Tilburg University

Last Time Buy and Control Policies With Phase-Out Returns

Krikke, H.R.; van der Laan, E.

Publication date:
2009

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Krikke, H. R., & van der Laan, E. (2009). *Last Time Buy and Control Policies With Phase-Out Returns: A Case Study in Plant Control Systems*. (CentER Discussion Paper; Vol. 2009-66). Organization.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

No. 2009–66

**LAST TIME BUY AND CONTROL POLICIES WITH PHASE-OUT
RETURNS: A CASE STUDY IN PLANT CONTROL SYSTEMS**

By Harold Krikke, Erwin van der Laan

August 2009

ISSN 0924-7815

Last Time Buy and control policies with phase-out returns: a case study in plant control systems

Harold Krikke^a and Erwin van der Laan

a. Corresponding author, CentER, Tilburg University,
PO box 90153, 5000 LE, Tilburg, The Netherlands, krikke@uvt.nl

b. RSM Erasmus University,
Department of Decision and Information Sciences,
PO box 1738, 3000 DR, Rotterdam, The Netherlands, elaan@rsm.nl

Abstract: This research involves the combination of spare parts management and reverse logistics. At the end of the product life cycle, products in the field (so called installed base) can usually be serviced by either new parts, obtained from a Last Time Buy, or by repaired failed parts. This paper, however, introduces a third source: the *phase-out* returns obtained from customers that replace systems. These returned parts may serve other customers that do not replace the systems yet. Phase-out return flows represent higher volumes and higher repair yields than failed parts and are cheaper to get than new ones. This new phenomenon has been ignored in the literature thus far, but due to increased product replacements rates its relevance will grow. We present a generic model, applied in a case study with real-life data from ConRepair, a third-party service provider in plant control systems (mainframes).

Volumes of demand for spares, defects returns and phase-out returns are interrelated, because the same installed base is involved. In contrast with the existing literature, this paper explicitly models the operational control of both failed- and phase-out returns, which proves far from trivial given the non-stationary nature of the problem. We have to consider subintervals within the total planning interval to optimize both Last Time Buy and control policies well. Given the novelty of the problem, we limit ourselves to a single customer, single-item approach. Our heuristic solution methods prove efficient and close to optimal when validated.

The resulting control policies in the case-study are also counter-intuitive. Contrary to (management) expectations, exogenous variables prove to be more important to the repair firm (which we show by sensitivity analysis) and optimizing the endogenous control policy benefits the customers. Last Time Buy volume does not make the decisive difference; far more important is the disposal versus repair policy. PUSH control policy is outperformed by PULL, which exploits demand information and waits longer to decide between repair and disposal. The paper concludes by mapping a number of extensions for future research, as it represents a larger class of problems.

Key words: spare parts, reverse logistics, phase-out, PUSH-PULL repair, non stationary, Last Time Buy, business case

1. Introduction

This paper discusses how service repair firms apply reverse logistics in supplying spare parts for servicing *aging* mainframe plant control systems. Plant control systems are used for automated process management of large industrial installations. In the final stages of the mainframe's product life cycle, spare parts become scarcer and expensive to produce. Repair is a good option, but *failed* parts cannot always be repaired. This paper therefore introduces a new source that can be tapped to achieve this: *phase-out* returns. These are returns retrieved from systems that are being abandoned (phased out) by one customer, but may still fit another customer in a second life. In particular for aging products such as mainframes, phase-out returns are used to guarantee availability of spare parts for products still in operation. Due to today's replacement rates we often face multiple so-called *phase-out occurrences* and supply of phase-outs can be abundant. Often these returns are in still good ('repairable') condition.

In the service process, three major actors play a role. Original Equipment Manufacturers (OEMs) fulfil targeted service levels to their customers at a certain price. Customers include oil refineries and platforms, offshore windmills, medical facilities and nuclear power plants, all high-value capital goods. Failures of plant control systems are rare, but incur tremendous downtime cost when they happen. So a high availability of spare parts is needed.

At some point, OEMs outsource the service to the already introduced third party repair firm, the focal company of this study. This service provider then supports customers on behalf of the OEM for the remaining product life time (usually several years). General terms and conditions (i.e., service levels and prices) on both servicing customers who keep the old (mainframe) systems and scheduling of phase-outs of others are formalized in an Extended Service Program (ESP), which is a deal between the OEM and the repair firm. The repair firm takes over full operational and financial responsibility for service repair of the mainframes. Based on the ESP the repair firm makes Service Level Agreements (SLAs) with individual customers, concerning similar but possibly differentiated terms and conditions.

The timing of outsourcing to the repair firm usually coincides with the conclusion of the OEMs production of new spares. It then offers the (one-time) opportunity to acquire a final batch of new service parts to the repair firm: the so-called *Last Time Buy (LTB)*. Alternatively, the repair firm may want to use future phase-out returns as a cheaper but more uncertain source for spares.

The repair firm balances between own (profit driven) interests and customer service. In this balancing act, often repair firms apply a combined *control policy* for repair and disposal decisions and LTB. A major issue is to decide to repair or scrap repairable returns immediately upon arrival (PUSH) or

on demand (PULL). In our opinion, decision support models can make the tradeoffs more visible. Figure 1 summarizes the problem at hand schematically.

Figure 1 about here

Triggered by a real-life business case (ConRepair), this paper investigates and evaluates the modelling and optimization of a combined LTB & control policy, with the additional option of phase-out returns. By documenting a real life case study we identify this new phenomenon, which has thus far been ignored in the literature, and whose relevance is increasing with shortening product life cycles. Phase-outs complicate the analysis considerably, since they reduce the installed base and with it, change the demand- and return rates over time. At the start of the planning period the demand for repairs will be relatively high and phase-out volumes low, but this changes over time. Clearly, the time-based matching of supply and demand is complicated as the problem is non-stationary. In optimizing LTB & control policies one cannot simply rely on net demand throughout the planning interval as most models do. Sub-intervals are necessary. We adopt a simplified (single item, single customer) heuristic approach as a first step. We use a cost driven model where lacking service is translated into backorder cost and penalties. The aim is to maximize overall profit for the (focal) repair firm, using generic LTB and control policies.

The case results in turn have consequences for the generic policies; as they are counterintuitive. In fact (endogenous) decisions of the focal firm turn out to have little impact on the overall result. ConRepair depends on exogenous variables to maximize its profits, while the endogenous policy variables serve the customers' interest. Optimizing LTB volume does not make the difference; rather, it is the repair/disposal control policy that is of major consequence. We show that PUSH is outperformed by PULL, which exploits demand information as it waits longer to decide.

Section 2 proceeds by developing the business case of ConRepair, and introduce the problem. Based on this, Section 3 reviews the relevant literature in order to obtain building blocks and motivate our modelling approach. Section 4 gives a (mathematical) problem description, including assumptions and delineations. We then develop an appropriate model to analyze the issues at hand. Section 5 validates the heuristics used and presents the results of an elaborate numerical study using ConRepair data. Finally, Section 6 contains a discussion and the conclusions. The results and limitations of our research are used to indicate possible follow-up research.

2. Business practice

The OEM in this case study is Honeywell Industrial Automation and Control (IAC), a global player in industrial automation. It produces and maintains Distributed Plant Control Systems, i.e. networks of intelligent (automated) mainframe stations that control an industrial (chemical) plant. Because of the high capital value of these plants involved and their dependency on control systems, the latter require prompt service in case of failure.

Nowadays, desktop-based plant control systems serve increasingly as an alternative to the existing mainframes in the field. Customers, of course, replace their mainframe systems at their own discretion, and some do that more quickly than others. Hence, the phase-out return of one customer (who introduces desktops) may well serve another customer as a spare part for mainframes that stay in operation. The installed base, i.e., all products in operation in the field, now provides *two sources of returns* (defects and phase-outs) and also serves as a *reuse market* (demand for spares). Due to the phase-outs, its size reduces over time.

Returned parts are repaired and made available again as spare parts, complementary to newly acquired spares. Phase-out returns represent larger volumes than failed parts, as both return rates and yield (the percentage of the return flow that is repairable) are higher.

Honeywell IPC has outsourced the service of controllers of -amongst others- HPEP, Excel EMC, DGP and Mikronik- to ConRepair. ConRepair was established by a management buy-out (MBO) from the former Honeywell Amsterdam Repair Center in 1997. The company oversees the financing, planning and execution of the entire support process, dealing directly with the final customer, based on the ESP.

ESPs concern the service process for “mission-critical” parts only. The repair firm commits it self to collecting and testing *all* defect and phase-out returns. We deal with expensive parts but failure incurs even higher collateral (downtime) cost. The repair company jockeys the pooling advantage using commonality between customers to both improve availability and reduce cost.

The time-period covered by an ESP is negotiated, but is usually 3 to 5 years. The contract period represents our planning interval. When the ESP agreement between OEM and repair firm ends all service is terminated. Based on the ESP, one develops individual SLA contracts between the final customers and the repair firm. An SLA defines minimal service requirements fulfilled by the repair firm and prices to be paid by customers. *Availability* from stock (either new- or repaired parts) and the *total reliability* (including delivery backorders within at most one week) is measured. New parts cost the full catalogue list price, while repaired parts generally cost less (about 80% of list price), but usually also have lower cost. Companies like ConRepair generate income by charging the prices in case of actual failure; thus, if no failures occur ConRepair has no income, but still all of the cost. The repair firm pays the OEM for the LTB, which represents a major investment at the start of the ESP planning interval. Although largely exogenous, phase outs scheduling is essential. To be of use in covering demand for spares, phase-out collection

must take place (just) before failures. Because phase-outs decrease the installed base' size (of mainframes) smaller, *demand* for spares decreases during the planning interval and phase-out volumes increase. This essential input of the control policy. As mentioned, the basic options for control policies in practice are PUSH and PULL. In the literature, e.g. van der Laan et. al (1998), it is described in the following manner:

1. PULL policy: all repair activities are done reactively as soon as the serviceable inventory position drops below the base stock level. This policy is expected to keep serviceable inventories small, but there is more dependency on the inbound repairables inventory.
2. PUSH policy: all repair activities are done proactively, upon arrival of a part return. This policy is expected to have larger outbound inventories (whereas there are no inbound inventories).

PUSH is more intuitive, and has more flexibility to schedule repairs conveniently as they are not directly needed for service. However, it also bears the risk of repairing parts that prove unneeded later on. PULL creates a time advantage by waiting for actual demand and delays financial investments (read repair cost) until actually needed. But it introduces additional lead-time for repair when serviceable inventory is too low. PULL operates in practice as a base stock level with one-for-one repair; repair starts whenever a demand arises. Please note that PUSH is not a special case of PULL, even if the base stock level is set arbitrarily high, since PUSH is triggered by product *returns* and PULL is triggered by product *demands*. Contrary to PUSH, PULL builds up an inbound inventory of repairables. Take-back of phase-out parts increases the possibility of oversupply and hence the need for *scrapping* repairables prematurely, i.e. scrap upon arrival. Similar decisions such as LTB-volume may have different optimal values for PUSH and PULL.

Next to the management of repair operations, a disposal policy controls the inflow of part returns. An (also optimized) *dispose-down level* stipulates the point at which repairable inbound returns are to be scrapped instead of repaired— namely, when the total number of parts in the system is sufficient to cover expected future net demand (demand minus future returns). Scrapping excessive returns averts the costs of keeping stock, but also removes the possibility to repair at a later stage. Non-repairables are scrapped immediately at all times.

To determine the Last Time Buy (LTB), the customers forecast the expected number of failed parts based on historic failure rates. The repair firm has to forecast the repair yields. Although so called planning interval is fixed in the ESP, uncertainty is caused by the demand for parts in the future, which depends on the failure rates and (on the supply side) the number as well as timing and repair yield of future phase-out and defect returns needed to cover future demand.

3. Literature

The literature on spare parts management is abundant and may be divided into two parts: scheduled maintenance and unplanned maintenance (Kennedy et al. 2002). For the latter, safety stocks need to be in place in order to be able to deal with future maintenance activities. Typical decisions that need to be made are whether to replace or repair on site, and where to strategically place inventories of spare parts. Since the Last Time Buy decision, the repairability of failed parts and the non-stationarity of demand due to phase-outs are relevant for our case, this review focuses on (combinations of) those aspects.

Papers in the field do not typically deal with drops in demand as the installed base size is stable. In that sense, our research is related to the literature on product obsolescence, where it is commonly assumed that, whereas demand is stable during the lifetime of a product, the remaining lifetime itself is a random variable (David et al. 1997). In our case, however, the remaining lifetime of a product is fixed and known due to ESP, although demand may drop due to planned phase-outs. Cattani and Souza (2003) acknowledge the importance of the Last Time Buy (in their case, 'end-of-life build') decision, and focus on its optimal timing for a fixed and known remaining number of periods of demand. Building on earlier results (Cattani 1997), Cattani and Souza (2003) quantify the benefits of delaying the end-of-life build. In our research, the timing of the Last Time Buy decision is fixed because it concurs with the moment of outsourcing of OEM to repair firm. Also the period which the repair firm must provide service to the customer is a given. But in our case demand for spare parts over the planning horizon is non-stationary, as the installed base changes over time, due to phase-outs. To achieve optimality, the control policy must deal with balancing supply and demand, which enhances the importance of the scrap versus repair decision.

Few studies combine Last Time Buy and repair/disposal policies, let alone phase-outs. Below we describe those who have come closest to our problem. Teunter and Klein Haneveld (1998) describe a situation in which an OEM stops supplying spare parts for a single machine. The operator of the machine is offered an opportunity to place final orders for a number of critical parts to keep the machine operational up to a fixed horizon. Based on this horizon, the failure rates of the components, the prices of spare components, holding costs and machine downtime costs, the size of the (near-)optimal final orders are determined. Unused parts are scrapped only at the end of the service period, which means that such a policy may lead to high stocking costs. In another paper, Teunter and Fortuin (1999) consider a more complicated situation in which failed spare parts can be repaired and reused, and unused parts can be removed from stock *before* the end of the horizon using a dispose-down-to level. Since the cost of repair is assumed to be negligible, all returned items are repaired and re-stocked. Some of the theoretical results were applied at Philips, as described in Teunter and Fortuin (1998). Our case is similar to the one in Teunter and Fortuin (1999), although we do not assume that repair costs are negligible nor that

all parts can be repaired. Repair cost reduce the viability of immediate repair (read PUSH policy), particularly when subject to a quality based yield factor as is the case in our study.

The above-mentioned papers constitute a category that resorts to newsvendor-type approaches to determine (near-) optimal final orders. In contrast, Spengler and Schröter (2003) take a system dynamics approach in order to take into account product life-cycle aspects and dependencies between sales and returns. Rather than focusing on finding optimal final orders, they investigate whether and how parts reuse can contribute to spare parts management after the final order has been placed. They also explore certain system dynamics.

The next category of relevance is the case of seasonal products with limited replenishment options. Although they lack a repair process, returns are explicitly modelled. Many products, such as apparel, are demanded during a short period of time only, while replenishment opportunities are very limited or non-existent. Mostard and Teunter (2006) analyze the case of a Dutch mail-order company. A single order is placed before the selling season starts. Purchased products may be returned by the customer for a full refund within a certain time interval. Returned products are resalable, provided they are returned before the end of the season and are undamaged (this equals a yield factor in the sense that the fraction of returns that can be reused is fixed and known). Products remaining at the end of the season are disposed of. All demands not met directly are lost. A simple closed-form newsvendor equation determines the optimal order quantity given the demand distribution, the probability that a sold product is returned, and all of the relevant revenues and costs. Although Teunter and Fortuin (1999) and Mostard and Teunter (2006) explicitly model product returns, and Spengler and Schroter (2003) control repair and disposal implicitly through recovery and recycling rates, the repair process is not very explicitly modelled.

A final relevant body of literature concerns models with combined new manufacturing and repair operations where control policies deal with an inbound repairables inventory and an outbound serviceable inventory (for an overview, see van der Laan et al. 2004). The latter consists of both new and repaired items, which are assumed to be identical. Convenient policies to control the repair process are so-called PUSH policies (part returns are repaired as soon as a certain quantity is reached) and PULL policies (a certain quantity from part-return inventory is repaired as soon as the inventory position drops below a certain value). Depending on the cost structure and lead times, either PUSH or PULL will perform better (van der Laan 1998, van der Laan 1999). No last time buy applies to these models.

All of the above papers take a cost-driven approach, meaning that falling short on service levels can be given a monetary value by means of backorder- and penalty costs. Service-level approaches (optimizing service levels regardless of the costs) are described in general terms (without returns) in Fortuin (1980, 1981), Sherbrook (2004) and Muckstad (2005). Service models are more intuitive to customers, but more difficult to use because a constant monitoring

and calculation is needed on the achieved performance levels. At each point in time we must check whether we should start repair operations. This aspect demonstrates the major drawback of a service-level approach, which is simultaneously a 'look back' and a 'look ahead' policy. Its mere look-ahead nature makes the cost approach more robust than service-level models, which are inclined to overcompensate past drops in performance. Also it allows for unambiguous interpretation of results as it is unilaterally monetary.

We do not know of any papers that combine the last time buy decision with explicit PUSH and PULL repair decisions with limited repair yield (of two streams). We are the first to introduce the phase-out phenomenon into the problem. The solutions as suggested in the past literature do not suffice to deal with this situation, because in order to exploit the phase-outs we must model volume, timing and yield of both return types explicitly because their behaviour is totally different. Also we must connect phase-out occurrences to installed base size and hence demand for spares. This paper develops control policies and heuristic control rules to facilitate the integral management of acquiring, repairing and disposing spare parts, while monitoring the performance throughout the remaining service time after the last time buy.

4. A modelling approach

4.1 Problem description, assumptions and simplifications

The installed base, of size $IB(t)$ at time t , is serviced through a pool (serviceable inventory) of identical spare parts that are either new or repaired. Demand fulfilled from serviceable stock, consists of new (unit cost c_m) and repaired (unit cost c_r) parts. Scrapping of unneeded or non-repairable parts is assumed to incur zero cost. New parts are used up first because of their higher capital value; the use of repaired parts takes place at an increasing pace over time. There is no quality difference between new and repaired parts. Although engineers have other duties, such as calibration of new systems, repair gets priority once repairables are available and repair is needed, so we linearize unit repair costs and to set repair capacity to be infinite.

The customers that operate the installed base are identical in the sense that they have negotiated the same conditions and prices for new (p_m) and repaired (p_r) parts as well as similar service levels. The presence of an ESP as a foundation for individual SLAs makes it feasible to model the total installed base as a single *meta-customer*.

In line with general ESPs, failures are fixed at most one week after the initial call—considering that, on average, redundancy in the mainframes keeps systems up and running for about that period of time. Repair lead-time is assumed to be one week and fixed. Available parts are delivered from the *serviceable inventory* immediately; otherwise the customer has to wait for the part during the (fixed) repair lead-time,

To monetarize service levels, low responsiveness is penalized. If no serviceable stock is available, then the demand is backordered, and a penalty fee c_b per part per week is paid to the customer to compensate for the delay. Backorders that are still outstanding at time $t=H$ cannot be fulfilled anymore and are penalized with c_f per part ($c_f \gg c_b$). Hence, penalty costs c_b provides an incentive to offer a certain availability of serviceable parts *throughout* the planning horizon through an optimized repair policy. Penalty cost c_f provides an incentive to guarantee a certain availability at *the end* of the planning horizon through an optimal choice of the LTB.

Part failures (a Poisson process with failure rate μ in the installed base both generate demand for spare parts and cause the return of the failed parts. The service provider has to place a final order for a specific part at time $t=0$, but it has contractual agreements to provide service for that part until time $t=H$. Then at time $t=0$ it needs to optimally choose Last Time Buy quantity, Q , and its future repair policy taking into account stochastic future demand and part failures and the deterministically scheduled phase-outs in order to optimize its profits during the time interval $[0, H]$. Since all demand is triggered by failed parts in the installed base, demand is itself a Poisson process with rate $\mu \cdot IB(t)$.

Standard procedure would be to determine the LTB based on projected *net* demand over the entire planning interval $[0, H]$. But with phase-outs this may not always lead to sufficient performance, as the following simplified example shows. Consider a deterministic world with an initial installed base of 100 parts, continuous demand for spare parts (4 parts per time unit) over the planning horizon (100 time units), but parts cannot be repaired. However, there is one phase-out occurrence of 50 parts after 50 time units (so halfway) and they are all repairable. After the phase-out occurrence the installed base halves and therefore the demand drops to 2 parts per time unit. The net demand for spare parts over the planning horizon equals demand prior to the phase-out plus demand after the phase-out minus the phase-out returns $= 50 \times 4 + 50 \times 2 - 50 = 275$. The LTB could therefore be chosen to be 250 and no stock-outs would occur. Note that the LTB is sufficient to handle the demand (200 parts) prior to the phase-out occurrence. Now suppose that instead of halfway, now after 75 time units (so later) a phase-out of size 75 (instead of 50) occurs. In that case we have the same net demand $75 \times 4 + 25 \times 1 - 75 = 250$, but now an LTB of 250 parts is no longer sufficient to cover the demand prior to the phase-out occurrence ($4 \times 75 = 300$ parts).

In other words: the *non-stationarity* of the problem and the fact that demand until the first phase-out occurrence may exceed net demand complicates the optimization of both LTB and operational control. Therefore we have to split the total planning interval into *subintervals* and formulate minimal availability for each one of them. Otherwise we bear the risk of (huge) temporary drops in service level. One way that is effective and tractable within our modelling framework is to specify a *maximum acceptable* stock-out probability, $1 - \beta_p$ just before a phase-out occurrence.

At ConRepair we observed that phase-outs are scheduled by customers. Therefore, in our model we assume that the timing of phase-out returns is deterministic and known at the beginning of the planning period. As these phase-outs reduce the size of the installed base, future demand for spare parts will also be reduced. Both phase-out returns and returns of failed parts can be repaired, but they may have different expected repair yields (y_P and y_F , respectively). To control the repairable inventory $N(t)$ and serviceable inventories of acquisitioned parts $M(T)$ and repaired parts $R(T)$, we employ either a PUSH policy (repair upon arrival) or a PULL policy (at any time t repair parts, as long as the inventory position— that is, serviceable inventory $M(t)+R(t)$ minus backorders plus repair work in-progress— drops below base-stock level $S(t)$). The fixed repair lead time equals L . As at some points in time the number of available parts may be deemed sufficient to service the projected future net demand, it may be cost-effective to dispose of some of the incoming repairable parts. To accommodate this, we introduce the decision variable $U(t)$: any incoming part will be disposed of if the echelon inventory position (Inventory of repairables plus inventory position) is at or above $U(t)$.

In order to maximize expected profits over the interval $[0, H]$ we need to develop rules that optimize decision variables Q , $S(t)$ and $U(t)$, which is the objective of the remainder of this section. The complete list of notations and base-case data is given in Table 1. Note that data are confidential and therefore normalized. A schematic representation of the model is given in Figure 2. The complexity makes a *heuristic* approach needed. As even the heuristic problem poses quite some challenges we limit ourselves to a *single-item*, which however does not harm the generic nature of the model when assuming that demand for individual spares is independent. We also assume perfect information on all parameters, even in sensitivity analysis, meaning that all decisions in LTB and control policy can be adapted when input values change. Exogenous variables such as failure rates, repair yields and variable (unit) operating costs and revenues are constant in time. To explain the model in a lucid manner, we start with a stationary situation without phase-out returns.

Figure 2 about here

4.2 No phase-outs

Even without phase-outs, the analysis is complicated by the fact that there are three different stocking points: for newly acquired parts, repaired parts and repairable parts with carrying charge (h_m , h_r and h_n , respectively). Schematically, if we ignore the repair lead-time L (which is reasonable, as L is very small compared to planning horizon H), then the inventory processes under PUSH and PULL control follow the behaviour depicted in Figure 3.

Figure 3 about here

Since there are no phase-outs, the installed base is fixed ($IB(t) = IB(0)$ for all t) and weekly demand is therefore stationary, with expectation $E(D) = \mu \cdot IB(0)$. Under PUSH control, the inventory of new parts decreases from the initial level Q with the average demand rate until all parts are used (say, at time X). Until that moment, repaired parts have accumulated to the expected value Qy_R (Q demands cause Q returns that are repaired upon arrival with expected yield Qy_R). As soon as the new parts have run out, demand is serviced through repaired parts. The serviceable stock decreases with the net demand (demand minus repairable returns) rate. Under PULL control, returned parts accumulate in the repairable inventory $N(t)$ until at some time t the inventory position $IP(t)$ falls below level $S(t)$. Then a sufficient number of repairable products, if available, is ordered for repair in order to restore the inventory position to level $S(t)$. If the stock of repairables is not sufficient to meet this level, then all of the available parts are repaired. To limit the number of decision variables, we assume zero fixed costs for repair, so that no batching takes place for economies of scale.

As long as repairables are available, the PULL policy operates as an $(S(t)-1, S(t))$ policy. If demand during lead-time follows a normal distribution with cdf $G(\cdot)$, which is reasonable since the underlying processes are Poisson processes, then in the long run the near-optimal base stock level $S^*(t)$ satisfies (Silver et al. 1998, p. 255)

$$S(t)^* - 1 = E[D(t, t+L)] + k\sqrt{\text{Var}[D(t, t+L)]},$$

where

$$k = G^{-1}\left(1 - \left(\frac{\phi}{\sqrt{\text{Var}[D(t, t+L)]}}\right)\right) \quad (1)$$

and $\phi = \frac{c_b}{c_b + h_r - h_n}$ (Silver et al. 1998, p. 266).

The purpose of the Last Time Buy quantity Q is to stock sufficient acquired parts at the beginning of the planning period such that stock-out costs at the end of the planning horizon are well balanced against the operational costs during the planning horizon. If there were no phase-outs to consider, we could have formulated the problem as a standard newsvendor problem. Given a certain realization, z , of net demand over the planning horizon (i.e. $ND(0, H) = z$), the profit function is given as (see Silver et al. 1998, p. 404)

$$P(Q, z) = \begin{cases} -Qv + pz + g(Q - z) & , z < Q \\ -Qv + pQ - B(z - Q) & , z \geq Q \end{cases}$$

where v denotes the cost per unit acquired, p denotes the revenue per unit sold, g denotes the salvage value for any unit not sold by the end of the planning horizon, and B denotes the penalty cost for demand not satisfied at the end of the planning horizon. In order to fit our problem in the standard newsvendor formulation, appropriate values of v , p , g , and B need to be developed. Note that v should somehow reflect the inventory carrying costs. Theory tells us that the expected profit function reads as

$$E(P(Q)) = (p - g)E(ND(0, H)) - (v - g)Q + (p - g + B)(1 - \Pr(ND(0, H) < Q)),$$

and is optimized for the Q that satisfies

$$\Pr(ND(0, H) < Q) = \frac{p - v + B}{p - g + B}. \quad (2)$$

Under PULL control, the repairable inventory reaches its maximum as soon as the inventory of newly acquired items, Q , is depleted at time $t = X$. Just before that time, at about $X - S / E(D)$, where $E(D)$ equals expected weekly demand, the repair facility starts processing repairable parts up to level $S(t)$ (see Figure 3b). Assuming that inventories linearly increase/decrease with the return rate and (net) demand rate, the total inventory carrying costs are approximately given by

$$\begin{aligned} & \frac{1}{2}Q \cdot X \cdot h_m + \frac{1}{2}(Q - S)y_F H h_n + \left(\frac{1}{2}S^2 / E(D) + S(H - x) \right) h_r \\ &= \frac{1}{2}Q \left(\left(\frac{X}{H} \right) h_m + y_F h_n \right) H + \left(\frac{S^2}{2E(D)} + S(H - X) \right) h_r - \frac{1}{2}y_F S h_n H \end{aligned} \quad (3)$$

Timepoint $X = Q / E(D)$ obviously depends on Q , which makes optimization problematic, since total carrying costs are quadratic in Q . However, the fraction of demand that is fulfilled by newly acquired parts, X / H , is approximately equal to 1 minus the fraction of demand that is potentially serviced through part repairs, or $X / H \approx 1 - y_F$. Using this approximation, the inventory carrying cost contributed through Q during time interval $[0, H]$ is then approximated by $Q((1 - y_F)h_m + y_F h_n)H / 2$ and should be included in the cost per acquired unit v .

Turning back to our newsvendor formulation, appropriate values for v , p , g , and B are as follows:

$$\begin{aligned}
 v &= c_m + ((1 - y_F)h_m + y_F h_n)H / 2 & - \text{unit acquisition cost} + \text{carrying cost per unit acquired} \\
 p &= p_m & - \text{unit sales price of new parts} \\
 g &= 0 & - \text{surplus is scrapped against zero profit/cost.} \\
 B &= c_f & - \text{shortage is penalized against } c_f \text{ per unit short}
 \end{aligned}$$

Equation (2) then becomes

$$\Pr(ND(0, H) < Q) = 1 - \frac{c_m + ((1 - y_F)h_m + y_F h_n)H / 2}{p_m + c_f}. \quad (4)$$

Under PUSH control, all repairables are ‘pushed’ to the serviceable inventory upon arrival, so that the total inventory carrying costs simply read

$$\frac{1}{2}Q \cdot X \cdot h_m + \frac{1}{2}Q y_F H h_r = \frac{1}{2}Q \left(\left(\frac{X}{H} \right) h_m + y_F h_r \right) H.$$

Through $X / H \approx 1 - y_F$ this is approximately equal to $Q((1 - y_F)h_m + y_F h_r)H / 2$, so that under PUSH control $v = c_m + ((1 - y_F)h_m + y_F h_r)H / 2$. This expression is similar to that under PULL control with h_n simply replaced by h_r . Hence, as the values of p , g , and B remain the same, under PUSH control equation (4) holds with h_n replaced by h_r .

In order to prevent the build-up of excessive stocks of repairables, we adopt the following disposal policy: do not accept incoming returns as long as the echelon inventory position $IP(t) + N(t)$ equals or exceeds level $U(t)$. In order to derive near-optimal values of $U(t)$, we analyze the impact of accepting an incoming return at time t through a marginal analysis. Suppose we accept an incoming return, and $IP(t) + N(t)$ is raised to $U(t)$. The marginal *investment* v_t of accepting the return should cover the marginal *benefit* of accepting it (sales value p + avoided penalty c_f times the probability of selling it), so that

$$v_t = (p + c_f)(1 - \Pr(ND(t, H) < U(t))),$$

or

$$\Pr(ND(t, H) < U(t)) = 1 - \frac{v_t}{p + c_f} \quad (5)$$

for appropriate values of v_t and p . Under PUSH control, the consequence of accepting an additional return at time t is repair cost c_r and total carrying cost $h_r(H-t)$, so $v_t = c_r + h_r(H-t)$. If the part is sold at the end of the planning horizon, it generates value $p = p_r$. Under PULL control you invest in repairing the part only when you expect to sell it, so $v_t = h_n(H-t)$ and $p = p_r - c_r$. This rule adjusts the desirable inventory of serviceable and repairable parts as more information (past demands, returns and repaired parts) becomes available. Note that one should dispose of all incoming returns whenever $H - L < t < H$, since they cannot be repaired before the end of the planning horizon. The values of v_t and p are summarized in Table 2.

Table 2 about here

4.3 Full model with phase-outs

Phase-outs complicate matters in two ways: 1) demand is no longer stationary over the planning interval $[0, H]$, as the demand rate decreases after each phase-out occurrence (demand is stationary though *in between* phase-out occurrences), and 2) it is no longer valid to determine Q solely on the basis of the net demand during the interval $[0, H]$. To see this, suppose that phase-outs are planned at times $\tau_i, i = 1, \dots, n$. If a phase-out is scheduled relatively early, then Q could be calculated solely on the basis of the net demand during the interval $[0, H]$ (see Figure 4a), since during the planning interval $[0, H]$ sufficient parts are available to ensure reliable operations. But if a phase-out is scheduled near $t=H$, this same Q may not be sufficient to satisfy demand up to the phase-out occurrence (see 5b).

Figure 4 about here

One way to get around this issue would be to specify the maximum allowable risk, i.e. stock-out probability $(1-\beta_p)$, that management is willing to accept just before a phase-out occurrence at time τ_i . Based on this maximum allowable risk we calculate Q_i for each interval $[0, \tau_i]$. Taking the maximum over all i ensures that at each $t = \tau_i$ the stock-out probability is at most $(1-\beta_p)$. Section 5.1 elaborates on the proper setting of β_p .

Although the inventory process is more complicated when compared to the case without phase-outs, we choose to use the same approach as in the previous section—that is, proceeding

from a particular timepoint X onwards, allow inventories to reduce gradually with the average net demand rate.

If we let $\tau_{n+1} = H$ and $\tau_0 = 0$, then over the complete planning horizon H the Last Time Buy quantity Q_{n+1} should satisfy

$$\Pr(ND(0, \tau_{n+1}) \leq Q_{n+1}) = 1 - \frac{v}{p_m + c_f}, \quad (6)$$

where under PULL control $v = c_m + (\rho h_s + y_F h_n)H/2$, and ρ is an approximation of the fraction of demand that is fulfilled by newly acquired parts:

$$\rho = 1 - \frac{y_F \sum_{j=0}^{n+1} \mu(\tau_{j+1} - \tau_j) IB(\tau_j) + y_P (IB(0) - IB(H))}{\sum_{j=0}^{n+1} \mu(\tau_{j+1} - \tau_j) IB(\tau_j)}.$$

Note that for $n=0$ we have $IB(0) = IB(H)$, so ρ reduces to $(1 - y_F)$ and equation (6) reduces to (4). For $1 \leq i \leq n$, we choose to satisfy the following conditions in order to maintain a minimum performance just before phase-out moments τ_i :

$$\Pr(ND(0, \tau_i) < Q_i) \geq \beta_p, \quad 1 \leq i \leq n.$$

Under PUSH control, the same expressions hold— with h_n simply replaced by h_r . The near-optimal value of the Last Time Buy quantity is then calculated as

$$Q = \max_i \{Q_i\}.$$

In the same vein we determine the dispose-down-to level at time t for planning horizon H :

$$\Pr(ND(t, H) < U_{n+1}(t)) = 1 - \frac{v_t}{p + c_f} \quad (7)$$

and

$$\Pr(ND(t, \tau_i) < U_i(t)) \geq \beta_p, \quad 1 \leq i \leq n,$$

where v_t and p are given, as in Table 2. The near-optimal value of $U(t)$, finally, is calculated as

$$U(t) = \max_{i|\tau_i - L > t} \{U(t)_i\}.$$

The above-mentioned heuristic decision rules applied together should minimize total net profit NP , defined as total sales minus total acquisition and repair costs, minus total inventory carrying costs, minus total penalty costs. The next section first assesses the quality of our approximate rules by comparing them with optimal solutions obtained through enumeration. Subsequently, we analyze our policies for the base case of Table 1 and various other scenarios.

5. Numerical study

The base-case scenario for our numerical study is based on actual data from ConRepair, collected over the period 2004-2007 for a particular part (see Table 1). For reasons of confidentiality, the cost data are normalized and the part name is not specified. We assume that parts fail according to a negative exponential distribution with failure rate μ . Hence, between two sequential phase-outs τ_i and τ_{i+1} the demand (and return) process is a Poisson process with mean $\mu \cdot IB(\tau_i)$. For the heuristics it is convenient to assume that net demand during some time interval is normally distributed. This assumption is justified, since between two sequential phase-outs τ_i and τ_{i+1} also net demand is a Poisson process, with mean $(1 - y_F) \cdot \mu \cdot IB(\tau_i)$.

We apply the PUSH and PULL heuristics as developed in the previous section, analyzing several variations of the base-case scenario through simulation. Each simulation run consists of a period of three years (150 weeks). For each scenario, the simulation was run 3000 times. The observed maximum relative error in the net profit over all reported simulations was below 1%. Section 5.2 reports the base case and describes the impact of exogenous changes in repair yield and the size and timing of phase-outs. Results are represented in terms of (net) profit, contribution of the LTB, disposal rates and availability of spares. First, however, we validate the developed heuristics in Section 5.1.

5.1 Model validation

As a result of model complexity, we had to resort to a number of approximations and heuristic procedures. Although optimal solutions (obtained, for instance, by extensive (enumerative) simulations) may be preferable, this can only be done for Last Time Buy quantity Q and never for all control policy decisions. Moreover, the derived formulas provide us an easy handle with which we can interpret and understand the numerical results, since they are more lucid and closer to the actual decision-making practice. Finally, there is the practical advantage of a considerable amount of time saved, compared to enumerative simulation. A drawback is that the approximation of the expected inventory may perform poorly, especially when phase-out volumes are large. The

heuristics should therefore be of sufficient quality. This we will check by comparing heuristic and optimal values of Q .

Figure 5 shows the actual behaviour of the inventory processes for the base case (see Table 1) under our heuristic policies and near-optimal decisions. It can be seen that these are very similar to the stylized pictures of Figure 3 that were used as a basis for the inventory cost approximations.

Figure 5 about here

In order to evaluate the accuracy of the Last Time Buy heuristics, we disable the disposal option and initially set $\beta_p = 0$. Table 3 shows that the heuristics for the Last Time Buy quantity are very accurate for the complete range of possible repair yields. The PULL heuristic slightly underestimates the Last Time Buy, but this hardly affects performance; all scenarios of Table 3 show heuristic Last Time Buy quantities that differ less than 2% from the optimal ones, while net profits differ less than 1% from optimal. Considering that this is comparable to simulation inaccuracies, the differences in net profit are probably not statistically significant. Increasing the size of the Last Time Buy decreases the length of the period that parts are repaired and thus the probability of stock-outs under PULL. The PULL heuristic does not take this into account, thereby underestimating the Last Time Buy. The performance of the PUSH and PULL policies depends on the availability of repairable products, and is therefore sensitive to stock-outs of repairable parts that may occur just before a phase-out. The performance will therefore depend on the appropriate setting of β_p (see Table 4). An alternative to the service-level approach in setting β_p is to follow exactly the same cost-based control rules that are used for $\tau_{n+1} = H$ for all τ_i , and then to take the maximum values of Q_i and $U_i(t)$. After some experimentation, we found that the qualitative results of such an approach are identical to the service-level approach. For ease of presentation, we therefore adhere to the service-level approach.

Table 3 also shows that the differences between PUSH and PULL regarding Last Time Buys are rather small, which suggests that inventory carrying costs are only a secondary determinant of the Last Time Buy quantity. Indeed, expression (4) claims that the main trade-off is between the unit acquisition cost, on the one hand, and the unit sales price and unit penalty cost at the end of the planning horizon, on the other— unless the planning horizon is very large.

Tables 3-4 about here

Evaluating the accuracy of the disposal policy is rather difficult, as we do not know the structure of the optimal policy. Its performance, however, can be compared against policies without

disposal. Table 5 shows the value of disposal for varied phase-out timing. In general, disposal appears to be beneficial, and its value increases as the phase-outs move further towards the end of the planning horizon. PUSH benefits much more from a disposal policy, as it is punished more severely for excessive stocks than PULL is. Expression (6) predicts that the disposal rate increases with higher inventory carrying charges, higher unit repair costs and lower unit sales costs of repaired parts. Based on (7), it is easily shown that for practically all relevant scenarios ($c_r < p_r$, and $h_n H \leq h_r H < c_f$) the disposal level is generally lower under PUSH than under PULL, which means that PUSH, on average, disposes of more product returns than PULL does.

Table 5 about here

5.2 Base-case results

Table 6 reports on the performance of the PUSH and PULL heuristics for the base-case scenario. In terms of net profit per week, the PULL policy is superior to the PUSH policy. This is due mainly to the difference in total inventory costs and penalty costs. PULL's emphasis on less inexpensive repairable inventory, rather than on repaired inventory, provides an important cost advantage. One might expect a firm to consider a PUSH policy in order to guarantee a higher service level, but results show that this is not the case. The explanation is that to reduce the investment in expensive repaired inventory, the Last Time Buy quantity Q is suppressed and/or more returns are disposed of. The probability of stock-outs then increases, resulting in service levels that are lower than under PULL. Actually, none of the many scenarios that we analyzed yielded a better performance for PUSH, as compared to PULL.

Table 6 about here

5.2.1 Impact of the unit repair cost

An important determinant of the performance of PUSH and PULL is the unit repair cost. Intuitively, if repair is cheap, PUSH may be a reasonable option. If repair is expensive, it may be better to use PULL. Figure 6 shows that while PUSH always performs worse than PULL does, its relative performance becomes better for smaller c_r . Here we modelled the carrying charge for repaired parts as $h_r = h_n + w \cdot c_r$, with w the opportunity cost of capital, in order to reflect the financial impact of stocking repaired parts. So, h_r increases linearly in c_r . It is easily shown that when $c_r=0$ (and thus $h_r = h_n$), PULL is equivalent to PUSH. As c_r grows it becomes increasingly attractive to temporarily stock returns before repair until they are really needed.

Figure 6 about here

5.2.2 Impact of the repair yield of failed parts

There is, of course, a time dependency between demands and returns (in the sense that a demand generates a return at the same moment), but the lead-time implies that one cannot benefit directly from that relation. In order to satisfy a demand directly from stock, a demand has to be fulfilled through a new part or a repaired part that had been returned in the past. It appears from Figure 7 that if the repair yield of failed parts rises, then the net profits increase—albeit less than linearly. Although the Last Time Buy volume naturally goes down with increased repairability, one has to rely on a return flow that is stochastic in timing and yield, which means that the Last Time Buy volume decreases less than linearly. It is important to note that serviceable inventory stock-outs can occur only once all new parts are used up. With decreasing Last Time Buy volume, the time t_r that all demand is satisfied directly from stock through new parts goes down too, so that the risk of stock-outs increases. This is also explained through relation (8), which shows that service level α goes down as t_r goes down. As the repair yield approaches 1.0, virtually all demand is satisfied from repairs, so there is no time for accumulating safety stocks of parts. Service level α therefore decreases sharply. The total fraction of demand delivered (β), however, remains high overall, so that customers can be confident—even if they are more dependent on returns. Note that the difference between PUSH and PULL diminishes as the repair yield approaches 1, since there is hardly any opportunity to stock parts.

Figure 7 about here

Analysis with the repair yield of phase-out shows similar results. Intuitively, one would expect that phase-outs, being a cheap source of supply, would enhance both service levels and net profits with increasing yield. Unfortunately, this is only partially true: too high yields lead to carelessness of the control policy regarding safety stock levels.

In case both the repair yields would be zero, the Last Time Buy would have to fully cover all demand for spares, and ConRepair's operations for this part would no longer be profitable (note that demand will still go down). Next to the repair yield of the phase-out returns, their volume and timing proves even more important as the next section shows.

5.2.3 Impact of phase-out volume

Next, we vary the volume of phase-out returns (Figure 8 and 10). As most of the demands are fulfilled through repairables, this is quite an important scenario. Observe, first of all, that net profit decreases with phase-out volume (Figure 8a). Reductions in the installed base decrease total demand for spare parts over the planning horizon, so that net profits decrease as well. The net profit *per unit demand* initially increases up to an annual installed-base reduction of 30% (Figure 8b)), since the returns coming from the phase-outs replace expensive new parts (hence, Q

decreases). At some point, however, the Last Time Buy cannot decrease further, since it has to protect against demand leading up to the first phase-out. An increasing volume of phase-out returns needs to be disposed of, and the relative contribution in fulfilling demand by returns goes down, forcing marginal profits down. This explains the convex curve of relative Last Time Buy contribution in Figure 9b. The diminishing role of returns also explains why the difference between PUSH and PULL decreases with higher annual reductions in the installed base. The superior performance of PULL is reconfirmed, however, as it waits for the right moment of repair and fewer returns are scrapped prematurely.

The managerial conclusion would be that phase-outs might be a source for returns, but that the effect of reduced turnover overrules all effects. In other words: no phase-outs may be best for ConRepair, since the installed base then remains of maximal size. But phase-outs are a given, so preventing or reducing them may not be an option. Next, we investigate the impact of the timing of phase-outs.

Figures 9, 10 about here

5.2.4 Impact of phase-out timing (and frequency)

Spreading the phase-outs (that is, increasing the frequency of phase-out moments into smaller batches, whilst keeping the total phase-out volume stable) hardly affects performance, so we will not report any further details here. A possible explanation is that inventory costs are relatively low, and if parts are returned in time and are not scrapped prematurely, there will be sufficient repairables available to meet demand. This result also strengthens the idea that phase-out returns make the difference mainly at the end of the period H , and failed parts fulfil much of the early demand, despite lower yield, because they are more equally spread. We therefore investigate the effect of changing the timing of an individual phase-out event.

In the scenario of Figure 10 there is one phase-out occurrence, and its timing is varied from week 25 to week 125. Phase-outs reduce the installed base, and thereby demand and net profits. Delaying phase-outs maintains the installed-base size and thereby total profit. It is even more instructive to look at net profits *per unit demand* to investigate the impact of phase-out timing. Delaying phase-outs at first increases net profits per unit demand, because the arrival of the phase-out returns allows the firm to become more and more aligned with the demand for repaired parts. Profits improve, due to lower holding costs. At some point, though, the phase-outs come in too late, more and more returns need to be disposed of and marginal profits start to decrease. Total profit keeps increasing, but the curve flattens. The managerial implication is that it pays off to delay phase-outs somewhat, although marginal profits decrease.

In this scenario the gap between PUSH and PULL is the largest when the role of returns is smallest— that is, when phase-outs occur either early (little demand, hence fewer failures) or late (excessive phase-out returns are disposed of) in the planning period.

Figure 10 about here

5.2.5 Wrap up

PULL is better at matching supply and demand because it waits longer to decide and uses actual demand information. PUSH often scraps too early or repairs too quickly. PULL, at the end of the day, leads to lower operational costs and due to better service has fewer problems with penalties and backorders. Last Time Buy quantity hardly differs between PUSH and PULL. The role of the scrap versus repair decision outweighs the Last Time Buy volume. So control policies should be PULL-driven repair combined with a dispose-down-to level geared optimally toward controlling the inbound returns.

A main paradox in our study is that customer vigilance leads to a service-level focus on their behalf, whilst repair firms want to maximize profits. Results show however that the profitability of phase-out use for ConRepair is low, because customers get a discount. Profits for ConRepair could be boosted by increased prices for repair. At the moment, all cost benefits are transferred to the customer: whereas the quality of repaired parts is equal to new parts, the price is only 80%. An alternative is to reduce repair cost leading to the same effect.

At first sight, therefore, there doesn't seem to be much point in take-back and repair of phase-outs. Lower phase-out volumes are actually beneficial to the firm because the installed base remains larger and the demand for spares stays high. Although an exogenous variable, this may be communicated to the customer (and OEM). Managerially this means that merely the scheduling of phase-outs can be adapted. It is essential that phase-outs are available as late as possible but in time to be reused.

There is another exogenous factor that has a big impact on profitability: the repair yield of both failed parts and phase-out returns. Once collection timing is optimized, collection quality must be ensured (by e.g. decent packaging), in order to avoid transport damage. Although endogenous variables can be optimized by PULL for the repair firm, the exogenous variables have a stronger impact on the profitability.

6. Conclusions and outlook

This paper studies the way in which (mainframe) plant control systems can be serviced using phase-out returns. The case study at hand represents a larger class of problems. Not only are

ESPs & SLAs of increasing importance in service logistics because downtime of installed bases is increasingly costly. Shortening product life cycles and replacement rates will boost phase-out returns. The global service and support market is growing. In 2010 it will represent a potential turnover of 90 billion US\$ per year (Blumberg, 1999, 2005). Improving performance in this field, as presented in this case, represents millions of dollars.

Practically all OEMs in capital goods at some point stop servicing and new spares production. Customers may abandon the mainframes phase out their systems, causing phase-out returns. At the same moment OEMs outsource to a 3rd party repair firm for customer who keep the old products, and offer Extended Service Programs (ESPs) for the remainder of the life cycle. At the start of these programs, OEMs generally offer the opportunity for the Last Time Buy (LTB) to the repair firm. An LTB of new parts can cover demand for spares for the given ESP period. Alternatively repairing both failed parts and particularly phase-out returns is possible as well. Repair is cheaper, but it introduces uncertainty. Service-level Agreements (SLAs) define the relationship between individual customers and repair firm.

Many parameters are written down in SLAs, but some trade-offs aspects remain difficult. This is clarified by introducing decision modelling. Contrary to existing literature, we model both types of returns as one of them (phase-outs) directly impacts the size of the installed base, leading to a non-stationary system. A dispose-down level is needed because oversupply of returns is likely to occur. We introduce subintervals to deal with net demand before phase-out occurrences. We also explicitly model the repair process with yield factors and optimize inbound and outbound inventories with PUSH –PULL policies. The model translates lacking performance into backorder and penalty costs. We prefer this to the service level approach because financial consequences of malperformance can be modelled unilateral with other cost. We develop an efficient heuristics model based on cost-level optimization.

Policies are optimized in a single customer, single-item situation with subintervals. The approximate decision rules prove to be both efficient and close to optimal. The resulting policies are not trivial, but in general PULL policies perform better, if well optimized, because they wait longer to decide. Postponing repair and scrapping of returns pays off, due to exploiting actual information on the dynamic installed base, related failure rates and the inbound inventory of (phase-out) repairables.

Customers strive for maximal certainty, whilst repair firms (such as ConRepair) aim for profit optimization. But the SLA and corresponding policies achieve the opposite effect. From the perspective of the repair firm, endogenous control policy variables primarily optimize customer service and exogenous variables mainly determine the company's profit. It is remarkable, that cost benefits are passed on fully to the customer. This calls for a reconsideration of the ESPs terms and conditions. Formalizing decision making in a model makes trade offs more visible, and we recommend that all parameters, such as penalty structures, should be part of an SLA.

Another suggestion for further research is to extend the model into multi-item situations. Also, one could investigate, for example, options for selective, condition-based take-back, applying local testing at the customer site. Moreover, all sensitivity analysis in this paper is carried out presuming *a priori* information about parameter changes. Under these circumstances, policies can be adapted in time. But suppose that information becomes available after the Last Time Buy has been acquired. Does the presence of phase-outs strengthen robustness of the system, or does it instead constitute another source of uncertainty itself? Another dimension is differentiation of service levels into e.g. platinum, gold and economy contracts. It would be interesting to investigate how control policies are influenced by such differentiation. Finally, optimization procedures could be extended with exact solutions, service-level models or hybrid PUSH-PULL. In this, a 'fair' calculation of penalty costs, based on collateral cost in the supply chain or insurance fees, is a fruitful area.

In general, the interaction and synergy of service - and reverse logistics deserves more exploration. This paper is a first step.

Acknowledgements

We would like to express our gratitude to dr M.C. van der Heijden of Twente University of Technology, Netherlands, for his critical comments and good advices.

This research is partially sponsored by Transumo ECO project, proj. no. GL05022b.

References

- Blumberg, D. F., 1999, Strategic examination of reverse logistics & repair service requirements, needs, market size and opportunities, *Journal of Business Logistics*, 20(2), pp. 141-159. Customer update 2005, obtained via Reverse Logistics Association.
- Cattani, K.D., 1997. Demand pooling effects of a universal product when demand is from distinct markets. Ph.D. Dissertation, Department of Industrial Engineering and Engineering Management, Stanford University.
- Cattani, K.D. and G.C. Souza, 2003. "Good buy? Delaying end-of-life purchases," *European Journal of Operational Research*, vol. 146, 216-228.
- David, I., E. Greenstein and A. Mehrez, 1997. "A dynamic programming approach to continuous-review obsolescent inventory problems." *Naval Research Logistics* 44, 757-774.
- Fortuin, L., 1980, 1981. "Five popular density functions: A comparison in the field of stock control models. *Journal of the Operations Research Society*, 31-10, 937-942, plus corrigenda 32-1, 425.
- Guide, D. and L.N. Van Wassenhove, 2006. "Closed-loop supply chains: introduction to the feature issue", *Production and Operations Management*, 15(4), pp.471-472
- Kennedy, W.J, J.W. Petterson and L.D. Fredendall, 2002. "An overview of recent literature on spare parts inventories," *International Journal of Production Economics*, 76, 201-215.
- Krikke, H.R., H.M. le Blanc and S. van de Velde, 2004. "Product modularity and the design of closed-loop supply chains." *California Management Review*, 46(2), 23-39.
- Mostard, J. and R.H. Teunter, 2006. "The newsboy problem with resalable returns: A single period model and case study." *European Journal of Operation Research*, vol. 169, n°1, 81-96.
- Muckstadt, J.A., 2005. *Analysis and Algorithms for Service Parts Supply Chains*. Springer, New York, ISBN 0-387-22751-6.
- Sherbrook, C.G., 2004. *Optimal inventory modeling of systems: Multi-echelon techniques*, Kluwer Academics, Deventer, ISBN 1-4020-7849-8.

- Silver, E.A., D.F. Pyke and R. Peterson, 1998. *Inventory Management and Production Planning and Scheduling*, 3rd ed., John Wiley & Sons, New York, ISBN 978-0-471-11947-0.
- Teunter, R.H. and W.K. Klein Haneveld, 1998. "The final order problem," *European Journal of Operational Research* 107, 35-44.
- Teunter, R.H. and L. Fortuin, 1998. "End-of-life service: A case study," *European Journal of Operational Research* 107, 19-34.
- Teunter, R.H. and L. Fortuin, 1999. "End-of-life service", *International Journal of Production Economics* 59, 487-497.
- Spengler, T. and M. Schroter, 2003. "Strategic management of spare parts in closed-loop supply chains— A system dynamics approach. *Interfaces* 33(6): 7–17.
- Van der Laan, E.A., R. Kuik, G. Kiesmüller, D. Vlachos and R. Dekker, 2004. "Managing product returns," in: R. Dekker, M. Fleischmann, K. Inderfurth, L.N. van Wassenhove (eds.), *Reverse Logistics: Quantitative Models in Closed-Loop Supply Chains*, Springer, Berlin, ISBN 3-540-40898-0
- Van der Laan, E.A., M. Salomon and R. Dekker, 1998. "An investigation of lead-time effects in manufacturing/remanufacturing systems under simple PUSH and PULL control strategies." *European Journal of Operational Research* 115, 195–214.
- Van der Laan, E.A., M. Salomon, R. Dekker and L.N. Van Wassenhove, 1999. "Inventory control in hybrid systems with remanufacturing." *Management Science*, 45(5): 733–747.

Table 1: Model parameters and base-case data (retrieved 2004-2007)

Model inputs		
H	Planning horizon (weeks)	150 weeks
$IB(t)$	Installed base at time t (parts)	$0 \leq t < 50$: 500 parts $50 \leq t < 100$: 400 parts ($\tau_1 = 50$) $100 \leq t < 150$: 300 parts ($\tau_2 = 100$)
μ	Failure rate (parts per week; total failures follow a piecewise homogenous Poisson process with rate $\mu IB(t)$)	0.02 parts per week
L	Repair lead time (weeks; deterministic)	1 week
h_m	Carrying cost for newly acquired parts	0.042 per part/per week
h_r	Carrying cost for repaired parts	0.026 per part/per week
h_n	Carrying cost for repairable parts	0.010 per part/per week
p_m	Price of newly acquired part	10 per part (normalized)
p_r	Price of repaired part	8 per part (normalized)
c_m	Unit cost of newly acquired part	8 per part (normalized)
c_r	Unit cost of repaired part	4 per part (normalized)
y_F	Probability that repair is successful for failed parts	0.7
y_P	Probability that repair is successful for phase-out parts	0.9
c_b	Backorder penalty per part	5 euro per part per week
c_f	Fail-to-deliver penalty per part	100 euro per part
$1 - \beta_p$	Acceptable probability of being out-of-stock just before phase-out occurrence, specified by management	0.01
Control decision variables		
Q	Last Time Buy Quantity (PUSH and PULL)	
$S(t)$	Repair-up-to level at time t (PULL)	
$U(t)$	Dispose-down-to level at time t (PUSH and PULL)	
System process variables		
$D(a,b)$	Demand during time interval $(a,b]$ (Poisson process)	
$R(a,b)$	Repairable returns during time interval $(a,b]$ (Poisson process)	
$ND(a,b)$	Net demand ($= D(a,b) - R(a,b)$) during time interval $(a,b]$ (Poisson process)	
$M(t)$	Inventory of newly acquired parts at time t	
$R(t)$	Inventory of repaired parts at time t	
$N(t)$	Inventory of repairable parts at time t	
$IP(t)$	Inventory position at time t ($= M(t) + R(t) + \text{repair in progress} - \text{backorders}$)	

Table 2: Parameters that determine the disposal quantity under PUSH and PULL

	PULL	PUSH
v_t	$h_n(H - t)$	$c_r + h_r(H - t)$
p	$p_r - c_r$	p_r

Table 3: Accuracy of the PULL and PUSH heuristics for varied repair yield and $\beta_p = 0$.

Disposal option is disabled; all other parameters are according to the base case.

Y_F	PULL				PUSH			
	Q (heuristic)	Q (optimal)	Net profit (heuristic)	Net profit (optimal)	Q (heuristic)	Q (optimal)	Net profit (heuristic)	Net profit (optimal)
0.3	701	707	13.10	13.28	700	700	9.38	9.38
0.5	457	466	20.52	20.69	456	456	17.55	17.55
0.7	212	222	27.06	27.27	210	211	24.83	24.84
0.9	49	50	23.54	23.54	49	49	23.18	23.18

Table 4: Performance of the PULL and PUSH heuristics as a function of service-level requirement β_p ; all other parameters are according to the base case.

β_p	PULL		PUSH	
	Q	Net profit	Q	Net profit
0.900	212	27.06	211	24.84
0.990	225	27.25	225	25.20
0.995	229	27.21	229	25.25
0.999	238	27.06	238	25.19

Table 5: The value of disposal for varied timing of the first (and only) phase-out occurrence τ_1 .

τ_1	PULL		PUSH	
	Net profit (disposal)	Net profit (no disposal)	Net profit (disposal)	Net profit (no disposal)
25	22.70	22.70	19.75	19.90
50	25.11	25.12	22.50	22.58
75	26.53	26.40	24.38	22.09
100	27.35	27.12	25.18	21.06
125	27.82	27.70	25.24	19.97

Table 6: Performance of PUSH and PULL in terms of costs and benefits per week.

	PULL		PUSH	
Sales new	15.00		15.00	
Sales repaired	51.94		51.87	
Production/purchase costs	12.00		12.00	
Repair costs	26.16		26.01	
New part. Inv. Costs	0.76		0.76	
Repaired part Inv. Costs	0.19		1.72	
Repairables Inv. Costs	0.35		-	
Backorder costs	0.04		0.17	
Failed to deliver costs	0.17		1.00	
Subtotal	66.94	39.68	66.87	41.67
Net profit	27.25		25.20	
Q	225		225	
Percentage of returns disposed	3.6%		4.3%	
Fraction of parts delivered directly from stock	99.65%		99.45%	
Fraction of parts delivered	99.98%		99.88%	

Figure 1: The service process and reverse logistics with Last Time Buy opportunity

Figure 2: Schematic representation of the model

Figure 3: Inventory processes as a function of time under PUSH (a) and PULL (b) control without phase-outs

(a)

(b)

Figure 4: The effect of phase-outs under PUSH control: an early phase-out (a) and a late phase-out (b).

Figure 5: Numerical example for base-case scenario under PUSH (a) and PULL (b) control

(a)

(b)

Figure 6: Performance of PULL and PUSH as a function of the unit repair cost c_r

Figure 7: Impact of repair yield

(a)

(b)

Figure 8: Financial impact of phase-out volume

(a)

(b)

Figure 9: Impact of phase-out volume with respect to Last Time Buy quantity and service levels

(a)

(b)

Figure 10: Impact of phase-out timing

(a)

(b)