

8-2012

Forecasting Product Returns for Remanufacturing Operations

Toyin Clotley

Iowa State University, tclotley@iastate.edu

W. C. Benton Jr.

Ohio State University

Rajesh Srivastava

Florida Gulf Coast University

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



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

The complete bibliographic information for this item can be found at https://lib.dr.iastate.edu/scm_pubs/81. 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.

Forecasting Product Returns for Remanufacturing Operations

Abstract

Driven by legislative pressures, an increasing number of manufacturing companies have been implementing comprehensive recycling and remanufacturing programs. The accurate forecasting of product returns is important for procurement decisions, production planning, and inventory and disposal management in such remanufacturing operations. In this study, we consider a manufacturer that also acts as a remanufacturer, and develop a generalized forecasting approach to determine the distribution of the returns of used products, as well as integrate it with an inventory model to enable production planning and control. We compare our forecasting approach to previous models and show that our approach is more consistent with continuous time, provides accurate estimates when the return lags are exponential in nature, and results in fewer units being held in inventory on average. The analysis revealed that these gains in accuracy resulted in the most cost savings when demand volumes for remanufactured products were high compared to the volume of returned products. Such situations require the frequent acquisition of cores to meet demand. The results show that significant cost savings can be achieved by using the proposed approach for sourcing product returns.

Disciplines

Business Administration, Management, and Operations | Management Information Systems | Operations and Supply Chain Management | Organizational Behavior and Theory

Comments

This accepted article is published as Clottey, T., Benton Jr., W.C., Srivastava, R., Forecasting Product Returns for Remanufacturing Operations. *Decision Sciences*. 2012 43(4);589-614. doi: [10.1111/j.1540-5915.2012.00362.x](https://doi.org/10.1111/j.1540-5915.2012.00362.x). Posted with permission.

Forecasting Product Returns for Remanufacturing Operations

Toyin Clotey

Department of Supply Chain and Information Systems, College of Business, Iowa State University, 3258 Gerding Business Building, Ames, Iowa 50011, e-mail: tclotey@iastate.edu

W. C. Benton, Jr. †

Department of Management Sciences, Fisher College of Business, The Ohio State University, 2100 Neil Avenue, Columbus, OH 43210, e-mail: benton_1@fisher.osu.edu

Rajesh Srivastava

Department of ISOM, Lutgert College of Business, Florida Gulf Coast University, Ft. Myers, FL 33965, e-mail: rsrivast@fgcu.edu

†Corresponding author.

ABSTRACT

Driven by legislative pressures an increasing number of manufacturing companies have been implementing comprehensive recycling and remanufacturing programs. The accurate forecasting of product returns is important for procurement decisions, production planning, and inventory and disposal management in such remanufacturing operations. In this study, we consider a manufacturer that also acts as a remanufacturer, and develop a generalized forecasting approach to determine the distribution of the returns of used products, as well as integrate it with an inventory model to enable production planning and control. We compare our forecasting approach to previous models and show that our approach is more consistent with continuous time, provides accurate estimates when the return lags are exponential in nature and results in fewer units being held in inventory on average. The analysis revealed that these gains in accuracy resulted in the most cost savings when demand volumes for remanufactured products were high compared to the volume of returned products. Such situations require the frequent acquisition of cores to meet demand. The results show that significant cost savings can be achieved by using the proposed approach for sourcing product returns. [Submitted: July 23, 2011. Revised: October 3, 2011; November 16, 2011; Accepted: November 29, 2011.]

Key words: Closed-Loop Supply Chains, Inventory Management, Operations Forecasting, Remanufacturing

INTRODUCTION

Remanufacturing is the process by which products are recovered, processed, and sold as like-new products in the same or separate markets. In a survey reported in Hauser and Lund (2003), remanufacturing operations accounted for total sales in excess of \$53 billion per year. The US Environmental Protection Agency (EPA) cites remanufacturing as an integral foundation of reuse activities and reports that less energy is used and fewer wastes are produced with these types of activities (EPA, 1998). Increasing legislation is one of the primary forces driving the remanufacturing revolution. At the same time many firms have realized that being “green” can actually be profitable. In the past decade a number of companies have implemented comprehensive remanufacturing programs. However, few guidelines are available to the practicing manager to aid in planning, controlling, and managing remanufacturing operations (Souza, 2008).

Production planning and control activities are more complex and difficult for remanufacturers partly because of the timing uncertainty of returned products. The returns that are used for remanufacturing are products that are sold to the customer and are returned when their useful life is over or when the customer wants to trade in the product for an upgrade or another unit of the product. Forecasting the proportions of product returns is important for procurement decisions, capacity planning, and disposal management. From an operational standpoint, accurate forecasting of the returned product quantities (in each period) is useful for production planning and inventory management.

Forecasting of product returns has been studied by Toktay, Wein and Zenios (2000) for disposable cameras, Kelle and Silver (1987) for reusable containers, and Goh and Varaprasad

(1986) for returnable bottles. In each of these environments, typically, only the sales and return volume in each period are known and the timing of returns and sales are unknown. Toktay (2004) defines this type of data as *period level information* and defines data in which the timing of sales and returns are tracked on an individual basis as *item-level information*. One method for forecasting returns with period level data is to use a time series consisting of past return volumes to forecast future return volumes; however such a method would ignore past sales data and therefore result in inaccurate results. The key to forecasting returns is the observation that returns in any one period are generated by sales in the preceding periods. Goh and Varaprasad (1986) proposed a Box-Jenkins transfer function model relating returns to previous sales. This model estimates the probability of a return after a certain number of periods, for a fixed set of data. However, in practice, the data is not fixed but is augmented in each period as new sales and return information become available. A model which is conducive to data augmentation is the distributed lag model (DLM) considered by Toktay et al. (2000). Another advantage of using this type of model is that it generally involves the estimation of fewer parameters than the transfer function method; therefore less data is required for the analysis. In the current study a Bayesian estimation approach is used with a DLM to forecast product returns. The model's key departure from the Toktay et al. (2000) approach is discussed in detail later.

While our approach can be applied in a variety of remanufacturing environments, the details of our model are rooted in the remanufacturing operations of an electronics original equipment manufacturer (OEM) located in the Midwest. The OEM designs and manufactures components for telecommunication applications, medical systems, commercial imaging, and

field-deployable defense applications. These components are in the form of embedded circuit boards and modular storage units. The OEM also remanufactures these components and is also a contract remanufacturer of some of the products that these components are placed in.

The OEM has service contracts with several of its customers containing clauses for repair as well as remanufacturing of the components and the products that the components are placed in. The service contract requires the OEM to be responsible for disposal or reuse of components and products. Also, a product returned to the OEM for repair or maintenance can itself be remanufactured and components can be replaced with remanufactured parts. Because of the requirements of the contractual agreement, the OEM accepts all products returned by the customer irrespective of the quality of the returned product. Therefore customers can return products to the OEM after their end-of-life or end-of-use and the OEM either disposes or remanufactures the products. The OEM's primary source of *cores* (a core is a returned product that can be remanufactured; Hauser & Lund, 2003) for remanufacturing is returned products from customers for repair, end-of-use, or end-of-life.

The OEM is a contract remanufacturer therefore the volume of products remanufactured by the OEM is determined by the monthly orders of the contractor(s). The terms of delivery are specified in the contract agreement.

At the start of each month the OEM takes stock of the cores in inventory. If there is an insufficient quantity of cores in-stock to meet the period's demand for remanufacturing (demand includes orders for remanufactured products and components for any repairs due in the next period) the OEM places an order for the balance of cores with core brokers. Ordering from brokers at the start of the month has three main advantages. First, the timing of product returns is uncertain and therefore the OEM's use of brokers early, as opposed to after

observing cores returned by customers during the month, leads to a more efficient scheduling of labor and materials for production. Second, by committing early to a bulk order at the start of the month, the OEM is often able to elicit discounts from the brokers. Third, by contacting the brokers early the OEM is able to give the brokers more time to allocate resources and effort toward fulfilling the order, thereby reducing the risk that the brokers will be unable to fulfill the order on time. If the brokers are unable to deliver the full order quantity on-time then as a last resort, the OEM will use new parts (either purchased or from stock) to meet the shortage. Even though new parts can be used for remanufacturing in the event of a shortage of cores, the reverse is not true in this case; remanufactured parts are not used as a substitute for new parts. This means that demand for new component parts and remanufactured parts are independent demand streams. Figure 1 below shows the setup for a product in the environment being considered when demand is greater than the on-hand inventory at the start of the period.

INSERT TABLE 1 ABOUT HERE

INSERT FIGURE 1 ABOUT HERE

A timeline of the acquisition decision described in Figure 1, with and without the use of forecasting, is shown in Figure 2 below.

INSERT FIGURE 2 ABOUT HERE

Currently, similar OEMs acquire cores according to the timeline shown in Figure 2a. In Figure 2a the OEM does not estimate the quantity, Q_t , and therefore the demand balance for cores is satisfied by acquired units and all returned products which can be remanufactured during the period are placed in inventory. The OEM is considering acquiring cores using the timeline shown in Figure 2b in which forecasts of the distribution of m_t are used to estimate the quantity Q_t . By accurately estimating the Q_t quantity at the start of the month (i.e., at time $t-1$), the OEM can better plan for the quantity of cores to acquire for production (i.e., A_t) and manage inventory more efficiently. However, both the over and under estimation of Q_t have cost implications. Overestimation of Q_t (i.e., less remanufacturable returns arrive in period t than predicted by the forecast of the distribution of m_t) leads to either an expedited order to the brokers at premium costs, or the use of new parts for remanufacturing, which are at a higher cost. Underestimation of Q_t (i.e., more remanufacturable returns arrive in period t than predicted by the forecast of the distribution of m_t) would result in either an accumulation of cores in inventory for the period, if there is sufficient space, or a removal of valuable usable parts from the cores followed by disposal of remaining parts if there is insufficient space to hold all the returns in inventory. For strategic reasons, the OEM prefers to remove valuable usable parts from the cores and then dispose of the remaining parts of the cores, when there is insufficient space, rather than simply sell the core to brokers.

While the key decision variable in Figures 1 and 2 is Q_t , estimating Q_t depends on knowledge about the distribution of m_t (i.e., the probability distribution of the number of remanufacturable products returned in each period). The quantity, Q_t^* , that minimizes the inventory and purchasing costs of cores can be determined with forecasts of the distribution

of product returns. We illustrate this approach in a later section. In the literature review section we provide an overview of existing methods that have been used to model the distribution of m_t , and show how these methods have been applied to production control for reuse activities. While the dual sourcing environment shown in Figure 1 is based on the service contract that the OEM has with its customers for disposal and reuse of returned products, the same environment arises because of environmental legislation. Laws mandating companies to be responsible for take-back of electronic wastes are becoming prevalent in the US, at the state level. The Remanufacturing Institute (www.reman.org) reports that as of May 2011, 26 states in the US had passed laws requiring OEMs to be responsible for taking back and reusing and/or disposing of electronic products at end-of-life. OEM- remanufacturers in states which have passed such laws may likewise face the forecasting problem described above.

LITERATURE REVIEW

Forecasting models for production are usually used for demand forecasting. An overview of the methods used to forecast demand for production can be found in Nahmias (2000, p. 59–108). A distinguishing characteristic of forecasting returns for remanufacturing are that the amounts forecasted depend on the volume of the original product sold. Failure to exploit this characteristic will result in inaccurate forecasts. Traditional time series methods used for forecasting demand such as exponential smoothing or ARIMA cannot exploit this relationship between returns and sales since the dependence of sales and returns cannot be represented in these models. Therefore the use of such models to forecast product returns would yield

inaccurate forecasts.

Goh and Varaprasad (1986) were one of the earliest to consider the relationship between sales and returns in a model they used for forecasting the returns of returnable bottles. They used a Box–Jenkins transfer function model that required large amounts of returns and sales data in order to obtain estimates of the probability that a product will be returned in each of the periods considered in the dataset. Once the estimates were obtained, they were used to forecast all future returns. A drawback to this approach is the data requirement. Chatfield (1996) evaluated Box–Jenkins transfer function models and noted that effective fitting of Box–Jenkins models required at least fifty observations in order to obtain reliable estimates.

Kelle and Silver (1989) developed a forecasting model for the acquisition of reusable containers which required less parameter estimations than the Goh and Varaprasad (1986) model. Like Goh and Varaprasad (1986) they also exploited the relationship between returns and past sales in their forecasting model. They developed a normal approximation to a multinomial distribution for returns, based on past sales, and they used this approximation to estimate the mean and variance of returns. These estimates were then used to compute a base stock level for a continuous review system used for remanufacturing. A drawback of their model is that it could not be easily updated and so new sales or return information would require re-estimation of the model parameters, involving multiple inference procedures and steps.

As previously mentioned, a characteristic of the remanufacturing environment is that the data is usually augmented in each period as new sales and return information becomes available. Bayesian estimation facilitates data updating and is therefore a natural

methodology for estimating product returns. Toktay et al. (2000) used Bayesian methods to estimate parameters for the distribution of product returns. They used a DLM to capture the dependence of returns on sales in previous periods.

DLMs have played a prominent role in numerous applications in the economics and agricultural economics literature. Early examples include the study on the response of capital investment to various aspects of the economic environment (Koyck, 1954), the studies by Nerlove (1958) on the response of agricultural supply to price, and the study on capital appropriations and expenditures by Almon (1965). Many econometric textbooks include a chapter on DLMs. As an example, when forecasting product returns let n_t and m_t denote the number of products sold and returned at time t , respectively; the general form of a DLM is as follows:

$$m_t = \sum_{k=1}^{t-1} \beta_k n_{t-k} + \varepsilon_t; \text{ For } t = 2, 3, \dots, T.$$

(1)

Equation (1) is known as the *finite* distributed lag model. The ε_t terms are usually assumed to be additive white noise (i.e., normally distributed, independent of the n_t values, independent of each other, with a constant variance given by σ^2). The term β_k in Equation (1) is the k th *reaction coefficient*, and it represents the proportion of n_{t-k} (e.g., sales in period $t-k$) that contributes units toward 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. One important aspect to be considered is the number of parameters involved in Equation (1). If the number of terms in Equation (1) is small (i.e., T is small) then the equation can be estimated by using ordinary least squares (OLS). Usually, when historical sales and return data is used for

estimation there are many terms and little is known about the form of the lag. In that case, direct estimation via OLS uses up a large degree of freedom and is likely to lead to imprecise parameter estimates because of multicollinearity (Pindyck & Rubinfeld, 1998). These difficulties can be avoided by assuming that the β_k coefficients are not all independent but functionally related (Zellner, 1971). There are different specifications for this functional relationship. Some of them are based on economic theory; others are based on expert opinion. Functions which have typically been used to represent these relationships include the geometric distribution first used by Koyck (1954), the negative binomial (or Pascal) distribution first used by Solow (1960), and a polynomial function first used by Almon (1965). The polynomial function developed by Almon (1965) offers the most flexibility, in terms of the various shapes that can be accommodated by the function, but requires several conditions to be met before it can be successfully used, and the parameters of the Almon model can be difficult to interpret (Pindyck & Rubinfeld, 1998). The Koyck (1954) and Solow (1960) methods do not have these issues.

In the case of remanufacturing, the functional relationship between the β_k coefficients is called the *delay function* and it represents the time for returns to be made. Toktay et al. (2000) used geometric and negative binomial delay functions to represent the β_k coefficients as follows:

$$\beta_k = pq(1-q)^{k-1}; \beta_k = p \binom{k+r-1}{r} q^r (1-q)^k, \quad (2)$$

where p is the probability that a sold product will ever be returned; q is the conditional probability that a product would be returned in the next instance of time given that it will be returned; r is a parameter pre-specified to represent the lag with the largest β_k coefficient.

The negative binomial delay function allows for more flexibility in the shape of the function compared to the geometric delay function.

The geometric and negative binomial distributions are discrete and assume that delay lags are of equal periods and are thus best suited for use with integer time periods. In practical terms, equal lags mean that the assumption is being made that sales and return volume information is being recorded at equal times each period (e.g., the same number of days between each record of returns or sales) and also that any sales used to obtain the forecast are end-of-period sales. There are situations when managers may want to make a forecast decision based on the most recent information available rather than wait until the end of the period to make the forecast. In such a situation a non-integer (fractional) lag will need to be used and therefore a model with a geometric or negative binomial delay function is not appropriate. To account for the possibility of unequal lags and improve the accuracy of the model, we consider using continuous functions to model the delay. In this study we will be using an exponential delay function of the form below:

$$\beta_k = p\lambda e^{-\lambda k}. \quad (3)$$

The exponential delay function is the continuous analog of the geometric delay function used in Toktay et al. (2000). Figure 3 shows an illustration of geometric and exponential delay functions with the parameter q for the geometric and λ for the exponential chosen so that the value of the delay functions are equal, to three decimal places, for the first lag.

INSERT FIGURE 3 ABOUT HERE

In Figure 3 the estimate of the delay function at lag 2 for the exponential delay function

(0.155) is approximately 7% larger than that of the geometric delay function (0.145), even though the two functions are equal at lag 1. This illustrates a bias associated with using the geometric delay function to estimate a continuous exponential delay. The cost implications of the bias are illustrated with a numerical example in a later section. Figure 3 also shows that the geometric delay function is undefined for fractional time periods, therefore forecasts using fractional time periods with the geometric delay would be an extrapolation. Another advantage of using an exponential delay function is that it is more consistent with the assumption of exponential inter-arrival times of returns used in returnable inventory systems (Buchanan & Abad, 1998; Toktay et al., 2000). A third benefit of using an exponential delay is in the interpretation of the parameters of the DLM. With an exponential delay the parameter λ is interpreted as the average *delay rate* which is the average number of lags per return period. Therefore it can be used to indicate, on average, how many months of previous sales contribute to a month's worth of returns. This rate interpretation may have more practical meaning to managers than the interpretation of q in the geometric delay as being the conditional probability of a product return. A key advantage of the methods developed in this study, for estimating the model with an exponential delay function, is that it can be extended to developing models with other types of continuous delay functions and therefore the contributions of this study extend to many areas, and provide a significant addition to the existing literature.

Once an appropriate forecasting mechanism has been developed for the product returns, a second concern from a production control and inventory management perspective is how to best use this information to minimize the production and inventory costs. Both deterministic

and stochastic models have been used in the remanufacturing literature to determine the quantity to remanufacture. Dekker, Fleischmann, Inderfurth, and van Wassenhove (2004) give an overview of inventory models that have been used for remanufacturing. The stochastic models used have treated demand and returns as stochastic processes, but many of these models assume that demand and returns are independent processes. Kiesmueller and Van der Laan (2001) showed that this assumption can lead to poor cost performances of inventory systems with product returns. Kelle and Silver (1987) were one of the earlier works to exploit the relationship between returns and past sales in an inventory model. They used their returns forecasting method to estimate parameters that were then used to compute a base stock level for a continuous review system. As previously mentioned there were drawbacks to their forecasting approach. Toktay et al. (2000) used a queuing model to determine the remanufacturable inventory quantity. Like Kelle and Silver (1987) they exploited the relationship between sales and returns volumes to estimate the rate of returns and compute an aggregate base stock level. However, the queuing model that was used required several exogenous parameters to be known or estimated including the service rates of suppliers and retailers, and the assumption that the retailer's inventory is observable. In addition the queuing model is reliant on the assumption that the parameters used for the model represent the steady state of the system. However, as noted by Toktay et al. (2000, p. 1416), the estimates of the parameters obtained from the DLM are transient and therefore the queuing model does not result in an optimal inventory policy. In this study, the use of the newsvendor model is proposed to determine the amount of returned product to remanufacture, which will minimize holding and production costs given the forecasted distribution of product returns. In

contrast to the queuing model, the newsvendor model does not require the estimation of several parameters and is not dependent on the estimation of steady state parameters for inventory decisions. This is because it is a single period model that allows for model parameters to be adjusted each period (Khouja, 1996, 1999). The single period nature of the decisions made in Figure 1 makes the newsvendor a natural modeling approach to employ, and it has been used in prior remanufacturing studies for acquisition decisions (e.g., Galbreth & Blackburn, 2006; 2010). While the use of a multi-period model (for an example see DeCroix & Zipkin, 2005) may result in an optimal decision, the added complexity of the model may outweigh the motivations and contributions of this study. Indeed, DeCroix and Zipkin (2005, p. 1252) noted that the consideration of product returns which are dependent on past sales would add additional complexity to their multi-period model which would yield little benefit. Further justification for the use of the newsvendor model is provided in the next section on model development.

MODEL DEVELOPMENT

With the exponential delay function, using the DLM in Equation (1), we can represent the number of products returned during a period that can be remanufactured as follows:

$$m_t = p\lambda e^{-\lambda} n_{t-1} + p\lambda e^{-2\lambda} n_{t-2} + p\lambda e^{-3\lambda} n_{t-3} \dots + p\lambda e^{-(t-1)\lambda} n_1 + \varepsilon_t; \text{ For } t=2, 3, \dots, T. \quad (4)$$

Using the Koyck transformation, subtracting $e^{-\lambda} m_{t-1}$ from both sides of Equation (4), we obtain:

$$m_t = e^{-\lambda} m_{t-1} + p\lambda e^{-\lambda} n_{t-1} + \varepsilon_t - e^{-\lambda} \varepsilon_{t-1}; \text{ For } t=2, 3, \dots, T. \quad (5)$$

Let $u_t = \varepsilon_t - e^{-\lambda} \varepsilon_{t-1}$. To estimate the parameters in Equation (5) we need to determine the likelihood function (i.e., a function of the data and parameters of interest; Rossi, Allenby, &

McCulloch, 2005). This is found by first determining the covariance matrix. For a given set of (T) time periods $\mathbf{u} = (u_2, u_3, \dots, u_T)$, the covariance matrix for \mathbf{u} is given by $\Sigma_u = \sigma^2 \mathbf{G}$ where \mathbf{G} is the $(T - 1) \times (T - 1)$ matrix:

$$\mathbf{G} = \begin{pmatrix} 1 + e^{-2\lambda} & -e^{-\lambda} & 0 & \dots & \dots & 0 \\ -e^{-\lambda} & 1 + e^{-2\lambda} & -e^{-\lambda} & \dots & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -e^{-\lambda} & 1 + e^{-2\lambda} \end{pmatrix}.$$

Given the set of returns ($\mathbf{m} = (m_2, m_3, \dots, m_T)$) and sales ($\mathbf{n} = (n_1, n_2, \dots, n_{T-1})$) for periods $t=1, \dots, T$, the likelihood (Rossi et al., 2005) for the parameters is therefore given by:

$$\ell(\lambda, p, \sigma^2) \propto \frac{|\mathbf{G}|^{-(T-1)/2}}{(\sigma^2)^{(T-1)/2}} \exp \left[-\frac{1}{2\sigma^2} (\mathbf{m} - e^{-\lambda} \mathbf{m}_{-1} - p\lambda e^{-\lambda} \mathbf{n})' \mathbf{G}^{-1} (\mathbf{m} - e^{-\lambda} \mathbf{m}_{-1} - p\lambda e^{-\lambda} \mathbf{n}) \right], \quad (6)$$

where $\mathbf{m}_{-1} = (m_1, m_2, \dots, m_{T-1})$, $|\mathbf{G}|$ is the determinant of the matrix \mathbf{G} and the symbol “ \propto ” indicates the absence of the normalizing constant. To estimate the parameters in Equation (6) using a Bayesian approach, we need to first specify *priors* for the parameters. We used the following *conjugate priors* for p , λ , and σ^2 :

$$\lambda \sim \text{Gamma}[\alpha_0, \beta_0],$$

$$p \sim \text{Uniform}(0,1),$$

$$\sigma^2 \sim \text{InvertedGamma}\left(\frac{v_0}{2}, \frac{v_0 s_0^2}{2}\right),$$

where s_0^2 (the sum of the squared deviations of returns) and v_0 are prior parameters to be specified; Rossi et al. (2005) provide details on how to specify v_0 .

Similarly α_0 and β_0 are prior parameters to be specified. The inverted gamma distribution

can also be represented as the inverse of a scaled chi-squared random variable with appropriate degrees of freedom (Rossi et al., 2005). Therefore an alternative representation of the prior distribution of the variance parameter is $\sigma^2 \sim \frac{\nu_0 s_0^2}{\chi_{\nu_0}^2}$. We make use of this representation later on.

In Bayesian analysis, the *prior* distribution represents the beliefs of the decision maker about the unknown parameters expressed in a probabilistic statement. Since the performance of successive updates depends on the prior, its choice is important (Rossi et al., 2005, Ravines, Schmidt, & Migon, 2006). Toktay et al. (2000) used the “non-informative” prior $\frac{1}{\sigma}$ for the variance parameter and the Uniform(0,1) prior for the parameters p and q . The advantages of using such priors are that they help to simplify the calculations and they are robust (Berger, 1985). The use of the exponential delay function means that the parameter space of λ (i.e., all positive real numbers) cannot be represented with the uniform distribution as considered in Toktay et al. (2000). Also, the $\frac{1}{\sigma}$ prior is an improper prior (i.e., it does not integrate to 1 over the possible parameter space of σ). Berger (1985) shows how the use of improper priors can lead to improper posterior distributions, and therefore an inaccurate forecasting distribution for returns. Gelfand and Sahu (1999) stress that the use of improper priors when employing Monte Carlo Markov Chain (MCMC) estimation techniques, as used in this study, often lead to biased results. The problem of improper priors is avoided by using the conjugate priors above. A conjugate prior for a parameter is a distribution for which the posterior is also of the same family. The inverse gamma distribution is the natural conjugate prior for σ^2 (Fink, 1995, Rossi et al., 2005, p. 24). The gamma distribution is the conjugate prior of λ and the uniform distribution is a conjugate prior for p (Fink, 1995). A criticism of

conjugate priors is that, in general, they are not as robust as non-informative priors. However, if different specifications of the conjugate priors lead to the accurate recovery of the actual parameters used in a simulation, then the prior is considered to be fairly robust and can therefore be used for estimation (Gelfand & Sahu, 1997). The simulation results in Appendix A show that this was indeed the case in this study; therefore robustness of the conjugate priors was not considered an issue. The specification of the conjugate priors and the use of the exponential delay make it difficult to solve analytically for the posterior distributions, therefore a MCMC solution approach was employed. The posterior distributions for p , λ , and σ^2 are estimated by making use of the random walk Metropolis–Hastings (M-H) algorithm to simulate draws from the posterior distribution of model parameters (Chib & Greenberg, 1995; Rossi et al., 2005, p. 86).

The details of the estimation procedure are provided below:

- (i) Start with initial point estimates \hat{p} , $\hat{\lambda}$, and $\hat{\sigma}^2$.
- (ii) Generate: $\hat{\lambda}^{new} = \hat{\lambda}^{old} + \varepsilon$; $\varepsilon \sim N(0, step^2)$; $\hat{p}^{new} = \hat{p}^{old} + \xi_2$; $\xi_2 \sim N(0, step_1^2)$;
 $step$ and $step_1$ are numerical values chosen to enable the algorithm to have sufficiently navigated the space where the posterior has high mass.
- (iii) Compute $\alpha = \min \left\{ 1, \frac{\ell(\hat{\lambda}^{new}, \hat{p}^{new}, \hat{\sigma}^2) \pi(\hat{\lambda}^{new})}{\ell(\hat{\lambda}^{old}, \hat{p}^{old}, \hat{\sigma}^2) \pi(\hat{\lambda}^{old})} \right\}$; where $\pi(\cdot)$ is the prior for λ .
- (iv) With probability α , $\hat{\lambda} = \hat{\lambda}^{new}$ and $\hat{p} = \hat{p}^{new}$, else $\hat{\lambda} = \hat{\lambda}^{old}$ and $\hat{p} = \hat{p}^{old}$.
- (v) Generate \mathbf{G} using $\hat{\lambda}$.
- (vi) Generate: $\hat{\sigma}_{new}^2 | \mathbf{m}, \mathbf{n}, \lambda, p \sim \frac{v_1 s_1^2}{\chi_{v_1}^2}$ with $v_1 = v_0 + (T - 1)$, $s_1^2 = \frac{v_0 s_0^2 + (T-1)s^2}{v_0 + (T-1)}$.
- (vii) Repeat (ii)–(vi) for a sufficient number of times (10,000 was used for this study).

The above yields the estimates $(\hat{p}, \hat{\lambda}, \hat{\sigma}^2)$ for the joint posterior distribution for p , λ , and σ^2 .

The estimate of the posterior distribution for m_t (i.e., $\hat{F}(m_t)$, which is the period forecast of the distribution for m_t) can then be obtained by making the relevant substitutions into the part of Equation (4) which does not include the error terms (ε_t). This forecasted distribution function can then be used to estimate the quantity Q_t^* which minimizes inventory costs in Figure 1, as explained below.

Estimation of the Quantity Q_t^* which Minimizes Inventory Costs

To estimate Q_t in the environment described in Figure 1, let c_o = cost of overestimation per unit per period and c_u = cost of underestimation per unit per period. In the context of the problem described in Figure 1, if less remanufacturable returns arrive than predicted by the forecast of the distribution of returns, in a period in which cores need to be acquired, then this leads to an expedited order to the brokers for the remaining balance. Thus, the number of units expedited is the units overestimated. Therefore, c_o is the cost differential between ordering a unit core from brokers later on in the period and at the beginning of the period (e.g., expedited order at premium costs, or the cost of a new part used for remanufacturing).

Similarly, if more remanufacturable products are returned in a period than are predicted by the forecast of the distribution of returns, then this would result in those additional cores remaining in inventory at the end of the period. Thus, the number of units underestimated is added to the end-of-period inventory. c_u is therefore the period cost of not remanufacturing a usable returned product during the period (e.g., period inventory cost per unit, or the cost of the removal of valuable usable parts then disposal, per unit). The following newsvendor objective can then be used to determine Q_t in each period t :

$$\min_{Q_t} \left\{ C(Q_t) = c_o \int_0^{Q_t} (Q_t - m_t) dF_t(m_t) + c_u \int_{Q_t}^{\infty} (m_t - Q_t) dF_t(m_t) \right\}, \quad (7)$$

where $F_t(\cdot)$ is the cumulative distribution function (cdf) for m_t . This leads to the well known result for the Q_t^* quantity that minimizes Equation (7):

$$F_t(Q_t^*) = \frac{c_u}{c_o + c_u}. \quad (8)$$

Thus, the estimate of Q_t^* can be obtained using the period forecast of the cdf of product returns (i.e., $\hat{F}(m_t)$). Note that the accuracy of the forecast, $\hat{F}(m_t)$, from the DLM in Equation (4) depends on the use of the most recent information meaning that the single period forecast based on the most recent sales and returns information will be the most accurate of any future period forecasts. The DLM, used in this way, is therefore consistent with a single period model such as the newsvendor model.

We coded the M–H algorithm, used to obtain $\hat{F}(m_t)$, with the **R** version 2.6.1 software language (a copy of the code is available from the authors upon request). Data is input as an Excel comma-separated values (csv) file and the results are output in the same format. An example with 40 periods of sales and product returns data required approximately 68 seconds to run on a Toshiba NB205 notebook with 2GB RAM and an Atom 1.6GHZ processor.

As long as proper priors are used, the M–H algorithm can successfully be used to obtain the posterior distribution of any properly specified likelihood (Chib & Greenberg, 1995, p. 330). Thus the above algorithm could be modified for estimating any type of delay function used in the DLM. We therefore used the M–H algorithm to also estimate the DLM with the geometric delay function used in Toktay et al. (2000) and evaluated this version along with our exponential delay model. Because of data confidentiality issues, we were unable to obtain the data used in Toktay et al. (2000) to evaluate our model. We therefore used simulated

versions of that data for our comparisons. The details and results of the simulation are shown in the next section.

MODEL VALIDATION AND COMPARISONS

Model Validation

To check that the M–H algorithm could correctly estimate the distribution of the rate parameter, λ , for our exponential delay model and that our conjugate prior choices were fairly robust, we ran a simulation. Details of the simulation are provided in Appendix A. We were able to accurately estimate all the parameters in each of the scenarios considered in the simulation. The results of the simulation therefore showed that the posterior distributions were fairly robust to our selection of prior specifications and also that the M–H algorithm was able to correctly estimate the various parameters.

We also estimated the posterior distribution for the geometric delay parameter (q) for the model used in Toktay et al. (2000), by using an inverted gamma conjugate prior for the variance parameter (instead of the non-informative prior used in Toktay et al., 2000) and constructing an M–H algorithm to perform the estimation. Similar to the exponential delay model, we performed a simulation to check that the M–H algorithm we constructed for the geometric delay could correctly recover actual parameters. Parameters for the simulation were chosen to be consistent with the simulation performed in Toktay (2004) (i.e., $q=0.125$, sales was generated with a Poisson random variable with parameter 200). The resulting 95% credible intervals and time series plots for the parameter q (not shown) indicated that the M–H algorithm was able to correctly estimate the parameter (q). Similar to the evaluation of the exponential delay model, we tried different prior (i.e., p_0, q_0 , and σ_0) specifications and

we were able to recover the actual parameters for all specifications.

Comparison of New (Exponential Delay) Model with Existing (Geometric Delay) Model

Our survey of the literature revealed that the DLM proposed by Toktay et al. (2000), for forecasting product returns, required the least amount of data for estimation out of all models considered in the literature in which random returns were allowed to depend on past sales.

Therefore the use of the DLM, along with Bayesian estimation, for modeling product returns has a significant estimation advantage over previous models. However, the use of a geometric or exponential delay in the DLM depends on the data. If the data is collected in discrete equal periods and the lags are more consistent with a geometric distribution then the geometric delay model can be used. However, if the lags of the data collected are consistent with the exponential distribution then our model provides an alternative to the geometric delay model.

A natural question that arises is: what is the bias involved in using a geometric delay for estimation when the data actually follows an exponential delay (and vice versa)? To answer this question we ran simulations to evaluate the performance of the two methods when the data was generated with the other type of delay function. First, an exponential delay function was used to estimate the parameters for data generated by a geometric delay function with equal lags. Next, a geometric delay function was used to estimate the parameters for data generated by an exponential delay function of equal and unequal lags. Performance was measured by the mean absolute percent estimation error (similar to the MAPE measure used for evaluating time series forecasts in Chatfield (1996)) for each delay function, as specified in Equation (9):

$$\text{Error} = \frac{\sum_k^T \frac{|\widehat{r}_D^{(k)} n_{t-k} - r_{D^{(k)}} n_{t-k}|}{r_{D^{(k)}} n_{t-k}}}{T} \times 100\%, \quad (9)$$

where $\hat{r}_D(k)$ is the point estimate of the delay function at lag k ; $r_D(k)$ is the actual value of the delay function at lag k (i.e., $r_D(k) = pq(1-q)^{k-1}$ or $r_D(k) = p\lambda e^{-\lambda k}$); n_{t-k} is the sales at lag k ; T is the maximum number of estimation periods in the dataset (40 in our simulation). We ran simulations to evaluate the performance of the methods, using Equation (9), with data specifications that were either matched or mismatched with model specifications. The unequal lags data for the simulation was generated with the first lag chosen to be 1.5 instead of 1. A practical interpretation of this is that this corresponds to making next month's forecast using sales information obtained in the middle of the current month as opposed to the end of the month. The sales at lag 1.5 was obtained using a Poisson random variable with parameter 10,000 (=20,000/2); all other sales were generated with a Poisson random variable with parameter 20,000. Returns were generated using a geometric delay function with parameters $q=1/t$ ($t=2, \dots, 12$) and an exponential delay function with parameter $\lambda = 1/t$ ($t=2, \dots, 12$). The values of t represent the average duration that a product stays with the customer before it is returned (e.g., minimum of two periods, maximum of 12 periods corresponding to a year). Note that the geometric delay is undefined for $q=1$ which is why $t=1$ was not used. For all cases $p = 0.5$ and $\sigma^2 = 1$. These parameters were chosen to be consistent with the simulations performed in Toktay (2004) for evaluating a DLM with a geometric delay. Table 2 below shows the simulation results.

INSERT TABLE 2 ABOUT HERE

Table 2 shows that there is negligible estimation error when the geometric delay is used for estimation with data generated with the geometric delay, with equal lags, and when the exponential delay is used for estimation with data generated with an exponential delay with

both equal and unequal lags. When the lag in the data followed an exponential pattern and it was estimated using a geometric delay (and vice versa), there was a bias of at least 6% on average, with maximum percentage errors exceeding 12%. On average, the estimation error was slightly lower when the exponential delay estimation was used for data generated with the geometric delay with equal lags than when the geometric delay estimation was used for data generated with an exponential delay with both equal and unequal lags. The percentages are not so different to indicate that one estimation method has a clear advantage over another when the data is generated using a distribution other than what is assumed in the model. However, our method does offer an alternative to the existing geometric delay function model. Table 2 shows that our method reduces the bias in estimation when the data is consistent with an exponential delay lag. As mentioned previously in the model development section, an exponential delay is more consistent with the assumption of exponential inter-arrival times of returns used in returnable inventory systems than a geometric delay. Percentage biases of the magnitude shown in Table 2 have cost implications. The next section provides a numerical example of our method for forecasting product returns, in the remanufacturing environment shown in Figure 1, which illustrates the cost implications.

NUMERICAL ANALYSIS

Parameter Values

We visited the plant for the OEM-remanufacturing operation depicted in Figure 1 and discussed the parameter designs for simulating the operation with company managers. The cost parameters used are either real data from the company or estimates that we derived after we consulted the manager. The average unit purchase price of the cores used by the

OEM-remanufacturer was estimated to be \$450. The average cost of a “bucket” of new parts, if used for remanufacturing as a replacement for a core, was estimated to be \$900. The unit cost of expediting an order for a typical core from brokers was determined to be \$45, on average. Therefore we set the unit cost of over estimating the volume of product returns, and having to expedite a unit from brokers, at \$45. In the Hauser and Lund (2003) survey, the average holding cost (i.e., \$/unit/per year) for cores in the electronics industry was estimated to be 15% of the cost of the core (when purchased from brokers). Therefore we set the unit holding cost in this study to be \$5.63. This unit holding cost represents the cost of under estimating the volume of product returns and therefore having to hold returned units in inventory. We used these costs as the basis for our numerical example. The monthly sales volume of new product for the company was estimated to be 2000 units, on average. Therefore, we modeled monthly sales of new products with a Poisson distribution with parameter 2000. To model product returns, the exponential delay was used with $\lambda = 0.125$ (i.e., the average duration that the product stays with a customer until it is returned is $= 1/0.125 = 8$ months for the simulation). The probability that a product will ever be returned (i.e., p) was set at 0.5. These values are consistent with previous simulations which used DLMS to represent the distribution of product returns (e.g., Toktay et al., 2000; Toktay, 2004). The results in Appendix A show that our estimation method is robust to different parameter specifications, therefore the use of these parameter values should not affect the qualitative conclusions of this simulation study. The standard deviation of the product returns was set at $\sigma = 100$, meaning that product returns had a low coefficient of variation (c.v.) of 10%. Demand for the remanufactured product was modeled as a Poisson distribution with

parameter 1000. This parameter was chosen so that average supply from the sale of new products would match average demand for the remanufactured product. We consider alternative c.v. and demand parameters in our sensitivity analysis.

Policies

The simulation considered the acquisition policies described earlier in Figures 2a and 2b. In the NF policy, shown in Figure 2a, no forecast of product returns is used for procurement decisions (i.e., the base case). The amount of cores in stock at the beginning of the period is used for remanufacturing and the remaining demand balance is satisfied by acquired cores from brokers. We assume that the brokers are able to supply all the quantities ordered (i.e., no shortages) and that there is no disposal of cores (i.e., all cores are remanufacturable).

Consideration of disposals and shortages should not affect the qualitative conclusions of the results. Using this policy, all remanufacturable product returns which arrive during the period are placed in inventory.

The remaining two policies, based on Figure 2b, use forecasts of the distribution of product returns to determine the quantity of cores returned during the period that should be remanufactured during the same period in order to minimize costs. It is assumed that all products returned during the period arrive at a time before the end of the period when they can be remanufactured to satisfy demand. In the G policy the geometric delay function estimates are used in the DLM to make the forecast, and for the E policy our proposed exponential delay function estimates are used to make the forecast. The amount of cores in stock at the beginning of the period is used for remanufacturing and, after accounting for Q_t^* , an order quantity (i.e., A_t) is placed with the brokers for the remaining demand balance. With

both G and E forecasting policies, if less remanufacturable returns arrive than predicted by the forecast of the distribution of returns (i.e., $m_t < Q_t^*$), in a period in which cores need to be acquired, then this leads to an expedited order to the brokers for the remaining balance (i.e., $Q_t^* - m_t$). Therefore, the number of units expedited would be the units overestimated. Similarly, if more remanufacturable products are returned in a period than predicted by the forecast of the distribution of returns (i.e., $m_t > Q_t^*$), then this would result in those (i.e., $m_t - Q_t^*$) additional cores remaining in inventory by the end of the period. Therefore, the number of units underestimated is added to the end-of-period inventory.

Implementation

For each of the three policies, the system begins with no cores in inventory. Each period in which acquisition decisions are made represents a month. It is assumed that all unsatisfied demand for remanufactured products is met with expedited cores. Thirty periods of sales were considered with forecasting started in the fifth period of sales (i.e., 26 periods of forecasts were considered). At the beginning of each period the following sequence of activities occur using the data to date: forecast the distribution of returns ($\hat{F}(m_t)$), estimate the quantity of cores returned during the period that should be remanufactured during the same period in order to minimize costs (Q_t^*), and then use the estimate to determine the quantity of cores to procure from core brokers (A_t). For the G and E policies, the procurement quantity represents the smallest integer greater than or equal to the amount suggested by using the forecast of the distribution of returns. Costs are computed over the 26 periods in which forecasts were used, and these include period costs for initial acquisition, expedited purchases, and inventory carrying. Thirty simulation runs were carried out for each

scenario enabling confidence intervals to be obtained for each performance measure. For each run, the same initial random number seed is used across all three policies in order to reduce the variance in pair-wise comparisons.

Results

The costs for the three simulated policies are shown in Table 3. The results from the table

INSERT TABLE 3 ABOUT HERE

show that the simulated average period costs when forecasts are used is 10.5% lower with the geometric delay and 19.3% lower with the exponential delay when compared to costs when no forecasts are used to determine the quantity of cores to acquire from brokers. The use of the exponential delay model for forecasting resulted in 11.1% lower costs than when the geometric delay was used. This translates to period (e.g., monthly) savings of \$8,031, on average, because of the improved forecast accuracy when using the exponential delay model with the data. Table 3 also shows that the average inventory over the 26 periods when no forecasting is used is 779 units. In contrast the average inventory when the geometric delay is used for forecasting is approximately 8 units and this is further reduced to approximately 5 units when the exponential delay is used for forecasting during the same 26 period forecasting horizon. These represent a 99.0% and 99.4% reduction, respectively, in units held in inventory, on average during the period, compared to when no forecasting is used. The use of the geometric delay model for forecasting data consistent with an exponential delay resulted in an average of 7 ($\approx 8.48 - 2.00$) additional units being expedited and an average of 4 ($\approx 7.39 - 4.21$) additional units held in inventory, compared to when the exponential delay model was used. These represent an 82.5% and 54.1% reduction, respectively, in units

expedited and units held in inventory on average during the period. These reductions are significant, as shown by the confidence intervals in Table 3, and indicate that there are significant cost savings to be made via the use of accurate forecasts of the product returns distribution when making purchasing decisions for cores.

Sensitivity Analysis

Realizing that some other factors may influence the results, we conducted sensitivity analyses by varying the cost structure, c.v. of product returns, and the demand rate for remanufactured products. For the cost structure, we varied the newsvendor critical fractile used in Equation (8) from the current 11% to 50% and 89%, respectively. The current 11% cost structure represents the situation whereby overestimation of Q_t^* is more expensive than underestimation (i.e., $c_o = \frac{11}{89} c_u$). The 89% critical fractile represents the opposite situation when $c_o = \frac{89}{11} c_u$. In practice, this type of cost structure may occur in the situation whereby constraints on storage space lead to per period unit inventory carrying costs, and/or disposable costs, which are high relative to unit core expediting or core replacement costs. The 50% cost structure represents the situation whereby the over and underestimation costs are equal (i.e., $c_u = c_o$).

For the c.v. of product returns, we varied the c.v. from the current low level of 10% to a medium level of 33% and a high level of 100%. These values are consistent with coefficients of variation used in other studies (e.g., Vlachos & Dekker, 2003) with newsvendor objective functions.

For the demand factor, we used demand rates for the remanufactured products that were representative of low and high rates that had previously been used in the literature. Kiesmuller

and Minner (2003) and Teunter, van der Laan, and Vlachos (2004) represented a low and high demand rate for remanufactured products as having a ratio of mean demand to mean product returns of 0.6 and 0.9, respectively. Therefore, we set our low demand level for our sensitivity analysis at 60% of the average demand in the base case (i.e., 600 units). In the base case described earlier, average demand for the remanufactured products was set equal to the average volume of products returned during the 30 periods of simulation; therefore the base case was set as our high demand level. Table 4 shows the cost ratios for various policies at the different factor levels used in the sensitivity analysis.

INSERT TABLE 4 ABOUT HERE

The results in Table 4 show that forecasting leads to a reduction in all the scenarios considered. The greatest percentage gain was in the low demand, low c.v., 89% cost structure scenario, in which cost savings in excess of 52% were achieved. This is because at low demand volumes, most of the demand can be met from cores already in inventory and/or products being returned during the period. This results in little or no expediting of cores when forecasting is used for acquisition decisions. Also, forecasting accounts for the expected quantity of products to be returned during a period and therefore results in infrequent acquisition of cores at low demand rates. Without forecasting, the balance of cores is purchased at the start of the period and any returned products during the period are placed in inventory. With high inventory related costs (e.g., a critical fractile of 89%) in addition to low demand, forecasting results in lower inventory levels and therefore significant cost savings versus no forecasting. At higher demand rates, acquisition and expediting costs increase

which result in cost savings from forecasting which are still positive but smaller than savings at lower demand rates.

In all the scenarios considered the largest percentage gains occur in the low c.v. scenarios. The lower variability in the distribution of returns, as measured by the c.v., results in less volatility in the forecast of the distribution of product returns and therefore less over and under estimation of the of the expected quantity of products to be returned during a period. This results in lower costs when forecasting is used with a low c.v. versus a medium or a high c.v.

The $\frac{G}{E}$ results in Table 4 show that at low demand rates there is little difference in the cost savings from using either the geometric or exponential model to make forecasts (i.e., at the low demand rate the percentage differences in Table 4 were insignificant). This is because at low demand rates overestimation does not necessarily result in expediting costs since acquisitions rarely occur and demand is mostly met with returned products or cores in inventory. Similarly, at low demand rates underestimation does not necessarily result in larger acquisition quantities leading to higher inventory levels. This means that any differences in forecast accuracy at low demand rates are unlikely to result in significant cost savings when forecasting is used.

At the high demand rate, Table 4 shows that there are cost savings at all three cost structures as a result of using our proposed exponential delay model for forecasting versus the geometric delay model. The cost savings were highest at a cost structure of 11%, which is when the cost of overestimation (e.g., expediting and/or using new parts for remanufacturing) is larger than the cost of underestimation.

CONCLUSIONS

We developed a model that can be used to accurately forecast the distribution of returns for remanufacturing in any period. For an OEM who also remanufactures, our analysis showed that the use of forecasts for sourcing cores from brokers can yield significant cost savings as compared to when cores are acquired without first forecasting the expected quantity of products to be returned during a period. The methodology used for estimation can be applied to the estimation of the parameters of any properly defined delay function in a DLM. This further generalizes the application of the DLM for forecasting product returns data with various lag patterns. We compared our method to an existing DLM which had been used for forecasting product returns and found that our approach provided more accurate estimates when the return lags were exponential in nature. The analysis revealed that these gains in accuracy resulted in the most cost savings when demand volumes were high compared to the volume of returned products. Such situations require the frequent acquisition of cores to meet demand. The results show that significant cost savings can be achieved by using the proposed approach for sourcing product returns in such situations. An area for future research is the evaluation of the cost savings as a result of using the proposed forecasting model in other (i.e., non newsvendor) inventory type settings. Another area for future research is the consideration of uncertainties in both the demand and supply of cores when using forecasts to make acquisition decisions.

REFERENCES

- Almon, S. (1965). The distributed lag between capital appropriations and expenditures. *Econometrica*, 33(1), 178–196.
- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*, 2nd ed. New York, NY: Springer-Verlag.
- Buchanan, D. J., Abad, P. L. (1998). Optimal policy for a periodic review returnable inventory system. *IIE Transactions*, 30(11), 1049–1055.
- Chib, S., & Greenberg, E. (1995). Understanding the Metropolis-Hastings algorithm. *The American Statistician*, 49(4), 327–335.
- Chatfield, C. (1996). Model uncertainty and forecast accuracy. *Journal of Forecasting*, 15(12), 495–508.
- DeCroix, A. G., & Zipkin, P. H. (2005). Inventory management for an assembly system with product or component returns. *Management Science*, 51(8), 1250–1265.
- Dekker, R., Fleischmann, K., Inderfurth, K., & van Wassenhove, L. N. (2004). *Reverse logistics, quantitative models for closed loop supply chains*. Berlin, Germany: Springer.
- Environmental Protection Agency, Office of Solid Waste (1998). *Macroeconomic importance of recycling and remanufacturing*.
- Fink, D. (1995). A compendium of conjugate priors. Technical report.
- Galbreth, M. R. & Blackburn, J. D. (2006). Optimal acquisition and sorting policies for remanufacturing. *Production & Operations Management*, 15(3), 384–392.
- Galbreth, M. R. & Blackburn, J. D. (2010). Optimal acquisition quantities in remanufacturing with condition uncertainty. *Production & Operations Management*, 19(1), 61–69.

- Gelfand, A. E., & Sahu, S. K. (1999). Identifiability, improper priors and Gibbs sampling for generalized linear models. *Journal of the American Statistical Association*, 94(445), 247–253.
- Goh, T.N., Varaprasad, N. (1986). A statistical methodology for the analysis of the life-cycle of reusable containers. *IIE Transactions*, 18(1) 42–47.
- Hauser, W. M., & Lund, R. T. (2003). *The remanufacturing industry: Anatomy of a giant*. Boston University, Boston, Massachusetts.
- Kelle, P., & Silver, E. A. (1987). Purchasing policy of new containers considering the random returns of previously issued containers. *IIE Transactions*, 21(4) 349–354.
- Khouja, M. (1996). The newsboy problem with multiple discounts offered by suppliers and retailers. *Decision Sciences*, 27(3), 589–599.
- Khouja, M. (1999). The single-period (news-vendor) problem: Literature review and suggestions for future research. *Omega*, 27(10), 537–553.
- Kiesmueller, G., & van der Laan., E. (2001). An inventory model with dependent product demands and returns. *International Journal of Production Economics*, 72(1) 73–87.
- Kiesmuller, G., & Minner, S. (2003). Simple expressions for finding recovery system inventory control parameter values. *Journal of the Operational Research Society*, 54(1), 83–88.
- Koyck, L. M. (1954). *Distributed lags models and investment analysis*. Amsterdam: New-Holland.
- Nahmias, S. (2000). *Production and operations analysis*, 4th ed. New York, NY: McGraw-Hill/Irwin.

- Nerlove, M. (1958). *Distributed lags and demand analysis agricultural and other commodities*. Agricultural Handbook No. 141, U.S. Department of Agriculture.
- Pindyck, R. S., & Rubinfeld, D. L. (1998). *Econometric models and economic forecasts (4th ed.)*. New York, NY: McGraw-Hill.
- Ravines, R. R., Schmidt, A. M., & Migon, H. S. (2006). Revisiting distributed lag models through a Bayesian perspective. *Applied Stochastic Models in Business & Industry*, 22(2), 193–210.
- Rossi, P., Allenby, G., & McCulloch, R. (2005). *Bayesian statistics and marketing (3rd ed.)*. Hoboken, NJ: Wiley.
- Solow, R. (1960). On a family of lag distributions. *Econometrica*, 28(2), 393–406.
- Souza, G. (2008). Closed-loop supply chains with remanufacturing. In Z. Chen & S. Raghavan (Eds.), *Tutorials in operations research*, Hanover, MD: INFORMS.
- Teunter, R. H., van der Laan, E., & Vlachos, D. (2004). Inventory strategies for systems with fast remanufacturing. *Journal of the Operational Research Society*, 55(5), 475–484.
- Toktay, B., Wein, L., & Zenios, S. (2000). Inventory management of remanufacturable products. *Management Science*, 46(11), 1412–1426.
- Toktay, B. (2004). Forecasting product returns. In V. D. Guide, Jr., & L. N. van Wassenhove (Eds.), *Business aspects of closed-loop supply chains*. Pittsburgh: Carnegie Mellon University Press.
- Vlachos, D., & Dekker, R. (2003). Return handling options and order quantities for single period products. *European Journal of Operations Research*, 151(1), 38–52.
- Zellner, A. (1971). *An introduction to Bayesian inference in econometrics*. New York: Wiley.

APPENDIX A: EVALUATION OF M–H ALGORITHM PARAMETER ESTIMATION

To check that the M–H algorithm could correctly estimate the distribution of the rate parameter, λ , for our exponential delay model and that our conjugate prior choices were fairly robust, we ran the following simulation. Parameters for this simulation were chosen to be consistent with those used by Ravines et al. (2006) to evaluate a DLM obtained using an MCMC technique. Five periods of returns (i.e., $T=6$, $m_1=0$) were simulated using code developed with the **R** version 2.6.1 software language. The parameters used for the data are as follows: $\lambda = 0.5$; $p = 0.5$; $\sigma^2 = 1$; the vector \mathbf{n} was generated using a Poisson random variable with parameter 200; \mathbf{m} was generated according to Equation (4). The step size of 0.7 was chosen to have an acceptance rate between 30%–40% which would enable the walk to have sufficiently navigated the space where the posterior has high mass (Rossi et al., 2005). The prior specifications were: $\alpha_0=2$, $\beta_0=1$, $\nu = 3$, and $s_0^2 = 1$. Other prior specifications were tested (results in the table below) and these did not have any influence on the resultant posterior. The algorithm was run for 10,000 iterations, with the first 1,000 iterations as burn-in (i.e., discarded). ACF (autocorrelation function) plots were constructed to check for convergence of the algorithm. The ACF plots (not shown) indicated that there was no autocorrelation between the draws after the burn-in (i.e., burn-in was at the 1000th draw and the autocorrelations were negligible after the 40th draw) and therefore the algorithm converges. The resulting 90% confidence interval contained the true parameter which showed that the code was able to recover the parameter of the simulation. Values of α_0 ranging from 1 to 10 and β_0 ranging from 0.1 to 3 were tried. These parameters were chosen to represent a variety of shapes and scales for the prior consistent with those considered in Fink (1995) for

evaluating the Gamma conjugate prior. These values did not have a strong effect on the posterior distributions of the parameters (i.e., the algorithm was still able to recover the parameters of the model with these alternate hyper parameter values). Alternate values of the parameters ($\lambda(= 0.2)$, $p(= 0.7)$, and $\sigma^2(= 0.3)$) were also tried and the model was also able to recover these parameters. Table A1 below shows the credible intervals for a subset of the values tried.

[INSERT TABLE A1 HERE.]

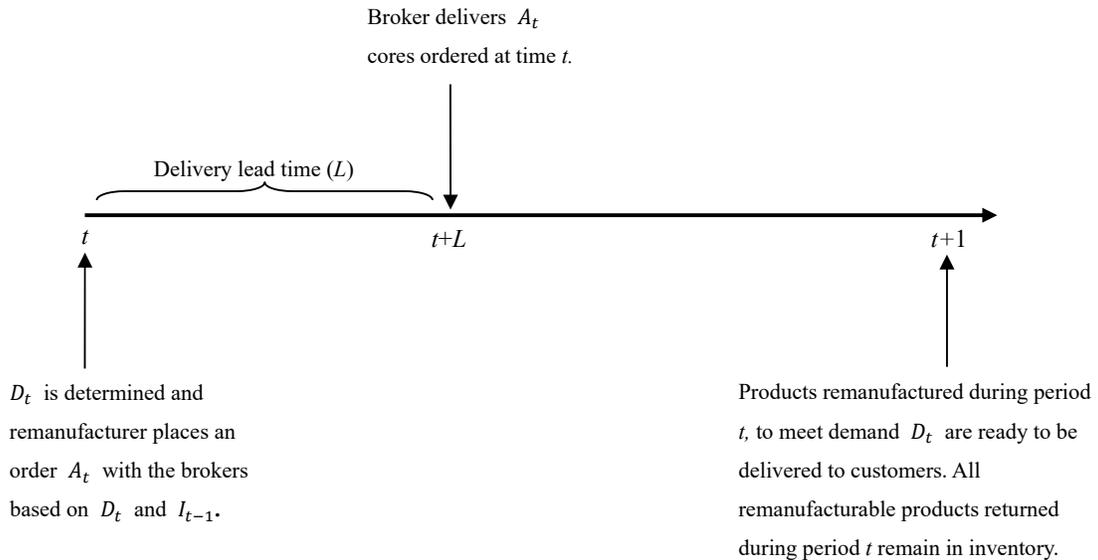
All of the 95% credible intervals reported in Table A1 capture the true parameters. This shows that the posterior distributions are fairly robust to our selection of prior specifications and also that the M–H algorithm is able to correctly estimate the various parameters.

Table A1: Credible intervals for different parameter and hyper parameter values.

α_0	β_0	λ	p	σ^2	95% Credible intervals for λ	95% Credible intervals for p	95% Credible intervals for σ^2
1	1	0.5	0.5	1.0	(0.38,0.77)	(0.31,0.67)	(0.51,5.8)
10	1	0.5	0.5	1.0	(0.48,1.2)	(0.35,0.81)	(0.38,5.1)
2	3	0.5	0.5	1.0	(0.38,0.88)	(0.41,0.58)	(0.8,5.8)
2	0.1	0.5	0.5	1.0	(0.32,0.65)	(0.45,0.72)	(0.3,5.2)
2	1	0.2	0.7	1.0	(0.18,1.3)	(0.56,0.89)	(0.2,2.8)
2	1	0.5	0.7	3.0	(0.48,0.78)	(0.61,0.93)	(0.8,5.4)
2	1	0.2	0.7	3.0	(0.05,0.35)	(0.77,0.98)	(0.2,4.8)

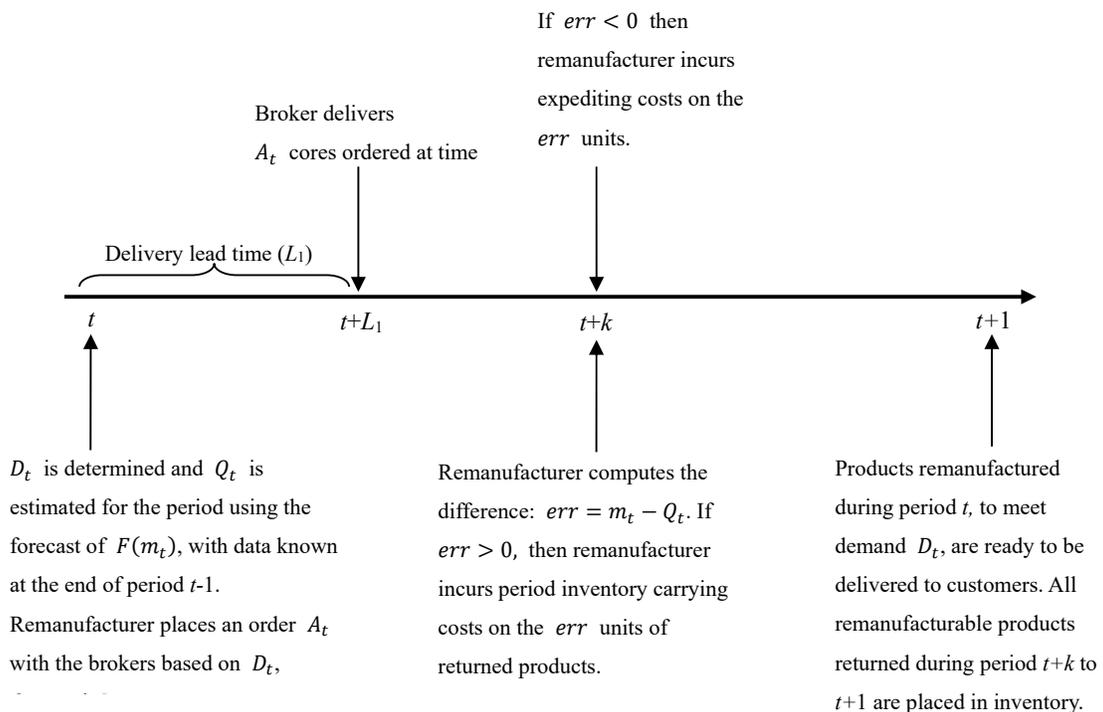
Figure 2: Timeline for the core acquisition decision by the OEM-remanufacturer.

2a) Without forecasting



Note: $0 < L < 1$.

2b) With forecasting



Note: $0 < L_1 < k < 1$; k is a cut-off determined by the remanufacturer such that any usable products returned during period t to $t+k$ can be remanufactured to meet demand D_t (e.g., when t is a month then if k is set at 0.75 this would represent 3 weeks).

Figure 3: Comparison of a discrete and a continuous delay function.

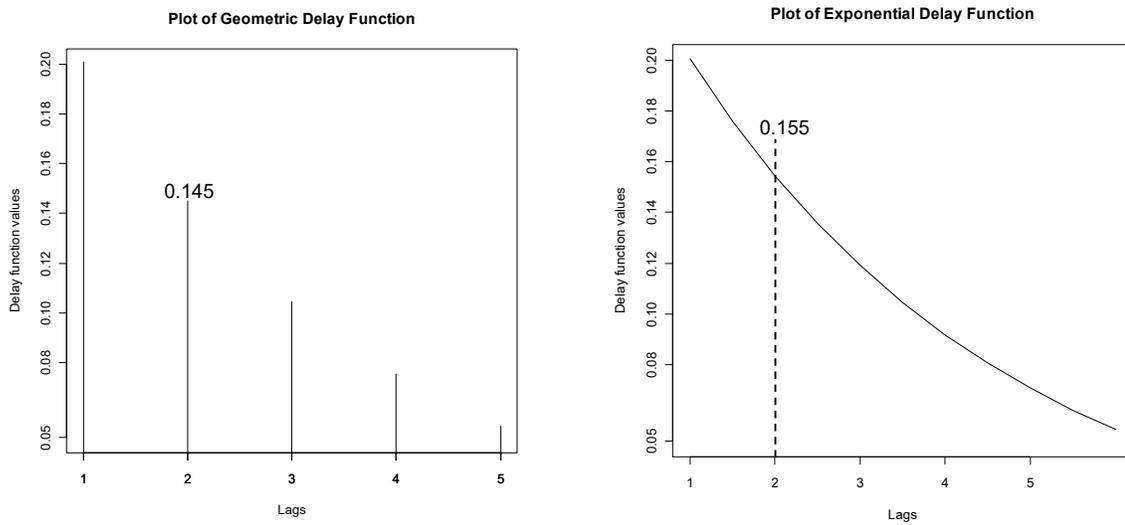


Table 1: Notations used in Figure 1.

Notation	Description
D_t	Demand for the remanufactured product by the end of period t , known at the end of $t-1$.
I_t	Volume of useable cores in inventory at the end of period t .
Q_t	Decision variable for quantity of cores returned during period t that should be remanufactured during period t to minimize costs.
n_t	Number of new components sold in period t .
A_t	Quantity of cores acquired from brokers by the end of period t .
m_t	Number of products returned during period $t-1$ to t that can be remanufactured.
$(a)^+$	$=\max(a,0)$.

Table 2: Simulation results when actual and estimated specifications are matched and mismatched.

Actual Specifications	Model Specifications	Minimum Percentage Error	Average Percentage Error	Maximum Percentage Error
Geometric delay	Geometric delay (priors: $q_{\theta}=0.2, \sigma_{\theta}=0.5$)	0.00%	0.006%	0.011%
	Exponential delay with equal lags (priors: $\lambda_{\theta}=0.2, \alpha_{\theta}=2, \beta_{\theta}=1, \sigma_{\theta}=0.5$)	1.82%	5.96%	12.13%
Exponential delay with equal lags (i.e., lag1=1)	Exponential delay with equal lags (priors: $\lambda_{\theta}=0.2, \alpha_{\theta}=2, \beta_{\theta}=1, \sigma_{\theta}=0.5$)	0.00%	0.017%	0.082%
	Geometric delay with equal lags (priors: $q_{\theta}=0.5, \sigma_{\theta}=0.5$)	2.28%	6.53%	18.28%
Exponential delay with unequal lags (i.e., lag1=1.5)	Exponential delay with unequal lags (priors: $\lambda_{\theta}=0.2, \alpha_{\theta}=2, \beta_{\theta}=1, \sigma_{\theta}=0.5$)	0.00%	0.007%	0.019%
	Geometric delay with unequal lags (priors: $q_{\theta}=0.5, \sigma_{\theta}=0.5$)	2.41%	6.71%	19.56%

Note: Percentage errors were calculated over 40 periods of estimation.

Table 3: Simulated averages (and 95% confidence intervals) for three acquisition policies.

Policy	Description	Total Costs (\$1,000s)	Units remaining in Inventory	Units expedited
1. NF	No forecasts	90 ± 0.91	779 ± 0.92	-
2. G	Forecasts using Geometric Delay	81 ± 0.91	7.39 ± 0.73	8.48 ± 0.30
3. E	Forecasts using Exponential Delay	72 ± 0.81	4.21 ± 0.51	2.00 ± 0.32

Table 4: Relative costs of acquisition policies under various scenarios.

Demand Rate	Coefficient of Variation (c.v.)	Policies Compared	Cost Structure $\left(\frac{c_u}{c_o + c_u}\right) \times 100\%$		
			11%	50%	89%
Low	Low	$\frac{NF}{G}$	142.9%	148.5%	151.9%
		$\frac{NF}{E}$	144.1%	148.1%	153.3%
		$\frac{G}{E}$	100.9%	99.7%	101.0%
	Medium	$\frac{NF}{G}$	139.0%	145.8%	149.6%
		$\frac{NF}{E}$	139.1%	145.6%	151.0%
		$\frac{G}{E}$	100.1%	99.9%	100.9%
	High	$\frac{NF}{G}$	129.2%	140.8%	147.3%
		$\frac{NF}{E}$	130.2%	140.0%	147.0%
		$\frac{G}{E}$	100.8%	99.4%	99.8%
High	Low	$\frac{NF}{G}$	111.7%	124.6%	131.0%
		$\frac{NF}{E}$	124.0%	136.3%	141.3%
		$\frac{G}{E}$	112.5%	109.3%	107.9%
	Medium	$\frac{NF}{G}$	110.8%	123.4%	127.8%
		$\frac{NF}{E}$	122.9%	132.1%	135.3%
		$\frac{G}{E}$	110.7%	107.1%	105.9%
	High	$\frac{NF}{G}$	105.5%	121.1%	124.2%
		$\frac{NF}{E}$	113.2%	128.4%	130.4%
		$\frac{G}{E}$	107.3%	106.1%	105.0%