


2-5-2016

Development and evaluation of a rolling horizon purchasing policy for cores

Toyin A. Clottey
Iowa State University, tclottey@iastate.edu

Follow this and additional works at: https://lib.dr.iastate.edu/scm_pubs

 Part of the [Business Administration, Management, and Operations Commons](#), [Business Analytics Commons](#), [Management Information Systems Commons](#), and the [Operations and Supply Chain Management Commons](#)

The complete bibliographic information for this item can be found at https://lib.dr.iastate.edu/scm_pubs/83. For information on how to cite this item, please visit <http://lib.dr.iastate.edu/howtocite.html>.

This Article is brought to you for free and open access by the Supply Chain and Information Systems at Iowa State University Digital Repository. It has been accepted for inclusion in Supply Chain and Information Systems Publications by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Development and evaluation of a rolling horizon purchasing policy for cores

Abstract

A number of companies utilise end-of-use products (i.e. cores) for remanufacturing or recycling. An adequate supply of cores is needed for such activities. Establishing a purchasing policy for cores, over a finite planning horizon, requires multi-step ahead forecasts. Such forecasts are complicated by the fact that the number of cores in any future period depends upon previous sales and recent returns of the product. Distributed lag models have been used to capture this dependency for single-period ahead forecasts. We develop an approach to use distributed lag models to make multi-period ahead forecasts of net demand (i.e. demand minus returns), and investigate the cost implications, at a prescribed service level, of using such forecasts to purchase cores on a rolling horizon basis. Our results indicate that the effects of errors in the sales forecasts are negligible if sales follow an autoregressive pattern but are substantial when sales are more random. Dynamic estimation of the parameters in a rolling horizon environment yielded the most cost savings at high prescribed service levels (i.e. >0.95). Collectively, our results demonstrate the conditions in which companies can best leverage the dynamic nature of distributed lag models to reduce the acquisition costs over a finite horizon.

Keywords

purchasing policies, forecasting, closed-loop supply chain

Disciplines

Business Administration, Management, and Operations | Business Analytics | Management Information Systems | Operations and Supply Chain Management

Comments

This accepted article is published as Clotey, T., Development and evaluation of a rolling horizon purchasing policy for cores. *International Journal of Production Research*, 2016 54(9); 2780-2790. doi: [10.1080/00207543.2016.1142133](https://doi.org/10.1080/00207543.2016.1142133). Posted with permission.

Development and Evaluation of a Rolling Horizon Purchasing Policy for Cores

A number of companies utilize end-of-use products (i.e., cores) for remanufacturing or recycling. An adequate supply of cores is needed for such activities. Establishing a purchasing policy for cores, over a finite planning horizon, requires multi-step ahead forecasts. Such forecasts are complicated by the fact that the number of cores in any future period depends upon previous sales and recent returns of the product. Distributed lag models have been used to capture this dependency for single-period ahead forecasts. We develop an approach to using distributed lag models to make multi-period ahead forecasts of net demand (i.e., demand minus returns), and investigate the cost implications, at a prescribed service level, of using such forecasts to purchase cores on a rolling horizon basis. Our results indicate that the effects of errors in the sales forecasts are negligible if sales follow an autoregressive pattern but are substantial when sales are more random. Dynamic estimation of the parameters in a rolling horizon environment yielded the most cost savings at high prescribed service levels (i.e., >0.95). Collectively, our results demonstrate the conditions in which companies can best leverage the dynamic nature of distributed lag models to reduce the acquisition costs over a finite horizon.

Keywords: purchasing policies, forecasting, closed-loop supply chain

1. Introduction

A number of organizations utilize end-of-use products (i.e., cores) for activities such as remanufacturing or recycling. An adequate supply of cores must be on hand to satisfy the demand for final products resulting from such activities. The acquisition of cores to support reuse activities requires careful planning in order to avoid the uncontrolled accumulation of inventory, or unacceptable levels of customer service (i.e., insufficient cores to meet demand). In order to establish a purchasing policy for cores, over a finite planning horizon, it is necessary to forecast several periods ahead. Such forecasts are complicated by the fact that the number of cores in any future period depends upon previous sales and recent returns of the product. A dynamic regression (distributed lag) model is often used to capture this dependency in order to

make single period forecasts. When future sales estimates are used to forecast product returns, then of course the standard error of this forecast is larger than that which would result if future sales were known without error. However, it has been shown (Pierce, 1975), under certain conditions, that even when the explanatory variables are forecasted in a distributed lag model, the single period forecast variance of the model's dependent variable is at most as large as the variance resulting from forecasting the dependent variable on the basis of its past history alone. Thus the use of distributed lag models for forecasting product returns holds promise for use in planning based on a rolling horizon. Previous applications of the distributed lag model to forecast product returns, for inventory management or production planning purposes, have been mostly limited to single period forecasts based on known sales. In this paper we explore the conditions in which to best leverage distributed lag models to make multi-period ahead forecasts of product returns, and investigate the cost implications, at a prescribed service level, of using such forecasts in a rolling horizon planning environment.

Planning (e.g., production or acquisition) decisions are usually made on a rolling horizon basis. The typical sequence of activities is as follows: a plan is made for a fixed number of periods, the first decision is implemented and the horizon is then rolled forward to the period where the next decision needs to be made. This results in a purchasing policy in which periodic acquisition plans are made, with each plan covering a certain fixed number of periods ahead. A purchase taking place in a period incurs a fixed setup cost; there is a marginal cost per unit purchased, and a holding cost for carrying inventory from one period to the next. Kelle and Silver (1989b) developed a model for purchasing returnable containers on a rolling horizon basis, using such cost considerations. Their model was based on forecasting net demand, that is, the consumer demand minus the number of returned products. Thus periods of negative net demand

would require replenishment via purchased cores. In a previous paper (i.e., Kelle and Silver 1989a), a normal approximation to a multinomial distribution for returns had been developed for use in forecasting the net returns. A similar approach to forecasting has been applied more recently by Krapp, Nebel and Sahamie (2013b) in an Economic Value Analysis (EVA) framework, using a Bayesian approach to perform the estimation. The distributed lag model approach to forecasting product returns is considered superior to the normal approximation approach due to its dynamic nature (Toktay, Wein and Zenios, 2000; Toktay, 2004; Clottey, Benton, and Srivastava, 2012), and small sample performance (Clottey et al., 2012). However, previous applications of the distributed lag model to forecasting product returns (e.g., Toktay et al. 2000, Clottey et al. 2012; Krapp, Nebel and Sahamie, 2013a) have not been on a rolling horizon basis. Toktay et al. (2000), developed a distributed lag model to forecast product returns in which the lags after which a sold product is returned were represented by discrete geometric and negative binomial delay functions. Krapp, Nebel and Sahamie (2013a) used a distributed lag type model with a discrete Poisson delay function to forecast product returns. Clottey et al. (2012) showed that a distributed lag model with a discrete delay function used to estimate continuous time, can lead to biases when used to forecast product returns. They developed a method for estimating a continuous exponential delay function in the distributed lag model. Clottey and Benton (2014) extended this approach to forecasting product returns with any properly specified continuous delay function. These previous studies typically employed the distributed lag model to make *single* period ahead forecasts. A rolling horizon requires the estimation of future sales in order to forecast product returns in multiple future periods. This leads to additional forecasting complications, which need to be resolved before the distributed lag model can be successfully employed on a rolling horizon basis. Krapp et al. (2013b), in a

multi-period ahead context, investigated the use of a forecasting model of product returns based on the “large sample” normal approximation for the sum of Bernoulli random variables. In this study we develop an approach to using the distributed lag model to make multi-step ahead forecasts in a rolling horizon purchasing context.

There have been other studies which have considered various methods for incorporating previous sales into the forecast of product returns, including simulation based approaches (e.g., Srivastava and Srivastava, 2006), fuzzy logic based approaches (Hanafi, Kara, Kaebernick (2009), and use of Markov chains (e.g., Kiesmuller and van der Laan, 2001). Krapp et al. (2013b, p990) noted that the distributed lag model “can be seen as the most accurate approach to estimate product returns” among these previous approaches. A methodological review of these forecasting approaches can be found in Toktay (2004). Similarly, there have been various methods used to include forecasts of product returns in models to facilitate inventory management, including queuing based approaches (e.g., Toktay et al. 2000), EVA based approaches (e.g., Guide and van Wasenhove, 2001; Krapp et al. 2013b), and mathematical optimization (e.g., Kelle and Silver, 1989b). A review of these approaches can be found in Souza (2013). There are however two issues which have not been investigated in these previous studies: 1) how much does the accuracy in the sales forecast affect the cost performance of the rolling horizon acquisition plan? and 2) how big a cost benefit is there in using a dynamic approach to forecasting (e.g., using Bayesian estimation with parameters of the net demand distribution updated over time), for the rolling horizon purchasing policy, versus the statistic approach (i.e., maximum likelihood estimation of the asymptotic distribution for the distributed lag model)? These are two important questions which we seek to address in this paper.

The next section of this paper describes methods developed by Kelle and Silver (1989b) to create a purchasing policy using forecasts of net demand. In section 3 we develop an approach to making multi-period ahead forecasts using the distributed lag model to forecast the net demand distribution, during the lead time, for a finite horizon acquisition plan. In section 4 we investigate the cost considerations, at a prescribed service level, of the distributed lag model approach for multi-period purchasing decisions. We also compare the accuracy of the distributed lag model approach to that of a widely used forecasting method. In the last section, we discuss the substantive and methodological implications of the results for acquiring cores in closed-loop supply chains.

2. Optimization Model

The optimization model is based on the formulation found in Kelle and Sliver (1989b) with notation as follows. The net demand to be satisfied in each period is the estimated demand for each period (i.e., d_t) minus the number of returned products (i.e., m_t) that can be reused. Since the product returns are random, the net demand is also random. A key objective when sourcing cores is to minimize the total of acquisition and expected inventory holding cost over the finite time horizon considered (e.g., H periods into the future), while at the same time ensuring a prescribed high service level in order to meet random net demand. Given t is the current period, the order quantities Q_{t+j} for $j=1,2,\dots,H$ are to be determined. The problem formulation is:

$$\min \sum_{j=1}^H \left[c_{t+j} Q_{t+j} + S_{t+j} \delta(Q_{t+j}) + h_{t+j} E(I_{t+j}^+) \right] \quad (1)$$

$$\text{Subject to} \quad \Pr(I_{t+j} \geq 0) \geq \xi_{t+j} \quad ; j = 1, \dots, H \quad (2)$$

$$I_{t+j} = I_{t+j-1} + Q_{t+j} - x_{t+j} \quad (3)$$

$$Q_{t+j} \leq \delta(Q_{t+j}) M \quad (4)$$

and $Q_{t+j} \geq 0$ (5)

where the following notation has been used:

c_{t+j} is the unit purchasing cost at time period $t+j$; S_{t+j} is the fixed setup cost for placing an order at time period $t+j$; h_{t+j} is the average inventory holding cost represented as a fraction of the unit purchasing cost at time period $t+j$; $x_{t+j} = (d_{t+j} - m_{t+j})$ is the net demand during period $t+j$; I_{t+j} is the net inventory level (i.e., on-hand minus backorders) at the end of the period; $I_{t+j}^+ = \max(I_{t+j}, 0)$ is the on-hand inventory level; ξ_{t+j} is the prescribed service level in period $t+j$; $\delta(Q_{t+j})$ is an indicator variable for when an order is placed (i.e., when $Q_{t+j} > 0$); M is a large number which ensures that $\delta(Q_{t+j})$ is linked to the decision variable Q_{t+j} .

Let μ_{t+j} denote $E(x_{t+j})$. Kelle and Silver (1989b) showed that when the service level ξ_{t+j} is at least 0.85, as typically prescribed in practice, then $E(I_{t+j}) = I_{t+j-1} + Q_{t+j} - \mu_{t+j}$ is a good approximation of $E(I_{t+j}^+)$.

The service level constraint (2) can be represented as:

$$F_{t+j}(I_{t+j-1} + Q_{t+j}) \geq \xi_{t+j}$$

or

$$I_{t+j-1} + Q_{t+j} \geq F_{t+j}^{-1}(\xi_{t+j}) \tag{6}$$

where F_{t+j} is the cumulative distribution function (i.e., CDF) of net demand, x_{t+j} , at time $t+j$.

Our study departs from the approach found in Kelle and Silver (1989b) in that we allow the CDF of net demand (i.e., F_{t+j}) to be dynamically approximated by the CDF of the forecast of net

demand, $\hat{x}_{t+j} = (d_{t+j} - \tilde{m}_{t+j})$, for each period in the planning horizon (i.e., $j=1, \dots, H$ periods),

where the distribution for \tilde{m}_{t+j} is determined via the methods described later in section 3. The estimates $\hat{F}_{t+j}^{-1}(\xi_{t+j})$ and \hat{x}_{t+j} are then utilized in constraints (2) and (3) to obtain solutions to the optimization problem. A static estimate of the CDF of net demand (i.e., $\hat{F}_{t+j} = \hat{F}$) can be used instead, by fitting an empirical distribution to the net demand data (i.e., $x_i = d_i - m_i ; i = 1, \dots, t$) known at time t and using the resulting parameter estimates for all future forecasts. In section 4 we perform an analysis comparing our dynamic approach to the static empirical distribution approach.

In a rolling horizon procedure, only the decision related to the first period of the solution to the H -period problem above is implemented. At the next period, the inventory status is revised and the multi-period problem is updated as more information becomes available. The solution procedure is repeated again, with only the first order being launched.

3. Forecasting Multi-period Ahead Product Returns with the Distributed Lag Model

Let n_t and m_t denote the number of products sold and returned at time t , respectively; the general form of the distributed lag model is as follows:

$$m_t = p\beta_1 n_{t-1} + p\beta_2 n_{t-2} + \dots + p\beta_{t-1} n_1 + \varepsilon_t ; \text{ For } t= 2, 3, \dots, T. \quad (7)$$

(7) is known as the *finite* distributed lag model. The ε_t (error) terms are usually assumed to be additive white noise (i.e., normally distributed, independent of the n_t 's, independent of each other and have a constant variance given by σ^2). The parameter p is the probability that a sold product will ever be returned. Each β_k ($k=1, \dots, t-1$) term in (7) is the k th *reaction coefficient*, and it represents the proportion of n_{t-k} (i.e., sales in period $t-k$) that contributes units towards m_t (i.e. the returns in period t). T is a finite period and represents the maximum number of periods of data available for estimation. When historical sales and return data is used for estimation, there are usually many terms in (7) and little is known about the form of the lag. In

that case, direct estimation via Ordinary Least Squares uses up a large number of degrees of freedom and is likely to lead to imprecise parameter estimates because of multicollinearity. This issue can be avoided by assuming that the β_k coefficients are functionally related, and an appropriate method is used to estimate them. In the context of product returns, the functional relationship between the β_k coefficients is called the *delay function* and it represents the time for returns to be made. Approaches that have been used to represent the delay include the Geometric and Negative binomial (Toktay et al. 2000), Exponential (Clotey et al., 2012), Poisson (Krapp et al., 2013a), and Gamma (Clotey and Benton, 2014) distributions. Using the distributed lag model approach and given a forecasting horizon of H periods into the future, an estimate of products to be returned j ($=1,2,\dots,H$) periods into the future, \tilde{m}_{t+j} , is given by:

$$\tilde{m}_{t+j} = p\beta_1\tilde{n}_{t+j-1} + p\beta_2\tilde{n}_{t+j-2} + \dots + p\beta_j n_t + \dots + p\beta_{t+j-1}n_1 + \varepsilon_t, \quad (8)$$

where \tilde{n}_{t+j} is a potential “non-optimal” sales forecast j periods into the future. Methods used to estimate \tilde{n}_{t+j} are described later in this section. Note that we use the tilde sign (i.e., \sim) for forecasts when that forecast is directly based on a sales forecast, and we use the hat (i.e., $\hat{}$) sign otherwise. Pierce (1975) noted that forecasts of the dependent variable in a distributed lag model can also be obtained by representing the dependent variable as a univariate linear process (e.g., a function of only product returns and not previous sales). By representing the coefficients of (8) as a polynomial function of lags (i.e., the Almon model), Pierce (1975) was able to analytically show that the mean squared error (i.e., MSE) of (8) is never greater than that of the univariate process of only product returns, since the information available in the univariate process, at the forecast origin (i.e., period t) is a subset of the information available in (8). His results suggest that it is possible to obtain more accurate forecasts of product returns using such forecasts in (8), compared to forecasting product returns based on historical returns data only. However, those

analytical results were based on a delay function representation which resulted in a closed-form representation of the MSE for (8). In this study we use the recent Gamma representation of the delay function which was shown by Clotey and Benton (2014) to be superior to the existing discrete geometric and negative binomial delay functions in terms of accuracy and flexibility, but results in an MSE which cannot be expressed in an analytic form. With the gamma delay function the distributed lag model in (8) becomes:

$$\tilde{m}_{t+j} = p \frac{\lambda^\alpha}{\Gamma(\alpha)} \left[1^{(\alpha-1)} e^{-\lambda} \tilde{n}_{t+j-1} + \dots + j^{(\alpha-1)} e^{-j\lambda} n_t + \dots + (t+j-1)^{(\alpha-1)} e^{-(t+j-1)\lambda} n_1 \right] + \varepsilon_t. \quad (9)$$

The parameter λ is the average *delay rate* (i.e., the average number of lags per return period), α is the *shape* parameter which, along with λ , determines the lag with the largest coefficient. To facilitate the dynamic updating of the parameters, we adopt a Bayesian approach (Rossi, Allenby and McCulloch, 2005). We first specify Gamma, Inverted Gamma and Uniform(0,1) conjugate priors and then use a modified version of the Monte Carlo Markov Chain (MCMC) algorithm of Clotey and Benton (2014) to obtain the estimates. The use of such conjugate priors, which allow the coefficients of the distributed lag model to integrate to unity across time, results in a posterior distribution obtained via the MCMC algorithm which is identifiable (Gelfand and Sahu, 1999). The period forecast of the distribution for \tilde{m}_{t+j} (i.e., the posterior distribution for \tilde{m}_{t+j}) is then obtained by making the relevant substitutions for α , λ and p into the part of (9) which does not include the ε_t term. The CDF of net demand (i.e., F_{t+j}) can then be approximated by the CDF of the forecast of net demand, $\hat{x}_{t+j} = (D_{t+j} - \tilde{m}_{t+j})$, for each period in the planning horizon. The point estimate is based on the posterior median (i.e., \tilde{m}_{t+j}). The median is used since the resulting posterior distribution of product returns is typically skewed (Clotey and Benton, 2014).

Details of the MCMC procedure for estimating the CDF of the net demand (i.e., F_{t+j}) for period $t + j$ are provided below:

- i) Start with initial point estimates $\hat{p} = \hat{p}^{old}$, $\hat{\alpha} = \hat{\alpha}^{old}$, $\hat{\lambda} = \hat{\lambda}^{old}$ and $\hat{\sigma}^2 = \hat{\sigma}_{old}^2$.
- ii) Generate: $\hat{\alpha}^{new} = \hat{\alpha}^{old} + \xi_1$; $\hat{\lambda}^{new} = \hat{\lambda}^{old} + \xi_2$; $\xi_i \sim N(0, step^2)$; $i = 1, 2$; $step$ is a numerical value chosen to enable the algorithm to have sufficiently navigated the space where the posterior has high mass.
- iii) Compute $\gamma_1 = \min \left\{ 1, \frac{\ell(\hat{\alpha}^{new}, \hat{\lambda}^{new}, \hat{p}, \hat{\sigma}^2) \pi(\hat{\alpha}^{new}, \hat{\lambda}^{new})}{\ell(\hat{\alpha}^{old}, \hat{\lambda}^{old}, \hat{p}, \hat{\sigma}^2) \pi(\hat{\alpha}^{old}, \hat{\lambda}^{old})} \right\}$; where $\pi(\cdot)$ is the prior for (α, λ) ; $\ell(\cdot)$ is the likelihood function.
- iv) With probability γ_1 , $\hat{\lambda} = \hat{\lambda}^{new}$ and $\hat{\alpha} = \hat{\alpha}^{new}$; else $\hat{\lambda} = \hat{\lambda}^{old}$ and $\hat{\alpha} = \hat{\alpha}^{old}$.
- v) Generate: $\hat{p}^{new} = \hat{p}^{old} + \xi$; $\xi \sim N(0, step_1^2)$; $step_1$ is chosen in a similar way as $step$ in ii).
- vi) Compute $\gamma_2 = \min \left\{ 1, \frac{\ell(\hat{\alpha}, \hat{\lambda}, \hat{p}^{new}, \hat{\sigma}^2) \pi(\hat{p}^{new})}{\ell(\hat{\alpha}, \hat{\lambda}, \hat{p}^{old}, \hat{\sigma}^2) \pi(\hat{p}^{old})} \right\}$; where $\pi(\cdot)$ is the prior for p .
- vii) With probability γ_2 , $\hat{p} = \hat{p}^{new}$; else $\hat{p} = \hat{p}^{old}$.
- viii) Generate: $\hat{\sigma}_{new}^2 | \mathbf{m}, \mathbf{n}, (\hat{\alpha}, \hat{\lambda}), \hat{p} \sim \frac{v_1 s_1^2}{\chi_{v_1}^2}$; with $v_1 = v_0 + (T - 1)$, $s_1^2 = \frac{v_0 s_0^2 + (T-1)s^2}{v_0 + (T-1)}$; where s_0^2 and v_0 are as defined in Clotey et al. (2012; p600).
- ix) Repeat steps ii-viii, 10,000 times [Rossi et al. (2005); Clotey et al. (2012)].
- x) Obtain \tilde{n}_{t+j} estimates for periods t to $t + j - 1$.
- xi) To obtain the empirical CDF of \tilde{m}_{t+k} , substitute the $\hat{p}, \hat{\alpha}, \hat{\lambda}$ estimates obtained from steps i-ix, along with the estimates obtained in step x, into the part of (9) which does not include the ε_t term.
- xii) The estimate of F_{t+j} is the resulting CDF for $(d_{t+j} - \tilde{m}_{t+j})$.

A key issue for managers employing the methods described thus far to make purchasing decisions is whether the accuracy in the non-optimal sales forecast (\tilde{n}_{t+j}) has a major effect on the performance of the proposed method for making acquisition plans. To investigate this issue we designed a study based on the framework, for evaluating the effect of forecast error in supply

chain planning, proposed by Fildes and Kingsman (2011).

3.1. Evaluating the effect of sales forecast accuracy on the distributed lag model

When there is monthly sales data with no outliers and non-seasonality, models of the form below have been shown (e.g., Makridakis and Hibon, 2000; Fildes and Kingsman, 2011) to adequately characterize situations commonly encountered in practice:

$$n_t = \delta + \rho n_{t-1} + e_t - \theta e_{t-1} \quad |\rho| \leq 1, |\theta| \leq 1, \quad (10)$$

where e_t is random white noise and δ is historical average sales. The equation in (10) can be used to represent a variety of sales situations such as sales as random noise (i.e., θ and ρ are both zero), sales as an ARIMA (0,1,1) process (i.e., $\rho = 1$ and $\theta < 1$), or sales as an AR[1] process (i.e., $\rho < 1$ and $\theta = 0$). The error variance for the k -step ahead forecast for an ARIMA (0,1,1) process remains constant with increasing k , which is similar to that of the random noise process, therefore the ARIMA(0,1,1) representation of sales was omitted from our analysis in favor of the random noise process. Fildes and Kingsman (2011) considered two cases in their study: the random noise scenario and the AR[1] process with $\rho = 0.9$, which is an appropriate series for when exponential smoothing is near-optimal. These two cases are considered to be representative extremes for illustrating the effect of forecast error (Fildes and Kingsman, 2011). The proposed methodology and insights from this paper should carry over to the analysis of other specifications of sales based on (10). The one-step ahead sales forecast for period t , based on (10) is given by:

$$\hat{n}_t = \delta + \rho n_{t-1} - \theta e_{t-1}, \quad (11)$$

where e_{t-1} is the observed forecast error in the previous period which is known as at time t . The white noise terms are assumed normal although the sales and the forecasts of sales are constrained to be positive. The k steps ahead forecast are given by:

$$\hat{n}_{t+k} = \delta + \rho \hat{n}_{t+k-1} - \theta e_t(k-1), \quad (12)$$

where $e_t(k-1)$ is the forecast error for future period $t+k-1$ estimated at time t . The general form for the forecast error variance $e_t(k)$ for an AR[1] model, k steps ahead is:

$$\text{var } e_t(k) = \frac{(1-\rho^{2k})}{1-\rho^2} \sigma^2,$$

where σ is the standard deviation for e_t at time t , which is also the standard deviation for the k steps ahead random noise sales model. Thus, the standard deviation for the AR[1] model increases with k while that of the random noise model stays constant. The forecast in (12) may be non-optimal due to a sub-optimal forecasting method being used to generate forecasts from (10). Thus the actual k -steps ahead sales forecast is given by:

$$\tilde{n}_{t+k} = \hat{n}_{t+k} + v_t \quad (13)$$

so that the actual forecast (i.e., \tilde{n}_{t+k}) equals the optimal forecast (i.e., \hat{n}_{t+k}) plus white noise (i.e., v_t). Fildes and Kingsman (2011) used an error standard deviation for v_t specified as $\kappa\sigma$, where κ was chosen to be 0%, 20% and 50% of σ . They also used values of σ which would result in coefficient of variations (CVs) ranging from 0 to 0.4, noting that these were plausible since surveys of forecasting accuracy for fast moving items, reported in the literature, suggested a CV of 0.25 as the norm (Fildes and Kingsman, 2011). The use of (13), with varying levels of σ and κ ensures that the effects of both sales (i.e., σ) and forecast error volatility (i.e., $\kappa\sigma$), can be evaluated on the performance of the proposed method for determining the purchasing policy for cores in a rolling horizon. In the next section we perform a numerical study that allows us to examine the cost consequences of acquiring cores using multi-step ahead forecasts in a rolling horizon. The analysis is performed at specified service levels over a variety of forecasting, sales

and demand situations. We also compare the accuracy of the distributed lag model approach to that of the Holt method which is a widely used forecasting approach.

4. Numerical Analysis

4.1. Parameter values

The fixed parameters for the numerical analysis are shown in Table 1.

To generate product returns, the value of the parameters for p , λ , and α in Table 1, were used in (9) with a random value drawn from the standard normal distribution being used for the error terms (i.e., ε_t) in (9). One hundred and twenty periods (e.g., 120 months worth) of returns data were generated with this approach. Demand data for these cores were generated by drawing from a uniform distribution in the interval zero and 3 times the average product returns during the 120 periods (i.e., this ensured that demands were on average 50% larger than the supply of product returns, thereby requiring the purchase of additional cores to make up any shortages). Sales data were generated from (10) either as a random noise [RN] series (i.e., $\theta = \rho = 0$) or as an autoregressive [AR] series (i.e., $\theta = 0; \rho = 0.9$), as previously explained. Average sales (i.e., δ) was set at 2,000 units with the standard deviation (i.e., σ) based on the CV levels (0, 0.1, 0.2, 0.3 and 0.4). Forecasts for sales were performed using (13) with $e_t(k)$ drawn from a normal distribution with zero mean and variance $\text{var } e_t(k)$, when sales were AR, and σ^2 when sales were RN. Since $\text{var } e_t(k) \leq \sigma^2$, in order to facilitate comparisons between AR and RN generation processes we chose different values of σ^2 , for generation using the AR process compared to that of the RN process, such that both generation processes had equal CV levels. Values for v_t were drawn from a normal distribution with zero mean and standard deviation $\kappa\sigma$, with values of κ chosen per the description in section 3. Service levels were set at 85%, 90%, 95% and 99% respectively. The CDF of net demand (i.e., F_{t+j}) was estimated using a dynamic

approach (i.e., using the Bayesian estimation approach described in section 3 to update the parameters each period) and static approach. The distributed lag model has an asymptotic normal distribution (Pindyck and Rubinfeld, 1998), so we used the `fitdistr()` function in the **R** software package (which uses maximum likelihood) to obtain parameter estimates for the static approach. Details about the `fitdistr()` function can be found in Ricci(2005). Table 2 provides a summary of the effects investigated in the analysis.

4.2. Implementation

For each of the forecasting approaches (i.e, RN or AR) twelve periods of sales, with the corresponding eleven periods of product returns data (e.g., each period representing a month), was used as the initial estimation data set. The system starts out with the products returned by the end of period 12 as the initial inventory. Purchase plans for three periods ahead, are made starting from period 13 till period 117 (i.e., 106 periods of forecasting were considered). At the beginning of each period the following sequence of activities occur using the data to date: (1) estimate the distribution of product returns as at the end of the previous period, (2) forecast the sales for the current period and two additional periods ahead using either the RN or AR approach, (3) use the sales forecasts to estimate the future net demand distributions as described in sections 2 and 3, (4) estimate the quantity of cores that should be acquired for the current period and the next two periods ahead using the approach described in section 2, and (5) acquire only the estimated quantity for the current period and update all inventory positions. Five replications, at each combination of levels, were carried out and detailed examination analysis of these runs showed that there was very little difference between the outcomes of each of the replications. The costs recorded for each combination of levels was the average total cost per unit over the five replications. Common fixed random number seeds were employed to reduce

variance (Law and Kelton, 2000). The analysis was created and performed in **R** version 3.2.0. Cplex was used for the optimization described in section 2, via the Rglpk library package in **R**. The demand and forecast generation processes, and the rolling schedule calculations were verified against manual calculations.

4.3. Results

We first analyze the results concerning the effect of sales forecast accuracy on the cost performance of the rolling horizon acquisition plan. We then analyze the results of using a dynamic versus static approach to estimating the net demand distribution for rolling horizon acquisition decisions. Relative performances are based on the “regret” of using one approach versus the other (e.g., sales forecast vs. actual sales; dynamic vs. static estimation of net demand) in the rolling horizon environment.

4.3.1. Effects of sales forecast accuracy

With dynamic parameter estimation, we found that the forecast error variance (i.e., EV) did not have any effect at all levels of the other factors. A possible reason for this unexpected result is that the dynamic updating of parameters via the Bayesian approach is more affected by high variability in historical sales (i.e., there are more historical sales data points used for estimation), which leads to less accurate returns forecast, than by the forecast error variance, which only affects the accuracy of the few k -step ahead forecasts used for net demand estimation. That is to say, that the effect of the sales forecast error variance (i.e., EV) could not be separated from the effect of sales volatility (i.e., CV), when forecasting using the Bayesian estimation approach described in section 3, therefore it was omitted from subsequent analysis.

Table 3 shows the relative costs of using sales forecasts in (9), to make the rolling horizon decision, versus using the actual sales amount when the CDF of net demand was estimated dynamically.

It can be seen that higher coefficients of variation, resulting in less accuracy in the sales forecasts from Equation (13), generally lead to higher costs when the sales forecasts were used as compared to using the actual sales amounts. These cost differences become less consequential at higher service levels. When sales have the autoregressive (AR) pattern, the percent regret in not using actual sales to make the rolling horizon acquisition is at most 1.3%, and therefore the accuracy of the sales forecast is unimportant in this situation. In contrast, sales which appear to have a random pattern (i.e., RN) result in significant cost differentials when sales forecasts are used instead of actual sales. Thus, the accuracy of the sales forecast is important in that situation.

A possible explanation of the above result is that updating the distributed lag model parameters via Bayesian estimation (i.e., dynamic estimation) is more consistent when the sales follow a particular pattern (which in this case was an autoregressive series). However, when the pattern of the data is random then dynamic estimation is likely to perform poorly with increasing volatility in the sales data (e.g., higher CV levels). Thus, while forecast accuracy is unimportant with the autoregressive model of sales, it resulted in significantly higher costs when sales were presented as a random noise series. The difficulties of Bayesian updating when a univariate series is random and highly volatile has been noted in previous studies (e.g., Fildes, 1986). Our results suggest that this also carries over to multivariate series (e.g., distributed lag models).

The same analysis was performed (not shown) for when the CDF of net demand was estimated via a static approach. The percent regret, in those results, were all found to be insignificant (i.e., sales forecast accuracy had no effect for all given service levels, when net

demand was estimated via the static approach). Because there is no updating with the static approach (and only a few periods of sales forecast were combined with actual sales in previous periods to estimate the parameters of the distribution) the forecasts had little to no effect on the resulting estimate of the distribution. Thus the inconsequential cost differences with the static approach to estimation, in this case.

4.3.2. Effects of estimating the net demand distribution dynamically

Table 4 shows the relative costs of using sales forecasts when dynamically estimating the net demand distribution versus static estimation of the net demand distribution.

It can be seen that higher coefficients of variation result in lower cost differences when the static approach to estimation of net demand is compared to the dynamic approach over all service level, sales model, and sales forecasting combinations. The dynamic approach resulted in a cost reduction at a high service level (e.g., 99%) for both autoregressive (AR) and random (RN) representations of sales. The cost reductions were higher for the autoregressive representation for sales, with inconsequential cost differences being experienced at the 95% service level in this case. When sales had a random representation then lower costs were experienced with the dynamic approach, compared to the static approach, at a 99% service level when actual sales (i.e., no sales forecasts) were used in the estimation. When forecasted sales were used in the same situation, lower costs were only experienced at the 99% service level when the sales forecast accuracy was high (i.e., corresponding to a sales CV of 0 or 0.1). At lower service levels (i.e., 85% and 90%), the dynamic approach resulted in higher costs than the (simpler) static approach to estimating the CDF of net demand.

A possible explanation of the relative performance of dynamic versus static estimation, in this case, is the nature of the net demand distribution obtained with dynamic estimation. Analysis

of the plots of the net demand distributions, obtained via the dynamic approach, showed that they often resulted in longer tails (i.e., skewness) than the distribution obtained with the static approach. Skewed posterior distributions for Bayesian estimation of product returns has been observed by Clotey and Benton (2014), and our results are consistent with this observation. The static approach used the large sample approximation (described earlier) to estimate the net demand, which results in a normally distributed net demand distribution. At high service levels (e.g., 99%), the skewness resulted in significantly lower service level units (i.e., $F_{t+j}^{-1}(\xi_{t+j})$) in each period than the service level units from the static approach (i.e., $F^{-1}(\xi_{t+j})$), and this resulted in lower costs. The lower service level units could also lead to an increase in the stockout risk; however stockouts were negligible for both the autoregressive and random noise representation of sales, although the latter representation resulted in a marginally greater number of backlogged units. The skewness had a negative effect at lower service levels, with dynamic estimation resulting in worse cost performance with the desired service levels set at 85% and 90%, respectively. The performance differences become inconsequential at all service levels as sales became more volatile (i.e., higher CVs) resulting in less accurate estimates of the net demand distribution. This was observed irrespective of the approach used to estimate the net demand distribution.

4.4. Accuracy comparison with existing methods

We compared our approach to the Holt's forecasting method, which is a widely applied forecasting procedure. To demonstrate the capability for predicting multi-step ahead product returns, we conducted three-step ahead forecasts using both methods. Our basis for comparison is the Mean Absolute Scaled Error (MASE) measure proposed by Hyndman and Koehler (2006). The MASE can be used for commonly occurring situations with data, namely data which

contains zeros and/or has occurrences of consecutive observations which take the same value. Such situations can lead previously recommended measures of forecast accuracy such as the Mean Absolute Percentage Error (MAPE) to result in undefined values. In addition, the MASE does not suffer from issues of scale dependence which plague other commonly used measures (e.g., Mean Square Error-MSE and Mean Absolute Deviation-MAD). The MASE scales the error of each forecast by the *in-sample* MAD of the naïve forecast method (i.e., the forecast based on the most recent observation). For the 3 step ahead naïve forecast, the in-sample MAD is calculated on a rolling basis. Thus, given an initialization sample of k periods and a total number of observations T for the data, the MAD of the three period ahead naïve forecast method is obtained by:

1. Selecting the observation at time $k+i-1$ as the forecast estimate for period $k+3+i-1$. The 3 step ahead error on the forecast for time $k+3+i-1$ is then computed.
2. Step 1 above is then repeated for $i=1,2,\dots,T-k-3+1$.
3. The in-sample MAD is then computed based on the 3 step ahead errors obtained as at period $k+i-1$.

Three step ahead forecasts, using both the dynamic distributed lag model approach and the Holt's method, were made in a rolling horizon manner (i.e., for a given i , data for periods $1,\dots,k+i-1$ were used to make the forecast for period $k+3+i-1$; the next forecast would use data for periods $1,\dots,k+i$, to make the next three period ahead forecast, and so on. Errors for each of the forecasts were calculated based on the deviation of the forecast from its corresponding out of sample value). We compare the out-of-sample performance (based on forecasting the data in the rolling three period hold-out sample using only information from the fitting period). The MASE is less than one if it arises from a better forecast than the average three period ahead naïve forecast computed in-sample. Conversely, it is greater than one if the forecast is worse than the average three period ahead naïve forecast computed in-sample. A portion of the 120 periods of

data generated in Section 4.2, were used in the analysis. The initialization sample was from period 1 to 12 (i.e., $k=12$). We evaluated the MASE from period 13 to period 36. A similar size dataset was used in Hyndman and Koehler (2006). Consistent with each approach, the Holt's method was only applied to the product returns data, while the distributed lag model approach was applied to data consisting of both the returns and sales in each period. Sales forecasts were used to make the three period ahead forecast estimates with the distributed lag model approach. Sales (historical and forecasts) were generated as RN or AR processes at the five CV levels described previously in Section 4.1.

Table 5 shows that the MASE of our proposed dynamic estimation approach is superior to the Holt's method when sales have the autoregressive (AR) pattern, at all CV levels. Conversely, the Holt's method had a MASE superior to that of our proposed method when sales were generated as random noise (i.e., RN). It can be seen that higher coefficients of variation, resulting in less accuracy in the sales forecast from Equation (13), yielded higher MASE values when sales followed an AR pattern. However the effect of CV levels on MASE values when sales followed an RN pattern, was not as conclusive (i.e., MASE values for CV levels of 0.2 and 0.3 were less than those at CV levels of 0, 0.1 and 0.4).

Collectively, these results support our earlier observation that our distributed lag approach to estimating the net demand on a rolling horizon basis may be best suited for use when sales follow an autoregressive pattern.

5. Discussion and Conclusions

In the last decade distributed lag models have become increasingly attractive for use in forecasting the returned products which are the basis of various reuse operations (e.g., Toktay et al., 2000; Toktay, 2004; Clotney et al. 2012; Clotney and Benton, 2014). Previous applications of

the model to forecasting product returns for inventory management or production planning purposes have been mainly limited to single period forecasts based on known sales. In this research we have developed an approach to making multi-period ahead forecasts via Bayesian estimation of the distributed lag model which dynamically approximates the net demand distribution for a finite horizon acquisition plan. The methodological issues we explored were aimed at evaluating: 1) the effect of sales forecast accuracy, and 2) the use of the dynamic estimation (i.e., Bayesian) approach versus a static (i.e., MLE) approach to estimating the net demand distribution, on the cost performance of a rolling horizon purchasing policy for cores.

When making multi-period ahead forecasts, using the distributed lag model, estimates of future sales are required; this means that the accuracy of those estimates can affect the cost performance of such forecasts. Our results indicate that the accuracy of the sales forecast is inconsequential if sales follow an autoregressive pattern, since then the Bayesian estimation approach is able to exploit the autoregressive nature of the sales to correctly update parameters as more sales information is obtained from period to period. Thus, the net demand estimates obtained using the sales forecast were comparable to those obtained using actual sales in that situation. This suggests that sales data which is consistent with an autoregressive series is the type of scenario where Bayesian estimation of the net demand distribution is *least* likely to be affected by the accuracy in the sales forecast. The use of autoregressive models to estimate sales is common (Pindyck and Rubinfeld, 1998). Monthly sales which are random, result in large cost discrepancies if using sales forecasts versus actual sales in the dynamic estimation of net demand. These cost discrepancies became large with highly volatile (i.e., sales with a high coefficient of variation) random sales. Thus, forecast accuracy has a significant negative impact

on costs if using the dynamic (i.e., Bayesian) estimation approach when sales are random and highly volatile.

While the inaccuracy of the sales forecast may be due to both volatility in sales and forecast error volatility (Fildes and Kingsman, 2011), we found that the effect of forecast error variance could not be separated from the effect of sales volatility when the parameters in our model were dynamically updated using the Bayesian approach. That is to say that forecast error variance had no effect on the cost performance while sales volatility did, with the dynamic approach. In future research, it may be possible to use a ‘prior’ to model the sales forecast error variance explicitly in the Bayesian framework, in which case the effect of sales forecast error variance may be measured and separated from sales volatility when estimating the net demand distribution.

The cost benefit of using the Bayesian estimation of net demand versus the static (asymptotic) approach, only occurred at high service levels (i.e., >95%) . This suggests that companies wanting to employ the proposed Bayesian approach to estimate the net demand, for use in determining the purchase quantities for cores on a rolling horizon basis, can best leverage the approach if high service levels are sought. The static approach resulted in lower costs than the Bayesian approach when service levels were in the 0.85 to 0.95 range. These lower service levels can be more practical for the cost structures of some companies.

A number of research questions remain, in particular whether the results generalize to other (more complex) closed-loop supply chains via the consideration of quality uncertainty, along with the timing and quantity uncertainty modeled with the distributed lag model, as well as other approaches to representing the net demand distribution. In addition, future research can evaluate the cost performance of various methods, including that proposed in this paper, for

determining the purchase lot-size of cores. A number of authors (e.g., Silver, 1976; Ferrer and Whybark, 2001; Galbreth and Blackburn, 2010) have proposed modifications/extensions to existing lot-sizing models to account for some or all of the quantity, quality and timing uncertainties inherent in the sourcing of cores. A future comparison of purchase lot-sizing policies for cores could yield fruitful managerial insights. Another avenue for future research is the investigation of the effects of both sales forecast model misspecification and sales forecast error when estimating the net demand distribution. Empirical analysis is also required for the types of returns generating processes observed in various reuse operations (e.g., recycling, remanufacturing, refurbishing of commercial returns, etc...), which may require models that can capture more complexity than that of the distributed lag model. A key question remaining is what improvements in the forecasting accuracy of net demand are realistically achievable. The results of this paper suggests that there are significant benefits that can result if effort is put towards achieving such improvements.

References

- Clottey, T., Benton, W.C., and Srivastava, R.S., 2012. Forecasting product returns for remanufacturing operations. *Decision Sciences*, 43(3): 589-614.
- Clottey, T. and Benton, W.C., 2014. Determining core acquisition quantities when products have long return lags. *IIE Transactions*, 46(9): 880-893.
- Ferrer, G. and Whybark, C.D., 2001. Material planning for a remanufacturing facility. *Production and Operations Management*, 10(2): 112-124.
- Fildes, R., 1986. An evaluation of Bayesian forecasting. *Journal of Forecasting*, 2: 137-150.
- Fildes, R. and Kingsman, B., 2011. Incorporating demand uncertainty and forecast error in supply chain planning models. *Journal of the Operational Research Society*, 62: 483-500.
- Galbreth, M.R. and Blackburn, J.D., 2010. Optimal acquisition quantities in remanufacturing with condition uncertainty. *Production and Operations Management*, 19(1): 61-69.
- Gelfand, A.E. and Sahu, S.K. 1999. Identifiability, improper priors and gibbs sampling for generalized linear models. *Journal of the American Statistical Association* , 94 (445): 247-253.
- Guide, Jr., V.D.R., and van Wassenhove, L.N. 2001. Managing product returns for remanufacturing. *Production and Operations Management* 10(2): 142-155.
- Hanafi, J., Kara, S., and Kaebnik, H. 2008. Reverse logistics strategies for end-of-life products. *International Journal of Logistics Management*, 19(3): 367-388.
- Hyndman, R.J. and Koehler, A.B. 2006. Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22 (1): 679-688.
- Kelle, P. and Silver, E.A., 1989a. Forecasting the Returns of Reusable Containers. *Journal of Operations Management*, 8(1): 17-35.

- Kelle, P. and Silver, E.A., 1989b. Purchasing policy of new containers considering the random returns of previously issued containers. *IIE Transactions*, 21(4): 349-354.
- Kiesmüller, G.P. and van der Laan, E.A. 2001. An inventory model with dependent product demands and returns. *International Journal of Production Economics*, 72(1): 73-87.
- Krapp, M., Nebel, J., and Sahamie, R., 2013a. Forecasting product returns in closed-loop supply chains. *International Journal of Physical Distribution and Logistics Management*. 43(8): 614-637.
- Krapp, M., Neble, J., and Sahamie, R. 2013b. Using forecasts and managerial accounting information to enhance closed-loop supply chain management. *OR Spectrum*, 35 (4): 975-1007.
- Law, A.M. and Kelton, D.W., 2000. *Simulation modeling and analysis (3rd ed.)*. New York, NY: McGraw Hill.
- Makridakis, S. and Hibon, M., 2000. The M-3 competition: results, conclusions and implications. *International Journal of Forecasting*, 16: 451-476.
- Pierce, D.K., 1975. Forecasting in dynamic models with stochastic regressors. *Journal of Econometrics*, 3(1): 349-374.
- Pindyck, R.S., and Rubinfeld, D.L., 1998. *Econometric Models and Economic Forecasts*, 4th ed., McGraw-Hill, New York, NY.
- Ricci, V., 2005. Fitting distributions with R
Available at <http://cran.r-project.org/doc/contrib/Ricci-distributions-en.pdf>.
- Rossi, P., Allenby, G., and McCulloch, R., 2005. *Bayesian Statistics and Marketing*, 3rd ed., Wiley, New York, NY.
- Silver, E. A., 1976. Establishing the reorder quantity when the amount received is uncertain. *INFORMS*, 14(1): 32-39.
- Souza, G.C. 2013. Closed-loop supply chains: A critical review, and future research. *Decision Sciences*, 44(1): 7-38.

Srivastava, S.K., and Srivastava, R.K. 2006. Managing product returns for reverse logistics. *International Journal of Physical Distribution and Logistics Management*, 36 (7): 524-546.

Toktay, B., Wein, L., and Zenios, S., 2000. Inventory management of remanufacturable products. *Management Science*, 46(4): 1412-1426.

Toktay, B., 2004. Forecasting product returns, in *Business Aspects of Closed-Loop Supply Chains*, Guide, Jr., V.D. and van Wassenhove, L.N. (eds) Carnegie Mellon University Press, Pittsburgh, PA, 203-209.

Table 1: Parameters for the numerical example

p	λ	α	ξ	$c_{t+j} = c$	$S_{t+j} = S$	$h_{t+j} = h$	H
0.2	0.0198	2.14	0.90	385	173	57.75	3

Table 2: Investigated effects

Name	Label	Number of Levels	Values
Coefficient of variation (sales)	CV	5	0.0,0.1,0.2,0.3,0.4
Sales model	SM	2	RN (Random noise), AR (Autoregressive)
Forecast error variance (sales)	EV	3	$\kappa = 0\%, 20\%, 50\%$
Estimate of net demand CDF (i.e., \hat{F}_{t+j}, \hat{F})	ND	2	Dynamic, Static
Service level	SL	4	85%, 90%, 95%,99%

Table 3: Percentage cost comparisons of sales forecast vs. actual sales¹ with dynamic estimation of the net demand CDF

<i>Service level</i>	<i>Sales model (SM)</i>									
	<i>Autoregressive (AR)</i>					<i>Random noise (RN)</i>				
	<i>Coefficient of variation (CV)</i>					<i>Coefficient of variation (CV)</i>				
	0	0.1	0.2	0.3	0.4	0	0.1	0.2	0.3	0.4
85%	0.0	0.8	0.9	1.2	1.3	0.0	4.1 [†]	6.2 [†]	6.6 [†]	8.2 [*]
90%	0.0	0.6	0.7	1.1	1.1	0.0	3.8 [†]	5.9 [†]	6.3 [†]	7.7 [*]
95%	0.0	0.5	0.6	0.7	0.8	0.0	3.5	5.6 [†]	5.8 [†]	7.2 [*]
99%	0.0	0.1	0.4	0.5	0.7	0.0	3.2	5.2 [†]	5.6 [†]	6.6 [†]

†Significant at less than 0.05; *Significant at less than 0.01

Note: ¹-Actual sales were generated using (13)(10), Forecasted sales were generated using (10)(13)

Table 4: Percentage cost comparisons of static vs. dynamic estimation of the net demand CDF

Sales model	Service level	Forecasted sales ¹					Actual sales ²				
		Coefficient of variation (CV)					Coefficient of variation (CV)				
		0	0.1	0.2	0.3	0.4	0	0.1	0.2	0.3	0.4
Autoregressive (AR)	85%	10.0*	9.6*	8.8*	8.4*	8.2*	10.0*	9.8*	9.1*	8.6*	8.3*
	90%	7.3*	6.6 [†]	6.4 [†]	5.9 [†]	5.7 [†]	7.3*	6.9*	6.7 [†]	6.1 [†]	5.5 [†]
	95%	2.6	1.6	1.5	1.3	0.5	2.6	2.2	1.5	1.4	0.7
	99%	-8.4*	-8.2*	-8.0*	-7.7*	-6.6 [†]	-8.4*	-7.9*	-7.8*	-6.9*	-6.6 [†]
Random (RN)	85%	16.5*	16.3*	13.1*	11.6*	10.1*	16.5*	20.3*	19.0*	18.7*	18.5*
	90%	11.7*	11.0*	9.0*	8.5*	6.8*	11.7*	16.3*	15.2*	14.8*	14.6*
	95%	5.5 [†]	4.6 [†]	4.2 [†]	3.8 [†]	1.4	5.5 [†]	10.6*	9.7*	8.5*	8.3*
	99%	-7.6 [†]	-7.3 [†]	-7.0 [†]	-6.7 [†]	-5.9 [†]	-7.6 [†]	-6.7 [†]	-3.5	-2.4	-1.1

[†]Significant at less than 0.05; *Significant at less than 0.01

Notes: ¹-Sales generated using (10)(13); ²-Sales generated using (13)(10)

Table 5: Mean Absolute Scaled Error (MASE)¹ comparisons of dynamic distributed lag model estimation versus Holt's method

<i>Estimation method</i>	<i>Sales model (SM)</i>									
	<i>Autoregressive (AR)</i>					<i>Random noise (RN)</i>				
	<i>Coefficient of variation (CV)</i>					<i>Coefficient of variation (CV)</i>				
	0	0.1	0.2	0.3	0.4	0	0.1	0.2	0.3	0.4
Holt's	0.37	0.47	0.60	0.71	0.83	0.61	0.67	0.49	0.52	0.75
Distr. Lag Model	0.14	0.15	0.15	0.16	0.17	0.83	0.87	0.77	0.78	0.99

Note: ¹-MASE<1 indicates that the proposed forecasting method gives on average smaller errors than the 3 step ahead errors from the naïve forecast. The smaller the MASE the better is the forecasting method (Hyndman and Koehler, 2006).