



MIT Open Access Articles

Inventory Management in a Consumer Electronics Closed-Loop Supply Chain

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	Calmon, Andre P. and Stephen C. Graves. "Institute for Operations Research and the Management Sciences (INFORMS)." <i>Manufacturing & Service Operations Management</i> (October 2017): 509-712 © 2017 INFORMS
As Published	http://dx.doi.org/10.1287/msom.2017.0622
Publisher	Institute for Operations Research and the Management Sciences (INFORMS)
Version	Author's final manuscript
Citable link	https://hdl.handle.net/1721.1/122357
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

Inventory Management in a Consumer Electronics Closed-Loop Supply Chain

Andre P. Calmon
INSEAD, andre.calmon@insead.edu

Stephen C. Graves
MIT, sgraves@mit.edu

This Version was created on October 18, 2016

The goal of this paper is to describe, model, and optimize inventory in a reverse logistics system that supports the warranty returns and replacements for a consumer electronic device. The context and motivation for this work stem from a collaboration with an industrial partner, a Fortune 100 company that sells consumer electronics. The reverse-logistics system is a closed-loop supply chain: failed devices are returned for repair and refurbishing; this inventory is then used to serve warranty claims or sold through a side-sales channel. Managing inventory in this system is challenging due to the short life-cycle of these devices and the rapidly declining value for the inventory. We examine an inventory model that captures these dynamics. We characterize the structure of the optimal policy for this problem for stochastic demand and introduce an algorithm to calculate optimal sell-down levels. We also provide a closed-form policy for the deterministic version of the problem, and we use this policy as a certainty-equivalent approximation to the stochastic optimal policy. Finally, using numerical experiments, we analyze the sensitivity of this system to changes in various parameters, and we also evaluate the performance of the certainty-equivalent approximation using data from our industrial partner.

Keywords: Inventory Management; Closed-Loop Supply Chains; Reverse Logistics; Sustainability

Electronic copy available at: <http://ssrn.com/abstract=2622405>

Inventory Management in a Consumer Electronics Closed-Loop Supply Chain

Andre P. Calmon

Technology and Operations Management, INSEAD, andre.calmon@insead.edu

Stephen C. Graves

A. P. Sloan School of Management, Massachusetts Institute of Technology, sgraves@mit.edu

The goal of this paper is to describe, model, and optimize inventory in a reverse logistics system that supports the warranty returns and replacements for a consumer electronic device. The context and motivation for this work stem from a collaboration with an industrial partner, a Fortune 100 company that sells consumer electronics. The reverse-logistics system is a closed-loop supply chain: failed devices are returned for repair and refurbishing; this inventory is then used to serve warranty claims or sold through a side-sales channel. Managing inventory in this system is challenging due to the short life-cycle of these devices and the rapidly declining value for the inventory. We examine an inventory model that captures these dynamics. We characterize the structure of the optimal policy for this problem for stochastic demand and introduce an algorithm to calculate optimal sell-down levels. We also provide a closed-form policy for the deterministic version of the problem, and we use this policy as a certainty-equivalent approximation to the stochastic optimal policy. Finally, using numerical experiments, we analyze the sensitivity of this system to changes in various parameters, and we also evaluate the performance of the certainty-equivalent approximation using data from our industrial partner.

Key words: Inventory Management, Closed-Loop Supply Chains, Reverse Logistics, Sustainability

History: This Version was created on October 18, 2016

1. Introduction

In this paper we examine an operational problem arising at a Wireless Service Provider (WSP) that is a Fortune 100 company. This WSP sells consumer electronics and offers voice and data plans to its customers. For this company, as for many other retail businesses, the management of warranty claims and regret returns is a key issue.

Our industrial partner operates with two warranty contracts in place: (i) the customer warranty, and (ii) the Original Equipment Manufacturer (OEM) warranty. The customer warranty is a contract between WSP and the customer, and protects the customer against any defects in the purchased product and also provides the customer a period for “regret returns.” In addition, the customer warranty has strict requirements: when a warranty claim is filed, a new or refurbished item is immediately shipped by our industrial partner to the customer together with a pre-paid shipping label so that the customer can return the original unit. Thus, a replacement item is sent

before the original item is received. The OEM warranty, on the other hand, is a contract between the OEM and the WSP, covering each device purchased from the OEM. This warranty is slow: the WSP must receive and then ship each defective product to the OEM for repair, which can then take several weeks before it (or a replacement) is returned to the WSP.

The WSP has a dedicated reverse logistics facility to process customer warranty claims and regret returns. Devices that are returned to this facility are typically repaired (usually by the OEM) and/or refurbished and then held in inventory. This inventory can then be used to serve as replacement devices for future customer warranty claims. The WSP can also sell refurbished devices into a side-channel. This reverse-logistics system acts as a closed-loop supply chain (CLSC), as returned devices get reused to satisfy customer warranty claims.

We study how to manage this device inventory to serve the customer warranty claims as cost-effectively as possible. There are two control decisions in each time period: how many new devices to buy, and how many refurbished devices to sell. Several important features characterize this problem: the demand for replacement devices is non-stationary, not all devices sent to the OEM can be repaired, and the device value degrades very rapidly over its short product life-cycle.

In this paper we address these challenges by introducing a finite-horizon discrete-time stochastic model for the inventory management problem faced by the WSP. This model incorporates the short life-cycle and the fast value depreciation of the device, and makes minimal assumptions on the demand and replenishment distributions. We assume that the replacement requests in each period follow a non-stationary stochastic process and we prove the structure of the optimal policy: in each period the optimal policy buys inventory myopically and then has a sell-down level for determining how much inventory to sell. Although, the optimal sell-down level in each period depends on the distributions of the demand and replenishment processes, as well as on the time-varying prices of new and refurbished devices, we present a simple algorithm for computing the optimal sell-down levels that can be implemented using a Monte Carlo simulation.

We then consider a deterministic version of this problem, and provide a closed-form expression for the optimal quantities to purchase and sell. We use the deterministic model to generate a worst-case bound on the inventory investment required when there is limited knowledge about the failure rates for a new device. We leverage these results to provide a certainty-equivalent heuristic for the stochastic problem where we replace the demand and replenishment processes by their average values.

Through numerical experiments, we analyze the sensitivity of the inventory management policy with respect to changes in different parameters such as seed-stock, OEM lead time, and warranty duration. We analyze the performance of the certainty-equivalent policy and observe that when

the number of devices is large (which is the case for our partner WSP) this approximation achieves a near-optimal performance and is sufficient for practical applications. Finally, we analyze the performance of the certainty equivalent policy using real-world data from two devices sold by the WSP. We observe that the certainty-equivalent policy captures over 95% of the clairvoyant profit.

The remainder of this paper is as follows. In Section 2 we present a review of the literature of inventory management in reverse logistics systems. In Section 3 we formulate the inventory management problem and discuss the dynamics of this system. In Section 4 the stochastic model for this problem is introduced and analyzed, the optimal policy is proved and an algorithm for calculating the optimal sell-down level is presented. In Section 5 we examine the deterministic version of the inventory management problem, present a worst-case bound, and introduce the certainty-equivalent heuristic. Finally, in Section 6 we present our numerical experiments.

2. Literature Review

This paper builds upon a vast literature on inventory management for closed-loop supply chains and for warranty replacements. Souza (2013), Guide and Van Wassenhove (2009), and Fleischmann et al. (1997) each present an overview of the literature on CLSCs. The main features that distinguish our work are as follows: (i) we consider time varying costs and prices over a finite horizon; (ii) given the context that we are analyzing, we do not model the repair process in any detail, since quality sorting and repair is done by the OEM, under a warranty agreement; (iii) we do not need any specific assumptions on the demand and replenishment (or arrival) processes for the CLSC, and can allow these processes to be correlated and non-stationary; (iv) in our CLSC set-up, there is no backlogging of demand and demand must be met before the return of a failed unit.

The benefits and drawbacks of different reverse channel structures are discussed by Savaskan et al. (2004). In fact, the structure of the reverse channel in our paper corresponds to the decentralized setting in Savaskan et al. (2004) where the retailer (in our case, the WSP) is responsible for collecting the failed devices.

Huang et al. (2008) offer a detailed overview of the literature in warranty management as it relates to inventory management. The paper develops and analyzes a multi-period single-item inventory model for which there is stochastic demand for both new and replacement items. However, unlike our model, Huang et al. (2008) assume that failed devices are not repaired and returned to inventory for reuse.

Our inventory management problem allows for both demand and returns, and has both replenishment and disposal control decisions. In this way it is related to the cash-management literature. Cash management entails deciding how much cash to keep in hand vs. invest in financial securities,

in light of uncertain future cash flows. A seminal paper in this area is Constantinides and Richard (1978), who characterize the form of the continuous-time, infinite horizon optimal policy when the demand for cash is modeled as a Weiner process. Feinberg and Lewis (2005) and Chen and Simchi-Levi (2009), build upon this literature to consider inventory management problems where demand can be positive or negative. This set-up is similar to ours, although time-varying costs are not considered.

The set-up of our model is similar to a dynamic lot size model where demands can be positive or negative, and disposals can occur. Beltran and Krass (2002) analyze a deterministic version of this problem, and present an algorithm to determine the optimal procurement and disposal strategies. Although we do not consider fixed disposal costs, and although our general model assumes a stochastic setting, where demand and arrivals are non-stationary stochastic processes, the deterministic model discussed in Section 5 is a sub-case of the model that Beltran and Krass (2002) study. However, the additional assumptions that we make, namely, no backlogging and non-increasing costs and prices, allow us to obtain a closed form solution to the optimal sell-down levels, which we believe to be of practical importance. In fact, a simplified version of this model was adopted by the WSP, as described in Petersen (2013).

Heyman (1977) also analyzes a similar set-up to ours. He analyzes a problem in which a firm leases equipment to a customer and then receives the equipment back when the customer terminates their service. If the returned item can be refurbished, it is stored in inventory to satisfy future demand. Given this set-up, Heyman (1977) addresses the question: *“when is the inventory level so high that retaining another returned item is uneconomical because the cost of repairing the item and holding it in inventory until it can be used is greater than the savings that can be obtained by satisfying a future demand with a repaired item rather than a new item?”*. He models this problem as a single-server queue, obtaining the exact optimal policy when return and demand processes are Poisson. Our work differs from this paper in that the demand and return processes are not stationary, and the relevant costs are time varying.

Pinçe et al. (2008) also consider a system with stochastic demand and supply, and with non-varying costs. As in Beltran and Krass (2002), they assume a fixed disposal cost. Their analysis assumes that the net-demand follows a stationary Brownian-motion process. This assumption allows for a detailed analysis of the benefits of the disposal option, in particular in the presence of fixed disposal costs.

We also build upon the literature on CLSCs for electronic products. Ferguson et al. (2011) consider the asset recovery process at IBM, in which returned devices can either be dismantled for parts or remanufactured. Each part of a product could have a different, stationary demand

distribution. Similar to our analysis, they show through simulation that simple heuristics work well. Khawam et al. (2007) present a simulation approach to obtain inventory management policies for Hitachi. Also, Toktay et al. (2000) propose a closed-loop queuing model that captures dynamics of recycling at Kodak, and discuss a control-policy for the procurement of components. More recently, Pınçe et al. (2016) consider a dynamic allocation problem where an OEM needs to dynamically choose if consumer returns that are refurbished will be used to fulfill warranty claims or if they will be remarketed as refurbished products. In this setting, the OEM also chooses the price of new and refurbished products.

Bayiz and Tang (2004) study the problem of managing the inventory of multiple types of thermoluminescent badges that are refurbished and leased to customers that work in environments subject to radioactive exposure (such as X-ray labs or nuclear facilities). The set-up in the paper is similar to ours in the sense that badges are refurbished and sent as replacements to customers before the used badges are returned. If not enough refurbished badges are available, new badges are acquired and sent to the customer. Unlike our set-up, excess inventory cannot be sold. The authors formulate the problem of determining inventory levels as a Mixed Integer Program. In Fleischmann et al. (2003) the authors examine the design challenges for managing a CLSC for spare parts at IBM. In their set-up, however, item returns do not immediately trigger demand for replacement items.

The results in this paper can also be related to literature on the control of repairable inventory. Allen and D'Esopo (1968) evaluate a CLSC where items fail according to a Poisson processes. However, not all items can be repaired; so new items might need to be procured into stock in order to satisfy demand. They characterize the performance of a reorder-point, reorder-quantity policy. Simpson (1978) considers a repairable inventory system with stochastic demand and with control decisions on the number of units to repair as well as the number of new units to buy in each period. For the case of instantaneous repair and purchase, the paper characterizes the form of the optimal policy under general assumptions on demand and return processes. A major difference with our work is that Simpson assumes that unmet demand is back-ordered.

More recently, Kiesmüller and van der Laan (2001) analyze a finite horizon refurbishment problem with predictable returns where backorders are allowed and demand is distributed as a Poisson random variable. Tao and Zhou (2014) also examine balancing policies for remanufacturing systems, and show that the two-approximation from Levi et al. (2008) holds in this set-up. Also, Zhou et al. (2011) consider a set-up where a single type of item is manufactured, but remanufacturing depends on the condition of the item once it is returned. In our set-up, we do not take product conditions into account; if the warranty applies to a failure, then the OEM must refurbish the

device or provide a replacement. Furthermore, we analyze a setting where prices and costs are non-stationary.

de Brito and van der Laan (2009) analyze the impact of imperfect information about regret returns on inventory management performance. Namely, they consider a model where demand for new products can be satisfied using returned products and that demand backlogging is allowed. Their numerical experiments indicate that, given their set-up, knowledge of the aggregate probability of return and of the maximum return horizon is sufficient to obtain a near-optimal cost performance if the return probability is low.

3. Problem Description

There are three actors in this reverse supply chain: the customer who purchases mobile devices, the Wireless Service Provider (WSP), and the Original Equipment Manufacturer (OEM). The forward supply chain between these actors is structured as a traditional supply chain. The customer purchases a mobile device from the WSP and usually subscribes to a wireless plan that includes voice and data services provided by the WSP. The WSP, whom we focus on, is both a retailer of mobile devices and a provider of wireless services to customers. Most of the revenue of the WSP comes from the services it provides, and one of its corporate goals is to ensure that the time customers spend “disconnected” from its network is minimized, since a disconnected customer means lost revenue and loss of customer goodwill. The OEM, at the top of the forward chain, acts as a wholesaler, and sells mobile devices in bulk to the WSP.

The reverse supply chain, however, is a CLSC. To understand the dynamics of this CLSC, we first need to describe the two warranty contracts that are in place in this system, namely, the customer warranty, offered by the WSP to the device user, and the Original Equipment Manufacturer warranty, offered by the OEM to the WSP.

The customer warranty is designed to minimize the time a customer spends without a working device. This warranty has a base length of 12 months, but can be extended if the customer purchases an extended warranty plan. If a device presents a problem, the user contacts the WSP’s call center where a technician provides assistance and tries to resolve the issue. If the customer’s issue cannot be resolved over the phone, a warranty claim is filed and a replacement device is immediately shipped to the user. In most cases the user receives a replacement device in less than 72 hours from the time a warranty claim is received by the WSP. In addition, the replacement device comes with a pre-paid shipping label for the user to use to mail the defective device to the WSP’s reverse logistics facility where the device goes through testing and triage. A key feature of this system is that, whenever possible, the replacement device is a remanufactured or refurbished device from

some previous warranty claim or regret return for the same model. If there are no refurbished devices available, the WSP will send the customer either a new device of the same model or will give the customer an upgrade.

The customer warranty also allows for regret returns, such that the user can return a new mobile device within a few weeks after purchase for a complete refund, net of a stocking fee. These returned devices are shipped from the retail site to the reverse logistics facility.

The OEM warranty is designed to protect the WSP from manufacturing defects. This warranty usually has a 12-month duration and starts when the WSP purchases the devices from the OEM. Unlike the customer warranty, the OEM warranty is not designed to minimize the time that the WSP spends without a working device. Rather, this warranty specifies a lead time for the OEM to remanufacture or replace the faulty device, usually a few weeks. For some OEM's, this contract also stipulates a seed-stock of new devices that the OEM provides to the WSP to be used as replacement devices for the initial warranty claims and to cover other losses in this CLSC.

There can be loss in the system due to the structure of the warranties and due to customer service considerations. An example is no trouble found (NTF) devices, which are devices where the customer claims that there is an issue with the device, but neither the WSP nor the OEM can replicate the problem. In this case, the OEM does not provide a replacement to the WSP, but the WSP does replace the user's device. Another source of loss is manufacturing defects that are detected by the WSP and disputed by the OEM. Loss can also occur due to out-of-warranty returns, i.e., defective devices with an expired OEM warranty that are returned by customers who are still covered by the customer warranty. In all of these cases, it may be too expensive for WSP to repair the device and hence there may be a loss to the CLSC.

3.1. Operational flows of the reverse logistics facility

The WSP reverse logistics facility processes the warranty claims and regret returns, and maintains an inventory of refurbished devices. This facility can also repair a failed device if the defect is small or not covered by the OEM warranty. Since a large volume of warranty claims are filed per day (usually over ten thousand), and regret returns can be as much as 5 to 10% of purchases, there are over one million devices in this facility at any given time. The dynamics of this system are depicted in Figure 1. Grey arrows denote the flow of items from customers into the system, dark arrows represent items leaving the system, and white arrows correspond to the flow of items within the system. The dashed box outlines the limits of the WSP's reverse-logistics facility that processes and stores returned items. The elements of the flow are as follows:

1. When a warranty claim is filed, the replacement is immediately shipped to the user from the inventory within the WSP's reverse logistics facility. Note that the customer receives a replacement *before* returning the original item.

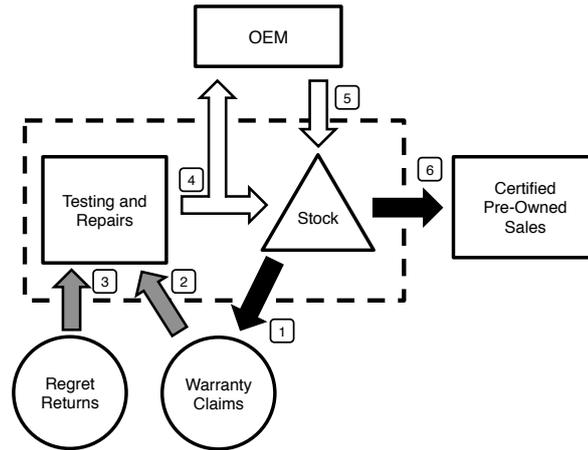


Figure 1 Dynamics of the CLSC. The dashed line outlines the reverse logistics facility.

2. Upon receipt of the replacement, the customer mails back the failed unit to the reverse logistics facility.

3. A second source of inventory is regret returns. Customers that return an item through this channel do not receive a new product.

4. Products received from warranty claims and regret returns go through a triage process that generally leads to one of four possible outcomes: (i) no problem is found and the product is sent to inventory after a refurbishment; (ii) there is a minor problem, and the device is repaired at the facility; (iii) there is a major problem that is covered by the OEM warranty, and a warranty claim is filed with the OEM; (iv) there is a major problem that is not covered by the OEM warranty, and the defective device is either repaired at the facility or disposed.

5. If the triage process determines that the defect is covered by the OEM warranty, the device is shipped to the OEM for repair and refurbishment. It usually takes a few weeks for the refurbished device (or a replacement) to return to the reverse logistics facility. Also, when a product is launched, the OEM provides an initial seed stock of new devices to the WSP, and these devices are kept in stock at the reverse logistics facility.

6. The refurbished devices held by the reverse logistics facility can be sold through side-sales channels. These channels include sales to other WSPs that sell these devices in other markets, to companies that recycle components of these devices, and also through certified pre-owned sales programs. Side-sales not only generate revenue, but also act as an inventory control mechanism, allowing the facility to reduce inventory levels, especially towards the end of a product's life cycle.

Since this is a closed-loop system, there is a correlation between the demand for refurbished devices (flow 1 in Figure 1) and the arrivals of refurbished devices to the inventory (from flow 4

and flow 5); furthermore, as noted earlier, there will be some loss as not all returned units are recoverable.

Due to the warranty contract, every customer warranty claim is fulfilled immediately, preferably using refurbished items. Backlogging is not allowed and if there are no refurbished products in stock, the WSP will send the customer either a new or upgraded device. Thus, the primary control that exists for managing inventory is the number of devices that are sold through side-sales channels, and the number of new devices brought into the system either through direct purchases from the OEM or as seed stock. Furthermore, the value of this inventory depreciates quickly over time; hence the holding cost for this inventory corresponds mainly to the depreciation rate for the side-channel price. The shortage cost for this inventory is the cost to acquire and ship a new or upgraded unit to the customer. Balancing these two costs in the face of the non-stationarity of the demand for replacement devices, the short product life-cycle of mobile devices, and the closed-loop nature of this system is a significant challenge.

3.2. Inventory dynamics and costs

There are two operational decisions involved in managing this system: (i) the decision of how many devices should be kept in the inventory at the reverse logistics facility, herein called the inventory management problem; and (ii) the decision of which devices in inventory should be matched to which customers, herein called the warranty matching problem. In order to obtain tractable policies for managing this supply chain, we uncouple these decisions and address them separately. Although, uncoupling these two decisions implies that the policies discussed in this paper are not optimal in a system-wide sense, it will lead to policies for which we can glean both theoretical and managerial insights. In this paper we consider the inventory management problem; the warranty matching problem is addressed in Calmon (2015). Furthermore, the numerical experiments in Calmon (2015) indicate that making these decisions independently does not lead to a significant additional cost.

To model the dynamics in Figure 1 we assume that time is discrete and that, at each time instant t , the demand for refurbished devices, equivalent to the number of warranty claims, is denoted by d_t . We specify the number of devices that arrive to stock and that are available to fulfill warranty claims by the process a_t . This process captures all inventory flows into the reverse logistics facility and includes regret returns that have been refurbished at the facility, as well as all units supplied by the OEM (i.e., seed stock plus repaired or replacement devices in response to earlier warranty claims). Thus, if the arrival process of devices from the OEM is w_t , and the arrival process for regret returns is r_t , we have that $a_t = r_t + w_t$. Finally, we denote the on-hand inventory at the end of period t as x_t .

The number of devices purchased at time t is denoted by u_t^+ , and devices sold at time t is denoted by u_t^- . We assume that the cost, c_t , of purchasing a new device and the revenue, p_t , obtained from selling a refurbished device in the side-sales channel are non-increasing, i.e., $c_t \geq c_{t+1}$ and $p_t \geq p_{t+1}, \forall t$. In addition, we assume that $c_t \geq p_t, \forall t$, such that there are no strategic buying opportunities. Finally, we assume that there is a non-increasing inventory holding cost of h_t per device per period, where $h_t \geq 0, \forall t$.

The justification for these cost assumptions comes from the application context at the WSP. The cost for obtaining new devices remains relatively stable over a product's life time, and does not increase. The price at which refurbished devices can be sold decreases over the life time. Indeed, the impetus for the research project came from the recognition that the effective holding cost for the refurbished inventory at the reverse logistics facility is very high, due to the rapid depreciation in its sales value. The WSP wanted to know at what point it was safe to sell refurbished devices into side channels, so as to avoid having excess devices relative to the remaining warranty claims. In addition to the holding cost due to value depreciation, there can be a physical holding cost for inventory, captured by h_t .

The inventory balance equation is given by

$$x_t = x_{t-1} + a_t - d_t - u_t^- + u_t^+$$

where we assume that the control decisions are made in each period after observing the arrival a_t and demand d_t . Furthermore, as no backlogging is permitted and $c_t \geq p_t, \forall t$, we can assert bounds on the purchase and sale amounts:

$$\begin{aligned} u_t^+ &\geq \max(0, d_t - a_t - x_{t-1}) \\ 0 &\leq u_t^- \leq \max(0, x_{t-1} + a_t - d_t). \end{aligned}$$

The objective of the inventory management problem is to maximize the net revenue over a T period horizon, starting at $t = 1$:

$$\max \sum_{t=1}^T E_{\{d_t, a_t\}} [p_t u_t^- - c_t u_t^+ - h_t x_t],$$

subject to the inventory balance equation, the constraints on the controls, and an initial inventory given by x_0 . The expectation is taken over the arrival and demand processes.

The demand and arrival processes are correlated, non-stationary, and their distributions are likely to be unknown at the launch of the device. With the notation and assumptions in hand, in the next section we analyze the stochastic version of this problem.

4. Stochastic Model and Optimal Policy

In this section, we characterize the optimal policy and how to determine it. We assume that we are given a non-stationary stochastic process $\{d_t, a_t | t = 1, \dots, T\}$ for both device demand and arrivals. The specification of this stochastic process, in general, can be quite complex, and will depend on some informational state that reflects the sales process of the device and its failure distribution, as well as the characteristics of the closed-loop inventory system. For instance, the arrival process will be highly correlated with both the demand process (reflecting warranty returns that are refurbished after an OEM lead time) and new sales, a percentage of which become regret returns. We formulate the inventory management problem as the following optimization problem where we assume that the initial inventory is 0 (this assumption is without loss of generality as any initial seed stock can be captured in a_1):

$$\begin{aligned}
 & \text{maximize} && \sum_{t=1}^T E_{\{d_t, a_t\}} [p_t u_t^- - c_t u_t^+ - h_t x_t] \\
 & \text{s.t.} && x_t = x_{t-1} + a_t - d_t - u_t^- + u_t^+, \forall t = 1 \dots T \\
 & && x_0 = 0 \\
 & && x_t, u_t^-, u_t^+ \geq 0, \forall t = 1 \dots T
 \end{aligned} \tag{1}$$

where $\{d_t, a_t | t = 1, \dots, T\}$ is a random process and each of the sequences of parameters $\{c_t\}$, $\{p_t\}$, and $\{h_t\}$ is non-increasing and non-negative, and $c_t \geq p_t, \forall t$. Here, in each period the state of the system is, in theory, very complex and includes (i) the current time t ; (ii) the on-hand inventory; (iii) the information available about the failure rate distribution; (iv) the realization of the net demand $d_t - a_t$; (v) the sales of new devices and the distribution of device ages in the market; (vi) the pool of devices currently under repair with the OEM (on-order inventory). We do not attempt here to model this complexity. Rather we will work with a simplification of the state space, which is aligned with the informational realities of practice. The intent is to get some insight into the structure of an optimal policy and to guide the development of implementable policies.

We will use the *net inventory* at time t defined by $\bar{x}_t = x_{t-1} - d_t + a_t$ as the system state; the optimal policy will be written as a function $\{\hat{u}_t^-(\bar{x}), \hat{u}_t^+(\bar{x})\}$. We note that there are no restrictions on the value of the net inventory, and in particular, it can be negative. We give the structure of the optimal policy in the following proposition.

PROPOSITION 1. *For some $v_t^* \geq 0$, the optimal policy for the stochastic inventory model in (1) is*

$$\begin{aligned}
 \hat{u}_t^+(\bar{x}_t) &= \max(-\bar{x}_t, 0), \\
 \hat{u}_t^-(\bar{x}_t) &= \max(\bar{x}_t - v_t^*, 0),
 \end{aligned}$$

where v_t^* depends on the distribution of demand and arrivals, and on the cost and price parameters.

All proofs are given in the Appendix.

For the optimal policy, in each period, we sell down to some level v_t^* that depends on the current informational state of the system. Because the cost of purchasing a new device is non-increasing, we only purchase items as is necessary to satisfy demand in the current period; that is, we purchase only when $\bar{x}_t < 0$. We note that the structure of the optimal policy is an artifact of the cost assumptions, namely that $\{c_t\}$, $\{p_t\}$, and $\{h_t\}$ are non-increasing and non-negative and that $c_t \geq p_t, \forall t$.

4.1. Calculating the Sell-Down Levels

The previous proposition does not provide any insight into how to calculate $\{v_t^*\}$.

Our goal is to calculate the sell-down level v_t^* at time $t < T$. We denote the expected profit in the horizon $[t, T]$ by $J_t(\bar{x}_t)$ as a function of the net inventory \bar{x}_t . For $\bar{x}_t \geq 0$ we can express $J_t(\bar{x}_t)$ by the recursion:

$$J_t(\bar{x}_t) = \max_{0 \leq v \leq \bar{x}_t} p_t(\bar{x}_t - v) - h_t v + E[J_{t+1}(v + a_{t+1} - d_{t+1})]. \quad (2)$$

Hence, the optimal sell down level, v_t^* can be found by solving

$$v_t^* = \arg \max_{v \geq 0} -p_t v - h_t v + E[J_{t+1}(v + a_{t+1} - d_{t+1})]. \quad (3)$$

We will use the above expression as the basis for an incremental analysis for finding the optimal sell-down level. In order for the incremental analysis to be valid, we need argue that the above expression within the maximization is concave in v . This is true since we can show by induction that J_{t+1} is concave and non-decreasing¹.

For the incremental analysis, we assume that we are given the optimal sell down levels $(v_{t+1}^*, \dots, v_T^*)$, a net inventory level \bar{x}_t , and that we have some candidate sell-down level v . We will analyze the marginal change in expected profit in periods $t + 1$ to T when the sell-down level is raised from v to $v + 1$, while keeping the sell down levels $(v_{t+1}^*, \dots, v_T^*)$ fixed.

First, if the net inventory level, \bar{x}_t , is less than or equal to v then increasing the sell-down level by one does not change the expected profits in $[t + 1, T]$. On the other hand, if $\bar{x}_t > v$, increasing the sell-down level to $v + 1$ might change the expected profit, since there will be one extra unit in inventory at the beginning of period $t + 1$. In this case, we observe that, subject to the same sell-down levels $(v_{t+1}^*, \dots, v_T^*)$, the inventory levels in the interval $[t + 1, T]$ for a sell down-level of v and $v + 1$ will differ by at most one unit for any sample path of device demand and arrivals.

¹The composition of a concave function with a concave non-decreasing function is concave (Boyd and Vandenberghe (2004)).

Furthermore, at some point the two inventory paths will be coupled, and will be identical until the end of the horizon T . This will occur in the first period after t where the inventory path for sell-down level v buys an item or when the inventory path for sell-down level $v + 1$ sells an item, whichever occurs first. One of these events is guaranteed to happen since $v_T^* = 0$; thus, the two inventory paths will meet, at the latest, in period T .

We will leverage this argument in order to calculate the optimal sell-down level. For some candidate sell-down level v let $\tau_t(v)$ be a random variable that represents the time period when the inventory path for sell-down level v couples with the inventory path for sell-down level $v + 1$. Namely, the inventory level for both sell-down levels will be the same in all periods after $\tau_t(v)$. Also, let the random variables $\bar{\tau}_t(v)$ and $\underline{\tau}_t(v)$ be defined as

$$\bar{\tau}_t(v) = \min \left\{ s \in [t+1, T] \left| \sum_{k=t+1}^s a_k - d_k \geq v_s^* - v \right. \right\},$$

and

$$\underline{\tau}_t(v) = \min \left\{ s \in [t+1, T] \left| \sum_{k=t+1}^s d_k - a_k > v \right. \right\}$$

Hence $\underline{\tau}_t(v)$ is the first time that the sample path with sell-down level v will need to buy an item to satisfy demand provided that no items are sold in the interval $[t+1, \underline{\tau}_t(v)]$; and $\bar{\tau}_t(v)$ is the first time that the sample path with sell-down level $v + 1$ will sell an item to reach a sell-down level provided that no items are purchased in the period $[t+1, \bar{\tau}_t(v)]$. By definition, it cannot be that $\bar{\tau}_t(v) = \underline{\tau}_t(v)$. We have that

$$\tau_t(v) = \min(\bar{\tau}_t(v), \underline{\tau}_t(v)).$$

If the net inventory level, \bar{x}_t is larger than v , the expected incremental profit in the interval $[t+1, T]$ by increasing the sell-down level by one will be denoted by $\delta_t(v)$ and is

$$\delta_t(v) = -h_t + \sum_{s=t+1}^T p_s \Pr(\bar{\tau}_t(v) = s, \tau_t(v) = s) + c_s \Pr(\underline{\tau}_t(v) = s, \tau_t(v) = s) - h_s \Pr(\tau_t(v) > s). \quad (4)$$

As explanation, the incremental profit from increasing the sell-down level by one depends on whether the incremental unit is subsequently sold (i.e., $\tau_t(v) = \bar{\tau}_t(v)$) or is used to avoid a purchase (i.e., $\tau_t(v) = \underline{\tau}_t(v)$). The total expected incremental profit in the interval $[t, T]$ will be positive if the marginal increase in future profits is higher than the revenue of selling a device in the current period, i.e.,

$$\delta_t(v) > p_t.$$

Hence, the optimal sell-down level will be

$$v_t^* = \min\{v \geq 0 | \delta_t(v) \leq p_t\}.$$

Thus, this formulation of the optimal sell-down level only depends on the future sell-down levels, on the structure of $\{c_t, p_t, h_t\}$ and on the distribution of arrivals and demand. Hence, the following algorithm can be used to find the optimal sell-down levels. This algorithm is a gradient search and is guaranteed to terminate due to the concavity of Equation 3.

Algorithm 1 Procedure for obtaining the optimal sell-down levels

```

 $v_T^* \leftarrow 0.$ 
for  $t = T - 1, \dots, 0$  do
  Set  $v_t = 0$ 
  while  $\delta_t(v_t) > p_t$  do
     $v_t \leftarrow v_t + 1.$ 
  end while
   $v_t^* \leftarrow v_t$ 
end for

```

The coupling probabilities in each time period can be calculated using a Monte-Carlo simulation, if the distribution of arrivals and demand are available, as well as the future sell-down levels. In fact, in Section 6, we will use this approach to calculate the optimal sell-down levels in our numerical experiments. Before reporting on our numerical experiments, we will analyze the case where the arrivals and demand are deterministic and known.

5. Deterministic Problem

In this section, we assume that both the demand process d_t and the arrival process a_t are deterministic and known over the T-period horizon, $t = 1, \dots, T$. This will lead to a simple closed-form solution for the optimal policy and we will leverage this in order to: (i) obtain an approximation for the optimal inventory management policy when demand and arrivals are stochastic; and (ii) obtain a worst-case bound for the amount of inventory to satisfy the demand for replacement devices. Beltran and Krass (2002) present a similar analysis to ours in the context of dynamic lot-sizing; however, we are able to obtain a closed-form solution to the optimal sell-down levels, based on our assumptions that $\{c_t\}$, $\{p_t\}$, and $\{h_t\}$ are non-increasing parameters, and $c_t \geq p_t, \forall t$.

In order to characterize the sell-down levels for each time period, we define

$$\tau_t^{\max} = \max \left\{ k \leq t \mid c_k - \sum_{s=t}^{k-1} h_s \geq p_t \right\},$$

and $\sum_{s=t}^{t-1} h_s = 0$. The definition of τ_t^{\max} is depicted in Figure 2. Thus, for any period $k \in [t, \tau_t^{\max}]$ we have that $c_k - p_t > \sum_{s=t}^{k-1} h_s$; that is, for $k \in [t, \tau_t^{\max}]$, it is more costly to sell a unit at time t

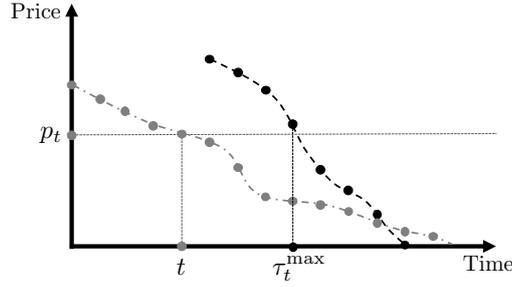


Figure 2 Definition of τ_t^{\max} . The gray line represents $\{p_t\}$ while the black line represents $\{c_k - \sum_{s=t}^{k-1} h_s\}$ for $k > t$.

and then purchase a replacement unit at time k , compared to just holding the unit in inventory over the time interval $[t, k]$. Similarly, for $k > \tau_t^{\max}$ it is cheaper to sell in period t and then replace in period k , compared to holding a unit from time t to time k . Consequently, an optimal policy will never sell a unit in period t and then purchase another unit in the interval $[t + 1, \tau_t^{\max}]$; and an optimal policy will never hold a unit from period t to period k , for $k > \tau_t^{\max}$. We use these observations to obtain a closed form expression for the optimal sell-down level, as stated in the next proposition.

PROPOSITION 2. For any $t_1 \leq t_2$, we define “net demand” $n(t_1, t_2)$ as

$$n(t_1, t_2) = \begin{cases} \max_{s \in [t_1+1, t_2]} \sum_{k=t_1+1}^s d_k - a_k, & \text{if } t_2 > t_1 \\ 0, & \text{if } t_2 = t_1. \end{cases}$$

Then, we have

$$v_t^* = \max(n(t, \tau_t^{\max}), 0), \quad \forall t \in [1, T].$$

For any time t , the sell-down level corresponds to the maximum net demand in the interval $[t + 1, \tau_t^{\max}]$. It is clearly uneconomic to sell below that level, as we would then need to buy at a higher cost in that time window. Similarly, if inventory is above the maximum net demand, then some items in stock will be held until some period k , $k > \tau_t^{\max}$, which is also uneconomic (as the unit could be sold at t and purchased in k).

We note that for this deterministic problem we do not need to know the future value of p_s for $s > t$ in order to determine the optimal sell down level in period t . We just need to know the future costs $\{c_s\}$ in order to determine τ_t^{\max} . With τ_t^{\max} in hand, we can find the sell down to level from Proposition 2.

For example, if the cost of purchasing a device is constant throughout its life-cycle, i.e., $c_t = c, \forall t$, and the holding cost is also constant, i.e., $h_t = h$, then $\tau_t^{\max} = \frac{c-p_t}{h} \forall t$, and the optimal solution in period t does not depend on the exact values of prices in the side-sales channel for periods after time

t ; we only need the fact that prices are non-increasing. This particular case is very useful from a managerial standpoint since estimating the cost of sourcing a new device and the revenue obtained from selling a refurbished device can be problematic in practice. Assuming that the purchase cost remains constant provides an upper bound on τ_t^{\max} , and hence, a conservative sell-down level that only depends on the assumption that the price and cost parameters are non-increasing.

5.1. Inventory Bound

In this section we develop a worst-case bound for the amount of initial inventory needed to purchase to satisfy demand. A strategic question that arises when managing the inventory to support warranty claims is: *If the failure rates are unknown, what is the maximum number of new devices that will have to be purchased to support the reverse chain?* The answer to this question is useful for two main reasons. First, we can use it to plan seed stock requirements, and guide operational decisions regarding refurbished device management. In addition, we can use it to bound the operational cost of supporting the warranty of a new device. This question is particularly salient for planning the reverse logistics systems for new product technologies, for which there might be little data available on failure distributions.

In order to analyze this issue, we assume that device failures are governed by a failure time distribution that only depends on a device's age, where age is defined as the time a device has been with a customer. We denote the estimate for the probability that a device will fail at age s by f_s .

For planning purposes, we model the demand process as being the expected number of failures. Thus, we treat the failure probabilities as fractions representing the percent of the device population that fails at a particular age. If we denote the number of devices sold at time t as y_t and assume that this is known and given, we can specify the number of warranty claims by

$$d_t = \sum_{s=1}^t y_s \cdot f_{t-s}, \forall t.$$

To model the arrival process we assume that there are no regret returns (since our goal is to obtain a worst-case bound) and that the OEM has a fixed lead time of l at the OEM and a process yield of $0 \leq \alpha \leq 1$, i.e., a fraction $1 - \alpha$ of the devices sent to the OEM cannot be fixed. We model the arrival process as deterministic, given by:

$$a_t = \alpha \cdot d_{t-l}.$$

Now if we suppose that “nature” can pick $\{f_t\}$ adversarially, and if we impose a bound on the total failure probability $\sum_s f_s = \beta \leq 1$, we obtain the following proposition.

PROPOSITION 3. If $\sum_k f_k = \beta \leq 1$, we have

the number of additional units needed to satisfy demand is bounded by $\beta \cdot \max_{t \in [1, T]} \sum_{s=1}^t y_s - \alpha \cdot y_{s-l}$.

Thus, the bound is independent of $\{f_t\}$.

To illustrate the utility of this proposition, assume that sales decay exponentially according to some rate γ , such that $y_{t+1} = \gamma \cdot y_t, \forall t \geq 1$, and $y_t = 0, \forall t \leq 0$. For some $0 \leq \alpha \leq 1$ and lead time l we have:

$$\sum_{s=1}^t y_s - \alpha \cdot y_{s-l} = \begin{cases} \frac{1-\gamma^t}{1-\gamma} \cdot y_1, & \text{if } t < l \\ \frac{1-\gamma^t}{1-\gamma} \cdot y_1 - \alpha \cdot \frac{1-\gamma^{t-l+1}}{1-\gamma} \cdot y_1, & \text{if } t \geq l. \end{cases}$$

The second expression will be decreasing in t if $\alpha > \gamma^l$. Thus, this leads to

$$\max_{t \in [1, T]} \sum_{s=1}^t y_s - \alpha \cdot y_{s-l} = \begin{cases} \frac{1-\gamma^l}{1-\gamma} \cdot y_1, & \text{if } \alpha \geq \gamma^l \\ \frac{(1-\alpha) - \gamma^T (1-\alpha/\gamma^l)}{1-\gamma} \cdot y_1, & \text{if } \alpha < \gamma^l. \end{cases}$$

The worst-case in terms of uncovered demand occurs either at $t = l$ because there are no items arriving from the OEM, or at time T , the end of the horizon, because the arrivals from the OEM continue to lag the failures. By letting $T \rightarrow \infty$, we obtain the bound

$$\max_{t \in [1, T]} \sum_{s=1}^t y_s - \alpha \cdot y_{s-l} \leq \frac{1 - \min(\gamma^l, \alpha)}{1 - \gamma} \cdot y_1$$

From the previous proposition, the maximum number of seed devices that are needed to satisfy warranty claims is $\frac{1 - \min(\alpha, \gamma^l)}{1 - \gamma} \cdot \beta \cdot y_1$. In this example, the total number of device failures is $\frac{1}{1-\gamma} \beta y_1$. Thus, $1 - \min(\alpha, \gamma^l)$ is an upper bound on the fraction of the failed devices for which new devices are needed as replacements. This fraction does not depend on β .

As an example, if we have a 85% yield at the OEM ($\alpha = 0.85$), a 5% decrease in sales per week ($\gamma = 0.95$), a 3 week lead time ($l = 3$), and that 15% of all devices fail ($\beta = 0.15$), the maximum number of seed devices needed is bounded by $0.43 \cdot y_1$, and these new devices will satisfy at most 15% of the failures.

5.2. Certainty-Equivalent Approximation

A *certainty-equivalent* approximation to a stochastic optimization problem is a way to obtain a suboptimal control policy by approximating the uncertainty in the problem by its average value. Ideally, this approximation leads to a control strategy that is easier to obtain (either numerically or analytically) than the optimal policy and has a near-optimal performance. This type of approximation has a long history in control theory, as discussed by Bertsekas (2005), originating from

linear quadratic control problems where in many cases the certainty-equivalent policy leads to the same control strategy as the optimal policy.

In our set-up, the certainty-equivalent approximation to the optimal policy still purchases items myopically, but uses an approximation to calculate the sell-down level in each time period. More specifically, we calculate the sell-down level assuming that the demand and arrival processes are deterministic and equal to their average values. When the volume of devices that fail in every period is large, we expect that this will lead to a satisfactory approximation of the optimal policy.

We assume that device failures are governed by age-dependent failure rates that are independent of a device's sales date. The deterministic approximation takes failure rates to be the fractions of total devices that fail. We then model the expected number of failures at time t to be

$$E[d_t] = \sum_{s=0}^t m_{s,t} \cdot \phi_s(t),$$

where $m_{s,t}$ is the number of customer devices that have survived to age s at time t and $\phi_s(t)$ is the estimate of the failure rate for devices of age s with the information available up to time t . With this in hand, along with a forecast of the number of new devices that are sold in every time-period, we can forecast the average number of devices that fail over the remainder of the horizon.

We use a similar approach to forecast the arrival process. We define $\epsilon_s(t)$ as the estimate at time t of the fraction of devices that return from the OEM s periods after the device failure. We denote by r_t the regret returns that arrive in the reverse logistics facility at time t ; we assume that the facility can refurbish these returns within the period so as to make these devices available for use. Then we model the arrival process by:

$$E[a_t] = E[r_t] + \sum_{s=0}^t E[d_{t-s}] \cdot \epsilon_s(t),$$

where $E[d_t]$ denotes the expectation for d_t in the future, but denotes the realization of d_t in the past. From Proposition 2, the certainty-equivalent sell-down level at time t , denoted by $v^c(t, \tau_t^{\max})$ will be

$$v^c(t, \tau_t^{\max}) = \max_{s \in [t, \tau_t^{\max}]} \sum_{k=t}^s E[d_k] - E[a_k],$$

where $\tau_t^{\max} = \min \left\{ k > t \mid c_k - \sum_{s=t}^{k-1} h_s \leq p_t \right\} - 1$ and $E[a_t]$ and $E[d_t]$ are computed above. From Jensen's Inequality we note that

$$\max_{s \in [t, \tau_t^{\max}]} \sum_{k=t}^s E[d_k] - E[a_k] \leq E \left[\max_{s \in [t, \tau_t^{\max}]} \sum_{k=t}^s d_k - a_k \right].$$

Thus, the sell-down level generated by the certainty-equivalent policy might be too low, which is a consequence of not accounting for the uncertainty of the demand and arrival processes.

6. Numerical Experiments

In this section, we investigate the sensitivity of the WSP's reverse logistics system to changes in various system parameters such as OEM lead time, seed-stock and warranty length. We will also examine the performance of the certainty-equivalent policy for different system parameters. As a benchmark, we compare the performance of this policy with the clairvoyant policy, the policy that knows ex-ante the sample path of the device failures and the arrivals from the OEM. We use the *total profit over the simulation horizon* as a performance metric in all our simulations. Thus, an inventory management policy that purchases a sequence of $\{u_t^+\}$ new devices and sells $\{u_t^-\}$ refurbished devices will generate a profit of

$$\sum_{t=1}^T p_t u_t^- - c_t u_t^+ - h_t x_t.$$

where x_t denotes the resulting inventory in each period. We also analyze the performance of the certainty-equivalent policy using real-world sales and failure data from two devices sold by the WSP.

In order to simulate these different scenarios, a large-scale discrete-time simulator was built using the Julia programming language (Lubin and Dunning (2013)). The simulator creates a virtual CLSC where each device and each customer is an individual object inside the simulation. Thus, each customer and device have a unique id, and failure times, warranty lengths and lead times can be individually realized for each customer/device pair. Furthermore, the simulator allows for each individual device to be in process at the OEM or held in inventory at the WSP's reverse logistics facility. At each time period, refurbished devices can be sold and new devices can be purchased from the OEM.

The simulation experiments were designed to be at the finest level of granularity as possible in order to be as realistic as possible and also to allow for an easy integration with databases and IT systems used to manage reverse operations in the real-world. In addition, by using the Julia language, the simulator was designed to run on a cloud-based environment and on multiple processors.

Furthermore, when using pre-defined distributions, we assume that when a simulated customer purchases a device, the OEM warranty and the customer warranty start simultaneously at the moment of purchase. We also assume that each device fails at most once. These assumptions allow us to ignore mismatches between the customer and OEM warranties, simplifying the simulation and allowing for a larger number of simulated sample paths. Unless specifically stated, one-period in all simulations corresponds to one week.

6.1. A note on state dependency of the optimal policy

As discussed in Section 4, the state of the system is very complex, and includes the number of devices of different ages in the market in each time-period. Hence, the total demand for replacements in each period $\{d_t\}$ and the number of arrivals in inventory $\{a_t\}$ are not simple Markov Processes and depend on the ages of devices currently in the market and on the ages of devices that are being refurbished.

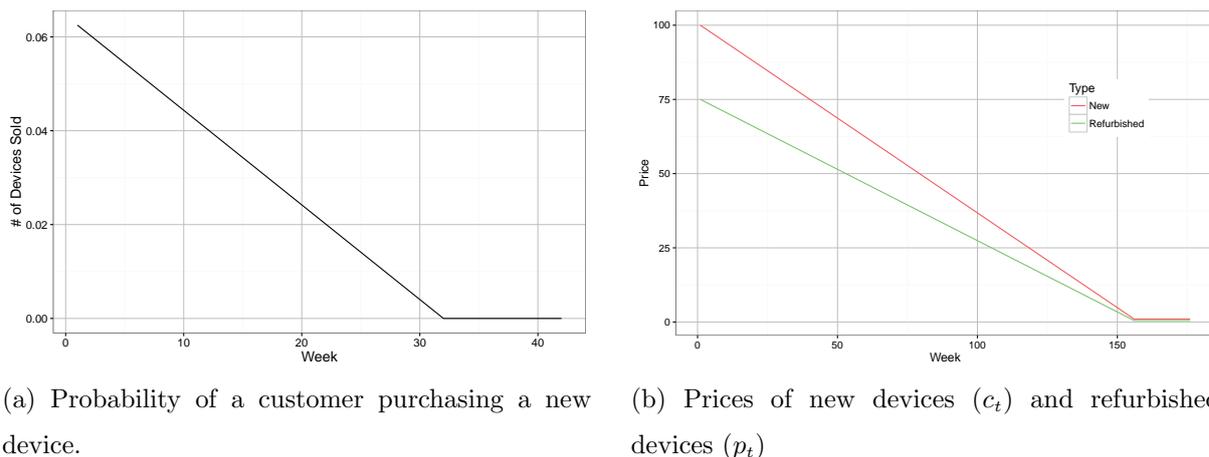
We deal with this by recomputing the sell-down levels in every period following the procedure described in Section 4.1. Hence, if the simulation is in period t , the sell-down levels (v_t^*, \dots, v_T^*) are computed using Algorithm 1 and the probabilities are calculated using a Monte Carlo simulation based on the current state of the system. Then, in period $t + 1$, we recompute the sell-down levels $(v_{t+1}^*, \dots, v_T^*)$ using the same procedure for the updated state of the system. This works well in practice since the number of devices in the system is large and since the failures in one period do not significantly impact the failure distributions in future periods.

6.2. Assumptions on regret-returns, lead time and warranty-length

In order to analyze the behavior of this closed-loop system to changes in different parameters, we assume that customer and OEM warranties have a 52-week length. If a customer has a device that fails and he/she is out of customer warranty, the faulty device will not be replaced. We also assume that the demand for new devices is a random variable with the probability distribution in Figure 3a; the simulation determines the purchase date for each consumer by sampling from this distribution. In this case, devices are sold during the 32 weeks after launch. Since we assume that the customer and device warranties start simultaneously when a customer acquires a new device, Figure 3a also represents the probability distribution for the activation date of both warranties for the customer-devices pair. In addition, the expected demand for new devices will have the same shape as the curve in Figure 3a. We assume a total simulation horizon of up to 150 weeks (allowing us to simulate different lead times), and that the prices of new devices and refurbished devices decrease linearly over time as depicted in Figure 3b. The holding cost per devices is assumed to be \$0.10 per week. Unless otherwise noted, the loss at the OEM is assumed to be 20%, i.e., each failed device has a 20% probability that it cannot be remanufactured.

Finally, we assume that the failure distribution of the device follows an exponential distribution with a mean failure time of 48 months. Thus, the failure rate is constant². Approximately 22% of newly purchased devices will fail under warranty. The expected demand for replacement devices is depicted in Figure 4. In our simulations, 20,000 devices are sold over a product's life-cycle. We calculate sell-down levels using Algorithm 1, and, for each v , we estimate $\delta_t(v)$ using a Monte-Carlo simulation with 100 samples.

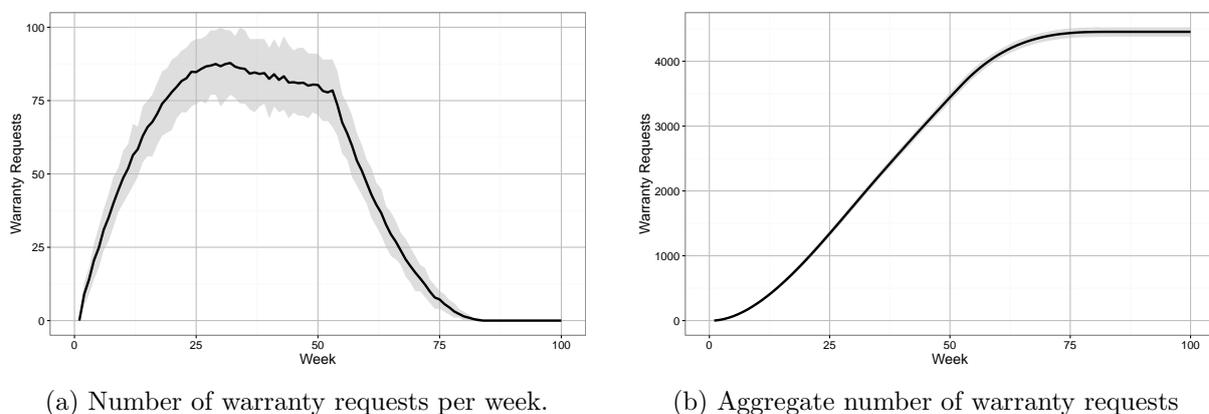
²This is assumption is consistent with previous studies done with our partner WSP, as discussed by Petersen (2013).



(a) Probability of a customer purchasing a new device.

(b) Prices of new devices (c_t) and refurbished devices (p_t)

Figure 3 Simulation parameters.



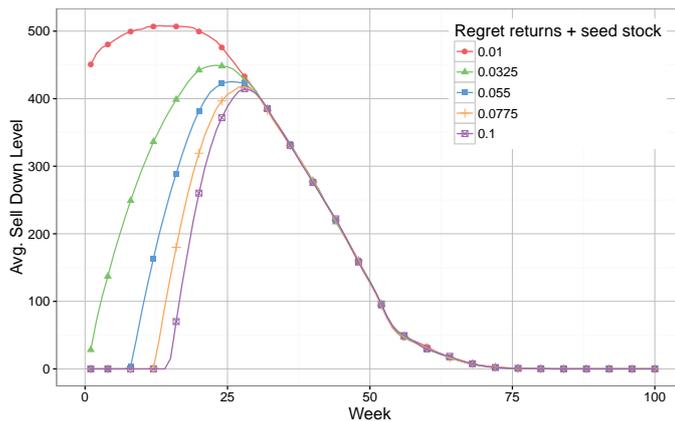
(a) Number of warranty requests per week.

(b) Aggregate number of warranty requests

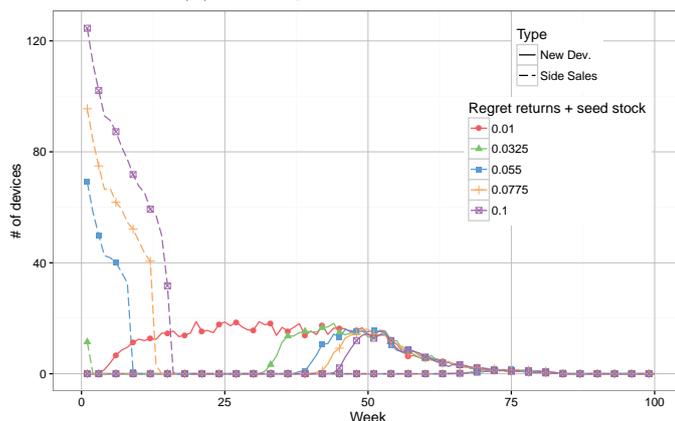
Figure 4 Weekly demand for replacement devices for 1000 sample paths of the system. The failure distribution is exponential with mean 208 weeks and 20,000 devices were sold. The dark line represents the average. The upper and lower limits of the grey area represent, respectively, the 90% and 10% quantiles of demand each week.

Although sales only occur during the first 32 weeks after launch, the demand for replacements peaks in around week 30. Furthermore, demand for replacement devices occurs until 84 weeks after the start of the selling season.

We assume that the repair or remanufacturing lead time is deterministic and equal to 3 weeks. This lead time covers from the time that a replacement device is shipped to a consumer until a remanufactured device is received from the OEM and is available at the reverse logistics facility. This corresponds to time between steps 2 and 5 in Section 3.2 and Figure 1. In the next several sections we report on the performance of the optimal policy, computed using algorithm 1, and its sensitivity to various parameter settings.



(a) Average sell-down levels.



(b) Avg. amount of new devices purchased and of side-sales

Figure 5 The figures are based on 100 simulations of the system. **Figure (a)** depicts average sell-down levels for different levels of regret returns and seed-stock. **Figure (b)** depicts the distribution of new devices used. In (b), the middle line in the box is the median, the upper and lower ends of the box are, respectively, the 25% and 75% quantiles, and the lines extend to 1.5 times the Interquartile (IQR) range. **Figure (c)** depicts the average paths of devices sold through side-sales (dashed lines) and new devices purchased into inventory (solid line).

6.2.1. Regret-returns and seed-stock. The WSP's reverse logistics facility has two sources of devices other than refurbished devices from the OEM: (i) regret-returns, and (ii) seed-stock from the OEM. Both regret returns and seed stock are functions of new device sales, as discussed previously.

For simplicity, in our simulations, we bundle together regret-returns and seed-stock as a fraction of sales. For example, if we assume a regret-return rate of 5% and seed-stock as 1% sales, then if 100 devices are sold in a week, this will correspond to 6 additional devices being provided to the reverse-logistics inventory.

We report the average sell-down levels, $\{v_t^*\}$, and new devices purchased, $\{u_t^+\}$, for different regret-returns and seed-stock levels in Figure 5. As expected, the average sell-down level decreases as the amount of regret-returns and seed-stock increase. However, there is a more significant decrease of the sell-down levels during the first 12 weeks after launch. This occurs due to holding costs and the non-increasing prices in this system. At the beginning of the device's lifecycle, it is more profitable to sell seed-stock and regret returns through side-sales channels, and potentially buy new devices in the future, than keeping regret-returns in inventory for future use. This can be further explored by comparing the average number of side-sales to the number of new devices purchased to fulfill demand, as depicted in Figure 5b. If the regret-returns increase beyond a certain level, the additional excess inventory will have a diminishing impact on the amount of new devices that will need to be purchased to sustain the closed-loop system. Furthermore, note that the sell-down levels in Figure 5 converge around week 30 since no new devices are being sold and, therefore, no more regret-returns and seed-stock are arriving into inventory.

6.2.2. Lead Time The WSP can influence the remanufacturing lead time, i.e., the time between shipping a replacement device to a customer and receiving a corresponding remanufactured device (steps 2 to 5 in Section 3.2) by either (i) reducing the time a customer has to return a faulty device, (ii) reducing the sorting and triage period of a device within the reverse logistics facility or (iii) negotiating a reduction in lead time with the OEM. The impact of different lead times on system performance is depicted in Figure 6. For this set of simulations, we assume that seed-stock and regret-returns correspond to 5% of sales.

An increase in lead time impacts inventory management throughout the life-cycle of the device, as depicted in Figure 6c. At the beginning of the life-cycle (during the first 20 weeks), an increase in lead time reduces the amount of regret-returns and initial seed-stock that is allocated to side-sales, as can be noted in week 1 to week 10 in Figure 6c. In the middle of the life-cycle (week 20 to week 60), an increase in lead time increases the number of new devices that are needed to sustain the system during the weeks of peak demand. Finally, towards the end of the life-cycle (week 60 until end of arrivals) an increase in lead time increases the amount of devices sold through side-sales. This is because faulty devices that are sent to the OEM during the peak of returns might arrive after the peak of the demand for replacements has past, when the demand for replacement devices is lower. In summary, for short lead times, side-sales mostly occur in the beginning of the device's life-cycle, while for long lead times, side-sales occur mostly towards the end of the device's life-cycle. These effects combined lead to the U-shaped curves observed in Figures 6a and 6b.

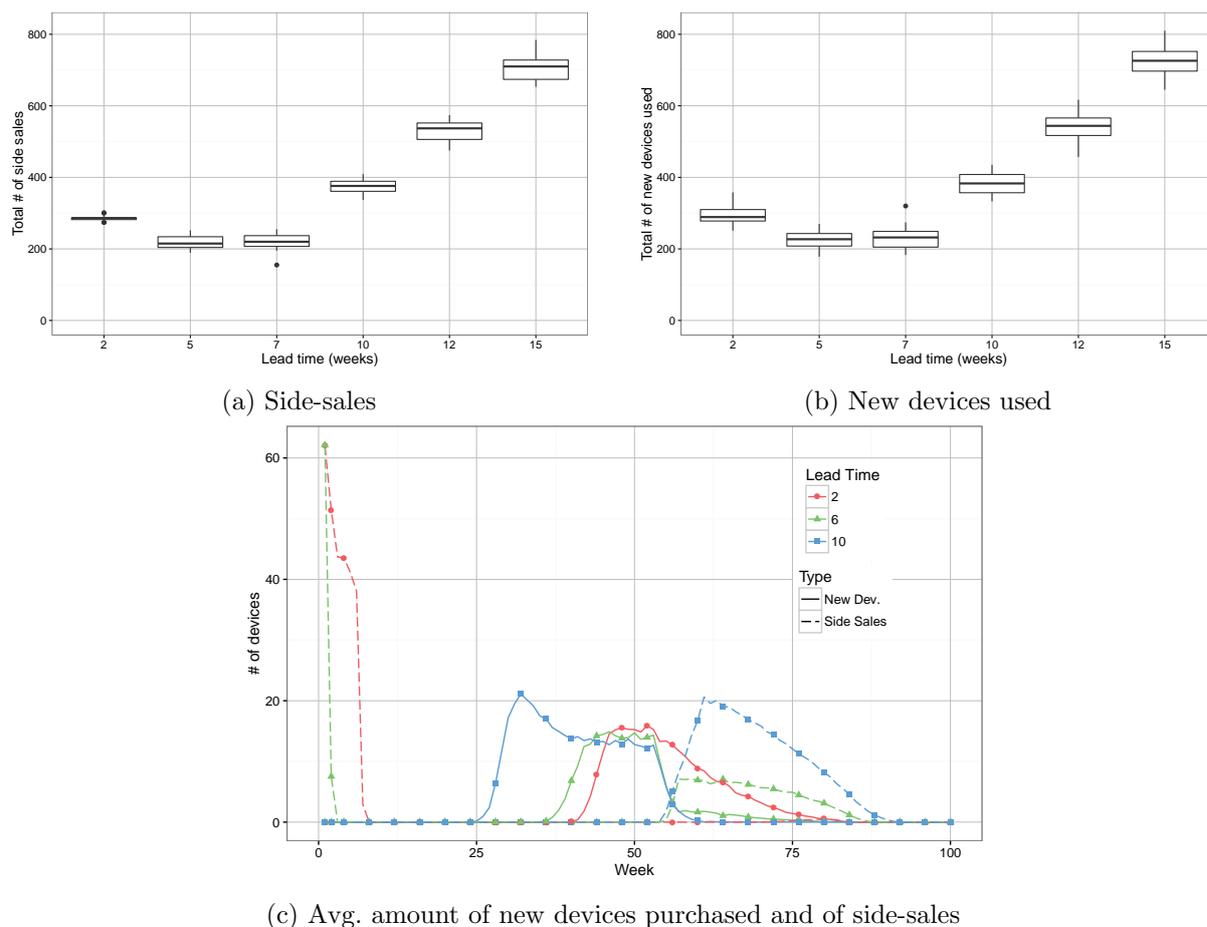


Figure 6 The figures are based on 100 simulations of the system. **Figure (a)** depicts the distribution of side-sales. **Figure (b)** depicts the distribution of new devices used. **Figure (c)** depicts the average paths of devices sold through side-sales (dashed lines) and new devices purchased into inventory (solid line). The seed-stock and regret-returns are 5% of total sales and each faulty device has 80% probability of being remanufactured.

6.2.3. Warranty Length The length of the customer warranty is the design parameter of this reverse logistics system that can influence the likelihood of a consumer purchasing a product or not. A longer customer warranty also changes the time that this device will have to be available at the reverse logistics system. The behavior of the system for different warranty lengths is depicted in Figure 7. We assume that the OEM warranty has the same length as the customer warranty.

An increase in warranty length increases the average sell-down level of the system and sell-down levels will also be non-zero for longer. This leads to an increase in the number of new devices needed to fulfill the demand for replacements. The amount of new devices required might be even larger if devices could fail multiple times, which we do not capture in our simulations. Identifying customer sensitivity to changes in warranty length is critical, since it has implications long after the sales of a device have ended (which occurs after 32 weeks). For warranty lengths larger than 46 weeks,

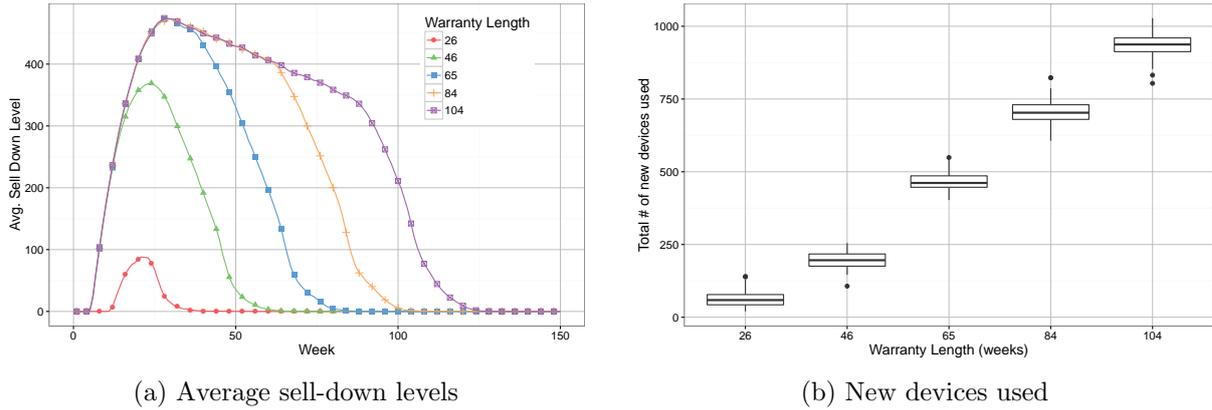


Figure 7 The figures are based on 100 simulations of the system. **Figure (a)** depicts the distribution of side-sales. **Figure (b)** depicts the distribution of new devices used.

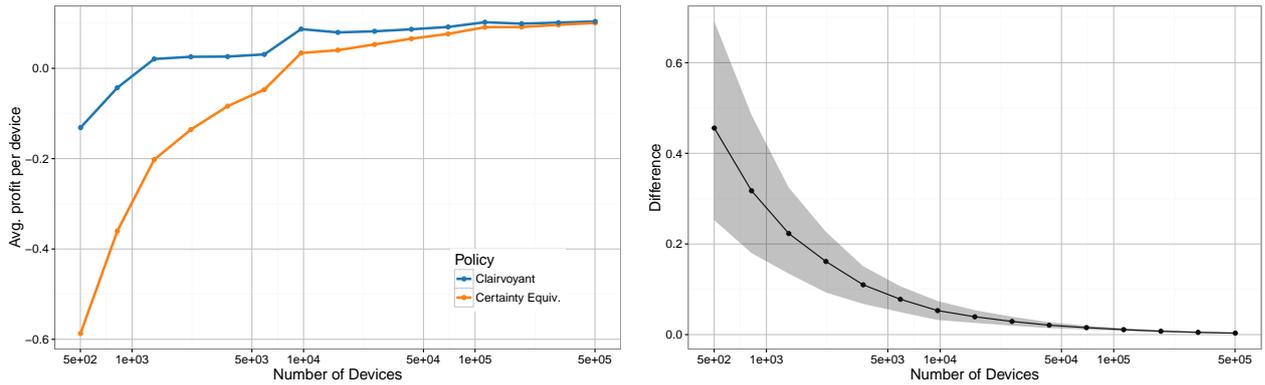
the sell-down levels are the same for the first 20 weeks the device’s life-cycle. This occurs because price depreciation and holding costs dominate at the beginning of the device’s lifecycle. Also, for a warranty length of 26 weeks, almost all demand for warranty replacements can be satisfied by seed-stock, regret-returns, and refurbished arrivals from the OEM. As the warranty increases, the average amount of weekly failures increases, and more refurbished arrivals need to be allocated into inventory.

6.3. The certainty-equivalent heuristic

In these experiments we simulate the performance of the certainty-equivalent heuristic. For all simulations, we assume that the customer and OEM warranties have a 12-month length and that one period corresponds to one day. Computing the optimal sell-down level at this level of granularity requires a large amount of computing time and memory while the certainty-equivalent sell-down level is simple to compute. Since computing the optimal policy is computationally taxing, we use the clairvoyant policy as a benchmark. The clairvoyant policy knows, for each simulation iteration, the realization of the demand and arrival processes and can then determine the optimal control policy by solving a large-scale linear program.

Once again we assume that the purchasing date of a device follows the probability distribution depicted in Figure 3a; the simulation determines the purchase date for each consumer by sampling from this distribution. We also assume a total simulation horizon of 2.5 years, and that the prices of new devices and refurbished devices decrease linearly over time as depicted in Figure 3b and that holding costs are \$0.10 per week. As before, the failure distribution of devices is exponential with a mean failure time of 48 months.

In Figure 8 we depict the performance of the certainty-equivalent policy for different numbers of devices sold. We assume that seed-stock corresponds to 5% of sales, that the total lead time



(a) Average profit per-device sold as a function of the number of devices sold. (b) Difference between the certainty equivalent and the clairvoyant approximations.

Figure 8 Comparison between the certainty equivalent and the clairvoyant policies as a function of the number of devices sold. The top and bottom of the shaded area in (b) represent the 90% and 10% quantiles of the difference distribution. The pictures are based on 300 repetitions of the simulation.

for the customer to return a failed device and the OEM to repair it is 28 days, and that the loss at the OEM is 20%. We normalize the profit by the number of devices so that we can compare the profit (cost) per device when the number of devices in the system changes. The averages are based on 300 repetitions of the simulation. When the number of devices is small (less than 3000), we observe that the performance of the certainty-equivalent policy seems much worse than the clairvoyant. This is not surprising, as we expect that when there is a small number of devices, there will be a relatively large coefficient of variation for the maximum total net demand which determines the sell-down level. As the number of devices in the system increases, this coefficient of variation decreases and the certainty equivalent approximations improve. The behavior of the coefficient of variation of the number of warranty requests in a week as a function of the number of devices sold is depicted in Figure 9. Weeks 5 and 65 are, respectively, at the beginning and the end of the device's life-cycle, where the number of weekly requests is low and the coefficient of variation is high, as was illustrated in Figure 4a.

In practice, the average number of devices of a given type sold by the WSP is on the order of hundreds of thousands to millions of devices. In this range the certainty equivalent heuristic is very close to the clairvoyant policy and should work well in practice. Simulations using real-world data from the WSP are presented in the next section.

6.4. Simulation using sales and failure data from the WSP

In this section we examine the performance of the certainty-equivalent heuristic compared to the clairvoyant policy when applied to data for two of the WSP's best-selling models, which we call

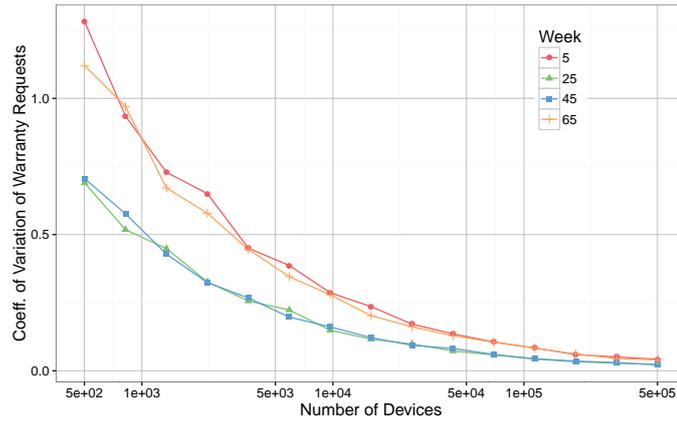


Figure 9 Coefficient of variation of the number of warranty requests in a week as a function of the total number of devices sold for weeks 5, 25, 45, and 65. The pictures are based on 300 repetitions of the simulation.

model A and model B. For every device of model A and model B that was sold we have the sales date as well as the failure date (if the device failed at all).

The number of devices sold and the failure rates for these two models are depicted in Figure 10. For confidentiality reasons, the data is aggregated per month and the peak sales and failure rates are normalized to one. In practice, between 10% and 25% of total devices sold from each model failed under warranty.

We assume that prices behave as in Figure 3, that the holding costs are \$0.10 per week, that there is a 70% probability that a faulty devices can be remanufactured and that regret returns and seed-stock correspond to 5% of sales. The OEM lead time was 4 weeks.

We also assume that the certainty-equivalent heuristic has a perfect forecast of the sales of new devices, knows the exact failure rate for each model, and also knows the probability that each faulty device will be successfully remanufactured by the OEM. The clairvoyant policy knows the exact failure date of every device and has a perfect forecast of arrivals from the OEM. We assume that both policies know the OEM lead time and that the total sales for each device are 100,000 units. The simulation has a time period of one day.

The average sell-down levels of the certainty-equivalent heuristic are depicted in Figure 11. The averages are taken over 100 sample runs of the system. We find that for both device models the certainty-equivalent heuristic is within 5% of the clairvoyant policy. More specifically, we find that the ratio between the profit of the certainty-equivalent heuristic and the profit of the clairvoyant policy has an average of 0.97 with a standard deviation of $0.6 \cdot 10^{-2}$ for model A, and an average of 0.96 with a standard deviation of $1.5 \cdot 10^{-2}$ for model B.

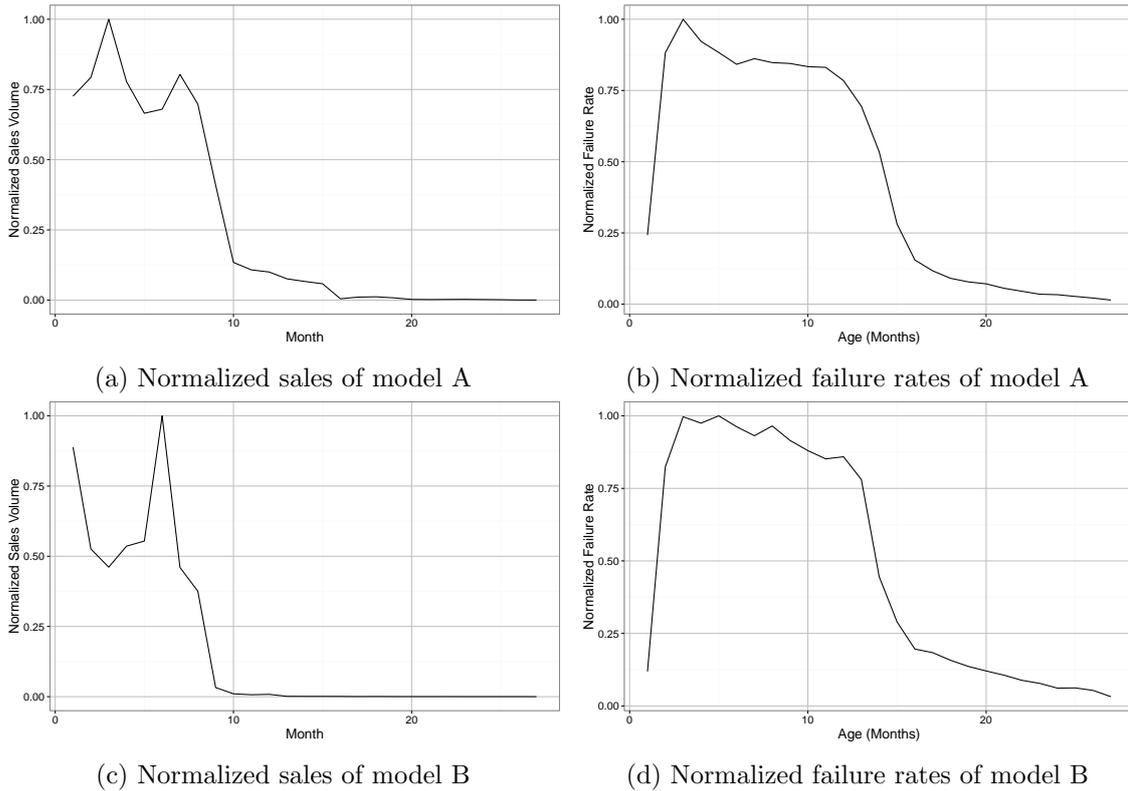


Figure 10 Sales and failure rates of two models sold by the WSP. Data was aggregated per month and the peak failure rates and sales were normalized to one.

As discussed in the previous section, the effectiveness of the heuristic is due to the large number of devices in the system. Furthermore, this highlights the practical value of this policy, since the sell-down levels in the certainty-equivalent policy are easy to compute, and can be implemented using spreadsheet software.

7. Conclusion

In this paper, we formulate and analyze a new inventory management problem found in the reverse logistics operations of an electronics retailer. We formulate the problem as a stochastic optimization problem that takes into account the closed-loop nature of the inventory system, as well as the short life-cycle and rapid value depreciation of the electronic devices.

We prove the structure of the optimal inventory in this set-up. The optimal sourcing strategy will be myopic in that we only buy enough items to satisfy the unmet demand for replacement devices in the current period. Conversely, the optimal selling quantity in each period depends on the future distribution of demand and arrivals, as well as on the behavior of prices and costs. We propose an algorithm that can be implemented using a Monte-Carlo simulation to calculate the optimal sell-down quantity. However, depending on the scale of the system, calculating optimal sell-down

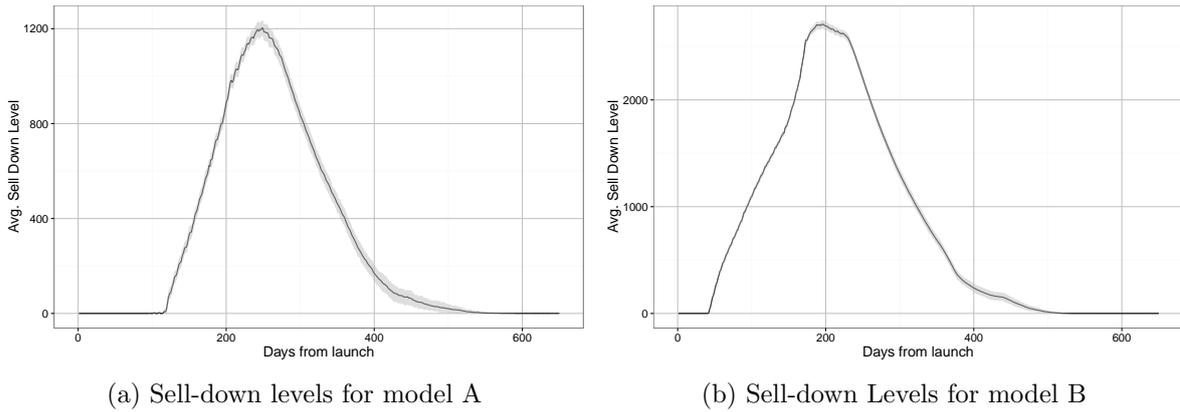


Figure 11 Figures (a) and (b) depict sell-down levels for each model. The black line corresponds to the average sell-down level and the top and bottom of the grey band correspond to the 90% and 10% range of the sell-down level over 100 simulations.

levels can be computationally intensive. Because of this, we analyze a certainty-equivalent version of this system, where demand and arrivals are assumed to be deterministic.

When the demand and arrival processes are deterministic, we find the optimal policy and also develop a worst-case analysis. The optimal selling quantity in each period is easy to compute and depends on the maximum total net demand in the interval between the current period and the time when the cost of sourcing a new device falls below the current price of a refurbished device in a side-sales channel. Thus, the maximum total net demand acts as a sell-down level. If inventory is above this level, items will be sold until the number of items in inventory is equal to this level. Conversely, if inventory is below this level, no items are sold.

Through numerical experiments, we analyze the behavior of this system when different system parameters change. Due to the non-stationarity of system parameters, changes in regret returns, lead times, and warranty lengths impact the inventory management policy throughout a device's life-cycle. Furthermore, an increase in regret-returns or seed-stock has a diminishing impact in overall performance of the system, due to a timing disparity: the peak for regret returns occurs well before the peak for replacement demand. We also observe that when the number of devices is large the certainty-equivalent policy achieves a near-optimal performance and is sufficient for practical applications. These findings were confirmed using real-world data from two devices sold by the WSP. We find that under practical settings the certainty-equivalent policy captures over 90% of the clairvoyant profit.

There are a few open problems that were not addressed in this paper. First, note that many issues in this supply chain are caused by the fact that the customer warranty and the OEM warranty have different specifications. Namely, the customer warranty requires a fast replacement, while the OEM

warranty allows for a comparatively large lead time. Studying how to redesign these contracts by taking into account the incentives and preferences of the customer, the WSP and the OEM might lead to novel insights about the management of warranty systems.

A second open problem is the analysis of inventory management policies that take into account how much time customers and devices in inventory have left in their respective warranties. Looking at closed-loop inventory management and matching problem together might lead to new policies for managing this system that could yield benefits.

A third issue is to extend the modeling and analysis to relax our assumption that the price at which refurbished devices can be sold in a side channel is non-increasing. If this were not true, then an optimal policy might hold excess inventory in the hopes of selling it in the future at a higher price. This would seem to require a different solution approach than the incremental analysis followed here, and we leave this for future research.

References

- Allen, Stephen G., Donato A. D'Esopo. 1968. An Ordering Policy for Repairable Stock Items. *Operations Research* **16**(3) 669–674.
- Bayiz, Murat, Christopher S. Tang. 2004. An Integrated Planning System for Managing the Refurbishment of Thermoluminescent Badges. *Interfaces* **34**(5) 383–393.
- Beltran, Jose Luis, Dmitry Krass. 2002. Dynamic lot sizing with returning items and disposals. *IIE Transactions* **34**(5) 437–448.
- Bertsekas, Dimitri P. 2005. *Dynamic Programming & Optimal Control, Vol. I*. 3rd ed. Athena Scientific.
- Boyd, Stephen, Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge university press.
- Calmon, Andre du Pin. 2015. Reverse Logistics for Consumer Electronics: Forecasting Failures, Managing Inventory, and Matching Warranties. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Chen, Xin, David Simchi-Levi. 2009. A new approach for the stochastic cash balance problem with fixed costs. *Probability in the Engineering and Informational Sciences* **23**(04) 545–562.
- Constantinides, George M., Scott F. Richard. 1978. Existence of Optimal Simple Policies for Discounted-Cost Inventory and Cash Management in Continuous Time. *Operations Research* **26**(4) 620–636.
- de Brito, Marisa P., Erwin A. van der Laan. 2009. Inventory control with product returns: The impact of imperfect information. *European Journal of Operational Research* **194**(1) 85–101.
- Feinberg, Eugene A., Mark E. Lewis. 2005. Optimality of four-threshold policies in inventory systems with customer returns and borrowing/storage options. *Probability in the Engineering and Informational Sciences* **19**(01) 45–71.

- Ferguson, Mark E., Moritz Fleischmann, Gilvan C. Souza. 2011. A Profit-Maximizing Approach to Disposition Decisions for Product Returns. *Decision Sciences* **42**(3) 773–798.
- Fleischmann, Moritz, Jacqueline M. Bloemhof-Ruwaard, Rommert Dekker, Erwin Van der Laan, Jo AEE Van Nunen, Luk N. Van Wassenhove. 1997. Quantitative models for reverse logistics: A review. *European journal of operational research* **103**(1) 1–17.
- Fleischmann, Moritz, Jo A. E. E. van Nunen, Ben Gräve. 2003. Integrating Closed-Loop Supply Chains and Spare-Parts Management at IBM. *Interfaces* **33**(6) 44–56.
- Guide, V. Daniel R, Luk N Van Wassenhove. 2009. OR FORUM: The Evolution of Closed-Loop Supply Chain Research. *Operations Research* **57**(1) 10–18.
- Heyman, Daniel P. 1977. Optimal disposal policies for a single-item inventory system with returns. *Naval Research Logistics Quarterly* **24**(3) 385–405.
- Huang, Wei, Vidyadhar Kulkarni, Jayashankar M Swaminathan. 2008. Managing the Inventory of an Item with a Replacement Warranty. *Management Science* **54**(8) 1441–1452.
- Khawam, John, Warren H Hausman, Dinah W Cheng. 2007. Warranty Inventory Optimization for Hitachi Global Storage Technologies, Inc. *Interfaces* **37**(5) 455–471.
- Kiesmüller, Gudrun P., Erwin A. van der Laan. 2001. An inventory model with dependent product demands and returns. *International Journal of Production Economics* **72**(1) 73–87.
- Levi, Retsef, Robin O. Roundy, David B. Shmoys, Van Anh Truong. 2008. Approximation Algorithms for Capacitated Stochastic Inventory Control Models. *Operations Research* **56**(5) 1184–1199.
- Lubin, Miles, Iain Dunning. 2013. Computing in Operations Research using Julia. *arXiv:1312.1431 [cs, math]* . ArXiv: 1312.1431.
- Petersen, Brian J. 2013. Reverse supply chain forecasting and decision modeling for improved inventory management. Thesis, Massachusetts Institute of Technology.
- Pinçe, Çerağ, Ülkü Gürler, Emre Berk. 2008. A continuous review replenishment disposal policy for an inventory system with autonomous supply and fixed disposal costs. *European Journal of Operational Research* **190**(2) 421–442.
- Pinçe, Çerağ, Mark Ferguson, L. Beril Toktay. 2016. Extracting Maximum Value from Consumer Returns: Allocating Between Remarketing and Refurbishing for Warranty Claims. *Manufacturing & Service Operations Management* .
- Savaskan, R. Canan, Shantanu Bhattacharya, Luk N. van Wassenhove. 2004. Closed-Loop Supply Chain Models with Product Remanufacturing. *Management Science* **50**(2) 239–252.
- Simpson, Vincent P. 1978. Optimum Solution Structure for a Repairable Inventory Problem. *Operations Research* **26**(2) 270–281.
- Souza, Gilvan C. 2013. Closed-Loop Supply Chains: A Critical Review, and Future Research. *Decision Sciences* **44**(1) 7–38.

- Tao, Zhijie, Sean X. Zhou. 2014. Approximation Balancing Policies for Inventory Systems with Remanufacturing. *Mathematics of Operations Research* **39**(4) 1179–1197.
- Toktay, L. Beril, Lawrence M. Wein, Stefanos A. Zenios. 2000. Inventory Management of Remanufacturable Products. *Management Science* **46**(11) 1412–1426.
- Zhou, Sean X., Zhijie Tao, Xiuli Chao. 2011. Optimal control of inventory systems with multiple types of remanufacturable products. *Manufacturing & Service Operations Management* **13**(1) 20–34.

Appendix A: Proof of Proposition 1

PROPOSITION 1 For some $v_t^* \geq 0$, the optimal policy for the stochastic inventory model in (1) is

$$\begin{aligned}\hat{u}_t^+(\bar{x}_t) &= \max(-\bar{x}_t, 0), \\ \hat{u}_t^-(\bar{x}_t) &= \max(\bar{x}_t - v_t^*, 0),\end{aligned}$$

where v_t^* only depends on the distribution of demand and arrivals, and on the cost and price parameters.

Proof. The proof of the proposition will be done by backwards induction on t . Since no backordering is allowed, we have

$$J_T(\bar{x}) = \max_{\substack{0 \leq \bar{x} + u^+ - u^- \\ u^+, u^- \geq 0}} -c_T u^+ + p_T u^- - h_T(\bar{x} + u^+ - u^-), \quad (5)$$

and for $t \in [1, T-1]$,

$$J_t(\bar{x}) = \max_{\substack{0 \leq \bar{x} + u^+ - u^- \\ u^+, u^- \geq 0}} -c_t u^+ + p_t u^- - h_t(\bar{x} + u^+ - u^-) + E_{\{d_{t+1}, a_{t+1}\}} [J_{t+1}(\bar{x} + u^+ - u^- + a_{t+1} - d_{t+1})]. \quad (6)$$

Note that the expectation in the second equation is taken over the demand and arrivals in period $t+1$.

The induction hypotheses are

1. The optimal ordering policy exists and is finite, and, for some $v_t^* \geq 0$, is

$$\hat{u}_t^+(\bar{x}) = \max(-\bar{x}, 0) \quad (7)$$

$$\hat{u}_t^-(\bar{x}) = \max(\bar{x} - v_t^*, 0) \quad (8)$$

2. $J_t(\bar{x})$ is non-decreasing and concave;

For $t = T$, we will buy \bar{x} items if $\bar{x} < 0$ or sell \bar{x} items if $\bar{x} > 0$. The optimal cost will be $\min(c_T \bar{x}_T, p_T \bar{x}_T)$, which is concave, and all the induction hypotheses hold.

Now, assume that the hypotheses hold for $t+1$. For ease of exposition, let u_t be defined as the net number of devices purchased and sold, i.e.,

$$u_t = u_t^- - u_t^+,$$

and let $\hat{u}_t(\bar{x}) = \hat{u}_t^-(\bar{x}) - \hat{u}_t^+(\bar{x})$. Then,

$$J_t(\bar{x}) = \max_{0 \leq \bar{x} - u_t} \{ \min(c_t u_t, p_t u_t) - h_t(\bar{x} - u_t) + E[J_{t+1}(\bar{x} - u_t + a_{t+1} - d_{t+1})] \}.$$

For the first induction hypothesis, we will analyze separately the case when $\bar{x} \leq 0$ and when $\bar{x} > 0$. Furthermore, we will use the transformation $y_t = \bar{x} - u_t$, such that the expression above becomes

$$J_t(\bar{x}) = \max_{0 \leq y_t} \{ \min(c_t(\bar{x} - y_t), p_t(\bar{x} - y_t)) - h_t y_t + E[J_{t+1}(y_t + a_{t+1} - d_{t+1})] \}.$$

If $\bar{x} \leq 0$, we have that

$$J_t(\bar{x}) = \max_{0 \leq y_t} c_t(\bar{x} - y_t) - h_t y_t + E[J_{t+1}(y_t + a_{t+1} - d_{t+1})].$$

Since, $J'_{t+1}(\bar{x}) \leq c_{t+1} + h_{t+1} \leq c_t + h_t, \forall \bar{x}$, the expression within the maximization above is non-increasing in y_t , and an optimal solution is $\hat{u}(\bar{x}, t) = \bar{x}$ will be an optimal solution. If $c_{t+1} < c_t$ and $h_{t+1} < h_t$ this solution will be unique.

We will now consider the case when $\bar{x} > 0$. Since $c_t(\bar{x} - y_t) - h_t y_t + E[J_{t+1}(y_t + a_{t+1} - d_{t+1})]$ is non-increasing in y_t , there will be an optimal solution where $y_t \leq \bar{x}$. Hence, the optimal policy and cost can be found by solving

$$\max_{0 \leq y_t \leq \bar{x}} \{p_t(\bar{x} - y_t) - h_t y_t + E[J_{t+1}(y_t + a_{t+1} - d_{t+1})]\}, \quad (9)$$

which has a finite optimum since y_t has to be within a finite interval. Let y^u denote the (potentially infinite) unconstrained minimum of this optimization problem and let $y^*(\bar{x})$ denote the constrained minimum. Then, we have that, since $\bar{x} > 0$,

$$y^*(\bar{x}) = \begin{cases} \bar{x} & \text{if } y^u \geq \bar{x} \\ y^u & \text{if } 0 \leq y^u \leq \bar{x} \\ 0 & \text{if } y^u < 0 \end{cases}$$

By defining $v_t^*(t) = \max(0, y^u)$, we can then write the optimal policy as

$$\hat{u}(\bar{x}, t) = \begin{cases} \bar{x} & \text{if } \bar{x} \leq 0 \\ \max(\bar{x} - v_t^*, 0) & \text{if } \bar{x} > 0 \end{cases}$$

which satisfies the first induction hypothesis.

For the second induction hypothesis, note that

$$\min(c_t(\bar{x} - y_t), p_t(\bar{x} - y_t)) - h_t y_t + E[J_{t+1}(y_t + a_{t+1} - d_{t+1})],$$

is concave ³ in (\bar{x}, y_t) . Furthermore, $y_t \geq 0$ is a convex set, maximization over $y_t \geq 0$ will preserve concavity.

Finally, using Equation 9, the optimal cost can be written as

$$J_t(\bar{x}) = \begin{cases} c_t \bar{x} + E[J_{t+1}(a_{t+1} - d_{t+1})] & \text{if } \bar{x} \leq 0 \\ p_t \bar{x} + \max_{0 \leq y_t \leq \bar{x}} \{- (p_t + h_t) y_t + E[J_{t+1}(y_t + a_{t+1} - d_{t+1})]\} & \text{if } \bar{x} > 0 \end{cases}$$

which is continuous and non-decreasing in \bar{x} , completing the proof. \square

Appendix B: Proof of Proposition 2

PROPOSITION 2 For any $t_1 \leq t_2$, we define “net demand” $n(t_1, t_2)$ as

$$n(t_1, t_2) = \begin{cases} \max_{s \in [t_1+1, t_2]} \sum_{k=t_1+1}^s d_k - a_k, & \text{if } t_2 > t_1 \\ 0, & \text{if } t_2 = t_1. \end{cases}$$

Then, we have

$$v_t^* = \max(n(t, \tau_t^{\max}), 0), \quad \forall t \in [1, T].$$

Proof. The proof will be done by backwards induction on t . The induction hypothesis is that

$$v_t^* = \max(n(t, \tau_t^{\max}), 0), \quad \forall t \in [1, T].$$

For $t = T$ we have that $v_T^* = 0$ since $\tau_T^{\max} = T$.

For $t < T$, assume that we are given the optimal $\{v_{t+1}^*, \dots, v_T^*\}$. Furthermore, let the candidate sell-down level v_t be $v_t = \max(n(t, \tau_t^{\max}), 0)$. If $n(t, \tau_t^{\max}) \leq 0$, then there will never be a need to purchase devices in the interval $[t, \tau_t^{\max}]$ and there is no point in carrying inventory to period $t + 1$. Hence, in this case, $v_t^* = v = 0$.

³The term $\min(c_t(\bar{x} - y_t), p_t(\bar{x} - y_t))$ is concave since $\min(c_t x, p_t x)$ is concave and the composition of a concave function with an affine function preserves concavity.

Note that if, in some period $s \in [t+1, \tau_t^{\max}]$ devices are sold, then it must be that no purchases occur in the interval $[s, \tau_t^{\max}]$ since $\tau_t^{\max} \leq \tau_s^{\max}$ and this would contradict the optimality of v_s^* (we could increase v_s^* by one and obtain a better outcome).

When $n(t, \tau_t^{\max}) > 0$, we will prove that: (i) $\delta(v_t - 1) > 0$; and (ii) that $\delta(v_t) \leq 0$. For such, let τ_t^* be defined as

$$\tau_t^* = \min \left\{ s > t \mid \sum_{k=t+1}^s d_k - a_k = n(t, \tau_t^{\max}) \right\}.$$

For case (i), if the inventory level at the end of period t is $v_t - 1$, then the inventory level at time τ_t^* will be upper bounded by $v_t - 1 + n(t, \tau_t^{\max})$ since some devices might be sold. Hence, since

$$v_t - 1 + n(t, \tau_t^{\max}) = n(t, \tau_t^{\max}) - 1 + n(t, \tau_t^{\max}) < 0,$$

we have that at least one unit will have to be purchased and, since $\tau_t^* \leq \tau_t^{\max}$, it must be that

$$c_{\tau_t^*} - \sum_{s=t}^{\tau_t^*-1} h_s > p_t,$$

and $\delta(v_t - 1) > 0$.

For case (ii), if the inventory level at the end of period t is v_t , then no items will be sold in the interval $[t+1, \tau_t^*]$, since $\tau_t^{\max} \leq \tau_s^{\max}, \forall s \geq t$ and, for all $s \in [t+1, \tau_t^*]$

$$\begin{aligned} v_t - \sum_{k=t+1}^s (d_k - a_k) &= \sum_{k=t+1}^{\tau_t^*} (d_k - a_k) - \sum_{k=t+1}^s (d_k - a_k) \\ &= \sum_{k=s+1}^{\tau_t^*} (d_k - a_k) \\ &\leq v_s^*. \end{aligned}$$

Thus, if the inventory at the end of period t is v_t , then the inventory at the end of period τ_t^* will be zero. If $\tau_t^* = \tau_t^{\max}$, then $\delta(v) = 0$ and we are done. If $\tau_t^* < \tau_t^{\max}$, it must be that no items are purchased in the interval $[\tau_t^*, \tau_t^{\max}]$. This is because $\sum_{k=\tau_t^*+1}^s a_k - d_k \geq 0, \forall s \in [\tau_t^* + 1, \tau_t^{\max}]$ since, from the definition of τ_t^* , for any $s \in [\tau_t^* + 1, \tau_t^{\max}]$,

$$\sum_{k=t+1}^{\tau_t^*} (d_k - a_k) - \sum_{k=t+1}^s (d_k - a_k) \geq 0 \implies \sum_{k=\tau_t^*+1}^s a_k - d_k \geq 0.$$

Items might be sold in the interval $[\tau_t^* + 1, \tau_t^{\max}]$, however, it cannot be that if an item is sold at time $s \in [\tau_t^* + 1, \tau_t^{\max}]$ would lead to an item being purchased in interval $[s+1, \tau_t^{\max}]$, since this would contradict the optimality of v_s^* (we could sell one less item at time s and increase overall profit by not having to purchase another item). Hence, since items will only be sold in the interval $[t+1, \tau_t^{\max}]$, and since $\{p_t\}$ is non-increasing, it must be that $\delta(v_t) \leq 0$ since increasing the inventory by one unit will lead to, at best, more devices being sold in the interval $[t+1, \tau_t^{\max}]$. Thus,

$$v_t^* = v_t = \max(n(t, \tau_t^{\max}), 0), \quad \forall t \in [1, T].$$

and the proof is complete. \square

Appendix C: Proof of Proposition 3

PROPOSITION 3 *Let f_t be the fraction of devices that fail at age t and y_t the number of devices sold at time t . Furthermore, let $d_t = \sum_{s=1}^t f_s y_{t-s}$ and $a_t = \alpha d_{t-l}$, for some lead time l and efficiency α . Then, if $\sum_k f_k = \beta \leq 1$, we have*

$$\text{the number of additional units needed to satisfy demand is bounded by } \beta \cdot \max_{t \in [1, T]} \sum_{s=1}^t y_s - \alpha \cdot y_{s-l}.$$

Thus, the bound is independent of $\{f_t\}$.

Proof. For failure probabilities $\{f_t\}$, and given a lead time of l at the OEM and an efficiency of α , the total number of new devices that will have to be purchased to satisfy demand is

$$\begin{aligned} n(1, T) &= \max_{s \in [1, T]} \sum_{k=1}^s d_k - a_k \\ &= \max_{s \in [1, T]} \sum_{k=1}^s d_k - \alpha \cdot d_{k-l}. \end{aligned}$$

The demand process can be written as

$$d_t = \sum_{s=1}^t f_s \cdot y_{t-s}.$$

Thus, we have that

$$\begin{aligned} n(1, T) &= \max_{t \in [1, T]} \sum_{s=1}^t \sum_{k=1}^s f_k y_{s-k} - \alpha \cdot \sum_{k=1}^{s-l} f_k y_{s-l-k}, \\ &= \max_{t \in [1, T]} \sum_{s=1}^t f_s \sum_{k=1}^{t-s} y_k - \alpha \cdot \sum_{s=1}^{t-l} f_s \sum_{k=1}^{t-l-s} y_k, \\ &= \max_{t \in [1, T]} \sum_{s=1}^t f_s \left(\sum_{k=1}^{t-s} y_k - \alpha \cdot y_{k-l} \right) \end{aligned}$$

In order to find the worst-case failure distribution, i.e., the failure probabilities that maximize $n(1, T)$, consider the problem

$$\begin{aligned} &\text{maximize}_{f_1, f_2, \dots, f_t} \sum_{s=1}^t f_s \left(\sum_{k=1}^{t-s} y_k - \alpha \cdot y_{k-l} \right) \\ &\text{subject to} \quad \sum_t f_t \leq \beta, \\ &\quad f_t \geq 0, \forall t. \end{aligned}$$

The optimal cost will be simply $\beta \cdot \max_{s \leq t} \left(\sum_{k=1}^{t-s} y_k - \alpha \cdot y_{k-l} \right)$ and the optimal failure distribution will have only one non-zero component. Thus, we obtain the inequality

$$\begin{aligned} n(1, T) &\leq \max_{t \in [1, T]} \beta \cdot \max_{s \leq t} \left(\sum_{k=1}^{t-s} y_k - \alpha \cdot y_{k-l} \right) \\ &= \max_{t \in [1, T]} \beta \left(\sum_{s=1}^t y_s - \alpha \cdot y_{s-l} \right). \end{aligned}$$

Which is the bound in the proposition. \square