

Empirical heuristics for improving Intermittent Demand Forecasting

Fotios Petropoulos^{1,*}, Konstantinos Nikolopoulos²,
Georgios P. Spithourakis¹, Vassilios Assimakopoulos¹

¹Forecasting & Strategy Unit, School of Electrical and Computer Engineering,
National Technical University of Athens, Greece

²The Business School, Bangor University, Bangor, UK

fotis@fsu.gr; k.nikolopoulos@bangor.ac.uk; giorgos@fsu.gr; vassim@fsu.gr

*corresponding author

Abstract

Purpose: Intermittent demand appears sporadically, with some time periods not even displaying any demand at all. Even so, such patterns constitute considerable proportions of the total stock in many industrial settings. Forecasting intermittent demand is a rather difficult task but of critical importance for corresponding cost savings. The current study examines the empirical outcomes of three heuristics towards the modification of established intermittent demand forecasting approaches.

Design/methodology/approach: First, optimization of the smoothing parameter used in Croston's approach is empirically explored, in contrast to the use of an a priori fixed value as in earlier studies. Furthermore, the effect of integer rounding of the resulting forecasts is considered. Lastly, we evaluate the performance of Theta model as an alternative of SES estimator for extrapolating demand sizes and/or intervals. The proposed heuristics are implemented into forecasting support system.

Findings: The experiment is performed on 3,000 real intermittent demand series from the automotive industry, while evaluation is made both in terms of bias and accuracy. Results indicate increased forecasting performance.

Originality/Value: The current research explores some very simple heuristics which have a positive impact on the accuracy of intermittent demand forecasting approaches. While, some of these issues have been partially explored in the past, the current research focuses on a complete in-depth analysis of easy to employ modifications to well established intermittent demand approaches. By this, we enable the application of such heuristics on an industrial environment, which may lead into significant inventory and

production cost reductions and other benefits.

Keywords: intermittent demand; smoothing parameters; rounding; theta method; empirical investigation

1. Introduction

Demand and inventory forecasts are required for virtually all decision making situations regarding future events, from short term forecasts dealing with inventories and scheduling to medium and long term ones needed for strategy and planning. Accurate forecasts are of great practical importance, linking inventory costs with revenues, customer satisfaction, stock-out costs and lead time (for example Huang et al., 2011). Intermittent demand patterns are characterized by infrequent demand arrivals coupled with variable demand sizes, whenever demand occurs. Intermittent demand items may be engineering spare parts or other items within the range of products offered by any organization and at any level of the supply chain.

Sparse demand creates significant problems in the manufacturing and supply environment as far as forecasting is concerned. It is not only the variability of the demand size, but also the variability of the interval between demands that make intermittent demand so difficult to forecast. If the fact that slow moving items may constitute up to 60% of the total stock in any industrial setting (Johnston et al., 2003) is also taken into account, it becomes obvious that small improvements can instigate substantial cost savings.

The current study examines three empirical heuristics used with established and commonly used forecasting approaches for intermittent demand (Croston, 1972; Syntetos and Boylan, 2001; Syntetos, 2001). Firstly, optimization issues, regarding the optimal smoothing parameters used, are looked into. Secondly, an intuitively attractive and practically indispensable heuristic is investigated: rounding of the final forecasts, as all demand sizes of SKUs are whole numbers. Lastly, an approach that combines Croston method for intermittent demand with Theta model (Assimakopoulos and Nikolopoulos, 2000) is thoroughly investigated.

All three heuristics result in very promising results when applied to data characterized by intermittent demand patterns, where the presence of zero demand is evident. Taking into account that such data arise in many industries, it comes without saying that the

current research aims at improving settings and procedures by easy-to-implement modifications on conventional techniques. Furthermore, to our knowledge it is the first time that the effectiveness of independently selecting the smoothing values for the numerator and the denominator of each one of the series, via a widely used cost function (MSE), is being empirically investigated. Lastly, the full potential of utilizing the Theta model in Croston's framework for the extrapolation of its components is analyzed and discussed.

The remainder of the current paper is structured as follows. Firstly, a short literature review on widely used intermittent demand forecasting techniques is presented at Section 2, followed by the experimental structure of our research and the description of the data set used (Section 3). Empirical results regarding the three heuristics examined are presented and discussed in Section 4. The implementation of the proposed three heuristics is explored in Section 5, through a dedicated forecasting support system. Finally, some managerial implications are mentioned (Section 6), while conclusions are summed up and avenues for future work are proposed in Section 7.

2. Intermittent demand forecasting approaches

Intermittent demand data or count data are frequently observed in industrial and inventory settings. Williams (1984) introduced a number of classification rules in order for a spare part to be characterized as slow moving, intermittent or lumpy. According to Teunter and Sani (2009), it is not an easy task to forecast intermittent demand, basically due to its erratic, and sometimes lumpy, nature. Nevertheless, it is very surprising that so little work has been done on forecasting intermittent demand data (Gooijer and Hyndman, 2005), with several industries and organizations relying on the single exponential smoothing (SES) method in order to forecast demand in a routine stock control system (Brown, 1959). As first shown by Croston (1972), the use of SES generally leads to inappropriate stock levels. As an alternative, Croston proposed the decomposition of the original intermittent series into two separate series. The first one includes all non-zero demand sizes, while the second series consists of the respective intervals between two consecutive non-zero demands. Each line is extrapolated separately, while the final forecast is simply calculated as a ratio of the two. Assuming \hat{z}_t and \hat{p}_t to be the forecasts of demand size and interval, respectively, for period t , Croston's forecast is given by:

$$\hat{Y}_t = \frac{\hat{z}_t}{\hat{p}_t}$$

In fact, later research in this field is heavily based on this single research by Croston. Willemain *et al.* (1994) and Johnston and Boylan (1996) have undertaken accuracy comparisons between SES and Croston's method, demonstrating the superiority of the latter, especially when the interval between demands exceeds 1.25 times the update period. Syntetos and Boylan (2001) proved that Croston's method is positively biased. Towards the correction of this behavior, they proposed a modification of the original Croston's method (Syntetos and Boylan, 2005), better known as Syntetos and Boylan Approximation (SBA). This new estimator is given by:

$$\hat{Y}_t = \left(1 - \frac{a}{2}\right) \frac{\hat{z}_t}{\hat{p}_t}$$

where a is the value of the exponential smoothing constant used in the extrapolation of the intervals series. Syntetos, in his PhD Thesis (2001), proposed another unbiased estimator, which can be obtained as follows:

$$\hat{Y}_t = \left(1 - \frac{a}{2}\right) \frac{\hat{z}_t}{\hat{p}_t - \frac{a}{2}}$$

Previous empirical studies (Syntetos and Boylan, 2001) have shown that the biased behavior of Croston's method is more apparent in the case of data with high intermittency (i.e. many periods with zero demand) when high smoothing values (a) are used. As a result, Croston's method is not recommended to be used with a values above 0.15. This result was verified by Teunter and Sani (2009), who analyzed the circumstances under which Croston's method and SBA approach tend to be biased. According to their findings, Croston's original method presents smaller bias if few demands are zero, whereas SBA modification has a smaller bias if many demands are zero. Moreover, it is argued that forecasts derived from Syntetos method are outperformed in terms of forecasting variance by the SB method (Syntetos, 2001; Teunter and Sani, 2009). Teunter and Sani (2009) suggest the use of Syntetos method as an alternative to Croston and SBA methods, but their findings are only based on simulated data.

In all three cases, the demand sizes and intervals are extrapolated using SES, while forecasts are updated only in periods with positive demand. The estimate of the demand under the SES method is given by (Makridakis *et al.*, 1998):

$$\hat{Y}_{t+1} = \hat{Y}_t + \alpha(Y_t - \hat{Y}_t) = \alpha \sum_{i=0}^{\infty} (1 - \alpha)^i Y_{t-i}$$

Although α -smoothing values in the range [0.05 – 0.2] are viewed as realistic (Croston, 1972; Willemain *et al.*, 1994; Johnston and Boylan, 1996), most of the empirical studies investigating intermittent demand assumed a constant value for α . The use of a smoothing constant is not a usual technique in fast moving series, where an optimization procedure takes place towards the selection of an ‘optimal’ smoothing parameter which minimizes the in-sample MSE. Moreover, Snyder (2002) argued for the use of different smoothing parameters for the demand sizes and intervals, as well as proposed the use of other compatible models and methods.

Recent research on intermittent demand forecasting has focused, other than forecasting performance in terms of accuracy, on the variability of intermittent demand estimates (Syntetos and Boylan, 2010), the importance of inventory obsolescence (Teunter *et al.*, 2011) and inventory performance under different types of information sharing (Ali *et al.*, 2012). Another new research route has been the temporal, non-overlapping aggregation of intermittent demand data into time buckets so that the resulting series are more likely to be non-intermittent. This technique has proven to be very promising in terms of forecasting accuracy (Nikolopoulos *et al.*, 2011) and customer service levels (Babai *et al.*, 2012). However, despite any research during the last 40 years, Croston’s method is very often applied in practice (Fildes *et al.*, 2008) and incorporated in commercial forecasting support systems. As a result, easy-to-apply modifications on the original framework that will lead in performance improvement are considered as beneficial.

3. Empirical Data & Experimental Structure

The empirical database used for the purposes of our research consists of the individual monthly demand histories of 3,000 SKUs from the automotive industry, over two years (24 consecutive monthly demand observations). The same database has been used in earlier studies (Syntetos and Boylan, 2005; Syntetos *et al.*, 2005). Detailed descriptive statistics (to the second decimal place) on the demand data series characteristics are presented in Table 1. It is worth mentioning that the data in hand are considered as “fast intermittent”, where the intermittent demand intervals are at any time less than 2, with

a median value at 1.26. The low degree of intermittence in this data set is coupled with low demand sizes with low degrees of variance. As a result, the data set is considered suitable as it contains series falling in all four classes of demand, according to Syntetos et al. (2005): erratic, lumpy, smooth or intermittent. At the same time, the empirical data are not to be related with issues regarding variability of the estimates and inventory obsolescence. This empirical dataset will not heavily affect the biased behavior of the Croston method, due to its low degree of intermittency.

3,000 SKUs	Demand Sizes		Demand Intervals		Demand per period	
	Mean	StDev	Mean	StDev	Mean	StDev
Min	1.00	0.00	1.04	0.21	0.54	0.50
25% ile	2.05	1.14	1.10	0.30	1.46	1.32
Median	2.89	1.76	1.26	0.52	2.33	1.92
75% ile	5.00	3.36	1.41	0.73	4.17	3.50
Max	193.75	101.42	2.00	1.60	129.17	122.75

Table 1. Demand data descriptive statistics

For the simulation purposes of the current research we held out the last 11 observations of each series, initializing all methods over the first 13 periods. We performed a sliding simulation (rolling evaluation) over the out-of-sample data via producing one-step-ahead forecasts; thus we calculated 11 one-step-ahead errors for each series, for each of the forecasting methods considered.

The evaluation of the results was performed by measuring the bias and accuracy of the examined methods. Mean Error (ME) offers a way to determine if an examined method is consistently positively or negatively biased, depending on the sign of the resulting value. ME can be calculated across all series using the following equation:

$$mean(ME) = \frac{1}{n \cdot h} \sum_{s=1}^n \sum_{t=1}^h (Y_t^s - \hat{Y}_t^s)$$

where Y_t^s and \hat{Y}_t^s are the actual and forecast values, respectively, of series s at time period t , n is the total number of series considered and h is the number of out-of-sample periods (horizon), thus $n=3,000$ and $h=11$. Accordingly, accuracy was calculated using average values of Mean as well as Median Absolute Scaled Error (MASE and MdASE respectively, Hyndman and Koehler, 2006). These two metrics are widely applicable, scale independent and easy to interpret: values of MASE greater than one indicate that

forecasts are worse, on average, than in-sample one step forecasts of the Naive method.

The average values of MASE and MdASE are given by:

$$mean(MASE) = \frac{1}{n \cdot h} \sum_{s=1}^n \sum_{t=1}^h \frac{|Y_t^s - \hat{Y}_t^s|}{\frac{1}{n-1} \sum_{i=1}^k |Y_i^s - Y_{i-1}^s|}$$

$$mean(MdASE) = \frac{1}{n} \sum_{s=1}^n median \left(\frac{|Y_t^s - \hat{Y}_t^s|}{\frac{1}{n-1} \sum_{i=2}^k |Y_i^s - Y_{i-1}^s|} \right)$$

where k is the number of in-sample periods, thus $k=13$.

Three estimators were used in terms of benchmarking: Naive, SES and Simple Moving Average (SMA). Naive forecasts are equal to the last actual demand, so:

$$\hat{Y}_{t+1} = Y_t$$

SES forecasts were generated using a constant level smoothing parameter, equal to 0.05. Last, SMA at length of 13 periods was used because this was the estimation procedure employed by the software manufacturer that provided the empirical data series used in this research (Syntetos and Boylan, 2005). The estimate of the demand under the SMA(13) is given by:

$$\hat{Y}_t = \frac{1}{13} \sum_{i=1}^{13} Y_{t-i}$$

4. Empirical Investigation & Discussion

Bias and accuracy results for the three benchmarks are reported in Table 2. Results for intermittent demand methods are also presented, when a smoothing constant in the range [0.05 – 0.2] is selected, the same for demand sizes and intervals. Overall, the most unbiased method for the examined dataset would be SBA method with $a=0.05$, followed by Naïve and SES. Both in Croston and SBA methods, an increase in value of the a smoothing constant leads to a greater absolute value of ME, resulting in more biased forecasts. ME for Croston's method is negative, denoting a positively biasing behavior, whereas in the case of SBA the bias has a negative direction. In terms of accuracy, SBA scores the lowest values for MASE and MdASE, at $a=0.15$ and 0.2, respectively. It is worth noting that, in contrast to bias, increased values of a smoothing constant have a positive effect on the measured (via MdASE) accuracy for both Croston

and SBA methods. Furthermore, we should also note the overall good performance of SES method, which scored the same accuracy level with SBA at $\alpha=0.05$. Empirical results of Table 2 indicate that Syntetos' method is not suitable for the examined data, being outperformed from Croston and SBA in terms of both bias and accuracy. As a result, further analysis of the current research is basically based on the performance of the latter methods.

Method	BIAS	ACCURACY	
	ME	MASE	MdASE
Naive	-0.054	1.127	2.86E-04
SMA(13)	-0.106	0.893	2.46E-04
SES(0.05)	-0.076	0.889	2.44E-04
Croston(0.05)	-0.108	0.896	2.46E-04
Croston(0.1)	-0.120	0.893	2.47E-04
Croston(0.15)	-0.136	0.896	2.46E-04
Croston(0.2)	-0.153	0.901	2.45E-04
SBA(0.05)	0.028	0.889	2.43E-04
SBA(0.1)	0.151	0.880	2.39E-04
SBA(0.15)	0.270	0.877	2.34E-04
SBA(0.2)	0.387	0.878	2.28E-04
Syntetos(0.05)	-0.383	0.912	2.54E-04

Table 2. Results of benchmarks and standard intermittent demand methods

4.1. Optimizing α smoothing parameter

The first heuristic of our research examines the optimization of α smoothing parameter, rather than using a constant value for demand sizes and intervals across all series. A linear optimization procedure takes place, where all values in the range [0.05 – 0.2] are examined separately using a step of 0.01, and the one minimizing the in-sample MSE is selected as the 'optimal', resulting in different α values for each series. The value of the in-sample MSE is given by:

$$MSE = \frac{1}{k} \sum_{t=1}^k (Y_t - \hat{Y}_t)^2$$

This linear optimization is a common practice for fast moving series, where smoothing parameters are selected in order to best fit the in-sample forecast model. In this case,

the optimization procedure is applied directly and exclusively to the decomposed series, i.e. the demands sizes and intervals, which may lead to different ‘optimal’ a values, as suggested by Snyder (2002).

Method	BIAS	ACCURACY	
	ME	MASE	MdASE
Croston(‘optimal’ a)	-0.122	0.895	2.46E-04
SBA(‘optimal’ a)	0.021	0.888	2.43E-04
Syntetos(‘optimal’ a)	-0.409	0.913	2.54E-04

Table 3. Optimizing a smoothing value

The results of this empirical heuristic for each intermittent demand method (Croston, SBA and Syntetos) are presented in Table 3. In terms of accuracy, the results are almost identical with those of the implementations of the methods where constant a smoothing value is set equal to 0.05 across all series. However, there is a significant reduction of the value of ME metric in the case of SBA method. The calculated bias drops to 0.021, which means a 25% error improvement. There is however not significant evidence that optimization benefits Croston and Syntetos methods.

	a -value	Demands (numerator)		Intervals (denominator)	
	13 in-sample observations	0.05-0.10	2897	96.5%	2547
0.11-0.15		17	0.6%	434	14.5%
0.16-0.20		86	2.9%	19	0.6%
18 in-sample observations	0.05-0.10	2620	87.3%	2546	84.8%
	0.11-0.15	60	2.0%	286	9.6%
	0.16-0.20	320	10.7%	168	5.6%
23 in-sample observations	0.05-0.10	2499	83.3%	2597	86.6%
	0.11-0.15	169	5.6%	193	6.4%
	0.16-0.20	332	11.1%	210	7.0%

Table 4. Distributions of optimal a -values

Table 4 presents the distributions of the optimal a -values for both demand and intervals in three instances of the rolling procedure. In more detail, the number of time series selecting as optimal a -values in the ranges [0.05-0.10], [0.11-0.15] and [0.16-0.20] along with the relevant

percentages are demonstrated. The three instances considered in this analysis were completed by 13, 18 and 23 in-sample observations respectively. A close observation of Table 4 makes it clear that smaller values of a are generally selected (inside the range [0.05-0.10]), especially in the case of short available history. As more observations become available, optimization enables a selective choice of greater a -values as well, for up to 17% of the time series.

4.2. The effect of rounding

When forecasting SKUs, providing decimal forecast values does not make much sense. This simple idea leads us to round the produced forecasts so that the reporting values would be whole numbers. Table 5 presents the results of the rounding effect, when indicative implementations of intermittent demand methods are used. In comparison with Table 2, the bias measured in all cases is almost at the same levels (if not even lower). Furthermore, there are notable improvements in terms of accuracy, as computed via MASE, where the calculated forecasts are approximately 2% more accurate for all methods tested.

Method	BIAS	ACCURACY	
	ME	MASE	MdASE
Croston(0.1, Round)	-0.112	0.878	2.50E-04
SBA(0.1, Round)	0.152	0.866	2.41E-04
SBA(0.2, Round)	0.379	0.861	2.35E-04
Syntetos(0.05, Round)	-0.382	0.897	2.50E-04

Table 5. Rounding SKUs forecasts

4.3. Combining Croston method with Theta model

The use of SES method in order to extrapolate the decomposed Croston's series has been criticized in many studies (see for example Snyder, 2002). We consider the use of an alternative, modern forecasting technique, the Theta model (the winner of M3 forecasting competition, Makridakis and Hibon, 2000), introduced by Assimakopoulos and Nikolopoulos (2000). Theta method decomposes the original series in two (or more) separate series (the so-called theta lines), whose primary qualitative characteristic is the better approximation of the long-term behavior of the data or the augmentation of short-term features, depending on the value of the Theta coefficient.

These theta lines are extrapolated separately. At the current study, we implement the Classic Theta model, as a three steps procedure:

1. Each time-series is decomposed into two Theta lines, the linear regression line (which is referred also as Theta Line ($\Theta=0$)) and the Theta Line ($\Theta=2$), which is calculated as follows:

$$ThetaLine(\Theta = 2)_t = 2Y_t - LRL_t$$

Where Y_t refers to the t -th actual observation of the raw data, while LRL_t denotes the t -th observation of the linear regression line, expressing the linear relationship between raw data and time.

2. The linear regression line is extrapolated in the usual way while the second line is extrapolated via Single Exponential Smoothing.
3. The forecasts produced from the extrapolation of the two lines are combined with equal weights.

Thus, we investigate the use of Theta model, which can replace SES method in either numerator or denominator of Croston's ratio, or even in both. Originally, the combination of Croston with Theta was proposed by Nikolopoulos *et al.* (2007), where Theta model was applied just for the extrapolation of the demand sizes (numerator). The intuition for using this combination was that Croston-Theta could pick up trends of non-stationary series, thus displaying more potent predictive power.

Method	BIAS	ACCURACY	
	ME	MASE	MdASE
Croston-Theta(0.05, Num & Denom)	-0.056	1.393	2.66E-04
Croston-Theta('optimal' a , Num & Denom)	-0.046	1.392	2.66E-04
Croston-Theta('optimal' a , Denom)	0.011	1.200	2.55E-04
Croston-Theta('optimal' a , Num)	-0.262	0.933	2.55E-04
Croston-Theta(0.05, Num)	-0.251	0.931	2.55E-04
Croston-Theta(0.05, Num, Round)	-0.253	0.916	2.50E-04

Table 6. Combining Croston with Theta

The results of Croston-Theta combination are presented in Table 6. Each row of the table displays the results for a separate implementation of the approach, in terms of the selected a smoothing value (0.05 or 'optimal', as discussed in Subsection 4.1), the level at which Theta was applied (numerator, denominator or both) and, last, the application

(or not) of the rounding heuristic (as discussed in Subsection 4.2) at the final forecasts. The results indicate significant improvements in terms of bias, when Theta model is used for the extrapolation of intervals (denominator). Specifically Croston-Theta('optimal' a , Denom) implementation has the best bias performance in comparison to all other implementations presented in the current study, with an improvement close to 50% from the second best implementation (SBA('optimal' a)). On the other hand, the Croston-Theta combination does not pay back in terms of accuracy. The results indicate worse out-of-sample accuracy performance than the in-sample accuracy under Naïve, when Theta model is applied to the denominator. Lastly, the rounding effect seems to work once again as a simple self-improvement heuristic, offering notable improvements in the resulting accuracy metrics, while keeping the bias level constant. The moderate performance of Croston-Theta combination in terms of accuracy could be interpreted as lack of the trend component in the examined empirical dataset. Even if the trend component is almost zero for the majority of intermittent demand series, the presence of a deterministic trend would considerably favor the use of Theta model over SES.

4.4. Discussion

The use of different smoothing parameters for the demand and intervals has been previously suggested (Snyder, 2002). However, appropriately choosing smoothing values independently for the numerator and the denominator of each one of the series, via a widely used cost function (MSE), has not been previously proposed nor investigated. In fact, this practice drives in considerable gains in terms of bias, while at the same time no negative impacts on accuracy are recorded. The most important observation, however, derives from comparing Tables 2 and 3. It is pretty clear, that as the static value of the smoothing parameter (a) raises from 0.05 to 0.20, there is a significant deterioration of the bias metric (measured as ME). However, selecting the most suitable smoothing value independently for each series, through minimizing in-sample error, leads to significantly better results (up to 25%). This practically means that the use of a cost function for selecting the best smoothing parameters per series can make a difference.

In many industrial applications, especially those involving time series of spare parts or SKUs, non-integer point forecasts are considered as non realistic. The impacts of rounding the final point forecasts derived from intermittent demand methods are

empirically examined in this research. The most significant result is that this technique results in better accuracy levels (up to 2%) while at the same time no deterioration in terms of bias is recorded. As a result, we strongly recommend the use of this heuristic, which is regarded as appealing in both empirical and practical terms.

Finally, an alternative to the traditional Croston's method was examined in Subsection 4.3. Originally, Croston proposed the use of SES as the extrapolation procedure for both decomposed series (demand and intervals). Given that both decomposed series represent sequences of non-zero values, we consider, for the first time, the full potential of using if the Theta model as the extrapolation technique of either the nominator, denominator or both. In fact, the choice of this model lies in its superior performance, as recorded in past major international forecasting competitions. This simple technique allows Croston's framework act almost in an unbiased way (improvement up to 90% from the original approach), when Theta model is used for the extrapolation of the intervals between non-zero demands.

5. The forecasting support system (FSS)

In many applications, statistical forecasts are produced via dedicated and autonomous forecasting support systems (FSS). This direction serves multiple purposes. Firstly, managers and practitioners may not be familiar with the necessary statistical background. Secondly, an automated FSS can handle, pre-process and forecast thousands of time series (sales or orders for SKUs) in just a few seconds. Lastly, many features of the modern FSSs, such as statistical analysis, handling the impacts of special periods and integration of judgmental interventions, are regarded as necessary to the forecasting process of any industry. The needs of the empirical analysis of the current research led us to the design and development of a unique and dedicated forecasting support system for handling data of intermittent nature.

The purpose of the current section is to give the general guidelines towards the implementation of a FSS that fully implements the proposed heuristics analyzed in the current research, so enabling the managers to have direct access to any gains derived through their practice in any manufacturing or industrial setting. Moreover, we aim to give insights to practitioners already utilizing customized software as to which directions of additional implementation or external modules should they target for exploiting these heuristics.

The Intermittent Demand Forecasting System (IDFS) was designed following a three-tier physical architecture (visualization, statistical/business and data). The main advantage of this architecture is that it is easily serviceable and expandable. The first layer of the architecture is the user interface, where both graphical and numerical interpretation of the data and the results are displayed. Furthermore, this layer enables users towards a detailed selection of the parameters related to the forecasting process (forecasting methods, horizons, hold-out samples, error metrics), along with the usage of the three heuristics presented in this paper. Figure 1 and Figure 2 show two typical screen displays of the proposed system. The second layer is the statistical/business tier. All the statistical and methodological procedures are modeled and implemented in the middle layer, which includes the functions related to the forecasting process. This layer includes all original forecasting methodologies (for example SES, Theta model, Croston, SBA and Syntetos) and also allows methods to interact with each other (for example, towards the formation of Croston-Theta). Additional adjustments, such as the use of optimized over contact smoothing values or the rounding of the final forecasts, are passing as external variables, through the interaction with the first layer. In addition, a middle tier is used to create the base for further extension with external software, encompassing interfaces, wrappers and web-services necessary for data exchange. Lastly, the third layer consists of the data tier of the application. The data base management system (database, views and relations) is lying on the data base server, a windows-based machine isolated from the internet in order to avoid threats and provide satisfactory response times. The data layer provides the statistics/business layer with the required historical data and stores any forecasts and accuracy results.




Figure 1. Graphical visualization of the input data and parameter initialization

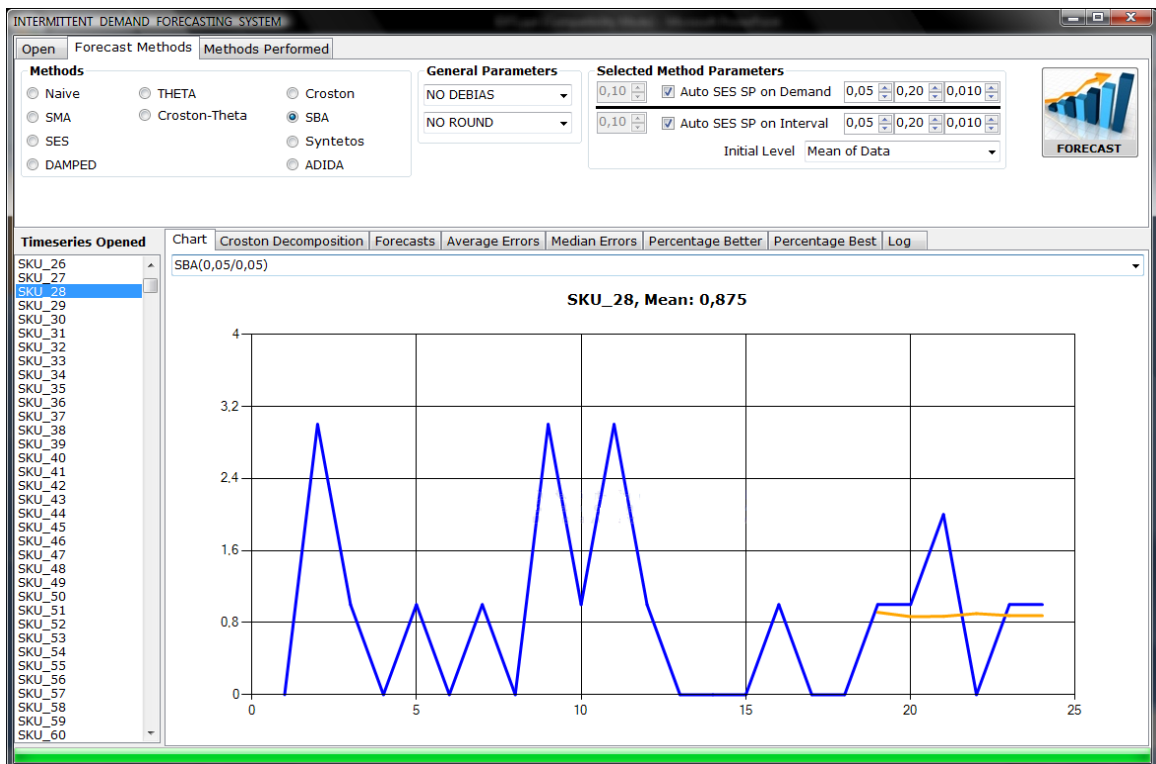


Figure 2. Forecasting methods/parameters and graphical interpretation of the empirical results

IDFS was developed using Microsoft's Visual Basic .NET 2008 while the Dundas Chart, for Visual Basic .NET, was employed for the system implementation, in regards to its advanced charting functionality and superior graphic options. Finally, the Microsoft SQL Server 2008 R2 database is utilized by IDFS to store and retrieve the required information for the data analysis and forecasting.

6. Managerial Implications

Heuristics linked with the simplicity of spreadsheets are considered managerial appealing and flexible, especially in the case of slow-moving items and when dealing with problems of different sizes (Hummel and Jesse, 1990). As a result, optimizing a smoothing parameter and rounding final forecasts seem also to be interesting from a managerial point of view. Both heuristics can easily be implemented by managers and practitioners, while offering notable gains regarding accuracy and/or bias.

To begin with, optimizing procedures for smoothing parameters of exponential smoothing methods are implemented and automated in major forecasting packages. In that sense, these procedures can easily be employed in a practical inventory setting, resulting in substantial improvements for the bias of SBA estimator (up to 25%), while scoring good accuracy levels (compared to SBA(0.05) implementation).

Moreover, the task of rounding the final forecasts derived from any intermittent demand method can be easily done by use of ordinary spreadsheets. This simple heuristic offers remarkable improvements in measured forecasting accuracy, keeping, at the same time, bias at lower levels. Forecasts for intermittent demand SKUs call for rounding, so as to make forecasts interpretable and directly usable for real supply chain management applications, such as order placement.

7. Conclusions & Perspectives

The current study examined the empirical effectiveness of three empirical heuristics towards the modification of commonly used forecasting approaches for intermittent demand. We proposed the use of non-constant a smoothing parameter, via selecting the best a value for each series through in-sample optimization. Moreover, these values may be different for the demand sizes and the intervals (Snyder, 2002). The use of SES method for the extrapolation of the decomposed series in Croston's framework is also examined. We considered as an alternative the Theta model, a technique that

outperformed all exponential smoothing methods in the M3 forecasting competition. Lastly, an intuitively appealing heuristic, concerning the rounding of the final forecasts, was proposed.

The results indicate that ‘optimal’ selection of α smoothing values results in almost identical accuracy levels, while in some cases there are significant improvements on the bias. Thus, model optimization is feasible and does pay back. Rounding seems to work surprisingly well, offering notable improvements in terms of accuracy and keeping bias constant. These results render this simple heuristic suitable in cases of data sets consisting of SKUs. Croston-Theta’s performance was moderate, as it performed well as far as bias is concerned, but seemed problematic in terms of accuracy. At last, a specialized FSS for intermittent demand data was proposed.

Further research should involve the interaction of the rounding heuristic from a theoretical point of view. Moreover, all experiments of this study could be replicated with different data sets, so as to reach more general conclusions about the three proposed heuristics. Specifically, the Croston-Theta combination should be tested thoroughly with simulated and field trended data. Finally, it is recently argued (Syntetos *et al.*, 2010) that in an inventory forecasting setting extrapolation methods should not only be evaluated with respect to their forecast accuracy but also in terms of their stock control implications, as measured through accuracy implication metrics (such as inventory costs and service levels achieved). Exploring the effects of the examined heuristics on stock control is an interesting line for further research and certainly worthwhile pursuing from a practitioner’s perspective.

References

Ali M.M., Boylan J.E. and Syntetos A.A. (2012), “Forecast errors and inventory performance under forecast information sharing”, *International Journal of Forecasting*, Vol. 28, pp. 830-841.

Assimakopoulos V. and Nikolopoulos N. (2000), “The theta model: a decomposition approach to forecasting”, *International Journal of Forecasting*, Vol. 16, pp. 521-530.

Babai M.Z., Ali M.M. and Nikolopoulos K. (2012), “Impact of temporal aggregation on stock control performance of intermittent demand estimators: Empirical analysis”, *Omega*, Vol. 40, pp. 713-721.

Brown R. (1959) *Statistical Forecasting for Inventory Control*, McGraw-Hill, New York.

Croston J.D. (1972), “Forecasting and Stock Control for Intermittent Demands”, *Operational Research Quarterly*, Vol. 23, pp. 289-303.

- De Gooijer J.G and Hyndman R.J. (2005), "25 years of time series forecasting", *International Journal of Forecasting*, Vol. 22, pp. 443-473.
- Fildes R., Nikolopoulos K., Crone S.F. and Syntetos A.A. (2008) "Forecasting and operational research: A review", *Journal of the Operational Research Society*, Vol. 59, pp. 1150-1172.
- Huang L.T., Hsieh I.C. and Farn C.K. (2011), "On ordering adjustment policy under rolling forecast in supply chain planning", *Computers and Industrial Engineering*, Vol. 60, pp. 397-410.
- Hummel J.W. and Jesse R.R. (1990), "A spreadsheet heuristic approach for the stocking and retention of slow-moving, obsolescent items", *Computers and Industrial Engineering*, Vol. 18, pp. 163-173.
- Hyndman R.J. and Koehler A.B. (2006), "Another look at measures of forecast accuracy", *International Journal of Forecasting*, Vol. 22, pp. 679-688.
- Johnston F.R. and Boylan J.E. (1996), "Forecasting intermittent demand: A comparative evaluation of Croston's method", *International Journal of Forecasting*, Vol. 12, pp. 297-298.
- Johnston F.R., Boylan J.E. and Shale E.A. (2003), "An examination of the size of orders from customers, their characterization and the implications for inventory control of slow moving items", *Journal of the Operational Research Society*, Vol. 54, pp. 833-837.
- Makridakis S., Wheelwright S.C. and Hyndman R.J. (1998), *Forecasting: Methods and Applications* (3rd ed.), Wiley, New York, NY.
- Makridakis S. and Hibon M. (2000), "The M3-Competition: results, conclusions and implications", *International Journal of Forecasting*, Vol. 16, pp. 451-476.
- Nikolopoulos K., Syntetos A.A. and Babai M.Z. (2007), "A new intermittent demand approach via combining Croston's method and the Theta model", paper presented at the *22nd European Conference on Operational Research EURO XXII*, July 8-11, 2007, Prague, Czech Republic.
- Nikolopoulos K., Syntetos A.A., Boylan J.E., Petropoulos F. and Assimakopoulos V. (2011), "An aggregate disaggregate intermittent demand approach (ADIDA) to forecasting: an empirical proposition and analysis", *Journal of the Operational Research Society*, Vol. 62, pp. 544-554.
- Snyder R. (2002), "Forecasting sales of slow and fast moving inventories", *European Journal of Operational Research*, Vol. 140, pp. 684-699.
- Syntetos A.A. and Boylan J.E. (2001), "On the bias of intermittent demand estimates", *International Journal of Production Economics*, Vol. 71, pp. 457-466.
- Syntetos A.A. (2001), *Forecasting for Intermittent Demand*. Brunel University: Unpublished Ph.D thesis.
- Syntetos A.A. and Boylan J.E. (2005), "The accuracy of intermittent demand estimates", *International Journal of Forecasting*, Vol. 21, pp. 303-314.
- Syntetos A.A., Boylan J.E. and Croston J.D. (2005), "On the categorization of demand patterns", *Journal of the Operational Research Society*, Vol. 56, pp. 495-503.

Syntetos A.A and Boylan J.E. (2010), "On the variance of intermittent demand estimates", *International Journal of Production Economics*, Vol. 128, pp. 546-555.

Syntetos A.A., Nikolopoulos K. and Boylan J.E. (2010), "Judging the judges through accuracy-implication metrics: the case of inventory forecasting", *International Journal of Forecasting*, Vol. 26, pp. 134-143.

Teunter R. and Sani B. (2009), "On the bias of Croston's forecasting method", *European Journal of Operational Research*, Vol. 194, pp. 177-183.

Teunter R.H., Syntetos A.A. and Babai M.Z. (2011), "Intermittent demand: Linking forecasting to inventory obsolescence", *European Journal of Operational Research*, Vol. 214, pp. 606-615.

Willemain T.R., Smart C.N., Shockor J.H. and DeSautels P.A. (1994), "Forecasting intermittent demand in manufacturing: A comparative evaluation of Croston's method", *International Journal of Forecasting*, Vol. 10, pp. 529-538.

Williams T.M. (1984), "Stock control with sporadic and slow-moving demand", *Journal of the Operational Research Society*, Vol. 35, pp. 939-948.