MONASH University

Australia

Department of Econometrics and Business Statistics

http://www.buseco.monash.edu.au/depts/ebs/pubs/wpapers/

**A View of Damped Trend as
Incorporating a Tracking Signal into a State Space Model**

**Ralph D. Snyder, Monash University and
Anne B. Koehler, Miami University**

September 2008

A View of Damped Trend as

Incorporating a Tracking Signal into a State Space Model

Abstract

Damped trend exponential smoothing has previously been established as an important forecasting method. Here, it is shown to have close links to simple exponential smoothing with a smoothed error tracking signal. A special case of damped trend exponential smoothing emerges from our analysis, one that is more parsimonious because it effectively relies on one less parameter. This special case is compared with its traditional counterpart in an application to the annual data from the M3 competition and is shown to be quite competitive.

Keywords: exponential smoothing, monitoring forecasts, structural change, adjusting forecasts, state space models, damped trend

1. Introduction

Damped trend exponential smoothing (Gardner and McKenzie, 1985) has a reputation for being a robust forecasting method (Gardner, 2006; Hyndman et al., 2008; and Fildes, 2008). In this paper we attempt to demonstrate why this should be the case. Using innovations state space models (Hyndman et al., 2008), we show that this method has close links with simple exponential smoothing when used in conjunction with a smoothed-error tracking signal to detect and respond to structural change.

The white noise model forms a convenient starting point for our study. It takes the form $y_t = \mu + \varepsilon_t$, where y_t is the series value in typical period t, and μ is a time invariant mean representing the global underlying level of the series. The error term ε_t is a normally distributed shock with a mean 0 and a constant variance σ^2 and is uncorrelated with the error terms in other time periods. The effect of a shock on the series is restricted to only one period, the one in which it occurs; it has no spillover effect on future series values.

The white noise model invariably fails in business and economic applications, one of a number of reasons being that it ignores the impact of structural change so often found in business and economic processes. Structural change is said to occur when a shock leads to changes to the underlying states of a process that in turn impact on future values of a time series. The white noise model can be adapted to allow for structural change; the mean μ is replaced by a time dependent random variable ℓ_t called the *local level*. The resulting innovations state space model is

$$y_t = \ell_{t-1} + \varepsilon_t \tag{1.1}$$

$$\ell_t = \ell_{t-1} + \alpha\varepsilon_t \tag{1.2}$$

The one period lag in the measurement equation (1.1) is used to indicate that underlying level references the *beginning* of period t.

The term $\alpha\varepsilon_t$ is the *change* in the underlying level from one period to the next. It is an underlying growth rate, but one which is uncertain because it is affected by the shock. Hence, Equations (1.1) and (1.2) define a *stochastic trend*. The following features of this underlying growth rates may be inferred:

2

1. they are normally distributed;

2. they have a common mean of zero and a variance $\alpha^2 \sigma^2$; and

3. they are temporally uncorrelated.

Models that are more general are conceivable, where the underlying growth has more enhanced features. To distinguish it from its more counterparts, we refer to it as a *simple stochastic trend*.

The parameter α determines the magnitude of the underlying growth. When $\alpha = 0$ there is no underlying growth, a situation corresponding to the white noise process. In most applications α is positive, its magnitude being a reflection of the amount of structural change in a time series. When $\alpha = 1$, the model corresponds to a *random walk*.

The simple stochastic trend model is estimated using *simple exponential smoothing* (Brown, 1959). When a time series value y_t has been observed, it may be compared with the one-step-ahead prediction ℓ_{t-1}, to give the error $\varepsilon_t = y_t - \ell_{t-1}$. The underlying level is then updated with the transition equation $\ell_t = \ell_{t-1} + \alpha \varepsilon_t$. For a time series of length T, the seed level ℓ_0 and the parameter α, both of which determine the trajectory of the underlying levels, may be chosen to minimize the sum of squared errors $SSE = \sum_{t=1}^{T} \varepsilon_t^2$. The point forecasts for periods beyond the prediction origin at the end of period T all equal the final underlying level ℓ_T.

The *structural change* embedded in the simple stochastic trend occurs from one period to the next. Being entirely reflected by the underlying growth, the properties 1 to 3 above ensure that this structural change is relatively stable. Yet occasionally a series can also be affected by a rapid structural change or a structural break, changes that are typically disproportionate in size to the stable structural change of the local level model. Although an extensive literature exploring the possibility of incorporating structural breaks into a model exists, the occurrence of these breaks is typically too infrequent to permit the reliable prediction of their timing and size. It makes more sense, from a practical perspective, to ignore the existence of structural breaks in the model itself, but to adopt instead measures to detect structural breaks when they occur so that commensurate adjustments can then be made to the forecasts.

The structure of the paper is as follows. In Section 2, we examine a tracking signal that uses the smoothed one-step-ahead prediction errors statistic as a monitoring statistic. Then, it is argued that the smoothed one-step-ahead prediction errors statistic represents the average deviation in the underlying level of the time series from its 'ideal' value. This deviation is incorporated as an adjustment to the underlying level in a simple stochastic trend model to yield a model that underpins a restricted version of classical damped trend exponential smoothing method. In Section 3, the new restricted damped trend model is compared with the conventional damped trend models in a study based on the annual time series from the M3 competition database (Makridakis and Hibon, 2000).

2. A model combining simple exponential smoothing with a tracking signal

The *smoothed-error statistic* b_t is defined recursively by the equation

$$b_t = \phi b_{t-1} + \beta \varepsilon_t \tag{2.1}$$

where β is a smoothing parameter such that $\phi + \beta = 1$. This statistic is an exponentially weighted average of the errors, the weights declining back through time. It may be used in conjunction with control limits to monitor forecasts and is known as the *smoothed-error tracking signal* (cf. Farnum and Stanton, 1989). Under the null hypothesis that the error terms are independent and normally distributed, b_t has a normal distribution with mean 0 and standard deviation $\sigma\beta/\sqrt{1-\phi^2}$. The smoothed-error statistic b_t can be employed as a tracking signal, in which case the out-of-control region is

$$|b_t| > z\sigma\beta/\sqrt{1-\phi^2} \tag{2.2}$$

where z is the value of the standard normal variable corresponding to a specified level of significance. This tracking signal is the well-known exponentially weighted moving average (EWMA) control chart (Roberts, 1959) that is applied to the error terms.

The choice of β is a challenge. If the simple stochastic trend model (i.e. simple exponential smoothing) is used for forecasting, β could be chosen to be equal to α.

However, there is no reason why the same smoothing parameter should be used for both smoothing and monitoring. When distinct smoothing parameters are used, it is incumbent on us to suggest a mechanism for determining a separate value for β.

It should be stressed that systematic patterns also emerge in the smoothed-error statistic when the process generating the time series changes or the model has been incorrectly specified. For example, suppose the change in the underlying level is given by $b + \alpha\varepsilon_t$ instead of just $\alpha\varepsilon_t$ where b is a constant termed the *drift*. A time series under this assumption typically displays an upward drift when $b > 0$. So, when this new regime applies, the smoothed-error statistic has a constant upward bias which can be shown to be well approximated by the quantity b/α. Other plausible deviations from the assumptions of the simple stochastic trend lead to other systematic patterns in the smoothed-error statistic.

A more robust approach to forecasting emerges if the predictions are corrected each period by the average error as depicted by the smoothed-error statistic b_t. At the beginning of period t, the expected value of b_t is ϕb_{t-1}. The Equations (1.1) and (1.2) of the simple stochastic trend can be augmented by this expected change to give the new equations

$$y_t = \ell_{t-1} + \phi b_{t-1} + \varepsilon_t \tag{2.3}$$

$$\ell_t = \ell_{t-1} + \phi b_{t-1} + \alpha\varepsilon_t \tag{2.4}$$

$$b_t = \phi b_{t-1} + \beta\varepsilon_t \tag{2.5}$$

This is a *damped stochastic trend model*, the statistical framework underpinning damped trend exponential smoothing. It differs, however, in one respect. Because the smoothed-error statistic is an exponentially weighted average, the parameters ϕ and β satisfy the additional restriction $\phi + \beta = 1$. The new model is a *restricted* version of the classical damped stochastic trend model.

The use of this new model enables us to resolve the problem of determining the appropriate value of β. Under the null hypothesis that there are no structural breaks, the value of β can quite simply be estimated in conjunction with the parameter α while fitting the restricted stochastic damped trend model to the time series data. It is possible

to still resort to the use of simple exponential smoothing in conjunction with the smoothed error statistic to detect structural breaks, using the value of β that emerges from this fitting exercise. Such an approach would successfully detect structural breaks provided the assumptions underpinning simple exponential smoothing apply.

Restricted damped trend exponential smoothing should be reasonably robust. Because it effectively adjusts the underlying level by the average error, it should automatically adapt to many of the possible deviations from the underlying conditions that make simple exponential smoothing the appropriate approach to forecasting. It should therefore yield robust forecasts.

A number of interesting questions now emerge. First, the restricted damped trend model effectively has one less parameter than the conventional damped trend model. This makes it inherently simpler. However, does it yield competitive forecasts? This question is explored in Section 3.

Second, is it more effective to use the restricted stochastic damped trend model than the simple stochastic trend model with monitoring? In general, the answer is 'maybe'. The parameters of the damped trend model are tuned to the stable structural change situation. When the prospect of a structural break is remote, it must then be better. Otherwise, the method that relies on the tracking signal is likely to be the better performer, provided that appropriate responses are made when the out-of-control situations are detected.
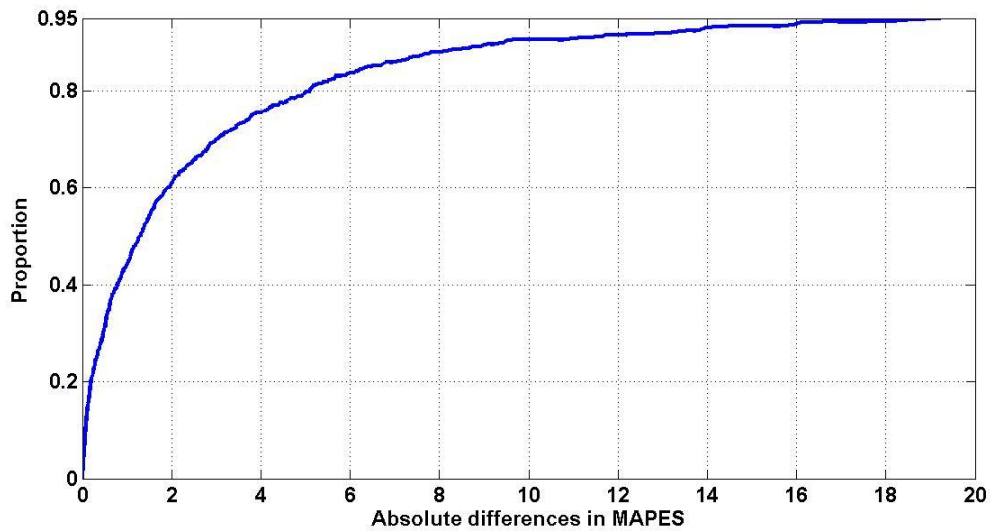
3. Comparison of Damped Trend Models

The two damped trend models were compared on the *annual* time series from the M3 competition database (Makridakis and Hibon, 2000). The final six years of values for each series were reserved for evaluating their forecasts; and the earlier part of each series was used to find minimum sum of squared error estimates of the model parameters and seed states. The one-step-ahead prediction errors needed for the evaluation of the likelihood function, for each trial set of values for the parameters and seed states, were obtained with damped trend exponential smoothing and its new restricted form. The smoothing parameters were restricted to the closed real interval [0,1]. The point predictions for the final six years from the fitted models were compared with the reserved

6

actual values using the mean absolute percentage error (MAPE) statistic. The outcomes of the study, obtained using Matlab, are summarized in Table 1. Intriguingly, the constrained form predicted better for slightly more than half of the series.

model	Wins		MAPEs	
	count	percent	mean	median
Unrestricted	312	48.4	22.66	10.92
Restricted	333	51.6	20.67	10.75

Table 1. Results for damped trends on M3 time series

Figure 1. Comparison of damped trend models: cumulative relative frequency distribution of the gaps between their MAPEs.

It might be conjectured that the differences between both models are small. The graph in Figure 1 shows the cumulative relative frequency of the magnitude of the differences in the MAPEs to determine the 'distance' between the two models. The MAPE's differed by less than two percent for sixty percent of the series; and, they differed by less than five percent for eighty percent of the series. Most of the results are remarkably close, consistent with the possibility that one might replace the conventional damped stochastic trend with its constrained analogue. There exist a minority of cases where the gap between the two models is unacceptably high; and an analysis of the data suggests that no particular model has a monopoly of wins in this particular circumstance. It is quite

7

possible that these larger gaps can be traced to computational problems caused by the distorting effects of sub-optimal solutions obtained when minimizing a multi-modal sum of squared errors function.

The exponential smoothing forecasting framework in Hyndman et al. (2002) already includes the classical damped stochastic trend model. These results indicate that this framework should be extended to include the restricted form of the damped stochastic trend model. Indeed, it might even make sense to drop the classical damped stochastic trend altogether.

4. Conclusions

The most significant contribution of this paper is in providing an explanation for why damped trend exponential smoothing has been a successful method for exponential smoothing: it adapts automatically to situations that may differ in a number of potential ways from the conditions needed for simple exponential smoothing to be the optimal form of algorithm. We have also demonstrated that a restricted form of damped trend exponential smoothing can be thought of as simple exponential smoothing with an embedded tracking signal that adapts forecasts to unanticipated structural change. Moreover, it was argued that the parameters of the damped trend could be restricted without seriously sacrificing its forecasting capacity, an advantage in the sense that it leads to a more parsimonious model.

References

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

Farnum, N.R. & Stanton, L.W. (1989). *Quantitative Forecasting Methods*, Boston:PWS-Kent.

Fildes, R., Nikolopous, K., Crone, S., & Syntetos, A. (2008). Forecasting and operational research: a review, *Journal of the Operational Research Society*, 59, 1-23.

Gardner, Jr., E.S. (2006). Exponential smoothing: the state of the art – part II, *International Journal of Forecasting*, 22, 637-666.

Gardner, Jr., E.S. & McKenzie, E. (1985). Forecasting trends in time, *Management Science*, 31, 1237-1246.

Hyndman, R. J., Koehler, A.B., Ord, J.K, & R.D. Snyder (2008). *Forecasting with Exponential Smoothing: The State Space Approach*, Berlin: Springer.

Hyndman, R. J., A. B. Koehler, R. D. Snyder, and S. Grose (2002). A state space framework for automatic forecasting using exponential smoothing methods, *International Journal of Forecasting,* 18, 439-454.

Makridakis, S. & Hibon M. (2000). The M3-competition: results, conclusions and implications, *International Journal of Forecasting,* 16, 451-476.

Roberts, S.W. (1959). Control chart tests based on geometric moving averages, *Technometrics*, 1, 239-250.