Periodic Review, Push Inventory Policies for Remanufacturing *

B. Mahadevan
Indian Institute of Management
Bangalore
560 076 India
mahadev@iimb.ernet.in

David F. Pyke**
The Tuck School of Business
Dartmouth College
Hanover, NH 03755
USA
603 646-2136
Fax: 603 646-1698
david.pyke@dartmouth.edu

Moritz Fleischmann
Erasmus University Rotterdam
Rotterdam School of Management / Faculteit Bedrijfskunde
PO Box 1738
3000 DR Rotterdam
The Netherlands

September 5, 2002

**Corresponding author
*The authors gratefully acknowledge the helpful comments and suggestions of Professor Ed Silver and of three anonymous referees.

Keywords: Inventory, Remanufacturing, Environment, Heuristics, Reverse Logistics

Abstract

Sustainability has become a major issue in most economies, causing many leading companies to focus on product recovery and reverse logistics. This research is focused on product recovery, and in particular on production control and inventory management in the remanufacturing context. We study a remanufacturing facility that receives a stream of returned products according to a Poisson process. Demand is uncertain and also follows a Poisson process. The decision problems for the remanufacturing facility are when to release returned products to the remanufacturing line and how many new products to manufacture. We assume that remanufactured products are as good as new. In this paper, we employ a "push" policy that combines these two decisions. It is well known that the optimal policy parameters are difficult to find analytically; therefore, we develop several heuristics based on traditional inventory models. We also investigate the performance of the system as a function of return rates, backorder costs and manufacturing and remanufacturing lead times; and we develop approximate lower and upper bounds on the optimal solution. We illustrate and explain some counter-intuitive results and we test the performance of the heuristics on a set of sample problems. We find that the average error of the heuristics is quite low.

1. Introduction

Sustainability has become a major issue for companies and countries as we enter the 21st Century. Several European nations have mandated stringent laws for "product take back" after the product's useful life ends, forcing companies to respond with product redesign, changes in packaging, and creative solutions to the problem of product recovery. Efforts in all these areas can be seen in the automotive, computer, copier, and other industries (VROM (2002); EU (2002)).

Product recovery is an attempt to reuse as much of the product as economically worthwhile, and it takes many forms as highlighted in Figure 1. Many companies currently can remanufacture their products, making them essentially as good as new. For these companies, the stream of returned products – also known as "carcasses" or "cores" – is uncertain. Therefore, they face a two-fold decision problem: First, should they remanufacture the carcasses they have in hand, and if so, when should these be released to the remanufacturing line? Second, should they manufacture new units because of low finished goods inventory and a trickle of carcasses, and if so, how many? In this paper, we utilize a particular policy – a periodic review, "push" policy – which addresses these two decisions.

Figure 1 about here.

Breeze-Eastern is a U.S. company that serves a niche market by producing rescue hoists and cargo hooks for helicopters, construction, logging and other applications.

(Examples of rescue hoists can be seen in the movie "The Perfect Storm.") Rescue hoists are extremely sophisticated and expensive, ranging from $60,000 to $120,000 per system. Cargo hooks are considerably less expensive, ranging from around $3000 to $20,000. Both categories are remanufactured at their FAA certified repair center – the more expensive products generally by prearranged schedule, and the less expensive ones generally whenever they are returned for repair or upgrade. Because of the prearranged schedule of hoists, the remanufacturing facility at Breeze-Eastern builds an inventory of cargo hook carcasses, waiting to release them to the shop floor. Remanufactured cargo hooks are as good as new, and are held in finished goods inventory waiting for sale. However, because the demand rate for finished goods is higher than the return rate of carcasses, Breeze-Eastern also manufactures new cargo hooks. Our analysis is focused on the Breeze-Eastern problem of remanufacturing inventory control. Specifically, we study a slightly simplified version of the remanufacturing of cargo hooks where we follow the Breeze-Eastern case in detail, except that, unlike the real case, we assume constant lead times for both remanufacturing and manufacturing, and we assume equal prices for products, regardless of their source. We explain these simplifying assumptions further in Section 3.

The form of the optimal policy is not known for the general problem we model. Nevertheless, we restrict consideration to a periodic review, push policy for several reasons. First, the pervasiveness of MRP systems, including at Breeze-Eastern, suggests that periodic review is a good fit with the behavior of practitioners. One such example is described in Section 3. Second, van der Laan, Salomon, Dekker, & Van Wassenhove

(1999) argues that push and pull systems are widely used in remanufacturing. And finally we build on analytical work by Inderfurth (1997) that addresses these policies.

Because the optimal policy parameters for our periodic review, push policy are difficult to find analytically, we develop heuristics based on traditional inventory policies. We also investigate the performance of the remanufacturing system as a function of return rates, backorder costs and manufacturing and remanufacturing lead times; and we develop approximate lower and upper bounds on the optimal solution. We illustrate and explain some counter-intuitive results and we test the performance of the heuristics on a set of sample problems.

This paper is organized as follows. In the next section, we review the relevant literature. In section 3, we describe the assumptions, notation and system in detail; then in section 4, we examine the behavior of the total cost function. In section 5, we develop three heuristics and then test them in section 6. Finally, we conclude and discuss future research in section 7.

2. Literature Review

Many authors have addressed inventory management in the context of product recovery and remanufacturing. In this section we summarize the main findings presented in the literature. The selected references highlight significant contributions but are not meant to be exclusive. For more detailed reviews we refer to Silver, Pyke, & Peterson (1998), Chapter 12, and Fleischmann (2001).

Although related models have been proposed as early as the 1960s, inventory control for product recovery and remanufacturing has been receiving growing attention in the past decade with the rise of environmental concern. In addition to numerous theoretical contributions, case studies have been reported on, e.g. for single-use cameras (Toktay, Wein, & Zenios (2000)), medical devices (Rudi & Pyke (2000)), automotive exchange parts (van der Laan (1997)), and electronic equipment (Fleischmann (2001)). The underlying inventory control models share two main distinctive characteristics, namely (i) an autonomous inbound item flow and (ii) two alternative supply options, i.e. product recovery versus 'virgin' procurement. While both of these elements as such are not new in inventory theory it is their specific interrelation that gives rise to novel issues as we discuss below.

Well-established models in the inventory control literature that may help understand the impact of the above characteristics include repairable-item-models (see e.g. Nahmias (1981)) and two-supplier-models (see e.g. Moinzadeh & Nahmias (1988); Moinzadeh & Schmidt (1991)). Yet, neither model class fully captures the setting of product recovery. Repair-models, such as the classical METRIC-model (Sherbrooke (1968)), essentially rely on a closed-loop system structure, where each (defective) item return triggers an immediate demand for a replacement item. In a product recovery setting the correlation between the two item flows tends to be much weaker and mainly reflects the dependence of returns on previous demand. Since the time lag between both processes may be large, many authors claim that, for inventory control purposes, one may even assume independence. Two-supplier-models address the trade-off between procurement costs and lead times. Typically, the models include a slow yet cheap supplier and a faster but also more expensive one. In a product recovery context the reasoning is different. Rather than lead time reduction it is the restricted availability of the (cheaper) recovery channel that calls for an alternative supply source.

Current literature comprises both deterministic and stochastic inventory control models for product recovery environments. Deterministic models can be further subdivided into stationary versus dynamic models. The former correspond to the mindset of the classical economic order quantity (EOQ). As early as in 1967 Schrady proposed an extension to this model that includes item returns (Schrady (1967)). His analysis seeks optimal lot sizes for the recovery channel and 'virgin' procurement, both of which involve fixed costs. More recently, variants to this model have been discussed e.g. by Richter (1996) and Teunter (2001). For the dynamic case, extensions to the classical Wagner-Whitin model have been presented. Beltran & Krass (2002) show that return flows increase the combinatorial complexity of this model. In particular the fundamental zero-inventory-property is lost.

Related stochastic models provide the basis for our investigation. Within this class one may distinguish between periodic review and continuous review approaches. Another important differentiation concerns single versus two-echelon models. In the single echelon case, the analysis is limited to end-item stock, while the two-echelon case involves a more detailed picture of the recovery channel, distinguishing end-item and recoverable stock.

A first stream of research dates back to Whisler (1967) who analyzes the control of a single stock point facing stochastic demand and returns. He shows the optimality of a two-parameter policy that keeps the inventory level within a fixed bandwidth in each period by means of disposal and new supply. Both actions are immediate and the costs are purely linear. Simpson (1978) extends this model to a two-echelon situation. The optimal policy then relies on three critical numbers that control the disposal, remanufacturing, and new supply decision, respectively. Inderfurth (1997) shows that both of the previous results still hold if both supply channels involve the same lead time. For different lead times, though, the growing dimensionality of the underlying Markov model inhibits simple optimal policy structures. Fleischmann & Kuik (1998) provide another optimality result for a single stock point. They show that a traditional (*s, S*) policy is optimal if demand and returns are independent, recovery has the shortest lead time of both channels, and there is no disposal option. Related models have also been analyzed by Kelle & Silver (1989) and Cohen, Nahmias, & Pierskalla (1980).

A parallel stream of research has evolved for continuous review models. Muckstadt & Isaac (1981) consider a single echelon model where the recovery process is modeled as a multi-server queue. They propose a heuristic (*s, Q*) policy for the 'virgin' supply channel and approximate the optimal control parameter values. Van der Laan, Dekker, & Salomon (1996) present an alternative approximation for this model and extend it with a disposal option. Finally, van der Laan et al. (1999) provide a detailed analysis of the corresponding two-echelon model. The authors develop Markovian formulations for several alternative heuristic policies, and compare their performance numerically. In particular, push and pull rules for the recovery channel are contrasted. In line with Inderfurth's results the authors emphasize that lead time differences between both supply channels may severely complicate inventory management in this setting.

In summary, most of the literature on product recovery focuses on the structure of optimal policies for specific cases, while computation of these policies is very time consuming as they involve evaluating large-scale Markov chains. This highlights the fact that practical implementation calls for more efficient evaluation of policy alternatives, and therefore for approximations to the optimal policy. Our paper answers this call. We provide accurate heuristics that can be evaluated almost instantaneously on a spreadsheet.

3. System Description

The system we model is based on the Breeze-Eastern case, but we employ simplifying assumptions to facilitate analysis and insights. We consider a single remanufacturing facility that receives returned carcasses according to a Poisson process with parameter λ_r. This facility maintains a stockpile of returned, but not yet remanufactured, carcasses and a stockpile of finished goods, or serviceable products. (Serviceable products are commonly defined as manufactured or remanufactured finished goods that are ready for sale.) We assume that all returned carcasses are fit for remanufacture. (Alternatively, we can assume that λ_r represents the return rate of carcasses that can be successfully remanufactured. However, we do not consider the time or cost of determining the status of carcasses.) Holding cost is charged on returned, but not yet remanufactured, products at a rate C_{hr}, and on serviceable goods at C_{hs}. Finished goods may be sourced from remanufacturing in lead time L_r, or from manufacturing in lead time L_m. These lead times are constant. We choose to employ this assumption, in spite of departing from the Breeze-Eastern case, to gain clarity on the effect of the two lead times on system performance.

Demand arrives at the serviceable stockpile according to a Poisson process with rate λ_d, and any unmet demand is backordered and charged at a cost of C_b \$/unit. (As an additional benchmark, we also ran tests with a backorder measure of \$/unit/unit-time.) Once again, we depart from the Breeze-Eastern situation by assuming that prices for new and remanufactured units are identical. This assumption allows us to treat demand as a single stream of customers and therefore isolate the remanufacturing process from marketing issues.

We assume that the population of products in the field is quite large so that a particular sale does not influence the rate of returns. This assumption is appropriate for many consumer products, and even for large durable consumer goods, but it is less appropriate for certain military aircraft applications that have a small number of very expensive parts in the field. We also assume that the cost to remanufacture is less than the cost to manufacture, so that there is an economic incentive to avoid scrapping all returned units.

The remanufacturing production process is controlled by a periodic review, push policy that operates as follows. Every R periods, release all carcasses from the returns stockpile into the remanufacturing facility. Let this (stochastic) quantity be denoted Q_r. Furthermore, let I_R be the inventory position of the finished goods stockpile, which we define as the serviceable inventory on hand, less backorders, plus any outstanding (manufacturing or remanufacturing) orders. If, after releasing the remanufacturing batch, I_R is less than the manufacturing order-up-to level S_m, order enough products, Q_m, from the manufacturing facility to bring inventory position up to S_m. (See Figure 2.)

Figure 2 about here.

The total relevant cost (TC) we are considering comprises the sum of holding costs for remanufacturable and serviceable items and backorder costs. Note that remanufacturable holding costs cannot be influenced under the above policy. Yet we include this term in order to compare inventory costs for different return rates. Moreover, note that we do not include variable (re)manufacturing costs, which are again fixed under the above policy. To assess the profitability of remanufacturing, production cost savings should be added to the above cost term.

The inventory management problem we are considering concerns minimizing TC by choosing an appropriate manufacturing order-up-to level S_m. The review period, R, is chosen to manage batch sizes and setup times, perhaps by using an economic lot scheduling problem (ELSP) algorithm or a similar approach. Other considerations, such as the schedule for MRP runs, also influence the choice of R. Therefore, we assume that R is given at this stage. (See McGee & Pyke (1996).)

Notation

λ_r average return rate (units/day)

λ_d average demand rate (units/day)

L_r remanufacturing lead time (days)

L_m manufacturing lead time (days)

n lead time multiplier; $L_m = n * L_r$

C_{hr} holding cost per unit per day of returned, but not remanufactured units ($/unit/day)

C_{hs} holding cost per unit per day of serviceable units, including newly manufactured units ($/unit/day)

C_b backorder cost per unit ($/unit)

R review period (days)

S_m order-up-to level for manufacturing (units)

$\overline{I}_s$ average daily serviceable inventory (units)

$\overline{I}_r$ average daily returned goods inventory (units)

$\overline{B}$ average number of new back orders per replenishment cycle (units)

Finally, for purposes of the experimental design, we define j as follows: $C_b = j*C_{hs}$. Roughly speaking, j is a multiplier, measured in days, that allows us to normalize the serviceable holding cost and vary only the backorder cost.

Solution Methodology

We know from Inderfurth (1997) that for the general inventory system above the complexity of the optimal control policy is prohibitive. From a practical perspective it therefore seems wise to resort to a simple heuristic policy instead. In this context, it should be noted that our suggested push policy reduces to a conventional (*R, S*) policy in the case of a vanishing return rate. Moreover, for the proposed remanufacturing strategy the manufacturing order-up-to policy can be shown to be optimal as long as $L_r \leq L_m$ (see Fleischmann & Kuik, 1998). In the case of $L_r = L_m$ this policy also coincides with Inderfurth's (1997) policy if we disregard the option of disposal. It is worth adding that several numerical studies suggest that disposal is a relevant option only for excessively high return rates (see, e.g. Teunter & Vlachos, 2002); Fleischmann, 2001)). All in all our suggested policy appears to be a natural extension of several well-grounded policies to a domain where optimality is beyond reach.

Once the policy is fixed the decision problem is reduced to choosing an appropriate order-up-to level S_m. As for conventional inventory models, determining optimal parameter values by means of a Markovian analysis does not appear to be attractive in a practical setting. Therefore, we develop several heuristics, which can easily be implemented on a spreadsheet. We test our heuristics by means of simulation, using the off-the-shelf simulation package, PROMODEL (www.promodel.com). Specifically, we set a value for S_m and find the total cost of the system for that policy using the PROMODEL simulation and the total cost equation

$$\text{TC (\$/day)} = C_{hs}\overline{I}_s + C_{hr}\overline{I}_r + C_b\overline{B} / R .$$

We then adjust the value of S_m and find the total cost for this value, searching for the optimal value of S_m. We use analytical work to narrow the simulation search and to develop bounds and heuristics. In this way, we aim at understanding the behavior of the system when the periodic review push policy is employed and to gain insight into the effect of return rates, backorder costs and lead times on system performance.

Parameters for the Experiments

For testing the heuristics, we employ the experimental design shown in Table 1. To test a variety of lead time settings, we vary the manufacturing lead time from ½ to 4 times the remanufacturing lead time. Likewise, to test a wide range of backorder/holding cost ratios, we vary the backorder cost multiplier from 5.7 to 50. We selected a range of 5.7 to 50 for the backorder multiplier on the basis of the behavior of such systems. Recall that our study employs \$/unit as the measure for backorder cost. A backorder multiplier value of 50 is quite high and therefore generates policies that have very few backorders. However, a multiplier below 5.0 generates an optimal safety stock of negative infinity in the traditional inventory model approximation. Similarly, if we use $/unit/unit-time as the measure for backorder cost, we experience similar results for any value of the backorder multiplier below 5.7. Clearly negative infinity safety stock is a nonsensical solution, and therefore is not valid for developing insights. (See Lau, Lau, & Pyke (2000) for further explanations of degenerate inventory policies like this.) Because we use this traditional model for our bounds and heuristics, we set 5.7 as the lower value for the backorder cost multiplier. With a backorder multiplier of 5.7, the solution is to hold very little inventory, allowing us to test both high (j = 50) and low (j = 5.7) inventory cases. The backorder costs in our experiment correspond to fill rates of 98%, 95%, 90% and 85%.

Table 1 about here.

Finally, note that the return rate is always less than the demand rate. In general, we observe three stages in a product life cycle pertaining to remanufacturing. Early in a product's life, few units are in the field, and therefore there are no carcasses to remanufacture. This is captured by the case of $\lambda_r = 0$. In the middle, and hopefully, longest stage of the life cycle, demand exceeds returns and the firm must manufacture and remanufacture to satisfy customers. Late in the life cycle, demand declines and returns increase, implying that the firm will not remanufacture every carcass. A new policy parameter must be introduced – the number of carcasses to dispose of. We leave this latter stage for further research.

4. Behavior of the System

From Figure 3 it is clear that the cost function is quasiconvex (i.e. level sets are convex) or even convex in S_m for the cases shown. This behavior was evident in every case we tested, so we were confident in employing simple linear search techniques for finding the optimal policy. Nevertheless, in a subset of cases we extended the search by significant amounts to be certain we had found the global optimum.

Figure 3 about here.

To further our understanding of the behavior of the cost function, we examined its shape using alternative backorder cost measures ($/unit and $/unit/unit-time). As can be seen from Figure 4, for high inventory, there are few shortages and the cost functions converge. As one would expect, for low inventory, costs increase dramatically, especially for the $/unit/unit-time measure. The $/unit measure has high, but flat, costs for low S_m because below a certain point, all units are backordered – there is no additional penalty cost and there is no additional holding cost savings by reducing S_m. (Figures 4 – 6 illustrate results for $j = 20$; however, conceptually the figures are similar for different backorder costs.)

Figure 4 about here.

Some interesting results are illustrated in Figures 5 and 6. Figure 5 shows the total cost of the system when $L_r < L_m$ and Figure 6 when $L_r > L_m$. First, consider the behavior of the optimal order-up-to level. From Figure 5, the optimal order-up-to level is the largest when the return rate, λ_r, is zero. This is because there are never any returns to process and therefore all demand is met from manufacturing. Recall that the manufacturing order-up-to level, S_m, generates an order from manufacturing for whatever gap is left *after* the remanufacturing batch has been released. When there are returns of 2 per day, and these are remanufactured in half the manufacturing lead time (5 days versus 10 days, as in Figure 5), the optimal order-up-to level decreases. This is because it is faster to meet demand from the remanufacturing process than the manufacturing process, and therefore it is not necessary to provide higher levels of safety stock from the manufacturing side. The pattern continues as the return rate increases. As expected, when the remanufacturing lead time is twice the manufacturing lead time, the pattern is the opposite (as in Figure 6) – the optimal S_m increases as the return rate increases. The push inventory policy pushes carcasses into the remanufacturing facility that takes a longer lead time to remanufacture. Therefore to avoid higher shortage costs the optimal value of S_m increases.

Figures 5 and 6 about here.

Second, consider the total cost for a given S_m. In a periodic push policy, a higher return rate implies that more units are held in both returned and serviceable stockpiles. Hence for high values of S_m the inventory costs in Figure 5 are increasing in λ_r. On the other hand, in the case of moderately low (below optimal) values of S_m, costs are decreasing in λ_r. In this case, there is not enough inventory to service the demand. Furthermore, returns inventory is held for R periods before being released to remanufacturing. Therefore shortage costs are significantly higher even though returned and serviceable inventory holding costs are lower. For very low values of S_m, every unit is backordered, so shortage costs are identical in all cases. Serviceable inventory costs are essentially zero because inventory is sold as soon as it becomes available.

When the remanufacturing lead time is twice the manufacturing lead time (Figure 6), the total cost pattern is exactly the opposite for high and moderately low (above and below the optimal) values of S_m. However, they are the same as before when a given S_m is very low, as shortage costs dominate the other costs in this case. Consider the case of high S_m. At higher return rates the push policy ensures that more demand is met from the slower remanufacturing facility. Consequently, serviceable inventory is less and the carcasses reside longer in a lower cost state. Furthermore, because S_m is so high, there are essentially no shortages. Hence, for high values of S_m, total costs are decreasing in the return rate. When S_m is moderately low, shortage costs are significantly higher when the return rate is higher because of the long remanufacturing lead time. These costs dominate both types of inventory cost, and total costs are higher for larger return rate.

One final graph illustrates the effect of lead time on the system at the optimal S_m. Figure 7 shows that when $\lambda_r = 0$ or when $L_r = 5$, total costs increase as manufacturing lead time increases; however, when $L_r = 2$ and $\lambda_r \neq 0$ total costs at optimality actually decrease when n (the manufacturing lead time multiplier) increases from 1 to 2. In other words, costs may decrease when manufacturing lead time increases. Several observations help explain this counter-intuitive behavior. First, note that when the return rate is zero, the system behaves like a traditional periodic review inventory system: increasing the value of n results in a larger manufacturing lead time and therefore higher costs. It is when the return rate is greater than zero that costs may decrease as n increases. The fundamental argument for why this may happen is a batching one: it is often more desirable to receive smaller batches more frequently than to receive larger batches less frequently. Smaller batches decrease holding costs, at the expense of slightly higher shortage costs.

Figure 7 about here.

Consider the case in which $L_r = 2$ and $\lambda_r = 8$. When $n = 1$, both remanufacturing and manufacturing batches arrive at the same time. In terms of a traditional inventory diagram such as Figure 2, the saw tooth diagram has large spikes when both orders arrive. Now when $n = 2$, the remanufacturing batch arrives in two days, while the manufacturing batch arrives two days later. Therefore, the saw tooth has more frequent, smaller spikes, such as in the second replenishment cycle of Figure 2. Of course, shortage costs increase with the longer manufacturing lead time, but most shortages occur at the end of the review period, R. Thus, from the perspective of reducing shortages, there is little benefit to receiving the manufacturing batch two days earlier. On the other hand, from the perspective of reducing serviceable inventory cost, there is great benefit to delaying the receipt of the manufacturing batch for two days. This same argument explains why the cost is monotonic in n when $L_r = 5$: the best situation for both holding and shortage costs is to have a fast lead time with evenly spaced batches, i.e. $L_r = 5$ and

$L_m = 2.5$ ($n = 0.5$). When $n = 1$, both batches arrive together, increasing both holding and shortage costs. The same is true for $L_m = 10$: the two batches arrive simultaneously, except that the manufacturing batch arrives one cycle later than the remanufacturing batch. We conclude by noting that a similar, counterintuitive lead time effect has been reported by van der Laan (1997) for a continuous review model. In that case, the difficulty is to define an appropriate inventory position for coordinating manufacturing and remanufacturing decisions.

5. Bounds on S_m^* and Heuristics for Solving for S_m^*

In this section we introduce upper and lower bounds on the optimal solution and then develop three heuristics for obtaining the approximately optimal order-up-to level. We test the tightness of the bounds and the accuracy of the heuristics in the next section. The purposes of the bounds are twofold. First, they provide intuition about the behavior of the system, and second, they help narrow the search for the optimal. Because of the use of approximate formulas in the bounds, however, they are sometimes violated, and they should be denoted *approximate* bounds. Therefore, if one is looking for the true optimal parameters, rather than employing a heuristic, it is wise to search several values of S_m beyond the bounds. Of course, the quasiconvexity of the cost function simplifies the search.

The upper bound on S_m is based on the most *pessimistic* assumptions: ignore returned units, thereby increasing net demand, and employ the maximum lead time. To quickly compute the bound, we approximate Poisson demand with a normal distribution, and we use an approximate traditional periodic review inventory model from Silver et al.

(1998). (Note that the normal distribution is a very accurate approximation for the Poisson when mean lead time demand is larger than about 10. In our experiments, the mean demand over the review period and lead time is at least 70.) The assumptions of larger demand and lead time will cause the approximate model to choose a larger S_m. More formally, to compute the upper bound:

Upper Bound

1. Define a pessimistic average lead time demand based on the largest lead time and no returns: Let $\lambda'_d = \lambda_d \max(R + L_r, R + L_m)$.
2. Compute the standard deviation of the approximate lead time demand distribution for use in the normal approximation: Let $\sigma'_d = \sqrt{\lambda'_d}$
3. Compute a safety factor based on an approximate traditional inventory model using he costs of our system: Use the continuous review formula 7.23 from Silver et al. (1998), $p_{u\geq}(k) = \dfrac{Qr}{DB_2}$, where Q is the order quantity, r is the annual carrying charge, D is the annual demand, and B_2 is the fractional charge per unit short. $p_{u\geq}(k)$ is the probability that a unit normal random variable is greater than or equal to k. Now in the notation of this paper, translated to a periodic review case, Q is the average order quantity, or 5 days of demand = $\lambda_d R$. Let v be the value of the item. Then, $B_2 = C_b/v$, while $C_b = C_{hs} * j$, and $C_{hs} = rv/250$, assuming 250 days per year. Therefore, in our case, $p_{u\geq}(k) = R / j$. Find k from unit normal tables, or in Excel, NORMSINV(1 – $p_{u\geq}(k)$).

4. Compute the optimal order-up-to level for the pessimistic parameters: $S_m^{UB} = \lceil \lambda_d' + k\sigma_d' \rceil$, where $\lceil x \rceil$ rounds x up to the next highest integer.

Note that formula 7.23 Silver, et al. (1998) is an approximation that degenerates as backorder cost decreases (Lau et al. (2000)). Therefore, for low backorder cost, the upper bound is more likely to be violated.

The lower bound is based on the most *optimistic* assumptions: a lower demand rate and a shorter lead time, implying that the choice of the order-up-to level will be smaller. Then use the normal approximation and the same approximate inventory model:

Lower Bound

1. Compute the optimistic average lead time demand based on the smaller of the two lead times and a reduced demand rate: Let $\lambda_d'' = [\min(R+L_r, R+L_m)]\max(\lambda_d - \lambda_r, \lambda_r)$. The max (·) term allows us to reduce the demand rate from the actual parameters but maintain a tighter bound than would min(λ_d, λ_r).
2. Compute the standard deviation: Let $\sigma_d'' = \sqrt{\lambda_d''}$
3. As in the upper bound: $p_{u\geq}(k) = R / j$. Find k from unit normal tables, or in Excel, NORMSINV(1 − $p_{u\geq}(k)$).
4. $S_m^{LB} = \lfloor \lambda_d'' + k\sigma_d'' \rfloor$.

One final note about the lower bound: when the lower bound is very tight and the lead time in step 1 is fractional, we observe that the optimal solution is actually lower than the lower bound. The reason is that the simulation treats days as discrete events, whereas the approximation does not. To remedy this situation, we actually modified step 1 as follows: Let $\lambda_d'' = \lfloor \min(R+L_r, R+L_m) \rfloor \max(\lambda_d - \lambda_r, \lambda_r)$, where $\lfloor x \rfloor$ is the integer part of x. Both bounds can be computed on a spreadsheet in seconds.

Narrowing the Search

As might be expected, total costs at the optimal S_m increase as backorder costs increase. These backorder cost effects, in conjunction with the insights from Figures 5 and 6, allow us to devise a strategy to improve the efficiency of the simulation experiments by dynamically updating bounds when running the entire experimental design, thereby saving significant computer time. The rules are as follows.

1. When moving from a lower backorder multiplier (j_1) to a higher one (j_2), the lower bound for the latter case, LB(j_2), equals the optimal S_m for the former case – $S_m^*(j_1)$.
2. When moving from a lower return rate (λ_{r1}) to a higher return rate (λ_{r2})
    a. If $L_r > L_m$, then LB(λ_{r2}) = $S_m^*(\lambda_{r1})$
    b. If $L_r < L_m$, then UB(λ_{r2}) = $S_m^*(\lambda_{r1})$
    c. If $L_r = L_m$, then $S_m^*(\lambda_{r2}) = S_m^*(\lambda_{r1})$

Heuristics

The first two heuristics employ the same traditional periodic review approximate inventory model. The first heuristic uses a weighted average lead time in conjunction with actual demand. Returned units, λ_r per period, are available for sale in L_r periods, while manufactured units represent the remainder (λ_d – λ_r) and are available in L_m

periods. The lead time for the heuristic is the demand-weighted average of these. The heuristic value of S_m is then computed in the same way as the bounds.

Heuristic 1

1. Let $\lambda_d''' = \left[R + \frac{L_m(\lambda_d - \lambda_r) + L_r\lambda_r}{\lambda_d} \right] \lambda_d$.
2. Let $\sigma_d''' = \sqrt{\lambda_d'''}$
3. $p_{u\geq}(k) = R\ /\ j$. Find k from unit normal tables, or in Excel, NORMSINV(1 – $p_{u\geq}(k)$).
4. $S_m^{\ 1}$ = ROUND($\lambda_d''' + k\sigma_d'''$, 0).

The second heuristic uses similar logic but uses the sum of two order-up-to levels – one based on returns only and the other based on the remaining demand met from manufacturing.

Heuristic 2

1. Let $\lambda_r' = (R + L_r)\lambda_r$.
2. Let $\sigma_r' = \sqrt{\lambda_r'}$
3. $p^1{}_{u\geq}(k^1) = R\ /\ j$. Find k^1 from unit normal tables, or in Excel, NORMSINV(1 – $p^1{}_{u\geq}(k^1)$).
4. $S_m^{\ 2r} = \lambda_r' + k^1\sigma_r'$.
5. Let $\lambda_m' = (R + L_m)(\lambda_d - \lambda_r)$.

6. Let $\sigma'_m = \sqrt{\lambda'_m}$

7. $p^2{}_{u\geq}(k^2) = R / j$. Find k^2 from unit normal tables, or in Excel, NORMSINV(1 $-\ p^2{}_{u\geq}(k^2)$).

8. $S_m^{\ 2m} = \lambda_m{}' + k^2\sigma'_m$.

9. $S_m^{\ 2}$ = ROUND($S_m^{\ 2r} + S_m^{\ 2m}$, 0).

Both of these heuristics rely on approximating the two alternative supply channels (i.e. manufacturing and remanufacturing) by a single aggregated channel. In a third alternative we model both channels separately. To this end, note that the inventory level in our system may have two 'peaks' per review cycle, corresponding to the arrival of a remanufacturing batch and a manufacturing batch (see Figure 2). The safety stocks corresponding to both of these epochs can be approximated in much the same way as in a traditional (R, S) system (see Silver et al. (1998)). To be specific, let SS_r and SS_m denote the expected net stock just before arrival of a remanufacturing batch and a manufacturing batch, respectively. Moreover, for the time being assume that $L_r < L_m \leq R$. Ignoring a potential excess of the order-up-to level S_m we can then approximate SS_r by

$$SS_r \ \approx\ \mathrm{E}[S_m - X_r], \tag{1}$$

where X_r is distributed as demand during a time interval of length $R + L_r$. Letting μ_r := $\mathrm{E}[X_r] = \lambda_d(R+L_r)$ and σ_r^2 := $\mathrm{Var}[X_r] = \lambda_d(R+L_r)$ and using a normal demand approximation, the expected shortage at this epoch is approximated by $\sigma_r\, G((S_m - \mu_r)/\sigma_r)$, where $G(.)$ denotes the usual unit normal loss integral. Analogously, we get

$$SS_m \ \approx\ \mathrm{E}[S_m - X_m + Q_r], \tag{2}$$

where X_m is distributed as demand during a time interval of length $R + L_m$ and Q_r as the number of items returned during one review period, and both distributions are independent. The expected shortage just before arrival of a manufacturing batch therefore approximately equals $\sigma_m G((S_m - \mu_m)/\sigma_m)$, with $\mu_m := E[X_m - Q_r] = \lambda_d(R+L_m) - \lambda_r R$ and $\sigma_m^2 := Var[X_m - Q_r] = \lambda_d(R+L_m) + \lambda_r R$. Putting all of these terms together yields the following approximation of the expected holding and shortage costs per review cycle (*ETCR*)

$$ETCR(S_m) \approx C_b [\ \sigma_r G((S_m-\mu_r)/\sigma_r) + \sigma_m G((S_m-\mu_m)/\sigma_m)] + C_{hs} (\ S_m + c)\ R\ , \qquad (3)$$

with some constant c that is independent of S_m. From this we get the following first order condition for S_m

$$p_{u\geq}((S_m-\mu_r)/\sigma_r) + p_{u\geq}((S_m-\mu_m)/\sigma_m) = R\ /\ j\ . \qquad (4)$$

This equation differs from the one for a standard (*R, S*)-system in that the left-hand side is a sum of two stockout probabilities, both depending on S_m. Therefore, (4) cannot be solved for S_m analytically. However, since both probabilities are strictly decreasing in S_m the equation can easily be solved numerically, e.g. using Excel's GOALSEEK function. Alternatively, (4) may be used as a basis for more simplistic heuristics, e.g. by replacing the two individual probabilities by a single value $p_{u\geq}(.)$ at some intermediate point or by ignoring one of the two terms.

The above reasoning goes through for the more general lead time conditions as introduced in Section 2 if we adjust X_r and X_m to the net demand during the corresponding relevant time intervals for both supply channels. We summarize our third heuristic in its general form as follows.

Heuristic 3

1. Let n denote the number of remanufacturing batches arriving before the manufacturing batch ordered in the current period, i.e. $n = \lceil L_m / R \rceil$ if remanufacturing orders arrive before manufacturing orders in each review cycle and $n = \lceil L_m / R \rceil$ -1 otherwise.
2. Let $\mu_r = \lambda_d (nR+L_r) - \lambda_r R (n-1)$.
3. Let $\sigma_r = \sqrt{\lambda_d (nR + L_r) + \lambda_r R(n-1)}$.
4. Let $\mu_m = \lambda_d (R+L_m) - \lambda_r Rn$.
5. Let $\sigma_m = \sqrt{\lambda_d (R + L_m) + \lambda_r Rn}$.
6. Find S_m numerically from $p_{u\geq}((S_m-\mu_r)/\sigma_r) + p_{u\geq}((S_m-\mu_m)/\sigma_m) = R / j$.

6. Tightness of the Bounds and Performance of the Heuristics

The experimental design for testing the bounds and heuristics was given in Table 1. In this section we report the results of the experiment, testing the tightness of the bounds and the performance of the heuristics. Table 2 contains the optimal values of S_m as well as the lower and upper bounds in the format (*UB*, *LB*) *Optimal*.

Table 2 about here.

First, observe that the lower bound is very tight for $n \leq 1$ and $\lambda_r = 0$. Otherwise, the bound is not particularly tight. When the return rate is zero, the “reduced” demand rate equals the actual demand rate, so we should expect a tight bound. However, if one lead time is considerably longer than the other, regardless of the return rate, using the lower lead time for the bound is quite optimistic and the bound is not tight. Nevertheless, it does restrict the search for the optimal by at least 50% over the case with no lower bound.

The upper bound reveals some interesting results. First, when $L_m = L_r$ and $j \geq 10$, the optimal solution is at the upper bound (within the error introduced by simulation and the approximation). When the lead times are identical, an order is placed at review epochs and either source of the product, manufacturing or remanufacturing, makes the product available $L_m = L_r$ days later. Therefore, returns have no influence on the choice of S_m, and the simple approximate inventory model used in the bound is close to optimal. Second, there are a number of cases in which the optimal solution is actually greater than the upper bound. Some of these are due to simulation error and minor errors introduced by the approximation. Note, for instance the case of $L_r = 2$, $\lambda_r = 0$, and $n = 4$. When j increases from 10 to 50, the difference between the optimal S_m and the upper bound decreases from 3 units to 1 unit. It is well known that the inventory approximation assumes relatively few backorders and that the approximation degrades as the backorder level increases. (See Silver et al. (1998), p. 253, and Lau et al. (2000).) Thus, when backorder cost is very low ($j = 5.7$), the inventory approximation degenerates rapidly and the optimal solution is actually greater than the upper bound.

Fortunately, we observed that the cost effect of violating the upper bound is minor. If one used the upper bound in place of the optimal S_m in cases in which the upper bound is violated, the average cost error is only 1.58%, although there was one case with an error of 11% and a second case with an error of 7%. All others had errors of 4% or less, and most were less than 1%. The lesson from this exercise is that, as in many cases in inventory theory, the models and bounds apply best when the parameters are such that the backorder level is low. As is clear from the table, we searched beyond the bounds to test their performance and to be certain of finding the optimal solution, subject to simulation error. However, managers should use the bounds with care when the optimal solution is to hold very little inventory due to a very small backorder cost. Fortunately, a very small backorder cost is quite rare in practice.

Accuracy of the Heuristics

The first two heuristics perform very well as can be seen from Table 3. Heuristic 1, based on weighted lead times, has an overall average cost error of 3.27%, measured as $\frac{TC(S_m^1)-TC(S_m^*)}{TC(S_m^*)}$, where $TC(S_m^*)$ is the total cost at the optimal order-up-to level. Heuristic 2, based on the sum of two order-up-to levels, has an overall average cost error of 5.96%, measured similarly. Heuristic 3 outperforms both other heuristics, achieving remarkably smaller average (0.44%) and maximum errors (see Tables 3 and 4). This performance is very good news indeed. It appears that managers can use simple approximate models from standard inventory theory to solve complex problems in the remanufacturing environment.

Table 3 about here.

Not all the news is good, however. Table 4 shows maximum errors for the three heuristics. Maximum errors can be quite high for the Heuristics 1 and 2, especially for a return rate of 4. The average error is still very small, however, *even for* $\lambda_r = 4$ – likely smaller than the effect caused by inaccuracies in most real input data. However, let us understand the reasons for the few cases of poor performance. An examination of the detailed data (available from the authors) shows that the largest errors occur when $L_r = 2$ and $L_m = 4$, or $L_r = 5$ and $L_m = 2.5$, and of course $\lambda_r = 4$. Recall the discussion in Section 4 explaining how total cost could decrease when manufacturing lead time increases. The ideal situation, from a cost perspective, is to have many small batches arriving frequently, which is precisely what happens in the large-error cases we are examining; i.e. roughly half of the demand is met from each source, and the batches arrive in regular intervals (about every 2 days, or every 2.5 days). Heuristics 1 and 2 employ averages, in effect assuming that the batches arrive together, thereby creating larger inventory spikes, and thus requiring a larger-than-optimal S_m. Therefore, the errors can become quite large. Heuristic 3 circumvents this pitfall by explicitly recognizing two batch arrivals per review period. The results in Tables 3 and 4 illustrate that this approach reduces the error significantly.

Table 4 about here.

When the return rate is 0 or 8, the errors of Heuristics 1 and 2 are smaller because either manufacturing or remanufacturing is the dominant source of inventory, so the actual inventory spikes are *not* of the same relative magnitude, and the approximate, average model is more accurate.

7. Conclusions and Future Research

In this paper we analyzed an inventory system with remanufacturing and manufacturing. Modeling the system using simulation, we observed the quasiconvexity of the objective function in the decision variable, and we saw some unusual behavior, such as costs decreasing when lead times increase. We developed bounds and heuristics based on traditional approximate inventory models that can be calculated easily on a spreadsheet. Each is based on simple, intuitive adjustments to the parameters of the traditional model. The two first approaches rely on an approximation of the manufacturing and remanufacturing sources by a single aggregate channel. The third approach explicitly considers the impact of both channels separately The performance of all the three heuristics is quite good on average, with average total cost errors of 3.27%, 5.96%, and 0.44% respectively. Maximum errors can be significant for the first two heuristics if both supply channels supply similar volumes and batches arrive equally spaced during the review period. In this case, modeling both channels separately yields better performance.

It seems valuable for future research to try and extend the analysis presented in this paper to other remanufacturable inventory models. In particular, addressing remanufacturing-pull models appears to be worthwhile. A pull model triggers a remanufacturing batch only if serviceable inventory falls below a certain threshold. In contrast to the push policy examined in this paper, a pull policy may postpone remanufacturing activities by keeping excess carcasses in stock beyond a single review period. Additional research should also focus on formulating a policy framework for the entire life cycle of the product, from the new product stage with few returns, to the end-of-life stage with returns outpacing demand. In the latter case, the inventory policy must consider disposal of returned units, thereby adding an additional parameter.

References

Beltran, J. L., & Krass, D., 2002. Dynamic lot sizing with returning items and disposals, IIE Transactions, 34, 437-448.

Cohen, M. A., Nahmias, S., & Pierskalla, W. P., 1980. A Dynamic Inventory System with Recycling, Naval Research Logistics, 27 (2) 289-296.

EU, 2002. Waste Electrical and Electronic Equipment, Communication 347 (2000). http://europa.eu.int/comm/environment/docum/index.htm. (01-16-02).

Fleischmann, M. 2001. Quantitative Models for Reverse Logistic, Lecture Notes in Economics and Mathematical Systems 501, Springer-Verlag, Berlin.

Fleischmann, M. & Kuik, R., 1998. On Optimal Inventory Control with Stochastic Item Returns, Management Report Series 21-1998, Erasmus University Rotterdam, The Netherlands.

Inderfurth, K., 1997. Simple Optimal Replenishment and Disposal Policies for a Product Recovery System with Leadtimes, OR Spektrum, 19 (2) 111-122.

Kelle, P., & Silver, E. A., 1989. Purchasing Policy of New Containers Considering the Random Returns of Previously Issued Containers, IIE Transactions, 21, 349-354.

Lau, A. H.-L., Lau, H.-S., & Pyke, D. F., 2000. Degeneracies in Order-Quantity, Reorder-Point Inventory Models, forthcoming, Naval Research Logistics.

McGee, V. E., & Pyke, D. F., 1996. Periodic Production Scheduling at a Fastener Manufacturer, International Journal of Production Economics, 46-47, 65-87.

Moinzadeh, K., & Nahmias, S., 1988. A Continuous Review Model for an Inventory System with Two Supply Modes, Management Science, 34 (6) 761-773.

Moinzadeh, K., & Schmidt, C. P., 1991. An (S-1, S) Inventory System With Emergency Orders, Operations Research, 39 (2) 308-321.

Muckstadt, J. A., & Isaac, M. H., 1981. An Analysis of Single Item Inventory Systems with Returns, Naval Research Logistics, 28 (2) 237-254.

Nahmias, S., 1981. Managing Repairable Item Inventory Systems: A Review. In: L. B. Schwarz (Ed.), Multi-Level Production/Inventory Control Systems: Theory and Practice, North-Holland, Amsterdam, The Netherlands, pp. 253-277.

Richter, K., 1996. The EOQ Repair and Waste Disposal Model with Variable Setup Numbers, European Journal of Operational Research, 95 (2) 313-324.

Rudi, N., & Pyke, D. F., 2000. Product Recovery at the Norwegian Health Insurance Administration, INTERFACES, 30 (3) 166-179.

Schrady, D., 1967. A Deterministic Inventory Model for Reparable Items, Naval Research Logistics, 14, 391-398.

Sherbrooke, C. C., 1968. METRIC - A Multi-Echelon Technique for Recoverable Item Control, Operations Research, 16 (1) 103-121.

Silver, E. A., Pyke, D. F., & Peterson, R., 1998. Inventory Management and Production Planning and Scheduling (3rd Edition.), John Wiley & Sons, New York.

Simpson, V. P., 1978. Optimum Solution Structure for a Repairable Inventory Problem, Operations Research, 26 (2) 270-281.

Teunter, R. H., 2001. Economic Ordering Quantities for Remanufacturable Item Inventory Systems, Naval Research Logistics, 48 (6) 484-495.

Teunter, R. H., & Vlachos, D., 2002. On the necessity of a disposal option for returned items that can be remanufactured, International Journal of Production Economics, 75, 257-266.

Thierry, M., Salomon, M., Van Nunen, J., & Van Wassenhove, L. V., 1995. Strategic Issues in Product Recovery Management, California Management Review, 37 (2) 114-135.

Toktay, L. B., Wein, L. M., & Zenios, S. A., 2000. Inventory Management of Remanufacturable Products, Management Science, 46 (11) 1412-1426.

van der Laan, E., 1997. The Effects of Remanufacturing on Inventory Control. Ph.D. Dissertation, Erasmus University, Rotterdam.

van der Laan, E., Dekker, R., & Salomon, M., 1996. Product remanufacturing and disposal: A numerical comparison of alternative control strategies, International Journal of Production Economics, 45, 489-498.

van der Laan, E., Salomon, M., Dekker, R., & Van Wassenhove, L., 1999. Inventory control in hybrid systems with remanufacturing, Management Science, 45 (5) 733-747.

VROM., 2002. Dutch Ministry of Housing, Spatial Planning, and the Environment, http://www.vrom.nl/pagina.html?id=1402 . (01-16-02).

Whisler, W. D., 1967. A Stochastic Inventory Model for Rental Equipment, Management Science, 13 (9) 640-647.

λ_r	0, 4, 8 units per day
λ_d	10 units per day
L_r	2, 5 days
L_m	$L_m = n * L_r$
n	0.5, 1, 2, 4
C_{hr}	0.4 $/unit/day
C_{hs}	0.8 $/unit/day
C_b	$C_b = j * C_{hs}$ $/unit
j	5.7, 10, 20, 50
R	5 days

Table 1: Experimental Design

LT mult (n)	Backorder mult (j)	L_r = 2 days			L_r = 5 days		
		$\lambda_r = 0$	$\lambda_r = 4$	$\lambda_r = 8$	$\lambda_r = 0$	$\lambda_r = 4$	$\lambda_r = 8$
0.50	5.70	(61,51) 56	(61,29) 56	(61,39) 61	(89,60) 65	(89,34) 73	(89,47) 85
	10.00	(70,60) 61	(70,36) 61	(70,48) 65	(100,70) 71	(100,42) 77	(100,56) 91
0.50	20.00	(76,65) 65	(76,40) 65	(76,52) 68	(107,75) 77	(107,46) 80	(107,61) 96
	50.00	(81,69) 71	(81,43) 71	(81,56) 72	(113,80) 82	(113,50) 85	(113,65) 102
1.00	5.70	(61,60) 65	(61,34) 65	(61,47) 65	(89,88) 95	(89,51) 95	(89,69) 95
	10.00	(70,70) 71	(70,42) 71	(70,56) 71	(100,100) 102	(100,60) 102	(100,80) 102
1.00	20.00	(76,75) 77	(76,46) 77	(76,61) 76	(107,106) 108	(107,65) 108	(107,86) 107
	50.00	(81,80) 82	(81,50) 82	(81,65) 81	(113,112) 114	(113,69) 114	(113,91) 113
2.00	5.70	(79,60) 85	(79,34) 74	(79,47) 66	(136,88) 145	(136,51) 125	(136,69) 105
	10.00	(90,70) 92	(90,42) 77	(90,56) 71	(150,100) 153	(150,60) 133	(150,80) 113
2.00	20.00	(97,75) 97	(97,46) 82	(97,61) 76	(159,106) 160	(159,65) 140	(159,86) 121
	50.00	(103,80) 103	(103,50) 86	(103,65) 82	(166,112) 166	(166,69) 147	(166,91) 129
4.00	5.70	(117,60) 125	(117,34) 100	(117,47) 75	(232,88) 245	(232,51) 185	(232,69) 124
	10.00	(130,70) 133	(130,42) 106	(130,56) 83	(250,100) 255	(250,60) 195	(250,80) 135
4.00	20.00	(138,75) 139	(138,46) 111	(138,61) 90	(261,106) 264	(261,65) 204	(261,86) 146
	50.00	(145,80) 146	(145,50) 117	(145,65) 97	(271,112) 272	(279,69) 213	(271,91) 155

Table 2: Bounds and Optimal Values

Legend: (UB, LB) Optimal

Back order multiplier	Pure manufacturing	Heuristic 1			Heuristic 2			Heuristic 3		
	λ_r = 0	λ_r = 4	λ_r = 8	Overall	λ_r = 4	λ_r = 8	Overall	λ_r = 4	λ_r = 8	Overall
5.7	4.86	2.91	1.70	2.31	8.79	5.79	7.29	0.77	1.55	1.16
10	1.01	3.10	1.00	2.05	3.10	1.00	2.05	0.28	0.18	0.23
20	0.74	4.56	2.45	3.51	7.70	3.77	5.74	0.32	0.08	0.20
50	0.67	6.10	4.31	5.21	11.71	5.78	8.75	0.25	0.11	0.18
All cases	1.82	4.17	2.37	3.27	7.83	4.09	5.96	0.41	0.48	0.44

Table 3: Performance of the Heuristics: Average Percentage Cost Error

Back order multiplier	Pure manufacturing	Heuristic 1		Heuristic 2		Heuristic 3	
	$\lambda_r = 0$	$\lambda_r = 4$	$\lambda_r = 8$	$\lambda_r = 4$	$\lambda_r = 8$	$\lambda_r = 4$	$\lambda_r = 8$
5.7	11.02	6.03	2.76	14.12	8.09	2.02	3.99
10	2.35	15.19	2.55	15.19	2.55	0.54	0.46
20	4.06	22.39	7.21	32.34	14.27	1.79	0.40
50	3.94	27.32	11.17	42.59	18.44	1.52	0.56

Table 4: Performance of the Heuristics: Maximum Percentage Cost Error

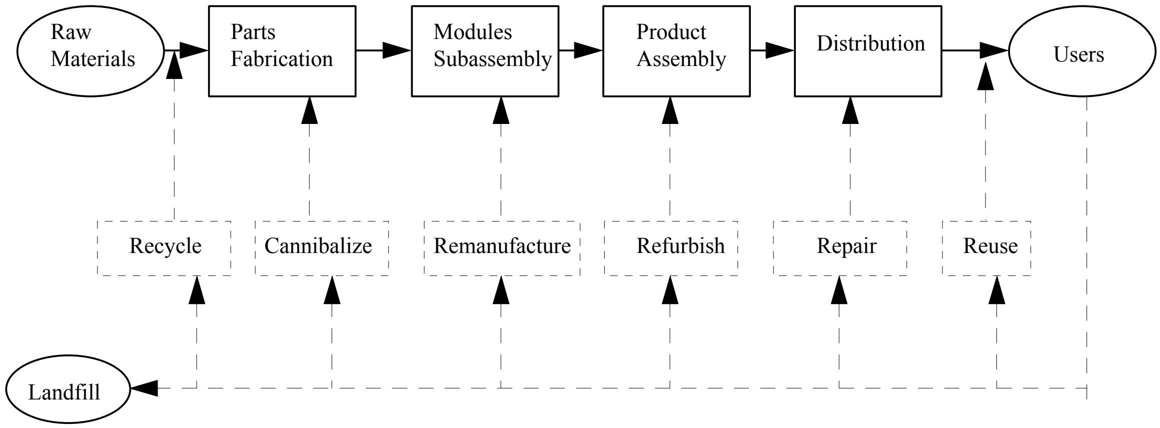

Figure 1: Product Recovery Options (adapted from Thierry, Salomon, Van Nunen, & Van Wassenhove (1995))

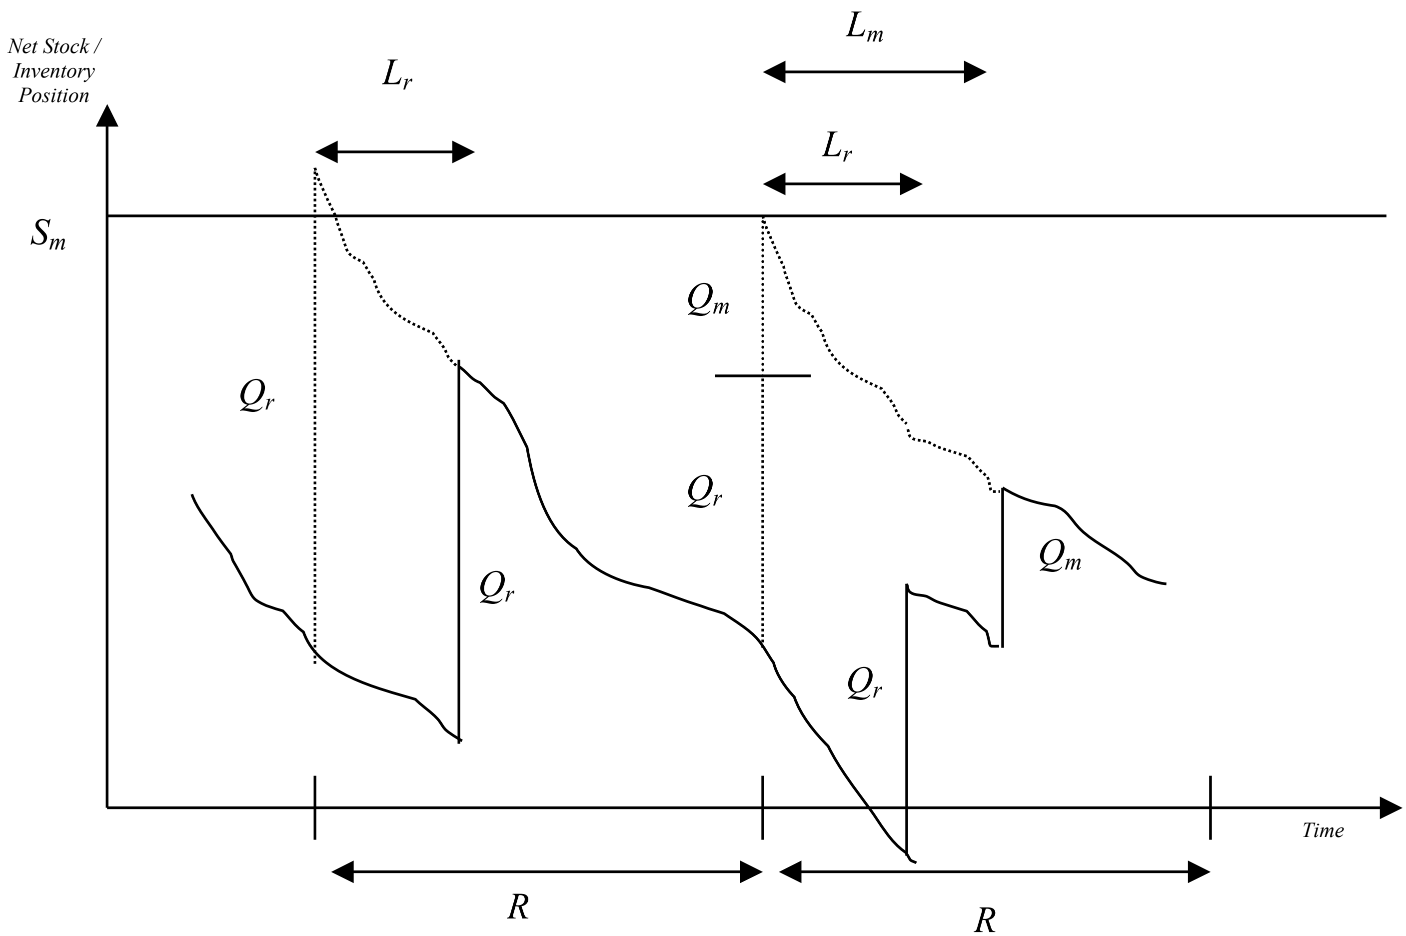

Legend:

.... Inventory position

———— Net stock or both the inventory position and the net stock (if they are equal)

Figure 2: Behavior of the System

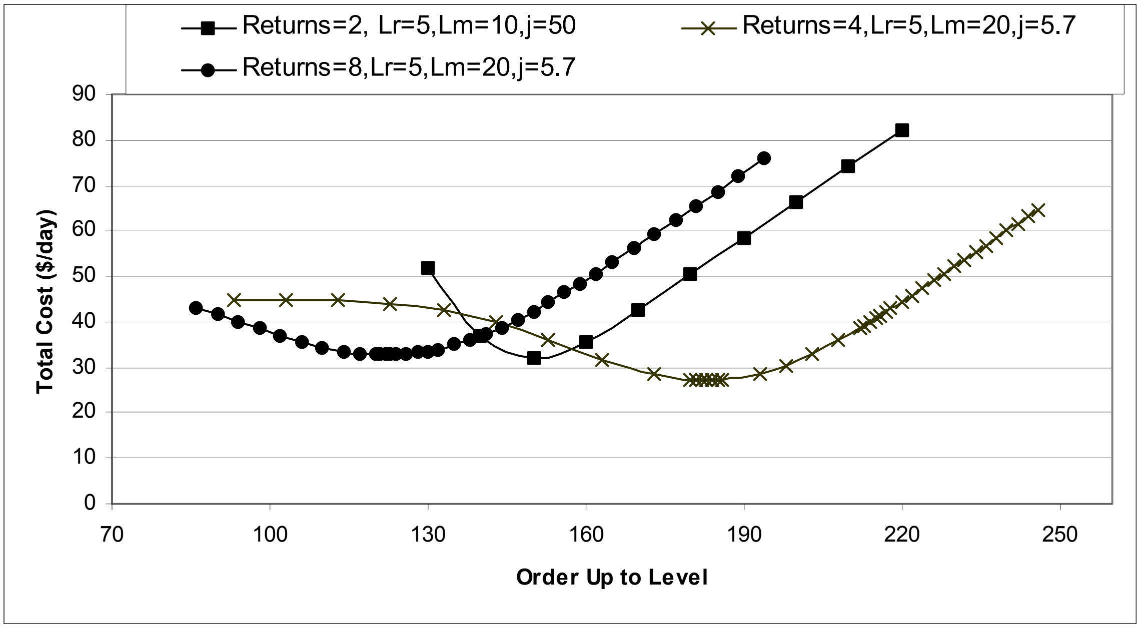

Figure 3: Shape of the cost function in S_m

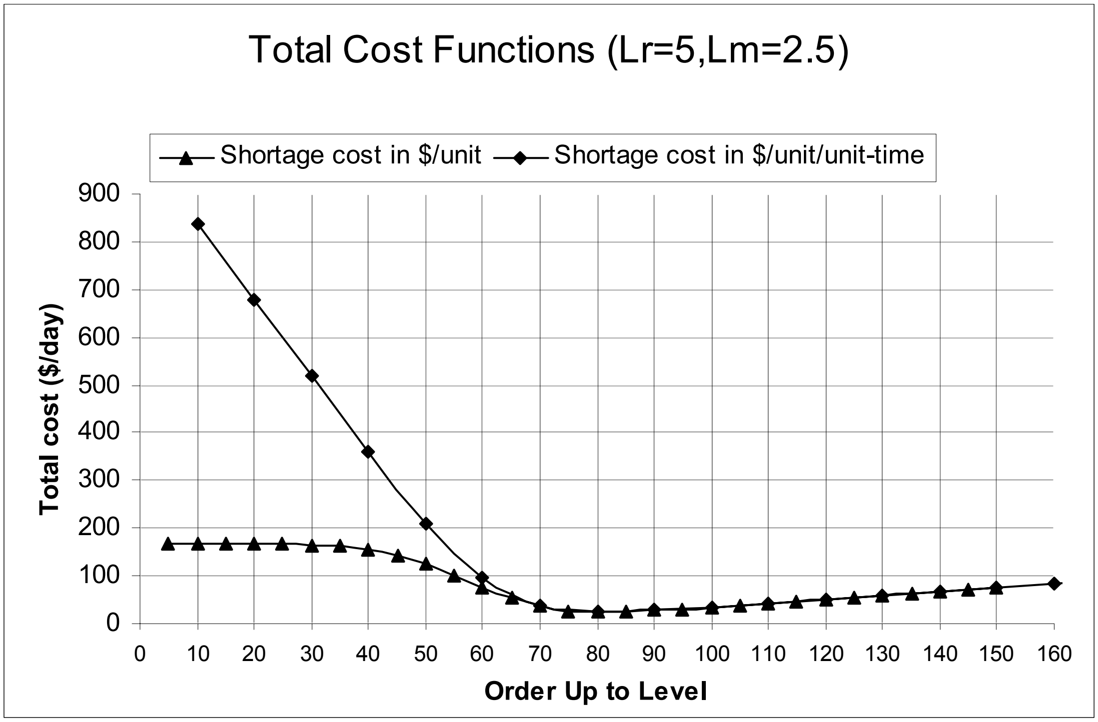

Figure 4: Cost Functions for the Two Shortage Cost Measures (j = 20)

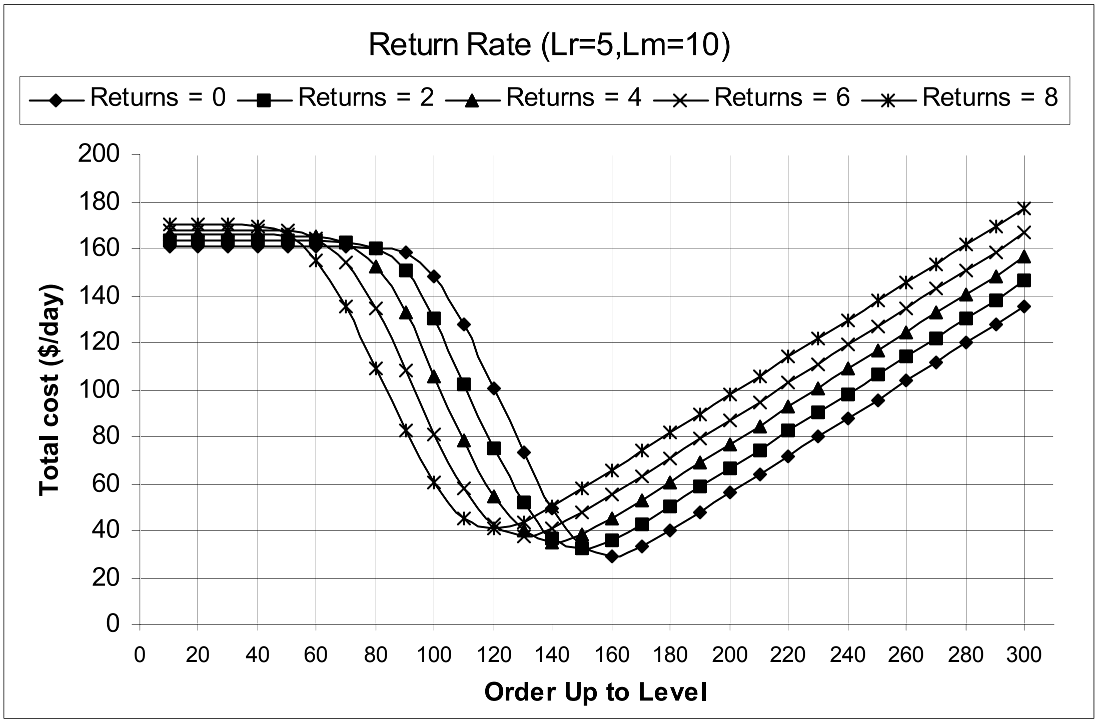

Figure 5: Effect of Return Rate – ($L_r < L_m$, $j = 20$)

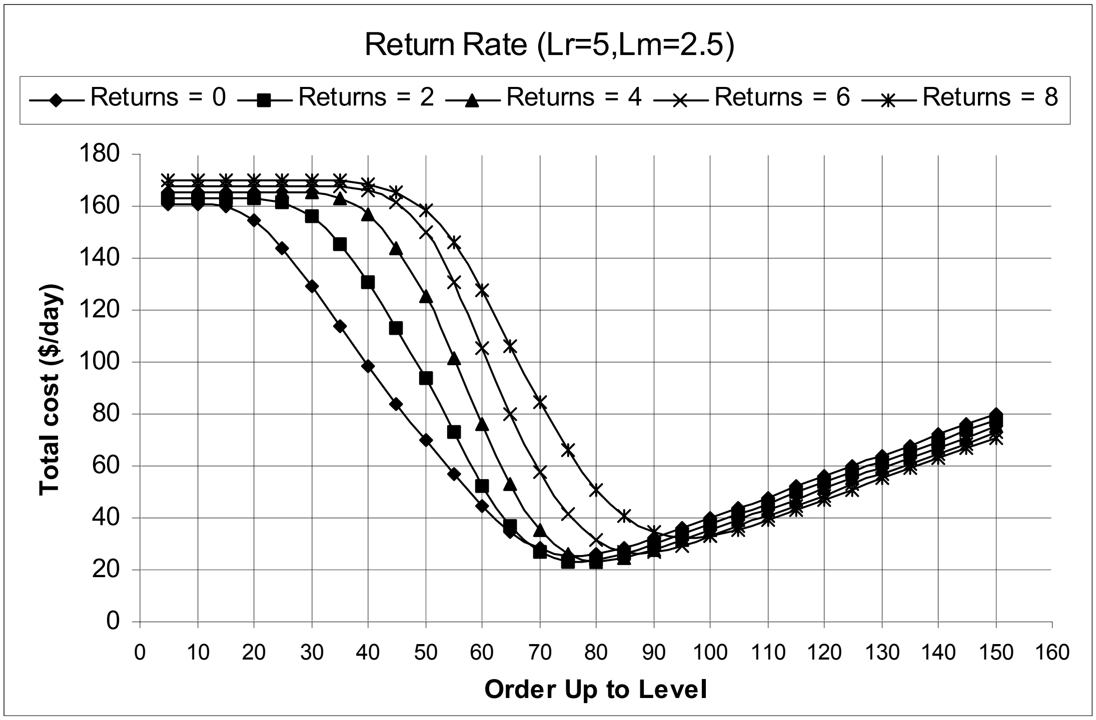

Figure 6: Effect of Return Rate – ($L_r > L_m$, $j = 20$)

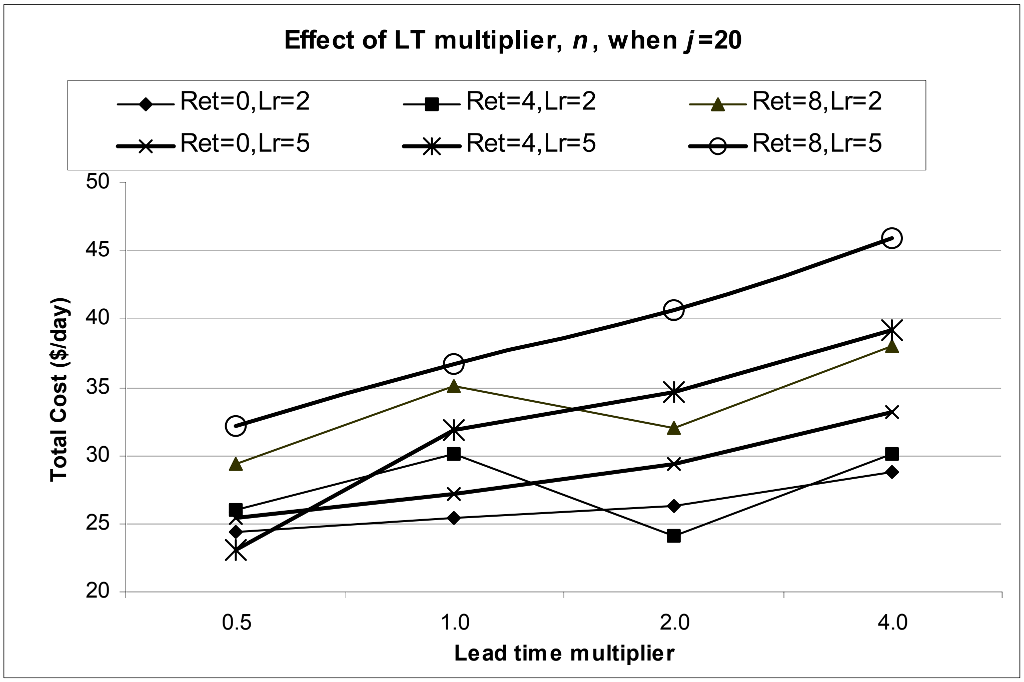

Figure 7: Effect of Lead Time Multiplier at the Optimal S_m