

INTERMITTENT DEMAND FORECASTING AND STOCK CONTROL: AN EMPIRICAL STUDY

Adriano O. Solis^(a), Letizia Nicoletti^(b), Somnath Mukhopadhyay^(c), Laura Agosteo^(d),
Antonio Delfino^(e), Mirko Sartiano^(f)

^(a) School of Administrative Studies, York University, Toronto, Canada

^{(b) (d) (e) (f)} MSC-LES, Mechanical Department University of Calabria, Rende, Italy

^(c) Information & Decision Sciences Department, The University of Texas at El Paso, El Paso, Texas, U.S.A.

^(a) asolis@yorku.ca, ^(b) letizia@msc-les.org, ^(c) smukhopadhyay@utep.edu, ^(d) agosteo.laura@libero.it,
^(e) delanto@hotmail.it, ^(f) mirko.sartiano@yahoo.it

ABSTRACT

Statistical accuracy measures are generally used to assess the effectiveness of demand forecasting methods. In the final analysis, however, these methods should be judged according to whether they actually lead to better inventory control performance. We empirically evaluate four methods (simple moving average, single exponential smoothing, Croston's method, and the Syntetos-Boylan approximation) in terms of statistical forecast accuracy and, more importantly, inventory system efficiency. We apply four forecasting methods to an industrial dataset involving more than 1000 stock-keeping units of a firm in the professional electronics sector. Demand is often intermittent, erratic, or both (i.e., lumpy). We devise and use a two-stage distribution involving uniform and negative binomial distributions to model the actual demand distribution, where possible. We then simulate the stock control performance of a (T,S) inventory system with respect to target customer service levels.

Keywords: intermittent/lumpy demand forecasting, forecast accuracy, inventory control, order-up-to periodic review system, simulation

1. INTRODUCTION

Demand for a stock-keeping unit (SKU) is said to be *intermittent* if there are periods in which demand is zero. When demand is intermittent and there are large variations in demand sizes, demand is said to be *lumpy*. Syntetos, Boylan, and Croston (2005) proposed a theoretically coherent scheme for categorizing demand into smooth, erratic, intermittent, or lumpy. In this categorization scheme, the average inter-demand interval (*ADI*) and the squared coefficient of variation (CV^2) of demand are compared with cutoffs of 1.32 for *ADI* and 0.49 for CV^2 , as follows:

- *smooth* demand: $ADI < 1.32$, $CV^2 < 0.49$;
- *erratic* demand: $ADI < 1.32$, $CV^2 > 0.49$;
- *intermittent* demand: $ADI > 1.32$, $CV^2 < 0.49$;
- *lumpy* demand: $ADI > 1.32$, $CV^2 > 0.49$.

This categorization scheme has been cited and applied by various researchers (e.g., Ferrari, Pareschi, Regattieri, and Persona 2006; Gutierrez, Solis, and Mukhopadhyay 2008; Altay, Rudisill, and Litteral 2008; Mukhopadhyay, Solis, and Gutierrez 2011).

In the intermittent demand forecasting literature, many papers have been published on the relative performance with respect to statistical measures of accuracy of various forecasting methods, most notably simple exponential smoothing (SES), Croston's method (Croston 1972), and an estimator proposed by Syntetos and Boylan (2005). Schultz (1987) suggested that separate smoothing constants, α_i and α_s , be used for updating the inter-demand intervals and the nonzero demand sizes, respectively, in place of Croston's single smoothing constant α . We note, however, that Mukhopadhyay, Solis, and Gutierrez (2011) investigated separate smoothing constants, α_i and α_s , in forecasting lumpy demand and did not observe any substantial improvement in forecast accuracy.

Syntetos and Boylan (2001) pointed out a positive bias in Croston's method arising from an error in his mathematical derivation of expected demand. They proposed (Syntetos and Boylan 2005) a correction factor of $\left(1 - \frac{\alpha_i}{2}\right)$ – where α_i is the smoothing constant

used in updating the inter-demand interval estimate – to be applied to Croston's estimator of mean demand. The revised estimator is now often referred to (e.g., Gutierrez, Solis, and Mukhopadhyay 2008; Boylan, Syntetos, and Karakostas 2008; Babai, Syntetos, and Teunter 2010; Mukhopadhyay, Solis, and Gutierrez 2011) in the intermittent demand forecasting literature as the Syntetos-Boylan approximation (SBA).

We also evaluate the 13-month simple moving average (SMA13) method, which is based upon dividing the 52 weeks in a year into 13 four-week “months”. SMA13 has been applied in a number of recent intermittent demand forecasting studies (e.g., Syntetos and Boylan 2005, 2006; Boylan, Syntetos, and Karakostas 2008) in view of its being built into some commercially available forecasting software.

This paper is organized as follows. In section 2, we discuss the forecasting methods under evaluation, the statistical measures of forecast accuracy that we use, and the nature of the industrial dataset and how data partitioning is performed. In the next section, we propose a two-stage approach to the modeling of demand distribution. We proceed to report on our empirical investigation of forecasting performance, based upon statistical accuracy measures, on the performance block of the actual data and on the simulated demand distribution. The performance of the forecasting methods in terms of inventory systems efficiency is reported in section 4. We present our conclusions in the final section.

2. FORECASTING METHODS AND DEMAND DATA

2.1. Forecasting Methods and Accuracy Measures

We evaluate four methods that are well-referenced in the intermittent demand forecasting literature: SMA13, SES, Croston's, and SBA. For the latter three methods which involve an exponential smoothing constant α , low values of α up to 0.20 have generally been suggested for lumpy demand (e.g., Croston 1972; Johnston and Boylan 1996). We evaluate four α values of 0.05, 0.10, 0.15, and 0.20 as used in a number of recent studies (e.g., Syntetos and Boylan 2005; Gutierrez, Solis, and Mukhopadhyay 2008; Mukhopadhyay, Solis, and Gutierrez 2011).

In the following sections where our results are presented, we only report those pertaining to SBA and not those for Croston's method, as we have found the former to consistently outperform the latter.

Gutierrez, Solis, and Mukhopadhyay (2008) and Mukhopadhyay, Solis, and Gutierrez (2011) found a neural network (NN) model, when applied to an industrial dataset exhibiting lumpy demand, to perform better overall than the SES, Croston's and SBA methods across different scale-free error measures. However, NN modeling requires a substantial number of time periods to 'train' or calibrate the model, which is not the case in the current study.

In addition to applying the more commonly used root mean squared error (RMSE) and mean absolute deviation (MAD) as forecast accuracy measures, we have also used mean absolute percentage error (MAPE), which is the most widely used accuracy measure for ratio-scaled data. The traditional MAPE definition, which involves terms of the form $|E_t|/A_t$ (where A_t and E_t , respectively, represent actual demand and forecast error in period t), fails when demand is intermittent. We applied an alternative specification (e.g., Gilliland 2002) of MAPE as a ratio estimate, which guarantees a nonzero denominator:

$$\text{MAPE} = \left(\frac{\sum_{t=1}^n |E_t|}{\sum_{t=1}^n A_t} \right) \times 100. \quad (1)$$

Eaves and Kingsman (2004) used the same three error statistics (MAPE, RMSE, and MAD) in comparing the performance of several methods (SES, Croston's, SBA, 12-month simple moving average, and the previous year's simple average) in forecasting demand for spare parts for in-service aircraft of the Royal Air Force (RAF) of the UK. They found SBA to provide the best results overall using MAPE, but the 12-month simple moving average yielded the best MADs overall. Willemain, Smart, Schockor, and DeSautels (1994) used median absolute percentage error (MdAPE), in addition to MAPE, RMSE, and MAD, as forecast accuracy measures to compare performance of SES and Croston's methods in intermittent demand forecasting. Noting that relative results were the same for all four measures, they reported only MAPEs.

2.2. Industrial Dataset and Partitioning

In the current study, we apply the four methods to an industrial dataset involving about 1500 items generally held in stock at a distribution center and a number of manufacturing plants of a firm operating in the professional electronics sector. The raw data consist of individual transactions as reported within the company's enterprise resource planning system, representing actual stock withdrawals. We initially aggregate the transactional data into weekly usage quantities, and further aggregate these usage quantities based on 13 four-week "months" in a calendar year. In this case, the monthly usage quantities do not constitute actual demand quantities, as the inventory on hand at the time of a transaction may not meet the required quantity. However, since demand is not traditionally tracked as well as actual usage in a transaction-based system, we treat monthly usage quantity as a surrogate measure of monthly demand.

This process yielded 66 months of "demand" data, which we broke down into initialization, calibration, and performance measurement blocks (as in Boylan, Syntetos, and Karakostas 2008) with each block consisting of 22 months in our study.

For each of the SES, Croston's and SBA methods, we selected α based upon the minimum MAPE attained in the calibration block for use as the smoothing constant in the performance block.

Given that the various SKUs represent end items, sub-assemblies, components, and spare parts that are used for building projects, retail sales, or servicing of professional electronic products, it is understandable that we found many of them to actually exhibit erratic or lumpy demand based on the earlier cited categorization scheme (Syntetos, Boylan, and Croston 2005).

In this paper, we report findings on a limited sample consisting of ten SKUs, with demand statistics presented in Table 1. These ten SKUs are not representative of our entire dataset. They were selected principally to demonstrate the approach we have taken, as well as to illustrate the results we have obtained thus far in both the empirical investigation of forecasting

performance and the empirical investigation of inventory control performance. The first nine SKUs are lumpy, while the last one is categorized as erratic (although its *ADI* of 1.27 is just below the cutoff of 1.32 as specified for lumpy demand).

Table 1: Sample of 10 SKUs

SKU #	1	2	3	4	5
Mean	10.97	1.44	0.71	9.74	2.82
Std Dev	13.30	2.23	1.76	13.82	7.19
CV^2	1.47	2.41	6.11	2.01	6.52
<i>ADI</i>	1.35	2.00	4.43	1.65	4.40
% of Zero Demand	27.3%	50.0%	80.3%	42.4%	80.3%
Demand Category	Lumpy	Lumpy	Lumpy	Lumpy	Lumpy

SKU #	6	7	8	9	10
Mean	3.47	4.85	9.27	3.03	6.37
Std Dev	5.01	6.76	19.35	8.03	7.85
CV^2	2.09	1.95	4.35	7.02	1.52
<i>ADI</i>	1.61	1.65	3.94	4.13	1.27
% of Zero Demand	37.9%	39.4%	77.3%	78.8%	24.2%
Demand Category	Lumpy	Lumpy	Lumpy	Lumpy	Erratic

2.3. Modeling of Intermittent/Lumpy Demand

A number of recent studies have referred to the use of a negative binomial distribution (NBD) to model the distribution of intermittent or lumpy demand items (e.g., Syntetos and Boylan 2006; Boylan, Syntetos, and Karakostas 2008; Syntetos, Babai, Dallery, and Teunter 2009). Syntetos and Boylan (2006) have argued that the NBD satisfies both theoretical and empirical criteria.

The NBD with parameters r and p , where $0 < p \leq 1$ and $r > 0$, is given by the discrete density function (e.g., Mood, Graybill, and Boes 1974):

$$f(x; r, p) = \binom{r+x-1}{x} p^r (1-p)^x I_{\{0,1,2,\dots\}}(x), \quad (2)$$

where the parameters p and r are a probability of “success” and a target number of successes, respectively. A realization x of the random variable X in this case represents a number of failures before the r th success is attained. The NBD has mean

$$\mu = E[X] = \frac{r(1-p)}{p} \quad (3)$$

and variance

$$\sigma^2 = V[X] = \frac{r(1-p)}{p^2}. \quad (4)$$

Since $V[X] = E[X]/p$, it follows that the variance of the NBD is greater than its mean.

When $r = 1$, the NBD reduces to a geometric (or Pascal) distribution with density

$$f(x; p) = p(1-p)^x I_{\{0,1,2,\dots\}}(x). \quad (5)$$

To generate an NBD to approximate the distribution of a random variable with mean μ and

variance σ^2 , we simultaneously solve (3) and (4) and obtain:

$$\hat{p} = \frac{\mu}{\sigma^2} \quad (6)$$

and

$$\hat{r} = \frac{\mu^2}{\sigma^2 - \mu}, \quad (7)$$

as initial estimates of the NBD parameters (using the mean \bar{x} and the variance s^2 of the 66-month actual demand time series as values of μ and σ^2 , respectively). These expressions for \hat{p} and \hat{r} , however, represent values in the set \Re of real numbers, while the NBD parameter r is supposed to be integer-valued. In applying (6) and (7), we generally obtain a non-integer value of \hat{r} . Thus, in seeking to simulate the actual demand distributions, we have investigated the rounded up and rounded down values of \hat{r} while adjusting the value of \hat{p} . However, we have found that, for many of the SKUs under study, it is not possible to obtain adjusted values of \hat{r} and \hat{p} that would lead to an NBD with mean, variance, CV^2 , and *ADI* that are reasonably close to those of the actual demand distribution.

For a SKU with a proportion z of periods with zero demand is relatively high, we simulate the demand distribution by way of a two-stage process: a uniform distribution in stage 1 and an NBD in stage 2. Stage 1 involves an appropriately determined probability z_1 of zero demand, taking into consideration both z as well as the NBD in stage 2. We determine the mean and variance of the nonzero demands in the actual distribution and use these to calculate first approximations of the NBD parameters in stage 2. We test rounded up and down values of the parameter estimate \hat{r} and refine the parameter estimate \hat{p} as the values of mean, variance, CV^2 , and *ADI* of the actual and simulated demand distributions are compared.

We attempted to simulate demand distributions of the SKUs under study applying our two-stage approach, performing 100 runs each consisting of 100 four-week “months”, for a total of 10,000 months in each experiment. We used AnyLogic as our simulation platform, but with some code written in Java to handle mathematical modeling which could not be readily undertaken within the standard AnyLogic library. We accordingly selected z_1 for stage 1 and the parameters r and p of the NBD in stage 2 based on what appeared to yield the best combination of values of mean, variance, CV^2 , and *ADI* of the simulated distribution in comparison with those of the actual distribution.

Generally, we found the percentages of zero demand and *ADIs* generated by our simulation approach to be reasonably close to those of the actual demand data. However, because the mean and standard

deviation of the simulated distribution using a combination of values of z_1 , r and p are usually not fully consistent with those of the actual distribution, we tended to favor standard deviation and CV^2 , which measure variability of demand sizes, in selecting the final parameter values.

Table 2 summarizes the parameters as determined and used for simulating the demand distributions of our sample of ten SKUs, applying our two-stage approach. In the currently reported sample of ten SKUs, the CV^2 values for the simulated distributions are greater than, but less than 120% of, CV^2 values for the actual data – except for SKUs 2 and 6 (with simulated values being 98% and 96% of actual CV^2). The percentages of zero demand and $ADIs$ are generally close to the actual demand distribution values.

Table 2: Simulated Demand Distributions for Sample of 10 SKUs Using Two-Stage Approach

SKU #	1	2	3	4	5
Actual Demand Distribution					
Mean	10.97	1.44	0.71	9.74	2.82
Std Dev	13.30	2.23	1.76	13.82	7.19
CV^2	1.47	2.41	6.11	2.01	6.52
ADI	1.35	2.00	4.43	1.65	4.40
% of Zero Demand (z)	27.3%	50.0%	80.3%	42.4%	80.3%
Simulated Demand Distribution					
<i>Stage 1: Uniform</i>					
z_1	21.9%	25.4%	79.1%	38.5%	80.0%
<i>Stage 2: NBD</i>					
r	1	1	7	1	2
p	0.0691	0.3300	0.6667	0.0640	0.1230
<i>Two-Stage Results</i>					
Mean	10.26	1.47	0.72	9.07	2.82
Std Dev	13.30	2.23	1.77	13.79	7.24
CV^2	1.71	2.36	6.27	2.34	6.77
ADI	1.38	1.98	4.89	1.72	4.96
% of Zero Demand	27.3%	49.5%	79.9%	42.2%	80.3%

SKU #	6	7	8	9	10
Actual Demand Distribution					
Mean	3.47	4.85	9.27	3.03	6.37
Std Dev	5.01	6.76	19.35	8.03	7.85
CV^2	2.09	1.95	4.35	7.02	1.52
ADI	1.61	1.65	3.94	4.13	1.27
% of Zero Demand (z)	37.9%	39.4%	77.3%	78.8%	24.2%
Simulated Demand Distribution					
<i>Stage 1: Uniform</i>					
z_1	25.1%	30.6%	77.3%	77.1%	14.2%
<i>Stage 2: NBD</i>					
r	1	1	5	1	1
p	0.1711	0.1263	0.1085	0.0740	0.1171
<i>Two-Stage Results</i>					
Mean	3.55	4.70	9.09	2.87	6.45
Std Dev	4.98	6.80	19.35	8.08	7.96
CV^2	2.01	2.14	4.80	8.16	1.53
ADI	1.61	1.66	4.38	4.62	1.32
% of Zero Demand	37.7%	39.7%	77.1%	78.8%	24.4%

3. FORECASTING PERFORMANCE

3.1. Forecast Accuracy: Performance Block

The exponential smoothing constant α selected from among the candidate values (0.05, 0.10, 0.15, or 0.20) for each of the SES and SBA methods, based upon the minimum MAPE in the calibration block, are shown in Table 3. When accordingly applying SMA13, SES, and SBA to actual demand data in the performance block, the resulting error statistics are likewise reported in the same table. There does not appear to be a method that exhibits superior performance across the ten SKUs.

Table 3: Error Statistics when Applying Various Methods to the Performance Block

SKU #	1	2	3	4	5
Smoothing Constants Selected in Calibration Block					
SES	0.20	0.20	0.20	0.10	0.20
SBA	0.05	0.05	0.05	0.10	0.05
MAPE					
SMA13	90.36%	160.84%	150.43%	153.09%	162.60%
SES	85.89%	157.94%	163.77%	144.16%	151.70%
SBA	82.83%	196.25%	122.93%	140.08%	169.21%
Best MAPE	SBA	SES	SBA	SBA	SES
MAD					
SMA13	9.570	0.804	1.846	10.647	4.287
SES	9.097	0.790	2.010	10.025	3.999
SBA	8.772	0.981	1.509	9.742	4.461
Best MAD	SBA	SES	SBA	SBA	SES
RMSE					
SMA13	11.486	0.946	2.665	12.026	5.792
SES	10.925	0.876	2.752	11.423	5.702
SBA	10.723	1.076	2.681	10.914	5.683
Best RMSE	SBA	SES	SMA13	SBA	SBA

SKU #	6	7	8	9	10
Smoothing Constants Selected in Calibration Block					
SES	0.20	0.05	0.05	0.15	0.20
SBA	0.05	0.05	0.05	0.05	0.20
MAPE					
SMA13	102.62%	85.84%	177.91%	284.23%	98.83%
SES	110.62%	89.14%	191.34%	261.61%	98.25%
SBA	122.19%	90.04%	201.22%	255.94%	100.45%
Best MAPE	SMA13	SMA13	SMA13	SBA	SES
MAD					
SMA13	2.332	6.594	9.490	5.168	5.301
SES	2.514	6.848	10.176	4.756	5.270
SBA	2.777	6.795	9.906	4.653	5.388
Best MAD	SMA13	SMA13	SMA13	SBA	SES
RMSE					
SMA13	3.312	9.064	15.419	6.718	5.977
SES	3.409	9.488	14.924	6.293	6.063
SBA	3.245	9.797	14.956	6.234	6.156
Best RMSE	SBA	SMA13	SES	SBA	SMA13

3.2. Forecast Accuracy: Simulated Demand

When applying the methods to the simulated demand distributions, however, we see in Table 4 that the SBA method outperforms SMA13 and SES overall across the three error statistics in nine out of the ten SKUs. This suggests the overall superiority of SBA over a sufficiently longer time frame.

4. INVENTORY CONTROL PERFORMANCE

Traditionally, demand forecasting and inventory control have been treated independently of each other (Tiacchi and Saetta 2009; Syntetos, Babai, Dallery, and Teunter 2009). However, demand forecasting performance, as assessed using standard statistical measures of accuracy may not necessarily translate into inventory systems efficiency (Syntetos, Nikolopoulos, and Boylan 2010).

Table 4: Error Statistics when Applying Various Methods to Simulated Demand Distributions

SKU #	1	2	3	4	5
Smoothing Constants Selected in Calibration Block					
SES	0.20	0.20	0.20	0.10	0.20
SBA	0.05	0.05	0.05	0.10	0.05
MAPE					
SMA13	97.67%	114.02%	165.42%	116.67%	163.50%
SES	99.39%	115.98%	163.89%	115.58%	163.96%
SBA	93.82%	111.92%	154.96%	113.61%	163.54%
Best MAPE	SBA	SBA	SBA	SBA	SMA13
MAD					
SMA13	10.118	1.681	1.203	10.551	4.576
SES	10.299	1.710	1.192	10.451	4.604
SBA	9.716	1.646	1.124	10.268	4.578
Best MAD	SBA	SBA	SBA	SBA	SMA13
RMSE					
SMA13	13.899	2.319	1.844	14.499	7.602
SES	14.136	2.363	1.825	14.293	7.643
SBA	13.544	2.259	1.800	14.087	7.264
Best RMSE	SBA	SBA	SBA	SBA	SBA
SKU #	6	7	8	9	10
Smoothing Constants Selected in Calibration Block					
SES	0.20	0.05	0.05	0.15	0.20
SBA	0.05	0.05	0.05	0.05	0.20
MAPE					
SMA13	104.36%	108.09%	156.34%	169.37%	93.82%
SES	106.08%	104.73%	156.68%	168.40%	94.92%
SBA	101.60%	102.67%	156.54%	163.98%	91.49%
Best MAPE	SBA	SBA	SMA13	SBA	SBA
MAD					
SMA13	3.728	5.106	14.209	4.8362	6.046
SES	3.791	4.944	14.199	4.8097	6.118
SBA	3.624	4.847	14.121	4.6655	5.897
Best MAD	SBA	SBA	SBA	SBA	SBA
RMSE					
SMA13	5.212	7.113	20.108	8.4369	8.262
SES	5.285	6.940	19.665	8.4355	8.383
SBA	5.046	6.935	19.529	8.1122	8.248
Best RMSE	SBA	SBA	SBA	SBA	SBA

A periodic review inventory control system has been recommended in dealing with intermittent demand (e.g., Sani and Kingsman 1997; Syntetos, Babai, Dallery, and Teunter 2009). Some recent intermittent demand forecasting studies (e.g., Eaves and Kingsman 2004; Syntetos and Boylan 2006; Syntetos, Babai, Dallery, and Teunter 2009; Syntetos, Nikolopoulos, Boylan, Fildes, and Goodwin, 2009; Syntetos, Nikolopoulos, and Boylan 2010; Teunter, Syntetos, and Babai 2010) that evaluate both forecasting and inventory control performance have used the order-up-to (T, S) periodic review system, where T and S denote the review period and the base stock (or ‘order-up-to’ level), respectively.

We assume in the current study a (T, S) inventory control system with full backordering. Inventory is reviewed on a monthly basis ($T = 1$). For most SKUs, the reorder lead time is more or less one month; we thus

set $L = 1$. Let I_t and B_t , respectively, denote the on-hand inventory and backlog at the time of review t . The literature on inventory control suggests a safety stock component in order-up-to levels to compensate for uncertainty in demand during the “protection interval” $T+L$. For each demand series, we calculated the standard deviation, s_{cal} , of monthly usage quantities during the calibration block. We apply a “safety factor” k as a multiplier of s_{cal} to obtain a safety stock level of $k \cdot s_{cal}$. This safety stock determination is more or less similar to the $z \cdot \sqrt{T+L} \cdot \sigma_d$ suggested when daily demand during the protection interval is assumed to be identically and independently normally distributed with standard deviation σ_d (e.g., Silver, Pyke, and Peterson 1998). With the safety stock component, the replenishment quantity to order is

$$Q_t = \sum_{j=t+1}^{t+T+L} F_j + k \cdot s_{cal} - I_t + B_t. \quad (8)$$

We simulate performance of the (T, S) inventory control system based on the two most commonly specified service level criteria (Silver, Pyke, and Peterson 1998). One service criterion is a target average *probability of no stockout* per review period. The other service measure is a target average fraction of demand to be satisfied from stock on hand, also called a *fill rate* (FR) which has considerably more appeal for practitioners.

With a 95% target FR, averages of inventory on hand arising in the simulation experiments from the use of SMA13, SES, and SBA are reported in Table 5. We find that, for all of the ten SKUs, the average inventory on hand is consistently lowest when SBA is the forecasting method applied.

Table 5: Average Inventory on Hand for a 95% Fill Rate

SKU #	1	2	3	4	5
SMA13	34.77	6.24	6.33	39.50	29.03
SES	35.47	6.40	6.16	38.73	28.14
SBA	33.61	6.07	6.15	37.56	27.54
SKU #	6	7	8	9	10
SMA13	13.51	18.32	59.00	37.85	20.02
SES	13.87	17.72	57.01	37.74	20.34
SBA	13.07	17.69	56.48	37.39	19.86

For the same 95% target FR, the means of total backlogs over 100 months (as averaged over 100 replications) are reported in Table 6. We find these 100-month means to be roughly equal across the forecasting methods.

Under a 95% target probability of no stockout, SBA likewise generally outperforms SMA13 and SES with respect to average inventory on hand. This is exhibited in Figure 1 where all indices (with SBA as

base) are above 100, with an index of 98.3 when using SES for SKU 8 being the only exception.

Table 6: Mean 100-Month Backlogs for a 95% Fill Rate

SKU #	1	2	3	4	5
SMA13	52.01	7.45	3.71	45.60	14.45
SES	52.10	7.58	3.80	45.78	14.36
SBA	52.12	7.58	3.85	45.79	14.43

SKU #	6	7	8	9	10
SMA13	17.82	24.09	45.43	15.22	31.84
SES	17.90	24.51	45.57	15.21	32.12
SBA	18.07	24.29	46.18	15.30	32.13




Figure 1: Average On-Hand Inventory Indices, with SBA as Base, for a 95% Probability of No Stockout

Likewise, with a 95% target probability of no stockout, the means of total backlogs over 100 months are reported in Table 7. None of the means arising from the use of SBA appears to be substantially higher than the mean corresponding to the use of either SMA13 or SES.

Table 7: Mean 100-Month Backlogs for a 95% Probability of No Stockout

SKU #	1	2	3	4	5
SMA13	75.49	14.72	12.21	80.99	47.66
SES	74.96	15.03	12.12	78.53	49.63
SBA	75.48	14.96	12.45	78.46	48.21

SKU #	6	7	8	9	10
SMA13	30.35	38.97	85.51	68.71	44.94
SES	30.40	39.25	84.69	68.76	45.62
SBA	30.76	39.51	79.06	70.06	45.49

5. CONCLUSION AND FURTHER WORK

We have devised a two-stage approach, involving uniform and negative binomial distributions, which allows modeling of the actual demand distribution, even when it is lumpy. Our work departs from earlier studies which have merely argued that the NBD satisfies both theoretical and empirical criteria, and accordingly assumed that an NBD adequately captures the behavior of intermittent demand. We believe that the simulated demand distributions arising from our two-stage modeling approach would more closely approximate the

actual demand distributions of the SKUs under consideration.

In empirically investigating the forecasting methods on the performance block (the final 22 months of the 66-month actual distribution) using three traditional statistical measures of forecast accuracy, we found none of the methods under consideration to be consistently superior to the others. However, when the methods are tested over considerably more time periods (100 replications of 100 months using our two-stage approach), SBA is found to be the best performing method overall in terms of statistical accuracy.

We then proceeded to apply the demand estimates arising from the different forecasting methods, on the basis of the simulated demand distribution generated for a given SKU. We assumed a (T,S) periodic review inventory control system with full backordering, with a one-month review period and a one-month replenishment leadtime. Using either a target FR or a target probability of no stockout as customer service level criterion, we have found SBA to yield the lowest average levels of inventory on hand in almost all cases. At the same time, the frequency of backorders under SBA is comparable to those using the other forecasting methods.

The observations reported here are based on a very limited sample of ten SKUs from the industrial dataset. At the time of the conference, we expect to report more robust findings – based upon our analysis of forecast accuracy and stock control performance over a larger number of SKUs as well as other customer service levels.

REFERENCES

- Altay, N., Rudisill, F., Litteral, L.A., 2008. Adapting Wright's modification of Holt's method to forecasting intermittent demand. *International Journal of Production Economics*, 111 (2), 389-408.
- Babai, M.Z., Syntetos, A.A., and Teunter, R., 2010. On the empirical performance of (T,s,S) heuristics. *European Journal of Operational Research*, 202 (2), 466-472.
- Boylan, J.E., Syntetos, A.A., and Karakostas, G.C., 2008. Classification for forecasting and stock control: a case study. *Journal of the Operational Research Society*, 59 (4), 473-481.
- Croston, J.D., 1972. Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, 23 (3), 289-304.
- Eaves, A.H.C. and Kingsman, B.G., 2004. Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society*, 55 (4), 431-437.
- Ferrari, E., Pareschi, A., Regattieri A., and Persona, A., 2006. Statistical management and modeling for demand of spare parts. In: H. Pham, ed. *Springer Handbook of Engineering Statistics*. Würzburg, Germany: Springer-Verlag, 905-929.

- Gilliland, M., 2002. Is forecasting a waste of time? *Supply Chain Management Review*, 6 (4), 16-23.
- Gutierrez, R.S., Solis, A.O., and Mukhopadhyay, S., 2008. Lumpy demand forecasting using neural networks. *International Journal of Production Economics*, 111 (2), 409-420.
- Johnston, F.R. and Boylan, J.E., 1996. Forecasting for items with intermittent demand. *Journal of the Operational Research Society*, 47 (1), 113-121.
- Mood, A.M., Graybill, F.A., and Boes, D.C., 1974. *Introduction to the Theory of Statistics*, 3rd ed. New York, USA: Mc-Graw-Hill.
- Mukhopadhyay, S., Solis, A.O., and Gutierrez, R.S., 2011. The Accuracy of Non-traditional versus Traditional Methods of Forecasting Lumpy Demand. *Journal of Forecasting* (forthcoming; DOI: 10.1002/for.1242).
- Sani, B. and Kingsman, B.G., 1997. Selecting the best periodic inventory control and demand forecasting methods for low demand items. *Journal of the Operational Research Society*, 48 (7), 700-713.
- Schultz, C.R., 1987. Forecasting and inventory control for sporadic demand under periodic review. *Journal of the Operational Research Society*, 38 (5), 453-458.
- Silver, E.A., Pyke, D.F., and Peterson, R., 1998. *Inventory Management and Production Planning and Scheduling*. New York, USA: John Wiley & Sons.
- Syntetos, A.A. and Boylan, J.E., 2001. On the bias of intermittent demand estimates. *International Journal of Production Economics*, 71 (1-3), 457-466.
- Syntetos, A.A. and Boylan, J.E., 2005. The accuracy of intermittent demand estimates. *International Journal of Forecasting*, 21 (2), 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*, 56 (5), 495-503.
- Syntetos, A.A. and Boylan, J.E., 2006. On the stock control performance of intermittent demand estimators. *International Journal of Production Economics*, 103 (1), 36-47.
- Syntetos, A.A., Nikolopoulos, K., Boylan, J.E., Fildes, R., and Goodwin, P., 2009. The effects of integrating management judgement into intermittent demand forecasts. *International Journal of Production Economics*, 118 (1), 72-81.
- Syntetos, A.A., Babai, M.Z., Dallery, Y., and Teunter, R., 2009. Periodic control of intermittent demand items: theory and empirical analysis. *Journal of the Operational Research Society*, 60 (5), 611-618.
- 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*, 26 (1), 134-143.
- Teunter, R.H., Syntetos, A.A., and Babai, M.Z., 2010. Determining order-up-to levels under periodic review for compound binomial (intermittent) demand. *European Journal of Operational Research*, 203 (3), 619-624.
- Tiacci, L. and Saetta, S., 2009. An approach to evaluate the impact of interaction between demand forecasting method and stock control policy on the inventory system performances. *International Journal of Production Economics*, 118 (1), 63-71.
- 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*, 10 (4), 529-538.

AUTHORS BIOGRAPHY

Adriano O. Solis is an Associate Professor of Logistics Management and Management Science at York University, Canada. After receiving BS, MS and MBA degrees from the University of the Philippines, he joined the Philippine operations of Philips Electronics where he became a Vice-President and Division Manager. He went on to obtain a PhD degree in Management Science from the University of Alabama. He was previously Associate Professor of Operations and Supply Chain Management at the University of Texas at El Paso. He has been a Visiting Professor in the Department of Mechanical, Energy, and Management Engineering, at the University of Calabria, Italy. He has published in *European Journal of Operational Research*, *International Journal of Simulation and Process Modelling*, *Information Systems Management*, *International Journal of Production Economics*, *Journal of Forecasting*, *Computers & Operations Research*, and *Journal of the Operational Research Society*, among others.

Letizia Nicoletti is currently PhD student at MSC-LES, Mechanical Department of University of Calabria. Her research interests include inventory management Modeling and Simulation. She is also applying Modeling & Simulation based approaches for training in complex systems, specifically marine ports and container terminals. She also actively supports the local organization of many international conferences in the modeling and simulation area.

Somnath Mukhopadhyay, an Associate Professor in the Information and Decision Sciences Department at the University of Texas at El Paso, has published in journals like *Decision Sciences*, *INFORMS Journal on Computing*, *Journal of Forecasting*, *Communications of the AIS*, *IEEE Transactions on Neural Networks*, *Neural Networks*, *Neural Computation*, *International Journal of Production Economics*, and *Journal of World Business*. He received MS and PhD degrees in Management Science from Arizona State University. He was a visiting research assistant in the parallel distributed processing

group of Stanford University. He has over 15 years of industry experience in building and implementing mathematical models.

Laura Agosteo, Antonio Delfino, and Mirko Sartiano are students in the *Laurea Magistrale* (master's degree) program in Management Engineering at the University of Calabria's Department of Mechanical, Energy, and Management Engineering. They obtained their first degrees in Management Engineering in 2010. They were recently exchange graduate students/visiting research scholars at the School of Administrative Studies, York University, Canada.