Forecasting and Inventory Planning for Parts with Intermittent Demand - A Case Study

Gokul Hariram Kottaram Venkitachalam, David B. Pratt, and Camille F. DeYong
School of Industrial Engineering and Management
Oklahoma State University
Stillwater, OK 74078, USA

Steven Morris and Michel Leonard Goldstein
School of Electrical and Computer Engineering
Oklahoma State University
Stillwater, Ok., 74078, USA

Abstract

Forecasting demand and developing inventory strategies for parts with an intermittent demand history presents a formidable challenge. We review the results of efforts to forecast part failures and determine inventory strategies for military aircraft parts. Applications of bootstrapping and Croston's method are summarized and the results contrasted with more traditional time series approaches.

Keywords
Intermittent Demand, Forecasting, Inventory Planning.

1. Intermittent Time Series
An intermittent time series is a time series of non-negative integers where some of the values are zero [1]. This paper explores forecasting and inventory planning of intermittent series characterized by a time series pattern that contains frequent and irregularly spaced zero values. Table 1 shows an example of a time series that exhibits this property. Time series of this type are typical in the early failure histories of aircraft repair parts.

Table 1. Intermittent Data Example

Time Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Demand	2	1	5	0	0	0	0	1	0	0	0	0	1	0	1

2. Traditional Techniques
When traditional forecasting techniques are applied to an intermittent series, the results are frequently unsatisfactory. As a typical example, consider the application of a 3-period moving average model with one period lead-time to the data in Table 1. The resulting forecasts and absolute forecast errors are shown in Table 2. The mean absolute error over the twelve time periods (period 4 to period 15) is 0.86. Contrast this with the naïve strategy of forecasting zero for each time period. This naïve strategy generates a mean absolute error over the twelve time periods of 0.25. This raises the question of whether there is any benefit in attempting to forecast at all.

Table 2. 3-Period Moving Average Forecast and Error

Time Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Demand	2	1	5	0	0	0	0	1	0	0	0	0	1	0	1
Forecast				2.7	2.0	1.7	0.0	0.0	0.3	0.3	0.3	0.0	0.0	0.3	0.3
Abs. Error				2.7	2.0	1.7	0.0	1.0	0.3	0.3	0.3	0.0	1.0	0.3	0.7

Table 3 shows the results of applying a simple exponential smoothing model to the example data. For this model, the parameter alpha is set to 0.2 and the initial forecast is set equal to the initial demand. The mean absolute error over the twelve time periods is 0.99, which is once again inferior to the naïve approach of forecasting zeroes.

It can legitimately be argued that the viability of traditional forecasting approaches is sensitive to the "intermittenness" of the series being forecast and the parameters selected for the model. However, the results shown above are typical and representative of those that were achieved in the case study presented in this paper.

Table 3. Simple Exponential Smoothing Forecast (alpha=0.10) and Error

Time Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Demand	2	1	5	0	0	0	0	1	0	0	0	0	1	0	1
Forecast	2.0	2.0	1.8	2.4	2.0	1.6	1.2	1.0	1.0	0.8	0.6	0.5	0.4	0.5	0.4
Abs. Error				2.4	2.0	1.6	1.2	0.0	1.0	0.8	0.6	0.5	0.6	0.5	0.6

3. The Case Under Study

The time series under consideration in this study are order and delivery data covering 163 periods for approximately 30,000 parts that are uniquely used on the KC-135 [2]. The ongoing project of which this study is a component aims to improve the forecasting of demand for these (and other) parts that have very low demand. Of particular interest are parts that experience demand "surges". Surge parts are defined as parts that have little or no demand over the history of the aircraft and then suddenly experience a large increase in demand [2]. Each part under consideration can be associated with a unique identification number known as an NIIN. Croston's method and a bootstrapping technique are used in conjunction with a clustering technique in an attempt to improve the forecast of demand for these intermittent parts.

3. Inventory Aspects

An important aspect of the problem being considered is to understand the inventory implications of the forecasts. If demand for a part is incorrectly forecast, then either an inventory excess or an inventory shortage will occur. If the part is critical to the safe and effective operation of an aircraft, then a shortage may well lead to a grounded aircraft. This aircraft could potentially be a fighter jet (not the example considered here) or a critical supplies transport, in which case, the strategic importance of the loss of capability incurred cannot easily be quantified in dollars. Thus, it is of paramount importance, where possible and economically viable, to have the fewest number of aircraft grounded. The implication of this aspect of the problem is that the impact of inventory shortages and inventory excesses may be very different. A shortage in inventory may be more harmful than an excess in inventory. The quantification of this difference is elusive and is handled in this study through the use of a sensitivity analysis. The sensitivity analysis explores a range of ratios (cost of shortage/cost of excess) to determine if this factor is significant in determining the preferred forecasting approach.

A second significant inventory aspect of this problem is the consideration of lead-time. The majority of the parts considered in this study have an average order lead-time of eight periods. That is, an order place in period 1 will arrive and be available for use in period 9. This long lead-time magnifies the impact of shortages and the impact of a forecasting model that is slow to respond to changes in the underlying demand pattern.

4. Measure of Error

Some forecast analysts believe that error measures should account for asymmetries in the cost of errors [3]. In other words, a positive forecast error is treated differently than a negative forecast error. Traditional measures of error, such as MSE (mean squared error), fail to effectively evaluate forecasting methods when the costs of errors are asymmetrical. Since the present case study presents a situation where asymmetry is important, this necessitates the development of a forecast evaluation method that considers the differing levels of impact due to shortages and excesses in inventory. This evaluation is handled through a sensitivity analysis. The first step in the analysis is to compute the error in each period generated from using a particular forecasting approach. On-hand inventories and order receipts for each period are accounted for in these calculations. In each period, either a shortage or an excess will occur. Tallies are kept of the cumulative number part shortages and part excesses. Summing the total number of shortages and the total number of excesses across the forecasted time periods determines the initial performance measure of the forecasting method. Table 4 illustrates this calculation. This calculation is repeated for each part and each forecasting methods under consideration.

Table 4.Calculation of Total Shortages and Excesses

Time Period	0	1	2	3	4	5	6	7	8	9	10	11
Inventory	0										18	
Fcst Error		-2	-1	-5	0	0	0	0	-1	0	0	0
Shortages		2	3	8	8	8	8	8	9	9		
Excesses											9	10

Shortages	63
Excesses	19

A sensitivity analysis is then conducted to determine the sensitivity of selecting a preferred forecasting method to the increasing weight (importance) placed on shortages. This is accomplished by multiplying the number of shortages by a weighting factor prior to calculating the final sum. This sensitivity analysis is illustrated in Table 5. Table 5 is then used to determine which technique dominates at each level of the sensitivity analysis. Dominance is established by the forecasting method that generates the lowest value of the weight sum of shortages and excesses. The highlighting in Table 5 illustrates identification of dominance. A forecasting method may dominate across all the assigned weights. Alternatively a forecasting method can dominate only beyond a threshold weight (as shown in Figure 5) or dominance can shift between the forecasting methods with no clear pattern.

Table 5. Sensitivity analysis and determination of the dominant forecasting technique

Part #1	Part Period		Increasing Weight on Shortages							
Forecasting Method	Shortages (S)	Excesses (E)	S+E	1.5S+E	2S+E	5S+E	7S+E	10S+E	20S+E	50S+E
Method 1	40	44	84	104	124	244	324	444	644	844
Method 2	44	106	150	172	194	326	414	546	766	986
Method 3	28	98	126	140	154	238	294	378	518	658

5. Croston's Method

Croston's method deals with the problem of forecasting demand levels when the demand patterns are not regular [4]. The method was proposed in 1972 and has since established itself as the standard approach to forecasting problems with irregular patterns. Croston observed that the use of the traditional exponential smoothing for intermittent demands is not suitable, since it tends to overestimate the demand levels. The alternative that Croston proposed involves breaking the intermittent demand time series into two constituent time series - one series for the non-zero demand values and other series for the time interval between the non-zero demand values. Traditional exponential smoothing is then used on each of the constituent parts separately. As an illustration, Table 6 shows the two Croston's constituent time series that would result from the data in Table 1.

Table 6. Croston's Constituent Time Series for the Data in Table 1

Non-zero element	2	1	5	1	1	1
Interval	1	1	5	5	2	

Stated briefly, Croston's algorithm works as follows:

If the actual demand at time t, $y_t = 0$, then $\hat{z}_t = \hat{z}_{t-1}$, $\hat{p}_t = \hat{p}_{t-1}$, where $\hat{z}_t$ is the forecast estimate made at time t and $\hat{p}_t$ is the time period estimate made in the same time t. This implies that forecasts are not updated when there is no demand at time t (y_t=0). Alternatively, if the actual demand at time t, $y_t \neq 0$, then $\hat{z}_t = \hat{z}_{t-1} + \alpha(y_t - \hat{z}_{t-1})$, $\hat{p}_t = \hat{p}_{t-1} + \alpha(q - \hat{p}_{t-1})$ where α is the smoothing constant and q is the difference between actual demand at time t and the previous non-zero demand. Croston assumes that the demands occur as a Bernoulli process and that the demand sizes are Normal IID (independent, identically distributed).

Given the inherent importance of the problem of intermittent demand, Croston's method has undergone considerable scrutiny, as evidenced by the continuing and sustained interest in his algorithm. Johnston and Boylan [5] studied the effect of uncertainty in the elapsed time between active periods. Willemain *et al*. [6] studied the performance of Croston's method compared to the exponential smoothing technique, when applied to (i) artificial data purposefully created to violate Croston's assumptions, and (ii) real world data from industrial sources. Their results demonstrate that Croston's method is robust and superior in forecast performance when compared to exponential smoothing.

6. Bootstrap Method

In a general inventory demand forecasting problem, the objective is to minimize the cost for the inventory system by making the best decision of what to purchase and when. One approach to making this decision is to determine the underlying probability distribution for part demand and then use knowledge of this distribution to provide a specified level of protection against shortages. Unfortunately, in a low-demand situation, the demand history of a part may not contain sufficient data to accurately estimate the demand distribution parameters. Moreover, standard distributions, such as Poisson and compound Bernoulli, may not effectively approximate the observed distribution.

To cope with these limitations, a clustering-bootstrap method was developed. Clustering is used to classify the NIIN's into groups that have similar patterns in their time series histories. Clustering provides the following benefits: (i) demand history for similar parts are pooled to provide better statistics for forecasting models and (ii) separate forecasting models may be constructed for each cluster rather than each part. This method uses the information from a group of parts to estimate the distribution without having to rely on distribution assumptions [7].

The grouping of the parts is based on the similarity in their order pattern history. The similarity chosen was the Euclidean distance between the cumulative history of orders up to the period of analysis. This choice increases the correlation of quarters that are close to one another and is also highly influenced by the total number of parts ordered. The grouping is made using a fuzzy clustering method. Fuzzy clustering assigns a membership value for each element to each cluster. This feature is also used for detecting outliers (parts that do not fall reasonably into any cluster).

Having grouped the similar parts, the application of bootstrapping is straightforward. Bootstrapping uses a sampling with replacement method to generate estimates of the distribution parameters based on the distribution of the sampled elements. Since the distribution that is being forecast is the total number of parts needed over the lead-time, the sampling is done on groups of periods of the same part to ensure that the time correlation is maintained.

The result of the bootstrap analysis is a probability distribution of the demand over the number of periods (lead-time) that are being analyzed. In a group there can be more than one lead-time, therefore more than one bootstrap analysis may be done within the same group to generate the distributions for the required number of periods. The forecast is generated from the probability distribution by determining a confidence level with respect to expected shortages, such that a specified proportion of the parts will have sufficient stock to fulfill the all the orders expected in a given period.

7. Results of Preliminary Work

A set of 72 NIIN's were chosen and forecasts were made for each of these NIIN's using Bootstrapping techniques with 25%, 50% and 75% confidence levels and Croston's method corresponding to α values of 0.1, 0.2 and 0.7. These forecasting methods were then evaluated using the sensitivity analysis method described in Section 4.0. A count of the number of times each of the forecasting methods dominated within each of the weight schemes is illustrated in Figure 1.

8. Conclusions and Continuing Work

Based on the evidence that resulted from the study of the performance of the forecasting methods on 72 NIIN's, it seems apparent that Bootstrapping technique based on 75% confidence level is preferred when the weight on the shortages exceeds 1.5. If shortages and excesses are equally weighted (weight = 1), no clearly preferred forecasting method emerges. Work on this project continues to explore the significance and ramifications of these results. Additional work is also being pursued to explore combining the Bootstrap method and Croston's method to study the performance of the combined forecast against the other forecasting methods.

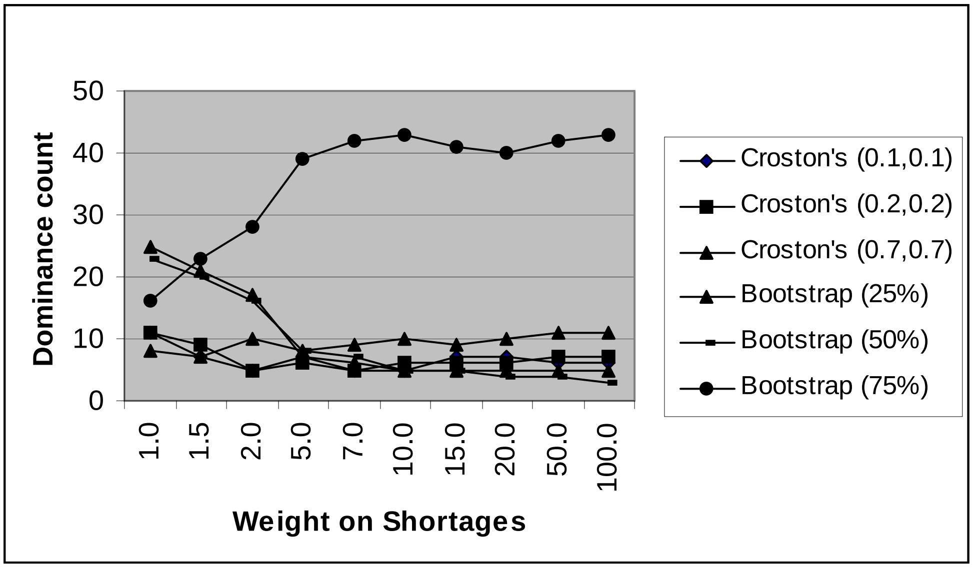

Figure 1. Comparative Evaluation of Croston's and Bootstrap forecasts

Acknowledgements

The generous support of the Department of the Air Force, Design and Engineering Support Program under Contract F42620-00-D-0036, Task Order 0005 is gratefully acknowledged. Also acknowledged for their contributions to this effort are the members of the project team who were not specifically listed as co-authors.

References

1. Armstrong, J.S. (2001c), "Extrapolation for Time-Series and Cross-Sectional Data," in J.S. Armstrong (ed.), *Principles of Forecasting.* Norwell, MA: Kluwer Academic Press.
2. Morris, S.A. (2002), "Phase II Final Report For Advanced Demand Forecasting", *Technical Report*, pp.2-14.
3. Armstrong, J.S. (2001d), "Evaluating Forecasting Methods," in J.S. Armstrong (ed.), *Principles of Forecasting.* Norwell, MA: Kluwer Academic Press.
4. Croston, J.D. (1972), "Forecasting and Stock Control for Intermittent Demands," *Operations Research Quarterly*, 23(3), pp. 289–203.
5. Johnston, F.R., and Boylan, J.E, (1996), "Forecasting for Items with Intermittent Demand," *Journal of the Operational Research Society*, 47, pp. 113–121.
6. 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, pp. 529–538.
7. A. C. Davison, A.C. and Hinkley, D.V., (1999), "Bootstrap Methods and Their Application", Reprinted with corrections. ed. Cambridge; New York: Cambridge University Press.