International Journal of Forecasting 30 (2014) 291–302

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Improving forecasting by estimating time series structural components across multiple frequencies

Nikolaos Kourentzes [a,*], Fotios Petropoulos [a], Juan R. Trapero [b]

[a] Lancaster University Management School, Department of Management Science, Lancaster, Lancashire, LA1 4YX, UK

[b] Universidad de Castilla-La Mancha, Departamento de Administracion de Empresas, Ciudad Real 13071, Spain

ARTICLE INFO

Keywords:
Aggregation
Combining forecasts
Exponential smoothing
M3 Competition
MAPA
ETS

ABSTRACT

Identifying the most appropriate time series model to achieve a good forecasting accuracy is a challenging task. We propose a novel algorithm that aims to mitigate the importance of model selection, while increasing the accuracy. Multiple time series are constructed from the original time series, using temporal aggregation. These derivative series highlight different aspects of the original data, as temporal aggregation helps in strengthening or attenuating the signals of different time series components. In each series, the appropriate exponential smoothing method is fitted and its respective time series components are forecast. Subsequently, the time series components from each aggregation level are combined, then used to construct the final forecast. This approach achieves a better estimation of the different time series components, through temporal aggregation, and reduces the importance of model selection through forecast combination. An empirical evaluation of the proposed framework demonstrates significant improvements in forecasting accuracy, especially for long-term forecasts.

1. Introduction

In forecasting, the selection and parameterisation of models are of critical importance, as they are tied to accurate and reliable predictions. Modern organisations have to produce large numbers of forecasts. It is desirable to identify and parameterise the appropriate model automatically for each series, in order to reduce the associated workload. This endeavour has met with mixed success in both research and practice, resulting in De Gooijer and Hyndman (2006) identifying it as an open research question. The lack of statistical expertise in organisations complicates the problem further (Hughes, 2001), and a relatively simple solution would be desirable for practical implementations.

In this paper, we propose a framework that mitigates the issue of model selection, while improving the forecasting accuracy, by taking advantage of temporal aggregation and forecast combination.

It is possible to emphasise different time series characteristics by transforming the original data to alternative time frequencies. We propose to aggregate a time series into multiple lower frequencies, i.e., a monthly time series becomes bi-monthly, quarterly and so on. At each aggregation level, different features of the time series are highlighted and attenuated (Andrawis, Atiya, & El-Shishiny, 2011). At lower aggregations (high frequency time series), periodic components such as seasonality will be prominent. As the aggregation level increases, high frequency signals are filtered out, and more importance is given to the lower frequency components, such as the level and trend of a time series. Intuitively, we expect to capture the seasonal elements of a time series better at lower aggregation levels (high frequency data). The opposite is true for the level and trend, which are highlighted at higher aggregation levels. Consequently, our motivation for such aggregation is to facilitate the identification, selection and parameter estimation of forecasting models.

* Corresponding author. Tel.: +44 1524 592911.
E-mail address: n.kourentzes@lancaster.ac.uk (N. Kourentzes).

0169-2070/$ – see front matter © 2013 International Institute of Forecasters. Published by Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ijforecast.2013.09.006

After temporal aggregation, we produce a forecast at each aggregation level, using the appropriate forecasting method. For this purpose we use the exponential smoothing family of methods, which provides a holistic framework for modelling all archetypal types of time series, whether level, trended or seasonal. The calculation of numerous forecasts at different frequencies inevitably leads to the use of combination schemes. The combination step offers further improvements in the robustness and accuracy of the resulting forecast.

Instead of combining the forecasts from the different aggregation levels directly, we propose to first break down the forecasts to their time series components, then combine these. This is motivated by potential differences in the ways in which the components appear or are transformed at different temporal aggregation levels. For example, if a seasonal monthly time series is aggregated into an annual one, it will have no seasonality. Simply combining monthly and annual forecasts would halve the seasonal signal. Therefore, the combination of the seasonal information must only be done for the aggregation levels at which it exists, whilst for the level and trend, the combination can be performed at any level. Moreover, combining forecasts instead of time series components does not allow us to discriminate between the desirable elements of each aggregation level that is to be combined. Further advantages of combining components instead of forecasts are discussed in the description of the proposed algorithm. Exponential smoothing provides such component information directly (Gardner, 2006). The combined components are subsequently joined in order to produce forecasts of the original time series.

The key advantages of the proposed framework are: (i) it lessens the importance of model selection, a principal problem in time series modelling; (ii) it takes advantage of gains in forecasting accuracy from both temporal aggregation and forecast combination, which have been demonstrated separately in the literature; (iii) it makes use of exponential smoothing, a widely established and well researched forecasting method that is implemented in several forecasting support systems; and (iv) it is simple, thus allowing transparency and direct implementation in organisations.

We conduct a large empirical evaluation of the proposed framework against exponential smoothing. We also use the direct combination of forecasts produced at different aggregation levels as a benchmark, in order to demonstrate the benefits of considering the time series components separately. We test the robustness of the proposed *Multi Aggregation Prediction Algorithm* (MAPA) by evaluating its performance, in terms of accuracy and forecasting bias, across datasets with different sampling frequencies. We find that MAPA has a promising forecasting performance, with major improvements for long-term forecasts.

The rest of the paper is organized as follows: in the next section we discuss the benefits of temporal aggregation and model combination in the literature, providing further motivation for our work. Section 3 describes the proposed MAPA, while Section 4 outlines the experimental setup and presents the evaluation results, followed by a discussion in Section 5. A further refinement of the model is investigated in Section 6, and concluding remarks are given in Section 7.

2. Background research

2.1. Aggregation and forecasting

There are two types of aggregation in the forecasting literature (Babai, Ali, & Nikolopoulos, 2012). On the one hand, temporal aggregation refers to time (non-overlapping) aggregation for a specific time series. This converts a high frequency time series into a lower frequency time series, according to the selected aggregation level. For example, we can derive the annual demand of a monthly series by setting the aggregation level equal to 12 periods. On the other hand, cross-sectional or hierarchical aggregation refers to the demand aggregation (bottom-up) of multiple items or stock keeping units (SKUs) for the formation of families of products.

Temporal aggregation has been investigated in the context of both ARIMA models and GARCH processes. For a thorough overview of the effects of temporal aggregation on time series models, the reader is encouraged to refer to the study by Silvestrini and Veredas (2008). Even though it has no theoretical support (Wei, 1990), aggregation has been shown empirically to work remarkably well for ARIMA models, in terms of both forecasting accuracy (for example, see Abraham, 1982; Amemiya & Wu, 1972) and bias reduction (Mohammadipour & Boylan, 2012; Souza & Smith, 2004).

The consequences of temporal aggregation and systematic sampling, in terms of lag lengths (Brewer, 1973) and dynamic relationships between variables (Weiss, 1984), have also been investigated, giving useful theoretical insights into the stability of ARMA and ARMAX models. Moreover, Drost and Nijman (1993) showed that, in the case of univariate GARCH models, the variance parameters of the low frequency model generally depend on the mean, variance, and kurtosis parameters of the high frequency model. Hafner (2008) derived results for the temporal aggregation of multivariate GARCH(1,1) processes, and concluded that the dynamics of the aggregated processes can be acquired in a very simple way. As a result, plenty of empirical and theoretical work has been done on ARIMA models and GARCH processes in regard to temporal aggregation. On the other hand, the literature on the effects of temporal aggregation with exponential smoothing models is much more limited, even though the method itself is applied widely in practice (Gardner, 2006).

Spithourakis, Petropoulos, Babai, Nikolopoulos, and Assimakopoulos (2011) examined the efficiency of a temporal aggregation framework for widely used forecasting techniques. The results, based on the monthly data from the M3-Competition (Makridakis & Hibon, 2000), demonstrated significantly improved forecasting accuracies for Naïve, SES and Theta (Assimakopoulos & Nikolopoulos, 2000). Athanasopoulos, Hyndman, Song, and Wu (2011) showed that, in the tourism industry, aggregated forecasts (produced from high frequency data) were more accurate than forecasts produced from low frequency data directly. They were the first to investigate the impact of temporal aggregation on the performance of the exponential smoothing family of methods empirically, though their study was limited to aggregating monthly and quarterly data to yearly. Their forecasting evaluation was performed