DETERMINING OPTIMAL ARCHITECTURE FOR DYNAMIC LINEAR MODELS IN TIME SERIES APPLICATIONS

Kathleen Mary Karlon

A Thesis Submitted to the
University of North Carolina Wilmington in Partial Fulfillment
Of the Requirements for the Degree of
Master of Science

Department of Mathematics and Statistics

University of North Carolina Wilmington

2006

Approved by

Advisory Committee

Accepted by

Dean, Graduate School

This thesis has been prepared in the style and format

consistent with the journal

American Mathematical Monthly.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

This work is focused on assessing the performance of one particular time series forecasting paradigm: Dynamic Linear Models (DLM). This research extends the M3 forecasting competition, a large-scale project to assess the efficacy of various forecasting methods and also that of the research done in [14]. This work provides insight into the performance of the DLM against the model architecture. Symmetric Mean Absolute Percentage Error and Linear Mixed Models are used to analyze the competition results, which showed that paradigm performance is dependent upon the class of time series. Furthermore, in some cases, the chosen DLM models from this work outperform optimal models from [14]. This work explores different DLM models and compares the results with previously chosen models to determine if the models from this work outperform other models.

ACKNOWLEDGMENTS

First and foremost, I'd like to thank my wonderful parents, Dot and George, without whom none of this would be possible. Thank you both for the continued support, encouragement and inspiration that you have given me throughout every moment of my life. Thank you to my siblings as well - Kim, Erik and Shane - for being such positive role models in so many aspects of my life. Many thanks to the rest of my family; each one of you has had a positive impact on my life.

Many thanks to the entire Mathematics and Statistics Faculty at UNCW. Most importantly, thanks to Dr. Boone and Dr. Ricanek for giving me this opportunity and agreeing to be my thesis advisors.

1 INTRODUCTION

Forecasting is an essential component in most company strategies. The ability to apply past experiences with the purpose of predicting future actions is a crucial part of effective business planning. The accuracy of such predictions is even more important to ensure the most successful outcome.

In the past, there have been numerous forecasting competitions in an attempt to compare the accuracy of various methods to determine which methods display optimal performance. This work familiarizes the reader with that of the M3 Forecasting Competition, a large-scale study to determine the forecast accuracy of various paradigms [9]. The M3 Competition consisted of 24 different paradigms, some of which were considered to be statistically sophisticated, such as Neural Networks and ARIMA models. One of the conclusions reached by the M3 Competition was that simpler models outperformed statistically sophisticated models. This work details a statistically sophisticated method that was overlooked by the M3 Competition.

The work contained herein describes results of analyses on data from the M3 Competition using a particular statistical paradigm for time series. The reader is introduced to the Dynamic Linear Model (DLM), as detailed in [13], which is a Bayesian method for time series analysis. Different DLM models are discussed, as well as the procedure for updating the models as new data is introduced.

The remainder of Section 1 introduces some basic forecasting concepts and examples, as well as the area of time series analysis and its role in forecasting. Also in Section 1, a brief introduction to the M3 Competition is given along with an overview of Bayesian Statistics. Section 2 presents the Dynamic Linear Model and how it is applied to time series forecasting. Section 3 introduces accuracy measures to be used for a comparison of the results, as well as the four candidate models to be tested. Section 4 gives an evaluation of the results obtained compared with those found in [14] and Section 5 concludes with a brief discussion of the conclusions drawn from this study.

1.1 Forecasting

> "I have seen the future and it is very much like the present, only longer."
>
> \- Kehlog Albran, *The Profit*

This is actually a concise description of statistical forecasting. Forecasters search for statistical properties of a time series that are constant in time - levels, trends, seasonal patterns, etc. The forecaster assumes that those properties will describe the future as well as the present [4].

Definition 1 *Forecasting is the act of calculating or estimating an event in advance.*

The purpose of forecasting is to be able to make statements (predictions) about the future based on past information, and to do so in a timely manner. A forecast is a set of probabilities attached to a set of future events. In order to understand a forecast, all one needs to do is interpret those two bits of information [5]. Moreover, forecasting is the process of analyzing current and historical data to determine future trends. Sometimes the forecaster has the ability to prevent undesirable events by forecasting the event, identifying the situations preceding the event and taking action to potentially lessen the impact of the event. Say, for example, based on past events a forecaster predicts a stock market crash, then stock investors have the chance to sell any stock that is owned instead of losing everything in the crash. The ability to predict such an event increases the chance to minimize the associated risk, otherwise known as the impact of the event. Forecasting is also beneficial if one is able to minimize loss from forecasting an event. For instance, forecasting the demand for a product enables one to control the inventory stock so that there are not too many products in a warehouse at any point in time, thus minimizing the loss of total money spent.

The appeal of forecasting lies in the fact that its application enables one to modify or control variables in order to alter, or prepare for, the future. The assumption is that the events responsible for creating the past will continue to operate similarly in the future. To forecast events, a forecaster must rely on information concerning events that have occurred in the past. The forecaster must identify a pattern (modeling and estimation) and then extrapolate the pattern into the future (forecasting). The main reason for forecast modeling is to provide efficient learning processes that will increase understanding and enable wise decisions [13].

Forecasting methods may be classified into two categories: Qualitative and Quantitative. Qualitative methods use experts' opinions to forecast future values of an event. Conversely, the quantitative methods use historical data to predict future values of a time series. Compared to qualitative methods, quantitative methods have the advantage of being supported by mathematical and statistical theory and can be reproduced by any forecaster [14]. The purpose of this work is to discuss, in detail, one particular quantitative method, the Dynamic Linear Model (DLM) and its efficacy against the performance results of [14].

1.2 Time Series Analysis

Time series forecasting is an area in which quantitative methods are widely applied. The assumption in a time series analysis is that successive values in the data series represent consecutive measurements taken at equally spaced time intervals. Referred to as a "time series," the historical values used are spaced equally over time to represent daily, monthly, quarterly and/or yearly data. Denoted by $\{Y_t : t \in \mathbb{Z}\}$, a time series is a sequence of observations, assumed to be a combination of a pattern and some random error. To fit a model to the data, the idea is to separate the pattern from the error. The pattern is the systematic component of the time series involving the trend (whether a long-term increase or decrease) and the seasonality (the change caused by seasonal factors such as fluctuations in use and/or demand).

Before deducing any statistical information from the data, stationarity of the process is often assumed. In the case of the Gaussian process, stationarity is the assumption that, through time, the process has a constant mean, variance and autocorrelation, the correlation between values of the same random process at different times. A process, $\{Y_t\}$, will only be stable if the parameters are within a certain range; otherwise, past circumstances would continue to accumulate, the values of $\{Y_t\}$ would continue to move towards infinity and the series would not be stationary. This concept of stationarity can be observed in the specification of the parameters in Section 2.

The main goals of time series analysis are to obtain an understanding of the underlying relationship of the observed data and develop a model that represents the relationship of the data in order to proceed on to forecasting. To explain these underlying relationships, there are three main time series components used to describe patterns of a data series: Level, Trend and Seasonality.

Trend is the long-term increase or decrease of a data series. The trend of a series may be a line or curve in an upward or downward direction, but not both. The level is the mean height of the data series.

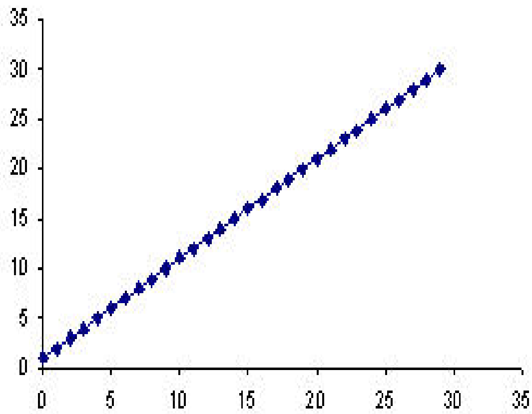

Figure 1: Trend Pattern

The simulated data illustrated in Figure 1 is at a constant incline with no repetitive nature to the pattern, thus, the figure demonstrates an upward trend.

Seasonality is the change in the data series caused by the repetition of seasonal factors. The seasonal component is similar to that of the trend component except for the fact that it repeats itself in systematic intervals over time. The simulated data in Figure 2 demonstrates the seasonality factor.

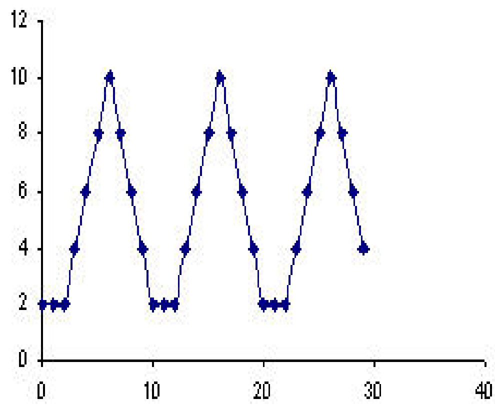

Figure 2: Seasonal Pattern

For example, the number of sales for retail stores increases significantly during the holiday season, thus, a seasonality pattern would be displayed in the sales data.

In many instances, the two components may also coexist in a pattern. Example 1 exemplifies this particular coexistence.

Example 1 *Sales of Product A*

Let t be the number of months after December 1992

y_t *= monthly sales totals of Product A*

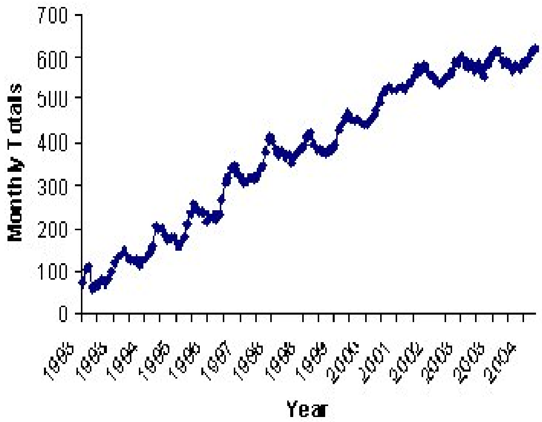

Seasonality is displayed in Example 1, as well as an increasing trend which explains the popularity growth of Product A.

Combining these components into one model results in the basic time series model which is represented as:

$$Y_t = \mu_t + \gamma_t + \varepsilon_t, \tag{1}$$

where μ_t,γ_t and ε_t represent the trend, seasonal and error components, respectively. The trend, μ_t, is modeled as a component with slowly fluctuating level and slope; the trend models the underlying direction. The seasonal component models any regular fluctuations between corresponding periods.

1.3 M3 Forecasting Competition

The first effort to compare a large number of quantitative forecasting methods across multiple time series was done in 1979 by Makridakis and Hibon [6]. The M3 Competition was an effort to respond to the criticisms of [6] and to incorporate the suggestions of the various commentators for improvements.

The M3 Competition was an empirical study that compared the performance of a large number of major time series methods using recognized experts who provided forecasts using their method of expertise [11]. Once the forecasts from each expert were obtained, they were evaluated and compared with those of the other experts as well as with some simple methods used as benchmarks [11]. Forecasting competitions, such as these, promised impartiality and expert knowledge.

The M3 Competition utilized a common database with 3,003 time series data sets of which most were business and economic. The 3,003 series of the competition, illustrated in Table 1, were selected on a quota basis to include various types of time series data (micro, industry, macro, etc.) and different time intervals between successive observations (yearly, quarterly, etc.) [11].

Time Intervals Between Successive Observations	Type of Time Series Data						
	Micro	Industry	Macro	Finance	Demographic	Other	Total
Yearly	146	102	83	58	245	11	645
Quarterly	204	83	336	76	57		756
Monthly	474	334	312	145	111	52	1428
Other	4			29		141	174
Total	828	519	731	308	413	204	3003

Table 1: The Classification of the 3003 Time Series

Seasonal Type	Data Detail					
	Total Series	Min Length	Median Length	Max Length	Average Length	Forecast Horizon
Yearly	645	14	19	41	22	6
Quarterly	756	16	44	64	41	8
Monthly	1428	48	115	126	99	18
Other	174	60	63	96	69	8

Table 2: Data Detail of all Seasonal Catalogs

Furthermore, the detail of the seasonal catalogs is shown in Table 2. The forecast horizon represents how far ahead that specific time series should be forecasted. A short-term forecast horizon usually has a length of 1-6 data periods; a medium-term forecast horizon has length around 3-12 data periods and a long-term forecast horizon has length of > 10 data periods. Based on previous competitions, short-term forecasting has been shown to be more accurate than long-term forecasting. For example in Table 2, the yearly time series is considered a short-term forecast when compared to a monthly forecast. Thus, the yearly results would most likely be more accurate than the monthly results. In Table 2, the different lengths represent the minimum number of observations to be used for each type of data.

The M3 competition reached four main conclusions which are as follows [9]:

1. Statistically sophisticated or complex methods do not necessarily produce more accurate forecasts than simpler ones.

2. The rankings of the performance of the various methods vary according to the accuracy measure being used.

3. The accuracy of the combination of various methods outperforms, on average, the specific methods being combined and does well in comparison with other methods.

4. The performance of the various methods depends on the length of the forecasting horizon.

The results of the M3 Competition confirmed the original conclusions of previous competitions [9]. Basically, the idea of parsimony, the simplicity of the structure of a model, was the main conclusion reached.

In order to compare forecast results, the M3 Competition made use of one specific accuracy measure known as the symmetric mean absolute percentage error (SMAPE). To be consistent, this work made use of the same accuracy measure. The SMAPE is defined as:

$$\frac{1}{n}\sum_{k=1}^{n}\frac{|X_k - F_k|}{(X_k + F_k)/2} \times 100, \tag{2}$$

where X_k is the actual value, F_k is the forecast value and n is the number of time series. The SMAPE is described as the average across all forecasts made for a given horizon in a specific type of time series data [14]. After computing the SMAPE at each horizon, the linear mixed model, which will be discussed later, was used to identify any differences among the forecasting results.

In an attempt to be comprehensive, the M3 Competition tested a large number of forecasting paradigms. However, due to the fact that there was a limited amount of resources and the experiment relied on external researchers to provide analyses of the series using independently chosen paradigms, some paradigms were omitted [14]. Many researchers chose to use commercially available implementations of various paradigms instead of standard textbook methods [14]. As a result, Dynamic Linear Models were omitted from the competition completely. For this reason, this work details the DLM and compares the results with those from Zhai [14]. The work done on the DLM in this research is an extension of the research conducted there. In his research, Zhai tested three different sophisticated paradigms including the Dynamic Linear Model. One of the results from his work was that the First-Order DLM is the recommended algorithm when dealing with yearly data due to the fact that the time series is short and has no seasonal pattern. This work goes deeper to test all possible DLM algorithms on each time series, determine the optimal DLM architecture for each time series and then determine whether or not the resulting DLM architecture outperforms the paradigms suggested in the M3 Competition and the research done by Zhai. First, though, to understand the idea surrounding the Dynamic Linear Model, the reader must be introduced to Bayesian Statistics.

1.4 Bayesian Statistics

Bayesian statistics is founded on the fundamental premise that all uncertainties should be represented and measured by probabilities [13]. In general, conditioning on what is known, in order to make statements about what is not known, is the basic idea behind Bayesian methodologies. These methods employ the use of prior distributions rather than parameters for the purpose of parameter estimation and hypothesis testing. In order to better understand the meaning behind this idea, the reader must be familiar with some frequently used terms.

Definition 2 *The probability that model A is true before any data are observed is known as the prior probability of A, $P(A)$.*

Definition 3 *The probability that model A is true after the observed data B have been taken into account is known as the posterior probability of A given B, $P(A|B)$.*

Definition 4 *The conditional probability of the data B, given a particular model A, is known as the likelihood of B given A, $P(B|A)$.*

With this being said, the discrete case of Bayes' theorem is as follows [12]. For two quantities A and B for which probabilistic beliefs are given, Bayes' Theorem states

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{3}$$

Furthermore, the continuous case of Bayes' theorem is represented in the following equation.

$$f(x|y) = \frac{f(y|x)f(x)}{\int_{-\infty}^{\infty} f(y|x)f(x)dx}, \tag{4}$$

where $f(x|y)$ denotes the posterior probability density function of X given $Y = y$; $f(y|x) = L(x|y)$ is the likelihood function of X given $Y = y$ and $f(x)$ is the prior density function of X.

Bayes' theorem basically explains how to update or revise beliefs after taking new evidence into account. The posterior is proportional to the likelihood times the prior probability. In general, this statement can be represented as

Posterior $\propto$ Observed likelihood * Prior.

Forecasts are conditional probability statements, the conditioning being on the existing state of knowledge [12]. Bayesian methodology offers a comprehensive way of routine learning that is not dependent on any particular assumptions, which is the foundation of the Dynamic Linear Model [14].

Throughout the literature, many techniques have been implemented to perform time series forecasting. The next section will familiarize readers with the Dynamic Linear Model.

2 DYNAMIC LINEAR MODEL

The Dynamic Linear Model is the Bayesian forecasting paradigm used throughout this work. The term dynamic refers to changes in processes due to the passage of time. Generally, the DLM can be defined by:

$$\{F_t, G_t, V_t, W_t\},$$

where:

$$\begin{aligned} F_t &\quad \text{is a known } (n \times r) \text{ matrix,} \\ G_t &\quad \text{is a known } (n \times n) \text{ matrix,} \\ V_t &\quad \text{is a known } (r \times r) \text{ covariance matrix,} \\ W_t &\quad \text{is a known } (n \times n) \text{ covariance matrix.} \end{aligned}$$

The representation of the DLM is as follows:

$$\begin{aligned} \text{Observation equation:} \quad & Y_t = F_t{}'\theta_t + \nu_t, & \nu_t \sim N[0, V_t], \\ \text{System equation:} \quad & \theta_t = G_t\theta_{t-1} + \omega_t, & \omega_t \sim N[0, W_t], \\ \text{Initial information:} \quad & (\theta_0 | D_0) \sim N[m_0, C_0], & \end{aligned}$$

where Y_t is the observation series at time t, $\mu_t = F_t{}'\theta_t$ is the mean response or level, θ_t is the state or system vector, ν_t is the observational error, ω_t is the evolution error, D_t is all information about time t, m_0 is the point estimate of this level and C_0 measures the associated uncertainty.

Based on the observation equation, given θ_t, Y_t is independent of all other observations and parameter values; likewise, given the present, the future is independent of the past.

2.1 First-Order Polynomial Model

The most commonly used DLM is the first-order polynomial model. Mostly used for short-term forecasting with no trend, this simple model is defined as follows [13]:

$$\{1, 1, V_t, W_t\},$$

such that,

$$\begin{aligned}
&\text{Observation equation:} \qquad && Y_t = \mu_t + \nu_t, && \nu_t \sim N[0, V_t], \\
&\text{System equation:} && \mu_t = \mu_{t-1} + \omega_t, && \omega_t \sim N[0, W_t], \\
&\text{Initial information:} && (\mu_0 | D_0) \sim N[m_0, C_0]. &&
\end{aligned}$$

At any time t, the only new information becoming available is the observed value Y_t so that $D_t = \{Y_t, D_{t-1}\}$.

As stated earlier, this model is mainly used for short-term forecasting. For example, in modeling demand for a particular product, μ_t represents true underlying demand at time t with ν_t describing random fluctuation, which arises in the actual placement of customer orders, about this level [13]. Typically, the underlying demand μ_t is considered roughly constant. Significant changes over longer periods of time are expected, but the zero-mean and independent nature of the ω_t series imply that the forecaster does not wish to anticipate the form of this longer term variation [13].

2.2 Second-Order Polynomial Model

Second-order polynomial models are useful for describing trends in time series data. Sometimes referred to as the linear growth model, the second-order polynomial model has also proven sufficient for short-term forecasting. This model is characterized as follows:

$$\left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, V_t, W_t \right\},$$

such that,

$$\begin{aligned} \text{Observation equation:} \quad & Y_t = \theta_{t,1} + \nu_t, & \nu_t \sim N[0, V_t], \\ \text{System equation:} \quad & \theta_{t,1} = \theta_{t-1,1} + \theta_{t-1,2} + \omega_{t,1}, & \omega_t \sim N[0, W_t], \\ & \theta_{t,2} = \theta_{t-1,2} + \omega_{t,2}, & \\ \text{Initial information:} \quad & (\theta_{t-1}|D_{t-1}) \sim N[m_{t-1}, C_{t-1}], & \end{aligned}$$

where:

$$\theta_t = \begin{pmatrix} \theta_{t1} \\ \theta_{t2} \end{pmatrix} = \begin{pmatrix} \mu_t \\ \beta_t \end{pmatrix},$$

$$\omega_t = (\omega_{t1}, \omega_{t2})',$$

$$m_{t-1} = \begin{pmatrix} m_{t-1} \\ b_{t-1} \end{pmatrix},$$

$$C_{t-1} = \begin{pmatrix} C_{t-1,1} & C_{t-1,3} \\ C_{t-1,3} & C_{t-1,2} \end{pmatrix}.$$

As usual, μ_t is the series level and now β_t represents incremental growth [13].

As an extension of the second-order polynomial model, the n^{th} order polynomial model may be generalized as follows:

$$\left\{ \begin{pmatrix} F_{1,t} \\ \vdots \\ F_{n,t} \end{pmatrix}, \begin{pmatrix} G_{1,1} & \dots & G_{1,t} \\ \vdots & \ddots & \vdots \\ G_{n,1} & \dots & G_{n,t} \end{pmatrix}, V_t, W_t \right\},$$

such that,

$$\begin{aligned}\text{Observation equation:}\quad & Y_t = \theta_{t,1} + \nu_t \\ \text{System equation:}\quad & \theta_{t,1} = \theta_{t-1,1} + \theta_{t-1,2} + \ldots + \theta_{t-1,n} + \omega_{t,1} \\ & \theta_{t,2} = \theta_{t-1,2} + \theta_{t-1,3} + \ldots + \theta_{t-1,n} + \omega_{t,2} \\ & \vdots \\ & \theta_{t,n} = \theta_{t-1,n} + \omega_{t,n}\end{aligned}$$

2.3 Seasonal Models

Seasonality is a term used to describe the cyclical or periodic fluctuations of a time series. For example, every year during the holiday season, sales for a particular retail store increase significantly. This is considered to be a seasonal factor that should be taken into account when forecasting. It is important to consider these seasonal factors when modeling certain time series data due to the fact that they may have important implications. To construct a seasonal model, superimposing the seasonal characteristics with that of the previously stated polynomial models is sufficient to observe this cyclical behavior. For the purpose of this work, seasonal characteristics were applied to the quarterly and monthly data. The following matrix representations illustrate a general seasonal pattern.

$$F_t = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}, G_t = \begin{pmatrix} 0 & 1 & 0 & \ldots & 0 & 0 \\ 0 & 0 & 1 & \ldots & 0 & 0 \\ \vdots & & & \ddots & & \vdots \\ 0 & 0 & 0 & \ldots & 0 & 1 \\ 1 & 0 & 0 & \ldots & 0 & 0 \end{pmatrix}$$

So, if F_t is an $n \times 1$ matrix, then G_t is an $n \times n$ matrix, with n representing the number of cycles within a time series. The quarterly matrices, for example, represent the four quarters in a year, thus, F_t and G_t are 4×1 and 4×4 matrices, respectively.

Example 2 *Quarterly Seasonal Representation*

$$F_t = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, G_t = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

Superimposition, to be applied when constructing seasonal models, is merging the seasonal characteristics with that of the polynomial models. For instance, superimposing the seasonal characteristics with that of the first-order polynomial would result in a model accounting for seasonality and no trend; likewise, superimposing the seasonal characteristics with that of the second-order polynomial would result in a model accounting for seasonality and trend. Once again, this idea applied to the quarterly data would yield the following matrices.

- No trend, Seasonal

$$F_t = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, G_t = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \mathbf{0} & \mathbf{1} & \mathbf{0} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{0} & \mathbf{1} & \mathbf{0} \\ 0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{1} \\ 0 & \mathbf{1} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix}$$

- Trend, Seasonal

$$F_t = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, G_t = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{0} & \mathbf{1} & \mathbf{0} & \mathbf{0} \\ 0 & 0 & \mathbf{0} & \mathbf{0} & \mathbf{1} & \mathbf{0} \\ 0 & 0 & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{1} \\ 0 & 0 & \mathbf{1} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix}$$

These seasonal matrices are designed to represent the seasonal cycle. For seasonal time series, the matrices consist of submatrices, known as the rotation matrices, which are in bold. These rotation matrices account for the seasonality in a specific time series.

2.4 Updating Process

At the current time period, the forecaster has an estimate of the process, consisting of the initial information and the observations up to that time period. At this point, the observation and system equations of the DLM will yield an estimate of the next observation, which is known as the forecast. This updating process continues as new observations are introduced.

The unknown parameters, which consist of all the information up to that time period, are expressed as probability distributions. These distributions, defined earlier and now applied to the updating process outlined in Table 3, are the prior distribution of θ_t, denoted $\theta_t|D_{t-1}$, the forecast distribution is denoted as $Y_t|D_{t-1}$ and the posterior distribution of θ_t as $\theta_t|D_t$. Table 3 displays this updating algorithm of the DLM.

Univariate DLM: unknown, constant variance $V = \phi^{-1}$		
Observation:	$Y_t = F_t'\theta_t + \nu_t,$	$\nu_t \sim N[0, V_t],$
System:	$\theta_t = G_t\theta_{t-1} + \omega_t,$	$\omega_t \sim N[0, W_t]$
Information:	$(\theta_{t-1}\|D_{t-1}) \sim N[m_{t-1}, C_{t-1}],$ $(\phi\|D_{t-1}) \sim G[\frac{n_{t-1}}{2}, \frac{n_{t-1}S_{t-1}}{2}].$	
Forecast:	$(Y_t\|D_{t-1}) \sim N[f_t, Q_t],$ $(\theta_t\|D_{t-1}) \sim N[a_t, R_t],$	
where	$R_t = G_tC_{t-1}G_t{}' + W_t,$ $Q_t = F_t{}'R_tF_t + S_{t-1},$	$a_t = G_tm_{t-1},$ $f_t = F_t{}'a_t.$
Updating Recurrence Relationships		
	$(\phi\|D_t) \sim G[\frac{n_t}{2}, \frac{n_tS_t}{2}],$ $(\theta_t\|D_t) \sim N[m_t, C_t],$	
with	$e_t = Y_t - f_t$ and $A_t = R_tF_t/Q_t,$ $n_t = n_{t-1} + 1,$ $S_t = S_{t-1} + \frac{S_{t-1}}{n_t}(\frac{e_t{}^2}{Q_t} - 1),$ $m_t = a_t + A_te_t,$ $C_t = \frac{S_t}{S_{t-1}}(R_t - A_tA_t{}'Q_t).$	
Forecast Distributions $k \geq 1$		
	$(\theta_{t+k}\|D_t) \sim N[a_t(k), R_t(k)],$ $(Y_{t+k}\|D_t) \sim N[f_t(k), Q_t(k)].$	

Table 3: Univariate DLM [13]

3 SELECTION CRITERIA

Once the forecaster has generated results for the particular DLM model, the next step is to compare the results with that of the M3 Competition. To do this comparison, a forecast accuracy measure must be defined.

3.1 Symmetric Mean Absolute Percentage Error

There are many different accuracy measures considered in empirical studies. One accuracy measure used in this work is the symmetric mean absolute percentage error (SMAPE). The SMAPE is also the accuracy measure used throughout the M3 Competition and is recommended for use on the M3 data due to the presence of zeroes. The SMAPE, previously defined as,

$$\frac{1}{n}\sum_{k=1}^{n}\frac{|X_k - F_k|}{(X_k + F_k)/2} \times 100, \tag{5}$$

is considered an improvement of the mean absolute percentage error (MAPE). The MAPE which is defined by

$$\frac{1}{n}\sum_{k=1}^{n}\frac{X_k - F_k}{X_k}, \tag{6}$$

has the disadvantage that there is a heavier penalty on positive errors than on negative; this led to the use of the "symmetric" measures [7]. The SMAPE is constrained to be between 0 and 200; so, the error is symmetric with respect to the scale of the errors. This idea is illustrated in Example 3.

Example 3 *Establishing Symmetry of Forecast Errors*

Given: $APE = \sum |\frac{x_t - f_t}{x_t}|$ *and* $SAPE = \frac{200|x_t - f_t|}{(x_t + f_t)}$

Case 1: Let $x = 100$ *and* $f = 150 \Rightarrow APE = 50\%$.

Case 2: Let $x = 150$ *and* $f = 100 \Rightarrow APE = 33\%$.

However, the $SAPE = 40\%$ *in either case, thus eliminating asymmetry.*

In this work, the SMAPE is used for an out-of-sample validation. As with most accuracy measures, a lower SMAPE value indicates a better fit to the data. The SMAPE as defined in Equation 5 will be displayed in the tables.

3.2 Mean Square Error

The next accuracy measure that needs to be defined is that of the mean square error (MSE). The MSE, written as,

$$\sum_{k=1}^{n} \frac{(X_k - F_k)^2}{n} \tag{7}$$

where n = *number of forecasts,*

is the sum of the squared forecast errors for each of the observations divided by the number of observations. In this work, the MSE is used for the purpose of an in-sample model selection.

Minimizing the MSE is, essentially, minimizing the errors and the variance of the errors; the estimator that minimizes the MSE is the mean. Thus, the objective is to find the model with the smallest MSE; the smallest MSE being the best fit.

3.3 Candidate Models

For the purpose of this work, the categories have been separated into four different models. The use of MATLAB was to design a study for each different time series and account for all possible models pertaining to that specific time series. The candidate models are as follows:

- Model 1: No Trend, No Seasonality

$$F_t = \begin{pmatrix} 1 \end{pmatrix}, G_t = \begin{pmatrix} 1 \end{pmatrix}$$

- Model 2: Trend, No Seasonality

$$F_t = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, G_t = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

- Model 3: No Trend, Seasonality (if applicable)

$$F_t = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, G_t = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix}$$

- Model 4: Trend, Seasonality (if applicable)

$$F_t = \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, G_t = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{pmatrix}$$

Models 3 and 4 are specifically designed for the quarterly data; the monthly data models are designed similarly except that the monthly models have to account for the twelve months in a year. The monthly and quarterly studies are written so that each study takes all four models into account in order to ensure that any seasonal and trend patterns are found. For yearly and other data, there are no obvious reasons for seasonal patterns; thus, only Models 1 and 2 are included in the studies for the yearly and other data.

The simulated data illustrated in Figure 3 represents the four candidate models displaying the patterns found in each model.

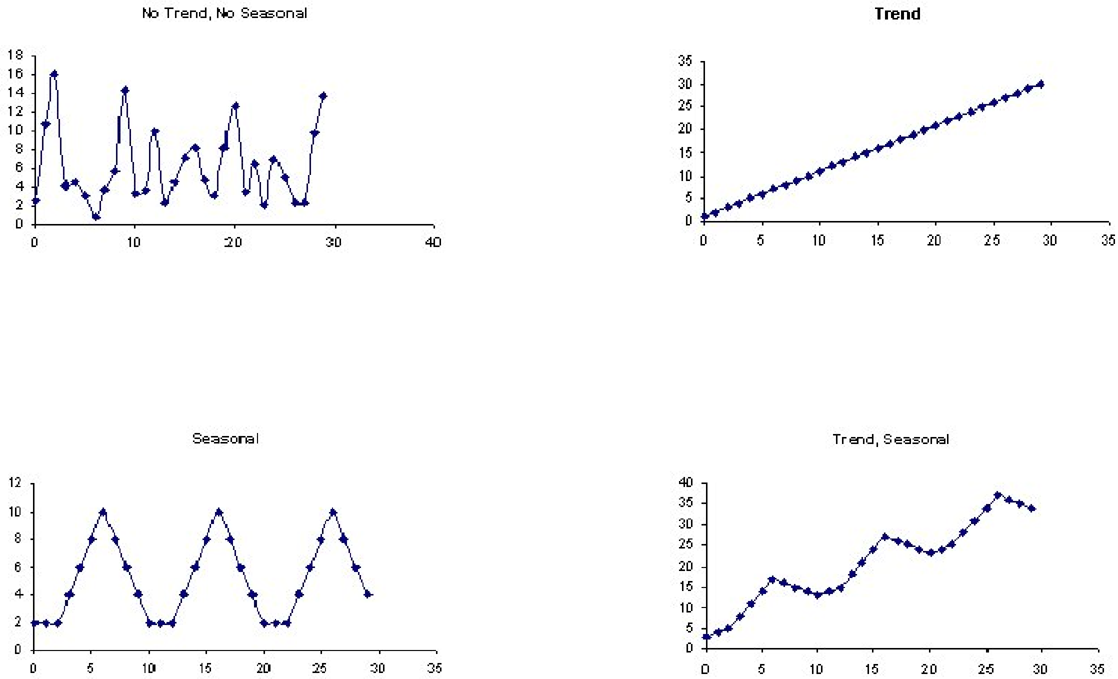

Figure 3: Model Patterns

Since each study is designed to account for the necessary combinations of trend and seasonal patterns, the next step is to select the best fit model. The method selection protocol involves calculating the SMAPE and MSE for each model of each time series. Each study automatically determines the lowest MSE value for the in-sample data and considers this model to be the best fit; each study also determines the lowest SMAPE for the out-of-sample data and considers this model to be the best fit. However, the two accuracy measures do not always choose the same best fit model.

3.4 Linear Mixed Model

After each study determines the best fit model for each set of time series data, the next step is to figure out if the chosen model differs in a statistically significant manner from the other models. In order to test the significance of the difference, this work utilizes a statistical tool known as the Linear Mixed Model (LMM). The linear mixed model is defined by:

$$y_{ij} = \mu + a_i + \beta_j + c_{ij}, \tag{8}$$

where y_{ij} is the accuracy measure associated with the i^{th} paradigm using the j^{th} time series, μ is the overall mean, a_i is the fixed effect of the i^{th} paradigm, β_j is the random effect of the j^{th} time series where $\beta_j \sim N(0, \sigma_\beta^2)$ and c_{ij} is the random error term for the interaction where $c_{ij} \sim N(0, \sigma^2)$ [14]. The linear mixed model, detailed in [10], is different from other tools in which it is a combination of fixed and random effects such that, fixed effects are used for modeling the mean of y while random effects govern the variance-covariance structure of y [10].

The linear mixed model tests the null hypothesis that all of the paradigms produce the same average SMAPE in every horizon; conversely, the alternative hypothesis is that at least one average SMAPE of one paradigm in a specific horizon is different from the others [14]. The mixed ANOVA, designed to test the differences between the means of two or more groups, enables one to determine if the paradigm differences in each average SMAPE at each horizon is statistically significant at the 0.05 confidence level. The next section is a detailed description of the ANOVA results for each paradigm performance.

4 RESULTS

Each type of time series is one study in itself. Each one is designed following the same template, but differ in the parameter specifications and models used. This section details the results of the studies conducted for this work.

While in certain instances results from previous studies such as [14] could not be duplicated, others showed to be an improvement from the results found in [14]. The tables listed in this section are for four types of time series data: Yearly, Quarterly, Monthly and Other. The tables include every horizon specific to that time series along with the average SMAPE at every horizon for the different types of models.

One conclusion that can be drawn from the results of this study is that the performance of a model depends on the length of the forecast horizon. For example, as illustrated in Table 4 the average SMAPE increases as the forecast horizon increases. Another notable conclusion is that the simple DLM tends to outperform the more complex DLM. For instance, the first-order polynomial model in Table 5 outperforms the other three models tested in the quarterly study. The remainder of this section details results specific to each time series study.

4.1 Yearly

As previously stated, the yearly study consisted of two models: the polynomial models. The yearly study began by initializing several parameters with the value of one, which was later to be adjusted in an attempt to optimize the MSE and SMAPE. The mean was initialized by calculating the mean of the first four periods. For the second-order polynomial model, the mean was a vector consisting of the mean of the first four periods and the trend, which was derived by subtracting the mean of the first two periods from the mean of the second two periods and dividing by two. An optimization loop was included to adjust certain parameters and determine the values that resulted in the lowest MSE and SMAPE. Table 4 displays the results of the yearly study; the table lists the average SMAPE at each horizon for each model.

		Forecast Horizon					
		1	2	3	4	5	6
No Trend, No Season	SMAPE	**14.86**	**18.65**	**20.38**	**24.40**	**27.57**	**30.34**
	StdDev	19.55	21.99	23.04	25.86	26.04	28.82
Trend, No Season	SMAPE	19.82	22.15	24.41	28.01	30.18	32.39
	StdDev	23.86	27.43	29.84	33.29	34.05	35.52

Table 4: Yearly DLM Results

Although it can be observed in Table 4, after performing the linear mixed model on the two different models, it was confirmed that the results for the first-order polynomial model were significantly different from the results for the second-order polynomial model for each horizon at the 0.05 confidence level. This is shown in Table 4; a SMAPE value in bold indicates a significant difference from other models at the corresponding horizon. For example, for Model 1 the SMAPE value at horizon 1 is significantly different from the SMAPE value for Model 2 at horizon 1. Thus, the model of choice for the yearly data from this study is the first-order polynomial model; however, this model does not perform better than that of the chosen DLM model from [14]. This yearly study was unable to duplicate the results found in [14], but did outperform all of the artificial neural networks that were tested. Furthermore, the second-order polynomial model from this work outperformed the one tested in [14], as well as one of the artificial neural networks. Lastly, the M3 Competition conclusion that the performance of the method depends on the length of the forecast horizon is a result of this work as well. The further ahead the forecast, the poorer the performance of the DLM.

4.2 Quarterly

The quarterly study involved all four paradigms for performance testing. This study initialized the mean by calculating the mean of the first eight periods. The trend was similar except that the calculation consisted of the mean of the first four periods, subtracted from the mean of the next four periods and divided by four. If modeling trend, the mean was a vector consisting of the mean and the calculated trend. There was also seasonality to account for in the quarterly study. This was done by subtracting the value of a certain period from the value of the corresponding period and dividing by four. For example, the mean is calculated by computing the mean for the first eight periods; thus, period 2 corresponds to period 6. All four models were tested on the quarterly data and those results are listed in Table 5.

		Forecast Horizon							
		1	2	3	4	5	6	7	8
No Trend, No Season	SMAPE	**13.93**	**13.43**	**14.51**	**14.45**	**15.36**	**15.53**	**17.21**	**18.31**
	StdDev	19.85	19.58	20.83	19.72	20.93	20.11	21.74	21.74
Trend, No Season	SMAPE	22.30	21.90	22.29	23.59	25.76	25.79	26.65	28.51
	StdDev	37.59	37.54	37.49	37.64	38.98	38.97	39.69	40.58
No Trend, Season	SMAPE	18.12	15.50	16.67	18.13	20.56	18.04	20.09	21.98
	StdDev	31.51	26.59	25.79	26.77	31.96	27.05	26.51	28.57
Trend, Season	SMAPE	19.74	20.33	20.53	21.43	22.24	22.56	23.21	26.03
	StdDev	35.65	37.22	35.76	35.46	36.08	36.35	36.10	37.32

Table 5: Quarterly DLM Results

The results displayed in [14] are results of two seasonal models of the DLM. However, this study included nonseasonal models, one of which was determined to be the model of choice for this data thus far. Compared to the other three models, the first-order DLM resulted in a lower average SMAPE at each horizon; thus, the model of choice for this study is the first-order DLM. After comparing these results with that of the model of choice from [14], the Linear Mixed Model determined that there were significant differences between the two models. These differences are displayed in Table 5; a SMAPE value in bold indicates a significant difference at the 0.05 confidence level. As shown in Table 5, the two models were significantly different at every horizon.

All of the DLM models from this work outperformed the seasonal DLM models used in [14]. This result is surprising because the quarterly data often demonstrates seasonality, so it would seem that a seasonal model would display the optimal performance. However, this work found the first-order polynomial model to be the model of choice for the quarterly time series.

The poor performance of the seasonal models may be attributed to the fact that any major jumps in the data caused the model to overcompensate for such jumps, making it more difficult to continue to forecast the data. This idea was illustrated after calculating the median SMAPE across each horizon. For example, in Model 4 of the quarterly study, the median SMAPE ranged from 7.84 to 15.16 for horizons 1 through 8 respectively. Comparing the median SMAPE with that of the mean SMAPE of this particular model resulted in a large discrepancy. This is evidence that certain jumps in the data had a huge influence on the mean, thus resulting in a poor DLM performance.

4.3 Monthly

As with the quarterly study, the monthly study included all four of the DLM paradigms for performance testing. The mean was initialized by calculating the mean of the first 24 periods and the trend and seasonality were accounted for using the same technique employed in the quarterly study. Table 6 illustrates the results of the monthly study.

		Forecast Horizon					
		1	2	3	4	5	6
No Trend, No Season	SMAPE	**13.17**	**13.93**	**14.24**	**14.63**	**15.01**	**15.24**
	StdDev	19.27	19.52	18.34	19.74	19.27	19.57
Trend, No Season	SMAPE	16.59	16.63	17.55	19.23	19.90	19.66
	StdDev	24.96	25.40	26.12	27.77	28.60	28.94
No Trend, Season	SMAPE	27.15	27.44	25.56	26.33	27.29	31.36
	StdDev	36.03	38.78	35.26	37.51	37.82	45.46
Trend, Season	SMAPE	23.43	26.12	24.78	23.42	24.94	29.75
	StdDev	32.44	35.98	35.41	33.05	37.45	43.52
		Forecast Horizon					
		7	8	9	10	11	12
No Trend, No Season	SMAPE	15.16	**16.72**	17.42	16.77	20.00	**18.35**
	StdDev	19.08	21.64	23.20	21.60	25.69	23.24
Trend, No Season	SMAPE	21.12	21.30	20.71	20.58	22.20	24.21
	StdDev	29.84	30.24	30.98	31.10	32.87	33.95
No Trend, Season	SMAPE	33.28	26.99	27.52	31.21	30.92	33.08
	StdDev	48.80	37.22	36.90	46.92	42.03	45.80
Trend, Season	SMAPE	32.35	26.07	26.48	31.01	27.85	32.10
	StdDev	47.23	37.53	39.95	46.52	40.34	46.58
		Forecast Horizon					
		13	14	15	16	17	18
No Trend, No Season	SMAPE	20.05	17.80	**17.83**	18.34	18.31	**17.79**
	StdDev	25.41	22.64	22.35	21.18	20.25	21.34
Trend, No Season	SMAPE	25.92	26.64	29.18	30.88	31.93	31.86
	StdDev	35.15	35.80	37.99	39.98	41.24	41.47
No Trend, Season	SMAPE	31.38	30.03	27.98	27.67	28.66	34.07
	StdDev	38.87	39.59	36.45	36.41	37.55	46.26
Trend, Season	SMAPE	28.30	30.26	29.59	26.16	29.55	35.30
	StdDev	38.77	40.98	41.40	36.62	42.73	49.22

Table 6: Monthly DLM Results

Compared to the other three models in this study, the first-order DLM resulted in a lower average SMAPE at every horizon; thus, the model of choice for this study is the first-order DLM. The Linear Mixed Model was performed on specific horizons to compare these results with that of the chosen model from [14]. The Linear Mixed Model concluded that there were significant differences between the two models. Shown in Table 6, the significant differences are illustrated by a bolded SMAPE value. Once again, the first-order polynomial model outperforms the other three models at every horizon. Only the seasonal DLM models were tested in [14] on the monthly data and, in turn, did not perform as well as the model of choice found in this study. The first-order DLM from this work also outperformed several of the artificial neural networks tested in [14].

Once again, it would seem that seasonal models would display the optimal performance for the monthly data, but this was not the case. After calculating the median at each horizon, the resulting median for Model 4 ranged from 11.60 to 14.66 for horizons 1 through 18 respectively. Again, this shows that the major jumps in the data to be modeled resulted in a poor performance of the seasonal models. The first-order polynomial model tends to perform better simply because it is not as sensitive to such drastic changes in the data.

Another observation to note is that, while the yearly forecast appears to worsen as the horizon increases, the monthly forecast appears to perform adequately as the horizon increases. For example, in Model 1 of Table 4, the average SMAPE continues to increase as the horizon increases, ranging from 14.86 to 30.34 across six horizons; while in Model 1 of Table 6, the average SMAPE ranges from 13.17 to 17.79 across 18 horizons. This is attributed to the fact that there was more past information used in the quarterly study.

4.4 Other

Since the other data does not have a specific pattern classification, the study for the other data is similar to that of the yearly study such that it only includes the nonseasonal models. Both of the DLM models perform very well on the other data, which is illustrated in Table 7.

		Forecast Horizon							
		1	2	3	4	5	6	7	8
No Trend, No Season	SMAPE	13.44	14.26	15.00	16.07	16.55	17.18	18.10	18.90
	StdDev	11.32	12.43	13.03	13.76	13.14	13.91	14.23	14.96
Trend, No Season	SMAPE	4.01	4.97	5.76	6.63	6.86	7.35	7.80	8.91
	StdDev	5.06	8.23	9.27	10.33	8.70	9.75	10.45	11.79

Table 7: Other DLM Results

Based on the results from this table, the second-order polynomial model is the best fit for the other data. The results from [14] show that the first-order polynomial model was the model of choice for the other data. While this work was unable to duplicate those results, it did, however, improve the results of the second-order DLM.

5 CONCLUSION

This work, an extension of the research conducted in the M3 Competition and in [14], assessed the performance of a time series forecasting paradigm known as the Dynamic Linear Model (DLM). The DLM originated from Bayesian Methodologies, which are based on a comprehensive way of routine learning. This work introduced four candidate models of the DLM and applied specific models to different time series data.

The goal of this work was to explore the DLM and determine if it outperformed the chosen DLM models in [14]. The previous section details the results found for each model in each time series, but one main conclusion that can be drawn from each study is that the performance of the model depends on the length of the forecast horizon. This conclusion was one of the four conclusions reached in the M3 Competition and this work confirms that conclusion. In some cases, the DLM models from this work outperformed the DLM models tested in [14], specifically, in the quarterly, monthly and other studies.

Furthermore, there is something to be said concerning the use of the MSE as an accuracy measure. Throughout each study, the model with the lowest MSE was not always the model with the lowest average SMAPE across each horizon. As stated earlier, the MSE was used as an in-sample model selection, whereas, the SMAPE was an out-of-sample validation. The two accuracy measures did not always agree on best fit models; the MSE had a tendency to choose a model that performed rather poorly based on the SMAPE. Therefore, this work does not recommend using the MSE for the in-sample model selection.

Opportunities exist to further the improvement of the DLM, such as, using more sophisticated optimization algorithms to improve the fit of the model. This work used a grid search method, which basically means starting with a fixed parameter and systematically varying the parameter to achieve optimality. Other algorithms, such as the numerical steepest descent and Gauss Newton methods, could be experimented with as well. However, restrictions need to be applied in order to ensure that certain parameters remain positive. While Dynamic Linear Models have yet to be included in any of the M Competitions, they tend to be a good competition for other time series paradigms. The field of Dynamic Linear Models is diverse and future research may find them to be the models of choice in specific time series.

APPENDIX A

1. MATLAB code using the Monthly time series as an example

```
nfiles = size(seriesfile,1);
horiz=18;
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Set Parameters for DLM;

%Trend Seasonal
Ft1 = [1;1;1;0;0;0;0;0;0;0;0;0;0;0];

Gt1 = [1 1 0 0 0 0 0 0 0 0 0 0 0 0;
       0 1 0 0 0 0 0 0 0 0 0 0 0 0;
       0 0 0 1 0 0 0 0 0 0 0 0 0 0;
       0 0 0 0 1 0 0 0 0 0 0 0 0 0;
       0 0 0 0 0 1 0 0 0 0 0 0 0 0;
       0 0 0 0 0 0 1 0 0 0 0 0 0 0;
       0 0 0 0 0 0 0 1 0 0 0 0 0 0;
       0 0 0 0 0 0 0 0 1 0 0 0 0 0;
       0 0 0 0 0 0 0 0 0 1 0 0 0 0;
       0 0 0 0 0 0 0 0 0 0 1 0 0 0;
       0 0 0 0 0 0 0 0 0 0 0 1 0 0;
       0 0 0 0 0 0 0 0 0 0 0 0 1 0;
       0 0 0 0 0 0 0 0 0 0 0 0 0 1;
       0 0 1 0 0 0 0 0 0 0 0 0 0 0];

%No Trend Seasonal
Ft = [1;1;0;0;0;0;0;0;0;0;0;0;0];

Gt =  [1 0 0 0 0 0 0 0 0 0 0 0 0;
       0 0 1 0 0 0 0 0 0 0 0 0 0;
       0 0 0 1 0 0 0 0 0 0 0 0 0;
       0 0 0 0 1 0 0 0 0 0 0 0 0;
       0 0 0 0 0 1 0 0 0 0 0 0 0;
       0 0 0 0 0 0 1 0 0 0 0 0 0;
       0 0 0 0 0 0 0 1 0 0 0 0 0;
       0 0 0 0 0 0 0 0 1 0 0 0 0;
       0 0 0 0 0 0 0 0 0 1 0 0 0;
       0 0 0 0 0 0 0 0 0 0 1 0 0;
       0 0 0 0 0 0 0 0 0 0 0 1 0;
       0 0 0 0 0 0 0 0 0 0 0 0 1;
       0 1 0 0 0 0 0 0 0 0 0 0 0];

%No Trend No Season
Ft2 = [1]; Gt2 = [1];

%Trend No Season
Ft3 = [1;0]; Gt3 = [1 1; 0 1];

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

baserow = size(Gt,1);
            wt = 1;
```

```
        W1 = 1;
        dt = 1;
        St = 1;
        Ct = 1;
        DLMMSEinit = 1e20; % Initial the MSE with a huge value

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%No Trend Seasonal

    rm1 = polyfit((1:24)',tseries(1:24),1);
    rm1f = [(1:24)',ones(24,1)]*rm1';
    mt12 = rm1f(1);
    mt11 = mean(tseries(1:4));
    m1 = (tseries(1) - rm1f(1) + tseries(13) -rm1f(13))/2;
    SSI2 = abs(tseries(1)-tseries(13))/4;
    m2 = (tseries(2) - rm1f(2) + tseries(14) -rm1f(14))/2;
    SSI3 = abs(tseries(2)-tseries(14))/4;
    m3 = (tseries(3) - rm1f(3) + tseries(15) -rm1f(15))/2;
    SSI4 = abs(tseries(3)-tseries(15))/4;
    m4 = (tseries(4) - rm1f(4) + tseries(16) -rm1f(16))/2;
    SSI5 = abs(tseries(4)-tseries(16))/4;
    m5 = (tseries(5) - rm1f(5) + tseries(17) -rm1f(17))/2;
    SSI6 = abs(tseries(5)-tseries(17))/4;
    m6 = (tseries(6) - rm1f(6) + tseries(18) -rm1f(18))/2;
    SSI7 = abs(tseries(6)-tseries(18))/4;
    m7 = (tseries(7) - rm1f(7) + tseries(19) -rm1f(19))/2;
    SSI8 = abs(tseries(7)-tseries(19))/4;
    m8 = (tseries(8) - rm1f(8) + tseries(20) -rm1f(20))/2;
    SSI9 = abs(tseries(8)-tseries(20))/4;
    m9 = (tseries(9) - rm1f(9) + tseries(21) -rm1f(21))/2;
    SSI10 = abs(tseries(9)-tseries(21))/4;
    m10 = (tseries(10) - rm1f(10) + tseries(22) -rm1f(22))/2;
    SSI11 = abs(tseries(10)-tseries(22))/4;
    m11 = (tseries(11) - rm1f(11) + tseries(23) -rm1f(23))/2;
    SSI12 = abs(tseries(11)-tseries(23))/4;
    m12 = (tseries(12) - rm1f(12) + tseries(24) -rm1f(24))/2;
    SSI13 = abs(tseries(12)-tseries(24))/4;

    mt = [mt11;m1;m2;m3;m4;m5;m6;m7;m8;m9;m10;m11;m12];
    Ct = diag([std(tseries(1:12))^2;SSI2;SSI3;SSI4;SSI5;SSI6;SSI7;SSI8;SSI9;SSI10;SSI11;SSI12;SSI13]);

    w1star=.01;
    W1star = diag([std(tseries(1:12)/10)^2;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star]);
    %W1star = Ct;

    MSENTSM = DLMFNMSE12(tseries(1:datavalid),Ft,Gt,mt,Ct,dt,St,W1star,W1star);


 %Begin DLM
      [f,f1] = DLMFN12(tseries(1:datavalid),Ft,Gt,mt,Ct,dt,St,W1star,W1star,horiz);
      finleng=size(f,2);
      smapein(i,:) = SMAPE_DLM(f1,trgnum);

      outputsmape1(i,:) = cat(2,horiz,smapein(i,:));
      f2 = f;
```

```
        f3 = f1;
        finleng1 = finleng;
        NoTrend(i,:)=cat(2,outputsmape1(i,:),MSENTSM);

        allfoc=cat(2,f(1:finleng-1),f1);
        subplot(2,2,1), plot(2:datasize,tseries(2:datasize),'b-',2:datasize,allfoc,'r--','linewidth',2);
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%Trend Seasonal

      rm1 = polyfit((1:24)',tseries(1:24),1);
      rm1f = [(1:24)',ones(24,1)]*rm1';
      mt12 = rm1(1);
      mt11 = rm1(2);
      s1tr = (std(tseries(13:24))^2+std(tseries(1:12))^2)/16;
      m1 = (tseries(1) - rm1f(1) + tseries(13) -rm1f(13))/2;
      SSI2 = abs(tseries(1)-tseries(13))/4;
      m2 = (tseries(2) - rm1f(2) + tseries(14) -rm1f(14))/2;
      SSI3 = abs(tseries(2)-tseries(14))/4;
      m3 = (tseries(3) - rm1f(3) + tseries(15) -rm1f(15))/2;
      SSI4 = abs(tseries(3)-tseries(15))/4;
      m4 = (tseries(4) - rm1f(4) + tseries(16) -rm1f(16))/2;
      SSI5 = abs(tseries(4)-tseries(16))/4;
      m5 = (tseries(5) - rm1f(5) + tseries(17) -rm1f(17))/2;
      SSI6 = abs(tseries(5)-tseries(17))/4;
      m6 = (tseries(6) - rm1f(6) + tseries(18) -rm1f(18))/2;
      SSI7 = abs(tseries(6)-tseries(18))/4;
      m7 = (tseries(7) - rm1f(7) + tseries(19) -rm1f(19))/2;
      SSI8 = abs(tseries(7)-tseries(19))/4;
      m8 = (tseries(8) - rm1f(8) + tseries(20) -rm1f(20))/2;
      SSI9 = abs(tseries(8)-tseries(20))/4;
      m9 = (tseries(9) - rm1f(9) + tseries(21) -rm1f(21))/2;
      SSI10 = abs(tseries(9)-tseries(21))/4;
      m10 = (tseries(10) - rm1f(10) + tseries(22) -rm1f(22))/2;
      SSI11 = abs(tseries(10)-tseries(22))/4;
      m11 = (tseries(11) - rm1f(11) + tseries(23) -rm1f(23))/2;
      SSI12 = abs(tseries(11)-tseries(23))/4;
      m12 = (tseries(12) - rm1f(12) + tseries(24) -rm1f(24))/2;
      SSI13 = abs(tseries(12)-tseries(24))/4;
      mt = [mt11;mt12;m1;m2;m3;m4;m5;m6;m7;m8;m9;m10;m11;m12];
      Ct = diag([std(tseries(1:12))^2;s1tr;SSI2;SSI3;SSI4;SSI5;SSI6;SSI7;SSI8;SSI9;SSI10;SSI11;SSI12;SSI13]);
      w1star=.01;
      W1star = diag([std(tseries(1:12)/1000)^2;s1tr/100000;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star;w1star]);
      %W1star = Ct;
      MSETSM = DLMFNMSE12(tseries(1:datavalid),Ft1,Gt1,mt,Ct,dt,St,W1star,W1star);

 %Begin DLM
    [f,f1] = DLMFN12(tseries(1:datavalid),Ft1,Gt1,mt,Ct,dt,St,W1star,W1star,horiz);
    finleng=size(f,2);
    smapein(i,:) = SMAPE_DLM(f1,trgnum);

    outputsmape2(i,:) = cat(2,horiz,smapein(i,:));
    f4 = f;
    f5 = f1;
    finleng2 = finleng;
    WithTrend(i,:)=cat(2,outputsmape2(i,:),MSETSM);
```

```
    allfoc=cat(2,f(1:finleng-1),f1);
    subplot(2,2,2) , plot(2:datasize,tseries(2:datasize),'b-',2:datasize,allfoc,'r--','linewidth',2);
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%No Trend No Season

      m1 = mean(tseries(1:24));
      mt = [m1];
      Ct = [std(tseries(1:24))^2];
      W1star = 10;
      MSENTNSM = DLMFNMSE12(tseries(1:datavalid),Ft2,Gt2,mt,Ct,dt,St,W1star,W1star);

   %Begin DLM
      [f,f1] = DLMFN12(tseries(1:datavalid),Ft2,Gt2,mt,Ct,dt,St,W1star,W1star,horiz);
      finleng=size(f,2);
      smapein(i,:) = SMAPE_DLM(f1,trgnum);

      outputsmape3(i,:) = cat(2,horiz,smapein(i,:));
      f6=f;
      f7=f1;
      finleng3=finleng;
      NoT_NoS(i,:)=cat(2,outputsmape3(i,:),MSENTNSM);

      allfoc=cat(2,f(1:finleng-1),f1);
      subplot(2,2,3) , plot(2:datasize,tseries(2:datasize),'b-',2:datasize,allfoc,'r--','linewidth',2);
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%Trend No Season

      rm1 = polyfit((1:24)',tseries(1:24),1);
      rm1f = [(1:24)',ones(24,1)]*rm1';
      mt12 = rm1(1);
      mt11 = rm1(2);
      s1tr = (std(tseries(13:24))^2+std(tseries(1:12))^2)/16;
      mt = [mt11;mt12];
      Ct = diag([std(tseries(1:8)/1000)^20;s1tr/100000]);
      W1star = Ct;

  %Begin DLM
      MSETNSM = DLMFNMSE12(tseries(1:datavalid),Ft3,Gt3,mt,Ct,dt,St,W1star,W1star);
      [f,f1] = DLMFN12(tseries(1:datavalid),Ft3,Gt3,mt,Ct,dt,St,W1star,W1star,horiz);
      finleng=size(f,2);
      smapein(i,:) = SMAPE_DLM(f1,trgnum);

      outputsmape4(i,:) = cat(2,horiz,smapein(i,:));
      f8=f;
      f9=f1;
      finleng4=finleng;
      Trend_NoS(i,:)=cat(2,outputsmape4(i,:),MSETNSM);

      allfoc=cat(2,f(1:finleng-1),f1);
      subplot(2,2,4) , plot(2:datasize,tseries(2:datasize),'b-',2:datasize,allfoc,'r--','linewidth',2);
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
      % Compare MSEs;

      MSETEST = [MSENTSM;MSETSM;MSENTNSM;MSETNSM];
```

```
    [MSEResult,MSEind] = min(MSETEST);

     if MSEind == 1
         outputsmape(i,:) = outputsmape1(i,:);
         %model = 'MSENTQ';
         f=f2;
         f1=f3;
         finleng=finleng1;
     end
     if MSEind == 2
         outputsmape(i,:) = outputsmape2(i,:);
         %model = 'MSETQ';
         f=f4;
         f1=f5;
         finleng=finleng2;
     end
     if MSEind == 3
         outputsmape(i,:) = outputsmape3(i,:);
         %model = 'MSENTNQ';
         f=f6;
         f1=f7;
         finleng=finleng3;
     end
     if MSEind == 4
         outputsmape(i,:) = outputsmape4(i,:);
         %model = 'MSETNQ';
         f=f8;
         f1=f9;
         finleng=finleng4;
     end

     model=MSEind;
     OPTSMAPE(i,:) = cat(2,model,outputsmape(i,:));
     fprintf(fid,'%s, %d, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4f, % 8.4

     csvwrite('T:\THESIS\THESIS_CODE\Monthly\DLM_Monthly_NoTrend1a.csv', NoTrend)
     csvwrite('T:\THESIS\THESIS_CODE\Monthly\DLM_Monthly_WithTrend1a.csv', WithTrend)
     csvwrite('T:\THESIS\THESIS_CODE\Monthly\DLM_Monthly_NoT_NoS1a.csv', NoT_NoS)
     csvwrite('T:\THESIS\THESIS_CODE\Monthly\DLM_Monthly_Trend_NoS1a.csv', Trend_NoS)

end
      allfoc=cat(2,f(1:finleng-1),f1);
      figure;plot(2:datasize,tseries(2:datasize),'b-',2:datasize,allfoc,'r--','linewidth',2); %grahp
      title({['Monthly DLM: ', file1]; 'Solid:Actual  Dash:Forecast'});

  toc;

  fclose(fid);
  t=toc
```

2. SAS code for the Linear Mixed Model

```
proc import out=work.yearly_h1
    datafile="Yearly_Model1_AllHorizons.xls"
    dbms=EXCEL2000 replace;
    range="Sheet1$";
    getnames=yes;
run;

proc mixed data=yearly_h1;
    class paradigm series;
    model smape=paradigm;
    random series;
    lsmeans paradigm/pdiff;
    title "Mixed ANOVA for yearly data, H=1";
run;
```

REFERENCES

[1] "Bayes' Theorem". Wikipedia.org 26 June 2006.
$< http : //en.wikipedia.org/wiki/Bayes'\_theorem >$

[2] Box, George E. and Jenkins, Gwilym M. *Time Series Analysis: forecasting and control*. Holden-Day, Inc., 1976.

[3] Campagnoli, Patrizia, Muliere, Pietro and Petrone, Sonia. "Generalized Dynamic Linear Models for Financial Time Series." *Applied Stochastic Models in Business and Industry* 17 (2001): 27 - 39.

[4] "Famous Forecasting Quotes". Duke.edu 20 March 2006.
$< http : //www.duke.edu/ \sim rnau/411quote.htm >$.

[5] Fischhoff, Baruch. "What Forecasts Mean." *International Journal of Forecasting* 10 (1994): 387-403.

[6] Makridakis, S. and Hibon, M. "Accuracy of Forecasting: An Empirical Investigation." *Journal of the Royal Statistical Society* Series A (1979): 97-145.

[7] Makridakis, Spyros, et. al. "The M2 Competition: A Real-Time Judgementally Based Forecasting Study." *International Journal of Forecasting* 9 (1993): 5-23.

[8] Makridakis, Spyros. "Forecasting: Its Role and Value for Planning and Strategy." *International Journal of Forecasting* 12 (1996): 513-537.

[9] Makridakis, Spyros and Hibon, Michèle. "The M3 Competition: Results, Conclusions and Implications." *International Journal of Forecasting* 16 (2000): 451-476.

[10] McCulloch, Charles E. and Searle, Shayle R. *Generalized, Linear, and Mixed Models*. New York: John Wiley & Sons, Inc., 2001.

[11] Pankratz, Alan. *Forecasting with Dynamic Regression Models*. New York: John Wiley & Sons, Inc., 1991.

[12] Pole, Andy, West, Mike and Harrison, Jeff. *Applied Bayesian Forecasting and Time Series Analysis*. Chapman & Hall, 1994.

[13] West, Mike and Harrison, Jeff. *Bayesian Forecasting and Dynamic Models*. New York: Springer, (1997).

[14] Zhai, Yusheng. "Time Series Forecasting Competition Among Three Sophisticated Paradigms." MS Thesis. University of North Carolina Wilmington, 2005.