Seasonal Adjustment and Dynamic Linear Models

1 Introduction

Predicting the future by inferring from the past requires that information is analyzed for its possible impact forward in time. This judgment can then be extrapolated to the prediction horizon. Several methods can help in integrating the information into a point estimate in order to learn about the past in an organized and replicable manner. This process is named modeling and can be described as something that "*organizes information and experiences providing a means of learning and forecasting*" (West & Harrison, 1989).

One approach to model building is in terms of a dynamic linear model, abbreviated DLM, and is due to Harrison & Stevens (1976). Based on the Kalman (1960) filter, DLMs are an adaptive Bayesian approach of model formulation in terms of state and space. In this formulation, the time series observation Y_t at time *t* is assumed to consist of an unobservable component θ_t and some noise ν_t. This component is in turn assumed to be a function of itself at a previous time point and some innovation noise ω_t. Following the notation used in West & Harrison (1989), the DLM system is stated as

Observation (Space): $$Y_t = \theta_t + \nu_t,$$
Component (State): $$\theta_t = G\theta_{t-1} + \omega_t.$$

A DLM uses Bayesian updating of the posterior distribution, so the unobserved component θ_t is updated by assimilating the prior knowledge θ_{t-1} and the data Y_t. This recursive filtering, which implies the Markov property that all necessary prior information is stored in the previous estimate, is shown by Kalman (1960) to be the optimal estimate of the unobserved component. It is based on partitioning the observation into a random component orthogonal to a linear component in the observation space.

A dynamic linear model carries this partitioning one step further and allows for more extensive modeling. The unobserved component is itself divided into components, e.g. a seasonal and a trend: $\theta_t = \tau_t + s_{t,j(p)}$. Here, a trend τ_t and a seasonal $s_{t,j(p)}$ at time *t* affect an observation at season *j* with a recurrence period of length *p* (say *p* = 4 for quarterly data).

A more commonly used model framework is the autoregressive and integrated moving average models (ARIMA) approach, as introduced by Box & Jenkins (1970). The technique is a way of modeling time series with linear functions and normally distributed random components. If the data follows an ARIMA model, the parameters are usually estimated by the Maximum Likelihood method and a model is fit to *past* observations. This approach often involves data filtering in terms of transformations (e.g. logarithmic or Box-Cox), outlier processing, dummy variable fitting (i.e. intervention variables) and integration (i.e. differencing) of various orders before the best (minimum squared error) model can be fitted. By construction, it is an averaging method that fits something that *on average* works well on *historical* data. Similar to DLM, these models may also be expressed in a state space representation, see e.g. Hamilton (1994).

A more algorithmic approach to model fitting, which also has state space representation, is the exponential smoothing methodology, originating from work by Holt (1957), Brown (1959) and Winters (1960). These methods were developed for forecasting and do not use statistical theory, unlike the ARIMA and DLM approaches. Yet, they offer similar possibilities as DLM in partitioning the observed time series into unobserved components such as the seasonal decomposition. This method requires parameters, but since it is not based on maximum likelihood inference, parameters can be estimated more flexibly than the ARIMA approach, without involving squared error functions to minimize.

2 Seasonal adjustments by DLM and exponential smoothing

Seasonal effects exhibited in many time series sometimes need to be cleared out. This is called seasonal adjustment and is often necessary for comparative analysis of time series. In official statistics production however, state space approaches are not commonly used for this purpose. Instead, two different kinds of signal filtering approaches are applied with the assistance of ARIMA modeling, namely the X11/X12-ARIMA and TRAMO/SEATS methods. These methods are well adapted to practical use but they are somewhat complicated to grasp. Put in that context, simpler state space approaches such as DLM (or exponential smoothing) should receive attention as they remain fairly unexplored for this specific purpose. Despite the rigorously supporting theory behind the Kalman filter, DLM are simple to intuitively understand and rather straightforward to apply. Seasonal estimation by X11/X12-ARIMA and TRAMO/SEATS have been compared in many studies, but often the comparison criteria can be difficult to understand, whereas to the authors knowledge, no comparisons of these standard methods have been made with DLM.

The first paper in this thesis (Section I) is an introductory and very basic attempt to compare seasonal adjustments made by DLM to those made by TRAMO/SEATS. Some diagnostics and comparison criteria are stated, such as mean squared errors of components, model fit, roughness etc., and applied to either method or to both methods when possible for comparative analysis.

Two varieties of DLM are used in the first study. First, a simpler seasonal DLM approach (similar to the Kalman filter) is applied to both simulated time series and empirical data. This model is found to be unsatisfactory with respect to the criteria stated. Second, for empirical data, a more elaborated DLM approach is applied and requires estimating parameters to use in the recursion. This is achieved by Gibbs sampling and turns out to be far more satisfactory than the basic model, although both DLM models are inferior to the benchmark TRAMO/SEATS and tend to smooth excessively, i.e. they appear to be too rigid. For the artificial data, it is seen that the DLM with informative priors can compare somewhat well with TRAMO/SEATS, but this result is limited to a few cases and occurs only for the irregular and trend components, although some cases are close to ties between the methods. The simple DLM with uninformative priors appears as unsatisfactory throughout. The overall conclusion is that this should be viewed as a window of opportunity for elaborating on DLM.

The data that are used in the study are time series of monthly Swedish foreign trade of goods, i.e. exports and imports. The particularity of this kind of data is the implicitly obtained trade balance (or net trade) series, which is exports minus imports. This kind of time series raises the issue of direct and indirect seasonal adjustments, which is briefly introduced in the first paper and is worked through in the two following papers on seasonal adjustment in Sections II and III.

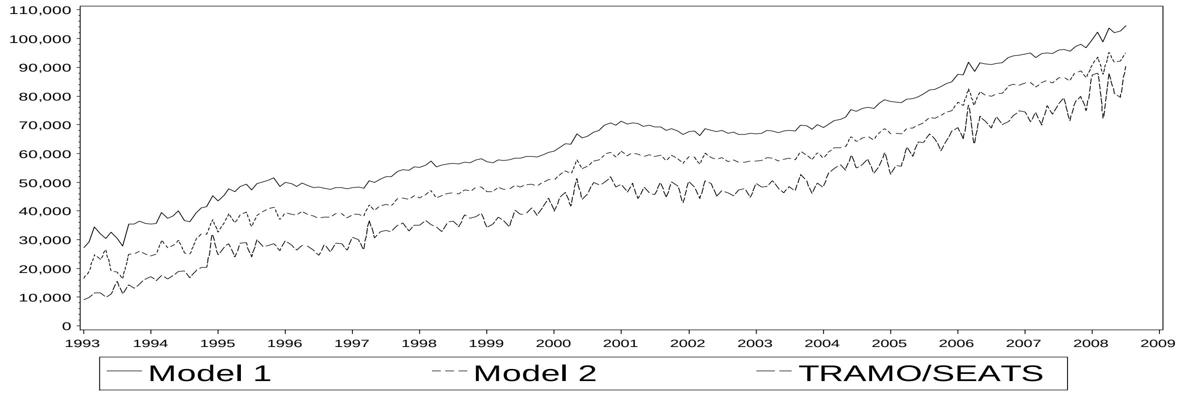

Figure 1. Seasonal adjustments of exports. Estimation through the simple Model 1 (on top), the elaborated Model 2 on minus 10 000 Million Swedish Kronor SEK of its actual value (in the middle) and TRAMO/SEATS on minus 20 000 Million Swedish Kronor SEK of its actual value (lowest line). January 1993-July 2009.

An example of seasonal adjustments of Swedish exports based on DLM and TRAMO/SEATS is shown in Figure 1, and the original time series is found in the first paper. Model 1 (the simpler DLM), gives the smoothest/stiffest line (on top), Model 2 (DLM with Gibbs sampling) renders a more fluctuating adjustment (in the middle) and TRAMO/SEATS (the lowest line) is more dynamic and is presumed to capture the underlying seasonal component more accurately, as concluded from simulated series.

3 Direct and indirect seasonal adjustments

Given that a time series is a function of two or more time series, such as the trade balance which is derived as exports minus imports, the issue of how to obtain the seasonal adjustment arises. One could either seasonally adjust both exports and imports prior to taking their difference, which gives the *indirectly* seasonally adjusted trade balance, or one could take their difference first and then *directly* seasonally adjust the obtained difference, i.e. the trade balance. Which way to go appears to be an unsolved issue in practice. In an early study, Geweke (1978) showed that in theory, a multivariate indirect approach was to prefer under the presumption that an optimal joint estimator was available, whereas e.g. Planas & Campolongo (2000) found that the problem depends on the spectral densities of the input series. Maravall (2005, 2006) argued for the benefits of the direct approach, which is a highly relevant point-of-view for various practical reasons. One of his arguments was that noisy subseries tend to aggregate to a less noisy total.

In the first paper, the DLM is found to render very small discrepancies between direct and indirect seasonal adjustments of the trade balance compared with TRAMO/SEATS. This is an expected result since an identical DLM filter is used for the three input series, whereas in TRAMO/SEATS, a unique adaptive filter is applied automatically for each of the series.

In the second paper (Section II), this issue is discussed more theoretically by applying a simple seasonal level model with no trend in state space. It is assumed that the underlying time series process is in a steady state, i.e. stationary. The innovation variance Var(ω_t)=W of the unobserved component, i.e. the signal, and the observation noise variance Var(ν_t)=V are derived for a system consisting of two artificial time series, rendering the variance matrices **W** and **V** for direct, indirect and optimal/multivariate models. Given each of the two series' signal-to-noise ratios (W/V), the ingoing covariance components are varied (by varying the correlations) in order to study the relative efficiency between the direct, indirect and optimal/multivariate estimations. The method is then applied to Swedish foreign trade data (after a data transformation) to infer which approach to prefer.

An example of the relative efficiency for a specific set of variance combinations (and thus signal to noise ratios), is given in Figure 2 and specified in the second paper. The vertical axis shows the relative efficiency, computed for a grid of correlations of the variance matrices W and V, where each correlation implicitly determines the covariance between the components. As can be seen, there is an area below unity where indirect seasonal adjustment is to prefer, otherwise direct seasonal adjustment is preferable for this specific combination of signal-to-noises. One conclusion drawn from the study is that for some correlations, it is always beneficial to apply direct seasonal adjustment, given that an optimal/multivariate method cannot be used. This is seen in Figure 2 as the areas where the correlations between component variances are strongly negative.

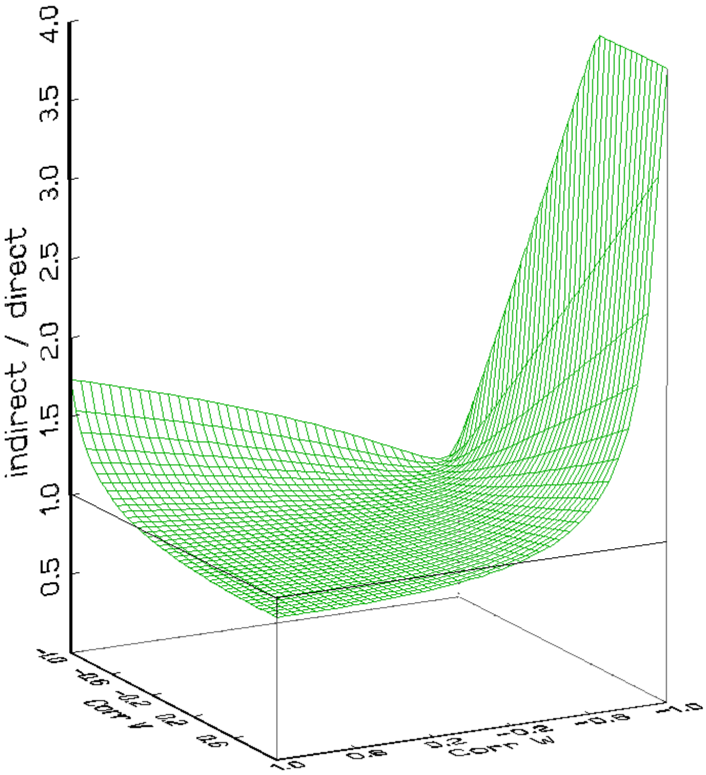

Figure 2. Relative efficiency of direct adjustment over indirect adjustment.

In the third paper (Section III), the issue is addressed with another presumption. In reality, using direct seasonal adjustment does not eliminate the possibility of obtaining the indirect seasonal adjustment since the included time series (i.e. the subseries) often have to be processed anyway. One might instead consider the possibility of seasonally adjusting the individual series so that their adjustments account for the aggregate as well. This is accomplished by formulating a total loss function that covers both the individual series and their aggregate. The loss function is minimized in a continuum from direct to indirect seasonal adjustments with some trade-off weights α reaching between the two approaches and conceptualizing the preference of minimum errors. This preference frontier formulation for an arbitrary loss function L(ε) for residuals ε is shown in Figure 3 below.

$$\min\left(\sum_{j=1}^{n}\left(\sum_{t=1}^{T} L_j(\varepsilon_{t,j})\right)\right) \qquad \min\left(\sum_{t=1}^{T} L_{Agg}\left(\sum_{j=1}^{n}\varepsilon_{t,j}\right)\right)$$

1 —— α —— 0

Figure 3. Trade-off frontier between indirect estimation ($\alpha = 1$) with *n* individual losses L_j and direct estimation ($\alpha = 0$) with aggregate loss $L_{Agg.}$.

An exponential smoothing method with a trend and a seasonal component is used, and necessary parameters are estimated so that the weighted total loss is minimized. The total loss is computed with either of two different loss functions: first, a squared error loss is applied, which is the standard loss function when normality is assumed; second, a Huber loss function known from robust estimation (Huber, 1964) is used. The method is applied to both quarterly Swedish gross domestic product (GDP) series, and to monthly Swedish foreign trade data used in the two previous papers (in Sections I and II). This results in guidelines on what kind a seasonal adjustment to apply and what not to do. It is also seen that differences between direct and indirect seasonal adjustments are smaller when the Huber loss function is applied compared to when the standard squared loss function is applied.

To summarize, the three papers presented in this thesis convey the issue of direct and indirect seasonal adjustments from different perspectives. In the first paper, some criteria are stated by which direct and indirect seasonal adjustments from two different methods are compared. In the second paper, the situations for which approach to prefer are discussed, given a model and evaluation criteria. In the third paper, the model is estimated so that a specific criterion is fulfilled. This overall process for determining what approach to prefer can be summarized as doing the following: 1) state some criteria for desired properties of the seasonal adjustment 2) set up a model to be estimated with minimum expected error and 3) estimate the model based on a target function with respect to the requirements in 1) and 2).

4 Using state space models to get the best estimate of opinions

Public opinion polls in Sweden are most often made by several survey institutes and are usually published on irregular dates, especially prior to elections. Even if each polling institute does a good job in their estimate of the political opinion, polling results tend to differ between institutes. In an attempt to reduce this uncertainty of the true opinion estimate, a model based approach can be used that combines the different polls in order to obtain a more accurate estimate of the underlying true opinion, see e.g. Silver (2008). This is an emerging "*Wisdom of crowds*"- approach (Surowiecki, 2005) based on Bayesian inference, which was pointed out by Harrison & Stevens (1976) as a Multi-Process modeling problem.

In the fourth paper (Section IV), a continuous time stochastic process model is applied to Swedish opinion polls to obtain the best estimate of the true political opinion. When new opinion polls are observed, often at non-equidistant time points and possibly with polling institute specific design effects and bias, the model is updated similar to DLM through the

Bayesian update of the posterior distribution. The changes in party preferences, which are the target parameter, are considered to follow a Wiener process, modeled by a stochastic differential equation without any auxiliary information since the observed opinion poll results are assumed to account for all available political information, similar to stock market models. In an attempt to reduce uncertainty even more, a trend is introduced and modeled as an Ornstein-Uhlenbeck process. We are able to fit a model both for the individual parties in the Swedish Parliament and for a block of parties, i.e. the Alliance (*Alliansen*), which has been the incumbent block since September 2006.

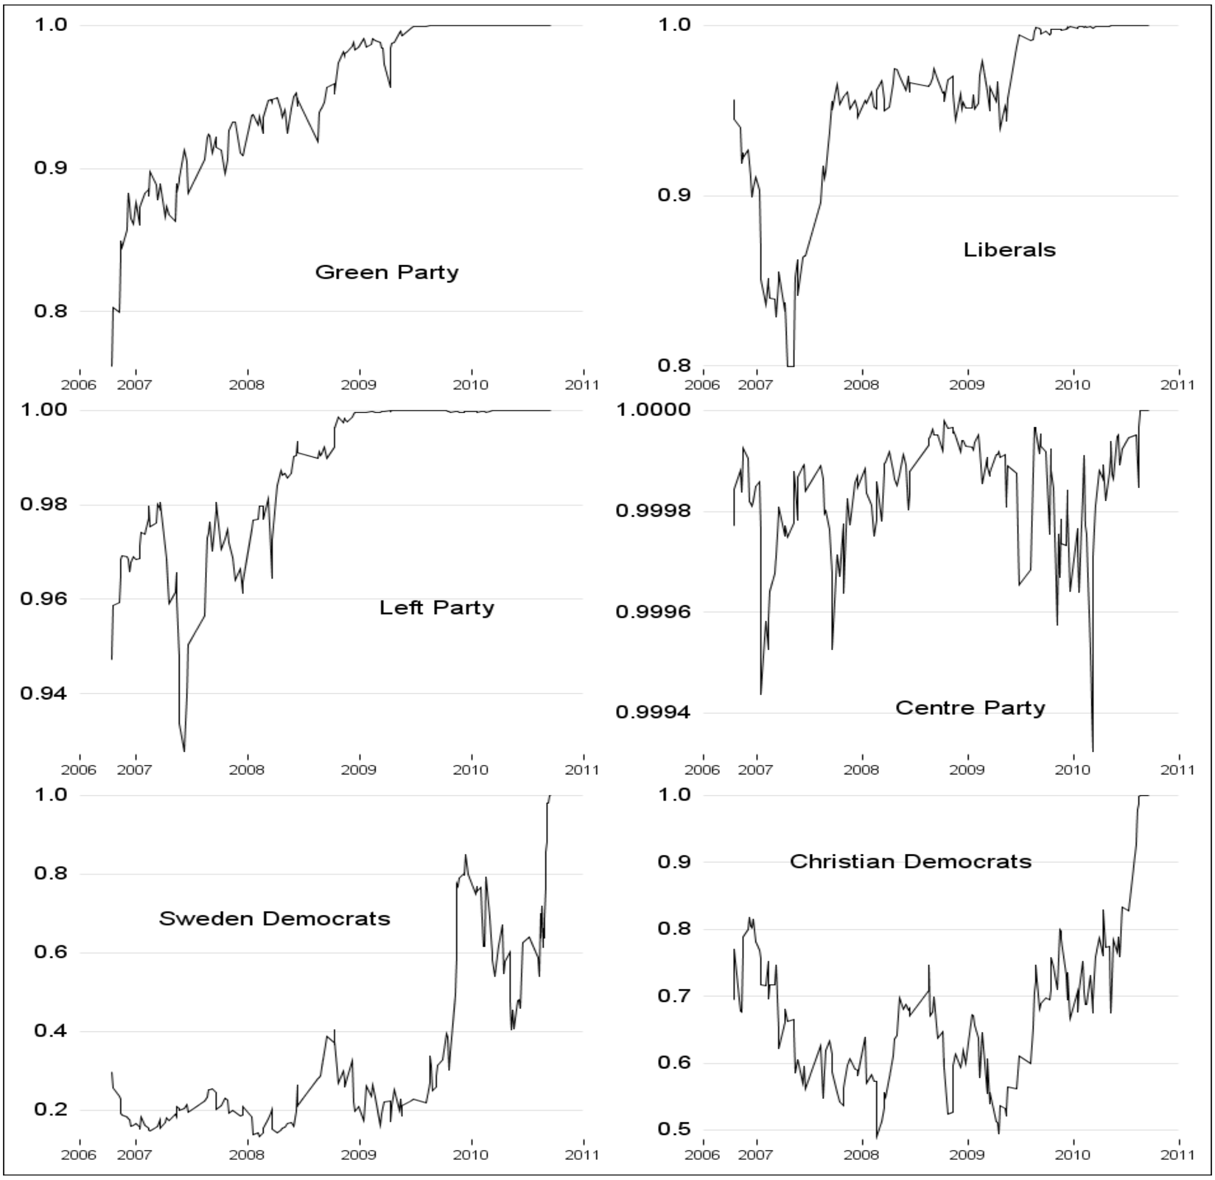

Figure 4. Probability of exceeding the parliamentary threshold of 4 % given all previous polls for all parties except the Social Democrats and the Moderates which have probability one. Note that the scales on the y-axes are different between the parties.

In Figure 4, the estimated probabilities of passing the election threshold in the elections in 2010 are shown for all sitting parties in the parliament except the two largest parties (the Social Democrats and the Moderates). As a fringe benefit of our model, we can predict the election outcome within a specific time window, e.g. three weeks prior to elections.

5 Issues that remain unexplored

5.1 Outliers in DLM

A remark made by an anonymous referee to the first paper in this thesis was on how to deal with outliers in the DLM context. When likelihood based inference is used, outliers can be identified by their contribution to the likelihood so that observations excessively influencing the likelihood function are accounted for in some sense. As for recursive estimations, like DLM, my proposition to dealing with outliers *ex ante* would be to consider the prior-to-posterior updating differences, which could be used as an *empirical* basis for determining when new observations should be treated (mechanically) as outliers. However, sometimes what appears to be an outlier is in fact a crucial turning point, which in practice is realized only *ex post*.

5.2 The variance discounting in DLM

One of the more hands-on actions when specifying a DLM is the choice of discounting value of the variances, which affects the persistence/duration of noise in the system. This choice has a substantial influence on the goodness of fit but has no apparent prior. In other model frameworks, discounting problems are sometimes viewed in relation to the frequency in data, i.e. depending on whether daily, quarterly or monthly data are used (see Öller, 1978). For DLM, a discounting strategy needs to be worked out and put in relation to some fitting parameters.

5.3 Exploring direct and indirect seasonal adjustments more deeply

Doing seasonal adjustments as a trade-off between direct and indirect seasonal adjustments is a fairly unexplored issue and has been examined here only through an exponential smoothing method. An extension of this trade-off formulation to other frameworks could be studied since the exponential smoothing method is practically not used at all for seasonal adjustments in official statistics.

6 Concluding the work

The ideas presented in this thesis have their basis in some unsolved issues. Direct or indirect seasonal adjustments remain as a practical dilemma but may be remedied by the ideas presented here. The dynamic linear models appear as a possible method for seasonal adjustments, but they require more elaboration to compete with standard tools. However, in terms of projections and predictions, the dynamic approach shows to be quite useful when applied to opinion poll results as they produce inference on the true party preferences and as they predict the election outcome.

References

Box, G.E.P. & Jenkins, G.M. (1970). *Time Series Analysis, Forecasting and Control*. Holden-Day, San Francisco.

Brown, R.G. (1959). *Statistical Forecasting for Inventory Control*. McGraw-Hill, New York.

Geweke , J. (1978). The Temporal and Sectoral Aggregation of Seasonally Adjusted Time Series. In *Seasonal Analysis of Economic Time Series*. A. Zellner (ed.). pp. 411-427. U.S. Department of Commerce, Bureau of the Census, Washington DC.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press, Princeton New Jersey.

Harrison, J. P. & Stevens, C. F. (1976). Bayesian Forecasting. *Journal of the Royal Statistical Society*, series B (Methodological), vol. 38, no. 3.

Holt, C.C. (1957). *Planning Production, Inventories and Work Force*. Prentice Halls, Englewood Cliffs, New Jersey.

Huber, P. J. (1964). Robust Estimation of a Location Parameter. *Annals of Mathematical Statistics*, vol. 35, number 1, pp. 73-101.

Kalman, R.E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* 82D, 33-45.

Maravall, A. (2005). An application of the TRAMO-SEATS automatic procedure; direct versus indirect adjustment. *Computational Statistics & Data Analysis* 50, pp. 2167-2190.

Maravall, A. (2006). An application of the TRAMO-SEATS automatic procedure; direct versus indirect adjustment, Documentos de Trabajo No. 0534, Banco de España.

Öller, L-E. (1978). *Time Series Analysis of Finnish Foreign Trade*. Ph.D. thesis, University of Helsinki.

Planas, C. & Campolongo, F. (2000). The Seasonal Adjustment of Contemporaneously Aggreagated Series. Joint research centre of European Commission.

Silver, N. (2008). Politics done right (Fivethirtyeight). Online resource available by 1 April 2013 at: http://www.fivethirtyeight.com/2008/03/frequently-asked-questions-last-revised.html

Surowiecki, J. (2005). *The Wisdom of Crowds*. Anchor Books.

West, M. & Harrison J. (1989). *Bayesian Forecasting and Dynamic Models*. Springer-Verlag New York Inc.

Winters, P.R. (1960). Forecasting Sales by Exponentially Weighted Moving Averages. *Management Science* vol. 6, pp 324-342.