WORKING PAPER NO: 540

Two-Stage Models for Forecasting Time Series with Multiple Seasonality

Anupama Lakshmanan
Doctoral Student
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 5600 76
anupama.lakshmanan14@iimb.ernet.in

Shubhabrata Das
Professor
Decision Sciences & Information Systems
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 5600 76
Ph: 080-26993150
shubho@iimb.ernet.in

Year of Publication – February 2017

Two-Stage Models for Forecasting Time Series with Multiple Seasonality

Anupama Lakshmanan
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore 560076, India
email: anupama.lakshmanan14@iimb.ernet.in

Shubhabrata Das
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore 560076, India
email: shubho@iimb.ernet.in

February 22, 2017

Abstract: Complex multiple seasonality is an important emerging challenge in time series forecasting. In this paper, we propose models under a framework to forecast such time series. The framework segregates the task into two stages. In the first stage, the time series is aggregated and existing time series models such as regression, Box-Jenkins or TBATS, are used to fit this lower frequency data. In the second stage, additive or multiplicative seasonality at the higher frequency levels may be estimated using classical, or function-based methods. Finally, the estimates from the two stages are combined. Detailed illustration is provided via energy load data in New York, collected at five-minute intervals. The results are encouraging in terms of computational speed and forecast accuracy as compared to available alternatives.
Keywords: ARIMA, Energy Load, Polynomial, Regression, TBATS, Trigonometric
AMS Classification: 37M10, 62M10

1 Introduction

In many contexts, e.g. energy load or demand, mobile network usage, data is recorded very frequently – at times, after every minute or every five minutes. Such data may exhibit not only annual and weekly seasonality but also within day and within hour seasonality that are often entangled with each other. Forecasting problems involving such high frequency time series data with complex multiple levels of seasonality are increasingly drawing the attention of both the academicians and practitioners (Hyndman and Fan [2015], Au et al. [2011], Livera et al. [2011], Gould et al. [2008], Soares and Souza [2006], Taylor [2003], among others). There is huge commercial interest due to the ability to capture data at rapid succession and the potential benefit that can be derived from successful forecast which leads to judicious resource planning.

Most traditional methods, e.g. Auto Regressive Integrated Moving Average (ARIMA), are not well-suited for dealing with long seasonal patterns or seasonality at multiple levels (such as within hour, within day, within week and within year). A relatively new methodology called BATS, and an extension called TBATS [Livera et al. [2011]], were developed primarily to address such problems. However, in the current implementation, TBATS is not designed to accommodate covariates. In addition, TBATS has high computational requirements and high instability in forecast (as seen subsequently in Tables 4 and 10). In this study, we propose a two-stage framework which can incorporate covariate information, is computationally much faster and yields better forecast accuracy than TBATS in the examples considered.

The paper is organized as follows. In Section 2, we describe the two-stage framework. Next, in Section 3, we demonstrate the implementation of the two-stage framework on New York energy load data, the primary case study that motivated this research. Further, this section presents the performance of the two-stage framework and compares it with TBATS. In Section 4, we select a couple of variations of the two-stage framework and compare their forecast accuracies with that of TBATS recursively with increase in the model period. In Section 5, we discuss various implementation aspects of the proposed two-stage methods. Specifically, in Section 5.1, we comment on model selection criteria that combines lack-of-fit with model complexity. In Section 5.2, we exhibit different choices in the level of separation between the two stages. In Section 5.3, we show results of forecast combinations on the New York energy load data. In Section 5.4, we explore the performance of the two-stage framework on another dataset. In Section 5.5, we narrate a few possible variations of the methods used in the two-stage framework. Finally, we wrap up with a summary in Section 6 and briefly describe future research.

2 The Two-Stage Framework

In the classical decomposition model, the time series Y to be modeled or forecast, is viewed as:

$$Y = T + S + \epsilon \quad \text{or} \quad Y = T \cdot S \cdot \epsilon, \tag{1}$$

depending on additive or multiplicative model, where T represents the trend or the long-term direction, S represents the seasonality or a pattern which repeats periodically, and ϵ represents the error which is random.

In this article, we study time series problems where the seasonality is not of a simple single periodicity, but is complex having multiple levels. The seasonality may be represented through $S_k, k = 1, ..., m$, if there are m-levels of seasonality in the time series. For example, when $m = 4$, S_1 could be the annual seasonality (pattern repeating every year), S_2 could be the weekly seasonality (pattern repeating every seven days), S_3 could be the daily seasonality (period equal to 24 hours), and finally S_4 could be the within-hour seasonality (period equal to one hour). Electricity consumption, gasoline consumption, and mobile usage are a few common examples of time series data that may similarly exhibit strong seasonality at multiple levels.

The two-stage framework proposed in this paper is inspired by (1), and separates the modeling of the time series into two stages – *low* and *high* frequency time series components. The final model in the two-stage additive framework looks like:

$$Y = Y_{\text{low}} + S_{\text{high}} + \epsilon, \tag{2}$$

while the same under a multiplicative framework is:

$$Y = Y_{\text{low}} \times S_{\text{high}} \times \epsilon. \tag{3}$$

Here Y_{low} represents the time series that incorporates seasonality at the low frequencies. The observed values of Y_{low} may be obtained by aggregating the original time series at a suitable level to be determined by the analyst. All possible time series models should be tried in the first stage in order to fit Y_{low}. S_{high} represents the seasonality at higher frequencies, possibly at multiple levels, but all beyond the determined threshold level of separation. Seasonality at high frequencies, S_{high}, is then estimated in the second stage. Finally, the two stages are combined to yield an eventual fit or forecast, as explained in Figure 1.

Figure 1: Proposed Two-stage Framework

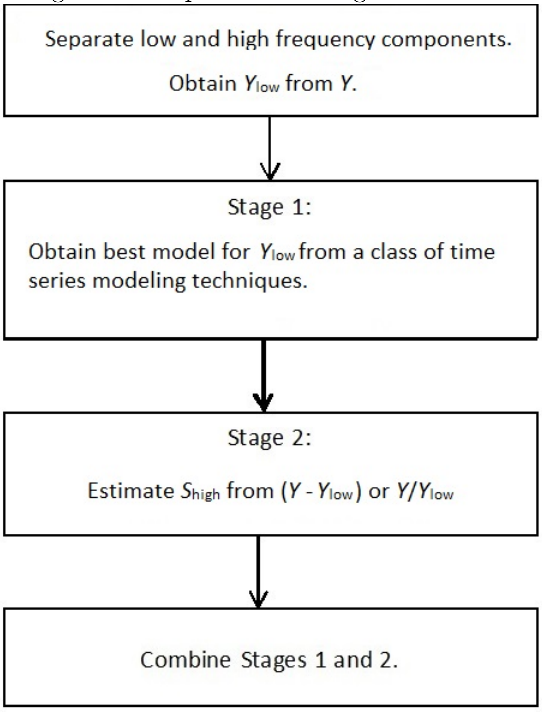

As an illustration, let us consider a time series data, recorded every minute, and having four-level seasonality. Let Y_{ijkl} denote the data corresponding to i-th week, j-th day of the week, k-th hour of the day, and l-th minute of the day, with $i = 1, \ldots, I$; $j = 1, \ldots, J = 7$; $k = 1, \ldots, K = 24$; and $l = 1, \ldots, L = 60$. Let us, as an illustration of the two-stage framework, group the daily and hourly seasonality as high-frequency and group the weekly and annual seasonality as low-frequency. (Note that a modeler may alternatively consider separation of high vis-à-vis low frequency at the hourly level or at the weekly level. In the former case, only the within-hour seasonality is treated as high-frequency, while in the latter only the annual seasonality is treated as low-frequency. This is further discussed in Section 5.2.)

Consequently, in the first stage, we consider various models for

$$Y_{\text{low}} = \bar{Y}_{ij\cdot\cdot} = \frac{1}{24 \times 60} \sum_{k=1}^{24} \sum_{l=1}^{60} Y_{ijkl}$$

that takes into account, among other factors, also weekly and annual seasonality. For example, in the regression framework, the weekly seasonality may be accounted by considering the dummy variables corresponding to (six) different days of the week (other than, say, Sunday). The various models that can and should be tried are discussed with specific illustrations in Section 3.

In the second stage, we estimate seasonality from the high frequency time series. Seasonality can be additive or multiplicative. Further, the additive or multiplicative seasonality is estimated using either the classical decomposition type method or function-based methods. A general schematic diagram of the second stage is provided in Figure 2.

Figure 2: Seasonality Estimation in Stage 2: Two-Stage Framework

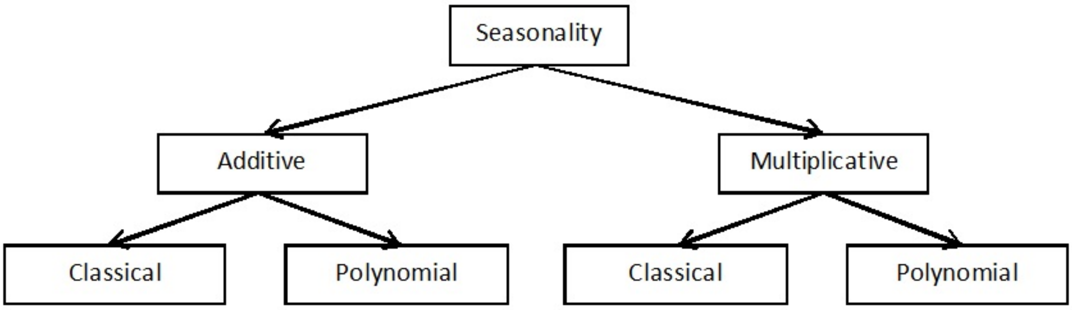

Finally, we combine Stage 1 and Stage 2 to arrive at the model.

We next describe the estimation of seasonality at higher frequency (Stage 2) using two alternative approaches.

Classical estimation of high-frequency seasonality

Going back to the example where data (Y_{ijkl}) is recorded every minute and has a four-level seasonality, in the second stage of accounting for high-frequency seasonality, we need to estimate daily and within-hour seasonality. In order to estimate the within-hour additive seasonality, $\hat{S}^{A,\text{hour}}$, we first obtain the deviation of per-minute data Y_{ijkl} from the hourly average $\bar{Y}_{ijk\cdot} = \frac{1}{60}\sum_{l=1}^{60} Y_{ijkl}$. We then average across weeks and days to get the within-hour seasonality; e.g. the additive seasonality for the l-th minute of the k-th hour of a day is computed as:

$$\hat{S}_{kl}^{A,\text{hour}} = \frac{1}{7I}\sum_{i=1}^{I}\sum_{j=1}^{7}\left(Y_{ijkl} - \bar{Y}_{ijk\cdot}\right). \tag{4}$$

Similarly, to obtain the effect of the different hours of the day, we use the deviation of average hourly data from that of the daily average:

$$\hat{S}_{k}^{A,\text{day}} = \frac{1}{7I}\sum_{i=1}^{I}\sum_{j=1}^{7}\left(\bar{Y}_{ijk\cdot} - \bar{Y}_{ij\cdot\cdot}\right). \tag{5}$$

For the additive model (2), S_{high} is estimated by the classical method through $\hat{S}_{\text{high}}$ which is the sum of (4) and (5), over suitable indices.

To find the multiplicative hourly and daily seasonality, instead of the deviations in equations (4) and (5), we consider the ratios:

$$\hat{S}_{kl}^{M,\text{hour}} = \frac{1}{7I}\sum_{i=1}^{I}\sum_{j=1}^{7}\left(\frac{Y_{ijkl}}{\bar{Y}_{ijk\cdot}}\right). \tag{6}$$

$$\hat{S}_{k}^{M,\text{day}} = \frac{1}{7I}\sum_{i=1}^{I}\sum_{j=1}^{7}\left(\frac{\bar{Y}_{ijk\cdot}}{\bar{Y}_{ij\cdot\cdot}}\right). \tag{7}$$

For the multiplicative model (2), S_{high} is estimated by the classical method through $\hat{S}_{\text{high}}$ which is the product of (6) and (7), over suitable indices.

The summations in (4) – (7) can be modified by carrying them out separately for different days of the week, or separately for weekdays and weekends, or for a multitude of other combinations. This may be dictated in part by the volume of data available, and in part by the assumption of the modeler regarding presence or absence of interaction between low-frequency and high-frequency seasonality. Objective model selection criteria may also be used in arriving at a choice. In our case study with New York energy load data, which is a very large data set, we find it beneficial to treat each weekday as different, even though there is a reasonably strong case for clubbing indices for weekdays together.

Function-based estimation of high-frequency seasonality

Note that, even if there are multiple levels of high frequency, these would typically be multiples of each other and hence in this approach, all the higher level frequencies may be integrated into a single high frequency. For example, with one minute frequency data separated at the daily level, both the within-hour and within-day seasonalities may be combined into a single seasonality of frequency equal to 24×60 entries.

In the aforementioned classical method of estimating additive/multiplicative seasonality, the number of parameters to be estimated is large. In the context of data collected at one-minute frequency data, we need 24 hourly values and 24×60 minute-level within-hour seasonality values to be estimated. Model prudence is desirable and simpler models are rewarded by model selection criteria such as AIC or BIC.

As an alternative, we propose to estimate the seasonality using special functions such as polynomial or trigonometric functions. In the following, we describe a polynomial-based estimation of high-frequency seasonality. Estimation based on trigonometric function is outlined in Section 5.5.2.

We start with the adjusted or residual data from stage 1, i.e. $\tilde{Y} = Y - Y_{\text{low}}$ or $\tilde{Y} = \frac{Y}{Y_{\text{low}}}$, depending on additive or multiplicative model (2) or (3). As per our model, $\tilde{Y}$ consists essentially of high-frequency seasonality of period upto the threshold separating the high and low frequencies; this is one day in our running example. Let us consider a re-parametrization, if necessary, of the time horizon so that $t = 1$ unit refers to the time period separating high and low frequencies; this would be one day in our illustration. Let $\tilde{t}$ be equal to $t - [t]$ (i.e. $\tilde{t} = t$ mod 1), and hence polynomial seasonality of degree k, refers to the model:

$$\tilde{Y}_t = \alpha_0 + \alpha_1 \tilde{t} + \ldots + \alpha_k \tilde{t}^k + \epsilon_t; \tag{8}$$

where ϵ_t is the usual regression residual and α's are the regression coefficients. However, periodicity due to (continuous form of functional) seasonality requires

$$E(\tilde{Y}_0) = \lim_{t \to 1} E(\tilde{Y}_t) \Rightarrow \alpha_0 = \sum_{i=0}^{k} \alpha_i \Leftrightarrow \sum_{i=1}^{k} \alpha_i = 0 \Leftrightarrow \alpha_k = -\sum_{i=1}^{k-1} \alpha_i.$$

Hence, it would be equivalent, and more convenient to run an unconstrained multiple regression model:

$$\tilde{Y}_t = \alpha_0 + \alpha_1 U_1 + \ldots + \alpha_{k-1} U_{k-1} + \epsilon_t, \tag{9}$$

with

$$U_i = \tilde{t}^i - \tilde{t}^k, \quad i = 1, \ldots, k-1. \tag{10}$$

The choice of k can be determined by model selection criteria.

In the next section, we demonstrate specific implementations of the proposed two-stage method in a case study on the New York energy load data set. In particular, we also examine how this framework can be leveraged to obtain better forecasts than comparable methods.

3 Illustration of the Two-Stage Framework: Primary Case Study

We demonstrate the new two-stage framework with a case study on energy load data from New York that is downloaded from [NYISO, 2014].

3.1 New York Energy Load Data

The data reflects the (spot) energy load recorded at five-minute (frequency) intervals. In addition to the electricity load, four weather-related variables (maximum and minimum dry and wet bulb temperature) are available, which are recorded daily. For the study, we consider September 5, 2008 – December 31, 2013 as the model period and January 1, 2014 – December 31, 2014 as the hold-out or validation period. The starting date for the model period is dictated by the availability of the weather variables.

A visual analysis of the data establishes strong complex multiple levels of seasonality. We discern four levels of seasonality, namely, annual, weekly, daily, and hourly. A small but illustrative set of diagrams is provided in Figures 3 – 6 to demonstrate this. Figure 3 plots the energy load data for four years from 2009 to 2012 and it is evident that the summer months from June to September tend to see higher usage of electricity than the other months.

Figure 3: Sample New York Electricity Load Data Exhibiting Annual Seasonality

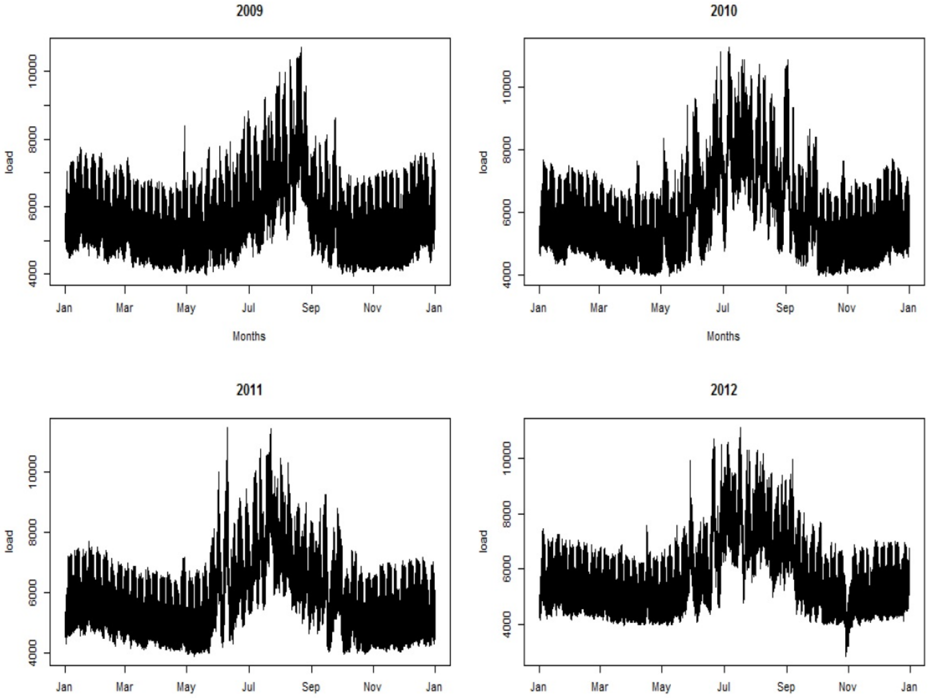

Next we zoom into weekly data. Picking three weeks in the year 2010, we again see strong seasonality, i.e. lower electricity usage during the weekend than on weekdays.

Figure 4: Sample New York Electricity Load Data Exhibiting Weekly Seasonality

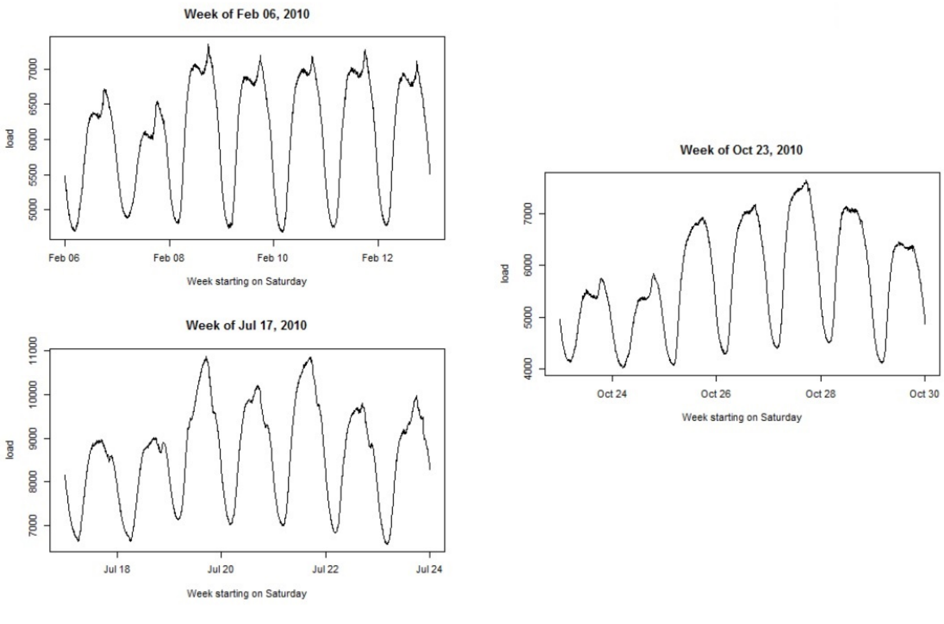

Further zooming into daily usage, we compare the load on a couple of Mondays and Tuesdays in June 2011 and 2013. It is not surprising that all four days show a similar usage pattern. Electricity usage starts climbing around 7:00 a.m., peaks near noon and begins to decline close to 5:00 p.m.

Figure 5: Sample New York Electricity Load Data Exhibiting Daily Seasonality

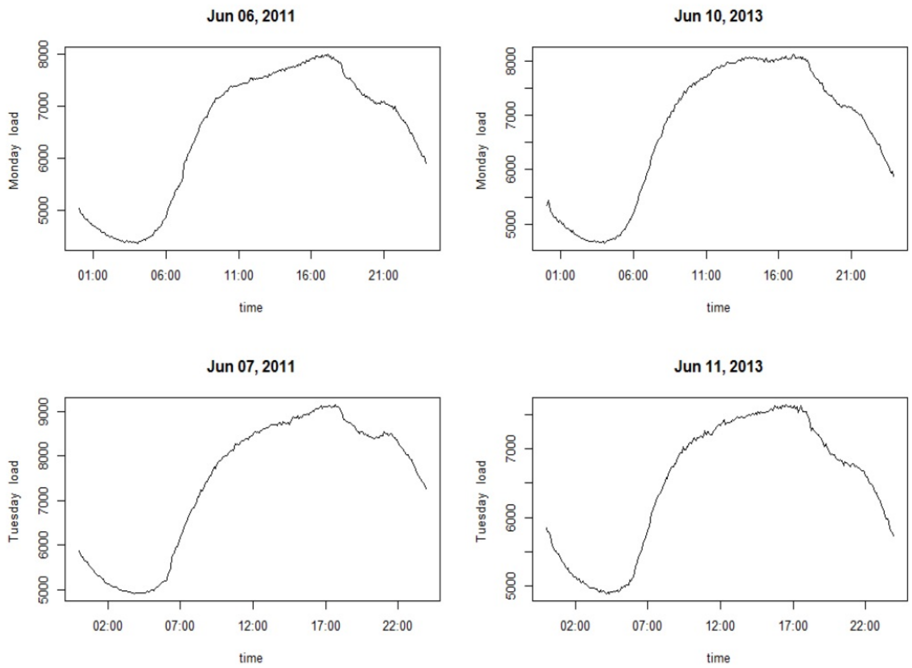

Finally, we examine seasonality within the hour. Looking at 8:00 - 9:00 a.m. usage on four consecutive days in June 2011, we see that all the five-minute blocks in that hour show rising energy usage.

Figure 6: Sample New York Electricity Load Data Exhibiting Within-hour Seasonality

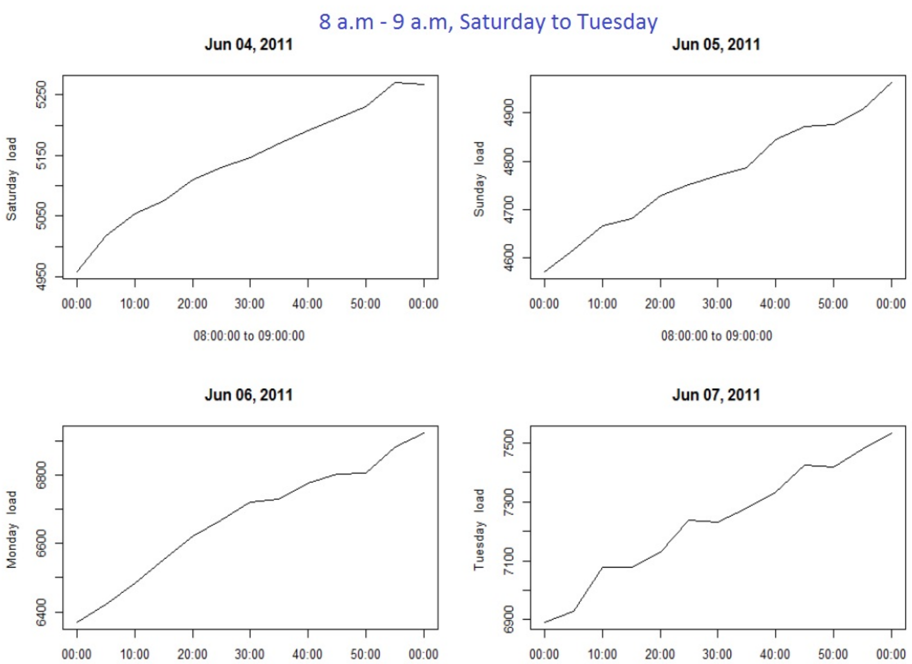

The above visual inspection confirms complex multiple levels of seasonality in this data set. In order to forecast such time series we need to perform a thorough and systematic study based on a model that accounts for such type of complex seasonality.

3.2 Implementation of Two-State Methods for New York Energy Load data

The two-stage framework is implemented in the software environment for statistical computing, R [CoreTeam [2014]]. For TBATS we use the `forecast` package in R [Hyndman [2016]].

3.2.1 Stage 1: Low Frequency Modeling

In the first stage, we consider suitable models for daily average load $\bar{L}_{day}$. The various classes of models considered towards this are depicted in Figure 7.

Figure 7: Options in Stage 1 : Two-stage Framework

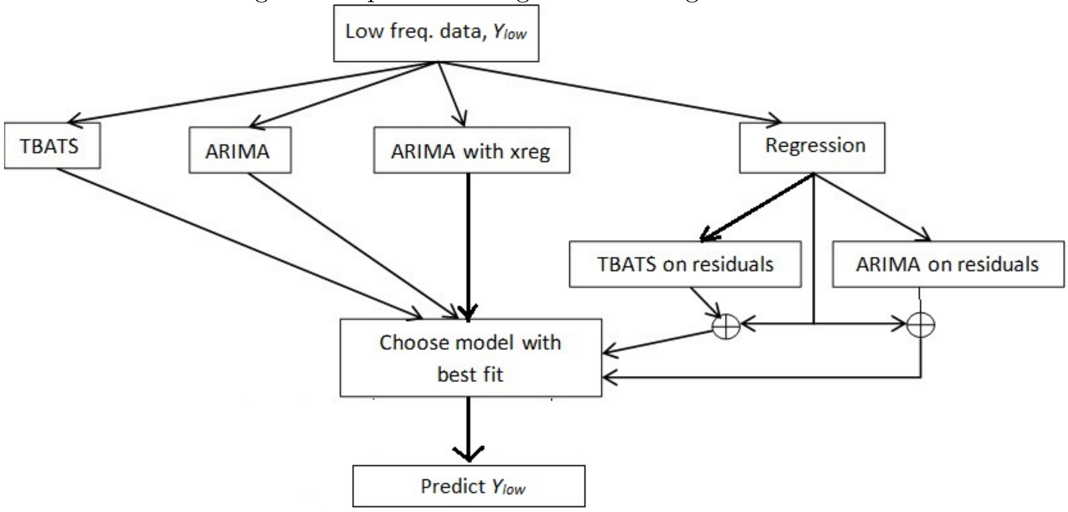

A natural approach is to use a regression framework, as the weekly seasonality can be captured through the choice of dummy variables corresponding to the days of the week, and we expect to capture a large part of the annual seasonality through weather-related explanatory variables. However, these explanatory variables may not be able to fully capture the weekly and annual seasonality. Hence, we also explore using TBATS or ARIMA on the residuals of the (best-fit) regression. One option towards this is to implement ARIMA with xreg in R. Typically this would be more efficient than running ARIMA on regression residuals.

Let us dwell on some additional details on the regression. The daily average load is regressed on all relevant independent variables that may include weather-related variables and dummy variables that account for different days of the week, or weekday, or other structural differences. Usual model selection procedures are adopted to arrive at an appropriate list of covariates which have significant impact. Certain non-linear transformation of variables are also considered; in particular, energy load is found to have quadratic relationship with temperature variables. Whenever the weather-related variables are available at higher frequency (e.g. hourly), the corresponding daily averages are considered. Model selection of independent variables is then carried out; using change point analysis we also determine whether the different days of the week have significant impact, or whether significant difference exists between weekdays and weekend, or between Saturday and Sunday. Subsequently, non-significant variables are dropped to reduce to the following regression model:

$$\bar{L}_{\text{day}} = \beta_0+\beta_1\text{DayNr}+\beta_2 T_{\max}+\beta_3 T_{\min}+\beta_4 b_{\max}+\beta_5 b_{\min}+\beta_6 T_{\max}^2+\beta_7 T_{\min}^2+\beta_9 b_{\min}^2+\beta_{10} D_{\text{wkday}}+\beta_{11} D_{\text{Sat}}, \quad (11)$$

where $T_{\max}$ and $T_{\min}$ denote maximum and minimum dry bulb temperature respectively, $b_{\max}$, $b_{\min}$ are the maximum and minimum wet bulb temperatures respectively, D_{wkday} is a dummy variable to denote weekday,

D_{Sat} is a dummy variable to denote Saturday and DayNr is the day number to account for linear trend. The parameter estimates are listed in the last row of Table 17 in the Appendix. A summary of the fit and prediction of the competing methods to forecast daily average load is provided in Table 1 by reporting mean square error (MSE), mean absolute deviation (MAD), and mean absolute percentage error (MAPE). The first row corresponds to the regression model in (11).

Note that $\bar{L}_{\text{day}}$ has two levels of seasonality – weekly and annual. ARIMA model, by itself, does not perform very well, as expected due to the long period of seasonality. ARIMA on the regression residual also did not improve the forecast and hence results from these variations have not been reported. The ARIMA with xreg option, which integrates the regression and ARIMA, performed best both in the model and in the forecast period and this corresponds to the last row of Table 1.

Next, we implement the TBATS model with three choices of seasonality – weekly, annual or both and the results are reported in rows 2 – 4. The errors from TBATS (with both weekly and annual seasonality) applied on the regression residuals are reported in the fifth row.

Table 1: Summary of Forecast Evaluation in Predicting Daily Average

Sl.		Model period 2008-13			Hold-out period 2014		
No.	Method	MSE	MAD	MAPE	MSE	MAD	MAPE
1	Regression as in (11)	101188.2	245.2	4.01%	341266.7	398.2	6.63%
2	TBATS with only weekly seasonality	134842.3	254.1	4.02%	806588	621.4	9.44%
3	TBATS with only annual seasonality	142833.7	279.2	4.44%	1537894	948.9	14.63%
4	TBATS with weekly & annual seasonality	99581.3	213.6	3.39%	667642.9	601.9	10.05%
5	Regression, then TBATS on regression residuals	55940.6	170.6	2.78%	343144.5	375.9	6.20%
6	Yearly seasonal ARIMA with xreg	45055.5	148.3	2.38%	284098	376.6	6.16%

3.2.2 Stage 2: Estimating Seasonality at Higher Frequency

In this step, we use high frequency data or the five-minute energy load data to estimate the multiple levels of seasonality. As described in Section 2, we broadly use two methods to estimate seasonality — the first is via a classical decomposition, while the second is by a function-based approach, more specifically in the form a polynomial function. In either scheme, the seasonality can be additive or multiplicative and consequently the estimation procedure is adjusted.

In order to estimate seasonality using the classical method, we use (4) - (7). The only difference is that in this data, the load is recorded every five minutes instead of every minute. In Figure 8, we plot the estimated daily multiplicative seasonality (corresponding to the different hours of the day) by the classical method. All the weekdays have almost identical pattern, while there is moderate difference between Sundays and Saturdays. However weekday patterns are quite different from the weekend.

Figure 8: Classical Estimate of Multiplicative Daily Seasonality

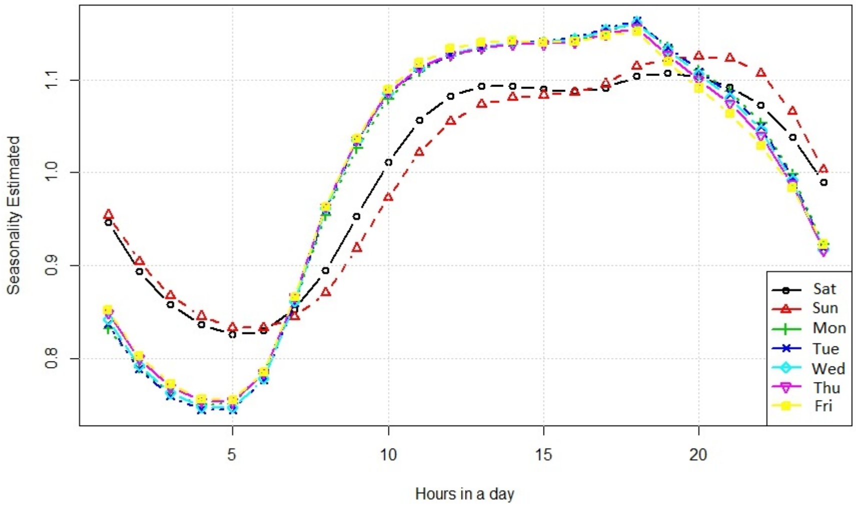

The estimated within-hour seasonality across days, as estimated by the classical method is shown in Figure 9. Again, the weekend shows a slightly different pattern compared to weekdays.

Figure 9: Classical Estimate of Multiplicative Within-hour Seasonality

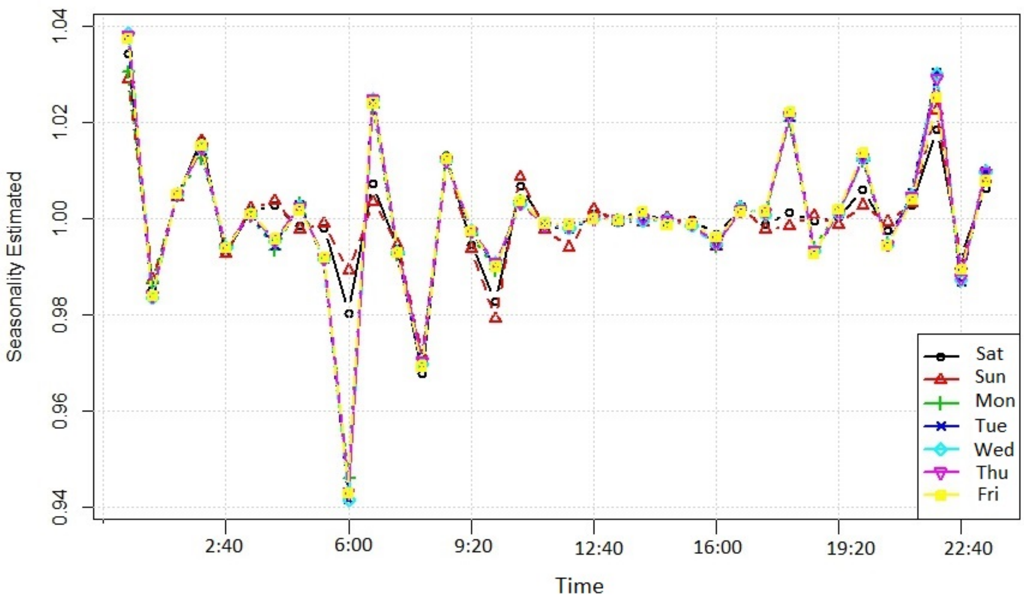

In Figure 10, we illustrate a fourth-degree polynomial fit that is used to estimate the within-day seasonality.

Figure 10: Polynomial Estimate of Multiplicative Daily Seasonality

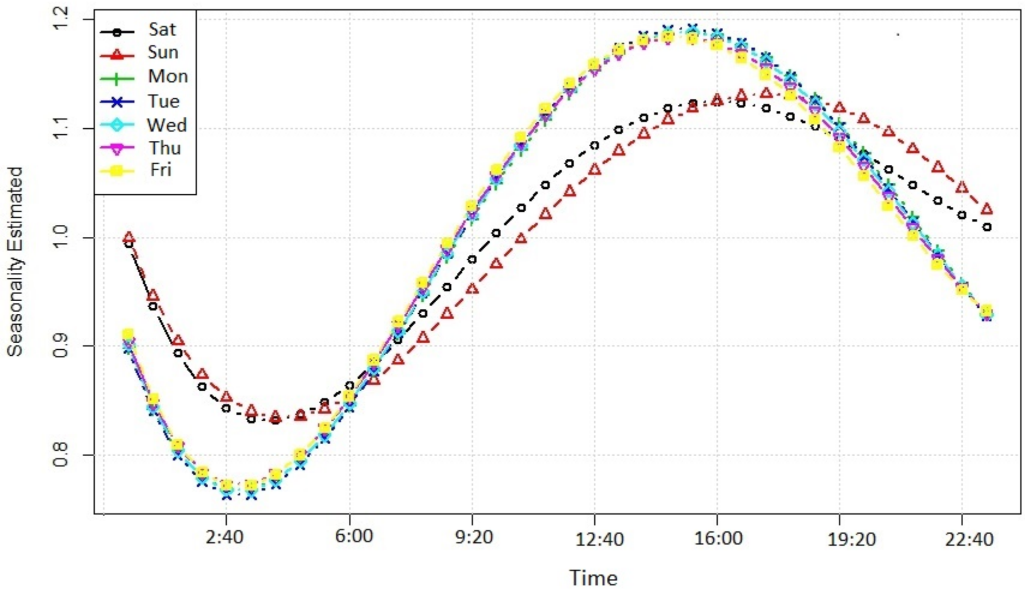

The multiplicative seasonality estimated by classical and polynomial methods, along with the adjusted data $\tilde{Y} = \frac{Y}{\hat{Y}_{low}}$, is shown in Figure 11. The plot shows the seasonality estimated for one week of data along with load data, adjusted by daily average.

Figure 11: Multiplicative Seasonality Estimated Using Classical and Polynomial Method

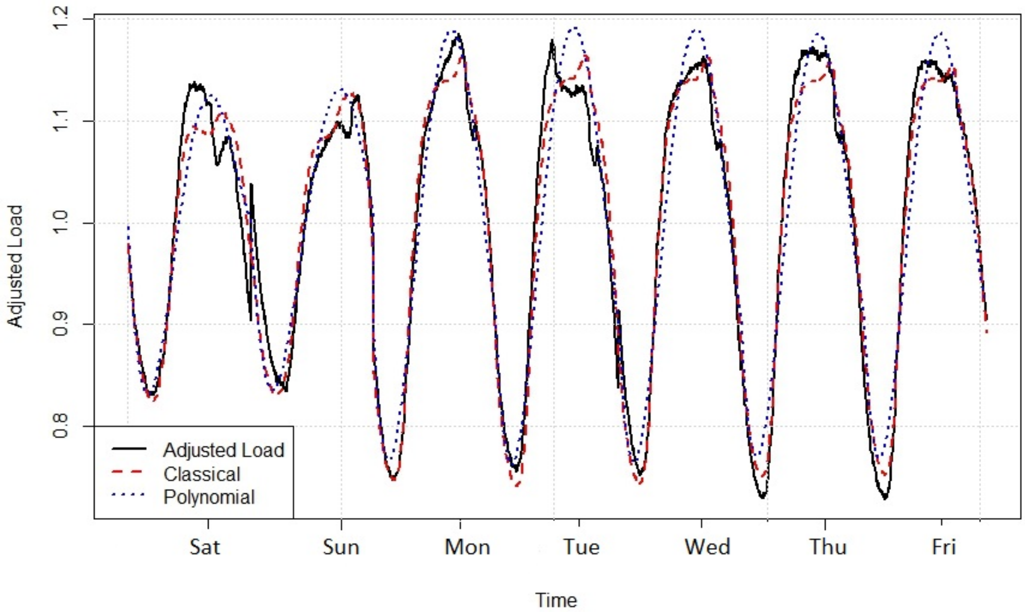

The plots corresponding to the additive models are largely similar to the multiplicative ones and hence omitted for brevity.

3.2.3 Combining Two Stages and Model/Forecast Comparison

In the first stage, we model the daily average load $Y_{\text{low}} = \bar{L}_{\text{day}}$ and come up with a suitable estimate $\hat{Y}_{low}$, based on model selected in this first stage (viz., best regression or ARIMA with xreg, etc). Daily seasonality is then estimated from the adjusted data $\tilde{Y}$ by classical or polynomial-based method and the two stages are combined to yield final fit or forecast. As an illustration, consider predicting load for March 3 (Tuesday), 2014 at 6:15 a.m. using regression (11) and classical additive seasonality. Based on the daily average load, the forecasted daily average load turns out to be 7155.87. Now from the adjusted data $\tilde{Y} = Y - Y_{\text{low}}$, the additive seasonality index for this block is found to be -1020.34, which is actually the sum of $S^{A,\text{day}}_{7,\text{Tue}} = -899.86$, the seventh hour effect for a Tuesday and $S^{A,\text{hour}}_{74} = -120.48$, the block effect for the fourth block of the seventh hour of a Tuesday. Thus, the eventual forecast for this period is $7155.87 - 1020.34 = 6135.53$ (actual observation is 6020.9). If on the other hand, the polynomial-based multiplicative seasonality estimate is used along with regression, the seasonal estimate of this block would be 0.8561, leading to a predicted value of $7155.87 \times 0.8561 = 6126.14$.

We now compare all the methods with specific focus on how the different models under two-stage framework fare as compared to the TBATS model with four levels of seasonality – hourly, daily, weekly and annual seasonality. Table 2 summarizes the performance of the key better-performing methods in fitting at five-minute intervals on the training/model period and in the hold-out period.

Table 2: Summary of Different Models in Predicting Load at Five-minute Intervals

			Model period			Hold-out period		
Method	Stage 1	Stage 2	MSE	MAD	MAPE	MSE	MAD	MAPE
M1	Regression	Classical additive	146252.1	288.9	4.80%	398837.8	438.9	7.40%
M2		Classical multiplicative	136952.4	277.3	4.58%	394285.7	425.2	7.09%
M3		Polynomial additive	176814.9	327.7	5.43%	423940.1	463.5	7.82%
M4		Polynomial multiplicative	165550.8	313.8	5.21%	417880.3	449.8	7.54%
M5	2 level TBATS	Classical additive	144644.7	260.1	4.12%	725218.1	627.7	10.60%
M6		Classical multiplicative	141849.8	254.0	3.99%	725928.9	621.1	10.40%
M7		Polynomial additive	175208.0	307.0	4.90%	750316.3	645.7	10.95%
M8		Polynomial multiplicative	170175.2	301.7	4.84%	748615.4	640.9	10.80%
M9	TBATS on regression residual	Classical additive	101004.4	231.2	3.80%	400715.6	422.2	7.05%
M10		Classical multiplicative	93030.6	219.1	3.56%	397679.4	406.5	6.70%
M11		Polynomial additive	131567.3	279.9	4.60%	425817.9	451.5	7.55%
M12		Polynomial multiplicative	121503.8	268.5	4.43%	421159.3	438.1	7.28%
M13	**ARIMA with xreg**	Classical additive	90119.2	211.9	3.41%	341668.0	409.9	6.79%
M14		**Classical multiplicative**	**82809.9**	**201.0**	**3.23%**	**332561.8**	**401.2**	**6.61%**
M15		Polynomial additive	120682.2	265.0	4.30%	366771.4	437.7	7.22%
M16		Polynomial multiplicative	111355.1	255.0	4.17%	357058.3	426.8	7.03%
M17	4 level TBATS		485.14	16.7	0.28%	707423.9	714.6	12.47%

Table 2 shows the utility of proposed seasonality estimation within a day (high frequency) and integrating it with daily average (low frequency) forecast based on TBATS, or regression, or ARIMA; this is the main idea behind the two-stage framework. Method M14 works the best. The TBATS model with all four levels of seasonality (M17) is not only computationally intensive, but it runs into the possible danger of over-fitting. Indeed, though it provides an superlative fit during the model period, its performance is worse (compared to several of the competing methods) during the hold-out period.

4 Recursive Forecast on New York Data

In Section 3, we observed that the two-stage framework performed favourably when compared with TBATS in terms of the forecast accuracy using data available for five years. It is fair to ask if this observation remains valid with change in model and/or forecast period. At any rate, a practitioner is often posed with the challenge of determining how much data is enough to place confidence in its forecast. Therefore, we consider *recursive* forecasts, where we first forecast using only (a little over) one year long data for modeling and then progressively increase the model period by three months. We also consider short, medium and long hold-out periods, by considering forecast for the next day, next week, 4-weeks (approximately a month), 13-weeks (approximately a quarter) and 52-weeks (approximately a year).

The error in the recursive fit for TBATS with four levels of seasonality, is provided in Table 3. The first observation in Table 3 is that the model period errors are startlingly small. A MAPE of 0.28% implies an almost perfect fit. However, the 1-day and 1-week forecast errors are occasionally large and fluctuate widely across the different model periods. Large variations in a single day forecast (or even 1-week) are perhaps to be expected, yet MAPE of 34% (for the 1-week forecast made on 01-Jul-10) is cause for concern.

Table 3: Recursive Forecast evaluation for TBATS: Model Period and Short Horizon Forecast

Forecast	Model period			1 day forecast			1 week forecast		
made on	MSE	MAD	MAPE	MSE	MAD	MAPE	MSE	MAD	MAPE
01-Oct-09	508.0	17.0	0.28%	411870.2	523.4	8.84%	776429.4	795.8	14.20%
01-Jan-10	516.4	17.2	0.29%	1363993.4	1042.3	19.25%	463520.8	545.6	9.44%
01-Apr-10	511.3	17.1	0.29%	1063118.6	998.5	18.26%	1324121.7	1070.1	19.49%
01-Jul-10	514.3	17.2	0.29%	835533.9	833.3	12.24%	11242330.2	2839.8	34.30%
01-Oct-10	495.9	16.9	0.28%	690406.0	653.6	9.97%	1680012.9	1235.9	22.64%
01-Jan-11	477.8	16.6	0.27%	418824.9	555.8	10.54%	191688.4	383.2	6.72%
01-Apr-11	484.6	16.7	0.28%	88895.1	264.1	4.32%	51724.2	188.5	3.43%
01-Jul-11	485.2	16.7	0.28%	51307.8	183.3	2.49%	400226.2	497.3	6.46%
01-Oct-11	498.8	16.9	0.28%	100003.3	269.0	4.92%	280771.1	500.2	9.50%
01-Jan-12	483.9	16.7	0.28%	712932.0	738.7	14.88%	563984.1	640.6	11.57%
01-Apr-12	480.3	16.6	0.28%	212723.1	365.0	7.40%	425973.1	573.5	10.88%
01-Jul-12	475.7	16.5	0.28%	3105969.0	1550.4	18.73%	3129791.2	1665.4	19.79%
01-Oct-12	486.4	16.7	0.27%	207146.7	369.3	6.01%	415712.8	575.5	10.07%
01-Jan-13	489.4	16.7	0.28%	1294129.2	1011.6	18.82%	1377382.4	1054.0	18.70%
01-Apr-13	488.1	16.7	0.28%	138162.1	331.0	5.81%	218413.9	411.7	7.66%
01-Jul-13	488.8	16.7	0.28%	16042.9	115.3	1.48%	336698.2	435.1	5.43%
01-Oct-13	480.6	16.6	0.27%	560158.8	647.9	9.81%	1032525.2	929.7	14.39%
01-Jan-14	485.1	16.7	0.28%	1456764.5	1052.4	18.81%	330967.2	426.1	7.19%
Mean	491.7	16.8	0.28%	707110.1	639.2	10.70%	1346792.9	820.4	12.88%
Std. Dev.	12.8	0.2	0.00%	763294.4	383.7	6.11%	2579493.5	622.7	7.68%

Table 4: Recursive Forecast Evaluation for TBATS: Medium and Long Horizons

Forecast	4 weeks forecast			13 weeks forecast			52 weeks forecast		
made on	MSE	MAD	MAPE	MSE	MAD	MAPE	MSE	MAD	MAPE
01-Oct-09	1013868	918.7	16.49%	2353445	1405.9	24.75%	3820479	1742.1	27.04%
01-Jan-10	543562.2	624.8	11.00%	974125.5	851.1	15.33%	1448650	1052.4	18.17%
01-Apr-10	1069644	937.1	17.11%	3091255	1470.0	22.21%	4303097	1851.9	28.55%
01-Jul-10	41085288	5934.9	71.02%	1.99E+08	12840.1	186.83%	2.65E+09	44997.8	778.74%
01-Oct-10	1494019	1175.3	21.61%	1232523	1025.1	18.42%	1201341	950.3	15.66%
01-Jan-11	88205.83	238.4	4.05%	104762.4	265.5	4.64%	361213.9	471.3	7.67%
01-Apr-11	62116.24	196.0	3.58%	536993.5	453.8	6.74%	478524	444.4	6.59%
01-Jul-11	1501301	951.5	11.14%	740675.6	648.5	8.68%	801327.6	771.5	13.05%
01-Oct-11	153536.4	341.6	6.43%	201377.5	374.1	6.97%	428429.2	517.2	8.53%
01-Jan-12	422858.7	579.0	10.13%	441930.7	593.6	10.65%	630577.5	653.4	11.15%
01-Apr-12	415496.2	571.5	10.76%	799863.7	650.7	10.39%	1075568	739.1	11.33%
01-Jul-12	1945752	1208.9	14.90%	1196185	897.5	11.67%	601043.2	568.8	8.73%
01-Oct-12	1197938	959.9	17.64%	1721923	1126.6	21.46%	1460747	1021.7	17.38%
01-Jan-13	1157233	917.4	16.31%	1486372	1048.8	18.70%	2260339	1282.3	21.82%
01-Apr-13	366810.1	548.8	10.57%	1077692	784.4	12.61%	1459426	825.8	12.36%
01-Jul-13	1106239	872.0	10.81%	1804118	1149.6	17.93%	3103622	1631.3	28.55%
01-Oct-13	335939	435.5	7.41%	231514.2	372.9	6.50%	453747.9	492.4	7.97%
01-Jan-14	280415.8	422.5	7.31%	317541.9	474.3	8.32%	707423.9	714.6	12.47%
Mean	3013345.7	990.8	14.90%	12095019.1	1468.5	22.93%	1.5E+08	3.4E+03	57.54%
Std. Dev.	9517805.4	1271.4	14.85%	46751428.0	2860.0	41.35%	6.2E+08	1.0E+04	180.13%

We examine whether the forecast errors stabilize over longer forecast horizons in Table 4. From Table 4, we note that forecasts made on 01-Jul-10 have a serious problem. Despite excellent model fitting and good 1-day forecast (MAPE around 12%), the MAPE exceeds 100% for the 13-week and the 52-week forecast horizons. The parameters for this model period do not differ significantly from the others in Table 16 in the Appendix, and do not provide any clues to its poor forecast performance. Forecasts made on some of the other dates also, for example 01-Oct-09, 01-Apr-10, 01-Jan-13 and 01-Jul-13, yield MAPE greater than 20% in the long-term forecast horizon (52-weeks). These corresponding forecasts for the short and medium term horizons are also poor. A few forecasts are good in between, such as the ones made on 01-Apr-11, 01-Jan-11 and 01-Oct-13.

Without a closer look at the coding of the TBATS package, it is difficult to ascertain the reason behind the poor performance of TBATS in some of the recursive forecasts, while the fit for the model period is excellent. However, over-fitting the data is plausible and the unreliability of TBATS with four levels of seasonality is a major concern.

We next study the performance of recursive forecast when a version of the two-stage framework is applied. The specific form of the two-stage framework that we adopt is M2 (as numbered in Table 2), i.e. 'best' regression (in stage 1) followed by multiplicative seasonal adjustment. The selected regression models are 'best' in the sense of having highest R-square with all significant regression coefficients, based on available data at the corresponding stage of prediction. The errors in fit for the different forecast horizons are reported in Tables 5 and 6.

Table 5: Recursive Forecast Evaluation for M2 (two-stage): Model Period and Short Horizon Forecast

Forecast made on	Model period MSE	Model period MAD	Model period MAPE	1 day forecast MSE	1 day forecast MAD	1 day forecast MAPE	1 week forecast MSE	1 week forecast MAD	1 week forecast MAPE
01-Oct-09	116697.5	258.1	4.25%	12183.7	85.1	1.76%	101249.5	262.8	4.99%
01-Jan-10	109873.0	248.8	4.13%	409203.7	549.0	10.33%	126920.1	285.5	4.95%
01-Apr-10	102799.1	238.5	3.97%	36116.9	169.3	2.98%	95104.2	261.3	4.86%
01-Jul-10	109928.7	248.0	4.12%	38271.7	166.8	2.57%	442555.4	520.1	6.54%
01-Oct-10	129503.0	273.9	4.45%	220575.1	423.4	6.78%	163822.7	363.7	6.64%
01-Jan-11	126890.5	269.1	4.40%	69560.7	237.0	4.66%	42528.8	166.8	2.98%
01-Apr-11	123317.0	265.6	4.35%	35456.9	172.1	2.97%	38630.5	160.1	3.00%
01-Jul-11	126505.0	269.5	4.43%	35291.7	162.3	2.39%	290855.9	387.3	5.43%
01-Oct-11	132603.7	275.9	4.48%	21729.7	96.2	1.78%	83297.9	226.9	4.30%
01-Jan-12	128719.5	271.0	4.43%	40804.3	176.8	3.64%	266723.8	410.2	7.06%
01-Apr-12	124614.6	265.7	4.36%	13064.4	98.8	2.02%	30323.1	131.3	2.41%
01-Jul-12	124370.0	265.1	4.36%	594232.1	720.5	8.88%	474050.2	578.9	7.11%
01-Oct-12	130168.0	272.8	4.45%	131428.0	354.4	6.63%	209937.5	388.1	6.68%
01-Jan-13	135074.5	274.0	4.52%	645929.5	662.5	12.54%	161379.4	298.5	5.45%
01-Apr-13	134055.3	273.3	4.52%	27314.3	123.7	2.13%	27029.1	139.2	2.61%
01-Jul-13	134773.3	274.5	4.55%	126573.2	314.5	4.35%	246493.2	413.9	5.23%
01-Oct-13	137757.9	278.5	4.58%	220236.8	428.4	7.88%	116977.0	283.9	4.87%
01-Jan-14	136952.1	277.3	4.58%	987659.5	844.3	15.18%	1045515.5	827.8	13.37%
Mean	125811.3	266.7	4.39%	203646.2	321.4	5.53%	220188.5	339.3	5.47%
Std. Dev.	10054.0	11.3	0.17%	277650.7	235.4	4.02%	244398.2	175.2	2.47%

Table 6: Recursive Forecast Evaluation for M2 (two-stage): Medium and Long Horizons

Forecast made on	4 week forecast MSE	4 week forecast MAD	4 week forecast MAPE	13 week forecast MSE	13 week forecast MAD	13 week forecast MAPE	52week forecast MSE	52week forecast MAD	52week forecast MAPE
01-Oct-09	72629.6	214.1	3.95%	114979.0	267.4	4.88%	235376.5	366.1	5.54%
01-Jan-10	82738.9	236.2	4.07%	71900.7	214.7	3.83%	186850.4	311.3	4.73%
01-Apr-10	50303.0	180.4	3.40%	176422.5	308.9	4.84%	191815.2	309.8	4.67%
01-Jul-10	539650.4	623.5	7.52%	342014.6	460.9	6.43%	170139.4	303.6	4.76%
01-Oct-10	123343.8	298.3	5.51%	102648.1	243.1	4.20%	153915.0	294.3	4.81%
01-Jan-11	162181.7	272.2	4.46%	100838.8	232.2	3.91%	139158.0	275.6	4.47%
01-Apr-11	52521.8	185.8	3.47%	169148.6	319.1	5.20%	135082.1	270.6	4.43%
01-Jul-11	378393.8	461.7	5.65%	214822.5	344.1	4.79%	124558.7	262.9	4.31%
01-Oct-11	126946.0	297.4	5.51%	80539.1	220.6	3.93%	122929.4	266.5	4.38%
01-Jan-12	134951.3	271.8	4.66%	70838.3	198.6	3.52%	160696.4	287.8	4.89%
01-Apr-12	44805.7	169.2	3.20%	135472.7	293.8	4.78%	170378.3	296.2	5.01%
01-Jul-12	345201.5	472.4	5.91%	226493.1	367.8	5.20%	175809.9	303.2	5.07%
01-Oct-12	131319.3	303.5	5.59%	219018.6	305.6	5.44%	171826.9	302.8	5.17%
01-Jan-13	224418.9	342.7	5.81%	123238.3	265.8	4.68%	146834.5	290.7	4.79%
01-Apr-13	40617.7	155.9	2.96%	151207.8	295.1	4.90%	185945.2	320.5	5.23%
01-Jul-13	343852.8	484.1	5.96%	202951.2	354.6	5.06%	207483.0	332.4	5.51%
01-Oct-13	121038.6	293.2	5.43%	118801.3	263.5	4.70%	311482.8	381.3	6.46%
01-Jan-14	562690.6	544.7	8.71%	290222.4	385.9	6.17%	394286.1	425.2	7.09%
Mean	196533.6	322.6	5.10%	161753.2	296.8	4.80%	188031.6	311.2	5.07%
Std. Dev.	166686.6	138.3	1.51%	75812.9	67.9	0.76%	67967.4	42.3	0.72%

The performance of our new proposed two-stage framework is very encouraging. Note that while M2 is not the best performing version in Table 2, it still outperforms TBATS in almost all the recursive forecasts across all forecast horizons considered (except 4-week 01-Jan-11, 01-Jan-14 and 13-week 01-Jan-11 where the forecasts are comparable). We also observe that the accuracy in longer horizon forecast is somewhat more stable than shorter horizon forecast, like one day – this is along expected lines.

We now examine results of another variation of the two-stage framework, M4, in Tables 7 and 8. M4 combines regression for the first stage and polynomial multiplicative seasonality for the second stage, and is attractive in model parsimony and forecast performance.

Table 7: Recursive Forecast Evaluation for M4 (two-stage): Model Period and Short Horizon Forecast

Forecast made on	Model period MSE	Model period MAD	Model period MAPE	1 day forecast MSE	1 day forecast MAD	1 day forecast MAPE	1 week forecast MSE	1 week forecast MAD	1 week forecast MAPE
01-Oct-09	146578.0	298.2	4.93%	63333.1	224.0	4.27%	141620.3	322.4	6.13%
01-Jan-10	141414.5	293.4	4.90%	387112.5	521.5	9.66%	160645.1	327.4	5.46%
01-Apr-10	135750.7	287.2	4.80%	69310.8	222.3	4.04%	121400.6	280.4	5.34%
01-Jul-10	141817.1	294.1	4.91%	39487.1	162.4	2.53%	443839.5	522.4	6.55%
01-Oct-10	158613.0	311.9	5.10%	236180.2	447.2	6.94%	197828.5	375.1	6.93%
01-Jan-11	156691.1	309.2	5.09%	67067.6	229.3	4.43%	82153.0	244.2	4.17%
01-Apr-11	154058.6	306.9	5.06%	78723.6	208.6	3.50%	83046.6	239.3	4.35%
01-Jul-11	157061.1	309.5	5.12%	42170.2	166.1	2.53%	286141.7	399.1	5.62%
01-Oct-11	161517.8	312.8	5.11%	31613.2	136.8	2.57%	116640.7	270.6	5.20%
01-Jan-12	158072.5	309.2	5.09%	43184.3	171.3	3.47%	302500.8	430.6	7.39%
01-Apr-12	154538.1	305.8	5.05%	26046.8	132.6	2.62%	66652.2	210.3	3.88%
01-Jul-12	154267.5	305.1	5.05%	596783.9	720.5	8.86%	458637.0	576.6	7.00%
01-Oct-12	158850.0	309.9	5.09%	166580.6	354.5	6.81%	244345.4	401.3	7.06%
01-Jan-13	163879.2	311.5	5.17%	618784.8	624.1	11.74%	191735.0	340.5	6.10%
01-Apr-13	163301.8	311.2	5.18%	74026.3	230.0	4.14%	67311.8	212.6	3.82%
01-Jul-13	163988.7	312.2	5.20%	153469.6	314.6	4.22%	248186.9	411.4	5.14%
01-Oct-13	166132.1	314.5	5.20%	252890.6	428.6	8.08%	136532.8	305.7	5.32%
01-Jan-14	165521.9	313.7	5.21%	959899.3	809.5	14.51%	1068127.0	841.2	13.60%
Mean	155669.7	306.5	5.07%	217036.9	339.1	5.83%	245408.1	372.8	6.06%
Std. Dev.	8889.8	7.9	0.12%	261817.8	209.0	3.50%	236423.7	154.4	2.18%

Table 8: Recursive Forecast Evaluation for M4 (two-stage): Medium and Long Horizons

Forecast	4 week forecast			13 week forecast			52week forecast		
made on	MSE	MAD	MAPE	MSE	MAD	MAPE	MSE	MAD	MAPE
01-Oct-09	113300.8	278.0	5.12%	153685.3	315.8	5.74%	264170.8	398.1	6.06%
01-Jan-10	123932.9	287.7	4.79%	112449.3	271.3	4.72%	214612.1	352.9	5.43%
01-Apr-10	84148.5	239.5	4.55%	200849.8	347.2	5.52%	219255.7	353.1	5.42%
01-Jul-10	543148.5	623.8	7.47%	348256.1	468.6	6.52%	198363.2	341.9	5.42%
01-Oct-10	160161.1	329.6	6.17%	139553.3	294.4	5.06%	182427.8	329.4	5.43%
01-Jan-11	198789.5	319.6	5.19%	141392.2	286.0	4.77%	167403.7	313.0	5.13%
01-Apr-11	88240.8	244.5	4.56%	196758.7	349.3	5.73%	162736.7	310.1	5.14%
01-Jul-11	379517.7	466.2	5.70%	222340.0	352.9	4.92%	152567.3	302.3	5.02%
01-Oct-11	163234.5	338.3	6.34%	116842.8	275.3	4.89%	151139.1	305.5	5.08%
01-Jan-12	172249.8	317.7	5.41%	109466.4	260.3	4.58%	187725.4	324.6	5.55%
01-Apr-12	78579.5	232.5	4.42%	164229.0	325.1	5.36%	197077.8	331.5	5.64%
01-Jul-12	340379.4	473.9	5.90%	234247.5	376.5	5.31%	202475.0	338.2	5.68%
01-Oct-12	168476.6	337.0	6.31%	251188.5	350.7	6.23%	199067.5	337.8	5.78%
01-Jan-13	261034.5	383.4	6.49%	161457.0	308.7	5.37%	174403.7	325.0	5.38%
01-Apr-13	78122.7	230.3	4.33%	179483.4	335.5	5.62%	213327.5	352.3	5.78%
01-Jul-13	347495.7	487.0	5.95%	212434.2	363.2	5.17%	233981.1	361.9	6.02%
01-Oct-13	151289.7	323.0	6.05%	152488.1	304.5	5.40%	337223.7	409.4	6.96%
01-Jan-14	596028.3	571.4	9.12%	328154.1	418.0	6.65%	418784.1	450.2	7.54%
Mean	224896.1	360.2	5.77%	190293.1	333.5	5.42%	215374.6	346.5	5.69%
Std. Dev.	156568.8	117.5	1.19%	68004.8	52.8	0.59%	67168.7	38.7	0.65%

While, in general, the forecast errors for the polynomial multiplicative seasonality are slightly worse when compared to classical multiplicative seasonality correction, we expect this because the polynomial method of seasonality estimation uses far fewer parameters when compared to the classical method. The consistency in forecast as captured through the standard deviation over the various periods is a testimony to the effectiveness of the method.

5 Discussions on Miscellaneous Aspects

5.1 Model Selection Based on Lack-of-fit and Number of Parameters

So far in this study, the effectiveness of the different methods have been compared in terms of forecast evaluation criteria such as MSE, MAD and MAPE, computed both during the model or training period as well as the hold-out or test period. On the basis of these forecast evaluation measures, a few methods in the two-stage framework proposed in this study are recommended for modeling time-series data with many levels of seasonality. We reiterate that such recommendations are based on consistent superior performance in the test period (in addition to model period) for varying horizons. The recommendations however do not take into account model complexity.

As an alternative, penalized likelihood methods such as AIC, BIC or AICC are often preferred as model selection criteria. However, since a penalized likelihood criteria depends on modeler's choice of the likelihood, its fairness as a model selection criterion across vastly different sets of methods is not beyond debate. In this article, we propose to use a penalized lack-of-fit yardstick ψ defined by:

$$\psi = n \ln(MSE) + 2p, \tag{12}$$

where p is the total number of parameters in the model and n is the number of data points. Note that (12) is indeed the AIC criterion corresponding to simple normal likelihood, and by replacing 2 with suitable constants, we get other equivalent information-based criteria such as BIC and AICC. However, irrespective of the likelihood presumed or the methods adopted, (12) can be conveniently interpreted in terms of penalized lack-of-fit, with the first term capturing how well the model explains the data, and the second capturing model complexity. A convenient trade-off then emerges when model selection is done by minimizing ψ. Computing ψ in (12) is also very easy, regardless of the complexity of the method, and may be carried out for model period as well as hold-out period.

Table 9: $\psi \times 10^{-6}$: NY energy load data

	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17
Model period	6.66	6.62	6.76	6.73	6.65	6.64	6.76	6.740	6.45	6.41	6.60	6.55	6.39	6.34	6.55	6.50	3.46
Forecast period	1.36	1.36	1.36	1.36	1.42	1.42	1.42	1.42	1.36	1.36	1.36	1.36	1.34	1.34	1.35	1.34	1.42

Figure 12: Comparing Penalized Lack-of-fit ψ: TBATS vs. Two-stage Methods: Recursive Forecast

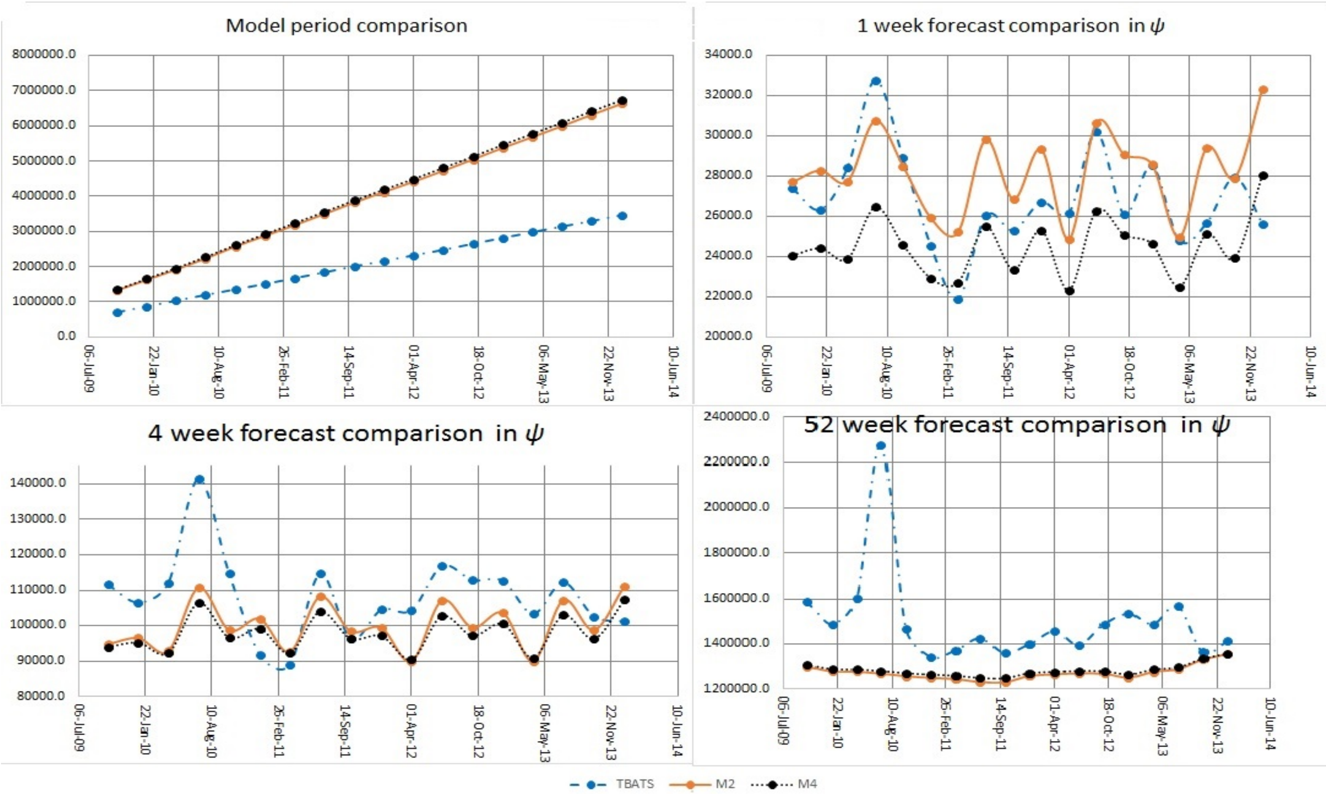

In Figure 12, we compare ψ for TBATS, M2 and M4 across the model period and three forecast horizons – 1-week, 4-weeks and 52-weeks. We are aware that typically AIC values are not calculated for the forecast period. However, since all our results seem to indicate that the TBATS method suffers from over-fitting, it is obvious that the ψ metric will be lower for TBATS in the model period. However, in the forecast period, where the two-stage framework does a better job of prediction, we see that ψ values are lower in short, medium and long term horizons. In the 1-week horizon forecast, M2 and TBATS have comparable ψ but M4 is clearly better. With the 4-week and 52-week forecasts, both M2 and M4 have lower and comparable ψ values. Recall that both M2 and M4 are not the best two-stage methods in terms of MSE. They were chosen to illustrate better performance both in terms of MSE and because they are easy to implement. In Table 10, a comparison of the time time taken for TBATS, M2 and M4 to run is presented. With M2 and M4, forecasts were available in a few minutes, whereas the standard implementation of TBATS took a few days.

Table 10: Time Taken (in seconds) in Recursive Forecast: TBATS vs. Two-stage Methods

	Forecast made on								
	1-Oct-09	1-Jan-10	1-Apr-10	1-Jul-10	1-Oct-10	1-Jan-11	1-Apr-11	1-Jul-11	1-Oct-11
TBATS	988.9	1290.9	1346.4	976.9	10479.0	20754.2	13539.6	17806.2	17374.1
M2	2.30	2.08	2.34	3.05	3.55	3.68	4.35	4.46	4.76
M4	0.82	1.03	1.11	1.28	1.47	1.64	1.7	1.97	2.16

	Forecast made on								
	1-Jan-12	1-Apr-12	1-Jul-12	1-Oct-12	1-Jan-13	1-Apr-13	1-Jul-13	1-Oct-13	1-Jan-14
TBATS	16064.4	24438.7	28932.8	31996.6	24434.8	32519.6	29625.7	41104.4	41105.1
M2	5.21	5.72	6.13	6.57	7.08	7.81	8.02	8.72	9.04
M4	2.31	2.47	2.65	2.8	2.92	3.11	3.31	3.5	3.72

5.2 Level of Separation between High and Low Frequencies

In the proposed two-stage framework, the high and low frequency seasonalities are treated differentially. While the seasonalities (consisting of one or more levels) in the low-frequency are proposed to be modeled using (a combination of) methods such as regression, ARIMA, and TBATS, the seasonalities in the high-frequency are estimated via a classical decomposition method or via a function-based method such as polynomial fitting or trigonometric fitting. Typically, although not as a requirement in all variations, the seasonalities at low and high frequencies would be assumed to be non-interactive as per the model and this may accordingly play an important role in deciding the suitable level of separation between the low and high frequencies. Naturally, errors in forecast and model evaluation, as captured during the model and hold-out periods would also provide guidance in selecting this suitable level of separation, as also would the criteria such as penalized likelihood or lack-of-fit, as described in Section 5.1.

For the NY electricity load data, the weekly separation between high and low frequencies, means that Y_{low} is load data averaged over the duration of a week. Thus, time series models such as TBATS, ARIMA will model only annual seasonality in the first stage. The seasonality estimation methods in the second stage capture the within-hour, within-day and weekly seasonalities. Table 11 shows that the model and forecast period errors are marginally worse than separation at the daily level.

Table 11: Forecast Evaluation with NY Energy Load: Two Stage Models with Weekly Separation in High/Low Frequency

Stage 1	Stage 2	Model period MSE	Model period MAD	Model period MAPE	Hold-out period MSE	Hold-out period MAD	Hold-out period MAPE
Regression	Classical additive	327845.0	423.3	6.97%	454418.8	488.5	8.28%
	Classical multiplicative	320074.8	413.6	6.71%	442567.8	473.1	7.93%
	Polynomial additive	244872.3	368.8	5.95%	435310.6	479.9	8.04%
	Polynomial multiplicative	233215.6	357.0	5.74%	423514.4	464.7	7.75%
Single level TBATS	Classical additive	378497.9	441.3	7.12%	538165.5	523.5	8.78%
	Classical multiplicative	372439.5	432.2	6.89%	526442.9	507.3	8.42%
	Polynomial additive	295528.7	387.9	6.10%	519056.6	513.5	8.51%
	Polynomial multiplicative	286172.4	382.0	6.03%	507136.4	501.0	8.27%
TBATS on regression residuals	Classical additive	312936.1	403.6	6.60%	440677.5	473.5	7.99%
	Classical multiplicative	305903.9	393.4	6.34%	429383.2	457.4	7.64%
	Polynomial additive	229959.9	348.4	5.57%	421569.9	464.0	7.73%
	Polynomial multiplicative	219239.5	339.2	5.42%	410308.5	449.0	7.45%
ARIMA with xreg	Classical additive	312681.1	405.8	6.65%	519328.0	535.6	9.21%
	Classical multiplicative	305119.8	395.3	6.39%	510624.5	523.3	8.94%
	Polynomial additive	229708.9	349.9	5.60%	500117.1	529.5	9.05%
	Polynomial multiplicative	218455.0	340.2	5.45%	490358.1	518.6	8.83%

Next, we consider separation at the hourly level for the NY energy load data. Here Y_{low} will be the energy load data averaged for every hour. Stage 1 models account for within-day, weekly and annual seasonalities. The TBATS in Table 12 has three levels of seasonality, daily, weekly and annual. Stage 2 estimates the within-hour seasonalities. We observe in Table 12 that the model period errors improve greatly compared to both daily and weekly separation. However, the forecast period errors are worse than those of daily separation (Table 2). This level of separation seems to suffer from the problem of over-fitting.

Table 12: Forecast Evaluation with NY Energy load:Two-stage Models with Hourly Separation in High/Low Frequency

Stage 1	Stage 2	Model period MSE	Model period MAD	Model period MAPE	Hold-out period MSE	Hold-out period MAD	Hold-out period MAPE
Regression	Classical additive	192539.1	334.9	5.65%	427653.7	465.7	7.96%
	Classical multiplicative	192552.2	334.9	5.65%	427694.1	465.7	7.96%
	Polynomial additive	197328.6	339.8	5.74%	432257.9	469.8	8.03%
	Polynomial multiplicative	197150.2	339.7	5.74%	432280.8	469.8	8.03%
3 level TBATS	Classical additive	10109.9	74.0	1.21%	1319803.3	912.6	16.27%
	Classical multiplicative	10106.1	73.8	1.21%	1320002.1	912.6	16.26%
	Polynomial additive	15290.9	88.7	1.47%	1317013.2	911.7	16.26%
	Polynomial multiplicative	15302.8	88.7	1.47%	1316375.4	911.5	16.25%
TBATS on regression residuals	Classical additive	17663.6	78.2	1.34%	433669.6	456.8	7.73%
	Classical multiplicative	17665.4	78.3	1.34%	433721.5	456.7	7.73%
	Polynomial additive	22846.6	97.1	1.66%	440825.7	463.4	7.85%
	Polynomial multiplicative	22833.1	97.0	1.66%	441316.4	463.7	7.85%
ARIMA with xreg	Classical additive	5454.3	51.3	0.85%	739595.3	677.3	11.37%
	Classical multiplicative	5451.5	51.3	0.85%	739599.9	677.3	11.37%
	Polynomial additive	10718.2	74.9	1.26%	744200.7	679.6	11.41%
	Polynomial multiplicative	10720.7	74.9	1.26%	744380.3	679.9	11.42%

5.3 Forecast Combination

Forecast combination is an established mechanism of further improving forecast performance [Clemen [1989], Elliott et al. [2006]]. Following this, we try to come up with improved forecast by combining the various methods listed in Table 2, i.e. M1, ..., M17 (of which, the first 16 are different variations of the two-stage methods and M17 represents the 4-level TBATS). We adopt one of the four principles of combining forecasts, viz. (a) simple average (b) variance-based, i.e. minimizing variance of the combined forecast, (c) basing on ordinary least square (OLS), and (d) basing on robust regression. Typically, a subset of methods out of these 17, are combined using these principles. A summary of error evaluation measures from some of these combinations is provided in Table 13.

Table 13: Results from Combining Forecasts: NY Energy Load

Methods combined	Criteria for combining	Model Period			Hold-out Period		
		MSE	MAD	MAPE	MSE	MAD	MAPE
M2, M4,M6,M8,M10,M12,M14,M16,M17	Simple	68021.3	191.5	3.11%	349450.2	403.7	6.78%
	Variance based	544.2	17.8	0.29%	681780.7	699.8	12.20%
	OLS	483.4	16.6	0.28%	703312.3	712.4	12.43%
	Robust Reg	483.6	16.6	0.28%	703548.8	712.6	12.43%
M2,M6,M10,M14,M17	Simple	51333.0	159.2	2.56%	350574.4	411.6	6.93%
	Variance based	498.8	17.0	0.28%	692496.9	706.1	12.32%
	OLS	483.8	16.6	0.28%	703431.6	712.5	12.43%
	Robust Reg	483.9	16.6	0.28%	703648.9	712.6	12.43%
M2, M6, M10, M14	Simple	80091.4	198.7	3.20%	350754.0	392.5	6.53%
	Variance based	79227.7	197.7	3.18%	338083.4	383.1	6.37%
	OLS	78214.5	194.9	3.12%	329798.4	377.7	6.28%
	Robust Reg	79255.2	193.5	3.08%	332610.6	379.3	6.31%
M4, M14, M17	Simple	42905.0	154.7	2.54%	347757.4	422.0	7.12%
	Variance based	485.4	16.7	0.28%	699298.8	710.0	12.39%
	OLS	483.6	16.6	0.28%	703353.7	712.4	12.43%
	Robust Reg	483.7	16.6	0.28%	703611.2	712.6	12.43%
M10, M14	Simple	81879.0	203.0	3.28%	327932.2	374.4	6.19%
	Variance based	81602.4	202.4	3.27%	326165.4	374.2	6.18%
	OLS	80753.4	200.3	3.22%	321984.5	379.0	6.26%
	Robust Reg	81106.4	199.8	3.21%	322671.7	384.3	6.34%

We note that typically if TBATS (M17) is one of the constituent forecasts, the performance in the forecast period continues to be poor while that in the model period is good. We notice a good improvement by considering variance-based combination of M10 and M14; this results in a reduction of MAPE in the forecast period from 6.70% to a combined 6.18% with negligible change in the MAPE in the model period.

5.4 Results from Implementing and Comparing Two-stage Methods vis-à-vis TBATS in other Data sets

In Livera et al. [2011], TBATS was introduced with illustration on three datasets. We implemented and compared the two-stage framework with TBATS for each of these data sets. As an illustration, in this subsection, we present results from the call center data [Hyndman [2011]], which contains five-minute call volume on weekdays between 7:00 am and 9:05 pm, from March 3, 2003, to May 23, 2003 in a large North American commercial bank.

TBATS is implemented with two levels of seasonality – hourly and daily. For the two-stage framework, we work with the average daily data in the first stage and exhaustively search for the best fitting model. Note that since there is no information on additional covariates, and there are only two levels of seasonality, the benefit of the two-stage method is naturally limited. In this context, regression simply accounts for a linear trend over time. Using classical and polynomial methods, both additive and multiplicative seasonalities are estimated in the second stage. Results in terms of the forecast evaluation in the model period (first 134 days) and hold-out period (next 30 days) are provided in Table 14.

Table 14: Forecast Evaluation for Call Center Data

Stage 1	Stage 2	Model period MSE	Model period MAD	Model period MAPE	Hold-out period MSE	Hold-out period MAD	Hold-out period MAPE
Regression	Classical additive	703.8	19.8	10.92%	642.0	19.7	12.86%
	Classical multiplicative	679.4	19.4	10.74%	646.6	19.7	12.63%
	Polynomial additive	970.4	24.1	14.50%	968.9	24.5	16.95%
	Polynomial multiplicative	973.4	24.0	14.42%	983.7	24.7	16.88%
TBATS with weekly seas.	Classical additive	608.3	18.7	11.22%	521.7	17.3	11.42%
	Classical multiplicative	546.6	17.3	9.97%	499.4	16.9	11.21%
	Polynomial additive	875.0	23.0	14.34%	848.6	22.6	16.00%
	Polynomial multiplicative	847.6	22.5	13.77%	836.8	22.4	15.69%
ARIMA	Classical additive	555.7	18.0	10.72%	529.7	17.5	11.61%
	Classical multiplicative	488.9	16.7	9.63%	509.7	17.1	11.39%
	Polynomial additive	822.3	22.3	13.99%	856.4	22.8	16.23%
	Polynomial multiplicative	788.5	21.8	13.49%	845.4	22.5	15.87%
TBATS with 2 levels of Seasonality		260.4	12.3	7.27%	779.6	21.2	12.62%

Table 14 shows that while TBATS offers a better fit, the performance of many of the two-stage methods is better in the test-period, even though the data context is not suited to derive fully the benefits of the two-stage methods. Changing the duration of the model and forecast period does not critically alter this observation and hence the corresponding results are not reported here.

5.5 Other Variations of Models and Methodology in the Two-stage Framework

5.5.1 Variation in Stage 1

In this study, we proposed to start the modeling in the second stage by working on $Y - \bar{Y}_{\text{day}}$ or $\frac{Y}{\bar{Y}_{\text{day}}}$, depending on additive or multiplicative framework. A possible alternative would be to replace $\bar{Y}_{\text{day}}$ by $\hat{Y}_{\text{day}}$, where the latter is the estimate obtained in the first stage. Thus, one could estimate the high frequency multiplicative seasonality by analyzing $\frac{Y}{\hat{Y}_{\text{day}}}$. We tend to favor the former (as exhibited in this article), not just because it would be computationally simpler (the data in the second stage would not be dependent on the method adopted in first stage), but we think this is also more in line with our starting philosophy of level of separation in the frequency. However, the alternative has its own appeal in terms of potentially working better with some data sets.

5.5.2 Alternative Methods of Seasonality Estimation at High Frequency: Trigonometric Estimation

In the aforementioned methods of estimating additive/multiplicative seasonality, the number of parameters to be estimated is large, especially in the case of the classical decomposition-like method. For one-minute data, with the classical method, we need 24 hourly values and 24×60 minute-level values to be estimated. Though the polynomial curve fitting method alleviates this issue by bringing the number of parameters to be estimated down to the order of the polynomial that is fitted, we tried yet another function-based approach to seasonality estimation – trigonometric functions.

In this method, we use the fact that any periodic function can be represented as a sum of simple sine/cosine waves. This sum is called Fourier series. We perform a Fast Fourier Transform (FFT) on the adjusted or residual data, $\tilde{Y} = Y - Y_{\text{low}}$ or $\tilde{Y} = \frac{Y}{Y_{\text{low}}}$, depending on additive or multiplicative model (2) or (3). The FFT on $\tilde{Y}$ returns a complex number for each frequency. Depending on the number of levels of seasonality in the data, we pick the appropriate number of highest amplitude frequencies in the FFT output. Suppose, we need one cosine function to represent $\tilde{Y}$. Let $a + ib$ be the complex number associated with the highest amplitude frequency f_0. The cosine function is $\hat{\tilde{Y}}_n = \sqrt{a^2 + b^2}\cos(2\pi f_0 n + \arctan(\frac{b}{a}))$, where $\sqrt{a^2 + b^2}$ is the amplitude and $\arctan(\frac{b}{a})$ is the phase.

The results of using this kind of trigonometric estimation of S_{high} are similar to the polynomial method, and are omitted for brevity.

6 Summary and Future Research

In this article, we proposed a two-stage framework which uses existing methods more efficiently to model and forecast time series data with *multiple levels* of seasonality. The split between high and low frequencies ensures that the modeling and forecast is carried out very quickly and with satisfactory accuracy. As shown in Tables 2 – 8, the performance of the selected few two-stage methods is consistently satisfactory in absolute terms, and relatively, better than TBATS in all these scenarios. Table 10 shows that the two-stage methods are way ahead of TBATS in terms of computational speed when data has more levels of seasonality. Since the two-stage methods take negligible amount of time in execution even for a very large time-series data set, recorded every (few) minute(s), the forecast can be carried out in real time. The ability to use information from covariates is another major advantage of this framework. Usage of all available information usually guarantees better forecasts. Compared to existing methods, the two-stage method does not suffer from over-fitting and also does not suffer from inconsistency. The predictions from the two-stage methods are typically more consistent when compared with TBATS.

In a follow-up study, we will examine the impact of frequency on forecast accuracy. In particular, we will examine whether data recorded at higher frequency leads to a better forecast accuracy. We will also examine the impact of time of recording within a period on the forecast accuracy. Preliminary results were reported in Lakshmanan and Das [2016].

For modeling and forecasting time series with complex, multiple seasonality, more challenges exist. Similar to the trigonometric method to estimate seasonality, other ways could be incorporated. Nested seasonality can also be exploited. The framework may also be expanded to include probabilistic forecasting [Hong et al. [2016]]. The research towards this is in a planning stage.

A package within R for the new two-stage framework would make it convenient to execute and possibly lead to more practitioners and academicians adopting it. The work towards this is currently under way.

Acknowledgement: Mr. Akshay Kumar Singh worked as a research assistant for a project which led to the present research. The second author would like to acknowledge the help rendered by Mr. Singh in terms of computation in R during the project.

References

T. Au, G.Q. Ma, and SN. Yeung. Automatic forecasting of double seasonal time series with applications on mobility network traffic prediction. 2011 Joint Statistical Meetings, Miami Beach, Florida, USA, July 30-August 4 2011.

R.T. Clemen. Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5:559 – 583, 1989.

CoreTeam. R: A language and environment for statistical computing. `http://www.r-project.org`, 2014.

G. Elliott, C.W. Granger, and A. Timmermann. *Handbook of Economic Forecasting*, chapter Forecast combinations, pages 135–196. Elsevier, 2006.

P.G. Gould, A.B. Koehler, J.K. Ord, R.D. Snyder, R.J. Hyndman, and F.V. Araghi. Forecasting time series with multiple seasonal patterns. *European Journal of Operational Research*, 191:207 – 222, 2008.

T. Hong, P. Pinson, S. Fan, H. Zareipour, A. Troccoli, and R.J. Hyndman. Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. *International Journal of Forecasting*, 32(3):896 – 913, 2016.

R.J. Hyndman. Call center data. `http://robjhyndman.com/data/callcenter.txt`, 2011. Accessed: 2017-01-30.

R.J. Hyndman. *forecast: Forecasting functions for time series and linear models.* `http://github.com/robjhyndman/forecast`, 2016.

R.J. Hyndman and S. Fan. Monash electricity forecasting model. May 2015.

A. Lakshmanan and S. Das. Effect of data frequency on forecast accuracy. International Symposium on Forecasting, Santander, Spain, 2016.

A.M.D. Livera, R.J. Hyndman, and R.D. Snyder. Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association*, 106(496):1513 – 1527, 2011.

NYISO. New york electricity load data. `http://www.nyiso.com/public/markets_operations/market_data/load_data/index.jsp`, 2014. Accessed: 2015-01-20.

L.J. Soares and L.R. Souza. Forecasting electricity demand using generalized long memory. *International Journal of Forecasting*, 2006.

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

Appendix

I. Methods M13–M16 in Table 2 use ARIMA in the first stage of their two-stage implementations. In Table 15, we provide the AIC values for different (p,d,q) values in the ARIMA with xreg function in R, that helps in model selection. The seasonal ARMA orders (P, Q) are fixed to zero, both d and D vary between 0 and 1, while p and q vary between 0 and 5. Based on the minimum AIC value in Table 15, the order is selected as (5,0,5) with seasonal $D = 1$. The corresponding BIC values also point towards the same model, and hence we omit reporting them for the sake of brevity.

Table 15: AIC Values for Seasonal ARIMA $(p, d, q) \times (0, D, 0)$ with xreg: New York Energy Load – Model Period 2008–13

		q					
	p	0	1	2	3	4	5
d = 0, D = 0	0	32191.34	30304.7	29649.51	29492.63	29366.43	29090.76
	1	29392.67	29205.2	29022.62	28951.58	28947.49	29068.94
	2	29304.55	29161.03	28954.58	28951.47	28906.62	29034.46
	3	29120.24	28961	28907.07	28906.61	28826.51	28535.22
	4	29121.29	29098.56	28737.38	28722.91	28716.19	28830.26
	5	28999.29	29082.85	28900.76	28795.46	28198.8	28573.76
d = 1, D = 0	0	29496.97	29399.31	29016.61	28942.83	28938.54	28895.68
	1	29462.27	29302.21	28945.72	28942.61	28897.46	28840.41
	2	29188.71	28952.93	28932.24	28901.22	28817.91	28525.58
	3	29182.16	28950.62	28898.18	28838.05	28578.45	28583.45
	4	29026	28823.42	28642	28390.49	28319.25	28232.83
	5	28700.06	28689.43	28508.27	28301.33	28387.41	28307.57
d = 0, D = 1	0	25532.69	24625.64	24577.46	24555.3	24535.74	24487.81
	1	24731.82	24559.65	24552.17	24545.09	24537.62	24525.24
	2	24634.51	24543.61	24541.07	24460.81	24458.5	24003.24
	3	24584.26	24543.79	24207.77	24108.34	24050.98	24092.13
	4	24529.61	24542.32	24062.19	24053.88	24020.49	24386.44
	5	24531.58	24527.57	24044.41	24053.94	24040.72	**23976**
d = 1, D = 1	0	25039.49	25039.71	24572	24548.58	24535.07	24528.97
	1	25041.03	24928.1	24540.71	24539.39	24533.25	24530.89
	2	24808.35	24808.23	24532.05	24525.88	24441.26	24096.3
	3	24809.57	24809.32	24528.46	24206.39	24102.89	24042.42
	4	24763.68	24496.43	24529.88	24044.2	24127.14	23987.8
	5	24477.78	24434.42	24259.06	24003.93	24037.75	24008.76

II. In Table 16, we provide the TBATS parameters for the recursive forecasts in Section 4. The reported parameters are Box-Cox transformation parameter ω, AR and MA orders (p, q), damping parameters η, in the seasonality, the periodicities and the corresponding number of trigonometric components.

Table 16: TBATS Parameters in Recursive Modeling: New York Energy Load

Forecast made on	TBATS(ω, {p,q}, η, seasonalities)
01-Oct-09	TBATS(1, {5,5}, 0.951, {(12,1), (288,1), (2016,1), (105120,1)})
01-Jan-10	TBATS(1, {5,0}, 0.874, {(12,1), (288,1), (2016,1), (105120,1)})
01-Apr-10	TBATS(1, {4,2}, -, {(12,1), (288,1), (2016,1), (105120,1)})
01-Jul-10	TBATS(1, {4,3}, 1, {(12,1), (288,1), (2016,1), (105120,1)})
01-Oct-10	TBATS(0.358, {0,0}, 0.951, {(12,2), (288,5), (2016,6), (105120,6)})
01-Jan-11	TBATS(0.027, {0,0}, 0.959, {(12,2), (288,10), (2016,6), (105120,7)})
01-Apr-11	TBATS(0.034, {0,0}, 0.959, {(12,2), (288,7), (2016,5), (105120,6)})
01-Jul-11	TBATS(0.084, {0,0}, 0.957, {(12,2), (288,8), (2016,6), (105120,5)})
01-Oct-11	TBATS(0.002, {0,0}, 0.946, {(12,3), (288,6), (2016,6), (105120,6)})
01-Jan-12	TBATS(0.009, {0,0}, 0.954, {(12,2), (288,7), (2016,6), (105120,5)})
01-Apr-12	TBATS(0.26, {0,0}, 0.955, {(12,2), (288,7), (2016,6), (105120,6)})
01-Jul-12	TBATS(0.012, {5,2}, 0.969, {(12,3), (288,8), (2016,6), (105120,7)})
01-Oct-12	TBATS(0.351, {0,0}, 0.96, {(12,2), (288,6), (2016,5), (105120,8)})
01-Jan-13	TBATS(0.399, {0,0}, 0.955, {(12,3), (288,5), (2016,1), (105120,7)})
01-Apr-13	TBATS(0.451, {0,0}, 0.957, {(12,3), (288,6), (2016,5), (105120,7)})
01-Jul-13	TBATS(0.286, {0,0}, 0.947, {(12,3), (288,6), (2016,6), (105120,6)})
01-Oct-13	TBATS(0.068, {0,0}, 0.966, {(12,2), (288,9), (2016,5), (105120,7)})
01-Jan-14	TBATS(0.11, {0,0}, 0.95, {(12,3), (288,7), (2016,6), (105120,7)})

III. In Table 17, we present the regression coefficients for the recursive forecasts in Section 4. The best fitting regression model varies in some iterations.

Table 17: Regression Coefficients in Recursive Modeling: New York Energy Load

Date	(Int.)	DayNr.	$T_{\max}$	$T_{\min}$	$b_{\max}$	$b_{\min}$	$T^2_{\max}$	$T^2_{\min}$	$b^2_{\max}$	$b^2_{\min}$	wkday	Sat	Mon	Tue	Wed	Thu	Fri
01-Oct-09	8770.7	-0.41	-104.64	-	30.05	-105.47	0.82	0.30	-	1.00	766.91	-	-	-	-	-	-
01-Jan-10	8842.9	-	-91.94	-	-	-88.14	0.73	0.23	0.25	0.88	-	96.59	800.09	850.15	823.12	745.44	695.69
01-Apr-10	8766.4	-	-78.66	-45.96	-	-48.19	0.63	0.64	0.19	0.53	-	94.64	785.68	843.00	820.02	747.62	696.87
01-Jul-10	9121.3	-	-108.50	-54.90	21.41	-43.19	0.87	0.77	-	0.43	-	97.27	789.28	856.18	839.25	765.12	714.76
01-Oct-10	9343.9	0.13	-103.08	-85.53	21.62	-35.22	0.83	1.13	-	0.34	-	103.49	816.55	885.37	881.35	809.97	752.41
01-Jan-11	9465.6	-	-96.24	-122.72	16.80	-	0.79	1.47	-	-	-	100.83	786.51	852.48	858.11	782.09	733.98
01-Apr-11	9308.0	-	-97.74	-111.28	15.68	-	0.80	1.37	-	-	-	93.77	754.94	829.47	837.84	760.43	708.18
01-Jul-11	9344.6	-	-101.01	-83.83	15.58	-29.12	0.83	1.11	-	0.30	-	106.28	755.04	831.19	850.55	778.11	722.81
01-Oct-11	9471.4	-	-108.43	-72.85	15.39	-38.51	0.90	0.99	-	0.41	-	100.43	763.39	834.59	862.67	792.14	736.66
01-Jan-12	9498.1	-	-106.75	-76.75	15.10	-37.25	0.89	1.01	-	0.41	-	110.14	756.10	827.80	854.39	784.76	738.64
01-Apr-12	9477.1	-	-109.02	-72.46	13.89	-35.36	0.91	0.97	-	0.40	-	115.72	748.38	820.70	841.12	781.34	737.75
01-Jul-12	9597.4	-0.07	-110.47	-78.24	14.15	-31.68	0.92	1.03	-	0.35	-	118.17	751.45	820.63	838.26	783.57	744.20
01-Oct-12	9731.7	-0.03	-103.97	-83.04	-	-27.65	0.88	1.07	0.14	0.31	788.78	113.40	-	-	-	-	-
01-Jan-13	9802.8	-0.05	-104.61	-83.44	-	-28.62	0.89	1.07	0.13	0.33	770.93	111.59	-	-	-	-	-
01-Apr-13	9618.2	-0.04	-108.95	-84.39	13.30	-28.00	0.92	1.08	-	0.33	765.88	118.20	-	-	-	-	-
01-Jul-13	9669.1	-0.05	-109.78	-80.17	13.78	-34.36	0.92	1.04	-	0.41	760.36	116.51	-	-	-	-	-
01-Oct-13	9728.3	-0.06	-110.44	-76.55	13.95	-41.08	0.92	1.00	-	0.49	766.97	113.83	-	-	-	-	-
01-Jan-14	9731.2	-0.06	-108.67	-83.45	14.17	-35.80	0.91	1.05	-	0.44	758.56	112.52	-	-	-	-	-