2012

Effects of correlation on intermittent demand forecasting and stock control

Nezih Altay, *DePaul University*
Lewis A Litteral, *University of Richmond*
Frank Rudisill, *University of South Carolina - Upstate*

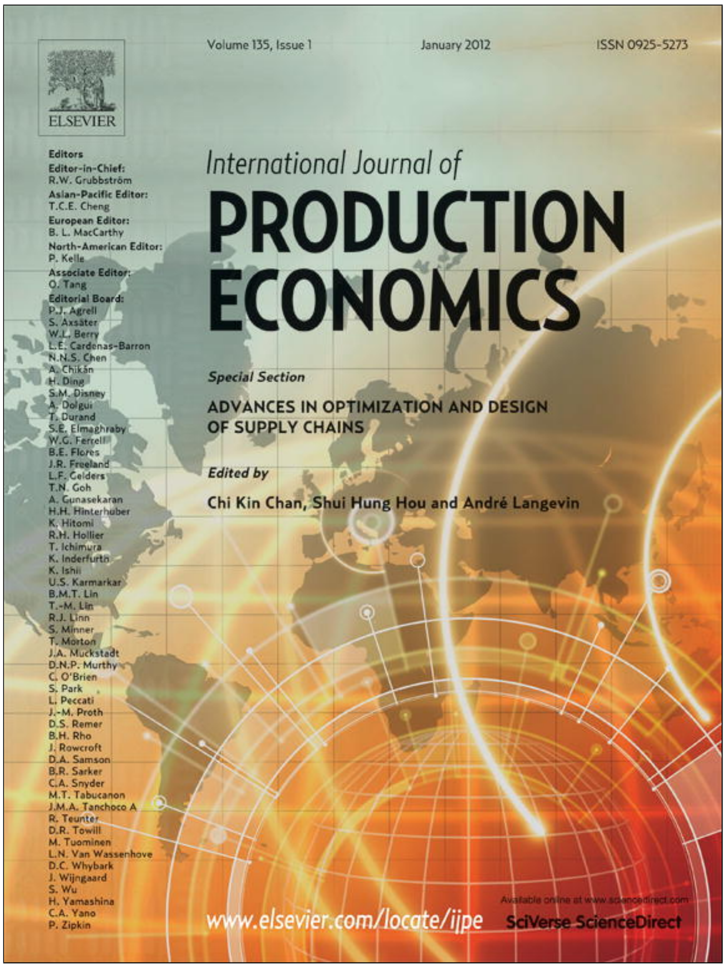
Volume 135, Issue 1
January 2012
ISSN 0925-5273
ELSEVIER
International Journal of
PRODUCTION ECONOMICS
Special Section
ADVANCES IN OPTIMIZATION AND DESIGN OF SUPPLY CHAINS
Edited by
Chi Kin Chan, Shui Hung Hou and André Langevin
Editors
Editor-in-Chief:
R.W. Grubbström
Asian-Pacific Editor:
T.C.E. Cheng
European Editor:
B. L. MacCarthy
North-American Editor:
P. Kelle
Associate Editor:
Q. Tang
Editorial Board:
www.elsevier.com/locate/ijpe
SciVerse ScienceDirect

Int. J. Production Economics 135 (2012) 275–283

Contents lists available at SciVerse ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

Effects of correlation on intermittent demand forecasting and stock control

Nezih Altay [a,*], Lewis A. Litteral [b,1], Frank Rudisill [c,2]

[a] DePaul University, Department of Management, 1 E. Jackson Blvd., Chicago, IL 60604, USA
[b] University of Richmond, Robins School of Business, Richmond, VA 23173, USA
[c] University of South Carolina Upstate, Johnson College of Business and Economics, 800 University Way, Spartanburg, SC 29303, USA

ARTICLE INFO

Article history:
Received 16 December 2010
Accepted 1 August 2011
Available online 19 August 2011

Keywords:
Forecasting
Time series
Intermittent demand
Correlation

ABSTRACT

This study investigates the effects of three different types of correlation on forecasting and stock control of intermittent demand items. Applying appropriate forecasting and stock control methods to theoretically generated compound Poisson demand data we show that correlation in intermittent demand does play a role in forecast quality and stock control performance. Negative autocorrelation levels lead to higher service levels than positive values, while cost does not significantly change. Our results also show that high intermittency levels intensify these changes in service level. We also show that cross-correlation produces results in the opposite direction of autocorrelation in size or intervals; that is, positive (negative) cross-correlation leads to higher (lower) service levels.

1. Introduction

Intermittent demand is random demand with a large proportion of zero values (Silver, 1981). Spare parts often exhibit intermittent demand and are particularly prevalent in the aerospace, automotive, military, and IT sectors (Johnston et al., 2003). In the military for example, spare parts make up such a big percentage of inventory the US Defense Logistics Agency (DLA) identified 24 strategic initiatives to better mitigate shortages of critical parts inventory (GAO, 2003). Intermittent time series present a special challenge because they display variability in demand size as well as demand arrival. Hence, forecasting and stock control of intermittent demand attracted a considerable amount of academic research (e.g. Altay et al., 2008; Syntetos et al., 2009a, 2009b, 2010).

Correlation in time series data also presents challenges. Most of the forecasting and inventory control theory is built on the assumption that successive demand values are independent. But auto-correlated demand sizes or inter-arrival times violate this assumption, requiring specially derived formulae. There is another type of correlation that is specific to intermittent demand items; correlation between demand size and the inter-demand interval of an item. In this paper, following the nomenclature of Willemain et al. (1994), we will refer to this type of correlation as *cross-correlation*. Cross-item correlation where one item's demand is correlated with the demand of another item has also been called cross-correlation in literature (Zhang, 1999; Liu and Yuan, 2000; Akkerman and Van Donk, 2009). In this research, we use *cross-correlation* to refer to the correlation within a single SKU's demand structure, rather than the dependence between different SKUs. Positive cross-correlation occurs when a long demand interval is followed by a high demand size, or a low demand size follows a short interval. Negative cross-correlation indicates that a short demand interval is followed by a high demand size or a long interval is followed by a small demand size.

Is auto- and cross-correlation common in spare parts? Willemain et al. (1994) argue that real data frequently exhibit auto and cross-correlation. Magson (1979) mentions that engineering spares in their dataset have negative autocorrelation with high demands frequently being followed by a series of low demands. Similarly, Eaves (2002) found significant number of items with auto- and cross-correlation in spare parts data from the Royal Air Force. In preparing for this study we obtained a dataset of 4588 aircraft service parts from the US DLA in which 35% of items had statistically significant cross-correlation at $\alpha=0.05$ level. Additionally, 27% of items displayed auto-correlated demand size and 22% of items auto-correlated intervals. In contrast, Nikolopoulos et al. (2010) tested 5000 SKUs from the Royal Air Force for autocorrelation but did not find significant signs of it. It is important to note, however, that they tested the whole intermittent time series for autocorrelation and not just the nonzero demand sizes as it was the case in Willemain et al. (1994) and Eaves (2002).

The effect of correlation on forecasting and stock control of intermittent time series has not been thoroughly studied. The only study we are aware of explores the effects of correlation on

* Corresponding author. Tel.: +1 312 362 8313; fax: +1 312 362 6973.
E-mail addresses: naltay@depaul.edu (N. Altay), llittera@richmond.edu (L.A. Litteral), frudisill@uscupstate.edu (F. Rudisill).
[1] Tel.: +1 804 289 8576; fax: +1 804 289 8878.
[2] Tel.: +1 864 503 5511; fax: +1 864 503 5583.

0925-5273/$ - see front matter
doi:10.1016/j.ijpe.2011.08.002

intermittent demand forecasting (Eaves, 2002). Using real data Eaves tests the effect of all three correlation types (auto-correlated sizes and intervals, and cross-correlation) on forecast accuracy. His results imply that best forecast accuracy is obtained when autocorrelation in demand sizes and intervals are negative, and cross-correlation is (in 55% of the cases) zero. However, further analysis shows no statistically significant difference between different types of correlations, leading him to conclude that the effects of correlation is mostly unknown.

The objective of this paper is to test for and characterize significant effects of demand size autocorrelation, interval autocorrelation, and size-interval cross-correlation on forecasting and stock control in the context of intermittent demand. Our review of the literature did not discover any published study investigating such effects on stock control or the integration of forecasting and stock control. In order to make this needed contribution to the literature we conduct a simulation study for which we use theoretically generated intermittent demand series. For the sake of simplicity, from this point on we will use the generic term correlation whenever we refer to any of the three versions of it.

The remainder of the paper is structured as follows: Section 2 summarizes the salient literature. Section 3 reveals the details of our experiments with regards to data, forecasting and stock control procedures and performance measures. Section 4 presents experimental results while Section 5 provides discussion and insights. Finally, in Section 6 we point to further research directions and conclude the paper.

2. Review of literature

2.1. Effects of correlation on non-intermittent demand

Proper planning and control of stock keeping units require a good understanding of their demand structure. Correlation in time series can be a part of this structure. But inventory control theory has been developed upon the assumption of independent demand. Johnson and Thompson (1975) show that when demand is auto-correlated a myopic policy is optimal. Toktay and Wein (2001) add that if demands are correlated and a myopic policy is used then positively correlated demands increase the base-stock level and the resulting cost, while negatively correlated demands decrease both.

Autocorrelation in demand also impacts safety stock levels. Under the conventional assumption of independent demand, positive autocorrelation in demand leads to underestimated safety stock levels while negative autocorrelation will cause overestimation (Ray, 1980; Lau and Wang, 1987; Graves, 1999). Charnes et al. (1995) look at this case using an order-up-to level inventory policy. They show that negative autocorrelation causes the stock-out percentage to be lower than it would otherwise be, while increasing positive autocorrelation causes the stock-out percentage to rise at an increasing rate. This is intuitively appealing because negative autocorrelation implies that we will experience demand that is somewhat less than before. This means that the order levels are set without considering this fact and we are over-estimating demand (Ray, 1981). Similarly, Miller (1986) implies that positive autocorrelation should reduce the order-up-to levels.

Heath and Jackson (1994) show that improved forecasts can reduce safety stocks without affecting service performance. Fotopoulos et al. (1988) also look at the impact of autocorrelation on safety stock and find that (i) positive autocorrelation has stronger impact on safety stock than negative autocorrelation; (ii) as autocorrelation increases safety stock increases and vice versa; (iii) the rate of change in safety stock is increasingly large when the value of autocorrelation is positive and increasing. One explanation for the change in safety stock is that autocorrelation does not affect the variance of demand directly, but rather, it exerts its effect on safety stock through its indirect effect upon the variance of lead time demand (Marmorstein and Zinn, 1993). Positive autocorrelation increases the standard deviation of demand during lead time resulting in larger amounts of safety stock, while negative correlation reduces it (Erkip et al., 1990).

Despite the depth of research on the effects of correlation on forecasting and stock control, we note that none of the papers mentioned above considered the case of intermittent demand. Therefore, the effects of autocorrelation on spare parts demand, where many periods have zero demand have not been yet assessed. Furthermore, the cases of auto-correlated demand intervals as well as the correlation between demand intervals and sizes have not been considered at all. In the following sections we aim to fill this gap in the forecasting and stock control literature through extensive simulation experiments.

2.2. Forecasting and stock control of intermittent demand items

What makes forecasting and stock control of intermittent demand items challenging is the variability in the timing of demand arrival in addition to its size. The seminal work on intermittent demand forecasting belongs to Croston (1972). However, Croston's estimator has been shown to be biased (Syntetos and Boylan, 2001), correction factors to overcome this bias have been presented (Syntetos and Boylan, 2005; Shale et al., 2006) and tested (Boylan and Syntetos, 2007; Teunter and Sani, 2009a). The Syntetos–Boylan Approximation (SBA) is generally found to be the most effective and efficient modification of Croston's method.

Control of inventories related to intermittent demand patterns is typically managed with periodic review systems. The appropriateness of periodic (s,S) systems in the context of intermittent demand have been shown through theoretical arguments (Porteus, 1985; Silver et al., 1998) simulations on real data (Sani and Kingsman, 1997; Babai et al., 2010) and case studies (Porras and Dekker, 2008). Since the calculation of the optimal (s,S) levels can be complex and the exact demand distribution is impossible to estimate in practice several heuristics were developed to estimate replenishment levels. Sani and Kingsman (1997) perform a comparative study of (T,s,S) heuristics and find that the performance differences between Power Approximation (Ehrhardt, 1979), Normal Approximation (Wagner, 1975), and Naddor's heuristic (Naddor, 1975) are generally small. They recommend the use of these three methods as best for intermittent and lumpy demand items. Babai et al. (2010) conduct a similar study but with a rather large dataset and find the Power Approximation and Naddor's heuristic give the best results in terms of average inventory cost.

Lately, the integration of forecasting and stock control has been gaining popularity in the academic literature. For example, Teunter and Sani (2009b) show the application of Croston's forecasts in calculating order-up-to levels. Their results indicate that the calculated order-up-to levels lead to service levels that are close to their expected targets. Syntetos et al. (2009c) investigate the case where the lead times are smaller than the average demand interval. Teunter et al. (2010) calculate order-up-to levels for compound binomial demand. Strijbosch et al. (2011) investigate the integration of forecasting and stock control when demand is non-stationary.

The concise review of salient literature presented above indicates that correlation research has not been extended to the intermittent demand context. We take on this challenge using well accepted and easy to implement forecasting and stock control procedures, which are described in the following section.

3. Experimental structure and data

Theoretically generated data is used in this study for three reasons. The first is to be able to remove the effect of initialization and observe only the effect of correlation. This requires long time series. Although we have access to US DLA data, the rather short length of the available time series (60 periods), given the intermittency of the demand (in some cases a little as 5 out of 60 periods have nonzero demand), makes it difficult to isolate the effect of correlation. Second, the real dataset does not lend to a controlled environment for the experiment because we cannot control the levels of parameters such as average demand interval (ADI) and squared coefficient of variation (CV^2). ADI and CV^2 are the usual criteria used in categorizing sporadic and/or irregular demand patterns (Syntetos et al., 2005). Our interest in this research lies in series with an $ADI > 1.32$, i.e. SKUs with erratic demand. This cut-off for ADI is higher than Johnston and Boylan's (1996) empirically found cut-off point 1.25 and thus assures an irregular demand pattern. Such series are further classified into intermittent ($CV^2 \leq 0.49$) and lumpy ($CV^2 > 0.49$) demand (Syntetos, 2001). Last, theoretically generated data allows us to focus our research on correlation alone since it comes without the noise that real data usually contains.

We study the forecasting and stock control implication of intermittent demand correlation through an extensive set of simulation experiments. We conduct three sets of experiments, each of which is developed upon three sets of theoretically generated demand data. As shown in Fig. 1, experiment (I) investigates the effect of correlation on forecasting. Experiment (II) focuses on the effect of correlation on inventory performance and represents a sensitivity analysis since the stock control policy uses the known parameters of the demand distribution. Finally, experiment (III) is designed to study the integration of forecasting and stock control since the stock control policy here uses the forecasted mean and variance. All three experiments are run using auto-correlated intermittent demand sizes (Case A), auto-correlated intervals (Case B), and cross-correlated demand sizes and intervals (Case C).

3.1. Theoretically generated data

Compound Poisson demand series with lognormal demand sizes, and negative exponential inter-arrival times (i.e. Poisson arrivals) were generated. The lognormal distribution was chosen for the purpose of our simulation experiment due to two reasons. First, Syntetos (2001) points out that considering its flexibility lognormal can be a good approximation for discrete demand. The lognormal distribution allows us to vary the coefficient of variation to observe effects of erraticity. Second, there is evidence in real world data that suggests the use of the lognormal distribution to represent demand sizes (Willemain et al., 1994; Syntetos, 2001).

Time series data (100 independent streams containing 1000 observations) for each condition of the three types of intermittent demand were obtained using similar methods based on Banks et al. (2005, pp. 337–344) to generate the autoregressive and cross-correlation components. The characteristics of the experimental data generated for each scenario are displayed in Table 1.

A similar algorithm was used for each condition since each condition we considered contains two correlated variables X_1 and X_3 (size–size, interval–interval, or interval–size). Each data series was developed by first generating 1000 bi-variate standard normal values, Z_1 and Z_2. The procedure followed includes:

Step 1: generate 1000 values each for Z_1 and Z_2, independent standard normal random variables.

Step 2: set $Z_3 = \rho Z_1 + \sqrt{(1-\rho^2)}Z_2$, where ρ is the desired correlation level. Z_1 and Z_3 are now correlated (Banks et al., 2005, pp. 340).

Table 1
Characteristics of the experimental data generated.

Case A — auto-correlated intermittent demand size				
Average demand	10	10	10	10
% Zeros	35	35	75	75
CV^2	0.25	1	0.25	1
Autocorrelation coefficients	−0.46	−0.48	−0.47	−0.48
	−0.22	−0.22	−0.22	−0.22
	0	0	0	0
	0.24	0.24	0.24	0.24
	0.47	0.48	0.47	0.47
	0.73	0.73	0.72	0.73
Case B — auto-correlated inter-arrival time				
Average demand	10	10	10	10
% Zeros	35	35	75	75
CV^2	0.25	1	0.25	1
Autocorrelation coefficients	−0.21	−0.21	−0.21	−0.22
	0	0	0	0
	0.21	0.22	0.21	0.21
	0.69	0.69	0.68	0.68
Case C — cross-correlated intermittent demand				
Average demand	10	10	10	10
% Zeros	35	35	75	75
CV^2	0.25	1	0.25	1
Autocorrelation coefficients	−0.48	−0.47	−0.59	−0.58
	−0.35	−0.34	−0.34	−0.34
	0	0	0	0
	0.32	0.32	0.34	0.34
	0.74	0.74	0.73	0.74
	0.87	0.87	0.87	0.86

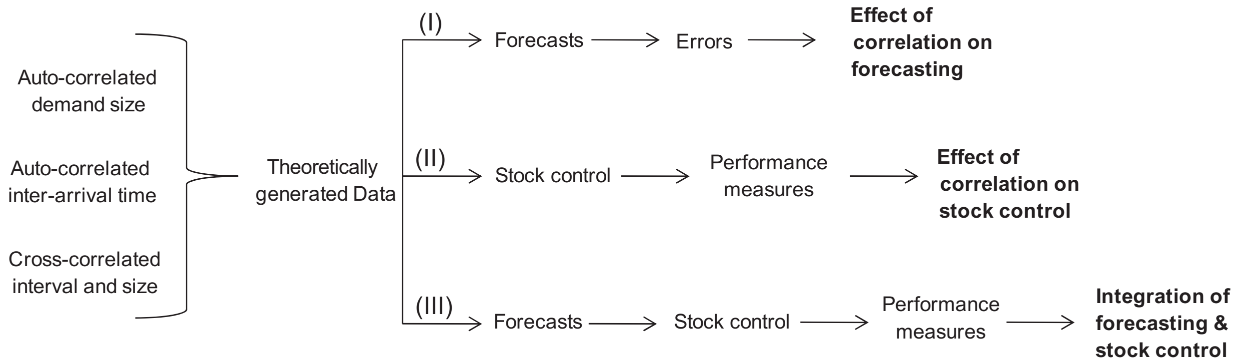

Fig. 1. Experimental structure.

Step 3: use the normal-to-anything transformation (Banks et al., 2005, pp. 342, 343) to calculate X_1 and X_3 where $X_1 = F_1^{-1}[\Phi(Z_1)]$ and $X_3 = F_2^{-1}[\Phi(Z_3)]$.

For each time series, the normal-to-lognormal transformation was used to generate the nonzero demands. Demand intervals were generated using the normal-to-exponential transformation. The function F_i^{-1} is the inverse cumulative distribution function for the desired distribution for X_i (lognormal or exponential in the cases we considered). The variables X_1 and X_3 will have correlation coefficients "close to" ρ. This is not always the case due to the nature of auto-correlated data, particularly in the high CV^2 scenarios.

Two levels for CV^2 (0.25 for intermittent and 1 for lumpy demand), and two levels for ADI (35% and 75% zeros corresponding to 1.54 and 4 periods ADI, respectively) were used to isolate the effect of correlation. Correlation levels change for each case as listed in Table 1. Since it is difficult to control the resulting level of correlation as explained in the procedure above we have uneven numbered scenarios between negative and positive correlation levels. However, the experimental structure is detailed enough to observe and understand any effect of correlation in intermittent demand.

Each generated series is 1000 periods long. Initial forecasts are assumed to be perfect (100% accurate). Once, the forecasting method is initialized, the first 250 periods are then used to optimize the smoothing constant. Once optimal smoothing parameters are identified we used the second 250 periods to initialize the estimates of level and variance of demand for the stock control process. The remaining (i.e. last) 500 periods are used for out-of-sample generation of results and evaluation of performance.

3.2. Forecasting procedure

To test the effects of correlation on forecasting and stock control we selected the corresponding technique as deemed appropriate in the literature. Consequently, the Syntetos–Boylan Approximation (SBA) method seems to be robust enough for forecasting intermittent demand and was therefore selected for our study (Syntetos and Boylan, 2005). The SBA method employs a correction factor to Croston's estimate of mean demand per period F_t (forecast made at the time t for period $t+1$).

$$F_t = \left(1-\frac{\alpha}{2}\right)\frac{z_t}{p_t}, \quad \text{where :} \tag{1}$$

$$z_t = z_{t-1}+\alpha(D_t-z_{t-1}) \quad \text{and} \tag{2}$$

$$p_t = p_{t-1}+\alpha(q_t-p_{t-1}) \quad \text{with} \tag{3}$$

D_t being the demand for an item at time t, q_t the actual demand interval at time t, z_t and p_t estimates of demand size and interval, respectively, and α, the smoothing constant.

Forecast performance is measured using the Relative Geometric Mean Absolute Error (RGMAE), and the Mean Absolute Scaled Error (MASE). RGMAE is selected based on the proposition of Syntetos and Boylan (2005) to use the relative geometric root mean squared error (RGRMSE) for measuring forecast accuracy when demand is intermittent. It is a relative measure; thus, in this paper we compared SBA generated forecasts to Single Exponential Smoothing (SES) forecasts. Later, Hyndman (2006) proves that RGMAE and RGRMSE are essentially the same for intermittent demand, but RGMAE is relatively simpler to calculate. Consequently, RGMAE is utilized as described in the following formula, where $F(A)_t$ represents SBA generated forecast while $F(B)_t$ represents SES forecasts:

$$\text{RGMAE} = \frac{(\prod_{t=1}^{n}|D_t-F(A)_t|)^{1/n}}{(\prod_{t=1}^{n}|D_t-F(B)_t|)^{a/n}} \tag{4}$$

We also utilized MASE, which is developed by Hyndman (2006) as a scale-free error measure. MASE uses naïve forecasts as a benchmark. Let e_t indicate forecast error. The scaled error, SE_t at time t is then calculated using Eq. (5) and MASE is the average of absolute values of SE_t. The denominator of Eq. (5) is simply the Mean Absolute Error for naïve forecasts. A MASE of less than one suggests better forecasts than the naïve method.

$$SE_t = \frac{e_t}{(1/n-1)\sum_{t=2}^{n}|D_t-D_{t-1}|} \tag{5}$$

For both, RGMAE and MASE a value less than one indicates that the SBA method is superior to the benchmark method. Increasing values for both measures indicate a deteriorating forecast performance.

3.3. Stock control procedure

A (T,s,S) type periodic stock control heuristic, namely Power Approximation with a fixed review period $T=1$ is employed in our research for simulation purposes. We chose Power Approximation because it was developed that demand can be represented by a compound Poisson distribution although it does not require knowledge of type of the distribution. More importantly, it does not require the knowledge of the item cost. It does, however, need a shortage cost per unit value short (i.e. backorder cost b). In this study, following Babai et al. (2010), a holding cost $h=1$ and a b/h ratio of 10 are used. This ratio results in an expected service level of 91%. We assumed a lead time of three periods (i.e. $L=3$). The Power Approximation procedure and its parameters are summarized as follows:

$\Psi(\cdot)$ denotes the cumulative standardized normal distribution.
F_t is the estimate of mean demand per time period.
K is the ordering cost (we assumed $K=0.5$ for our experiments).
σ_t is the estimate of the standard deviation of demand per period.
σ_{L+1} is the standard deviation of demand over $L+1$ periods.
μ_{L+1} is mean demand over $L+1$ periods.

Let $\lambda = b/(b+h)$ denote the availability index. The order quantity Q is calculated using

$$Q = 1.3F_t^{0.494}\left(\frac{K}{h}\right)^{0.506}\left(1+\frac{\sigma_{L+1}^2}{F_t}\right)^{0.116} \tag{6}$$

The tentative re-order and order-up-to levels are then calculated with the following two formulas:

$$s_p = 0.973\mu_{L+1}+\sigma_{L+1}\left(\frac{0.183}{z}+1.063-2.192z\right) \tag{7}$$

$$S_0 = F_t(L+1)+\Psi^{-1}(\lambda)\sigma_t\sqrt{L+1} \tag{8}$$

where

$$z = \sqrt{\frac{hQ}{\sigma_{L+1}b}} \tag{9}$$

The (s,S) levels are then set using the following rule:

$$\text{IF}\left(\frac{Q}{F_t} > 1.5\right)\text{THEN } s = s_p \text{ and } S = s_p+Q$$

$$\text{ELSE } s = min\{s_p,S_0\} \text{ and } S = min\{s_p+Q,S_0\} \tag{10}$$

Following the recent trend in literature in experiment (III) we calculate the stock control parameters needed (mean and variance of per-period-demand) from the forecasts. The variance of the demand per period is estimated by using the Mean Square Error MSE_t, where

$$\mathrm{MSE}_t = \gamma(F_{t-1} - D_t)^2 + (1-\gamma)\mathrm{MSE}_{t-1} \tag{12}$$

with γ being the smoothing constant. Initial stock is assumed to be equal to the first order-up-to level S that was calculated, based on the study of Iyer and Schrage (1992). They focus on the application of a deterministic (T,s,S) model with constant lead times and shortage costs. Probabilistic demand is assumed in this study. But because of the rather lengthy initialization period and the similarity of the remaining procedure, this assumption seems reasonable. It is also assumed that no orders are due to arrive at the beginning of the simulation.

The performance measures are Average Service Level (AvgSERV) and Average Cost (AvgCOST). With service level we actually refer to fill rate (i.e. the percentage of demand satisfied from stock on hand). Fill rate is calculated for every period with nonzero demand. The average fill rate for the last 500 periods then gives the AvgSERV. An inventory cost per period is calculated by summing up the costs of ordering, backordering, and carrying stock in each period. AvgCOST is the average of cost per period over the last 500 periods.

4. Simulation results

The experimental results confirm the general consensus in the literature that correlation in intermittent demand time series has a significant impact on forecasting and stock control performance. In general, as positive cross-correlation improves forecasting performance, positive autocorrelation in demand size deteriorates it. When stock control is of concern, overall averages of cost and service level do not show significant differences between different sources of correlation. The overall results for the integration of forecasting and stock control (i.e. mean and variance are estimated), however, show that while service levels are not significantly different, inventory cost decreases from the cross-correlation case,

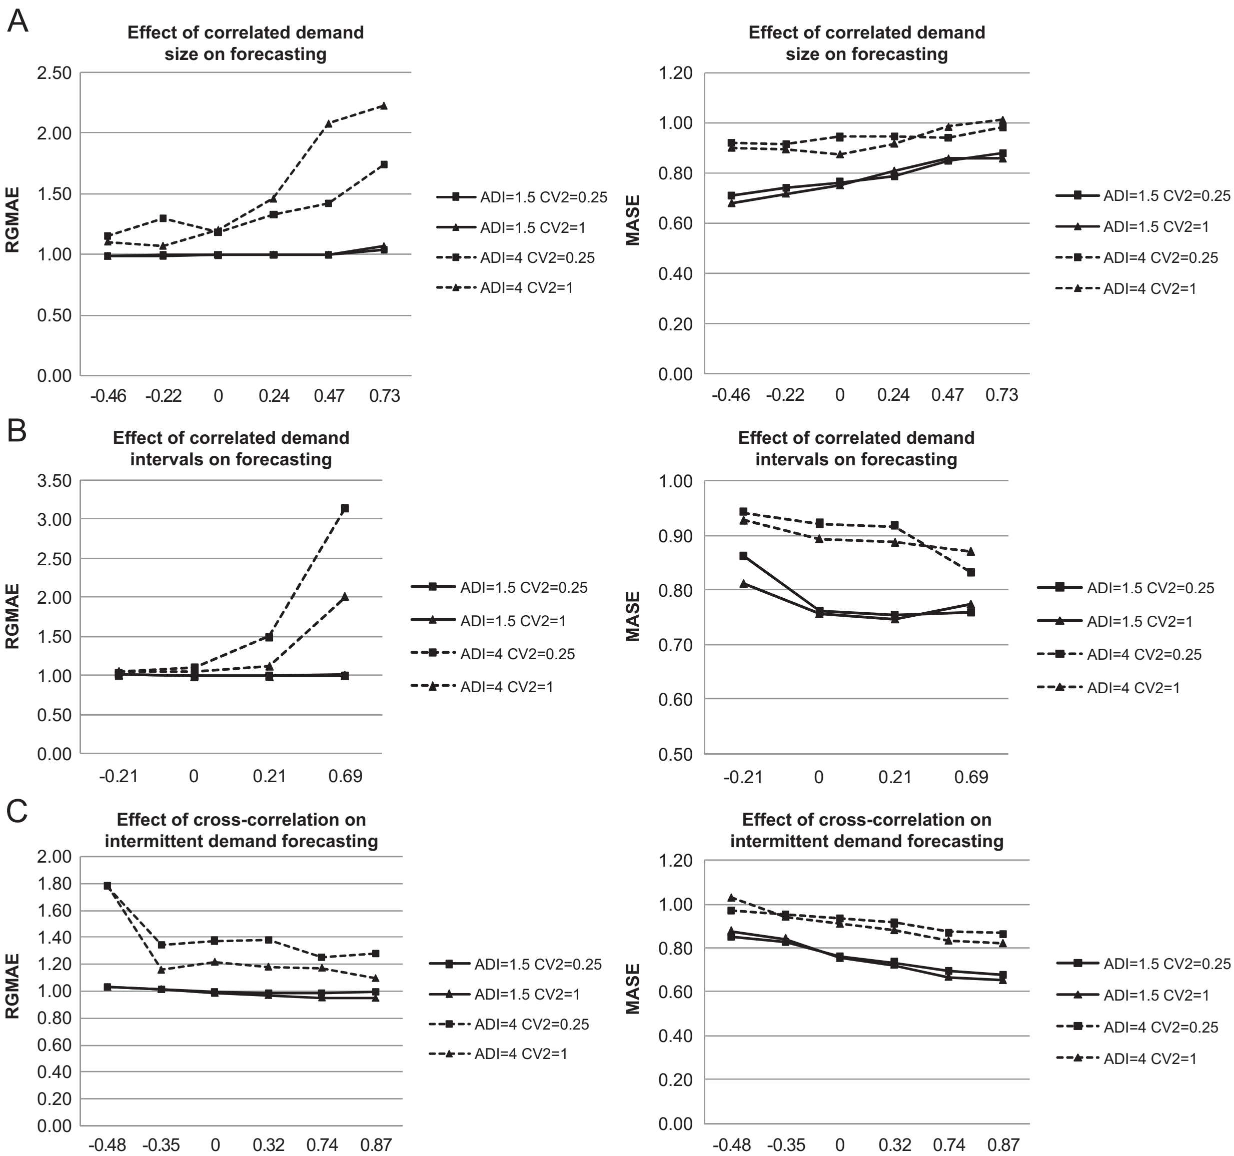

Fig. 2. Experiment I—effect of correlation on forecasting. Case A—auto-correlated demand size, Case B—auto-correlated demand interval, and Case C—cross-correlated intermittent demand.

to the auto-correlated demand size, to the auto-correlated interval case.

Detailed results of the simulation experiments are presented in Figs. 2–4. Interested readers can contact the corresponding author for more details on the analysis of variance results for all three experiments and for each case of correlation. The following sub-sections summarize the results and Section 5 discusses the managerial insights.

4.1. Experiment I: effect of correlation on forecasting

Correlation clearly affects forecasting performance. For auto-correlated demand size, Case A, RGMAE shows a difference only for the high ADI case, where RGMAE increases for positive correlation levels. For these high ADI time series the drop in forecast accuracy is worse when CV^2 is also high. For low ADI levels (i.e. ADI$=$1.5) the change in RGMAE is not statistically significant. For the high ADI time series the results for MASE are similar. Autocorrelation in demand sizes seems to make forecasting more difficult from negative to positive correlation. MASE results also show that, for a given correlation coefficient, it is the ADI level that decides on the forecasting performance rather than CV^2. For MASE, the majority of variation is due to ADI (77%) and the autocorrelation coefficient is statistically significant contributing 17% of the variability.

For Case B, auto-correlated intervals, RGMAE results are similar to Case A but this time the high ADI scenario has a bigger impact on forecast performance. For autocorrelation values between $r=-0.21$ and $+0.21$ RGMAE shows no substantial change. Results for MASE, on the other hand, generally imply that forecasts get relatively better as autocorrelation of intervals moves from negative to positive.

For cross-correlated intermittent demand, Case C, results for RGMAE do not show a substantial change over the range of correlation levels tested, except for the high ADI scenario at $r=-0.48$. For this specific case forecast performance drops drastically. On the other hand, MASE results indicate, that as cross-correlation increases (from negative to positive) MASE decreases. That means, with respect to MASE, it is easier to forecast positively cross-correlated demand compared to negatively cross-correlated demand. About 27% of the variation in MASE is due to the cross-correlation effect.

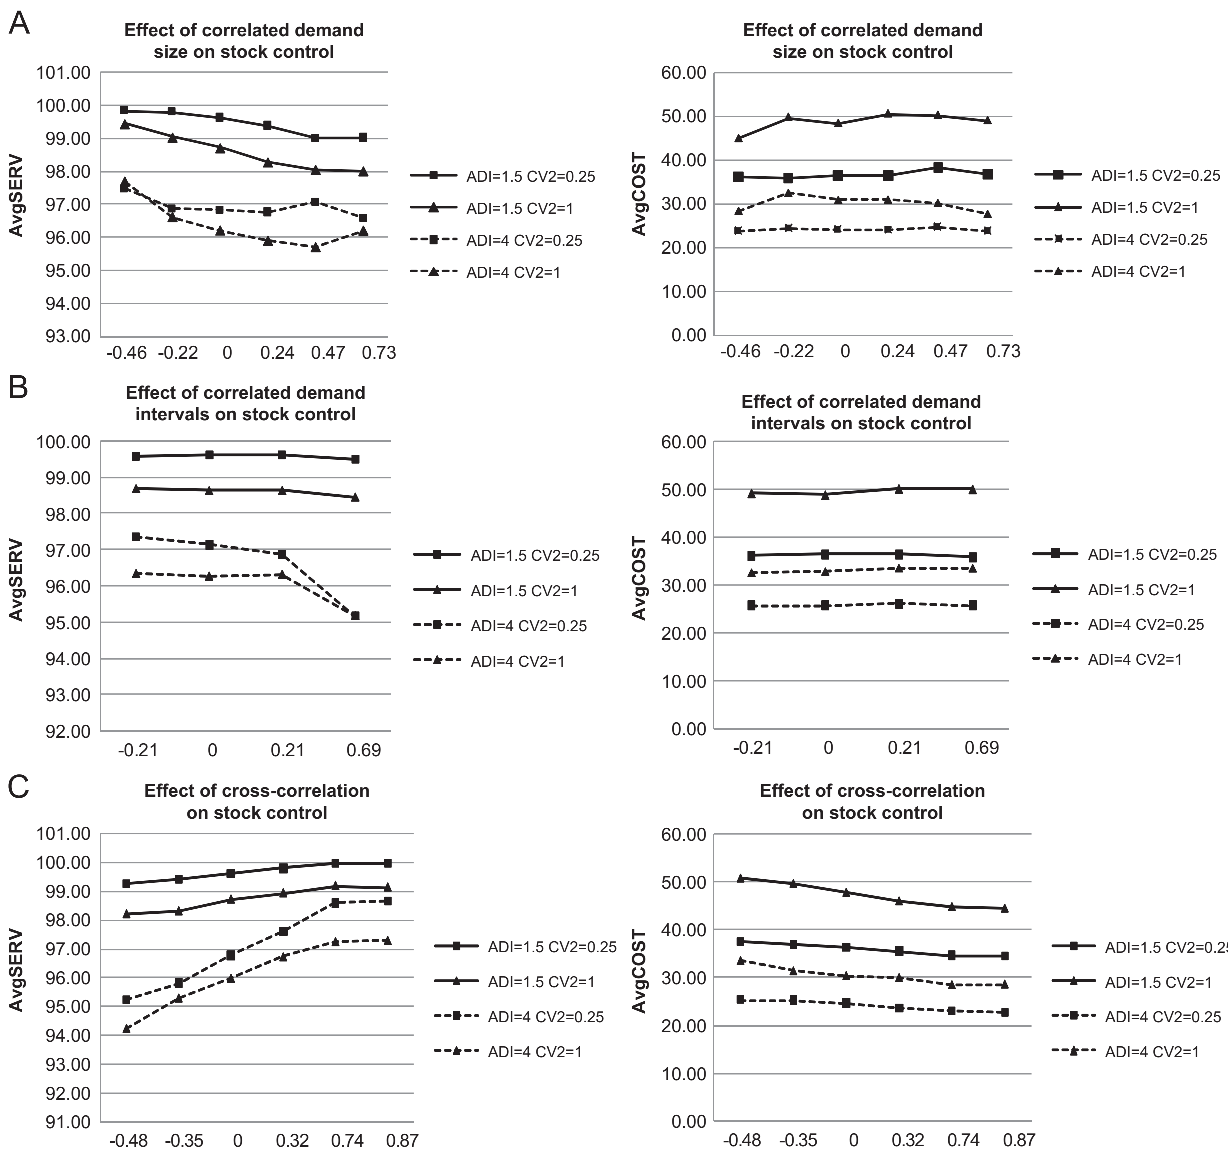

Fig. 3. Experiment II—effect of correlation on stock control. Case A—auto-correlated demand size, Case B—auto-correlated demand interval, and Case C—cross-correlated intermittent demand.

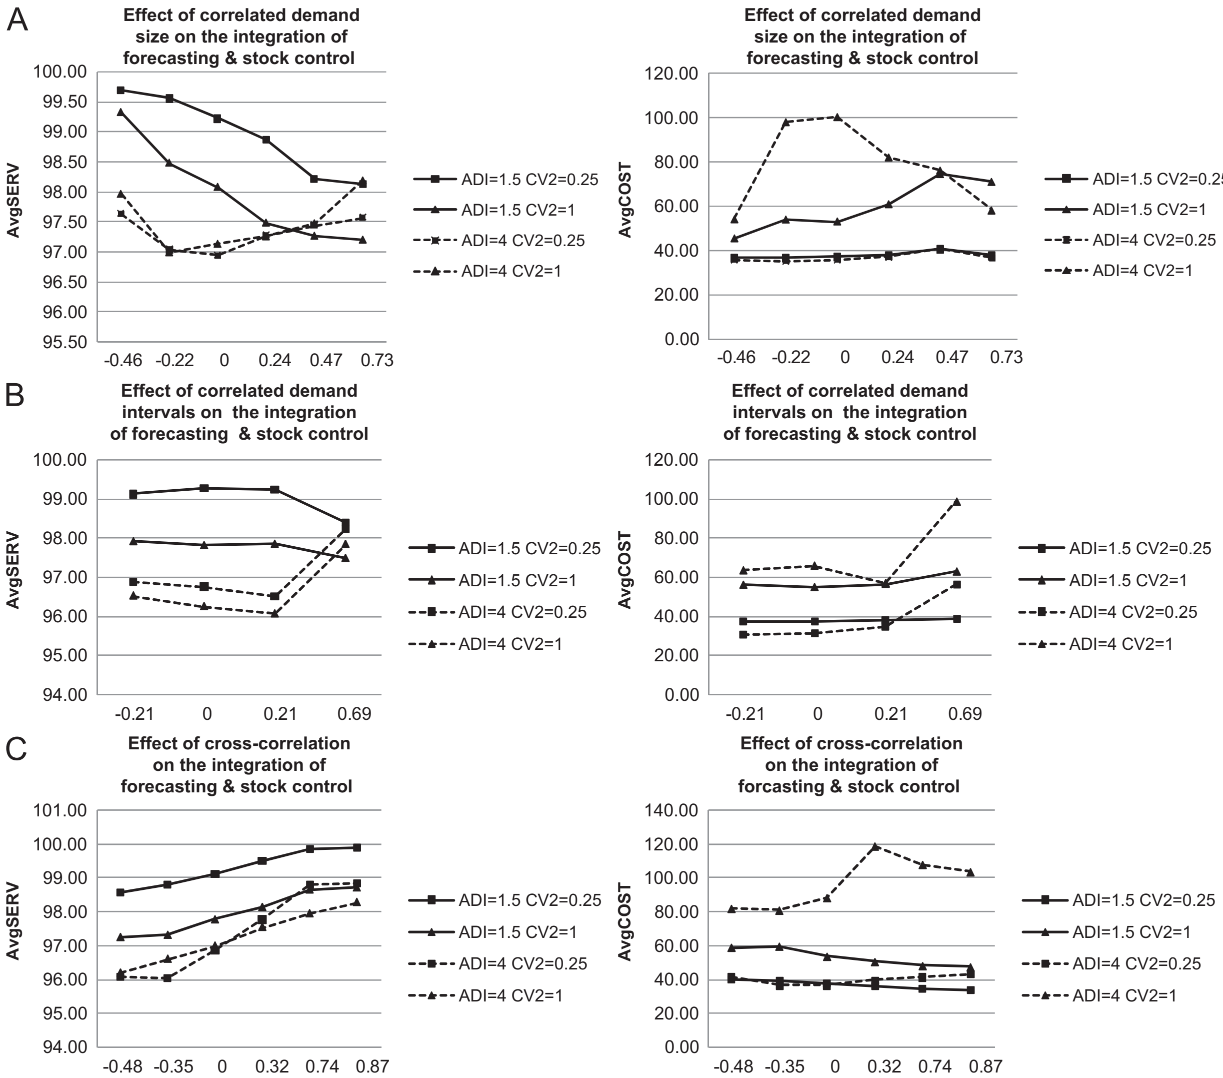

Fig. 4. Experiment III—effect of correlation on the integration of forecasting and stock control. Case A—auto-correlated demand size, Case B—auto-correlated demand interval, and Case C—cross-correlated intermittent demand.

It is the dual variability in demand size and demand timing that makes intermittent demand difficult to forecast. Consequently we expected that as ADI and CV^2 change so would the difficulty of forecasting as represented by the error measures RGMAE and MASE. Fig. 2 indicates that high ADI values generally make forecasting difficult.

4.2. Experiment II: effect of correlation on stock control

The graphs in Fig. 3 indicate that for both types of autocorrelation tested service level generally decreases when correlation is largely positive. This effect is greater for the high level of ADI when the autocorrelation coefficient is 0.69 for the case of autocorrelated intervals; Case B. The opposite is however observed for cross-correlation, Case C. Increasing cross-correlation increases AvgSERV. As cross-correlation increases, there is a statistically significant decrease in AvgCOST for each combination of ADI and CV^2. Theoretical data results suggest that autocorrelation (size or interval) has negligible impact on AvgCOST. For cross-correlation, cost decreases as the correlation coefficient moves from negative to positive.

For all cases in Experiment II it can be said that high ADI levels and high CV^2 levels affect AvgSERV negatively. The reduction in service level is more evident for lumpy demand, i.e. when CV^2 is high. But while average cost is lower with the high ADI level, it is higher for the high CV^2 level. These results are not surprising. When ADI increases, intermittency in the time series increases, hence the total number of units demanded in a given time span decreases (simply because there are less nonzero demand periods). This reduces average inventory carried, and therefore reduces cost, but also the service level. On the other hand, increasing variability in demand sizes as measured by the coefficient of variation increases uncertainty, therefore reducing service level, and increasing cost due to backorders.

4.3. Experiment III: integration of forecasting and stock control in the presence of correlation

For these experiments the mean and variance of demand are drawn from the forecasts to calculate the re-order and order-up-to levels. Fig. 4 indicates that the integration of forecasting and stock control produces interesting results especially for the high intermittency and high CV^2 scenarios. Generally speaking, cases with high ADI result in lower service levels. However, in each of the three correlation types there is a cross-over effect at highly positive correlation coefficients, where service levels for high ADI scenarios catch or surpass low ADI service levels. When ADI is low, low CV^2 leads to better service levels. But when ADI is high the results are not consistent across all three correlation types. In the case of AvgCOST, high ADI and high CV^2 result in the highest cost for all correlation types and levels.

Another interesting observation is that for the low ADI scenarios highly positive correlation coefficients reduce AvgSERV for both types of autocorrelation, but increase it for cross-correlation. For auto-correlated demand sizes, Case A, increasing correlation levels reduces service level when ADI is low. But the same scenario produces more of a U-shaped curve for high ADI cases, where the lowest point of the curve is close to zero.

For auto-correlated intervals, Case B, highly positive correlation coefficients reduce AvgSERV for low ADI cases, while increasing service levels for high ADI cases. The increase in AvgSERV for the high ADI cases is associated with a cost increase for the same data. However, the same pattern is not evident when ADI is low. When the autocorrelation coefficient is 0.69, AvgCOST increases substantially in each case but more so for high ADI.

Probably the clearest results in this section are for service levels for cross-correlated intermittent demand, Case C. AvgSERV increases as we move from negative cross-correlation to positive. Here correlation level contributes 29% of the variation in AvgSERV. The effect of cross-correlation on AvgCOST on the other hand, is not statistically significant.

5. Discussion and insights

In this simulation study we focused on the effects of three different types of correlation in intermittent demand series on forecasting and stock control performance. Results for the forecasting experiment show that the two accuracy measures used – RGMAE and MASE – are consistent regarding the finding that positive autocorrelation reduces forecasting performance, while the opposite is true for positive cross-correlation. This result although in line with Eaves' (2002) findings is counterintuitive because positive autocorrelation implies that data has the tendency to remain the same from one observation to another. This means positively auto-correlated data should be easier to forecast. One explanation for this phenomenon may come from intermittency (i.e. ADI). Altay et al. (2008) showed that the ADI level complicates forecasting intermittent demand more than the CV^2. The drastic difference between the two ADI levels in forecasting performance (Cases A and B in Fig. 2) may indicate that the effect of intermittency counters the impact of autocorrelation on forecast accuracy. As a matter of fact this negative impact of intermittency on forecasting and stock control is persistent and visible in Figs. 3 and 4.

Changes in service level have forecast implications. It is reasonable to expect service level to suffer when underestimated forecasts are used as inputs to a stock control system. Consequently, RGMAE does not seem to be a robust accuracy measure for intermittent demand. This connection between forecast accuracy and inventory service level is clearly observed when MASE is used as a measure of forecast accuracy. In every case MASE increases in Fig. 2, service level decreases in Fig. 4, and vice versa. Even though both of these measures have been recommended for intermittent demand forecasting in the literature, the experimental results leads the reader to conclude that the search for error measures appropriate for intermittent demand is not over yet.

The effect of correlation on stock control was measured by isolating the order-up-to inventory policy from the forecasting procedure using the mean and variance of the theoretically generated data directly. Since the (T,s,S) policy we utilized is not specifically designed to account for correlation in data, these experimental results give an indication of the consequences of an inventory manager's behavior who does not consider that correlation exists in the demand data when setting stocking levels. Inventory control literature on auto-correlated demand suggests that not considering the existence of correlation in demand should lead to lower (higher) average inventory levels and thus higher (lower) shortages for positive (negative) autocorrelation. Our experimental results for auto-correlated demand sizes confirm this behavior. Since most inventory theory is built on the assumption of independent demand, the existence of negative autocorrelation means that we will experience demand that is somewhat less than before. Consequently, negative correlation coefficients produce higher service levels than positive autocorrelation. Meanwhile, there was no significant change in cost. The same experiment set showed that positively auto-correlated intervals also lead to lower service levels, especially when intermittency is high.

Inventory managers should be aware that to the contrary of autocorrelation literature cross-correlation of intermittent demand actually produces results in the opposite direction; that is, negative cross-correlation leads to lower service levels and higher cost, while significantly positive cross-correlation values lead to high service levels and lower cost. Positive cross-correlation means that a rather long interval will be followed by a large demand and a short interval by a small demand. This means that when a large demand figure is recorded the re-order and order-up-to levels will increase; increasing the inventory level. Consequently, there will be either adequate or abundant items in stock regardless of how large or small the next demand figure is. However, in the case of negative cross-correlation long intervals are followed by small demand values and short intervals with large. That means the stock control policy will not have enough time to react to a large demand figure following a small one causing expensive backorders. Hence service level deteriorates while cost increases. This behavior is highly significant when intermittency is high. Furthermore, this conclusion also holds even when the mean and variance of demand is estimated through a forecasting method.

Forecasting and stock control (and their relationship) are among the most significant operational issues for many firms. Understanding the central tendency, variability, and the distribution of demand are crucial to developing effective systems and most researchers and practitioners are well aware of their importance. When the demand is intermittent, the need to consider the additional factor of many periods with zero demand makes the problem even harder. Within this larger context the potential effects of correlation in any of the forms described in this paper are often not considered. So one may conclude, based on the results provided here, that these effects should be considered and systems should be modified to improve performance.

6. Conclusions and further research

Intermittent demand items such as spare parts have been shown to have auto- or cross-correlation, with the latter indicating that demand sizes are correlated with demand intervals. However, the effects of auto and cross-correlation in intermittent demand items have not been sufficiently explored. This is, to the best of our knowledge, the first study investigating such effects by considering three different types of correlation; auto-correlated demand sizes, auto-correlated intervals, and cross-correlated demand sizes and intervals. Using theoretically generated compound Poisson demand data the impact of correlation on forecasting, stock control, and the integration of these two was simulated.

Using the Syntetos–Boylan Approximation to estimate demand mean and variance, and Power Approximation to calculate the re-order and order-up-to levels of a periodic replenishment policy, we show that correlation in intermittent demand does play a role in forecast accuracy and stock control performance. Negative autocorrelation levels lead to higher service levels than positive autocorrelation, while cost does not significantly change. Our results also show that high intermittency in demand intensifies these changes in service level. Cross-correlation on the other hand, produces results in the opposite direction of autocorrelation in size or intervals; that is, positive (negative) cross-correlation leads to higher (lower) service levels.

Our research for this paper highlights three major future research directions. First, results for the forecasting experiment indicate that the search for a good error measure for intermittent demand forecasting should continue. Results using MASE and RGMAE are not consistent in some cases. Even though both measures show promise, more research is needed to find robust error measures for intermittent demand forecasting. Secondly, this study, due to its exploratory nature, utilizes simulation to investigate the impact of correlation on forecasting and stock control of intermittent demand items. The obvious next step now is to take an analytical approach to develop closed form solutions for stock control. Such a study would also bridge the theoretical results with the experimental. Lastly, this experiment should be replicated on other experimental settings as a validation and confirmation of the results presented here. Effects of correlation should also be tested on real data series due to its practical implications.

References

Akkerman, R., Van Donk, D.P., 2009. Product mix variability with correlated demand in two-stage food manufacturing with intermediate storage. International Journal of Production Economics 121, 313–322.

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, 389–408.

Babai, M.Z., Syntetos, A.A., Teunter, R., 2010. On the empirical performance of (*T*,*s*,*S*) heuristics. European Journal of Operational Research 202, 466–472.

Banks, J., Carson, J.S., Nelson, B.L., Nicol, D.M., 2005. Discrete-Event System Simulation, fourth ed. Pearson Prentice Hall, Upper Saddle River, NJ.

Boylan, J.E., Syntetos, A.A., 2007. The accuracy of a modified Croston procedure. International Journal of Production Economics 107, 511–517.

Charnes, J.M., Marmorstein, H., Zinn, W., 1995. Safety stock determination with serially correlated demand in a periodic-review inventory system. Journal of the Operational Society 46 (8), 1006–1013.

Croston, J.D., 1972. Forecasting and stock control for intermittent demands. Operational Research Quarterly 23, 289–304.

Eaves, A.H.C., 2002. Forecasting for the Ordering and Stock-holding of Consumable Spare Parts. Doctoral Dissertation, Lancaster University.

Erkip, N., Hausman, W.H., Nahmias, S., 1990. Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands. Management Science 36 (3), 381–392.

Ehrhardt, R., 1979. The power approximation for computing (s, S) inventory policies. Management Science 25, 777–786.

Fotopoulos, S., Wang, M., Rao, S.S., 1988. Safety stock determination with correlated demands and arbitrary lead times. European Journal of Operational Research 35, 172–181.

General Accounting Office, 2003. Defense Inventory: Several Actions are Needed to Further DLA's Efforts to Mitigate Shortages of Critical Parts. Report to the Chairman, Subcommittee on Defense, Committee on Appropriations, House of Representatives. GAO-03-709.

Graves, S.C., 1999. A single-item inventory model for a nonstationary demand process. Manufacturing and Service Operations Management 1 (1), 50–61.

Heath, D.C., Jackson, P.L., 1994. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. IIE Transactions 26 (3), 17–30.

Hyndman, R.J., 2006. Another look at forecast-accuracy metrics for intermittent demand. Foresight: The International Journal of Applied Forecasting 4, 43–46.

Iyer, A.V., Schrage, L.E., 1992. Analysis of the deterministic (*s*,*S*) inventory problem. Management Science 38, 1299–1313.

Johnson, G.D., Thompson, H.E., 1975. Optimality of myopic inventory policies for certain dependent demand processes. Management Science 21 (11), 1303–1307.

Johnston, F.R., Boylan, J.E., 1996. Forecasting for items with intermittent demand. Journal of the Operational Research Society 47 (1), 113–121.

Johnston, F.R., Boylan, J.E., Shale, E.A., 2003. An examination of the size of orders from customers, their characterization and the implications for inventory control for slow moving items. Journal of the Operational Research Society 54, 833–837.

Lau, H., Wang, M., 1987. Estimating the lead-time demand distribution when the daily demand is non-normal and autocorrelated. European Journal of Operational Research 29, 60–69.

Liu, L., Yuan, X., 2000. Coordinated replenishments in inventory systems with correlated demands. European Journal of Operational Research 123, 490–503.

Magson, D.W., 1979. Stock control when the lead time cannot be considered constant. Journal of the Operational Society 30 (4), 317–322.

Marmorstein, H., Zinn, W., 1993. A conditional effect of autocorrelated demand on safety stock determination. European Journal of Operational Research 68, 139–142.

Miller, B.L., 1986. Scarf's state reduction method, flexibility, and a dependent demand inventory model. Operations Research 34 (1), 83–90.

Naddor, E., 1975. Optimal and heuristic decisions on single and multi-item inventory systems. Management Science 21, 1234–1249.

Nikolopoulos, K., Syntetos, A.A., Boylan, J.E., Petropoulos, F., Assimakopoulos, V., 2010. An aggregate–disaggregate intermittent demand approach (ADIDA) to forecasting: an empirical proposition and analysis. Journal of the Operational Research Society (Advance online publication April 28. doi:10.1057/jors.2010.32).

Porras, E.M., Dekker, R., 2008. An inventory control system for spare parts at a refinery: an empirical comparison of different reorder point methods. European Journal of Operational Research 184, 101–132.

Porteus, E.L., 1985. Numerical comparisons of inventory policies for periodic review systems. Operations Research 33, 132–152.

Ray, W.D., 1980. The significance of correlated demands and variable lead times for stock control policies. Journal of the Operational Society 31 (2), 187–190.

Ray, W.D., 1981. Computation of reorder levels when the demands are correlated and the lead time random. Journal of the Operational Society 32 (1), 27–34.

Sani, B., Kingsman, B.G., 1997. Selecting the best periodic inventory control and demand forecasting methods for low demand items. Journal of the Operational Society 48, 700–713.

Shale, E.A., Boylan, J.E., Johnston, F.R., 2006. Forecasting for intermittent demand: the estimation of an unbiased average. Journal of the Operational Research Society 57, 588–592.

Silver, E.A., 1981. Operations research in inventory management: a review and critique. Operations Research 29, 628–645.

Silver, E.A., Pyke, D.F., Peterson, R., 1998. Inventory Management and Production Planning and Scheduling, third ed. John Wiley & Sons, Inc, New York.

Strijbosch, L.W.G., Syntetos, A.A., Boylan, J.E., Janssen, E., 2011. On the interaction between forecasting and stock control: the case of non-stationary demand. International Journal of Production Economics 133, 470–480.

Syntetos, A.A., 2001. Forecasting of Intermittent Demand. Unpublished Ph.D. Thesis, Buckinghamshire Business School, Brunel University.

Syntetos, A.A., Boylan, J.E., 2001. On the bias of intermittent demand estimates. International Journal of Production Economics 71, 457–466.

Syntetos, A.A., Boylan, J.E., 2005. The accuracy of intermittent demand estimates. International Journal of Forecasting 21, 303–314.

Syntetos, A.A., Boylan, J.E., Croston, J.D., 2005. On the categorization of demand patterns. Journal of the Operational Research Society 56, 495–503.

Syntetos, A.A., Keyes, M., Babai, M.Z., 2009a. Demand categorisation in a European spare parts logistics network. International Journal of Operations and Production Management 29, 292–316.

Syntetos, A.A., Nikolopoulos, K., Boylan, J.E., Fildes, R., Goodwin, P., 2009b. Integrating management judgement into intermittent demand forecasts. International Journal of Production Economics 118, 72–81.

Syntetos, A.A., Babai, M.Z., Dallery, Y., Teunter, R., 2009c. Periodic control of intermittent demand items: theory and empirical analysis. Journal of the Operational Research Society 60, 611–618.

Syntetos, A.A., Babai, M.Z., Davies, J., Stephenson, D., 2010. Forecasting and stock control: a study in a wholesaling context. International Journal of Production Economics 127, 103–111.

Teunter, R., Sani, B., 2009a. On the bias of Croston's method. European Journal of Operational Research 194, 177–183.

Teunter, R., Sani, B., 2009b. Calculating order-up-to levels for products with intermittent demand. International Journal of Production Economics 118, 82–86.

Teunter, R., Syntetos, A.A., Babai, M.Z., 2010. Determining order-up-to levels under periodic review for compound binomial (intermittent) demand. European Journal of Operational Research 203, 619–624.

Toktay, L.B., Wein, L.M., 2001. Analysis of a forecasting–production–inventory system with stationary demand. Management Science 47 (9), 1268–1281.

Wagner, H.M., 1975. Principles of Management Science (With Applications to Executive Decisions). Prentice-Hall, Inc, Englewood Cliffs, NJ.

Willemain, T.R., Smart, C.N., Shockor, J.H., DeSautels, P.A., 1994. Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. International Journal of Forecasting 10, 529–538.

Zhang, R.Q., 1999. Expected time delay in multi-item inventory systems with correlated demands. Naval Research Logistics 46, 671–688.