

On Replenishing Items with Seasonal Intermittent Demand

¹Gary Mitchell and ²Meike Niederhausen

¹Pamplin School of Business Administration,
University of Portland, Portland, OR 97203-5798

²Department of Mathematics, University of Portland, Portland, OR 97203-5798

Abstract: Problem statement: There are numerous difficulties associated with replenishing intermittent demand items and these are compounded when the demand distribution(s) vary seasonally. Excess inventories during an “off” season are typical, while shortages frequently occur during the “in” season, especially at the transition points between “seasons”. **Approach:** Evaluate the extent to which items characterized by non-stationary (seasonal) intermittent demand can be managed with commonly used forecasting and replenishment methods, including existing “intermittent demand” methods. Extensive simulation studies were conducted using combinations of two commonly used forecasting methods and two replenishment policies to evaluate the impact of non-stationary intermittent demand on key inventory performance measures, including average inventory and net profits, including the extent to which combinations of forecasting and replenishment models is adversely impacted by the non-stationary demand distributions. **Results:** No combination of forecasting and replenishment methods tested consistently outperformed the others and all methods demonstrated a propensity to replace demanded units an average of 10 weeks before the inventory was required to avert a shortage. **Conclusion:** The inventory performance of the policies tested was consistent with our expectations and offered evidence of the need for further development of forecasting and replenishment models addressing the special characteristics of items with non-stationary intermittent demand.

Key words: Forecasting, inventory, intermittent demand

INTRODUCTION

The difficulties associated with managing inventory for items with intermittent demand are widely recognized and this problem characterizes a large percentage of the inventory management decisions faced by many companies. Spare parts, accessories, slow selling items in large assortments and low volume big ticket items are all representative of classes of items likely to be typified by intermittent demand. An August 2000 Business Wire article (Business Editors, 2000) quotes a study estimating spare parts spending of \$700 billion (8% of US GDP) in the US alone. Approximately 90% of the items stocked at retail by a national specialty retailer known to one of the authors sell fewer than one unit per week. These statistics help to illustrate the scope of this problem and emphasize the need for an efficient and effective method for making intermittent demand item inventory decisions.

It is this particular specialty retailer who motivated this study. In this case, the retailer’s business is of a

highly seasonal nature such that both the frequency of demand events and the size of demand events may differ substantially between the retailer’s “in” season and “off” season. Although initially motivated by a single company, the nature of the demand described 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.

In this study, we use a large-scale simulation to evaluate the impact of non-stationary intermittent demand on inventory performance. We assume that the time between demand events is variable, such that demand frequently does not occur in consecutive periods and that both the inter-demand event times and the number of units demanded during a demand event can be characterized by probability distributions. We further assume, consistent with the motivation, that both the time between demand events and the size of demand events are non-stationary. Specifically, we assume that the company’s selling year can be

Corresponding Author: Gary Mitchell, Pamplin School of Business Administration, University of Portland, Portland, OR 97203-5798




Fig. 1: Demand plots for four items showing non-stationary intermittent demand

separated into two intervals, one “in” season and one “off” season. Within each of these intervals, the time between demand events and the unit size of demand events can be characterized by probability distributions

The demand shown in Fig. 1 illustrates this “in” and “off” season behavior. For these four items, it is easy to see that demand events are more common during fifteen to twenty weeks in the middle of the year. Moreover, demand events occurring during this time frame tend to be of greater magnitude than those that occur during the remainder of the year.

We assume that a rational inventory manager is concerned with increasing net profits, reducing holding costs, limiting inventory time supplies, increasing realized service levels and fill rates and improving other common profit-oriented metrics of inventory performance. This may not always be the case. For example, Teunter and Duncan (2009) consider the intermittent demand problem in the context of service parts for the UK Royal Air Force (RAF) where profitability clearly is not the primary objective. In such a context, a rational inventory manager may wish to minimize total relevant inventory costs incurred for an intermittent demand item while satisfying an acceptable fraction of the demand.

It is without doubt that the performance objectives of the inventory manager are frequently in conflict with one another (e.g., increasing service levels may decrease net profits due to corresponding increases in holding costs). Arguably, then, an inventory manager

may accept a small percentage decrease in net profits, for example one percent, if accompanied by a large reduction in average inventory, say twenty percent, which could be invested in other inventory with better performance characteristics.

We assume that lead times are deterministic and constant and are sufficiently short such that the probability of a non-zero demand event before a resupply operation can be completed is zero. Additionally, we assume that the unit holding cost rate is deterministic and that 100% of the gross margin for the demanded item is lost in the event of a shortage. This latter assumption may be easily relaxed by retaining a portion of the gross margin, in the event of a shortage, where the retention percentage $\delta < 1.00$ indicates the margin retained on a less profitable substitute item and $\delta < 1.00$ represents the margin retained on a more profitable substitute item. Order placement and setup costs are assumed to be negligible.

This study makes four contributions. First, we highlight a new class of intermittent demand replenishment problems, where the distributions of time between demand events and the size of demand events are not stationary. This is an important practical problem faced by many companies in supply chain management, manufacturing and service operations. Second, we evaluate the performance of commonly used methods and highlight their weaknesses with regard to the intermittent demand replenishment problem when both the time to a demand event and the

size of a demand event are non-stationary. To the best of our knowledge, this is the first paper to look at this important and relevant variation of this problem. Third, we show that a simple policy combining a forecasting approach suggested by Croston (1972) with a non-optimal (s,S) reordering policy can substantially reduce average inventory and improve net profits, as compared to the simple base stock policy in Croston's seminal paper. And finally, we introduce a new metric for evaluating a policy's ability to recognize when a replenishment order can be deferred in the face of non-stationary intermittent demand.

This study is related to previous research on demand forecasting and inventory replenishment. Croston (1972) develops a method for improving intermittent demand item forecasting and replenishment in which he measures the time between demand events separately from the quantity of an item demanded. In this study, Croston assumes that demand occurs in any period with probability p and that demand occurring in a period is independent of demand occurring in any other period, so that the number of periods between demand events follows a geometric distribution. He assumes that the size of a demand event (units) follows a normal distribution and is independent of the time between events. Using these assumptions, Croston develops a method for updating the average number of periods between demand events \bar{p}_n and the mean \bar{z}_n and mean absolute deviation m_n of demand per demand event, compared to the average per unit time seen in typical time series methods. He then suggests using a base stock policy where the order up to level is calculated as $\bar{z}_n + km_n$, with $k = 3$ corresponding to a target service level of approximately 99%. Croston further asserts that an unbiased estimator of the average demand per period can be calculated as $\hat{y}_n = \frac{\bar{z}_n}{\bar{p}_n}$.

Shenstone and Hyndman (2005) write that Croston's method is the most widely used approach for intermittent demand forecasting. However, there are a number of papers, several recent, revisiting this topic. Many of these focus on whether Croston's method yields a biased estimator of expected weekly demand (Syntetos and Boylan, 2001), while a number of papers propose bias adjusting corrections (Syntetos and Boylan, 2006; Syntetos *et al.* 2005; Eaves and Kingsman, 2004; Boylan and Syntetos, 2007; Gutierrez *et al.*, 2008; Teunter and Sani, 2009). Johnston and Boylan (1996) compare the performance of exponentially weighted moving average intermittent demand forecasts to those developed from a model generating demand forecasts from order size

mean and variance and average inter-order times. They find that their method yields superior results when average inter-demand times exceed 1.25 forecast review periods. Shenstone and Hyndman (2005) point out that Croston's method is an ad-hoc procedure with no properly formulated underlying stochastic model. Hua *et al.* (2007) use logistic regression to predict non-zero demand events and bootstrapping to estimate lead time demand. Syntetos *et al.* (2009) provide a complete review of the literature in this area.

Instead of demand forecast accuracy and bias, Schultz (1987; 1989) focuses on restocking policies that schedule replenishments a specified number of time periods after a demand event, rather than replenishing immediately after an event as is typical of a base stock policy. Schultz (1987) introduces the concept of scheduling replenishment orders k time periods after a replenishment under the assumption of stationary demand sizes and stationary time between demand events. If a stockout occurs before the replenishment order is placed, then an emergency order is placed to meet the shortage and the remainder of the scheduled order quantity is placed as scheduled. Under the added assumptions of unit demands, high unit item costs relative to order processing costs, constant lead times that are short relative to the average time between demand events and large holding costs relative to shortage costs, Schultz (1989) finds optimal replenishment delays for distributions of time between demand events with increasing failure rates.

Present research is also related to recent "bootstrapping" approaches to estimate lead time demand for items with intermittent demand (Willemain *et al.*, 2004; Porras and Dekker, 2008) as well as traditional approaches to forecasting non-stationary demand (Makridakis *et al.*, 1997) and traditional replenishment models (Hadley and Whitin, 1963; Silver *et al.*, 1998). Bookbinder and Lordahl (1989) use a bootstrapping procedure for determining reorder levels when the lead time demand distribution cannot be adequately represented by a standard distribution. Teunter and Duncan (2009) evaluate the performance of Croston's method and "bootstrapping" and find that both methods outperform single exponential smoothing and moving averages. They note that replenishment orders are triggered only in periods with demand events and propose an alternative method for estimating lead time demand.

The rest of the study is organized as follows. First, we explicitly define the general problem, including assumptions and develop our approach, using graphical examples to highlight key problem considerations.

Next, we simulate inventory levels and ordering activity, evaluating the resulting inventory performance on a variety of performance measures. We compare and contrast performance across forecasting methods and reordering policies, drawing insights into which methods are best suited to address particular intermittent demand characteristics. We conclude by summarizing our findings, discussing implications for practicing inventory managers, and identifying potential extensions of present research.

MATERIALS AND METHODS

Method formulation:

Problem definition: Croston (1972) establishes the basic problem setting which consists of a routine stock control system in which updating occurs at fixed unit time intervals, which are typically much shorter than the times between successive demands for the product. Thus there will be frequent occasions for which the demand will be zero, although the average demand can be greater than one since it is calculated over multiple periods from demands of varying size.

We assume the demand year for an item consists of two intervals defined as “in” season and “off” season. “In” season, demand is generated by a Bernoulli process with probability p_I that a demand event will occur in a period. During the “off” season, demand is generated by a Bernoulli process with probability p_O that a demand event will occur. The distribution of the size of a demand event is also non-stationary and assumed to follow a shifted Poisson distribution with average demand per event λ_I “in” season and λ_O “off” season. Demand generated according to this process will exhibit characteristics similar to the demand shown in Fig. 1.

Under these assumptions about demand, we consider the performance of four policies consisting of the possible combinations of the commonly used forecasting and reordering methods listed in Table 1.

Forecasting models: Single Exponential Smoothing (SES) is a simple method that can approximate the results of weighted moving averages with minimal data requirements (Brown, 1959). Let y_t and \hat{y}_t be the actual and expected (i.e., forecasted) demand for period t , with \hat{y}_t calculated at the end of period $t-1$, $e_t = y_t - \hat{y}_t$ and α , $0 \leq \alpha \leq 1$, the exponential smoothing factor. Then:

$$\hat{y}_{t+1} = \alpha y_t + (1 - \alpha) \hat{y}_t = \hat{y}_t + \alpha e_t \quad (1)$$

Table 1: Forecasting and reordering methods

Forecasting approaches	Reorder policies
Single Exponential Smoothing (SES)	Base stock (order up to S)
Croston’s Method (CM)	(s,S) ¹

¹: This type of reordering policy is frequently referred to in practice as an OP/OUTL policy. Although it is theoretically possible to simultaneously determine s and S to minimize (maximize) a cost (profit) objective, it is more common in commercial practice to select the value of s to satisfy some form of target service objective, while S is generally selected to reflect constraints in ordering frequency and holding capacity

The Mean Absolute Deviation (MAD) (Given a MAD, m , for a normally distributed RV, the standard deviation, σ , can be found using the relation $\sigma = m\sqrt{\frac{\pi}{2}}$), a common measure of demand variability, is often calculated using SES:

$$m_{t+1} = \alpha |e_t| + (1 - \alpha) m_t \quad (2)$$

The difficulties associated with using exponential smoothing with intermittent data are well-known, including overstatement of demand variability. Croston’s method attempts to address these by separately estimating the average size of demand and time between demand events, assuming size and time between events are independent. In periods when $y_t > 0$:

$$\begin{aligned} \hat{z}_{t+1} &= \alpha y_t + (1 - \alpha) \hat{z}_t \\ m_{t+1} &= \alpha |y_t - \hat{z}_t| + (1 - \alpha) m_t \\ \hat{p}_{t+1} &= \alpha q + (1 - \alpha) \hat{p}_t \\ q &= 1 \end{aligned} \quad (3)$$

and when $y_t = 0$:

$$\begin{aligned} q &= q + 1 \\ \hat{z}_{t+1} &= \hat{z}_t \\ m_{t+1} &= m_t \\ \hat{p}_{t+1} &= \hat{p}_t \end{aligned} \quad (4)$$

Finally, the estimate of single period demand is given by $\hat{y}_t = \frac{\hat{z}_t}{\hat{p}_t}$.

Reorder policies: A base stock policy generally is used to calculate a replenishment order quantity, Q , each period when the desired system behavior is to return the inventory position, $IP = OH + OO - BO - CS$, to a base stock level S . OH is the current quantity on hand

(typically in units), OO is the current quantity on order (in systems where orders may be placed for a delivery more than a lead time in advance, for special purposes, then OO should represent the units currently on order for delivery within a lead time), BO represents the number of units currently backordered to customers and CS represents committed stock units (i.e., units on hand but reserved for special purpose and therefore not available for satisfying customer demand). In the remainder of this study we assume for simplicity and without loss of generality, $BO = CS = 0$ so that:

$$Q_t = S_t - IP_t = S_t - OH_t - OO_t \quad (5)$$

$$S_t = (1 + c)\hat{y}_t + ss_t \quad (6)$$

Where:

l = The lead time

c = The cycle time (the desired sell-through time of the on hand inventory)

ss_t = The safety stock necessary to buffer against uncertainty during the lead time (8)

There are various methods for determining ss_t , depending on the demand distribution and the nature of shortage costs (Silver *et al.*, 1998). Croston's paper uses a simplification of (6) consistent with an assumption of short lead times relative to the time between demand events. Using Croston's recommendation of $k = 3$, a base stock level calculated using (7) would be expected to satisfy 99.17% of demands completely from stock on hand. To provide a consistent basis for comparison, we use $S_t = \hat{y}_t + km_t$ for a base stock policy used in conjunction with SES, also setting $k = 3$:

$$S_t = \hat{z}_t + km_t \quad (7)$$

To calculate the replenishment order quantities when using an (s,S) policy, we use:

$$Q_t = \begin{cases} S_t - IP_t & IP_t \leq s_t \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

$$s_t = (r + 1)\hat{y}_t + ss_t$$

where, r represents the review interval (i.e., the time between order placement opportunities). We calculate safety stock to satisfy a target probability that demand will be met from stock:

$$ss_t = F_t^{-1}(P) - 1 \cdot \hat{y}_t \quad (9)$$

where, $F_t^{-1}(P)$ is the inverse cumulative demand distribution function for the item given information at

time t . Equation 9 is useful as it is easily applied using any demand distribution.

The introduction of the order point, s_t , provides a mechanism inhibiting the automatic replacement of inventory when the remaining on hand inventory may be sufficient for satisfying demand. This would seem particularly true in the case of intermittent demand, where there may be significant positive probability that a demand event will not occur during one or more periods following a demand event.

The non-stationary nature of our demand makes it even less reasonable to assume that the occurrence of a demand event should automatically trigger a replenishment order. Consider the demand shown in Fig. 2. It seems reasonable to assume that the "in" season starts at approximately week 14 and concludes at about week 35. Assume a base stock level of 8 units and consider the sequence of events that would occur beginning in week 4. Using a base stock policy, the demand events in periods 31 and 35 would generate replenishment orders for four and two units respectively and the "off" season would begin with 8 units of inventory, which would be carried with very little probability of a sale for roughly thirty weeks. Considered another way, the remaining inventory of four units after the demand is satisfied in period 31, would be sufficient to handle the demand for the next thirty-three weeks (nearly eight months), without a replenishment order. This would seem to suggest that the (s,S) policy should perform better than the base stock in terms of idle inventory.

The actual inventory performance is a function of the forecasting method and inventory policy. Figure 3 shows the actual performance given the demand in Fig. 2 when using Croston's forecasting and a base stock reorder policy. This evidences a definite bias toward providing a very high level of service performance, which illustrates why new methods are required under non-stationary intermittent demand. The circles show the replenishment order quantity, clearly calculated to return IP to S. The ending inventory position (denoted by the inverted triangles) in every period is significantly larger than the demand observed in the period. The elevated IP levels in the initial weeks are clearly the result of one or more very large demand events in the item's history prior to the plotted weeks. We can observe the result of the Croston's forecast update, lowering the ending inventory position each time a demand event occurs prior to the beginning of the "in" season. Of greatest note, however, is that the ending inventory at the entry into the "off" season is adequate to meet demand during the next thirty-seven periods. Similarly, the ending inventory in week 32 is also sufficient to meet demand in the next thirty-seven demand periods. The reordering decisions made in




Fig. 2: Sample demand showing “off” season




Fig. 3: Sample demand with order quantities and order up to levels

weeks 32 and 35 resulted in carrying 225% more inventory through the “off” season.

Figure 4 shows the order quantities resulting from the use of an (s,S) order policy in conjunction with Croston’s forecasting method. A simple count shows that the (s,S) orders with less than half the frequency as the base stock policy and, as a result, the ending inventory and consequently the ending inventory position, are allowed to assume lower levels, resulting in holding costs savings.

Figure 4 suggests the need for non-stationary demand intermittent forecasting methods and update policies even more clearly than Fig. 3 does. Although the (s,S) policy places orders with only 39% of the frequency of the base stock policy, it still places a

large order in week 32, resulting in high on hand “off” season inventories, relative to demand.

Simulation studies: We wished to evaluate the ability of the methods listed in Table 1 to manage inventory:

- How well do the methods perform in general?
- Does any single method consistently perform better than the others?
- How significantly is a method’s performance affected by the degree of non-stationarity in the time between demand events and the size of demand events given a fixed set of policy parameters?

Simulation overview: We conducted extensive numerical simulations in order to analyze the inventory performance for the four combinations of the two




Fig. 4: Sample demand, order quantities and order up to levels using (s,S) policy

forecasting methods and the two order policies (collectively referred to as replenishment policies) and also for varying inter-demand event time, demand size and gross margin ratio characteristics. Inventory performance was simulated 100 times for each replenishment policy over a ten year period (Each combination of forecasting method and reordering policy was evaluated for 100, 10-year periods against 54 different demand datasets. Effectively, each method was evaluated over a total of 10,000 years on 54 different classes of non-stationary intermittent demand). We calculated inventory performance on a large number of metrics:

- Average on hand inventory (period ending)
- Total units shorted
- Total gross profit
- Total holding costs
- Total net profit
- Total opportunity costs
- Average order periods delay
- Maximum order periods delay
- Average inventory position time supply
- Maximum inventory position time supply
- Percentage of order delay periods exceeding the maximum desired time supply (max DTS)
- Percentage of inventory position time supplies exceeding max DTS
- Average service level (all periods)
- Average service level in periods with demand events (i.e., |demand > 0)
- Average fill rate (all periods)
- Average fill rate |demand > 0

Most of these are self-explanatory, while a few probably merit a description. Average (and maximum) Order Periods Delay (OPD) is calculated any time an order is placed and indicates the number of weeks an order could have been delayed without incurring a shortage (i.e., while satisfying demand from stock). OPD is an accurate calculation, based on demand realized after the order placement rather than on estimated demand at the time the order was placed.

Average (and maximum) Inventory Position Time Supply (IPTS) is calculated using the inventory position at the end of a period (i.e., including any order quantities placed during the period) and accurately indicates the number of periods' demand that could be satisfied without placing an additional order. As with OPD, IPTS is calculated based on demand realized after the order placement.

The average service levels and fill rates over all periods are predictably high (due to the many instances of zero demand periods in which both service level and fill rate would be 100%). Of significantly more interest, especially in conjunction with average inventory levels, are service levels and fill rates given a demand event. These provide a much more informative assessment of service level and fill rate performance, calculated only from periods in which demand > 0.

Averages and confidence intervals were calculated on each of the metrics for each dataset (described previously). In particular we were interested in observing the differences in average inventory, average service levels and fill rates and the order delay periods across the replenishment policies. The following parameters were kept constant in all of simulations: the initial on hand amount was established as the average demand from the first five demand events, initial order amount was 0 units, unit holding cost per period was 0.48% of the item's value and lead time was 1 week.

Table 2: Inter-demand event time and demand size parameters for generated data 1-9

Inter-demand event time			
Demand size			
Distribution	$p_i = 0.25$	$p_i = 0.5$	$p_i = 0.5$
Parameters	$p_o = 0.02$	$p_o = 0.02$	$p_o = 0.25$
$\lambda_i = 1.5$	Data 1	Data 2	Data 3
$\lambda_o = 1.02$			
$\lambda_i = 5.0$	Data 4	Data 5	Data 6
$\lambda_o = 1.02$			
$\lambda_i = 8.0$	Data 7	Data 8	Data 9
$\lambda_o = 1.02$			

The exponential smoothing factor used for both Croston (1972) method and SES was $\alpha = 0.1$. The safety factor used in the base stock policy was $k = 3$. A Target Service Level (TSL) of 0.999 was used for the (s,S) policy calculations, roughly approximating the same level of target service as $k = 3$ for the base stock policy. Desired time supply for the (s,S) policy was held constant at 45 days when using Croston’s forecasting method and 14 days when using simple exponential smoothing (the only real accommodation made for the different replenishment methods). All other parameters were consistent across the methods (i.e., no effort was made to optimally select parameters for each method).

Generating random demand data: We used the probability distributions described in Method Formulation to randomly generate the demand data; the inter-demand event times followed the non-stationary geometric distribution and the demand sizes were assumed to follow a shifted Poisson distribution. We chose nine different quadruples $(p_i, p_o, \lambda_i, \lambda_o)$ to vary the expected inter-demand times and demand sizes both in and out of season and created a dataset for each. These values are summarized in Table 2.

In addition, for each of these nine data types we simulated data for six different gross margin ratios varying from 0.25-1.5, in increments of 0.25. The gross margin ratios were calculated as (sales cost-unit cost)/(unit cost). Thus in total we generated fifty-four data sets with different characteristics. Inventory performance for each of these data sets was simulated over 520 periods (weeks) and replicated 100 times (independently).

RESULTS

Although this study is more concerned with how well these replenishment policies handle non-stationary intermittent demand, the results of the simulation appear to confirm findings of prior researchers (e.g., Teunter and Duncan, 2009) that no method outperforms the others consistently for all variations of intermittent demand characteristics.




Fig. 5: OH Averages from 100 10-year simulations

Comparing the average on hand, total units shorted and total net profit: Figure 5 shows the average on hand inventory values calculated over the 10-year simulation horizon from 100 independent replications (Bar height represents average inventory levels. The “whiskers” bracketing the top of each bar indicate 95% confidence intervals, which are very narrow due to the large scale of the simulation). On hand inventories for Croston’s/base stock are the lowest across all methods only for Data 1-3, where the size of the average demand event is small and the difference between the “in” season and “off” season demand size distributions is small. In the remaining six cases, Croston’s/base stock inventory levels exceed those of at least one other replenishment method, while combining Croston’s forecasting method with an (s,S) reorder policy yielded reduced inventory levels in 5 of those 6 cases and the lowest levels overall in 3 of those cases.

On hand inventory levels alone yield an incomplete picture. With the exception of the Data 1-3 (low volume, small demand rate differences “in” and “off” season), the Croston’s/base stock replenishment policy consistently provides the best performance (with shortages considerably below those of other methods). However, this result is interesting when considering the average inventories in Fig. 5 in conjunction with Total Units Shorted in Fig. 6. For example, the average inventory for the SES/base stock policy is considerably higher than the Croston’s/base stock policy for Data 8, yet the SES/base stock policy clearly incurs higher shortage costs than Croston’s method.




Fig. 6: Average total shortages from 100 10-year simulations

The average total net profit results shown in Fig. 7 are important. Although the shortage costs associated with the Croston's/base stock policy are very low, the cost of achieving the higher fill rates is an increase in holding costs. When calculating the net profit (net of holding costs), we see that the net profit performance across the four methods is fairly consistent.

Figure 7 indicates that while the Croston's/base stock policy may be a good choice for management of critical supplies where shortages cannot be easily born by the system, any of the policies is capable of providing comparable performance in terms of net profits. By choosing a different policy, a firm could earn roughly the same profits on an item while gaining inventory dollars to invest in another item, thereby increasing its total earnings from inventory operations.

Comparing the order weeks delay and the IP time supply: Figure 8 shows the average number of weeks an order could have been delayed without incurring shortages. The values are high for all methods, with the lowest being for the Croston's/(s,S) policy in all cases except Data 1 and 3. This graph shows the significance of the (s,S) policy's ability to postpone reordering until a subsequent demand event occurs. Figure 8 shows that the (s,S) policy outperforms the base stock policy in six of nine cases. Moreover, a more detailed evaluation shows that the (s,S) policy consistently outperforms the base stock policy when the degree of demand size non-stationarity is greater (Data 4-9).

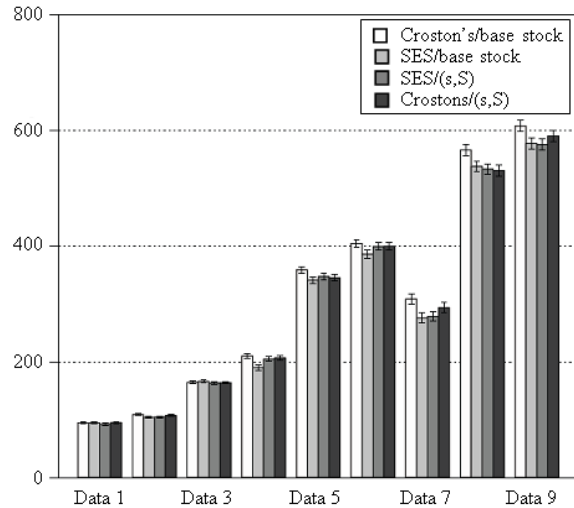


Fig. 7: Average total net profit from 100 10-year simulations




Fig. 8: Average order weeks delay from 100 10-year simulations

However, Fig. 9 also gives clear indication that none of the methods' performances is beyond criticism. With the majority of cases showing maximum order delays of 30 or greater for all policies, all of the policies are seen to raise inventory levels well above that required to meet demand.

Figure 10 and 11 show average and average maximum inventory position time supplies respectively for all methods. They all fare poorly, with minimum average time supplies of two months and average maximum time supplies of nine months or more. These large inventories could be mitigated, to some extent, by reducing the safety stock for all of the methods.




Fig. 9: Average maximum order weeks delay from 100 10-year simulations




Fig. 12: Average % order delay periods exceeding maximum desired time supply




Fig. 10: Average IP time supplies from 100 10-year simulations




Fig. 11: Average maximum IP time supplies from 100 10-year simulations

This would certainly reduce the time supplies, but would almost certainly increase shortages as well. Because the existing methods do not address the non-stationary characteristics of the intermittent demand, they are incapable of maintaining the service performance while simultaneously reducing inventory levels (and thereby improving the rate at which an inventory investment generates net profits).

DISCUSSION

Most organizations that manage inventory employ rules of thumb for quickly analyzing inventory levels. Maximum time supplies, while problematic for maximizing inventory performance, nonetheless have value in high-level analyses. Figure 12 shows the results of applying a maximum time supply to order delay periods (i.e., the number of periods an order could be delayed without incurring shortages). Specifically, the graph shows the frequency with which orders were placed even though the existing inventory exceeded a maximum time supply benchmark. All methods display a propensity to order when the existing time supply would exceed the threshold. This does not imply that a simple order cutoff of this nature should be implemented. On the contrary, doing so would very likely jeopardize service performance. Rather, high percentages give evidence to the need for replenishment methods that explicitly account for the non-stationary nature of the intermittent demand.

Figure 13 shows further evidence that non-stationary intermittent demand specific methods could yield superior performance. The central tendency for the base stock policy used with both forecasting methods clearly indicates a greater incidence of unnecessary ordering behavior. Moreover, the dispersion of the base stock policy methods is larger, especially for the Croston's/base stock replenishment policy.




Fig. 13: Order weeks delay frequency distributions from 100 10-year simulations. (a) Croston's/base stock; (b) SES/base stock; (c) Croston's/(s,S); (d) SES/(s,S)




Fig. 14: Average service levels when demand > 0

Finally, Fig. 14 and Fig. 15 show the methods' performances in terms of realized service levels and fill rates, respectively, measured only in periods when demand events occur (both measures would be 100% in other periods by definition). Service levels measure the percentage of periods when demand is completely satisfied from stock. In other words, demand is satisfied only when all units are filled directly from stock. Fill rates, in contrast, measure the fraction of unit demand satisfied from stock. From an aggregate perspective, the performance of all methods is quite good, with fill rates nearing 100%. Greater variation is observed in the datasets with higher degrees of non-stationarity of inter-demand event time and demand event size.

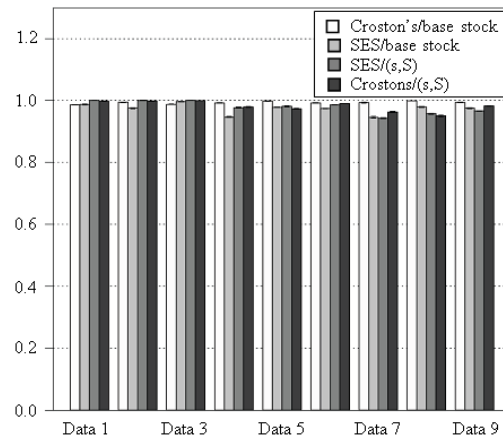


Fig. 15: Average fill rates when demand > 0

Owing largely to the higher inventory levels, the Croston's/base stock policy yielded the highest service levels on all but the three cases with the least non-stationarity in demand size.

CONCLUSION

In this study, we introduce an important problem faced by many inventory managers; specifically, we consider the problem of replenishing items with intermittent, non-stationary demand. While this problem is related to previous study on intermittent demand replenishment, our study shows that prevailing methods cannot adequately manage inventory when intermittent demand exhibits non-stationary characteristics. In general, we show that a manager who wants to attain good net profitability with less inventory investment should adopt a policy incorporating explicit order triggers (rather than the implicit trigger of the base stock policy). Thus, in choosing from the existing methods, managers in industries requiring high service performance at nearly any cost (i.e., military parts and supplies and hospitals) could be well-served by a Croston's/base stock (Croston, 1972), or similar, policy while those in profit taking operations would do well to consider the adoption of a Croston's/(s,S) policy.

Our results seem to corroborate those of previous researchers that no one method consistently outperforms the other methods. When demand levels are very low, a traditional Croston's method application (i.e., Croston's forecasting/base stock policy) yields the lowest average inventory levels. However, for higher average demand sizes, this method yields average inventories that are among the highest of the methods. In some cases we observed that the shortages yielded by the methods were not consistent with the average

inventories suggesting, not surprisingly, that improved fill rates can be attained by holding inventories judiciously. We demonstrated that service level and fill rate performance was very good for all methods despite the sometimes substantial differences in average inventories. We note this high service level and fill rate performance is the result of high inventory levels and that unnecessarily high inventory levels are frequently maintained during the “off” season, reducing net profit and offsetting the gross margin earned on sales.

This study represents the first of a number of research projects related to the problem of replenishing non-stationary intermittent demand items. We are currently working on a project to develop a non-stationary intermittent demand replenishment policy, as well as a project to characterize demand size and time to demand event distributions. Due to the scarcity of the intermittent demand data, commonly used goodness of fit tests frequently do not yield acceptable results. In another extension, we are working to develop an order policy that more directly targets the rate at which inventory investments generate profits. A fourth related research project involves developing an efficient and effective method for updating the rates at which demand events occur “in” and “out” of season.

The opportunity exists to conduct a comprehensive study evaluating and comparing the performance of the four replenishment policies studied here under base case performance. This would involve selecting policy control parameters (e.g., α, k) that yield the best performance on some inventory metric. A difficulty would be in selecting a metric as the outcomes will vary considerably by choice of metric. Finally, this study has focused primarily on the problem of replenishing non-stationary intermittent demand items with short lead times. Extending this study to items with long lead times would yield an important practical extension for organizations managing international supply chains and extensive outsourcing.

REFERENCES

- Bookbinder, J.H. and A.E. Lordahl, 1989. Estimation of inventory re-order levels using the bootstrap statistical procedure. *IIE Trans.*, 21: 302-312. DOI: 10.1080/07408178908966236.
- Boylan, J.E. and A.A. Syntetos, 2007. The accuracy of a modified Croston procedure. *Int. J. Prod. Econ.*, 107: 511-517. DOI: 10.1016/j.ijpe.2006.10.005.
- Brown, R.G., 1959. *Statistical Forecasting for Inventory Control*. 1st Edn., McGraw-Hill Book Company, New York, ISBN: 13: 978-0070081451, pp: 240.
- Business Editors, 2000. Spare Parts, Equipment and Other Fixed-Asset Purchasing Requires Special Functionality Not Provided in Today’s Indirect Purchasing Platforms. *Business Wire*. <http://www.thefreelibrary.com/Spare+Parts,+Equipment,+and+Other+Fixed-Asset+Purchasing+Requires...-a064817452>.
- Croston, J.D., 1972. Forecasting and stock control for intermittent demand items. *Operat. Res. Q.*, 23: 289-303. DOI: 10.1057/jors.1972.50
- Eaves, A.H.C. and B.G. Kingsman, 2004. Forecasting for the ordering and stock-holding of spare parts. *J. Operat. Res. Soc.*, 55: 431-437. DOI: 10.1057/palgrave.jors.2601697
- Gutierrez, R., A. Solis and S. Mukhopadhyay, 2008. Lumpy demand forecasting using neural networks. *Int. J. Prod. Econ.*, 111: 409-420. DOI: 10.1016/j.ijpe.2007.01.007
- Hadley, G. and T.M. Whitin. 1963. *Analysis of Inventory Systems*. 1st Edn., Prentice-Hall, Englewood Cliffs, New Jersey, ISBN: 13: 978-0130329530, pp: 528.
- Hua, Z.S., B. Zhang, J. Yang and D.S. Tan. 2007. A new approach of forecasting demand for spare parts inventories in the process industries. *J. Operat. Res. Soc.*, 58: 52-61. DOI: 10.1057/palgrave.jors.2602119
- Johnston, F.R. and J.E. Boylan, 1996. Forecasting for items with intermittent demand. *J. Operat. Res. Soc.*, 47: 113-121. DOI: 10.2307/2584256
- Makridakis, S., S. Wheelwright, and R. Hyndman, 1997. *Forecasting: Methods and Applications*. 3rd Edn., John Wiley and Sons, New York, ISBN:13: 978-0471532330, pp: 656.
- Porras, E. and R. Dekker, 2008. An inventory control system for spare parts at a refinery: An empirical comparison of different re-order point methods. *Eur. J. Operat. Res.*, 184: 101-132. DOI: 10.1016/j.ejor.2006.11.008
- Schultz, C., 1987. Forecasting and inventory control for sporadic demand under periodic review. *J. Operat. Res. Soc.*, 38: 453-458. DOI: 10.2307/2582735
- Schultz, C., 1989. Replenishment delays for expensive slow-moving items. *Manage. Sci.*, 35: 1454-1462. DOI: 10.1287/mnsc.35.12.1454
- Shenstone, L. and R. Hyndman, 2005. Stochastic models underlying Croston’s method for intermittent demand forecasting. *J. Forecast.*, 24: 389-402. DOI: 10.1002/for.963
- Silver, E., D. Pyke and R. Peterson, 1998. *Inventory Management and Production Planning and Scheduling*. 3rd Edn., John Wiley and Sons, New York, USA., ISBN: 13: 978-0471119470, pp: 784.

- Syntetos, A. and J. Boylan, 2001. On the bias of intermittent demand estimates. *Int. J. Prod. Econ.*, 71: 457-466. DOI: 10.1016/S0925-5273(00)00143-2
- Syntetos, A. and J. Boylan, 2006. On the stock control performance of intermittent demand estimators. *Int. J. Prod. Econ.*, 103: 36-47. DOI: 10.1016/j.ijpe.2005.04.004
- Syntetos, A., J. Boylan and J.D. Croston, 2005. On the categorization of demand patterns. *J. Operat. Res. Soc.*, 56: 495-503. DOI: 10.1057/palgrave.jors.2601841
- Syntetos, A.A., J.E. Boylan and S.M. Disney, 2009. Forecasting for inventory planning: A 50-year review. *J. Operat. Res. Soc.*, 60: S149-S160. DOI: 10.1057/jors.2008.173
- Teunter, R. and B. Sani, 2009. On the bias of Croston's forecasting method. *Eur. J. Operat. Res.*, 194: 177-183. DOI: 10.1016/j.ejor.2007.12.001
- Teunter, R.H. and L. Duncan, 2009. On the bias of Croston's forecasting method. *J. Operat. Res. Soc.*, 60: 321-329. DOI: 10.1057/palgrave.jors.2602569
- Willemain, T., C. Smart and H. Schwarz, 2004. A new approach to forecasting intermittent demand for service parts inventories. *Int. J. Forecast.*, 20: 375-387. DOI: 10.1016/S0169-2070(03)00013-X