European Journal of Operational Research 263 (2017) 412–418

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

Integrated hierarchical forecasting

Clint L. P. Pennings*, Jan van Dalen

Rotterdam School of Management, Erasmus University, T9-08, Burgemeester Oudlaan 50, Rotterdam 3062 PA, The Netherlands

ARTICLE INFO

Article history:
Received 8 December 2015
Accepted 18 April 2017
Available online 27 April 2017

Keywords:
Forecasting
Hierarchical
Top-down
Bottom-up
Decision-making

ABSTRACT

Forecasts are often made at various levels of aggregation of individual products, which combine into groups at higher hierarchical levels. We provide an alternative to the traditional discussion of bottom-up versus top-down forecasting by examining how the hierarchy of products can be exploited when forecasts are generated. Instead of selecting series from parts of the hierarchy for forecasting, we explore the possibility of using all the series. Moreover, instead of using the hierarchy after the initial forecasts are generated, we consider the hierarchical structure as a defining feature of the data-generating process and use it to instantaneously generate forecasts for all levels of the hierarchy. This integrated approach uses a state space model and the Kalman filter to explicitly incorporate product dependencies, such as complementarity of products and product substitution, which are otherwise ignored. An empirical study shows the substantial gain in forecast and inventory performance of generalizing the bottom-up and top-down forecast approaches to an integrated approach. The integrated approach is applicable to hierarchical forecasting in general, and extends beyond the current application of demand forecasting for manufacturers.

1. Introduction

For organizations, demand forecasting is essential as it drives production, inventory and planning decisions. Supply has to match demand as well as possible to avoid excess inventory and stockouts. Large manufacturers often have SKUs ranging in the thousands, spanning several product categories, each of which requires forecasts. Several decision makers are involved from operations, marketing, sales and finance, who require forecasts at various levels of aggregation. Forecasts are more easily discussed at an aggregated product level, but for production these forecasts have to be available at the SKU level.

SKUs naturally group together in hierarchies with individual sales per product at the base line, followed by several intermediary levels, denoting sales for groups of related products at increasingly general aggregation levels, such as product groups and categories, and with a top level that lists total sales. Two commonly used approaches in practice and research start from opposite ends of this hierarchy to generate forecasts for all series: bottom-up forecasting and top-down forecasting (Widiarta, Viswanathan, & Piplani, 2009). In bottom-up forecasting, base forecasts are generated for product demand at the lowest level in the hierarchy (Gordon, Morris, & Dangerfield, 1997). Subsequently, these are aggregated to determine forecasts at higher hierarchical levels. Bottom-up forecasting is commonly contrasted with top-down forecasting, in which aggregated demand forecasts are disaggregated downwards to determine forecasts at lower levels in the hierarchy (Kahn, 1998). Research stretches over three decades with mixed results as to a preference for either bottom-up or top-down forecast approaches (Syntetos, Babai, Boylan, Kolassa, & Nikolopoulos, 2016).

Both approaches generate forecasts for a selected part of the hierarchy, aggregated upwards or allocated downwards to obtain forecasts for the remaining series. This aggregation and allocation imply a potential loss of information, as the ignored series can only be recovered under stringent conditions. The loss of information is exacerbated as the selected series are forecasted separately. Product dependencies, such as complementarity of products and product substitution, are explicitly ignored. Yet product dependencies motivate combining similar products in groups and the existence of hierarchies.

Hyndman, Ahmed, Athanasopoulos, and Shang (2011) introduce a combination approach that uses forecasts of all series in the hierarchy. By taking a linear combination of the bottom-up and top-down forecasts at various hierarchical levels, their approach offers an ensemble of the bottom-up and top-down approaches. The combination entails a post-hoc revision of forecasts to ensure that forecasts add up consistently throughout the hierarchy. More forecasts are involved than in either the bottom-up and top-down approaches alone, but the initial forecasts are still generated independently.

* Corresponding author.
E-mail addresses: cpennings@rsm.nl (C.L.P. Pennings), jdalen@rsm.nl (J. van Dalen).

http://dx.doi.org/10.1016/j.ejor.2017.04.047

The bottom-up, top-down and combination approaches use the hierarchy of products only after the initial forecasts are generated. By incorporating the hierarchical structure at an earlier stage, i.e., during the generation of forecasts, we introduce an integrated approach, superseding the traditional discussion of bottom-up versus top-down forecasting. This has at least two advantages. First, instead of selecting isolated series for forecasting, all the available data in the hierarchy can be used. Second, product dependencies can be explicitly incorporated, such as complementarity of products and product substitution, while these are otherwise ignored.

An empirical application evaluates the forecasting approaches for one of the largest manufacturers of consumer products, which has hundreds of brands spanning fourteen categories of food products, home and personal care. The empirical study shows a substantial gain in forecast and inventory performance of generalizing the bottom-up and top-down forecast approaches to an integrated approach.

The remainder of this paper is organized as follows. In Section 2, we present an overview of the relevant literature on hierarchical forecasting and the bottom-up, top-down, and combination approaches for forecasting. We especially focus on the use of the hierarchical structure, product dependencies and demand heteroscedasticity, and we critically evaluate several approaches. In Section 3, we introduce multiple state space models that are used as an integrated approach for hierarchical forecasting and outline the empirical study. For the empirical study, we compare approaches in terms of forecasting and inventory performance and use the company's own forecast as a benchmark. Section 4 lists the results and their implications, while Section 5 concludes and gives suggestions for future research.

2. Theoretical background

Hierarchical forecasting has different forms pertaining to temporal and contemporaneous aspects. Here, we exclusively focus on contemporaneous hierarchies, specifically on products aggregated in groups and categories. This section summarizes the relevant theoretical background on hierarchical forecasting and the approaches of bottom-up, top-down, and the combination approach of Hyndman et al. (2011) for forecasting. We especially focus on the use of the hierarchical structure, product dependencies and heteroscedasticity in product demand, and critically evaluate approaches.

Over three decades of forecasting literature show mixed results as to a preference for either top-down or bottom-up forecasting (Syntetos et al., 2016). This is not surprising as the performance of the approaches depends on the underlying demand process of products (Lütkepohl, 1984). Due to the additive nature of the hierarchy, in which sums of product sales determine group sales, which, in turn, add up to determine category sales, the underlying demand process is transformed at various levels of the hierarchy. Aggregation can lead to substantial information loss, which makes bottom-up forecasting seem favorable (e.g., Edwards & Orcutt, 1969; Orcutt, Watts, & Edwards, 1968; Zellner, 1969). However, if no important information is lost, benefits can be gained if random noise cancels out (Fliedner, 1999), which makes top-down forecasting seem more favorable. A wide variety of performance is seen as the nature and extent of differences between top-down and bottom-up are dependent upon context and the assumed demand processes (Wei & Abraham, 1981). Examples show that conclusions may revolve around differences in demand processes or parameter settings (Widiarta, Viswanathan, & Piplani, 2007; 2009).

Dependencies between the demand for different products are a key characteristic of the demand process, and hence a main driver of differences in performance between top-down and bottom-up approaches (Chen & Boylan, 2009; Kohn, 1982; Schwarzkopf, Tersine, & Morris, 1988; Tiao & Guttman, 1980). A particular type of demand dependency does not unequivocally make either bottom-up or top-down more favorable (Fliedner & Mabert, 1992; Fliedner, 2001; Sohn & Lim, 2007). Stronger negative cross-correlations between individual demand series lead to less variation at an aggregate level, but imply differences between individual product sales. In contrast, stronger positive correlations between individual demand series lead to more variable aggregate sales, but imply that differences at the individual product level are smaller.

This explains why empirical studies are unable to consistently show one approach outperforming the other. Dangerfield and Morris (1992) compare bottom-up and top-down approaches on empirical data and conclude that bottom-up forecasting is more accurate, especially when products are highly correlated. By contrast, Fliedner (1999) concludes that stronger positive and negative correlations improve the forecast at the aggregate level to such an extent that the top-down approach is more accurate.

An important difference between the bottom-up and top-down approaches is that the latter requires additional measures to allocate an aggregate forecast downwards to lower levels in the hierarchy. Gross and Sohl (1990) compare various ways of determining allocation proportions. A common allocation is based on averaging historical sales proportions, where the unweighted proportion p_j for each product j is determined as its sales y_j relative to the total sales in the product category y over time period T.

$$p_j = \frac{1}{T}\sum_{t=1}^{T}\frac{y_{j,t}}{y_t} \tag{1}$$

A common alternative is based on a single, total proportion observed over all time periods, leading to a weighted allocation:

$$p_j = \sum_{t=1}^{T} y_{j,t} \Big/ \sum_{t=1}^{T} y_t \tag{2}$$

Both allocations perform well in practice (Gross & Sohl, 1990).

The two approaches of top-down and bottom-up can also be combined at intermediary levels in the hierarchy, known as the middle-out approach. Forecasts are generated at a particular level and then aggregated upwards using the bottom-up approach, and allocated downwards using a top-down approach.

Recently, Athanasopoulos, Ahmed, and Hyndman (2009) and Hyndman et al. (2011) introduced a different approach, labeled the combination approach, which uses the hierarchical structure to create revised forecasts. This forecasting approach follows two steps: (1) generate independent forecasts for each series in the hierarchy, (2) weight these forecasts according to the hierarchical structure to determine the final forecasts. These final forecasts adhere to the hierarchical structure in the sense that aggregates of the forecasts at the bottom level exactly match forecasts at higher levels in the hierarchy.

The combination approach proposed by Hyndman et al. (2011) is a continuation of earlier work on revising measurements of macro-economic indicators (e.g., Byron, 1978; Solomou & Weale, 1991; 1993; 1996; Stone, Champernowne, & Meade, 1942; Weale, 1985; 1988). A salient difference is that Hyndman et al. (2011) have underlying time series of sales available for each forecast. We introduce notation for hierarchical series to discuss the combination approach, focusing on sales without loss of generality. We have a large vector $\boldsymbol{y}_t$ which contains the n sales series at all levels of the hierarchy. Sales at higher levels are determined by aggregating sales of m products at the lowest level $\boldsymbol{b}_t$. $\boldsymbol{y}_t$ is an $n \times 1$ matrix determined by linear combinations of the $m \times 1$ vector $\boldsymbol{b}_t$ containing sales at the base product level, using an $n \times m$ design matrix $\boldsymbol{S}$ to link