Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/79860/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Rostami-Tabar, Bahman ORCID: <https://orcid.org/0000-0002-3730-0045>, Babai, Mohamed Zied, Ducq, Yves and Syntetos, Argyrios ORCID: <https://orcid.org/0000-0003-4639-0756> 2015. Non-stationary demand forecasting by cross-sectional aggregation. International Journal of Production Economics 170 (Part A) , pp. 297-309. 10.1016/j.ijpe.2015.10.001 file

Publishers page: <http://dx.doi.org/10.1016/j.ijpe.2015.10.001>
<<http://dx.doi.org/10.1016/j.ijpe.2015.10.001>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.

See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.

Non-stationary Demand Forecasting by Cross-
Sectional Aggregation

Bahman Rostami-Tabar, Mohamed Zied Babai,
Yves Ducq, Aris Syntetos

PII: S0925-5273(15)00381-3
DOI: <http://dx.doi.org/10.1016/j.ijpe.2015.10.001>
Reference: PROECO6250

To appear in: *Intern. Journal of Production Economics*

Received date: 3 July 2013
Revised date: 23 September 2015
Accepted date: 2 October 2015

Cite this article as: Bahman Rostami-Tabar, Mohamed Zied Babai, Yves Ducc and Aris Syntetos, Non-stationary Demand Forecasting by Cross-Sectional Aggregation, *Intern. Journal of Production Economics* <http://dx.doi.org/10.1016/j.ijpe.2015.10.001>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting galley proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Non-stationary Demand Forecasting by Cross-Sectional Aggregation

Bahman Rostami-Tabar^a, Mohamed Zied Babai^{b,*}, Yves Ducq^c, Aris Syntetos^d

^a *Coventry University, School of Strategy and Leadership, Priory Street, Coventry, CV1 5FB, UK*

^b *Kedge Business School, 680 cours de la libération, 33400 Talence, France*

^c *Univ. Bordeaux, IMS, UMR 5218, 33400 Talence, France*

^d *Cardiff Business School, Cardiff University, Aberconway Building, Cardiff CF10 3EU, UK*

Abstract:

In this paper the relative effectiveness of top-down (TD) versus bottom-up (BU) approaches is compared for cross-sectionally forecasting aggregate and sub-aggregate demand. We assume that the sub-aggregate demand follows a non-stationary Integrated Moving Average (IMA) process of order one and a Single Exponential Smoothing (SES) procedure is used to extrapolate future requirements. Such demand processes are often encountered in practice and SES is one of the standard estimators used in industry (in addition to being the optimal estimator for an IMA process). Theoretical variances of forecast error are derived for the BU and TD approach in order to contrast the relevant forecasting performances. The theoretical analysis is supported by an extensive numerical investigation at both the aggregate and sub-aggregate level, in addition to empirically validating our findings on a real dataset from a European superstore. The results demonstrate the increased benefit resulting from cross-sectional forecasting in a non-stationary environment than in a stationary one. Valuable insights are offered to demand planners and the paper closes with an agenda for further research in this area.

Keywords: Demand Forecasting; Cross-sectional aggregation; Non-Stationary Processes; Single Exponential Smoothing

* Corresponding author: Mohamed Zied BABAI; e-mail: mohamed-zied.babai@kedgebs.com

E-mail addresses: ac1055@coventry.ac.uk (Bahman ROSTAMI-TABAR), SyntetosA@cardiff.ac.uk (Aris A. SYNTETOS), yves.ducq@ims-bordeaux.fr (Yves DUCQ).

1. INTRODUCTION

Demand forecasting is the starting point for most planning and control organizational activities. A considerable part of the forecasting literature has been dedicated to strategies and methods for single time series, but in reality there are often many related time series that can be organized hierarchically and aggregated at several different levels in groups based on products, customers, geography or other features. The hierarchical level at which forecasting is performed then it will depend on the function the forecasts are fed into. With regards to products (or Stock Keeping Units - SKUs) in particular, forecasting at the individual SKU level is required for inventory control, product family forecasts may be required for Master Production Scheduling, forecasts across a group of items ordered from the same supplier may be required for the purpose of consolidating orders, forecasts across the items sold to a specific large customer may determine transportation and routing decisions etc.

One intuitively appealing approach to obtain higher level forecasts is by *cross-sectional (also referred to as hierarchical) aggregation*, which involves aggregating different items (i.e. aggregating the requirements for different items usually in one specific time period) to reduce variability. Existing approaches to cross-sectional forecasting usually involve either a bottom-up (BU) or a top-down (TD) approach (or a combination of the two). When forecasting at the aggregate level is of interest, the former involves the aggregation of individual SKU forecasts to the group level whereas the latter relates to forecasting directly at the group level (i.e. first aggregate requirements and then extrapolate directly at the aggregate level). When the emphasis is on forecasting at the sub-aggregate level, then bottom-up relates to direct extrapolation at the sub-aggregate level whereas top-down involves the disaggregation of the forecasts produced directly at the group level (Gross and Sohl, 1990; Widiarta et al., 2007). An important issue that has attracted the attention of many researchers as well as practitioners over the last few decades is the (relative) effectiveness of such cross-sectional forecasting approaches.

TD and BU are extremely useful towards improving the accuracy of forecasts and plans when leveraged within an S&OP (Sales and Operations Planning) process (Lapide, 2006). The S&OP is a multi-functional process that involves managers from all departments (Sales, Customer Service, Supply Chain, Marketing, Manufacturing, Logistics, Procurement and Finance), where each department requires different levels of demand forecasts (Lapide, 2004). For example, in marketing, forecasting of revenues by product groups and brands is needed, sales departments deal with sales forecasts by customer accounts and/or sales channels,

supply chain managers request SKU level forecasts, while finance requires forecasts that are aggregated into budgetary units in terms of revenues and costs (Bozos and Nikolopoulos, 2011).

In this paper, we study analytically the relative effectiveness of the BU and TD approach when the underlying series follows a non-stationary Integrated Moving Average process of order one, $ARIMA(0,1,1)$, and the forecasting method is the Single Exponential Smoothing (SES) which is the optimal estimator for the process under consideration (Box et al., 2008). Both assumptions bear a significant degree of realism. There is evidence to support the fact that demand often follows non-stationary processes (please refer also to subsection 2.1). Moreover, SES is a very popular forecasting method in industry (Acar and Gardner, 2012; Gardner, 1990, 2006; Taylor, 2003). In terms of the practical relevance of our research we refer to a set of SKUs where a large proportion of them follow an $ARIMA(0,1,1)$ process; this is not an untypical scenario as demonstrated by analysis of empirical datasets including our own empirical investigation.

The question is whether it is appropriate to use sub-aggregate data or one should rather aggregate data to derive the individual and aggregate forecasts. In addition, we analyse the case of non-stationary processes to reveal whether there is an increased benefit resulting from cross-sectional forecasting when departing from the stationarity assumption. To do so we compare the variance of the forecast error obtained based on the aggregate demand (V_{TD}) to that of the sub-aggregate demand (V_{BU}). Comparisons are performed at both the aggregate and sub-aggregate level, in the former case using both a theoretical and a numerical analysis while in the latter case only by means of a numerical simulation (since the mathematical results in that case are intractable). Our analysis is consistent with the fact that companies are often using both levels of forecasting to support different decision-making processes. In addition, it renders the comparison between the two approaches a more fair exercise since one might expect that BU provides more accurate forecasts at the sub-aggregate level and TD works better at the aggregate level (Zotteri et al., 2005).

We mathematically show that the ratio of the variance of forecast error of the top-down to that of the bottom-up approach is equal to one for identical process parameters when compared at the aggregate level. The mathematical analysis is complemented by a numerical experiment to evaluate in detail the conditions under which one approach outperforms the other. Such an experiment also allows the introduction of non-identical process parameters of the sub-aggregate series (a condition that cannot be considered mathematically) and the comparison at the sub-aggregate level. In addition, an empirical investigation is also

conducted to assess the validity of the results on real data from a European superstore. Important managerial insights are derived based on the above research and tangible suggestions are offered to practitioners dealing with inventory forecasting problems.

To the best of our knowledge, the only papers directly relevant to our work are those by Widiarta et al. (2007, 2008, 2009) and Sbrana and Silvestrini (2013). The researchers evaluated analytically at the aggregate level the effectiveness of the TD and BU approaches under the assumption of AR(1) (Auto-Regressive process of order 1), MA(1) and IMA(1,1) processes. Our additional contribution to the literature is threefold: (i) we analyse the superiority conditions of BU and TD approaches at both the aggregate and the sub-aggregate levels of forecasting by means of both analytical and simulation work, (ii) through a more detailed sensitivity analysis using simulation, we investigate the impact of all the process and control parameters on the comparative performance of the two approaches, and (iii) we analyse and validate empirically our theoretical results on real data noting that none of the previous theoretical work comparing BU and TD approaches has been validated empirically.

With regards to this last point, it is important to note that rather recently Athanasopoulos et al. (2009) and Hyndman et al. (2011) have proposed a new approach to handling hierarchical time series forecasting. This approach does not emphasise the estimator being used to extrapolate requirements (i.e. any forecasting method may actually be used) but rather the weighted contribution of the forecasts produced at all levels of a given hierarchy (to appropriately retain important information that may be available at any hierarchical level) for the purpose of producing a required forecast at a particular level. Despite the fact that this approach lacks analytical insights it has been shown to perform well in practice and thus it is further considered (in addition to the BU and TD approaches) in the empirical part of our investigation.

The remainder of our paper is structured as follows. In Section 2 we provide a review of the literature on demand aggregation related issues. In Section 3 we describe the assumptions and notations used in this study and we conduct an analytical evaluation of the variance of the forecast error related to both the BU and TD approaches, followed by a simulation study performed in Section 4. We conduct an empirical investigation in Section 5 and the paper concludes in Section 6 with the implications of our work for real world practices along with an agenda for further research in this area.

2. LITERATURE REVIEW

Demand forecasting for sales and operations management often concerns many items, perhaps hundreds of thousands, simultaneously. The conventional forecasting approach is to extrapolate the data series for each SKU individually. However, most businesses have natural groupings of SKUs; that is, the SKUs may be aggregated to get higher levels of forecasts across different dimensions such as product families, geographical areas, customer types, supplier types etc. (Chen and Boylan, 2007). Such an approach enables the potential identification of time series components such as trend or seasonality that may be hidden or not particularly prevalent at the individual SKU level. Group approaches for example are known to offer considerable benefits towards the estimation of seasonal indices (Chen and Boylan 2008). Most of the forecasting literature in this area has looked at the comparative performance of the top-down (TD) and the bottom-up (BU) approach. The findings with regards to the performance of these approaches are mixed.

Many researchers have provided evidence in favour of the TD approach. Gross and Sohl (1990) for example, numerically found that the TD approach (in conjunction with an appropriate disaggregation method) provided better estimates than BU forecasting in two out of three product lines examined. Flidner (1999) evaluated by means of simulation the forecast system performance at the aggregate level resulting from varying degrees of cross correlation between two sub-aggregate time series. The sub-aggregate items were assumed to follow a Moving Average process of order one, MA(1), and the forecasting methods considered were SES and the Simple Moving Average (SMA). This research showed the forecast performance at the aggregate level to benefit from the TD approach. Barnea and Lakonishok (1980) examined the effectiveness of BU and TD on forecasting corporate performance. They reported that positive cross-correlation contributes to the superiority of forecasts based on aggregate data (TD).

On the other hand, Orcutt et al. (1968) and Edwards and Orcutt (1969) argued that information loss is substantial when aggregating and therefore the bottom-up approach provides more accurate forecasts. Dangerfield and Morris (1992) and Gordon et al. (1997) used a subset of the M-competition^b data (Makridakis et al., 1982) to examine the performance of TD and BU approaches on sub-aggregate demand forecasting. They found that forecasts by the BU approach were more accurate in most situations especially when

^b The M Competition is an empirical forecast accuracy comparison exercise introduced by Prof. Makridakis.

items were highly correlated or when one item dominated the aggregate series. Weatherford et al. (2001) evaluated the performance of BU and TD approaches to obtain the required forecasts for hotel revenue management. The data they considered was perceived as very typical within the hotel industry. They experimented with four different approaches (fully disaggregated, aggregating by rate category only, aggregating by length of stay only, and aggregating by both rate category (i.e. the price per night) and length of stay) to get detailed forecasts by day of arrival, duration of stay and rate category. The results of their study showed that a purely sub-aggregate forecast strongly outperformed even the best aggregate forecast.

Some authors take a contingent approach and analyse the conditions under which one method produces more accurate forecasts than the other. Shlifer and Wolff (1979) evaluated analytically the superiority of BU and TD on forecasting sales for specific and entire market segments. They specified the conditions under which BU is preferred to TD and vice versa. Such superiority was found to be a function of the number of markets, market size and forecast horizon. They mentioned that BU is preferable for the purpose of forecasting the aggregate series based on their observations in real situations. However, when the comparison was performed at the sub-aggregate level, they found that BU performs better for small marker segments, while both BU and TD perform equally well for large segments. Lütkepohl (1984) evaluated the performance of BU and TD approaches for forecasting at the aggregate level by using the mean squared error. It was shown that it might be preferable to forecast aggregate time series using a BU strategy when the data generation process is known. However, if the ARIMA processes used for forecasting, including the order of process and parameters, were estimated from a given set of time series data then the TD approach outperformed BU. Widiarta et al. (2007) studied analytically the conditions under which one approach outperforms the other for forecasting the item level demands when the sub-aggregate items follow a first-order autoregressive [AR(1)] process with the same autoregressive parameter for all the items and when SES is used to extrapolate future demand requirements. They found that the superiority of each approach is a function of the autoregressive parameter. Widiarta et al. (2008, 2009) also evaluated analytically the effectiveness of TD and BU approaches at the sub-aggregate and aggregate level. They showed that when all sub-aggregate items follow an MA(1) process with identical moving average parameters, there is no difference in the relative performance of TD and BU forecasting as long as the optimal smoothing constant is used in both approaches. Subsequently, they conducted a simulation analysis considering non-identical process

parameters for sub-aggregate items and concluded that there is significant difference between the two approaches. The superiority of each approach was a function of the moving average parameter, the cross-correlation and the proportion of a sub-aggregate component's contribution to the aggregate demand. Sbrana and Silvestrini (2013) evaluated the effectiveness of BU and TD approaches when forecasting the aggregate demand using a multivariate exponential smoothing framework. They established the necessary and sufficient condition for the equality of mean squared errors (MSEs) of the two approaches. In addition, they showed that the relative forecasting accuracy of TD and BU depends on the parametric structure of the underlying framework.

In summary, both BU and TD approaches appear to be associated with more accurate forecasts depending on the level of comparison, structure of the series and cross-correlation related assumptions. It is easy to observe that most of the literature dealing with the issue of aggregation for forecasting purposes focuses only on stationary series and it does not consider the most realistic case of non-stationary processes. There is considerable evidence to suggest that inventory demand is non-stationary and thus relevant processes could be eventually assumed for representing their underlying structure. In the following subsection we provide an overview of the literature on the validity of the nonstationary assumption for real world applications.

Before we close this section it should be noted that recently Athanasopoulos et al. (2009) and Hyndman et al. (2011) have proposed a new approach, referred to as 'the optimal method', to handling hierarchical time series forecasting. As discussed in the previous section, the approach under concern is based on independently forecasting the series at all levels of a given hierarchy and then using a regression model to optimally combine and reconcile these forecasts at a particular desired forecast output level. By means of a simulation study using ARIMA type series and an empirical investigation using Australian tourism demand data, Hyndman et al. (2011) have shown that the optimal method outperforms both the TD and the BU approaches and as such we further consider it as a benchmark in the empirical part of our work.

2.1. The validity of the non-stationary demand assumption

Compared to stationary demand processes, nonstationary processes have received less attention in the academic literature (Bijvank and Vis, 2011) although there is evidence that most of the forecasting and inventory control problems occur in situations where demand is

nonstationary (and partially observed) (Treharne and Sox, 2002). Naturally this may be attributed to the fact that the nonstationarity assumption complicates the relevant analyses and limits the theoretical results that may be obtained making it very difficult to determine an optimal forecast and stocking levels (Shang, 2012).

Nonstationary demand is the rule rather than the exception in most industries nowadays. The nonstationarity may arise due to many reasons such as: (1) product life cycles with multi-stages, (2) technological innovation and reduced product life, (3) seasonal effects, (4) volatile customer preferences, (5) changes in economic conditions, (6) exchange rate fluctuations, etc. (Li et al., 2011). Companies in all markets are introducing new products at a higher frequency with increasingly shorter life cycles. For instance, in the high-tech industry, the products have relatively short life cycles and their demand patterns are generally considered as nonstationary (Chien et al., 2008; Graves and Willems, 2000, 2008; Raghunathan, 2001).

Furthermore, nonstationary demand processes have been observed in the wholesaling and retailing industry. Martel et al. (1995) argued that in the grocery distribution, because of the various promotion mechanisms such as weekly special promotions, national television advertising campaigns, etc., demand gets clearly nonstationary. Erkip et al. (1990) and Lee et al. (1997) empirically found that demands of consumer products are nonstationary and highly autocorrelated. Lee et al. (2000) used panel data to examine the weekly sales patterns of 165 SKUs at a supermarket. They found that 150 out of the 165 SKUs analysed demonstrated nonstationary behaviour with high autocorrelation. Ali et al. (2011) experimented with a demand dataset of 1798 SKUs from a major European supermarket in Germany. They found that around 30% of the SKUs follow a nonstationary process and further an 80% of them follow an ARIMA(0,1,1) process. Moreover, Mitchell and Niederhausen (2010) noted that the nature of a nonstationary demand processes is consistent with the nature of retail demand for a wide variety of merchandise including apparel, consumer electronics, toys and other holiday items, patio furniture and other summer seasonal merchandise and school supplies.

Another sector where it was reported that demand follows a nonstationary process is the tourism (Goh and Law, 2002). Finally, Tunc et al. (2011) confirmed that nonstationary stochastic demands are very common in all industrial settings associated with seasonal patterns, trends, business cycles, and limited-life items.

There is also evidence that demand may follow an ARIMA(0,1,1) process in particular (which is the process considered in this study). This process has often been found to be useful in inventory control problems and econometrics (Box et al., 2008). More generally, Mahajan

and Desai (2011) argued that retailers often face a nonstationary demand that follows an ARIMA(0,1,1) process.

In this study we compare the performance of BU and TD approaches on demand forecasting under the assumption of a nonstationary ARIMA(0,1,1) process. In the next section we analyse theoretically the forecasting effectiveness of these approaches.

3. THEORETICAL ANALYSIS

In this section we derive the variance of the forecast error associated with the TD and BU approaches. These approaches work as follows:

The top-down approach consists of the following steps: i) sub-aggregate demand items are aggregated; ii) the forecast of aggregate demand is produced by applying SES at the aggregate level, and iii) the forecast is disaggregated back to the original level by applying an appropriate disaggregation method, if a sub-aggregate forecast is needed. Various proportional approaches may be used to disaggregate the TD forecasts. The reader is referred to Gross and Sohl (1990) for more details about such approaches.

In the bottom-up approach: i) sub-aggregate demand forecasts are produced directly for the sub-aggregate items; ii) the aggregate forecast (if needed) is obtained by combining individual forecasts for each SKU, i.e. potentially a separate forecasting model is used for each item in the product family (Zotteri et al., 2005). These approaches are presented schematically in Figure 1. The presentation style follows that adopted by Mohammadipour et al. (2012).

Figure 1. Graphical representation of the TD (left) and BU (right) approach

Comparisons may be performed at both the aggregate and sub-aggregate level although in our theoretical analysis the comparisons are performed only at the former level since (analytical) results regarding the latter are intractable. However, in the simulation study that follows the theoretical analysis we relax various assumptions and we also present results for both levels of comparison.

3.1. Notation and Assumptions

We denote by:

$d_{i,t}$: Sub-aggregate item demand i in period t

$\rho_{i,j}$: Correlation between the error term of sub-aggregate item i and j (cross-correlation)

D_t : Aggregate demand in period t

$\varepsilon_{i,t-k}$: Independent random variable for sub-aggregate item demand i in period t , normally distributed with zero mean and variance σ^2

ε'_{t-k} : Independent random variable for aggregate demand in period t , normally distributed with zero mean and variance σ'^2

$f_{i,t}$: Forecast of sub-aggregate demand in period t , the forecast produced in $t-1$ for the demand in t .

F_t : Forecast of aggregate demand in period t , the forecast produce in $t-1$ for the demand in t .

α_i : Smoothing constant used in the Single Exponential Smoothing method for each sub-aggregate item i in the BU approach, $0 \leq \alpha_i \leq 1$

α_{TD} : Smoothing constant used in the Single Exponential Smoothing method for aggregate demand in TD approach, $0 \leq \alpha_{TD} \leq 1$

p_i : the relative weight of sub-aggregate item i 's contribution to the aggregate family, where $\sum_{i=1}^N p_i = 1$.

V_{BU} : Variance of Forecast Error of the BU approach

V_{TD} : Variance of Forecast Error of the TD approach

V_{OP} : Variance of Forecast Error of the optimal method

θ_i : Moving average parameter of sub-aggregate item demand i , $|\theta_i| < 1$

θ' : Moving average parameter of aggregate demand, $|\theta'| < 1$

C_i : Constant value of sub-aggregate item demand i

C' : Constant value of aggregate demand.

N : the total number of sub-aggregate items.

We assume that all the sub-aggregate demand series $d_{i,t}$ follow an Integrated Moving Average process of order one, ARIMA(0,1,1), that can be mathematically written by (1).

$$\begin{cases} d_{i,1} = C_i + d_{i,0} & i = 1, 2, \dots, N \\ d_{i,t} = d_{i,t-1} + \varepsilon_{i,t} - \theta_i \varepsilon_{i,t-1} & i = 1, 2, \dots, N; t = 2, 3, \dots \end{cases} \quad (1)$$

From (1) it is obvious that the demand in the next period is the demand in the current period plus an error term. By writing and expanding (1) in a recursive form we have:

$$d_{i,t} = (C_i + d_{i,0}) + \varepsilon_{i,t} + \alpha_i \varepsilon_{i,t-1} + \alpha_i \varepsilon_{i,t-2} + \dots + \alpha_i \varepsilon_{i,1} \quad (2)$$

where $\alpha_i = 1 - \theta_i$. We note that only under this condition on α_i , SES is optimal as it provides the minimum mean square error forecasts for the ARIMA(0,1,1) process. Here we consider the smoothing constant values as a control parameter determined by forecasters which varies between 0 and 1. Obviously, under this condition since $0 \leq \alpha_i \leq 1$, θ_i will only take values between 0 and 1 and does not cover the whole range of $-1 \leq \theta_i \leq 1$.

However, the theoretical analysis is still valid for the whole range of $-1 \leq \theta_i \leq 1$. In addition, in the simulation analysis we will relax this assumption to cover the whole range of $-1 \leq \theta_i \leq 1$ when the value of the smoothing constant is fixed.

We assume that all the sub-aggregate demand process parameters are identical ($\theta_1 = \theta_2 = \theta_3 = \dots = \theta_N$). This assumption is considered only for the purpose of the theoretical analysis and, as above, it is also relaxed in the simulation part of our work. The assumption under concern implies that the aggregate demand also follows an ARIMA (0,1,1) process. If $\theta_1 \neq \theta_2 \neq \theta_3 \neq \dots \neq \theta_N$ then the sum of the sub-aggregate items is not necessarily an ARIMA(0,1,1) process (Granger and Morris, 1976).

The aggregate demand in period t , D_t can be expressed as the sum of the demands of the sub-aggregate items, i.e. $D_t = \sum_{i=1}^N d_{i,t}$.

The forecasting method considered in this study is the Single Exponential Smoothing (SES); this method is being applied in very many companies and in particular in an inventory production planning environment due to its simplicity (Gardner, 1990). At this point we should mention that the MSE is formally defined as $MSE = Var(d_t - f_t) + Bias^2$, where $Bias$ is defined as the expected forecast error and equals to $E(d_t - f_t)$ (Syntetos, 2001).

Note that under the condition that $\alpha_i = 1 - \theta_i$, SES is an unbiased estimator of the demand process considered in this study. Therefore $Bias = 0$ which means that the variance of the forecast error is equal to the mean square error. Using SES, the forecast of sub-aggregate demand i in period t produced at the end of period $t-1$ is

$$f_{i,t} = (C_i + d_{i,0}) + \alpha_i \varepsilon_{i,t-1} + \alpha_i \varepsilon_{i,t-2} + \dots + \alpha_i \varepsilon_{i,1} \quad (3)$$

We further assume that the standard deviation of the error term in (1) is significantly smaller than the expected value of the demand, so should demand be generated the probability of a negative value is negligible.

3.2. Comparison of the Variance of Forecast Error

We calculate the ratio of the variance of the forecast error corresponding to the TD approach (V_{TD}) to the variance of the forecast error associated with the BU approach (V_{BU}). A ratio that is lower than one implies a benefit in favour of the TD approach. Conversely, if the ratio is greater than one then the BU approach performs better (and if the ratio is equal to one both strategies perform the same).

We begin the analysis by deriving the V_{BU} , which is defined as follows:

$$V_{BU} = Var\left(D_t - \sum_{i=1}^N f_{i,t}\right) = Var\left(\sum_{i=1}^N d_{i,t} - \sum_{i=1}^N f_{i,t}\right) = Var\left(\sum_{i=1}^N (d_{i,t} - f_{i,t})\right) \quad (4)$$

by substituting (2) and (3) in (4) we have:

$$V_{BU} = Var\left(\sum_{i=1}^N \varepsilon_{i,t}\right) \quad (5)$$

Since $Var(\varepsilon_{i,t}) = \sigma_i^2$ and $Cov(\varepsilon_{i,t}, \varepsilon_{j,t}) = \rho_{i,j} \sigma_i \sigma_j$ we have:

$$V_{BU} = \sum_{i=1}^N \sigma_i^2 + 2 \sum_{i=1}^{N-1} \sum_{j=i+1}^N \rho_{i,j} \sigma_i \sigma_j \quad (6)$$

We now derive the variance of the forecast error for the TD approach. As discussed above, it has been shown that when the sub-aggregate items follow an IMA (1,1) process, the aggregate family demand also follows an IMA (1,1) process (Granger and Morris, 1976). The family aggregate process is defined as follows:

$$\begin{cases} D_1 = C' + D_0 \\ D_t = D_{t-1} + \varepsilon'_t - (1 - \alpha_{TD}) \varepsilon'_{t-1} \quad t = 2, 3, \dots \end{cases} \quad (7)$$

where $\theta' = 1 - \alpha_{TD}$.

Considering $\theta_1 = \theta_2 = \theta_3 = \dots = \theta_N = \theta$ results in the same theta also in the aggregate demand so, $\theta' = \theta$. Now by considering $\theta' = 1 - \alpha_{TD}$ and $\theta = \theta'$, it is obvious that the optimal smoothing constant for the aggregate demand is $\alpha_{TD} = 1 - \theta$, which is equal to the optimal smoothing constant for the sub-aggregate process.

The aggregate demand and its forecast can be expressed as a function of the error terms, so we have:

$$D_t = (C' + D_0) + \varepsilon'_t + \alpha_{TD} \varepsilon'_{t-1} + \alpha_{TD} \varepsilon'_{t-2} + \dots + \alpha_{TD} \varepsilon'_1 \quad (8)$$

Knowing that $\varepsilon'_t = \sum_{i=1}^N \varepsilon_{i,t}$, we obtain

$$Var(\varepsilon'_t) = \sum_{i=1}^N Var(\varepsilon_{i,t}) + 2 \sum_{i=1}^{N-1} \sum_{j=i+1}^N Cov(\varepsilon_{i,t}, \varepsilon_{j,t})$$

The aggregate forecast is

$$F_t = (C' + D_0) + \alpha_{TD} \varepsilon'_{t-1} + \alpha_{TD} \varepsilon'_{t-2} + \dots + \alpha_{TD} \varepsilon'_1 \quad (10)$$

The variance of the TD forecast error is defined as:

$$V_{TD} = Var[D_t - F_t] \quad (11)$$

By substituting (8) and (10) into (11), we have:

$$V_{TD} = Var[\varepsilon'_t] \quad (12)$$

By substituting (9) into (12) we have:

$$V_{TD} = \sum_{i=1}^N \sigma_i^2 + 2 \sum_{i=1}^{N-1} \sum_{j=i+1}^N \rho_{i,j} \sigma_i \sigma_j \quad (13)$$

Proposition. If all the sub-aggregate demand items follow an ARIMA(0,1,1) process with identical moving average parameters ($\theta_1 = \theta_2 = \theta_3 = \dots = \theta_N$) and the optimal smoothing constant value is used to forecast both the sub-aggregate and aggregate demand, then the performance of the TD and BU approaches for forecasting aggregate demand is identical ($V_{TD} = V_{BU}$). The proof of the proposition follows directly from (6) and (13).

This finding is in agreement with the results reported by Widiarta et al. (2009) which theoretically show that there is no significant difference between the TD and BU approaches on forecasting aggregate demand when all sub-aggregate items follow an MA(1) process with identical process parameters.

4. SIMULATION STUDY

In this section, we perform a simulation study to evaluate the relative performance of the TD over the BU approach under some more realistic assumptions. In particular we consider the following scenarios: *i*) a simulation study at the aggregate level for non-identical ($\theta_1 \neq \theta_2 \neq \dots \neq \theta_N$) process parameters; *ii*) a simulation investigation to discuss the effectiveness of the BU and TD approach compared at the sub-aggregate level for non-identical ($\theta_1 \neq \theta_2 \neq \dots \neq \theta_N$) process parameters. In both cases, the search procedure has been performed in the whole range of $-1 \leq \theta_i \leq 1$.

When the underlying process follows an ARIMA(0,1,1) representation, as θ_i moves from +1 toward -1 the resulting underlying structure changes considerably. When θ_i is negative, the autocorrelation parameter exhibits a smooth exponential decay with positive values and the autocorrelation spans all time lags (not only lag 1). For example for $\theta_i = -0.9$ the autocorrelation is very close to +1. As we move up towards $\theta_i \approx +1$ the autocorrelation reduces but still remains positive and for high positive values of θ_i it becomes close to zero meaning that the series are random.

By considering many SKUs in the simulation experiment, the presentation of results and the evaluation of the impact of different parameters on the ratio of V_{TD} / V_{BU} becomes complex. Therefore, we restricted the simulation analysis to a family of two SKUs to obtain meaningful insights. This is in concordance with most of the earlier papers using simulation approaches as they have also restricted the number of items to two (Dangerfield and Morris, 1992; Fliedner, 1999; Widiarta et al., 2008, 2009).

The parameter values for our simulation experiment are presented in Table 1.

Table 1. Parameters of the simulation experiment

μ_i	σ_i^2	θ_i	ρ_{ij}	N° Sub-Aggregate	N° Replications	N° Time Periods
400	900	-0.9 : +0.9	-0.9 : +0.9	2	100	1000

The sub-aggregate demands in each period are generated randomly subject to the parameters described in Table 1.

The value of σ_i is set to be quite smaller than μ_i to avoid the generation of negative sub-aggregate demand values. Experiments have also been conducted with other values of μ_i and σ_i but these are not reported here as they lead to the same insights.

To generate the demands in each period t , we first generate randomly the error terms $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ with a cross-correlation coefficient of ρ_{12} and then we use (1) to generate the correlated sub-aggregate demands. We initialise the generated demand at the value of the mean plus an error term. The simulation experiment has been designed and run in the forecast package in R. For each parameter combination described in Table 1, a demand series of 1000 observations is generated and we introduce 100 replications.

We split the generated demand for each series at both the sub-aggregate and aggregate level, into two parts. The first part (within sample) consists of 700 time periods and is used in

order to fit the model and estimate the parameters. The smoothing constant and the initial value are found based on Maximum Likelihood estimation (Hyndman et al., 2008; Hyndman et al., 2002). Note that for the BU approach, the smoothing constants are optimized for each item individually. Finally, in order to evaluate the performance of the two forecasting approaches, we calculate the value of the variance of the forecast error for the last 300 periods of the simulation (out-of-sample).

The relative benefit of one forecasting approach over the other is measured by V_{TD} / V_{BU} . As previously discussed, a ratio lower than one implies that the TD approach outperforms the BU one whereas a ratio greater than one implies the opposite.

4.1. Comparison at the Aggregate Level

We first analyse the relative performance of the two forecasting approaches at the aggregate level when the sub-aggregate process parameters are not necessarily identical. For each experiment, the ratio of the variance of the forecast error is calculated as

$$\text{Var}(D_t - F_t) / \text{Var}\left(D_t - \sum_{i=1}^2 f_{i,t}\right).$$

The simulation results show that when the process parameters are identical there is no difference between the BU and the TD approach. Whereas, when the process parameters are not identical, which is more realistic, the results are different. The results for the latter case are presented in Figure 2.

We see that as the cross-correlation coefficient changes from -0.9 toward $+0.9$ the ratio of V_{TD}/V_{BU} is being reduced. The ratio is higher than or equals to 1 when the cross-correlation is negative, equals to zero or takes low positive values. The ratio is lower than 1 only if the cross-correlation is (highly) positive.

The detailed results show that when the moving average parameters, θ_1 and θ_2 , take negative values, the performance of BU and TD approaches is always identical regardless of the values of the cross-correlation. One possible explanation for this result is that when the MA parameters of sub-aggregate series take negative values, the optimal value of the smoothing constant is set at the highest value in the considered range which is equal to 0.99 for both approaches. As the smoothing constant for both BU and TD approaches is equal and the same procedure of forecasting is used for the sub-aggregate items and the aggregate one, the aggregate forecast under both BU and TD approaches is the same.

Figure 2. Relative performance of the TD and BU approach in forecasting aggregate demand under different combinations of θ_1 , θ_2 and ρ_{12}

When the cross-correlation is positive the superiority of each approach depends on the value and the sign of the moving average parameters, θ_1 and θ_2 . The TD approach

outperforms the BU one only when the cross-correlation is (highly) positive and the moving average parameters take high values and have opposite signs, i.e. either $\theta_1 < 0$ and $\theta_2 > 0$ or $\theta_1 > 0$ and $\theta_2 < 0$.

Note that as the cross-correlation decreases the superiority of the TD approach decreases too. For highly positive cross-correlation, TD outperforms BU with a forecast error variance reduction that can go up to 15%. By decreasing the cross-correlation to 0.5, the maximum benefit of the TD approach decreases to 5% and it tends toward zero when the cross-correlation tends towards zero as well. Under negative cross-correlation, BU outperforms TD.

When the two moving average parameters take opposite signs under the ARIMA(0,1,1) process, this means that one series has positive autocorrelation while the other has a low autocorrelation (series with random fluctuations). In addition, when the cross-correlation is positive there is a tendency for the pair of series to move together in the same direction, so the demand series have the same pattern. When using TD, we sum up all sub-aggregate series to get an aggregate one, so the fluctuations from one series may be cancelled out by those of another resulting in a less random series associated with a lower forecast error.

When the cross-correlation coefficient is negative, for all values of θ_1 and θ_2 , with the exception of the case when both are negative, the BU approach performs better. Performance differences are further inflated when the moving average parameters have opposite signs in which case the variance reduction achieved by the BU approach can be as high as 500% for highly negative cross-correlation. For negative cross-correlation, the pair of series moves in the opposite direction, (i.e. if one increases the other decreases), so the sub-aggregate demand series have different patterns of evolution. Combination of different patterns of variation and opposite autocorrelation values lead to a large forecast error for the TD approach and consequently large values of V_{TD} / V_{BU} for very high negative cross-correlation. In these cases it is better to forecast sub-aggregate requirements separately and then aggregate them to get the aggregate forecast.

When the θ_1 and θ_2 values are positive, the ratio is almost equal to 1 for highly positive cross-correlation and greater than 1 for less positive and negative cross-correlation. In the latter case, the ratio of V_{TD} / V_{BU} is increased as θ_1 takes low values and θ_2 is high and vice versa.

In summary, when the sub-aggregate items follow an ARIMA(0,1,1) process and the goal is to forecast at the aggregate demand level, then: *i*) if the autocorrelation of all items is highly positive, the performance of BU and TD is always identical; *ii*) if items have different

autocorrelation patterns, one has a very high positive autocorrelation while the other has an autocorrelation close to zero, the superiority of each approach is affected by the cross-correlation between items; for a highly positive cross-correlation, TD outperforms BU and for a highly negative cross-correlation BU outperforms TD; *iii*) when the autocorrelation for all items is low, BU generally dominates TD, although for highly positive cross-correlation the difference is very low.

Our findings are somehow in agreement with some of the earlier studies in this area by Barnea and Lakonishok (1980) and Fliedner (1999) (although we do note that our results are not directly comparable to these studies as we analyse a non-stationary case). The analysis of Barnea and Lakonishok (1980) based on an empirical evaluation showed that positive cross-correlation contributes to the superiority of forecasts based on aggregate data (TD), which is also the case in our study. Fliedner (1999) used a simulation study to compare the performance of TD and BU in forecasting aggregate series where the two sub-aggregate items follow an MA(1) process. He found that TD dominated BU regardless of the values of the cross-correlation coefficient. They have not reported the values of θ_1 and θ_2 used in their study, so our interpretation is that this work considered only the opposing signs for θ_1 and θ_2 . Should this be the case then these findings are in agreement with ours.

4.2. Comparison at the Sub-Aggregate Level

In this subsection we evaluate the relative performance of the TD and BU approaches in forecasting sub-aggregate demand when the moving average parameters are not necessarily identical. The simulation structure in terms of within and out-of-sample arrangements is as discussed in the previous sub-section. Under the BU approach, we generate 300 one step-ahead forecasts for each item individually using the optimal smoothing constant. Under the TD approach, we first sum up the demand of all sub-aggregate items to obtain aggregate series, we then produce the aggregate forecast and finally we break down (disaggregate) that to sub-aggregate forecasts by using proportional factors based on the historical contribution of each series. For each experiment, the ratio of the variance of forecast error is calculated as: $\sum_{i=1}^2 \text{Var}(d_{i,t} - p_i * F_t) / \sum_{i=1}^2 \text{Var}(d_{i,t} - f_{i,t})$.

Figure 3 shows the ratio of the variance of forecast error of the TD over the BU approach at the sub-aggregate level for different values of θ_1 , θ_2 , ρ_{12} and p_1 when the sub-aggregate items follow an ARIMA(0,1,1) process with non-identical moving average parameters ($\theta_1 \neq \theta_2$). Different levels of item proportion, p_i , are used to reflect the cases where the sub-

aggregate items contribute almost equally to the aggregate forecast and cases where one item dominates the aggregation process. The results show that the BU approach always outperforms TD in forecasting the sub-aggregate items regardless of the ρ_{12} and p_i .

Figure 3. Relative performance of TD and BU approaches in forecasting sub-aggregate items under different values of θ_1 , θ_2 , ρ_{12} and p_i .

In Figure 3 we show that by moving from a cross-correlation of -0.9 toward +0.9 the ratio of V_{TD}/V_{BU} is generally reduced but it always remains greater than 1 regardless of the cross-correlation coefficient and the items' contributory power.

When the cross-correlation and the moving average parameters, θ_1 , θ_2 , are highly positive, *i.e.* $\theta_1 \approx 0.99$, $\theta_2 \approx 0.99$ and $\rho_{12} \approx 0.99$, the ratio of V_{TD}/V_{BU} becomes close to one.

Figure 3a shows also that BU outperforms TD by a maximum of about 50% for highly negative cross-correlation; the rate of superiority of BU becomes very high when θ_1 and θ_2 are not highly positive (see Figure 3b, c, d). Widiarta et al. (2009) reported that when demand follows an MA(1) process, BU outperformed TD by a maximum of 4% when $\theta_1=0.3$ and $\theta_2=-0.8$ and the cross-correlation is negative. The superiority of BU at the sub-aggregate level can be attributed to the potentially high positive autocorrelation between demand periods. When the series follow an ARIMA(0,1,1) process, the autocorrelation is (highly) positive unless the moving average parameter takes high positive values, in which cases autocorrelation becomes close to zero. Generally, as θ_i moves from positive toward negative values, the autocorrelation between two consecutive observations $d_{i,t}$ positively increases, in addition to spanning higher time lags (not only a lag of 1). This makes it much more difficult to apportion the resulting aggregate forecast, Ft , to each item in the family based on the historical demand proportion, p_i . As a result, the performance of the TD approach is affected adversely. The performance of the BU approach, however, is not affected as it forecasts the demand for each item individually.

Our findings are in accordance with those previously reported in the academic literature. Widiarta et al. (2007) argued that when the sub-aggregate time series follow an AR(1) process and the value of the autocorrelation is high, there is a sharp deterioration in the relative performance of TD. Gordon et al. (1997) and Dangerfield and Morris (1992) used empirical data from the M-competition database and stated that BU dominated TD when forecasting the sub-aggregate time series. Weatherford et al. (2001) have shown that a purely disaggregate forecast (BU) strongly outperformed even the best aggregate forecast (TD) at the sub-aggregate level.

These results generally confirm our findings although we must note (as we did in the previous sub-section) that there is not a direct comparison between these studies and ours due to the consideration of a non-stationary ARIMA(0,1,1) time series process. Contrasting our results with those reported by Widiarta et al. (2007, 2009) on stationary MA(1) and AR(1) processes, we observe that the rates of superiority of the BU approach when the process is

non-stationary is much higher than the stationary case. When the demand follows a stationary AR(1) process, the maximum ratio of V_{TD}/V_{BU} equals to 2 and is obtained with series with high positive autocorrelation (high positive autoregressive parameter values ϕ_1 , ϕ_2) and highly negative cross-correlation, while this ratio for the IMA(1,1) process is higher than 50.

5. EMPIRICAL ANALYSIS

In this section, we assess the empirical validity of our results on the comparative performance of the TD and BU approaches. In addition, the empirical performances of these approaches are compared to that of the optimal approach proposed by Hyndman et al. (2011), at both the aggregate and subaggregate levels. We first provide details of the empirical data available for the purposes of our investigation along with the experimental structure employed in our work. We then present the actual empirical results.

The demand dataset available for the purposes of our research consists of 103 weekly sales observations (i.e. it spans a period of two years) for 1,798 SKUs from a European grocery store. The *auto.arima* function of the forecast package in R has been used to identify the underlying ARIMA demand process for each series and to estimate the relevant parameters. This function uses a variation of the Hyndman and Khandakar (2008) algorithm which combines unit root tests, minimization of the Akaike's Information Criteria (AIC_c) and maximum-likelihood estimation (MLE) to identify an ARIMA(p,d,q) model. First, the number of differences d is determined using unit-root tests by applying repeated KPSS tests (Kwiatkowski et al., 1992). Then, the value of process orders, p and q , are chosen by minimizing the AIC_c after differencing the data d times. Please refer to Hyndman and Khandakar (2008) for a discussion on ARIMA identification methodology related issues.

Based on the identification process discussed above, it was found that around 24% of the series (424 series) may be represented by the process considered in this research, ARIMA(0,1,1). In Table 2 we summarize the characteristics of the SKUs relevant to our study by indicating the estimated parameters for the ARIMA(0,1,1) process. It is important to note that these results are sensitive to the modelling methodology being used to identify the series in the first place. If the methodology employed by the *auto-arima* function potentially identifies ARIMA(0,1,1) series incorrectly then our results will be obviously subject to relevant errors.

To facilitate a clear presentation, the estimated parameters are grouped in intervals and the corresponding number of SKUs is given for each such interval. The average θ value per

interval is also presented. This categorisation allows us to compare the empirical results with the theoretical findings. We must remark that the θ parameter values are all positive, except for two SKUs, and most of them take highly positive values. As such, the data do not cover the entire theoretically feasible range of the parameters.

Table 2. Processes present in the empirical data set

Group	θ intervals	Average of θ	No. of SKUs
1	[0.1,0.3[	0.2097	4
2	[0.3,0.4[	0.3652	8
3	[0.4,0.5[	0.4656	17
4	[0.5,0.6[	0.5591	32
5	[0.6,0.7[	0.6561	67
6	[0.7,0.8[	0.7503	108
7	[0.8,0.9[	0.8467	141
8	[0.9,1]	0.9534	47
Total number of SKUs:			424

The data series have been divided into two parts. The first part (within sample) consists of 70 time periods and is used in order to estimate the *SES* parameters. The second part consists of 33 time periods which are used to evaluate the performance of each approach (out-of-sample). The geometric mean (across SKUs) of the V_{TD}/V_{OP} and V_{BU}/V_{OP} ratios is considered for comparison purposes at the disaggregate level (where V_{OP} is the variance of the forecast errors resulting from the implementation of the optimal approach). Note that the ratio V_{TD}/V_{BU} can be directly deduced from the two variance ratios as $(V_{TD}/V_{OP}) / (V_{BU}/V_{OP})$.

The empirical results presented in Table 3 are shown for the same θ intervals considered in Table 2. With regards to the comparative performance of BU and TD, we can see that when the smoothing constant values are optimised for both approaches, the variance ratio is greater than 1 when the comparison is undertaken at the sub-aggregate level, whereas when the comparison is undertaken at the aggregate level, the difference between BU and TD is insignificant. This means that overall one can consider that the BU approach provides more accurate forecasts. Furthermore, when the smoothing constants used for BU and TD are equal, the ratio of V_{TD}/V_{BU} equals to 1 in the case of disaggregate demand forecasting. As discussed above, the moving average parameter θ is highly positive for most SKUs considered in this research. More than 85% of the SKUs have a moving average parameter greater than 0.6 (see Table 2). In addition, the sub-aggregate cross-correlation coefficients between SKUs vary

between -0.5 and +1; however most of those coefficients are positive. By referring to the detailed results of the simulation study we see that for this range of moving average parameter values, $0 < \theta < 1$, the BU approach performs better than TD at the subaggregate level. However, for the comparison at the aggregate level, the ratio of V_{TD}/V_{BU} is close to one, and the superiority of each approach depends on the cross-correlation. TD may outperform BU for highly positive cross-correlation. In an empirical context, the average of the variance of forecast error reduction may be as high as 1.01% when the comparison is performed at the aggregate level, while 70% variance error reduction may be achieved for the comparison at the sub-aggregate level.

Table 3. Empirical variance ratios for an ARIMA(0,1,1) process

Group	θ intervals	Comparison Level			
		Aggregate		Disaggregate	
		V_{TD}/V_{OP}	V_{BU}/V_{OP}	V_{TD}/V_{OP}	V_{BU}/V_{OP}
1	[0.1,0.3[	0.998	1.020	2.391	0.990
2	[0.3,0.4[	1.001	1.022	2.078	0.993
3	[0.4,0.5[	1.000	1.005	1.880	0.997
4	[0.5,0.6[	1.000	1.002	1.727	0.994
5	[0.6,0.7[	1.000	1.020	1.629	0.990
6	[0.7,0.8[	1.000	1.008	1.428	1.003
7	[0.8,0.9[	1.000	1.007	1.290	0.996
8	[0.9,1]	1.001	1.000	1.162	0.999
Average		1.000	1.011	1.613	0.995

With regards to the implementation of the optimal approach, the hierarchy structure of our data consists of two levels. At the top level we have aggregated all series to get one single series while at the bottom level we have 424 series, so we have 425 series in total for forecasting purposes. First, we generated forecasts for all 425 series using SES. Next, we have used the *combinef* function in the *hts* package (Hyndman et al., 2014) of R to reconcile these forecasts using the optimal method to obtain forecasts at both aggregate and sub-aggregate levels.

With regards to the comparative performance of the optimal approach (as reported based on its variance of forecast errors, V_{OP}), Table 3 shows that for comparison at the aggregate level, the optimal method performs better than BU; however there is no significant difference

between the performance of the optimal method and the TD approach. When the comparison is undertaken at the sub-aggregate level, the optimal method significantly outperforms TD, whereas BU works slightly better than the optimal method (the difference is less than 1%). Generally, for the empirical data used in this study, either the TD or the optimal method can be used for forecasting at the aggregate level, while at the sub-aggregate level, BU is preferable, although the optimal approach may also be used as the difference is less than 1%.

In Table 3, we report ratios of the variance of forecast errors in specific moving average parameter interval. In Table 4, we report collective performance across different possible (ranges of) moving average parameter values and we evaluate the impact of such values on the superiority of each approach. To do so we create a category containing groups 1, 2 and 3 that includes 29 SKUs; this is regarded as a category with the lowest values of θ . As we move from this category to groups 4, 5 and 6 the value of θ increases. We aggregate all these groups with group 8 that represents the highest values of θ . The ratios V_{BU}/V_{OP} and V_{TD}/V_{OP} are then presented for both levels of comparison.

Table 4. Empirical variance ratios by reporting performance across different groups (intervals of θ values)

Comparison Level		Group	1,2,3	4	5	6
Aggregate	V_{TD}/V_{OP}	8	1.001	1.001	1.000	1.000
	V_{BU}/V_{OP}		0.855	0.902	0.977	0.998
Subaggregate	V_{TD}/V_{OP}		1.484	1.435	1.399	1.366
	V_{BU}/V_{OP}		0.956	0.976	0.987	0.998

With regards to the comparative performance of BU and TD, the results indicate that when two groups with different moving average parameters are considered (Group 1,2,3 with 8) then the variance ratio V_{TD}/V_{BU} is high and as the θ values increases (tending towards the values covered by group 8) the ratio decreases. This implies that when we aggregate groups of SKUs with low and high θ values there is a greater benefit of using the BU approach in terms of accuracy. This is exactly what we have observed in the simulation results for 2 SKUs (one associated with a small and one with a high θ value). These empirical results generally confirm the findings of the simulation study. With regards to the performance of the optimal approach, Table 4 shows that for comparisons at the sub-aggregate level, the optimal method outperforms TD. However, BU performs better than the optimal approach. When

performances are contrasted at the aggregate level, both optimal and TD approaches perform equally well and BU outperforms them.

6. DISCUSSION, CONCLUSION AND FURTHER RESEARCH

In this paper we have evaluated analytically the effectiveness of the bottom-up (BU) and top-down (TD) approaches to forecasting aggregate and subaggregate demand when the subaggregate series follow a first order integrated moving average [ARIMA(0,1,1)] process. Forecasting was assumed to be relying upon a Single Exponential Smoothing (SES) procedure and the analytical results were complemented by a simulation experiment at both the aggregate and sub-aggregate level as well as experimentation with an empirical dataset from a European superstore.

Admittedly, the current fast changing market environments result in many demand processes being non-stationary in nature. Some empirical pieces of work discussed in subsection 2.1 confirm such a statement and provide support for the frequency with which ARIMA(0,1,1) processes are encountered in real world applications. In addition, SES is a most commonly employed forecasting procedure in industry and its application implies a non-stationary behaviour (SES is optimal for an ARIMA(0,1,1) process). Both BU and TD approaches are very useful in practice when dealing with Sales and Operations Planning systems in which forecasting is required at both aggregate and subaggregate levels. In summary, we feel that the problem setting we have considered is a very realistic one. Analytical and simulation developments were based on the consideration of the variance of forecast error for the TD and BU approaches and comparisons were undertaken at the aggregate level in the theoretical part of this work and at both the subaggregate and aggregate level in the simulation investigation. The conditions under which one approach outperforms the other were identified and the main findings can be summarized as follows:

- When the moving average parameter for all the subaggregate items is identical ($\theta_1 = \theta_2 = \dots = \theta_N$), there is no significant difference between TD and BU in forecasting the aggregate level, as long as the optimal smoothing constant is used for both approaches. When the smoothing constant used for all the subaggregate items and the aggregate level is set to be identical ($\alpha = \alpha_{TD}$), TD and BU perform the same in forecasting the demand at the aggregate level regardless of the values of the moving average parameters and the cross-correlation between items. In addition when the

observations of the subaggregate items are highly auto-correlated (negative θ), the performance of BU and TD is also the same for all autocorrelation values.

- TD performs better than BU at the aggregate level when the subaggregate moving average parameters take relatively high values of an opposing sign and the cross-correlations between sub-aggregate items are (highly) positive. Otherwise, when cross-correlation is positive low or takes negative values, BU is preferable. Therefore, using the aggregate data to produce top level forecasts is preferable only if the sub-aggregate items follow similar series evolution with combination of high vs. low autocorrelation. The TD appears not to be very accurate when the sub-aggregate items consist of different patterns of fluctuation.
- BU outperforms TD when forecasting at the sub-aggregate level and when the smoothing constant is set to its optimal value for both approaches, regardless of the cross-correlation, the disaggregation weight and the values of the process parameters. It's not preferable to use the aggregate data to derive the individual forecasts, when the autocorrelation of subaggregate items is highly positive, in which cases subaggregate data provide more accurate forecasts. The degree of superiority of the BU approach for non-stationary processes is much higher compared to stationary ones.
- The performance of BU improves as the cross-correlation decreases, moving from positive toward negative values. For highly negative cross-correlation values BU is always preferable; this is generally true for comparisons at both the aggregate and sub-aggregate level.
- The benefit achieved by BU and TD for non-stationary demands in terms of forecast accuracy is higher than that associated with stationary cases.
- The optimal approach performs well at both levels of comparison as indicated by the empirical results. When considering ratios of the variance of forecast errors in particular moving average parameter intervals or different possible (ranges of) moving average parameter values, the optimal approach is superior as it performs as well as the BU at both levels but significantly better than the TD at the disaggregate level. It should be noted though that the optimal approach bears considerable relevance to many realistic cases when: i) more than two levels of hierarchy need to be considered, and ii) (more than one) various forecasting methods need to be employed.

Please also note that since the optimal approach is not always superior to the BU and TD approaches, the comparative performance between BU and TD needs to be carefully

considered in order to inform real world applications. In addition, and even when the optimal approach outperforms both BU and TD, the latter approaches still remain of explicit interest to practitioners (because of the simplicity characterising their implementation, their intuitive appeal and their support by most inventory software) but also to academics (because of the insights that the mathematical analysis of those cases may offer – in contrast with the optimal approach where the optimisation procedure hinders any explicit messages as to what is going on with the underlying properties of the series.) The major difficulty associated with the optimal approach is the computation of the reconciliation weights used to form a weighted average of forecasts at an individual node and the non-transparent nature of the regression analysis taking place. Another difficulty is the fact that forecasts from all levels need to be taken into account when producing the final (reconciled) forecasts which obviously increases the computational effort and managerial involvement beyond that required by either BU or TD.

If the practitioners require demand forecasts at the SKU level when demand is non-stationary (and highly autocorrelated) then it would always be preferable to use the BU approach. If a higher level demand forecast is needed then the BU approach should be considered when the individual items are associated with different patterns of evolution, and the TD or the optimal approach when they have the same patterns but are associated with different autocorrelation values. In addition, if one uses the same value of smoothing constants for both BU and TD, then both approaches perform the same in forecasting aggregate demand.

In this paper we have considered the case of non-stationary demand processes in conjunction with the SES forecast method to evaluate the comparative performance of TD and BU in forecasting aggregate and item level demand. Naturally, there are many other avenues for further research and the following possibilities should be very important in terms of advancing the current state of knowledge in this area.

- The consideration of more extensive datasets that cover the whole range of the process parameters should allow a better understanding of the comparative benefits of the TD and BU approach.
- The extension of the work described here to cover inventory/implication metrics would allow a linkage between forecasting and stock control. Cross-sectional aggregation is known to be very helpful in inventory applications and it is indeed being covered by relevant software packages. Further work into the interactions

between forecasting and stock control in a cross-sectional setting should add value to real world practices.

- The interface between (and the potential of combining) temporal and cross-sectional aggregation has received minimal attention both in academia and industry and is an issue that we are to take further in the next steps of our research. A *unified approach* – one that simultaneously considers choices of aggregation levels and frequency along multiple dimensions – would seem to be a valuable step in the right direction. The problem with the separation of the cross-sectional and temporal dimensions is that the *right level* of cross-sectional aggregation may vary across time frequencies and vice versa. Procedures that combine forecasts for a cross-sectional hierarchy, such as the optimal approach discussed in this paper, and procedures that combine forecasts over time frequencies, such as the multiple temporal-aggregation technique discussed by Kourentzes et al. (2014), may conceivably be *pooled* to form a holistic strategy for forecasting hierarchies (see also the Introduction of Tashman et al. (2015)).
- Finally, the analytical and empirical consideration of Integer ARMA (INARMA) processes offers a great opportunity for advancements in the area of aggregation. Such processes bear a considerable relevance to intermittent demands where the benefits of aggregation may be even higher due to the reduction of zero observations (Mohammadipour and Boylan, 2012).

REFERENCES

- Acar, Y., Gardner, E.S., 2012. Forecasting method selection in a global supply chain. *International Journal of Forecasting* 28, 842-848.
- Ali, M.M., Boylan, J.E., Syntetos, A.A., 2011. Forecast errors and inventory performance under forecast information sharing. *International Journal of Forecasting*.
- Athanasopoulos, G., Ahmed, R.A., Hyndman, R.J., 2009. Hierarchical forecasts for Australian domestic tourism. *International Journal of Forecasting* 25, 146-166.
- Barnea, A., Lakonishok, J., 1980. An Analysis of The Usefulness of Disaggregated Accounting Data For Forecasts of Corporate Performance. *Decision Sciences* 11, 17-26.
- Bijvank, M., Vis, I.F.A., 2011. Lost-sales inventory theory: A review. *European Journal of Operational Research* 215, 1-13.
- Box, G., Jenkins, G.M., Reinsel, G., 2008. *Time Series Analysis: Forecasting & Control*. John Wiley and sons.
- Bozos, K., Nikolopoulos, K., 2011. Forecasting the value effect of seasoned equity offering announcements. *European Journal of Operational Research* 214, 418-427.
- Chen, H., Boylan, J.E., 2007. Use of Individual and Group Seasonal Indices in Subaggregate Demand Forecasting. *The Journal of the Operational Research Society* 58, 1660-1671.
- Chen, H., Boylan, J.E., 2008. Empirical evidence on individual, group and shrinkage seasonal indices. *International Journal of Forecasting* 24, 525-534.

- Chien, C.-F., Chen, Y.-J., Peng, J.-T., 2008. Demand forecast of semiconductor products based on technology diffusion, Simulation Conference, 2008. WSC 2008. Winter, pp. 2313-2322.
- Dangerfield, B.J., Morris, J.S., 1992. Top-down or bottom-up: Aggregate versus disaggregate extrapolations. *International Journal of Forecasting* 8, 233-241.
- Erkip, N., Hausman, W.H., Nahmias, S., 1990. Optimal Centralized Ordering Policies in Multi-Echelon Inventory Systems with Correlated Demands. *Management Science* 36, 381-392.
- Flüedner, G., 1999. An Investigation of Aggregate Variable Time Series Forecast Strategies with Specific Subaggregate Time Series Statistical Correlation. *Computers and Operations Research* 26, 1133-1149.
- Gardner, E.S., Jr., 1990. Evaluating Forecast Performance in an Inventory Control System. *Management Science* 36, 490-499.
- Gardner, E.S., Jr., 2006. Exponential smoothing: The state of the art-Part II. *International Journal of Forecasting* 22, 637-666.
- Goh, C., Law, R., 2002. Modeling and forecasting tourism demand for arrivals with stochastic nonstationary seasonality and intervention. *Tourism Management* 23, 499-510.
- Gordon, T., Dangerfield, B., Morris, J., 1997. Top-down or bottom-up: Which is the best approach to forecasting? *The Journal of Business Forecasting* 16, 13-16.
- Granger, C.W.J., Morris, M.J., 1976. Time series modeling and interpretation. *Journal of the Royal Statistical Society. Series A*, 246-257.
- Graves, S.C., Willems, S.P., 2000. Optimizing Strategic Safety Stock Placement in Supply Chains. *Manufacturing & Service Operations Management* 2, 68-83.
- Graves, S.C., Willems, S.P., 2008. Strategic Inventory Placement in Supply Chains: Nonstationary Demand. *Manufacturing & Service Operations Management* 10, 278-287.
- Gross, C.W., Sohl, J.E., 1990. Disaggregation methods to expedite product line forecasting. *Journal of Forecasting* 9, 233-254.
- Hyndman, R.J., Ahmed, R.A., Athanasopoulos, G., Shang, H.L., 2011. Optimal combination forecasts for hierarchical time series. *Computational Statistics & Data Analysis* 55, 2579-2589.
- Hyndman, R.J., Khandakar, Y., 2008. Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software*, 26, 1-22.
- Hyndman, R.J., Koehler, A.B., Ord, J.K., Snyder, R.D., 2008. *Forecasting with Exponential Smoothing: The State Space Approach*. Springer Verlag, Berlin.
- Hyndman, R.J., Koehler, A.B., Snyder, R.D., Grose, S., 2002. A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of Forecasting* 18, 439-454.
- Hyndman, R.J., Wang, E., Lee, A., 2014. Package 'hts': Methods for analysing and forecasting hierarchical and grouped time series.
- Kourentzes, N., Petropoulos, F., Trapero, J.R., 2014. Improving forecasting by estimating time series structural components across multiple frequencies. *International Journal of Forecasting* 30, 291-302.
- Kwiatkowski, D., Phillips, P.C., Schmidt, P., Shin, Y., 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics* 54, 159-178.
- Lapide, L., 2004. Sales and Operations Planning Part I: The Process. *The Journal of Business Forecasting*, 17-19.
- Lapide, L., 2006. Top-Down & Bottom-Up Forecasting In S&OP. *The Journal of Business Forecasting* 25, 14-16.
- Lee, H.L., P., P., Whang, S., 1997. Bullwhip effect in a supply chain. *Sloan Management Review* 38, 93-102.
- Li, B., Wang, H.-W., Yang, J.-B., Guo, M., Qi, C., 2011. A belief-rule-based inventory control method under nonstationary and uncertain demand. *Expert Systems with Applications* 38, 14997-15008.
- Lütkepohl, H., 1984. Forecasting contemporaneously aggregated vector ARMA processes. *Journal of Business & Economic Statistics* 2, 201-214.
- Mahajan, S., Desai, O., 2011. Value of Information in a Serial Supply Chain under a Nonstationary Demand Process, Working Paper Series. Indian Institute of Management (IIMB), Bangalore
- Makridakis, S., Andersen, A., Carbone, R., Fildes, R., Hibon, M., Lewandowski, R., Newton, J., Parzen, E., Winkler, R., 1982. The accuracy of extrapolation (time series) methods: Results of a forecasting competition. *Journal of Forecasting* 1, 111-153.

- Martel, A., Diaby, M., Bector, F., 1995. Multiple items procurement under stochastic nonstationary demands. *European Journal of Operational Research* 87, 74-92.
- Mitchell, G., Niederhausen, M., 2010. On Replenishing Items with Seasonal Intermittent Demand. *American Journal of Economics and Business Administration* 2, 90-102.
- Mohammadipour, M., Boylan, J., Syntetos, A., 2012. The Application of Product-Group Seasonal Indexes to Individual Products. *Foresight: The International Journal of Applied Forecasting*, 18-24.
- Mohammadipour, M., Boylan, J.E., 2012. Forecast horizon aggregation in integer autoregressive moving average (INARMA) models. *OMEGA* 40, 703-712.
- Orcutt, G.H., Watts, H.W., Edwards, J.B., 1968. Data aggregation and information loss. *The American Economic Review* 58.
- Raghunathan, S., 2001. Information Sharing in a Supply Chain: A Note on its Value when Demand Is Nonstationary. *Management Science* 47, 605-610.
- Sbrana, G., Silvestrini, A., 2013. Forecasting aggregate demand: Analytical comparison of top-down and bottom-up approaches in a multivariate exponential smoothing framework. *International Journal of Production Economics*.
- Shang, K.H., 2012. Single-Stage Approximations for Optimal Policies in Serial Inventory Systems with Nonstationary Demand. *Manufacturing & Service Operations Management*.
- Shlifer, E., Wolff, R.W., 1979. Aggregation and Proration in Forecasting. *Management Science* (pre-1986) 25, 594.
- Syntetos, A.A., 2001. Forecasting of intermittent demand, Ph D, Business School, Buckinghamshire Chilterns University College, Brunel University, London, UK.
- Tashman, L., 2015. Guide on 'Techniques for forecasting product and temporal hierarchies. *FORESIGHT: The international journal of applied forecasting*.
- Taylor, J.W., 2003. Exponential smoothing with a damped multiplicative trend. *International Journal of Forecasting* 19, 715-725.
- Trehan, J.T., Sox, C.R., 2002. Adaptive Inventory Control for Nonstationary Demand and Partial Information. *Management Science* 48, 607-624.
- Tunc, H., Kilic, O.A., Tarim, S.A., Eksioglu, B., 2011. The cost of using stationary inventory policies when demand is non-stationary. *Omega* 39, 410-415.
- Weatherford, L.R., Kimes, S.E., Scott, D.A., 2001. Forecasting for hotel revenue management: Testing aggregation against disaggregation. *The Cornell Hotel and Restaurant Administration Quarterly* 42, 53-64.
- Widiarta, H., Viswanathan, S., Piplani, R., 2007. On the effectiveness of top-down strategy for forecasting autoregressive demands. *Naval Research Logistics (NRL)* 54, 176-188.
- Widiarta, H., Viswanathan, S., Piplani, R., 2008. Forecasting item-level demands: an analytical evaluation of top-down versus bottom-up forecasting in a production-planning framework. *IMA Journal of Management Mathematics* 19, 207-218.
- Widiarta, H., Viswanathan, S., Piplani, R., 2009. Forecasting aggregate demand: An analytical evaluation of top-down versus bottom-up forecasting in a production planning framework. *International Journal of Production Economics* 118, 87-94.
- Zotteri, G., Kalchschmidt, M., Caniato, F., 2005. The impact of aggregation level on forecasting performance. *International Journal of Production Economics* 93-94, 479-491.