International Journal of Forecasting 32 (2016) 1151–1161

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

# Models for optimising the theta method and their relationship to state space models

Jose A. Fiorucci[a,b], Tiago R. Pellegrini[c], Francisco Louzada[d], Fotios Petropoulos[e,*], Anne B. Koehler[f]

[a] Department of Statistics, Federal University of São Carlos, Brazil
[b] Department of Statistics, University of Connecticut, USA
[c] Department of Mathematics and Statistics, University of New Brunswick, Canada
[d] Department of Applied Mathematics and Statistics, University of São Paulo, Brazil
[e] Cardiff Business School, Cardiff University, UK
[f] Professor Emerita, Miami University, OH, USA

ARTICLE INFO

*Keywords:*
Time series forecasting
Theta method
State Space Models
M3-Competition
Combination

ABSTRACT

Accurate and robust forecasting methods for univariate time series are very important when the objective is to produce estimates for large numbers of time series. In this context, the Theta method's performance in the M3-Competition caught researchers' attention. The Theta method, as implemented in the monthly subset of the M3-Competition, decomposes the seasonally adjusted data into two "theta lines". The first theta line removes the curvature of the data in order to estimate the long-term trend component. The second theta line doubles the local curvatures of the series so as to approximate the short-term behaviour. We provide generalisations of the Theta method. The proposed Dynamic Optimised Theta Model is a state space model that selects the best short-term theta line optimally and revises the long-term theta line dynamically. The superior performance of this model is demonstrated through an empirical application. We relate special cases of this model to state space models for simple exponential smoothing with a drift.



## 1. Introduction

The development of accurate, robust and reliable forecasting methods for univariate time series is very important when large numbers of time series are involved in the modelling and forecasting process. In industrial settings, it is very common to work with large lines of products; thus, efficient sales and operational planning (S&OP) depend heavily on accurate forecasting methods.

Despite the advantages of automatic model selection algorithms (Hyndman & Khandakar, 2008; Hyndman, Koehler, Snyder, & Grose, 2002; Poler & Mula, 2011), there is still a need for accurate extrapolation methods. Forecasting competitions have played an important role in moving toward the forecasting of large numbers of time series, with the objective of identifying high-performing methods. The Theta method attracted the attention of researchers by its simplicity and surprisingly good performance (Koning, Franses, Hibon, & Stekler, 2005; Makridakis & Hibon, 2000), and has been one of the benchmarks in more recent forecasting competitions (Athanasopoulos, Hyndman, Song, & Wu, 2011).

* Correspondence to: Logistics and Operations Management Section, Cardiff Business School, Cardiff University, Cardiff, CF10 3EU, UK.
*E-mail address:* petropoulosf@cardiff.ac.uk (F. Petropoulos).

http://dx.doi.org/10.1016/j.ijforecast.2016.02.005
0169-2070/

The Theta method (Assimakopoulos & Nikolopoulos, 2000, hereafter A&N) is applied to non-seasonal or deseasonalised time series, where the deseasonalisation is usually performed via the multiplicative classical decomposition. The method decomposes the original time series into two new lines through the so-called theta coefficients, denoted by $\theta_1$ and $\theta_2$ for $\theta_1, \theta_2 \in \mathbb{R}$, which are applied to the second difference of the data. The second differences are reduced when $\theta < 1$, resulting in a better approximation of the long-term behaviour of the series (Assimakopoulos, 1995). If $\theta$ is equal to zero, the new line is a straight line. When $\theta > 1$, the local curvatures are increased, magnifying the short-term movements of the time series (A&N). The new lines produced are called theta lines, denoted here by $Z(\theta_1)$ and $Z(\theta_2)$. These lines have the same mean value and slope as the original data, but the local curvatures are either filtered out or enhanced, depending on the value of the $\theta$ coefficient.

In other words, the decomposition process has the advantage of exploiting information in the data that usually cannot be captured and modelled completely through the extrapolation of the original time series. The theta lines can be regarded as new time series and are extrapolated separately using an appropriate forecasting method. Once the extrapolation of each theta line has been completed, recomposition takes place through a combination scheme in order to calculate the point forecasts of the original time series. Combining has long been considered as a useful practice in the forecasting literature (for example, Clemen, 1989; Makridakis & Winkler, 1983; Petropoulos, Makridakis, Assimakopoulos, & Nikolopoulos, 2014), and therefore its application to the Theta method is expected to result in more accurate and robust forecasts.

The Theta method is quite versatile in terms of choosing the number of theta lines, the theta coefficients and the extrapolation methods, and combining these to obtain robust forecasts. However, A&N proposed a simplified version involving the use of only two theta lines with prefixed $\theta$ coefficients that are extrapolated over time using a linear regression (LR) model for the theta line with $\theta_1 = 0$ and simple exponential smoothing (SES) for the theta line with $\theta_2 = 2$. The final forecasts are produced by combining the forecasts of the two theta lines with equal weights. In the M3-Competition, this simplified version of the Theta method was applied only to the monthly time series (Nikolopoulos, Assimakopoulos, Bougioukos, Litsa, & Petropoulos, 2011).

The performance of the Theta method has also been confirmed by other empirical studies (for example Nikolopoulos, Thomakos, Petropoulos, Litsa, & Assimakopoulos, 2012; Petropoulos & Nikolopoulos, 2013). Moreover, Hyndman and Billah (2003), hereafter H&B, showed that the simple exponential smoothing with drift model (SES-d) is a statistical model for the simplified version of the Theta method. More recently, Thomakos and Nikolopoulos (2014) provided additional theoretical insights, while Thomakos and Nikolopoulos (2015) derived new theoretical formulations for the application of the method to multivariate time series, and investigated the conditions under which the bivariate Theta method is expected to forecast better than the univariate one. Despite these advances, we believe that the Theta method deserves more attention from the forecasting community, given its simplicity and superior forecasting performance.

One key aspect of the Theta method is that, by definition, it is dynamic. One can choose different theta lines and combine the produced forecasts using either equal or unequal weights. However, A&N limit this important property by fixing the theta coefficients to have predefined values. Thus, the Theta method, as implemented in the M3-Competition, is limited in the sense that it focuses only on specific information in the data. On the other hand, if the selection of the appropriate theta lines had been carried out through optimisation, the method could focus on the information that is actually important.

The contributions of this work are fourfold. First, we extend the A&N method by the optimal selection of the theta line that describes the short-term movements of the series best, maintaining the long-term component. The forecasts derived from the two theta lines are combined using appropriate weights, which ensures the recomposition of the original time series. Second, we provide theoretical and practical links between the newly proposed model, the original Theta method and the SES-d model. Third, we also perform a further extension of the model that allows the regression line (the long term component) to be revised at every time period. An empirical evaluation using the M3-Competition database is undertaken in order to obtain insights into the performances of the proposed models. The results reveal improvements in the forecasting accuracy when using the model with both extensions. This model outperforms several benchmarks as well as the A&N simplified version of the Theta method. Fourth, we reproduce the results for the Theta method, as applied to the monthly data in the M3-Competition, very closely.

The paper is organised as follows. Section 2 reviews the original Theta method of A&N and its relationship with the SES-d model. Section 3 presents different models for optimising the Theta method. Section 4 presents the forecasting performances of the proposed models, compared to a list of widely used benchmarks. The evaluation includes more than 3000 time series. Lastly, Section 5 presents our final comments and directions for future research.

## 2. Theta method and SES-d

### 2.1. The original Theta method

Originally, A&N proposed the theta line as the solution of the equation

$$\nabla^2 Z_t(\theta) = \theta\, \nabla^2 Y_t, \quad t = 3, \ldots, n, \tag{1}$$

where $Y_1, \ldots, Y_n$ is the original time series (non-seasonal or deseasonalised) and $\nabla$ is the difference operator (i.e., $\nabla X_t = X_t - X_{t-1}$). The initial values of $Z_1$ and $Z_2$ are obtained by minimising $\sum_{t=1}^{n} [Y_t - Z_t(\theta)]^2$. However, an analytical solution to compute the $Z(\theta)$ was obtained by H&B, which is given by

$$Z_t(\theta) = \theta Y_t + (1 - \theta)(A_n + B_n t), \quad t = 1, \ldots, n, \tag{2}$$