

ISSN 1440-771X
ISBN 0 7326 1085 0

Unmasking the Theta Method

Rob J. Hyndman and Baki Billah

Working Paper 5/2001

2001

**DEPARTMENT OF ECONOMETRICS
AND BUSINESS STATISTICS**

MONASH
UNIVERSITY
AUSTRALIA

Unmasking the Theta method

Rob J Hyndman¹, Baki Billah¹

13 August 2001

Abstract: The “Theta method” of forecasting performed particularly well in the M3-competition and is therefore of interest to forecast practitioners. The description of the method given by Assimakopoulos and Nikolopoulos (2000) involves several pages of algebraic manipulation and is difficult to comprehend. We show that the method can be expressed much more simply; furthermore we show that the forecasts obtained are equivalent to simple exponential smoothing with drift.

Keywords: exponential smoothing, forecasting competitions, state space models.

1 Introduction

The “Theta method” of forecasting was introduced by Assimakopoulos and Nikolopoulos (2000), hereafter referred to as A&N. Their description of the method is complicated and confusing and involves several pages of algebra. However, the method performed particularly well in the M3-competition (Makridakis & Hibon, 2000), and is therefore of interest to forecast practitioners.

We examine the Theta method and show that it can be expressed much more simply than in A&N; furthermore we show that the forecasts obtained are equivalent to simple exponential smoothing (SES) with drift. Using this equivalence, we derive appropriate prediction intervals for the method based on a state space model underlying SES with drift. Finally, we show that SES with drift can produce better forecasts than the Theta method if the parameters are optimized using a maximum likelihood approach.

Section 2 reproduces the main results from A&N using a different (and much simpler) notation. We obtain an explicit expression for point forecasts in Section 3 and show that these are equivalent to the point forecasts from SES with drift. In Section 4, we describe a state space model with equivalent forecasts, thus enabling the computation of prediction intervals and likelihood estimates. Finally, in Section 5 we compare the Theta method with fully optimized SES with drift to the annual data from the M3-competition.

¹Department of Econometrics and Business Statistics, Monash University, VIC 3800, Australia.
Correspondence to Rob.Hyndman@buseco.monash.edu.au.

2 Theta Method

Let $\{X_1, \dots, X_n\}$ denote the observed univariate time series. From this series A&N construct a new series $\{Y_1(\theta), \dots, Y_n(\theta)\}$ such that

$$Y_t''(\theta) = \theta X_t'' \quad (1)$$

where X_t'' denotes the second difference of X_t and $Y_t''(\theta)$ denotes the second difference of $Y_t(\theta)$. We note that (1) is a second-order difference equation and has the solution (see, e.g., Box, Jenkins & Reinsel, 1994, pp.120–125)

$$Y_t(\theta) = a_\theta + b_\theta(t-1) + \theta X_t \quad (2)$$

where a_θ and b_θ are constants. Thus $Y_t(\theta)$ is equivalent to a linear function of X_t with a linear trend added. A&N call $Y_t(\theta)$ a “theta line”.

For a fixed θ , A&N find the values of $Y_1(\theta)$ and $Y_2(\theta) - Y_1(\theta)$ which minimize the sum of squared differences

$$\sum_{i=1}^t [X_t - Y_t(\theta)]^2 = \sum_{i=1}^t [(1-\theta)X_t - a_\theta - b_\theta(t-1)]^2.$$

This is equivalent to minimizing the above sum of squares with respect to a_θ and b_θ . Thus, it is a simple regression of $(1-\theta)X_t$ against time $t-1$. Therefore the solution is simply

$$\begin{aligned} \hat{b}_{\theta,n} &= \frac{6(1-\theta)}{n^2-1} \left(\frac{2}{n} \sum_{t=1}^n tX_t - (n+1)\bar{X} \right) \\ \text{and} \quad \hat{a}_{\theta,n} &= (1-\theta)\bar{X} - \hat{b}_{\theta,n}(n-1)/2. \end{aligned}$$

The equivalent results are derived in more than two pages of algebra by A&N.

Note that the mean value of the new series is

$$\bar{Y}(\theta) = \hat{a}_{\theta,n} + \hat{b}_{\theta,n}(n-1)/2 + \theta\bar{X} = \bar{X},$$

the same result A&N derived in one-third of a page using their notation. Furthermore, it is easy to see that $\frac{1}{2}[Y_t(1+p) + Y_t(1-p)] = X_t$ since $\hat{a}_{1+p,n} + \hat{a}_{1-p,n} = 0$ and $\hat{b}_{1+p,n} + \hat{b}_{1-p,n} = 0$; this result takes up about half a page in A&N.

Forecasts from the Theta method are obtained by a weighted average of forecasts of $Y_t(\theta)$ for different values of θ . However, A&N only explain how to get forecasts for $\theta = 0$ and $\theta = 2$, the set-up they used in the M3-competition. In this case they define

$$\hat{X}_{n+h} = \frac{1}{2}[\hat{Y}_{n+h}(0) + \hat{Y}_{n+h}(2)]$$

where $\hat{Y}_{n+h}(0)$ is obtained by extrapolating the linear part of (2) and $\hat{Y}_{n+h}(2)$ is obtained using simple exponential smoothing on the series $\{Y_t(2)\}$. Hence,

$$\hat{Y}_{n+h}(0) = \hat{a}_{0,n} + \hat{b}_{0,n}(n+h-1) \quad (3)$$

and (see Makridakis, Wheelwright & Hyndman, 1998, p.149)

$$\hat{Y}_{n+h}(2) = \alpha \sum_{i=0}^{n-1} (1-\alpha)^i Y_{n-i}(2) + (1-\alpha)^n Y_1(2) \quad (4)$$

where α is the smoothing parameter for the SES.

So far we have simply shown how A&N's results can be replicated much more easily using our notation. The rest of our paper gives new results concerning this forecasting methodology.

3 Point forecasts

The above results can be combined to obtain a simple expression for the forecasts \hat{X}_{n+h} . From (4) we obtain

$$\begin{aligned} \hat{Y}_{n+h}(2) &= \alpha \sum_{i=0}^{n-1} (1-\alpha)^i \left[\hat{a}_{2,n} + \hat{b}_{2,n}(n-i-1) + 2X_{n-i} \right] + (1-\alpha)^n (\hat{a}_{2,n} + 2X_1) \\ &= \hat{a}_{2,n} + \hat{b}_{2,n} \left[n - \frac{1}{\alpha} + \frac{(1-\alpha)^n}{\alpha} \right] + 2\tilde{X}_{n+h} \end{aligned}$$

where \tilde{X}_{n+h} is the SES forecast of the series $\{X_t\}$. Noting that $\hat{a}_{2,n} = -\hat{a}_{0,n}$ and $\hat{b}_{2,n} = -\hat{b}_{0,n}$, we obtain

$$\hat{X}_{n+h} = \tilde{X}_{n+h} + \frac{1}{2}\hat{b}_{0,n} \left(h - 1 + \frac{1}{\alpha} - \frac{(1-\alpha)^n}{\alpha} \right). \quad (5)$$

For large n , this can be written as

$$\hat{X}_{n+h} = \tilde{X}_{n+h} + \frac{1}{2}\hat{b}_{0,n}(h - 1 + 1/\alpha).$$

Thus it is SES with an added trend where the slope of the trend is half that of the fitted trend line through the original time series.

4 Underlying stochastic models

A&N do not give an underlying stochastic model for their forecasting method. However, it is possible to find such a model using a state space approach. Let $X_1 = \ell_1$ be fixed and for $t = 2, 3, \dots$,

$$X_t = \ell_{t-1} + b + \varepsilon_t \quad (6)$$

$$\text{and } \ell_t = \ell_{t-1} + b + \alpha\varepsilon_t \quad (7)$$

where $\{\varepsilon_t\}$ is Gaussian white noise with mean zero and variance σ^2 .

Then X_t follows a state space model which gives forecast equivalent to SES with drift. This is a special case of Holt's method with the smoothing parameter for the slope set to zero. Note that X_t can also be written as

$$X_t = X_{t-1} + b + (\alpha - 1)\varepsilon_{t-1} + \varepsilon_t,$$

that is an ARIMA(0,1,1) process with drift (see Box, Jenkins & Reinsel, 1994, pp.125–126).

Now, the point forecasts for $\{X_t\}$ are given by (see, e.g., Hyndman et al., 2001)

$$\hat{X}_{n+h} = \ell_n + hb.$$

Further, we note that

$$\hat{X}_{t+1} | X_1, \dots, X_t = X_t + b + (\alpha - 1)\varepsilon_t \quad (8)$$

$$\text{and that } \varepsilon_t = X_t - X_{t-1} - b + (1 - \alpha)\varepsilon_{t-1}. \quad (9)$$

By repeatedly substituting (9) into (8), we obtain

$$\begin{aligned} \hat{X}_{n+1} | X_1, \dots, X_n &= \alpha \sum_{i=0}^{n-1} (1 - \alpha)^i X_{n-i} + (1 - \alpha)^n X_1 + \frac{b}{\alpha} [1 - (1 - \alpha)^n] \\ &= \tilde{X}_{n+1} + \frac{b}{\alpha} [1 - (1 - \alpha)^n] \end{aligned}$$

where \tilde{X}_{n+1} is the SES forecast. Similarly, the h -step ahead forecast is:

$$\hat{X}_{n+h} | X_1, \dots, X_n = \tilde{X}_{n+1} + b \left[h - 1 + \frac{1}{\alpha} - \frac{(1-\alpha)^{n+1}}{\alpha} \right].$$

Thus, we obtain identical point forecasts as for the Theta method (5) if $b = \hat{b}_{0,n}/2$.

Furthermore, our state space approach enables us to obtain maximum likelihood estimates of b and it provides prediction intervals. For example, 95% prediction intervals for h period ahead forecasts are given by

$$\hat{X}_{t+h} \pm 1.96\sigma \sqrt{(h-1)\alpha^2 + 1}.$$

(Equivalent results are obtained using the ARIMA(0,1,1) model.)

5 Application to Annual M3 Competition Data

The preceding analysis suggests we may be able to obtain better forecasts if we optimize the value of b rather than setting it equal to $\hat{b}_{0,n}/2$. To evaluate this idea, we apply the model to the 645 annual series from the M3 competition (Makridakis and Hibon, 2000). We computed forecasts up to 6 steps ahead and then we computed the symmetric mean absolute percentage error (SMAPE) as in Makridakis and Hibon (2000).

In these comparisons, we use a different initialization to that described above. Rather than fixing ℓ_1 , we fix ℓ_0 and then optimize the likelihood of the state space model (6)

Table 1: *Average SMAPE for the annual M3 competition data*

Methods	Forecasting Horizons						Average	
	1	2	3	4	5	6	1 to 4	1 to 6
(1) A&N Theta method	8.0	12.2	16.7	19.2	21.7	23.6	14.02	16.90
(2) Recalculated Theta method	8.2	12.3	16.4	18.6	21.2	23.0	13.89	16.62
(3) SES with drift	7.9	12.1	18.6	17.2	20.6	22.9	13.95	16.55

and (7) over the parameters ℓ_0 , b and α as described in Ord, Koehler and Snyder (1997). To initialize the optimization, we use the same procedure as outlined in Hyndman et al. (2001). The value of α was constrained to lie between 0.1 and 0.99. The results are presented in Table 1.

Table 1 shows the average SMAPE for: (1) the original A&N forecasts (as given in Makridakis & Hibon, 2000); (2) our forecasts using the Theta method described in A&N; and (3) forecasts based on the state space model (6) and (7). Our results differ slightly from those of A&N, probably because we initialized the SES differently, and possibly also because we estimated the smoothing parameter α differently (we used a likelihood approach).

Note that the state space model performs better than the Theta method for all forecasting horizons except $h = 3$. This is because the state space model performs relatively badly for series N0529, particularly for forecasting horizon 3. If series N0529 is omitted, the state space SMAPE for $h = 3$ becomes 16.3.

6 Conclusion

We have demonstrated that the Theta method proposed by A&N is simply a special case of SES with drift where the drift parameter is half the slope of the linear trend fitted to the data. We have also demonstrated that prediction intervals and likelihood-based estimation of the parameters can be obtained using a state space model.

References

- ASSIMAKOPOULOS, V. and NIKOLOPOULOS, K. (2000) The theta model: a decomposition approach to forecasting. *International Journal of Forecasting* **16**, 521–530.
- BOX, G.E.P., JENKINS, G.M., and REINSEL, G.C. (1994) *Time series analysis: forecasting and control*, 3rd ed., Prentice Hall: New Jersey.
- HYNDMAN, R.J., KOEHLER, A.B., SNYDER, R.D., and GROSE, S. (2001) A state space framework for automatic forecasting using exponential smoothing methods, *International Journal of Forecasting*, to appear.
- MAKRIDAKIS, S., AND HIBON, M. (2000) The M3-Competitions: results, conclusions and implications, *International Journal of Forecasting*, **16**, 451–476.
- MAKRIDAKIS, S., WHEELWRIGHT, S.C. and HYNDMAN, R.J. (1998) *Forecasting Methods and Applications*, 3rd edition, John Wiley & Sons: New York.
- ORD, J.K., A.B. KOEHLER and R.D. SNYDER (1997) Estimation and prediction for a class of dynamic nonlinear statistical models, *J. Amer. Statist. Assoc.*, **92**, 1621–1629.