

where $y_t = A_1 y_{t-1} + \ldots + A_p y_{t-p} + v_t \qquad v_t : i.i.d(0, \sigma_v^2)$

They proposed the following statistics:

$$\text{T-statistic}\,(t_k) = \frac{1}{T^2} \sum_{t-1}^{T} \frac{S_t^2}{\hat{\sigma}_\infty^2}$$

Where S_t^2 is the partial sum of the residuals, $\hat{\sigma}_\infty^2$ is the long run variance. Accept H_0 when the t-statistics is less than the critical value, that is, X_t is stationary. Reject H_0 for large values of t_k (i.e. $t_k > c_r$), X_t has a unit root.

? **Seasonal Unit root (Hegy test)**

This test has been proposed by hylleberg *et al* (1990) to check for seasonal unit root. For monthly time series, Frances (1990) discussed the test for seasonal unit root based on the model

$$\begin{aligned}\Delta_{12} X_t = {} & \pi_1 z_{1,t-1} + \pi_2 z_{2,t-1} + \pi_3 z_{3,t-1} + \pi_4 z_{3,t-2} + \pi_5 z_{4,t-1} \\ & + \pi_6 z_{4,t-2} + \pi_7 z_{5,t-1} + \pi_8 z_{5,t-2} + \pi_9 z_{6,t-1} \\ & + \pi_6 z_{4,t-2} + \pi_7 z_{5,t-1} + \pi_8 z_{5,t-2} + \pi_9 z_{6,t-1} \\ & + \pi_{10} z_{6,t-2} + \pi_{11} z_{7,t-1} + \pi_{12} z_{7,t-2} + \sum_{j=1}^{p} \alpha_j^* \Delta_{12} X_{t-j} + u_t\end{aligned}$$

The number of lagged seasonal differences $\Delta_4 X_{t-j}$ has to be chosen before the HEGY test can be performed. The process X_t has a regular (zero frequency) unit root if $\pi_1 = 0$ and it has seasonal unit root if any one of the other π_i (i = 2,3,…,12) is zero. If all the π_i (i =1,…,12) are zero, then a stationary model for the monthly seasonal differences of the series is suitable .

Model Estimation

The parameters of the selected $SARIMA(p,d,q)(P,D,Q)s$ model can be estimated consistently by least-squares or by maximum likelihood estimation methods. Both estimation procedures are based on the computation of the innovations e_t from the values of the stationary variable.

Model Diagnostic test

Once we have identified and estimated the SARIMA models, we assess the adequacy of the selected models to the data. This model diagnostic checking step involves both parameter and residual analysis by the use of ACF and PACF residuals plot, Ljung-Box Statistics and Normality test.

If the univeriate modeling procedure is utilized for forecasting purposes then this step can also form an important part of the diagnostic checking. This involves short forecast, middle forecast and long forecast statistics of the fitted models.

3. Modeling and Forecasting Evaluations

The focus is to use the seasonal autoregressive integrated moving average (SARIMA) techniques based on Box and Jenkins (1994) methodology to build models (Modelling) for the monthly average temperature of Sokoto city using data set for the period January 1995 to December 2003. The SARIMA model is then used to perform an out of sample forecast for January 2004 to December 2004. The data sets were obtained from the Metrological department, Sokoto State International Air port.

3.1 Identification of the seasonal models.

Time plot Fig 4.01 displays the time plot of the monthly average temperature series. A noticeable feature is the persistent recurrence of the pattern variability in all the periods, suggesting that the series has a pronounced seasonal pattern and hence is not stationary. In this case a formal test has to be carried out to test the presence or absence of seasonal unit root.

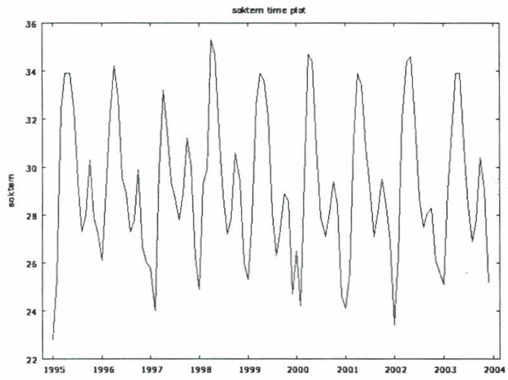

Fig 4.01: Time plot of monthly average temperature(X_t) for Sokoto from 1995:1 to 2003:12.

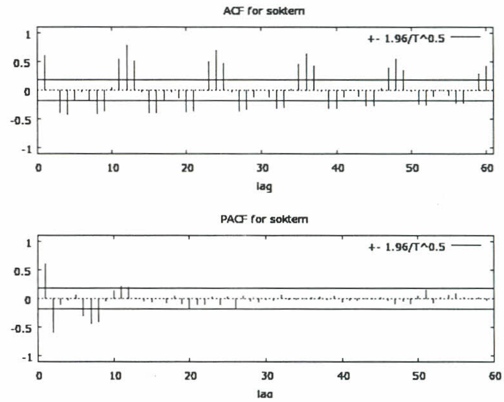

Figure 4.02: ACF and PACF for the time series, 1995:1 to 2003:12.

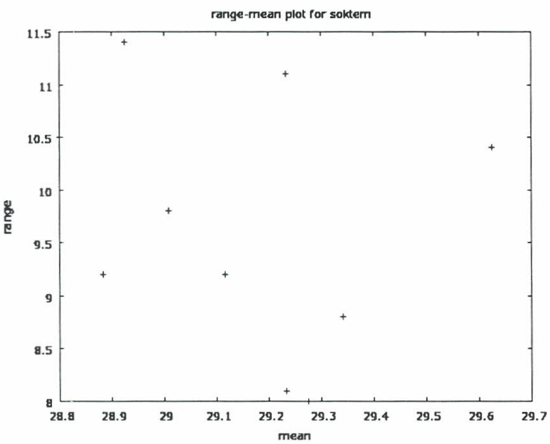

Fig 4.03: Range-Mean plot for Sokoto monthly average temperature, 1995:1 to 2003:12

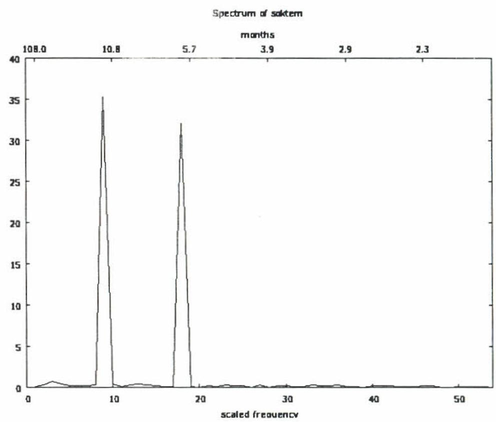

Fig 4.04: Sample periodogram of Sokoto temperature, 1995:1 to 2003:12

ACF and PACF

Consider the ACF plot of Fig 4.02 in which the highest spikes always occur at lags 12, 24, 36, etc., this indicates that the series is seasonal with period 12. Also the series is highly autocorrelated and the correlation is very persistent. Since the autocorrelation at seasonal periods are positive we expected that the fitted model should have seasonal autoregressive (SAR) component. On the other hand the PACF shows that the model is a mixed model with both AR and MA components.

Range-mean plot

We observe in the Fig 4.03 that the ranges are not increasing or do not tend to increase with the means. This means that there is no strong positive relationship between the sample mean and the sample variances for each period in the data. Finally, this indicates that there is no need for a log transformation.

Spectral analysis:

Spectral analysis is a useful frequency domain tool for detecting the existence of periodicity in a time series (Hamilton, 1994). This can be achieved by plotting the periodogram or spectral density of the series against either period or the frequency.

It can be seen in Fig 4.04 that there is a large-scale component at a frequency of nine cycles, precisely. In this case, there were 108 samples (9 years of data). Therefore, a frequency of nine is nine cycles every 108 months, or one cycles every 12 months (108/9). There is also another spike at a frequency of 18, which corresponds to a period of 6 months (108/18). The frequency spectrum clearly shows that there are both seasonal (12 month) and monthly (6 months) cycles in the sokoto temperature data .the height of the spikes tell you how much each spectral component contribute to the original data.

Unit root test

We use two methods to determine the order of non-seasonal integration of the series: ADF (Augmented Dickey-Fuller) and KPSS tests. The ADF test checks the null hypothesis of unit root against the alternative of stationarity for the data generating process. The KPSS test checks the null hypothesis of stationarity against the alternative of a unit root for the data generating process. The results for the ADF and KPSS tests are in Table 4.01. At the 5% significant level, the ADF test rejects the null hypothesis of unit root and KPSS test does not rejects the null hypothesis of stationarity. Therefore conclusively the time series does not required non-seasonal differencing.

HYGY Test

The HYGY statistic tests the null hypothesis there is no seasonal unit root against the alternative seasonal unit root. The p-value in table 4.01 is 0.006. Hence the null hypothesis of no seasonal is rejected at 5% significance level, confirming our expectation that the time series is seasonally integrated.

Table 3.01: Summary results for different tests

TEST	*t-statistics*	*p-values*
ADF	-4.6388	0.00011
KPSS	0.0189	0.146
HYGY for levels	0.5567	0.006
HYGY test for seasonal difference	0.0756	0.405

Penalty function criteria

To specified the range of values of the SARIMA parameters . The values of three of the parameters are known now: =12, = 0, and = 1; we have shown that the order of nonseasonal integration is zero; the order of seasonal integration is 1 and the periods of seasonality 12.

For the parameter space = 0,1,2,...,5; = 0,1,2,...,4; = 0,1,2,...,6; =0,1,2,...,4, the most parsimonious models given by the two information criteria AIC and BIC using ASTSA are:

1. SARIMA (1, 0, 0)(0,1,1) 12
2. SARIMA (2, 0, 1)(2,1,0) 12
3. SARIMA (2, 0, 2)(3,1,2) 12
4. SARIMA (3, 0, 1)(4,1,0) 12
5. SARIMA (4, 0, 2)(5,1,1) 12

An extension of the search to any wilder parameter space produced the same results. This confirms the optimality of the five models above.

Estimation of Models

The parameter estimation results show that all the models parameters are significant by using their standard error with their P values. The Table below represents the estimates:

Diagnostic checking

We test whether or not the residuals are generated by a white noise process by using (i) the ACF and PACF plots using the Ljung- Box test to check whether or not the residuals are uncorrelated, (ii) normal probability plots and the Anderson-Darling test to test the normality of the residuals.

Table 4.03 shows the results for Ljung-Box test, .The tests reveals that only the residuals for the models SARIMA(2,0,1)(2,1,0)12 and SARIMA (2,0,2)(3,1,2)12 are not uncorrelated, using the 5% significance level; these two cases are identified by the symbol *.

Table 4.02: Models Estimations

Model	*Predictor*	*Coefficient*	*Std. Error*	*T- ratio*	*P- value*
SARIMA	AR(1)	0.24	0.096	2.53	0.13
(1,0,0)(0,1,1)	SMA(1)	0.78	0.073	10.72	0.000
Model	***Predictor***	***Coefficient***	***Std. Error***	***T- ratio***	***P- value***
SARIMA	AR(1)	0.95	0.0017	546.9	0.000
(4,0,2)(5,1,1)	AR(2)	-0.38	0.0015	-248.6	0.000
	AR(3)	0.34	0.0017	198.3	0.000
	AR(4)	-0.22	0.0015	-148.6	0.000
	MA(1)	-0.36	0.0017	-208.8	0.000
	MA(2)	1.33	0.0017	774.3	0.000
	SAR(1)	-0.06	0.0017	-38.1	0.000
	SAR(2)	-0.52	0.0015	-350.7	0.000
	SAR(3)	-0.52	0.0016	-324.3	0.000
	SAR(4)	-0.28	0.0017	-162.9	0.000
	SAR(5)	-0.17	0.0017	-98.3	0.000
	SMA(1)	0.96	0.0017	551.9	0.000
Model	***Predictor***	***Coefficient***	***Std. Error***	***T- ratio***	***P- value***
SARIMA	AR(1)	0.03	0.0068	3.9	0.000
(3,0,1)(4,1,0)	AR(2)	-0.28	0.0097	-29.3	0.000
	AR(3)	0.04	0.0093	4.3	0.000
	MA(1)	-1.18	0.0097	-122.5	0.000
	SAR(1)	-0.38	0.0082	-46.7	0.000
	SAR(2)	-0.65	0.0089	-73.0	0.000
	SAR(3)	-0.61	0.0095	-64.7	0.000
	SAR(4)	-0.31	0.0095	-32.3	0.000

TABLE 4.03: Ljung-Box statistics (to test the residual autocorrelation as a set rather than individuals

Seasonal ARIMA Models	Lung-Box test	$\chi_{(1-\alpha\ ;k-p)}$	P value
SARIMA (1,0,0)(0,1,1)	$Q_{LB}(30)=22.4286$	41.3	0.0404
SARIMA(2,0,0)(2,1,0)	$Q_{LB}(30)=27.3162$	18.9	0.8406*
SARIMA(2,0,2)(3,1,2)	$Q_{LB}(15)=14.6718$	12.6	0.0498*
SARIMA(3,0,1)(4,1,0)	$Q_{LB}(15)=9.4458$	14.1	0.0424
SARIMA(4,0,2)(5,1,1)	$Q_{LB}(15)=11.7841$	12.6	0.006

The results for the normality test are in Table 4.04.The residuals of the entire five models pass the normality test.

Table 4.04: Results for the normal probability plot and the Anderson-Darling test

Models	*Normal plot value*	*P-Value (base on Anderson-Darling)*
(1,0,0)(0,1,1)	0.98800	0.057
(2,0,0)(2,1,0)	0.99171	0.196
(2,0,2)(3,1,2)	0.99602	0.865
(3,0,1)(4,1,0)	0.98416	0.344
(4,0,2)(5,1,1)	0.99170	0.870

On the basis of the results of the diagnostic checking the following three models were selected:

1. SARIMA (1,0,0)(0,1,1)12

1. SARIMA (3,0,1)(4,1,0)12
2. SARIMA (4,0,2)(5,1,1)12

Forecast evaluation

The Table below represents the lower, middle and long forecast statistics for the sokoto temperature series.

Table 4.05: Short, Middle and Long Forecast statistics for the three fitted SARIMA models

	Short forecast statistics			Middle forecast statistics			Long forecast statistics		
SARIMA Models	Mean Error	Mean Abs. Error	Mean Sq. Error	Mean Error	Mean Abs. Error	Mean Sq. Error	Mean Error	Mean Abs. Error	Mean Sq. Error
SARIMA (1,0,0)(0,1,1)	0.175	0.925	1.253	-0.75	0.9	1.125	0.95	0.95	1.01
SARIMA (4,0,2)(5,1,1)	-0.2	0.9	1.03	-0.3	0.45	0.505	0.225	0.625	0.478
SARIMA (3,0,1)(4,1,0)	-0.2	0.85	0.885	-0.35	0.75	0.813	0.825	0.825	0.986

The short forecast statistics (from 2004:01 to 2004:04) in table 4.05 show that the mean errors of the models are significantly smaller than their mean absolute errors. This implies that the forecast are neither systematically over forecasting nor under forecasting the temperature. SARIMA (3,0,1)(4,1,0)12 model is optimal forecast, since it has the lower mean square error of short forecast statistics. Also the Middle forecast statistics (from 2004:05 to 2004:08) explained that the mean errors of the models are significantly smaller than their mean absolute errors. This implies that the forecast are neither systematically over forecasting nor under forecasting the temperature. The middle forecast of SARIMA (4,0,2)(5,1,1)12 model is significantly more appropriate to forecast temperature, since it provides optimal forecast statistics by having minimum mean error and mean absolute error. Similarly for the Long forecast statistics (from 2004:09 to 2004:12) the models forecast are neither systematically over forecasting nor under forecasting the temperature, since the mean error of the models are significantly lower than their mean absolute error and mean square error. SARIMA (4, 0, 2)(5, 1,1)12 model has optimal long forecast statistics, since it mean error , mean absolute error and mean square error are less than that of the rest of the models.

CONCLUSION

This paper has considered the seasonal autoregressive integrated moving average (SARIMA) modeling and forecasting of sokoto monthly average temperature. Five seasonal models were chosen, by using model selection criteria. Only three models have passed the diagnostic test while the rest failed one or more of the tests.

From the out-of-sample forecast statistics analysis, we conclude that the short forecast statistics (2004:1 to 2004:4) of the three fitted models have shown that SARIMA (3,0,1)(4,1,0) 12 model minimizes the mean square error statistics of the forecasted models. This model was then found to be adequate and significantly better than the rest of the fitted models, and is adequate for a good description and forecasting of temperature pattern of the city. The middle forecast statistics (2004:5 to 2004:8) and long forecast statistics (2004:9 to 2004:12) have both shown that, SARIMA (4,0,2)(5,1,1)12 model has optimal forecast to the temperature by minimizing the mean squared errors of forecasts than the rest of the models.

Therefore, conclusively the best seasonal model among the forecasted models that is adequate to describe and forecast the seasonal dynamics for Sokoto city temperature is

SARIMA (3,0,1)(4,1,0) 12 for the first quarter and first month in the second quarter of year and SARIMA (4,0,2)(5,1,1) 12 for the rest of the year.

REFERENCES

Box, G.E. and Jenkins, (1976). '*Time series Analysis*' Forecasting and Controls, San Francisco, USA.

Box, Jenkins and Reinsel. (1994) ' *Time series analysis*' ARIMA Models, San Francisco, USA

Dickey, D.A and Fuller, W.A (1979) 'Estimators for autoregressive time series with a unit root', *Journal of the American Statistical Association* 74: **427-431**

Frances, P.H (1990) 'Testing for seasonal unit roots in monthly Data', *Econometric Institute Report* 9032A, Erasmus University Rotterdam.

Hamilton. J., (1994). '*Time Series Analysis*', Princeton University press, New Jersey, USA

Hylleberg, S. and Granger (1990) 'Seasonal Integration and co integration', *Journal of Econometrics* 44: **215-238**.

Kwiatkowski, D., Phillips, P.C.B., Schmidt, P. and Shin (1992) 'Testing the null of stationarity against the alternative of a unit root: How sure are we that the economic time series have a unit root? *Journal of Econometrics* 54: **159-178**

Meyler, A., Kenny, G. and Quinn, T. (1998) 'Forecasting Irish Inflation using ARIMA Models' Research paper Economic Analysis, Research and Publications Department, Central Bank of Ireland, Dublin.