Reverse Logistics for Consumer Electronics: Forecasting Failures, Managing Inventory, and Matching Warranties

by

Andre du Pin Calmon

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

Signature redacted

Author ....
Sloan School of Management
February 1, 2015

Signature redacted

Certified by ....
Stephen C. Graves
Abraham J. Siegel Professor of Management Science & Engineering Systems
Thesis Supervisor

Signature redacted

Accepted by ....
Dimitris Bertsimas
Boeing Professor of Operations Research
Co-Director, Operations Research Center

Reverse Logistics for Consumer Electronics: Forecasting Failures, Managing Inventory, and Matching Warranties

by

Andre du Pin Calmon

Submitted to the Sloan School of Management
on February 1, 2015, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

The goal of this thesis is to describe, model, and optimize reverse logistics systems commonly used in the Consumer Electronics industry. The context and motivation for this work stem from a collaboration with an industrial partner, a Fortune 500 company that sells consumer electronics and is one of the top retailers in its sector. The thesis is divided into three parts.

In the first part of the thesis we model and analyze the problem of forecasting failures of new products. When a new device is introduced to the market there is limited information available about its failure time distribution since most devices have yet to fail. However, there is extensive failure time data for prior devices, as well as evidence that the failure time distribution for new devices can be forecast from the data for prior devices. In this setting, we propose two strategies for forecasting the failure distribution of new products that leverages the censored failure observations for the new devices as well as this massive amount of data collected for prior devices. We validate these strategies using data from our industrial partner and using data from a social enterprise located in the Boston area.

The second part of the thesis concerns inventory management in a reverse logistics system that supports the warranty returns and replacement for a consumer electronic device. This system is a closed-loop supply chain since failed devices are refurbished and are kept in inventory to be used as replacement devices or are sold through a side-sales channel. Furthermore, managing inventory in this system is challenging due to the short life-cycle of this type of device and the rapidly declining value for the inventory that could potentially be sold. We propose a stochastic model that captures the dynamics of inventory of this system, including the limited life-cycle and the declining value of inventory that can be sold off. We characterize the structure of the optimal policy for this problem. In addition, we introduce two heuristics: (i) a certainty-equivalent approximation, which leads to a simple closed form policy; and (ii) a dual balancing heuristic, which results in a more tractable newsvendor type model. We also develop a robust version of this model in order to obtain bounds for the overall performance of the system. The performance of these heuristics is analyzed using data from our industrial partner.

The final part of the thesis concerns the problem faced by a consumer electronics retailer when matching devices in inventory to customers. More specifically, we analyze a setting where there are two warranties in place: (i) the consumer warranty, offered by the retailer to the consumer, and (ii) the Original Equipment Manufacturer (OEM) warranty, offered by the OEM to the retailer. Both warranties are valid for a limited period (usually 12 months), and once warranties expire, the coverage to replace or repair a faulty device ends.

Thus, a customer does not receive a replacement if he/she is out of consumer warranty, and the retailer cannot send the device to the OEM for repairs if it is out of OEM warranty. The retailer would ideally like to have the two warranties for a device being matched, i.e., the customer would have the same time left in his consumer warranty as the device would have left in the OEM warranty. A mismatch between these warranties can incur costs to the retailer beyond the usual processing costs of warranty requests. Namely, since a device can fail multiple times during its lifecycle the replacement device sent to customers that file warranty requests can lead to out-of-OEM-warranty returns. In order to mitigate the number of out-of-OEM-warranty returns, we propose an online algorithm to match customers that have filed warranty claims to refurbished devices in inventory. The algorithm matches the oldest devices in inventory to the oldest customers in each period. We characterize the competitive ratio of this algorithm and, through numerical experiments using historical data, demonstrate that it can significantly reduce out of warranty returns compared to our partner's current strategy.

Thesis Supervisor: Stephen C. Graves
Title: Abraham J. Siegel Professor of Management Science & Engineering Systems

Acknowledgments

My time at MIT was a truly inspiring and enlightening experience. In large part, this is due to the never failing support, mentorship, and guidance from my advisor Steve Graves. Working and teaching with Steve was both inspiring and transformative, and I will always be proud of having been his student. Thank you for helping me broaden my horizons.

I would also like to thank the members of my committee, Karen Zheng and Vivek Farias, for their support throughout my thesis research. Both Karen and Vivek played an important role during my Ph.D. studies, and I am grateful for their interest, helpful comments they provided, and all the discussions that we had.

Many other MIT faculty were also very important during my time at MIT. Retsef Levi provided crucial support and guidance during my first year, and many of the ideas and concepts that I learned working with him are reflected in this thesis. Dimitris Bertsimas played a key role both as a professor and as the ORC co-director, and was always available when I needed advice or mentorship. Georgia Perakis was also very supportive since the beginning of my graduate studies, and it was always a pleasure to interact with her. I am thankful to all of them.

At the Operations Research Center, I had the privilege to study with brilliant colleagues, who also became some of my closest friends. They provided invaluable support both academically and personally, each in their unique way. I would like to acknowledge my close friends Adam Elmachtoub, Vishal Gupta, Allison O'Hair, Forin Ciocan, Ross Anderson, Gonzalo Romero, Angie King, Fernanda Bravo, Kris Johnson, and Nicholas Howard for all the academic discussions, support throughout the thesis, and for making my time in Boston truly special.

I am also thankful for many other colleagues at the ORC who helped foster the stimulating environment in which this thesis was developed. Maxime Cohen, Paul Grigas, and Leon Valdes provided great times and discussions at Sloan. Also, Jason Acimovic, Rong Yuan, and Annie Chen, students from Steve Grave's group, provided invaluable feedback for this work.

I was also privileged to have friends and mentors outside of MIT that helped me develop the skills used to create the practical side of this work. In particular, I would like to thank Marcelo Ballestiero, for his mentorship, Hugh Barrigan, for the business adventures, and

Christin Sander, for all the support during the last few years.

I am extremely grateful to my twin brother and MIT colleague Flavio Calmon. His support, advice, and twinship were key for the success of this research. Words cannot express how blessed I feel to have him as a brother.

I am also forever grateful to my wife, Raquel Machado. Raquel's never failing love and support has now expanded over three continents and three different universities. Thank you for exploring the world together with me.

Finally, I would like to thank my parents Paulo and Katya Calmon for all the sacrifices and investment that they made to ensure that Flavio and I had the opportunity to be the best we could be. None of this would be possible without them.

Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

Chapter 1

Introduction

In 2014, U.S. consumers paid over $20 billion dollars buying and extending warranties for electronic devices they purchased[1]. In addition, costs related to warranty management are a significant expense for most companies in the consumer electronics sector. For example, Apple Inc., one of the largest players in this segment, spent $2.9 billion dollars in 2013 alone with warranty related costs [2]. Furthermore, the repair and refurbishment of electronic devices has become a key issue for government and industry, particularly because of the environmental concerns related to disposing and managing the end-of-life of these devices.

Within this setting, the goal of this thesis is to describe, model, and optimize a type of reverse logistics system commonly used in the consumer electronics industry for managing device failures and customer warranty claims. We study three different facets of this type of system and develop a theoretical framework that addresses key issues in each one of them. Namely, the issues that we address are:

- **Forecasting warranty claims:** We model and analyze the problem of forecasting failures of new products. When a new device is introduced to the market there is limited information available about its failure time distribution since most devices have yet to fail. However, many companies usually collect extensive failure time data for prior devices, as well as evidence that the failure time distribution for new devices can be forecast from the data for prior devices. In this setting, we propose two strategies for forecasting the failure distribution of new products that leverages the censored failure observations for the new devices as well as the potentially massive

[1] http://www.warrantyweek.com/archive/ww20141009.html
[2] http://goo.gl/iIYFwg

amount of data collected for prior devices. We validate these strategies using data from our industrial partner and using data from a social enterprise located in the Boston area.

- **Inventory Management:** we address inventory management in reverse logistics systems that support the warranty returns and replacement for consumer electronic devices. This type of system is a closed-loop supply chain since failed devices are refurbished and are kept in inventory to be used as replacement devices or are sold through a side-sales channel. Furthermore, managing inventory in this system is challenging due to the short life-cycle of this type of device and the rapidly declining value for the inventory that could potentially be sold. We propose a stochastic model that captures the dynamics of inventory of this system, including the limited life-cycle and the declining value of inventory that can be sold off. We characterize the structure of the optimal policy for this problem. In addition, we introduce two heuristics: (i) a certainty-equivalent approximation, which leads to a simple closed form policy; and (ii) a dual balancing heuristic, which results in a more tractable newsvendor type model. We also develop a robust version of this model in order to obtain bounds for the overall performance of the system. The performance of these heuristics is analyzed using data from our industrial partner.

- **Warranty Matching:** we analyze the problem faced by a consumer electronics retailer when matching devices in inventory to customers. More specifically, we analyze a setting where there are two warranties in place: (i) the consumer warranty, offered by the retailer to the consumer, and (ii) the Original Equipment Manufacturer (OEM) warranty, offered by the OEM to the retailer. Both warranties are valid for a limited period (usually 12 months), and once warranties expire, the coverage to replace or repair a faulty device ends. Thus, a customer does not receive a replacement if he/she is out of consumer warranty, and the retailer cannot send the device to the OEM for repairs if it is out of OEM warranty. The retailer would ideally like to have the two warranties for a device being matched, i.e., the customer would have the same time left in his consumer warranty as the device would have left in the OEM warranty. A mismatch between these warranties can incur costs to the retailer beyond the usual processing costs of warranty requests. Namely, since a device can fail multiple times during its lifecycle the replacement device sent to customers that file warranty requests can lead to out-of-OEM-warranty returns. In order to mitigate the number of out-of-OEM-warranty returns, we propose an online algorithm to match customers that have filed warranty claims to refurbished devices in inventory. The algorithm matches the oldest devices in inventory to the oldest customers in each period. We characterize the competitive ratio of this algorithm and, through numerical experiments using historical data, demonstrate that it can significantly reduce out of warranty returns compared to our partner's current strategy.

Even though this work was developed with practical applications in mind, the theoretical framework developed throughout this work stands on its own as a contribution to the theory of closed-loop supply chains. Furthermore, since several companies, and e-tailers in particular, use the same type of reverse operations strategy that we analyze, our findings can find a broader application in the retail industry.

The remainder of this chapter will be structured as follows. In Section 1.1, we describe the reverse operations at an industrial partner that played a central role in the development of this work. In Section 1.2 we contextualize our contributions through a brief literature review. Finally, Section 1.3 serves as an outline of the thesis and we summarize the main results in each chapter.

1.1 Reverse-Logistics at our Industrial Partner

The practical backdrop for this work stems from a collaboration with an industrial partner, a Fortune 100 company that is a Wireless Service Provider (WSP) and that is also one of the top retailers in its sector. For this company, as for many other retail businesses, the management of warranty claims and regret returns is a key issue. In fact, the volume of warranty claims for products commercialized by our industrial partner is substantial (in the order of thousands per day), and a significant portion of sold items are returned. Coupled with the short life cycle of their products, usually less than one year, this leads to large levels of inventory of refurbished products that suffer fast depreciation and that are expensive to dispose.

This company uses a reverse logistics model that is similar to the one adopted by other retailers, especially on-line retailers. In this model, there are two warranty contracts in place:

(i) the consumer warranty, and (ii) the Original Equipment Manufacturer (OEM) warranty. The consumer warranty protects the consumer against any defects in the purchased product and also provides the consumer a period for "regret returns". In addition, the consumer warranty has strict requirements - when a warranty claim is filed, a new or refurbished item is immediately shipped by our industrial partner to the consumer together with a pre-paid shipping label so that the customer can return his original unit. Thus, a replacement item is sent *before* the original item is received. The OEM warranty, on the other hand, is offered by the OEM to the WSP, covering every device purchased from the OEM. This warranty is slow - a defective product sent to the OEM takes weeks or months to be fixed and a replacement device is not shipped immediately.

Because of the differences between the OEM and consumer warranty contracts, our partner WSP has a reverse logistics facility that is dedicated to processing customer-regret returns and customer warranty claims. This facility also holds inventory of refurbished devices and can execute repairs if the defect in the returned product is small or not covered by the OEM warranty. However, if the returned product has a defect that is covered by the OEM warranty, the device is sent to the OEM for repair and refurbishment. Since this reverse logistics facility uses repaired and refurbished devices to satisfy consumer warranty claims, this system fits in the context of closed-loop supply chains (CLSCs).

If at any point in time the WSP deems that it has too much inventory of refurbished devices, it can sell excess inventory through a side-sales channel. Thus, the key tradeoff for managing inventory in this system is deciding between keeping a device in stock to satisfy some future warranty claim, or selling the device through a side-sales channel. Managing this trade-off is challenging due to three issues: (i) the non-stationary nature of the demand for replacements; (ii) the short life-cycle of these devices; and (iii) the fast value depreciation faced by the refurbished device in inventory that could potentially be sold.

Both the consumer warranty and the OEM warranty are valid for a limited period (usually 12 months), and once warranties expire, the coverage to replace or repair a faulty device ends. Namely, a customer does not receive a replacement if he/she is out of consumer warranty, and the retailer cannot send the device to the OEM for repairs if it is out of OEM warranty. In addition, the OEM warranty is associated to a specific device, while the consumer warranty is specified to the consumer.

The WSP would ideally like to have the two warranties for a device being matched, i.e., the customer would have the same time left in his consumer warranty as the device would have left in the OEM warranty. A mismatch between these warranties can incur costs to the retailer beyond the usual processing costs of warranty requests. Namely, this extra-cost is incurred when a customer still covered by the consumer warranty has a device that fails, and this device is not covered by the OEM warranty. In this case, the WSP will then either pay for the OEM to repair the device, which incurs additional costs to the system, or it will scrap the device. At our partner WSP, these out-of-OEM-warranty devices are a significant source of cost for their reverse operations. Furthermore, since a device can fail multiple times during its lifecycle, the replacement device sent to customers that file warranty requests can lead to out-of-OEM-warranty returns. Also, the OEM warranty does not restart once a device is remanufactured and it is not paused while a device is in stock at the WSP, such that "old" devices, with little OEM warranty left, can potentially be sent to customers as replacements.

1.2 Literature Review

This thesis builds upon a vast body of knowledge that addresses closed-loop supply chains (CLSC), reverse logistics, forecasting, and on-line algorithms. Although each of the following chapters will have its own literature review, we provide a broad overview of prior work in this area.

Two early examples of papers that consider the management of inventory of repaired and refurbished items that influence this work are Simpson (1978) and Allen and D'Esopo (1968). A more recent summary of research in this field can be found in Guide and Van Wassenhove (2009), which presents an overview of the literature in CLSC and discuss future directions of research.

Forecasting with censored or truncated observations has been an active area of research for the last 50 years since the introduction of the Kaplan-Meier Estimator in Kaplan and Meier (1958). Using the EM algorithm for incomplete data was introduced in Dempster et al. (1977), and our strategy builds upon this work. Taylor (1995) introduces a semi-parametric approach for maximum-likelihood information with censored data when failures are exponentially distributed. A detailed discussion of the role of regularization for obtaining sparse representations was done in Tibshirani (1996), and an example of the use of regularization for quantile regression with censored data can be found in Lindgren (1997). In the operations literature, Karim and Suzuki (2005) present a comprehensive literature review on statistical methods for analyzing warranty claim data. To the best of our knowledge, our work is the first to apply regularized regression for hazard rate estimations, and to use this completely data-driven approach to forecasting failures in a closed-loop supply chain context.

A detailed analysis of the importance and challenges related to forecasting warranty claims at our partner WSP is explored in Petersen (2013). The forecasting part of this work benefited from the same partnerships, discussions, and data as Petersen (2013) and, because of this, shares many of its core ideas. However, while Petersen (2013) presents results and strategies tailored to the WSP, this chapter frames the discussion in more generic terms, and we believe our approach has a wide range of applications.

With regard to inventory management for warranty replacements, Huang et al. (2008) offers a detailed overview of the literature in warranty management, and analyzes inventory management when there is demand for both new and replacement items,without taking into account remanufacturing. Khawam et al. (2007) use a simulation approach to obtain inventory management policies for Hitachi. From a conceptual level, product warranty management is discussed by Murthy and Blischke (1992). Also, the connection between the warranty and logistics literature is discussed by Murthy et al. (2004), and the relationship between warranty service and customer satisfaction is discussed. The impact of regret returns on inventory management is analyzed in de Brito and van der Laan (2009) where the authors highlight the effect of imperfect information about returns on inventory management. Different models for remanufacturing products in a CLSC are analyzed in Savaskan et al. (2004). Finally, our work fits in the wider field of perishable inventory systems, and a review of the results in this area can be found in Nahmias (2011).

As for the theoretical tools that we use in this paper, Bertsekas (2005) presents an overview of Dynamic Programming in the finite and infinite horizon setting, including many examples in inventory management. Balancing policies have also been an active area of study, and Levi et al. (2008) contains an application of this approach in inventory management.

Research in matching items and individuals has also been very active during the last 50 years. One of the seminal works in this area is on the Hungarian algorithm, and can be found in Kuhn (1955). More recently, the problem of matching under preferences has played a central role in economics and operations research. In particular stable matching problems, which is discussed in Gusfield and Irving (1989), has been an area of extensive research. A review of tools for analyzing on-line algorithms, which encompass the type of policies that we use in the third part of the thesis, can be found in Albers (2003). Finally, the type of matching problem that we consider in this thesis has yet to be covered in the literature, and we believe that our application provides an interesting context for assignment and matching algorithms.

1.3 Outline and Summary of Main Contributions

In this section we will briefly summarize three parts of the thesis. Each part corresponds to a different aspect of managing the closed-loop reverse logistics system we examine in this work. Chapter 2 addresses the problem of forecasting failures and estimating failure age distributions in the context of our partner WSP. In Chapter 3, we present and analyze a control model that captures the dynamics and key decisions of inventory in the reverse supply-chain. Finally, in Chapter 4 we study the problem of assigning refurbished devices in inventory to customers, and we analyze different assignment strategies

Each chapter is, for the most part, independent and can be read as an individual work. The reader should feel free to dive into the chapter that is the most relevant to their interests.

Chapter 2: Data-Driven Failure Age Estimation

The closed loop nature of this reverse logistics system makes forecasting failures important in order to determine the amount of inventory of replacement devices that should be kept in stock. Furthermore, an effective forecasting strategy can help the WSP identify if a device recently introduced into the market is having a higher than expected failure rate, allowing for operational and strategic adjustments such as fixing the hardware/software of a device or changing the way that it is marketed. However, forecasting failures at the beginning of the life-cycle of a device can be challenging, since the number of observations is limited and failure observations might be censored or truncated. For example, if t weeks have passed after the launch of a device, then no customer has a device that is more than t weeks old, and failures after this period are still unobserved. Because of this, traditional non-parametric approaches, such as the Kaplan-Meier estimator, are not sufficient for estimating the failure distribution of a new device. On the other hand, the WSP has a large amount of historical data available, since it collects detailed information about sales and failure times from the different devices that it sells. In fact, since our partner WSP is one of the largest players in this market, it has data from millions of customer purchases and failures.

In this setting, we propose and analyze two methods for estimating failure distributions of newly launched devices that leverages the historical data of failures from other devices. The proposed strategies are based on a *hazard rate model* developed under the assumption that customers in the same *cohort* have devices that fail according to the same age-dependent failure distribution. A cohort is a pool of customers that share similar features (e.g. phone model owned, data plan, etc.).

The first estimation strategy uses an Expectation-Maximization (EM) type algorithm to estimate the parameters of a mixture model. Here, we assume that hazard rates of devices in a new cohort are drawn from a mixture of scaled hazard rate distributions built from historical data. Furthermore, since maximizing the likelihood function of a set of observations is intractable, we use an EM approach to maximize a lower bound of the likelihood and obtain an estimate of the parameters of the mixture model.

The second estimation strategy, which we call *hazard rate regression*, uses a model selection method, where we assume a "basis" set of hazard rate distributions determined from historical data. We then use a regularized regression to identify and fit the relevant hazard rates distributions from the basis to the observed failures from the new cohort. This allows for a sparse representation of the estimated hazard rate distribution, which can be useful depending on the context of the estimation problem. For example, a sparse solution can help identify which cohorts are the most similar and can help with an investigation of the features that these cohorts have in common.

Both of these estimation strategies assume that the hazard rate distribution that is being estimated can be described by a mixture of scaled versions of the hazard rate distributions in the basis. Using data from our partner WSP we argue that this assumption holds in practice. Furthermore, to the best of our knowledge, this type of approach is novel and has not been studied before in the Operations Management community.

In the final part of this chapter, we describe how these estimation strategies can be used to create a forecast of the volume of warranty requests received by our partner WSP, and introduce different metrics to measure the quality of the forecast. We also examine both estimation strategies through a series of numerical experiments using data from our partner WSP and using data from Project Repat, a Boston-based social enterprise that transforms old t-shirts into quilts. Through these experiments, we observe that both the EM algorithm and the Regression approach have a similar average performance, but the performance of the regression approach has a lower variance. Furthermore, the regression approach leads to sparser representation of the hazard rate distribution, while the estimate produced by the EM algorithm is dense in the mixture parameters.

Chapter 3: Inventory Management in a Closed-Loop Supply Chain

The short life-cycle of devices sold by our partner WSP coupled with the fast value depreciation of these devices, makes managing inventory at the reverse logistics facility a challenging problem. Furthermore, the demand for replacement devices is non-stationary, and not all devices sent to the OEM can be refurbished. In this system, devices arrive in inventory either as seed-stock from the OEM (usually 1% of sales) or as refurbished devices corresponding to previous failures or regret returns. Devices leave inventory when they are sent as replacements to customers that filed warranty claims, or through a side-sales channel that the WSP uses to sell excess inventory.

The second part of the thesis proposes and analyzes two versions of the inventory management problem. First, we discuss a discrete-time deterministic version, for which hazard rates are fixed fractions of sold devices that fail. We prove the optimal policy for this case and also present a worst case analysis. The second version introduces a discrete-time stochastic model, for which we prove the structure of the optimal policy and discuss a heuristic for managing this system. These models depart from the other inventory management models in the literature since it incorporates the short life-cycle of devices, the fast value depreciation of the devices as well as by making no assumptions on demand and arrival distributions.

In the deterministic model, we assume that both the demand and the arrival processes are deterministic and known. Since this problem has a finite horizon corresponding to the life-cycle of the device, finding the optimal buying and selling quantities in each period is equivalent to solving a linear optimization problem. Despite being an optimization problem that can be solved efficiently using a numeric solver, we prove and discuss the structure of the optimal solution for this problem. There are a three main reasons why we do this: (i) By proving the structure of the optimal policy we can make explicit the relationship between the optimal policy and the dynamic cost structure of the problem; (ii) The results from this analysis will be used to prove the optimal policy of the stochastic version of the problem; (iii) The Certainty-Equivalent approximation of the stochastic problem is based on the deterministic model.

We assume that the costs of purchasing a new device into inventory and the revenue of selling a device in a side-sales channel are both non-increasing. Furthermore, we also assume that the revenue obtained selling a refurbished device in a side-sales channel is lower than the cost of sourcing a new device into inventory. Thus, in this setting, the optimal sourcing strategy will be myopic in the sense that we only buy enough items to satisfy the unmet demand for replacement devices in the current period. Conversely, the optimal selling quantity in some time period will depend on the inventory level at the beginning of the period and 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. The maximum total net demand in some interval is defined as the maximum cumulative difference between the demand for replacement devices and the amount of devices that arrive from the OEM. Thus, the maximum total net demand acts as a sell-down-to 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.

We also address the question: *If the hazard rate distribution is 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, it can be used to bound the operational cost of supporting warranty of a new device, which is useful for planning considerations when releasing a new device into the market. We answer this question by examining a model where an adversary can choose the hazard rate distribution, and prove a tight worst case bound in this setting.

In the stochastic model, we assume that replacement requests are generated by the hazard rate distribution discussed in the first part of the thesis. We prove the structure of the optimal policy in this case, and show that it has the same sell-down-to structure as the optimal policy in the deterministic case. However, the optimal sell-down-to level in every period depends on the distributions of the demand and arrival processes and cannot be computed in closed-form. Because of this, we introduce two heuristics for managing inventory.

The first heuristic is a *certainty-equivalent* approximation where we obtain a suboptimal inventory control policy by approximating the uncertainty in the problem by its average value. Since the volume of devices in the WSP's reverse logistics system is very large, this heuristic works well in practice. We construct the certainty-equivalent approximation using the hazard rate model discussed the first part of the thesis. In this case, the certainty-equivalent approximation to the optimal policy still purchases items myopically, but uses an approximation to calculate the sell-down-to-level in each time period.

The second heuristic is the *cost-balancing policy* that takes into account the uncertainty of the demand and arrival process. More specifically, we solve a newsvendor-type problem that balances the costs of selling too few or too many items. We find the optimal sell-down-to level through a Sample Average Approximation (SAA), a well studied approach in the Operations Management literature.

Through numerical experiments, we analyze the performance of these policies and simulate their sensitivity with respect to changes in different parameters of the system such as number of devices, failure distribution, OEM lead time, seed-stock and the loss at the OEM. As a benchmark, we compare the performance of these policies with the *clairvoyant policy*, the policy that knows ex-ante the sample path of the device failures and the arrivals from the OEM.

Next we compare the certainty equivalent policy with the *cost-balancing* policy for different distributions of device failures. We observe that the cost-balancing policy usually leads to a better performance than the certainty-equivalent approximation. However, when the number of devices is large (which is the case for our partner WSP) the certainty-equivalent approximation achieves a near-optimal performance and is sufficient for practical applications.

Finally, we analyze the performance of the different policies using real-world data from a device sold by the WSP. This simulation incorporates learning, i.e., the methodology discussed in the previous chapter is employed and the hazard rate distribution of the device is updated as new information on failure rates becomes available. We observe that in practical settings both policies capture over 90% of the clairvoyant profit.

Chapter 4: Warranty Matching in a Closed-Loop Supply Chain

As mentioned in the introduction of this chapter, there are two warranties in place in this system: (i) the consumer warranty (offered by the WSP to its consumers), and the (ii) OEM warranty (offered by the OEM to the WSP). Ideally the two warranties would be matched, i.e., the customer would have the same time left in his consumer warranty as the device would have left in the OEM warranty. A mismatch between these warranties incurs costs to the retailer beyond the usual processing costs of warranty requests. Namely, this extra-cost is incurred when a customer still covered by the consumer warranty has a device that fails, and this device is not covered by the OEM warranty. In this case, the WSP will then either pay for the OEM to repair the device, which incurs additional costs to the system, or it will scrap the device and the device leaves the system. At our partner WSP, these out-of-OEM-warranty devices are a significant source of cost for their reverse operations.

Since a device can fail multiple times during its lifecycle, the replacement device sent to customers that file warranty requests can lead to out-of-OEM-warranty returns. Also, the OEM warranty does not restart once a device is remanufactured and it is not paused while a device is in stock at the WSP, such that "old" devices, with little OEM warranty left, can potentially be sent to customers as replacements. At the WSP's reverse logistics facility, devices in stock were matched at random to consumers that placed warranty claims. More specifically, refurbished devices received from the OEM were not sorted by time left in OEM warranty, and customer requests were also not sorted according to the time left in their customer warranty. This would lead to "old" devices being sent to "young" customers, creating a scenario where a customer with a few months left in its consumer warranty receives a device with an expired OEM warranty. Conversely, this would also lead to cases where "young" devices were sent to "old" customers, effectively wasting OEM warranty coverage time.

Given this setting, in the third part of the thesis we model the problem of matching devices to customers and analyze different assignment strategies and how they impact mismatch costs and out-of-OEM-warranty returns. Note that the assignment strategy is crucial in mitigating out-of-warranty returns, and simple strategies that do not take into account the closed-loop nature of the system can lead to a large average mismatch. For example, a first-in-first out strategy for assigning replacement devices to customers, not taking into account customer ages, might lead to large mismatches. Similarly, randomly assigning devices to customers, which was the practice used by our partner WSP, might also lead to large mismatches and, consequently, to poorly matched assignments.

The three assignment strategies that we focus on are:

- The *Random assignment policy*: where devices in inventory are randomly assigned to customers that require replacement devices, ignoring the time remaining in both the customer and device warranties. This was the policy used by the WSP at the time that our collaboration began;

- The *Youngest-Out-First (Myopic) policy*: where in every time period devices in inventory are assigned to customers as to minimize the mismatch in that specific period. We prove that this is the optimal single-period assignment strategy in our formulation;

- The *Oldest-Out-First policy*: a policy that always assigns the oldest devices in inventory to the oldest customers that require replacements.

Our first analysis involves assuming that the activation date of customers that need replacements devices is random and we assume that failure ages of devices are i.i.d. Also, we assume that the total number of device failures is constant in every period and given by n, and that there are $m \geq n$ devices in inventory at the reverse logistics facility in every period. In this context, we prove distribution-free upper and lower bounds on the expected mismatch cost for the random assignment policy. These bounds have a practical interpretation and can help a plant manager decide if it is worth investing in a matching policy other than random assignment.

We then consider the Youngest-Out-First policy, where customers and devices are sorted by age and matched from youngest to oldest. We prove that, in the long-run and for a lead time of l, the mismatch cost of the YOF policy will be very close to l. In fact, the distribution of the mismatch cost will have an exponentially decreasing tail. However, this policy has a major drawback. If we allow m and n to be random variables that change over time, this policy may lead to an accumulation of "old" devices in the system, since it uses the youngest first, and when n fluctuates, devices that are out of OEM warranty might be sent to customers.

In order to address this, we propose the Oldest-Out-First policy. This policy also sorts devices and customers by age but, instead of matching them from youngest to oldest, it matches them from oldest to youngest. The intuition behind this policy is that in the long-run it is not worth allowing devices to "age" in inventory, even though using them immediately is not the optimal short-term thing to do. We analyze this policy as an online algorithm and, by assuming certain conditions for the system's behavior, we prove the competitive ratio of this policy. The competitive ratio allows for a comparison between the *clairvoyant policy*, i.e., the policy that "knows" all the information of the system, and the oldest-out-first policy.

We evaluate these policies through numerical experiments that use data from our partner WSP and also using a simulated scenario where failures are chosen from a pre-set distribution. We compare the performance of these policies using two metrics:

- *Average uncovered time per replacement device shipped*: if a refurbished device of age j is sent as a replacement to a customer of age i, the *uncovered time* of the customer will be $\max(j-i,0)$. Since we assume that both the customer warranty and the OEM warranty have the same length, this represents the amount of time that a customer is still covered by the customer warranty while the device he/she owns is not covered by the OEM warranty. This is a measure of exposure of the WSP with regards to out-of-warranty returns;

- *Percentage of failures that are out-of-warranty*: the percentage of all the failures that happened when the customer was covered by the customer warranty but the device that failed was not covered by the OEM warranty. In this case, the device is either scrapped or the WSP has to pay for its repair/refurbishment.

In our experiments, we observe the OOF significantly decreases the average number of uncovered weeks with respect to random matching, and that it performs better than the Youngest-Out-First policy since "old devices" do not accumulate in stock over time. We also observe that both the OOF and the Youngest-Out-First policy present significant improvements over random matching due to the *power of sorting* requests and devices.

Chapter 2

Data-Driven Failure Age Estimation in a Closed Loop Supply Chain

2.1 Introduction

One of the main challenges in supply chain management is forecasting demand, particularly when introducing a new product into the market. In this chapter, we propose and analyze a strategy for forecasting the demand for replacement devices in a large Wireless Service Provider (WSP). More specifically, this WSP offers a warranty to their clients (usually 12 months in duration), and clients covered by the warranty that have devices that fail are entitled to receive a replacement from the WSP. For the WSP, estimating the demand for replacement devices is critical for planning the operations that support the customer warranty.

If a customer covered by the WSP's customer warranty has a device that fails, he/she receives a replacement which is shipped overnight from the WSP's reverse logistics facility. The replacement device is usually not a new device, but a device that failed at an earlier time and was refurbished. When a customer receives a replacement device, he/she ships the broken device to the WSP (usually within one or two weeks), which then proceeds to refurbish/repair the device (if possible) and stores it in inventory in order to use it as a replacement device in the future. If the WSP finds that, at some point, it has too many refurbished devices in inventory, excess devices can be sold through a side-sales channel.

The closed loop nature of this reverse logistics system makes forecasting failures important in order to determine the amount of inventory of replacement devices that should be kept in stock. Furthermore, an effective forecasting strategy can help the WSP identify if a device recently introduced into the market is having a higher than expected failure rate, allowing for operational and strategic adjustments such as fixing the hardware/software of a device or changing the way the devices are marketed. However, forecasting failures at the beginning of the life-cycle of a device can be challenging, since the number of observations is limited and failure observations might be censored or truncated. For example, if t weeks have passed after the launch of a device, then no customer has a device that is more than t weeks old, and failures after this period are still unobserved.

On the other hand, the WSP has a large amount of historical data available, since it collects detailed information about sales and failure times from the different devices that it sells. In fact, since our partner WSP is one of the largest players in this market, it has data from millions of customer purchases and failures. Leveraging this information will play a key role in the estimation strategies discussed in this chapter. Also, this company launches around a dozen to two dozen new devices per year and, since the life-cycle of these devices is between one and two years, it is usually managing replacement request from around 40 different devices.

In this chapter, we propose and analyze methods for estimating failure distributions of newly launched devices that leverages the historical data of failures from other devices. The proposed strategies are based on a *hazard rate model* developed under the assumption that customers in the same *cohort* have devices that fail according to the same age-dependent failure distribution. A cohort is a pool of customers that share similar features (e.g. phone model owned, data plan, etc.). For example, a device model could be iPhone 5's sold in August in Boston to customers in a certain type of plan.

We propose two different strategies for estimating the hazard rates of a new cohort of customers. Both of these strategies use the hazard rates determined using historical data from other cohorts containing different customers and/or devices. The first estimation strategy uses an Expectation-Maximization (EM) type algorithm to estimate the parameters of a mixture model. Here, we assume that hazard rates of devices in the new cohort are drawn at random from a set of scaled hazard rate distributions built from historical data.

The second estimation strategy uses a model selection method, where we assume a "basis" set of hazard rate distributions determined from historical data. We then use a regularized regression to identify and fit the relevant hazard rates distributions from the basis to the observed failures from the new cohort. The practical performance of these two approaches will be analyzed through numerical experiments that use data from our partner WSP.

Note that both these strategies are parametric, although the basis functions used are all built from data. In fact, a pure non-parametric estimation strategy that deals with censoring (such as the Kaplan-Meier estimator) would not be effective in this case, since all failure observations of a recently launched device are truncated in the same time period.

A detailed analysis of the importance and challenges related to forecasting warranty claims at our partner WSP is explored in Petersen (2013). This chapter benefited from the same partnerships, discussions, and data as Petersen (2013) and, because of this, shares many of its core ideas. However, while Petersen (2013) presents results and strategies tailored to the WSP, this chapter frames the discussion in more generic terms, and we believe our approach has a wide range of applications.

In fact, in Section 2.5.4, we present an application of our estimation strategy to Project Repat, a social enterprise that transforms old t-shirts into quilts. In this case, the company needed to forecast the volume of t-shirts that were sent by their customers in order to make staffing decisions in their textile plant. Since customers first purchase the quilts on-line and then mail their t-shirts to Project Repat using a pre-paid shipping label, the problem is similar to the one faced by our partner WSP.

The remainder of this chapter is structured as follows. In Section 2.2 we introduce the model for device failures and briefly discuss the Kaplan-Meier estimator, a non-parametric approach for estimating hazard rates. In Section 2.3 we formalize our estimation problem and its challenges. In Section 2.4 we present two distinct estimation strategies. The first strategy is based on the EM algorithm, while the second strategy involves a regularized regression. In Section 5 we present a few numerical experiments including one utilizing data from our partner WSP. Finally, in Section 2.5.4 we discuss how this strategy was used at Project Repat.

2.2 Estimating Failure Rates

In this section, we will present a review of the literature on estimation with censored and truncated data. We then describe the device failure process and introduce the notation that will be used throughout this chapter. In the final part of this section, we review and discuss the Kaplan-Meier (KM) estimator. As our goal is to estimate the distribution of failure ages of devices at our partner WSP, we will frame our discussion within this context.

2.2.1 Literature Review

The history of methods for estimating failure or survival distributions of products, machines, and subjects in clinical trials has a long history, dating back to seminal work of Greenwood in the early 20th century, in Greenwood and others (1926). These methods attempt to build a distribution for the occurrence time of an event (such as age of failure of a device or death of a patient) based on a set of observations of the said event. As in most of the literature, we will call the time at which an event occurs the *failure time*. In many practical settings, these observations can be *censored*, i.e., there is no information available on the exact time that a failure occurs, only that it is outside of some interval. For example, if we are trying to estimate the failure time distribution of electronic devices sold at different times during the last few months, we have censored observations in that we have yet to observe the failure times for the devices that have yet to fail. All we can say is that their failure times are at least as long as the devices' current ages.

The most popular non-parametric approach for estimating failure time distribution is the Kaplan-Meier estimator, introduced in Kaplan and Meier (1958), and widely used in practice. We will cover the KM estimator at the end of this section section.

There are also many parametric strategies for estimating failure distributions. One classic parametric approach is the Cox Proportional Hazards Model (Cox and Oakes (1984)), where the hazard rate is assumed to be a scaled version of some baseline hazard function, and the scaling parameter is a function of certain covariates. This approach is commonly used in medical applications when analyzing the survival rate of a patient as a function of characteristics the patient has, such as genetic factors or pre-existing conditions. Other parametric strategies involve the Expectiation-Maximation (EM) type algorithms, and have been used to estimate parameters of failure time distributions (McLachlan and Krishnan

(2007)). An EM approach will be the cornerstone of one of the estimation strategies proposed in Section 2.4.

Taylor (1995) introduces a semi-parametric approach for maximum-likelihood information with censored data when failures are exponentially distributed. A detailed discussion of the role of regularization for obtaining sparse representations was done in Tibshirani (1996), and an example of the use of regularization for quantile regression with censored data can be found in Lindgren (1997). Also, the use of $L1$ penalized estimation is used in the context of the Cox proportional hazards model to estimate hazard rates in Tibshirani and others (1997) and in Goeman (2010). In the operations literature, Karim and Suzuki (2005) present a comprehensive literature review on statistical methods for analyzing warranty claim data. To the best of our knowledge, our work is the first to apply regularized regression for hazard rate estimations that uses historical data to build a basis of hazard rate distributions.

2.2.2 The device failure process

We assume that the failure age of a device can be described by a discrete failure distribution, where, for $t \in \mathbb{Z}_+$,

$$\Pr(\text{failure of device at age } t) \triangleq p_t.$$

It is also useful to describe the failure process in terms of hazard rates. The hazard rate at age t is the probability that, conditioned on the non-failure up until the beginning of age t, the device fails at age t. Thus,

$$\Pr(\text{failure at age } t \mid \text{survived to age } t) \triangleq h_t.$$

The relationship between the hazard rate and the failure distribution is

$$h_t = \frac{p_t}{1 - \sum_{k=1}^{t-1} p_k}.$$

Let the complementary cumulative distribution function (CCDF) be denoted by $\bar{F}_t$. Then, $\bar{F}_t = \Pr(\text{failure} > \text{age } t)$ and the relationship between the hazard rate and the CCDF is

$$\bar{F}_t = \prod_{k=1}^{t} (1 - h_t).$$

Thus, note that $p_1 = h_1$ and that, for $t > 1$,

$$p_t = h_t \cdot \prod_{k=1}^{t-1} (1 - h_k). \tag{2.1}$$

We also assume that a failure observation can be of two types: (i) an *uncensored* observation, i.e., the exact age that the device had when it failed is observed; (ii) a *censored* observation, i.e., if a failure observation is censored at age t, we know that the failure age of this device is strictly larger than t. Another critical assumption that we make is that *failure age and censoring are independent.* In the WSP case, this is equivalent to assume that sales date and failure date are independent. This assumption will be discussed in more detail in Section 2.3.

2.2.3 The Kaplan-Meier estimator

The Kaplan-Meier (KM) estimator is the maximum likelihood non-parametric estimator of a failure distribution when failure observations are censored. We will derive the KM estimator for the case where failure ages are discrete. The extension to continuous time is straight forward and can be found in Kaplan and Meier (1958).

We assume that we have two sets of samples in hand: (i) a set of uncensored failure observations $y = (y_1, \ldots, y_T)$, where y_t is the number of devices that failed at age t, and (ii) a set of censored failure observations $z = (z_1, \ldots, z_T)$, where z_t is the number of observations *censored* at age t, i.e. all that is known is that the failure of these devices will happen at an age strictly larger than t. The total number of observations is $\sum_{i=1}^{T} y_i + z_i$. We also assume that whether an observation is censored or not is independent of the failure age. Under these assumptions, the probability of observing this set of samples for a given failure distribution $p = (p_1, \ldots, p_T)$ is given by

$$\begin{aligned} L(p; y, z) &= \Pr(\text{observing } (y, z) \mid p) \\ &= \prod_{i=1}^{T} p_i^{y_i} \cdot \bar{F}_i^{z_i}. \end{aligned}$$

Where $L(p; y, z)$ is the likelihood (in this case, also the probability) of observing the sample y, z if the failure distribution were described by the vector p. Rewriting this expression in terms of the hazard rates, the likelihood becomes

$$\begin{aligned} L(h;y,z) &= \prod_{i=1}^{T} \left(h_i^{y_i} \cdot \prod_{j=1}^{i-1} (1-h_j)^{y_i \cdot} \right) \left(\prod_{j=1}^{i} (1-h_j)^{z_i} \right) \\ &= \left(\prod_{i=1}^{T} h_i^{y_i} \right) \cdot \left(\prod_{i=1}^{T} (1-h_i)^{z_i + \sum_{j>i} y_j + z_j} \right) \end{aligned} \tag{2.2}$$

Let $r_i \triangleq \sum_{j \geq i} y_j + z_j$ be the population of devices at *risk of failure* at time i, i.e., the number of observations with failure time greater or equal to i. Then the maximum-likelihood hazard rate vector h^* will be

$$h^* = \arg\max_h \sum_{i=1}^{T} y_i \log h_i + (r_i - y_i) \log(1 - h_i) \tag{2.3}$$

where this expression comes from taking the logarithm of the likelihood function. By taking the derivative of the log-likelihood function, it is straightforward to see that the right-hand side of (2.3) is maximized when

$$h_i^* = \frac{y_i}{r_i}.$$

Thus, we have that the maximum-likelihood CCDF is

$$\bar{F}_i^* = \prod_{j=1}^{i} \frac{r_i - y_i}{r_i},$$

which is the commonly known Kaplan-Meier estimate of the CCDF.

If the probability distribution has a discrete support, the hazard rates can be interpreted as transition probabilities in a Markov Chain with an absorbing state. This leads to a natural interpretation of the KM estimator, where the estimate of the hazard rate for some age t is the ratio between the transitions observed from the state "working at age t" to the "failure" state and the total number of transitions observed out of the state "working at age t".

For the estimation problem faced by our partner WSP, the KM estimator is not very useful for estimating the failure distribution of a new device. When a new device is launched to market, all failure observations are truncated, since the oldest device sold is no older than the time past since the launch date. Thus, if the KM estimator were applied to this case, it would not be possible to estimate the right tail of the failure distribution.

However, historical failure data indicates that the failure distribution of devices sold by the WSP have a similar shape. For example, consider the hazard rate distributions estimated from four devices sold by the WSP depicted in Figure 2-1[1]. These devices are from different manufacturers and have different specifications: devices A and B have keyboards while C and D do not; devices C and D have the same operations system, while A and B have different ones; all devices have distinct CPUs; they use different wireless technologies.

Despite the differences between the devices in Figure 2-1, their failure distributions are somewhat similar, resembling scaled versions of each other. There is a spike at the beginning of the hazard rate distributions due to regret returns and dead-on-arrival failures. Afterwards, the hazard rate is fairly constant up until one year, the time when most customer warranties expire. After one year, there is a sharp decrease since only customers that purchased extended warranties can exchange their devices.

Since failure distributions of different devices are similar, we propose estimation strategies that leverage these similarities. Thus, when a new device is introduced into the market, the "prior information" available from other devices will be used to estimate the right-tail of the failure distribution of the new device. Before presenting our estimation strategies, we will first set up the estimation problem in the next section.

2.3 Problem Set-Up

In our set-up, we assume the that customers are divided into *cohorts*. A cohort of customers is a pool of customers that share similar features and whose devices fail according to the same hazard rate. For example, features that define a cohort could include the type of device that a customer owns,customer location, or the month that the customer purchased a device. For our partner WSP, a pool of customers is given by all the customers that own the same type of device, regardless of when the device was purchased. This is illustrated in Figure 2-2. In the case of Project Repat, the other company we collaborated with, a cohort is defined as all customers that purchased a product in a given week.

We assume that there is a maximum failure age T for devices, such that, for all practical purposes, devices of age larger than T will never fail. This comes from the fact that the

[1]The failure dates in the data we obtained from the WSP corresponded to the dates when the failed device was received from the customer, not the date when failure claim was filed. Since the customer has a few weeks to return the broken device, the real failure date was a few weeks earlier.

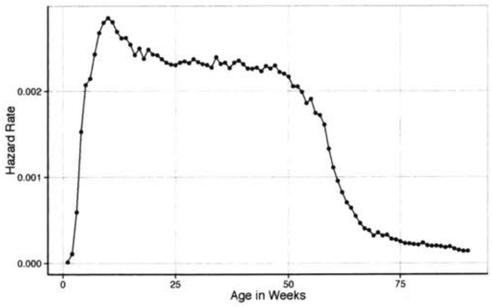

(a) Device A

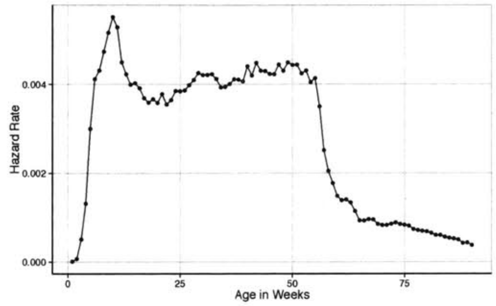

(b) Device B

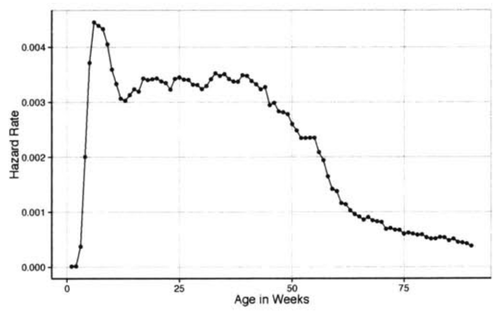

(c) Device C

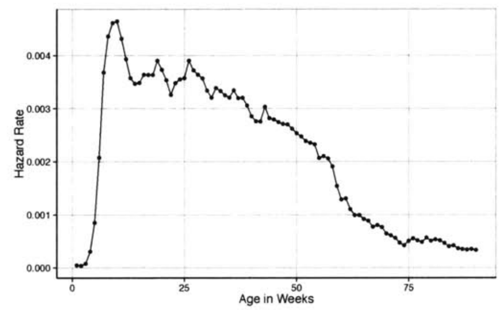

(d) Device D

Figure 2-1: Hazard Rates of four different devices sold by the WSP. The age of failures is given in weeks.

WSP offers a warranty of at most two years, and customers that are not covered by a warranty are not entitled to a replacement device. In practice, the choice of T depends on the context of the estimation problem. We denote the "true" (initially unknown) hazard rate, failure distribution, and CDF by h^*, p^*, and F^*, respectively. Also, let $\bar{t}$ be the age of the oldest device in the system. If $\bar{t} < T$, there are no observations for failure times in the interval $[\bar{t}, T]$.

For example, if a new device is launched in January and the current month is April, then $\bar{t} = 4$ months, and there are no observations of failure times larger than 4 months. This is problematic in the context of the WSP since, around this time, the WSP needs to start making decisions as to how much stock of this device to carry in its reverse logistics facility, as well as wether to take counter-measures if the failure rate of the device is too high. If there is no information about failures larger than 4 months, making tactical and strategic decisions might be difficult. Furthermore, the KM estimator does not provide solace when $\bar{t} < T$, since it is a non-parametric estimator and cannot provide information on the right tail of the failure distribution. We overcome this issue in a data-driven way by

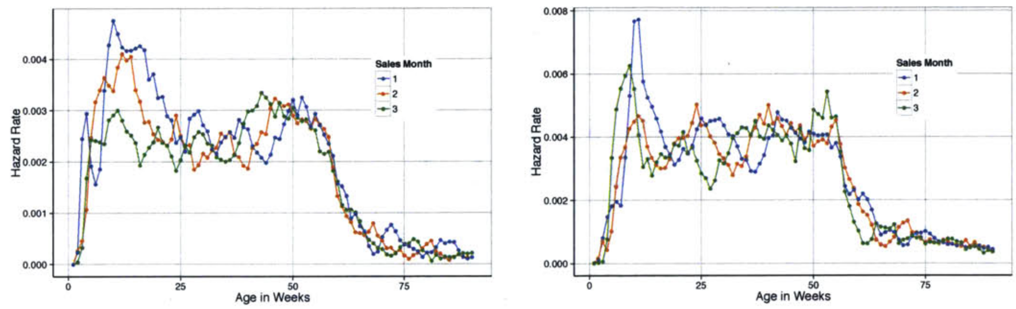

(a) Device A (b) Device B

Figure 2-2: Hazard Rates estimated from sales in three different months for two devices.

using the failure distributions of devices already in the market as priors for the estimated failure distribution.

We denote the estimate of the hazard rate of a cohort of customers given a set of uncensored and censored failure observations $y = (y_1, \ldots, y_{\bar{t}})$ and $z = (z_1, \ldots, z_{\bar{t}})$, respectively, by a vector $h(y, z) = (h_1(y, z), \ldots, h_{\bar{t}}(y, z))$. The estimate of the discrete failure distribution from the observations y and z is given by $p(y, z) = (p_1(y, z), \ldots, p_{\bar{t}}(y, z))$. The corresponding estimate of the CDF is then $F(y, z) = (F_1(y, z), \ldots, F_{\bar{t}}(y, z))$, where

$$F_t(y, z) = \sum_{i=1}^{t} p_i(y, z).$$

and

$$p_t(y, z) = h_t(y, z) \cdot \prod_{k=1}^{t-1} (1 - h_k(y, z)). \tag{2.4}$$

Ideally, we would like an estimation strategy such that if no observation is truncated, i.e., $\bar{t} = T$, and if y and z are sampled from the failure distribution defined by the hazard rates h^*, we have

$$\sup_s |h_s(y, z) - h_s^*| \to 0 \text{ almost surely as } \sum_{i=1}^{T} y_i + z_i \to \infty.$$

The KM estimator satisfies this criteria when $\bar{t} = T$ and the number of observations goes to infinity. This is clear from the Markov-Chain interpretation of the discrete hazard model, since the horizon is finite and thus all failure "states" will be visited an infinite number of times as the number of observations goes to infinity. However, as mentioned if $\bar{t} < T$ there are no observations in the interval $[\bar{t}, T]$. Also, at the WSP, censoring occurs since devices in a cohort are sold in different time periods and might not have failed up to the current time period. More specifically, if we are in period t and a device sold in period $s < t$ has not yet failed, we only know that this device might fail with a failure age greater than $t - s$. In addition, we assume that the time that a device is sold and its failure time are independent, since all devices in a cohort fail according to the same hazard rate distribution.

Let $\{\hat{h}^1, \ldots, \hat{h}^m\}$ be a collection of m different hazard rate distributions, such that $\hat{h}^j = (\hat{h}_1^j, \ldots, \hat{h}_T^j)$ represents the hazard rates for some other cohort or device or a hazard function built using expert opinion. We view this collection as a basis set for modeling the population of possible hazard rate distributions. Furthermore, we allow for each element j of the basis to be scaled by some positive parameter λ_j, such that $\lambda_j \hat{h}^j = (\lambda_j \hat{h}_1^j, \ldots, \hat{\lambda}_j \hat{h}_T^j)$, and $\lambda_j \hat{h}_i^j < 1, \forall i, j$. We use the set of distributions as the basis in a mixture model for the estimation of the true hazard distribution h^*. Namely, we assume that each device in the cohort will fail according to scaled basis element j with probability θ_j^*. Hence, we have that

$$\Pr(\text{failure at age } t | \text{survived} \geq t) = h_t^* = \sum_{j=1}^{m} \theta_j^* \cdot \lambda_j^* \cdot \hat{h}_t^j.$$

Also,

$$\Pr(\text{failure at age} > t | \text{survived} \geq t) = 1 - h_t^* = \sum_{j=1}^{m} \theta_j^* \cdot (1 - \lambda_j^* \hat{h}_t^j).$$

Note also that $\sum_{j=1}^{m} \theta_j^* \cdot (1 - \lambda_j^* \hat{h}_t^j) = \sum_{j=1}^{m} \cdot (1 - \lambda_j^* \theta_j^* \hat{h}_t^j)$, as $\sum_j \theta_j^* = 1$. The expression for the failure probabilities also becomes slightly different than before. The probability of failure of a device at age t in this case is

$$\Pr(\text{failure at age } t) = \sum_{j=1}^{m} \theta_j^* \cdot \left(\lambda_j^* \hat{h}_t^j \cdot \prod_{i=1}^{t-1} (1 - \lambda_j^* \hat{h}_t^j) \right).$$

In this setting, our goal is to estimate $\lambda^* = (\lambda_1^*, \cdot, \lambda_m^*)$ and $\theta^* = (\theta_1^*, \ldots, \theta_m^*)$. With sufficient data, this allows for a greater flexibility in the estimate. From a practical standpoint, this also allows a practitioner to identify if the failure distribution of devices in some cohort recently launched into the market is a more intense or subdued version of the hazard rate distributions in the basis.

After observing failures described by the vectors y and z, let $\theta(y, z)$ be the estimate for θ^* and $\lambda(y, z)$ be our estimate for λ^*. We consider an estimation strategy to be effective if, given a set of samples y and z of observations, we have

$$\theta_j(y,z) \to \theta_j^* \text{ and } \lambda_j(y,z) \to \lambda_j^*, \forall j \text{ almost surely as } \sum_{i=1}^{\bar{t}} y_i \to \infty \text{ and } \sum_{i=1}^{\bar{t}} z_i \to \infty,$$

and also, for any (y, z) we want $0 \leq \lambda_j(x, y) \cdot \hat{h}_i^j \leq 1, \forall i, j$, so that the resulting hazard rates are feasible. Note that there is a straight forward equivalence between the hazard rates and failure probabilities as shown in Equation 2.4.

Although this is a parametric approach, we make no explicit assumptions on the underlying shape of the failure distribution. This is a departure from other models, such as the Cox Proportional Hazards model and models that assume a specific underlying distribution. In our case, the hazard rate distributions in the basis are completely determined by historical data or are defined by the modeler. Although generic, this approach is only adequate when there is an abundance of data and, of course, historical failure distributions or expert opinion available that can be used as priors.

In many applications, such at our partner WSP, estimating a hazard rate distribution as a mixture of the hazard rate distribution of other devices is also useful to identify which manufacturers and/or features lead to large number of failures. For example, this estimation strategy can help identify if devices with similar operating systems have similar hazard functions. Additionally, this strategy can quickly help identify if a recently launched device has an unusually high (or low failure rate).

In the next section, we will describe our estimation strategies and their strengths and weaknesses.

2.4 Estimation Strategies

With the basis formed by the collection of hazard rate distributions $\{\hat{h}^1, \dots, \hat{h}^m\}$ that will be used in our mixture model and the failure observations in hand, we are ready to present the two strategies for estimating the hazard rates of devices in a cohort. Following a similar notation as in Section 2.2.3, assume that the failure age of a cohort of devices has a finite discrete support $[1, T]$. Also, assume that observations are truncated at some age $\bar{t} \leq T$. As before, we observe $y = (y_1, \dots, y_{\bar{t}})$, where y_i represents the number of device failures at age i, and we observe $z = (z_1, \dots, z_{\bar{t}})$, where z_i is the number of observations of age i yet to fail. Furthermore, let $r_i = \sum_{j\geq i} y_j + z_j$ be the number of devices of age i that are at risk, i.e., devices that were observed to have a failure age greater or equal to i.

The first strategy we present is an Expectation-Maximization (EM) algorithm tailored to this problem, where an estimate of the parameters θ^* and λ^* is obtained by maximizing a lower bound of the likelihood function for a set of observations (y, z). The EM algorithm is a classic approach in statistical estimation, and a survey of this method can be found in Dempster et al. (1977).

The second approach is a regularized regression approach, where we estimate the scaling parameters by solving an optimization problem. This approach has the advantage that the estimate of $\theta_j^*\lambda_j^*$ will converge as the number of samples goes to infinity.

2.4.1 Estimation using an Expectation-Maximization Approach

As a first step towards developing the Expectation-Maximization approach for estimating the failure rates, we will derive the log-likelihood function for a set of observations y, z. Following Equation 2.2, and recalling that

$$\Pr(\text{failure at age } t) = \sum_{j=1}^{m} \theta_j^* \cdot \left(\lambda_j^* \hat{h}_t^j \cdot \prod_{i=1}^{t-1} (1 - \lambda_j^* \hat{h}_t^j) \right),$$

we can write the likelihood of some sample y, z for mixing probabilities $\theta = (\theta_1, \ldots, \theta_m)$ and scaling parameters $\lambda = (\lambda_1, \ldots, \lambda_m)$ as

$$\begin{aligned} L(\theta, \lambda; y, z) &= \Pr(y, z | \theta, \lambda), \\ &= \prod_{i=1}^{\bar{t}} \left(\sum_{j=1}^{m} \theta_j^* \cdot \lambda_j \hat{h}_t^j \cdot \prod_{k=1}^{i-1} (1 - \lambda_j \hat{h}_k^j) \right)^{y_i} \cdot \left(\sum_{j=1}^{m} \theta_j^* \cdot \prod_{k=1}^{i} (1 - \lambda_j \hat{h}_k^j) \right)^{z_i}. \end{aligned}$$

The log-likelihood then becomes

$$\log(L(\theta, \lambda; y, z)) = \sum_{i=1}^{\bar{t}} y_i \log \left(\sum_{j=1}^{m} \theta_j^* \cdot \lambda_j \hat{h}_t^j \cdot \prod_{k=1}^{i-1} (1 - \lambda_j \hat{h}_k^j) \right) + z_i \log \left(\sum_{j=1}^{m} \theta_j^* \cdot \prod_{k=1}^{i} (1 - \lambda_j \hat{h}_k^j) \right). \tag{2.5}$$

This expression is not necessarily concave and is analytically intractable. There are no guarantees that there are unique vectors θ and λ that maximize this expression, nor that optimal solutions can be described in closed form. In fact, if m is large, i.e., there is a large number of hazard rate distributions in the basis, there may even be a subspace of parameters that maximizes the log-likelihood. In this case, we say that the model is non-identifiable. Also, note that even if the parameters θ were known, maximizing the log-likelihood with respect to λ is still challenging since the expression is not concave.

On the other hand, instead of maximizing the likelihood in Equation 2.5 directly, assume that we also had access to observations of two other sets of variables: $\{\alpha_{i,j}\}$ and $\{\beta_{i,j}\}$, where $\{\alpha_{i,j}\}$ corresponds to the fraction of uncensored failures observed for an age i that were generated by basis element $\hat{h}^j$, while $\{\beta_{i,j}\}$ corresponds to the fraction of censored observations for an age i that were generated by basis hazard rates $\hat{h}^j$. Note that $\sum_{j=1}^{m} \alpha_{i,j} = 1$ and $\sum_{j=1}^{m} \beta_{i,j} = 1$. In this case, the likelihood of y, z, $\{\alpha_{i,j}\}$ and $\{\beta_{i,j}\}$ in terms of λ and θ is denoted by $L(\{\alpha_{i,j}, \beta_{i,j}, y, z|\theta, \lambda\})$ and is

$$L(\theta, \lambda; \{\alpha_{i,j}, \beta_{i,j}\}, y, z) = \prod_{i=1}^{\bar{t}} \prod_{j=1}^{m} \left(\theta_j \left(\lambda_j \hat{h}_t^j \cdot \prod_{k=1}^{i-1} (1 - \lambda_j \hat{h}_k^j) \right) \right)^{\alpha_{i,j} y_i} \cdot \left(\theta_j \prod_{k=1}^{i} (1 - \lambda_j \hat{h}_k^j) \right)^{\beta_{i,j} z_i}$$

where the first term inside the product is the probability of $\alpha_{i,j} y_i$ non-censored observation from basis j and the second term is the probability of observing $\beta_{i,j} z_i$ censored observations from basis j.

Since we cannot observe $\alpha_{i,j}$ and $\beta_{i,j}$ we formulate the estimation problem as maximizing the expectation of the likelihood above. Namely, we can estimate θ and λ by solving

$$\max_{\lambda, \theta} E \left[L(\theta, \lambda; \{\alpha_{i,j}, \beta_{i,j}\}, y, z) | \theta, \lambda \right].$$

where the expectation is taken over $\{\alpha_{i,j}, \beta_{i,j}\}$. Unfortunately, this expression is still intractable. In order to obtain a tractable formulation note that, from Jensen's inequality,

$$\log \left(E \left[L(\theta, \lambda; \{\alpha_{i,j}, \beta_{i,j}\}, y, z) | \theta, \lambda \right] \right) \geq E \left[\log \left(L(\theta, \lambda; \{\alpha_{i,j}, \beta_{i,j}\}, y, z) \right) | \theta, \lambda \right],$$

and we can attempt to maximize the right-hand-side of the expression above, obtaining a lower bound for the expected likelihood $E \left[L(\theta, \lambda; \{\alpha_{i,j}, \beta_{i,j}\}, y, z) | \theta, \lambda \right]$. The right hand side of the expression above is

$$E[\log(L(\theta,\lambda;\{\alpha_{i,j},\beta_{i,j}\},y,z))|\theta,\lambda] = E[\bar{L}_\theta(\theta;\{\alpha_{i,j},\beta_{i,j}\},y,z)|\theta,\lambda]+E[\bar{L}_\lambda(\lambda;\{\alpha_{i,j},\beta_{i,j}\},y,z)|\theta,\lambda] \tag{2.6}$$

where we define $\bar{y}_{i,j} = E[\alpha_{i,j}y_i|\theta,\lambda]$, $\bar{z}_{i,j} = E[\beta_{i,j}z_i|\theta,\lambda]$, and where $\bar{L}_\theta$ and $\bar{L}_\lambda$ are

$$E[\bar{L}_\theta(\theta;\{\alpha_{i,j},\beta_{i,j}\},y,z)|\theta,\lambda] = \sum_{i=1}^{\bar{t}}\sum_{j=1}^{m}(\bar{y}_{i,j}+\bar{z}_{i,j})\log(\theta_j)$$

and

$$\begin{aligned} E[\bar{L}_\lambda(\theta;\{\alpha_{i,j},\beta_{i,j}\},y,z)|\theta,\lambda] &= \sum_{i=1}^{\bar{t}}\sum_{j=1}^{m}\bar{y}_{i,j}\log\left(\lambda_j\hat{h}_t^j\cdot\prod_{k=1}^{i-1}(1-\lambda_j\hat{h}_k^j)\right)+\bar{z}_{i,j}\log\left(\prod_{k=1}^{i}(1-\lambda_j\hat{h}_k^j)\right) \\ &= \sum_{i=1}^{\bar{t}}\sum_{j=1}^{m}\bar{y}_{i,j}\log\left(\lambda_j h_i^j\right)+\left(\sum_{k=i}^{\bar{t}}\bar{y}k,j+\bar{z}_{k,j}-\bar{y}_{i,j}\right)\log\left(1-\lambda_j h_i^j\right). \end{aligned}$$

If the expected values $\{\bar{y}_{i,j},\bar{z}_{i,j}\}$ were fixed, maximizing Equation 2.6 would be straightforward, since it is separable in θ and λ and is also convex in both arguments. In addition, given θ and λ, we have that

$$\bar{y}_{i,j} = E[\alpha_{i,j}y_i|\theta,\lambda] = y_i\cdot\Pr(\text{haz. rate is }\hat{h}^j)|\lambda,\theta,\text{failure age } i,\text{ survived}\geq i).$$

Using Bayes rule, we know that the probability that the basis is $\hat{h}^j$ given a device that survived until the beginning of age i and fails with age i can be written as

$$\Pr(\text{haz. rate is }\hat{h}^j)|\lambda,\theta,\text{failure age } i,\text{ survived}\geq i) = \frac{\lambda_j\hat{h}_i^j\theta_j}{\sum_{s=1}^{m}\lambda_j\hat{h}_i^j\theta_j}.$$

Similarly, we have

$$\bar{z}_{i,j} = E[\beta_{i,j}z_i|\theta,\lambda] = z_i\cdot\Pr(\text{haz. rate is }\hat{h}^j)|\lambda,\theta,\text{failure age } i,\text{ survived}\geq i).$$

and

$$\Pr(\text{haz. rate is }\hat{h}^j)|\lambda,\theta,\text{fail} > i,\text{ survived}\geq i) = \frac{(1-\lambda_j\hat{h}_i^j)\theta_j}{\sum_{s=1}^{m}(1-\lambda_j\hat{h}_i^j)\theta_j}.$$

Thus, given the expected fractions $\{\bar{y}_{i,j},\bar{z}_{i,j}\}$ it is straightforward to maximize Equation

2.6 and, given θ and λ, it is easy to calculate $\{\bar{y}_{i,j}, \bar{z}_{i,j}\}$. This suggests a procedure where we alternate between maximizing Equation 2.6 and calculating the expectations $\{\bar{y}_{i,j}, \bar{z}_{i,j}\}$. This is the basis of the Expectation Maximization (EM) algorithm. This procedure is outlined below:

1. Given y and z, define some initial values for θ^{old} and λ^{old} and stopping criteria ϵ_θ and ϵ_λ;

2. (Expectation Step) Calculate $\bar{y}_{i,j}$, $\bar{z}_{i,j}$ as

$$\bar{y}_{i,j} = y_i \frac{\lambda_j^{old} \hat{h}_i^j \theta_j^{old}}{\sum_{s=1}^m \lambda_s^{old} \hat{h}_i^s \theta_s^{old}}, \quad \bar{z}_{i,j} = z_i \frac{(1 - \lambda_j^{old} \hat{h}_i^j)\theta_j^{old}}{\sum_{s=1}^m (1 - \lambda_s^{old} \hat{h}_i^s)\theta_s^{old}}$$

3. (Maximization Step) Define $r_{i,j} = \sum_{k=i}^{\bar{t}} \bar{y}_{k,j} + \bar{z}_{k,j}$. Also, let $\hat{h}^j_{\min} = \min_i\{\hat{h}_i^j\}$. Calculate λ_j^{new} for every j by solving

$$\lambda_j^{new} = \arg \max_{\lambda_j \le 1/\hat{h}^j_{\min}} \sum_{i=1}^{\bar{t}} \bar{y}_{i,j} \log\left(\lambda_j h_i^j\right) + (r_{i,j} - \bar{y}_{i,j}) \log\left(1 - \lambda_j h_i^j\right).$$

   Also, calculate θ^{new} defined as

$$\theta_j^{new} = \frac{\sum_{i=1}^{\bar{t}} \bar{y}_{i,j} + \bar{z}_{i,j}}{\sum_{i=1}^{\bar{t}} y_i + z_j}.$$

4. If $||\theta^{new} - \theta^{old}|| \le \epsilon_\theta$ and $||\lambda^{new} - \lambda^{old}|| \le \epsilon_\lambda$ stop and return the current estimates of θ^{new} and λ^{new}. If not, make $\theta^{old} = \theta^{new}$ and $\lambda^{old} = \lambda^{new}$ and repeat from Step 2.

In the procedure above, note that θ^{new} maximizes Equation 2.6. Also, we restrict the maximum value of λ_j^{new} in order to ensure feasibility of the resulting scaled hazard rate distribution. The convergence of the previous algorithm follows directly from the convergence of EM algorithms Wu (1983).

This proposed estimation method has a few limitations. First, even though the procedure itself converges, it does not necessarily guarantee convergence to the maximum-likelihood estimate of θ and λ. Furthermore, it does not directly address potential identifiability problems in this model. Finally, the estimates for θ and λ may be dense, i.e., with most elements of both vectors being non-zero, which can make identifying the most important elements in the basis difficult.

2.4.2 Hazard Rate Regression

In this subsection, we will present and discuss an alternative approach for estimating the parameters of the mixture model. Namely, we will calculate the empirical quantiles, use a model selection procedure to identify relevant basis and, finally, use a regression to calculate the weights in our model. From Section 2.2.3, we have that the non-parametric estimate of the hazard rates for non-censored observations y and censored observations z is

$$h_i^{KM} = \frac{y_i}{r_i},$$

where $r_i = \sum_{j \geq i} y_j + z_j$. We denote $h^{KM} = (h_1^{KM}, \ldots, h_{\bar{t}}^{KM})$. Also, from our probabilistic model, we have that the hazard rate at some age i for a device is given by a scaled combination of basis hazard rates $\{\hat{h}^1, \ldots, \hat{h}^m\}$, where each vector in this set has dimension T. More specifically, each basis j is scaled by some factor λ_j and each device in the cohort will fail according to basis j with probability θ_j. In this case, the probability of failure of a device that survived up until the beginning of age i fails with age i will be $\sum_{j=1}^m \theta_j \lambda_j \hat{h}_i^j$.

A first approach to estimate θ and λ would be to find parameters that minimizes the distance

$$\left\| \sum_{j=1}^m \theta_j \lambda_j \hat{h}^j - h^{KM} \right\|_{p,\bar{t}}, \tag{2.7}$$

where $\|\cdot\|_{p,\bar{t}}$ is the p-norm of the first $\bar{t}$ components of the vector. Namely, for some vector x,

$$\|x\|_{p,\bar{t}} = \left(\sum_{i=1}^{\bar{t}} |x_i|^p \right)^{1/p}.$$

Equation 2.7 is not convex due to the bilinear term. However, since in practice the goal is often to identify relevant bases and then use the weighted model for forecasting, it might not be necessary to learn θ and λ individually. If this is the case, we can define $w = (w_1, \ldots, w_m)$, where $w_j = \theta_j \cdot \lambda_j$. If we also consider feasibility constraints in order to ensure the resulting estimated hazard vector is feasible, we can write the regression problem as

$$\begin{aligned} \text{minimize} \quad & \left\| \sum_{j=1}^m w_j \hat{h}^j - h^{KM} \right\|_{p,\bar{t}} \\ \text{s.t.} \quad & \sum_{j=1}^m w_j \hat{h}_i^j \leq 1, i = 1, \ldots, T, \\ & w_j \geq 0, j = 1, \ldots, m. \end{aligned} \tag{2.8}$$

Simply put, this approach finds a point in the cone generated by the basis hazard rates that is a feasible hazard rate distribution and that minimizes the distance between the non-parametric estimates of the hazard rates. Furthermore, if $\bar{t} = T$ as the number of samples goes to infinity, $\sum_{j=1}^{m} \lambda_j^* \theta_j^* \hat{h}^j$ will converge to h^{KM}. Since we assume that the original samples were drawn from the mixture model in Section 2.3, we have that the solution of the above, denoted by w^*, will converge as well, such that $\sum_{j=1}^{m} w_j^* \hat{h}^j$ converges to $\sum_{j=1}^{m} \lambda_j^* \theta_j^* \hat{h}^j$.

One major problem of this approach is that if $\bar{t}$ is small and m is large, the problem could have multiple optimal solutions. In fact, there might be a subspace of solutions that minimizes the norm, indicating that this model is unidentifiable. Furthermore, even if we obtain a solution, the vector w might mostly have non-zero components, making it difficult to identify the truly important basis in our estimation problem.

We can address the identifiability issue by regularizing the regression. This is done by adding a term to the objective function, ensuring that it is strictly concave. Furthermore, we can try to obtain a sparse solution w by penalizing or limiting the cardinality of this vector. Obtaining a sparse solution can be useful depending on the context of the estimation problem. For example, a sparse solution can help identify which cohorts are the most similar and can help with an investigation of the features that these cohorts have in common. Ideally, we would have a constraint of the type $\text{card}(x) \leq c$, where c is some number chosen by the practitioner. However, adding this term is impractical since it destroys the convexity of the problem.

One alternative that addresses both issues is using the $l1$ norm as a penalty function[2]. The $l1$ norm is the convex envelope of the cardinality function and, in practice, tends to lead to sparse solutions in optimization problems as discussed in Tibshirani (1996). This approach is commonly referred to as the LASSO regression. Thus, the penalty function then becomes

$$\begin{aligned} \text{minimize} \quad & \left\| \textstyle\sum_{j=1}^{m} w_j \hat{h}^j - h^{KM} \right\|_{p,\bar{t}} + \gamma \|w\|_1 \\ \text{s.t.} \quad & \sum_{j=1}^{m} w_j \hat{h}_i^j \leq 1, i = 1, \ldots, T, \\ & w_j \geq 0, j = 1, \ldots, m. \end{aligned} \tag{2.9}$$

[2]The $l1$ norm of some vector x is simply $\sum_i |x_i|$.

where γ is some positive weight that is an input to the optimization. The larger the weight, the larger the penalty on the norm of w. Note that the problem is now strictly convex and has a unique optimal solution.

Finally, the optimization in Equation 2.9 is useful for *model selection*, i.e., for identifying which of the basis hazard rates are the most relevant. Thus, the optimization is done through the following steps:

1. Define a tolerance $\epsilon \geq 0$ and let d^* be the optimal cost of the problem in Equation 2.8.

2. For some $\gamma > 0$ we solve the optimization problem and select the relevant basis, which will usually be the basis elements for which the corresponding weight is non-zero.

3. With the relevant basis elements in hand, we discard the non-used basis and resolve the problem in Equation 2.8 with $\gamma = 0$ using only the new set of basis. The new cost is denoted by d^r.

4. If $d^r \leq (1+\epsilon)d^*$, stop. If not, decrease γ and return to Step 2.

The solution of the second step will be taken as our estimate of $\theta_j^* \lambda_j^*$. In the next section we will analyze through numerical experiments both of these approaches.

2.5 Forecasting and Numerical Experiments

In this section we will examine the performance of the EM algorithm and of the hazard rate regression procedure proposed in the previous section through three sets of experiments. The first set of experiments uses data from our partner WSP and the goal is to estimate the amount of weekly customer warranty claims it receives for two different devices. The second set of experiments corresponds to an artificial set-up where we compare, in a controlled setting, the performance of the two estimation strategies. The final set of experiments is based on data from Project Repat, a social enterprise in the Boston area that transforms old t-shirts into quilts.

Before presenting the experiments, we will first discuss, using the context of our partner WSP, how the estimate of the hazard rate distribution was used to forecast the number of device failures and, therefore, the number of warranty claims.

2.5.1 Forecasting

As discussed in Section 2.3, if the maximum warranty of a device is T periods in length, and if we are $\bar{t}$ periods after the launch of the device, the oldest device in the market will be of age $\bar{t}$ and there will be no failure observations in the interval $[\bar{t}, T]$. The estimate of the hazard rate distribution using information up until time $\bar{t}$ will be denoted by $h^{\bar{t}} = (h_1^{\bar{t}}, \ldots, h_T^{\bar{t}})$.

Furthermore, assume that sales of devices happen in some interval $[1, T_s]$, such that there are no sales after time T_s. Let the sales in each period be described by a vector $s = (s_1, \ldots, s_{T_s})$. For our experiments, we assume that this vector is known. This is not true in practice, but the estimation of sales volume is beyond the scope of this chapter.

In this setup, we expect the last period in which failures occur to be $T + T_s$. Hence, we define $T_{max} = T + T_s$ to be the length of time that the WSP has to manage warranty claims for a device. Also, let the age distribution of devices in a cohort be given by $x(t) = (x_1(t), \ldots, x_T(t))$ where $x_i(t)$ represents the amount of devices of age i in the cohort of devices at the beginning of period t. Then, we have that $x_1(t) = s_t, \forall t$ and that

$$E[x_{i+1}(t+1)|h^{\bar{t}}] = (1 - h_i^{\bar{t}}) \cdot x_i(t), \forall i = 2, \ldots, T; \forall t = 1, \ldots, T_{max-1}.$$

Also, in the first period, $x(1) = (s_1, 0, \ldots, 0)$. Furthermore, let the number of failures at time t be given by a vector $f(t) = (f_1(t), \ldots, f_T(t))$, where $f_i(t)$ is the number of failures of age i at time t. Then, given an estimate $h^{\bar{t}}$ of the hazard rate distribution, we have

$$E[f_i(t)|h^{\bar{t}}] = h_i^{\bar{t}} \cdot x_i(t), \forall i = 1, \ldots, T; \forall t = 1, \ldots, T_{max-1}, \tag{2.10}$$

and the expected number of failures at time t will be $\sum_i h_i^{\bar{t}} \cdot x_i(t)$. In our experiments, we take this expectation as the estimate of the failures.

One metric for the performance of an estimation strategy is to compare the estimates of $(f(1), \ldots, f(T_{max}))$ to the true failures. A second metric that we use is comparing the maximum distance between the estimated Cumulative Distribution Function (CDF) and the true CDF. We define estimated CDF with information up until time $\bar{t}$ by $F^{\bar{t}} = (F_1^{\bar{t}}, \ldots, F_T^{\bar{t}})$, where $F_i^{\bar{t}}$ is

$$F_i^{\bar{t}} = 1 - \prod_{k=1}^{i} (1 - h_k^{\bar{t}}),$$

and if the true CDF is $F^* = (F_1^*, \ldots, F_T^*)$, we have that the maximum distance is

$$\max_i |F_i^{\bar{t}} - F_i^*|.$$

We call this distance the *Kolmogorov-Smirnov (KS) distance*, since it has the same form of the Kolmogorov-Smirnov statistic for an empirical distribution. From a practical standpoint, if all devices were sold on a single day, this distance would represent the maximum error of the estimate of the cumulative number of failures of a device.

2.5.2 Forecasting Failures at the WSP

We consider the problem of estimating the weekly number of failures for two models of devices sold by our partner WSP. The first device, which we call device A, was a device marketed towards business customers and was one of the best-selling devices in this segment. The second device, which we call device E, was a device marketed towards regular customers and was one of the best-selling models in its launch year. Device E was also a device that had an unusually high failure rate, higher than other models in the segment, and which led to a stock-out of replacement devices at our partner WSP.

We take the hazard rate distributions estimated from 5 other devices as the set of basis hazard rate distributions in our estimation. We have access to all failures of these devices so that the hazard rate distributions of these devices were estimated using the KM estimator. The 5 devices used as basis were made by 4 different manufacturers and had different operating systems and physical characteristics than devices A and E. Since the number of elements in the basis is small, we did not include a regularization term in the regression estimation procedure. For the EM algorithm, the stopping criterium of the algorithm was that the parameters did not change by more than 1%.

Our experiments consist of estimating the weekly number of failures for both of these devices, considering that there are different number of weeks of information available, i.e., different values of $\bar{t}$. We take T as 100 weeks and $T_{\max}$ as 150 weeks, since more than 95% of sales happen during the first 50 weeks from launch. Recall that if we are $\bar{t}$ weeks after the launch of the device, the oldest device in the market is of age at most $\bar{t}$ and we have no failure age observations in the range $[\bar{t}, T]$. Thus, in the case, the information in the interval $[1, \bar{t}]$ is used to estimate $h^{\bar{t}}$ and, by assuming that the weekly sales of this device

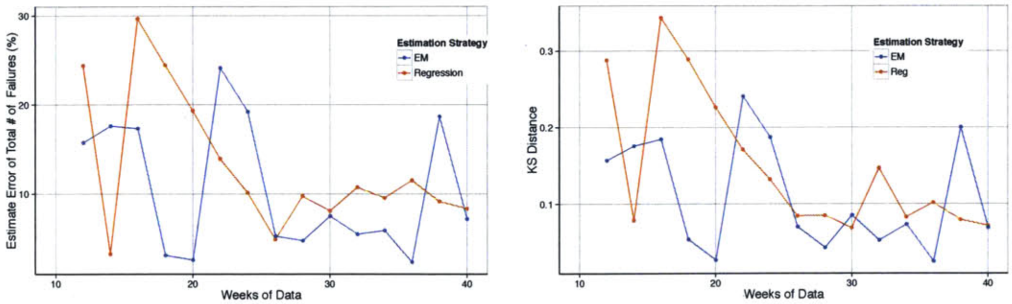

(a) Est. error of total # of failures.

(b) KS Distance

Figure 2-3: Estimate error of the total number of failures after 150 weeks as a % of the real total failures and the KS distance for different amounts of data for device A.

is known, we create a forecast for the expected weekly failures from week 1 to week T_{max} using Equation 2.10.

The results for the EM algorithm and for the hazard rate regression for device A are depicted in Figure 2-3 and in Figures 2-8 and 2-9 in the Appendix of this chapter. Figure 2-3a shows the absolute error in the forecast of the total number of failures of device A for $\bar{t} \in [1, 40]$. Thus, if we are $\bar{t}$ weeks after launch, we have $\bar{t}$ weeks of information, and we use this information to estimate the total volume of failures. In Figure 2-3b we depict the KS distance between the estimate of the CDF and the true CDF of the device for different amounts of information available. After receiving 25 weeks of failure data, both estimators produce good results, having a maximum error around 10% and a KS distance of 0.1. Note that 25 weeks is still at the beginning of the warranty life-cycle of the device, which is usually around 100 weeks long. The performance of both estimation strategies for different amounts of information can be visually compared in Figures 2-8 and 2-9 in the appendix of this chapter.

The performance of the estimation strategies for device E, which had a higher than expected failure rate, is depicted in Figures 2-4, 2-10 and 2-11. Once again we assume that the device is launched on week 1 and we make the estimate assuming we are $\bar{t}$ weeks after launch and only have information up until week $\bar{t}$. In this case, the EM algorithm failed to produce good forecast of the failures. This happened for a few reasons. First, the hazard rate of device E was unusually larger than the hazard rates of other devices in the basis. Furthermore, the EM algorithm does a local search, and does not guarantee a quality estimate of the mixture parameters. On the other hand, the hazard rate regression

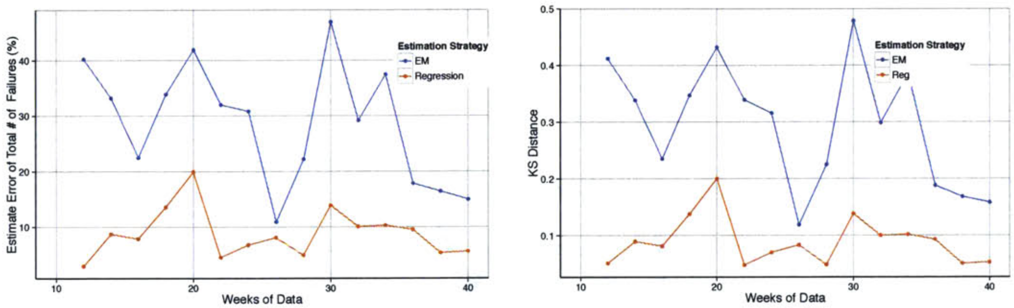

(a) Est. error of total # of failures.

(b) KS Distance

Figure 2-4: Estimate Error of the total number of failures after 150 weeks as a % of the real total failures and the KS distance for different amounts of data for device E.

estimation strategy produced estimates with KS distances that were consistently around 0.1, with only 22 weeks of data. In addition, with more than 22 weeks of data, the estimate of the total number of failures is within 10% of the true total number of failures. This indicates that it would have been indeed possible to detect early on the unusual failure rate of this device, allowing for tactical decisions such as better management of inventory of refurbished devices. Figures 2-10 and 2-11 illustrate the performance of the two estimation strategies.

A tailored version of the estimation strategy was implemented at our partner WSP, and the implementation is described in Petersen (2013). In addition, a plug-in was developed for Microsoft Outlook and Excel that allowed managers at the reverse logistics facility to forecast the amount of failures and also provided an estimate of inventory needs.

2.5.3 Simulation using computer generated data

The second set of experiments consists of comparing the EM algorithm and the hazard rate regression in a controlled setting. More specifically, we assume that $T = 200$ and that all members of a cohort enter the market at the same time. We consider a basis set of 30 hazard rate distributions that will be used to estimate the hazard rate of a target new cohort. The failure distributions of cohorts in the basis and the (initially unknown) failure distribution of the new cohort are generated as follows:

1. Sample a and b from a Uniform distribution with parameters $[1, T]$. Furthermore, sample some value p from a Uniform distribution with parameters $[0, 1]$;

2. The failure age distribution of the cohort is set to be a mixture of a Uniform random variable with parameters $[1, a]$ and an Exponential random variable with mean b. The mixture probability is p.

The mixture of a uniform and exponential random variables are chosen to represents two populations inside a cohort. It also generates a non-trivial failure distribution for each cohort. We also assume that failure observations are censored, and that the censoring random variable is uniform with parameter $[0, T]$. Thus, for each failure sample we generate a corresponding censored variable such that, if the sampled failure age is x and the sampled censoring variable is y, we observe $\min(x, y)$. We also assume that we know if an observation is censored or not.

We assume that we do not have direct access to the hazard rate distributions in the basis, and that the hazard rate distributions in the basis are estimated using the KM estimator for 100 (potentially censored) samples of the corresponding cohort's failure distribution. For the new cohort, the target of the estimation problem, we also assume that there are 200 (potentially censored) failure observations.

We evaluate the performance of the two estimation strategies for different values of $\bar{t}$. Thus, if a sampled failure age for a member of the new cohort is x and the censoring variable is y, we will observe $\min(x, y, \bar{t})$. For each value of $\bar{t}$, we generate 100 different bases and 100 new cohorts using the sampling strategy described in the beginning of this subsection. We then estimate the hazard rate distribution of the new cohorts and corresponding value of $\bar{t}$ using both the EM algorithm and the regression approach. For the EM approach the algorithm was interrupted if the parameters did not change by more than 1%. For the regression approach we included a regularization component, and we set the tolerance to be 5% of the non-regularized cost.

The results are summarized in Figure 2-5. In the box plot, the black line in the middle of the box is the median of the sample, the top line is the 75% quantile and the bottom line is the 25% quantile. The lines extend to 1.5 times the inter-quartile range. Note that that there is no improvement for $\bar{t}$ larger than 10 (10% of the horizon), and both approaches produce failure distribution estimates with a KS distance lower than 0.2. Thus, at 10 periods after launch it would be possible to have a reasonable estimate of the failure distribution. Note that the KS distance of the regression approach has a much lower variance than the EM algorithm. This is because of the sensitivity of the EM algorithm to the initial parameters

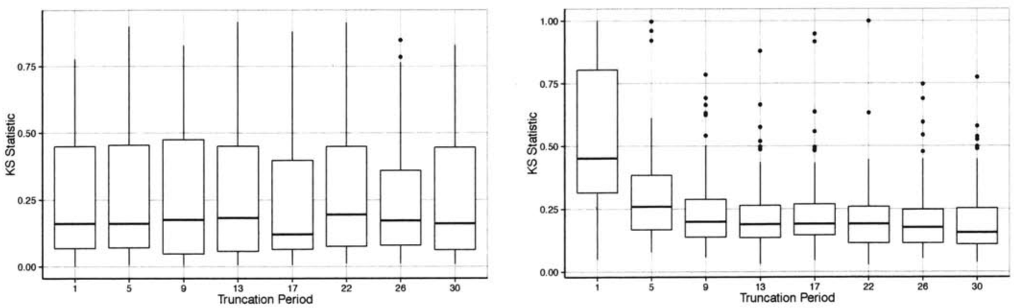

(a) EM Algorithm (b) Hazard Rate Regression

Figure 2-5: Box-plot of the KS distance between the true hazard rate and the estimated hazard rate for different values of $\bar{t}$. The distance is calculated for the horizon [1,100].

and the fact that it does not necessarily converge to the maximum-likelihood estimate of the model parameters.

The regression approach, which included regularization with a 5% tolerance, also led to sparser representations. None of the observations in the simulation used more than 22 elements of the basis and, the average number of basis selected was 7. On the other hand, in the EM approach, in the majority of the simulations all elements of the basis had non-zero weights.

2.5.4 Estimating returns at Project Repat

The third set of numerical experiments uses data from Project Repat, a social enterprise in the Boston area that transforms old t-shirts into quilts. Their product is popular among college students and recent graduates that want to preserve their college (or fraternity/sorority) t-shirts, and it is also popular among athletes, particularly runners, who collect t-shirts from races. Depending on the season of the year, Project Repat can sell from hundreds to thousands of quilts per week.

The dynamics of Project Repat's customer-facing operation is as follows: (i) customers "purchase" and pay for a quilt on Project Repat's website; (ii) Project Repat registers the order and send the customer a pre-paid envelope; (iii) The customer puts old t-shirts in the envelope and sends it to Project Repat; (iv) The t-shirts are received, cut, sewn into a quilt; and (v) the quilt is shipped to the customer.

Project Repat puts a high value on the social and environmental impact of their work. Besides being a company that *upcycles* old t-shirts, Project Repat contracts all of its sewing

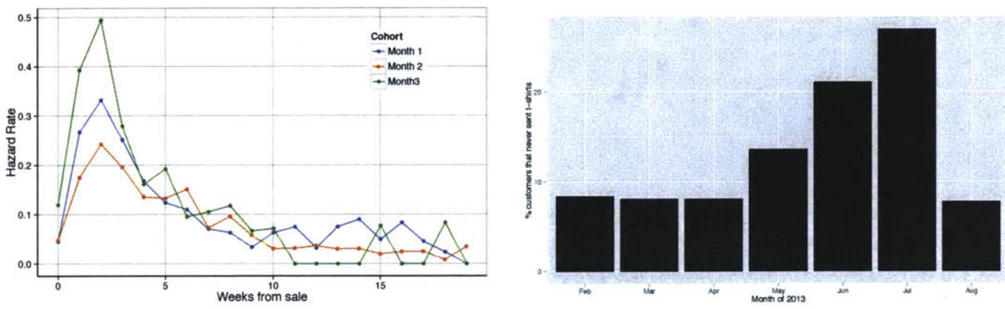

(a) Hazard Rate Distribution

(b) Customers that never return

Figure 2-6: Figure (a) is the hazard rate distribution of the time until the customer sends the t-shirts. Figure (b) is the fraction of customers that never sent their t-shirts.

to textile plants in the United States as an attempt to " repatriate the textile industry". Finally, this company actively works with NGOs that employ individuals with disabilities and that have limited employment opportunities.

A major issue in Project Repat's operations is forecasting the volume of envelopes with t-shirts that they receive from the customers that purchased a quilt on-line. More specifically, they use these forecasts to decide how many working-hours they should contract from textile plants; and if the volume of work needed exceeds the minimum contracted, they have to pay overtime.

In this context, we use the hazard rate regression strategy to estimate the lead time between the customer receiving an envelope from Project Repat, and sending back their old t-shirts. We model this lead-time as a random variable, having a similar interpretation as the failure age. While before we were estimating the failure age of devices, here we are estimating the *customer lead time*, i.e., how many days (or weeks) after purchasing a product do customers send in their t-shirts.

The hazard rate of the amount of time until customers send in their t-shirts is depicted in Figure 2-6a for three sample months. Note that these hazard rates appear to be scaled versions of each other. Also, from the return data, we have that customers take between 2 and 3 weeks on average to send their t-shirts, if they send it in at all. However, the lead-time is heavy tailed, and over 20% of customers that eventually send their t-shirts take more than 5 to weeks send them.

In Figure 2-6b, we have the fraction of customers that never sent their t-shirts. This seems to be the driving factor that makes hazard rates dissimilar. The fraction of orders

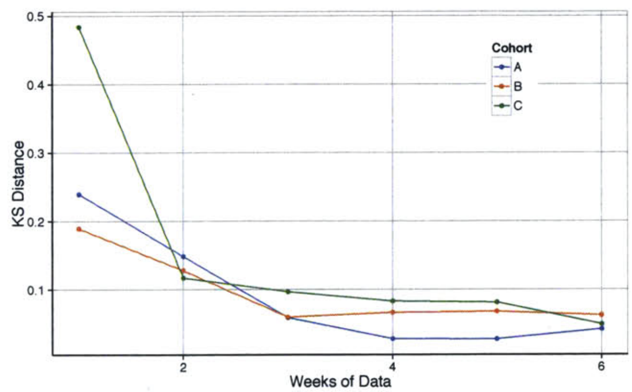

Figure 2-7: KS Distance for estimates from different cohorts for different amounts of information available.

that never send t-shirts are driven by two factors: (i) seasonality effects - quilts purchased as gifts have a lower percentage of returns; (ii) the promotion in place - coupons and discounts attract customers with a lower return rate.

For our estimation, we define a cohort as the customers that purchase devices in a given week. The goal is to estimate the weekly hazard rate of a customer sending in the t-shirts. In our experiment, we use $T = 24$ weeks and use the empirical hazard rate from 35 different weeks as the basis set. We chose 3 cohorts to estimate - each cohort is from a different month, in order to ensure that they are not too similar, and they were all chosen from sales after the cohort weeks in the basis.

We observe the performance of the estimation strategies for different amounts of information available, i.e., different values of $\bar{t}$. Thus, we assume we are $\bar{t}$ weeks after a cohort purchased their quilts, and use the information from the interval $[1, \bar{t}]$ to estimate the lead-time distribution. Note that, for all 3 cohorts, with two weeks of information (out of a horizon of 140 weeks) the KS distance is less than 0.15 and, with three weeks of information, the distance is less than 0.1. Since all customers in a cohort purchase a quilt at the same time, the KS distance is the same as the maximum error of the aggregate estimate of the number of t-shirts sent per week. The estimation procedure also led to sparse representations, and for all cohorts and values of $\bar{t}$ no estimate used more than 7 elements of the basis.

This estimation procedure was built into a cloud-based forecasting tool that was given to Project Repat. Through the tool, the company can forecast the volume of t-shirts received given open pending orders of each cohort in the system. Also, by analyzing the basis selected by the estimation procedure, Project Repat can identify which weeks best represent a new cohort, and use this to try to identify what influences the customer lead-time.

2.6 Conclusion

We proposed and analyzed two methods for estimating failure distributions of newly launched devices that leverages the historical data of failures from other devices. The proposed strategies are based on a *hazard rate model* developed under the assumption that customers in the same *cohort* have devices that fail according to the same age-dependent failure distribution. A cohort is a pool of customers that share similar features (e.g. phone model owned, data plan, etc.).

The first estimation strategy used an Expectation-Maximization (EM) type algorithm to estimate the parameters of a mixture model. Here, we assumed that hazard rates of devices in a new cohort are drawn from a mixture of scaled hazard rate distributions built from historical data. Furthermore, since maximizing the likelihood function of a set of observations is intractable, we use an EM approach to maximize a lower bound of the likelihood and obtain an estimate of the parameters of the mixture model.

The second estimation strategy, called hazard rate regression, uses a model selection method, where we assumed a "basis" set of hazard rate distributions determined from historical data. We then used a regularized regression to identify and fit the relevant hazard rates distributions from the basis to the observed failures from the new cohort. This allows for a sparse representation of the estimated hazard rate distribution, which can be useful depending on the context of the estimation problem.

In the final part of this chapter, we described how these estimation strategies can be used to create a forecast of the volume of warranty requests received by our partner WSP, and introduced different metrics to measure the quality of the forecast. We also examined both estimation strategies through a series of numerical experiments using data from our partner WSP and using data from Project Repat, a Boston-based social enterprise that transforms old t-shirts into quilts. Through these experiments, we observed that both the EM algorithm and the Regression approach have a similar average performance, but the performance of the regression approach has a lower variance. Furthermore, the regression approach leads to sparser representation of the hazard rate distribution, while the estimate produced by the EM algorithm is dense in the mixture parameters.

There are a few open problems that we have yet to examine. First, a more thorough theoretical characterization of the hazard rate regression procedure, and an analysis of its connections with other estimation strategies may lead to a deeper understanding of its advantages and disadvantages. A second problem is the connection between estimation and inventory management. For example, in this setting, it is not clear if the presence of censored information leads to policies that oversell items, or a policy that undersells items. Finally, investigating how the basis in the estimation impacts overall estimation quality can lead to a more precise guidelines for defining cohorts and selecting the basis used for estimation.

2.7 Appendix - Figures

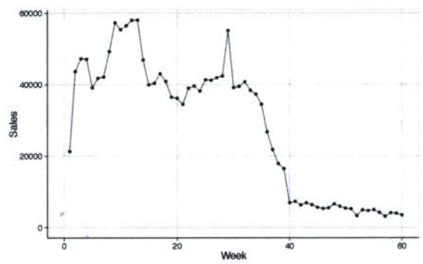

(a) Sales of device A

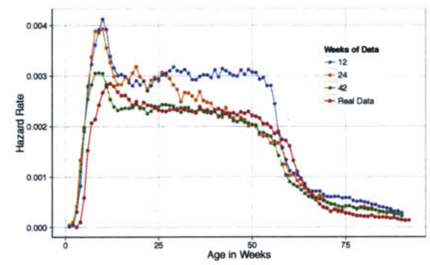

(b) Hazard rate dist. - reg. approach.

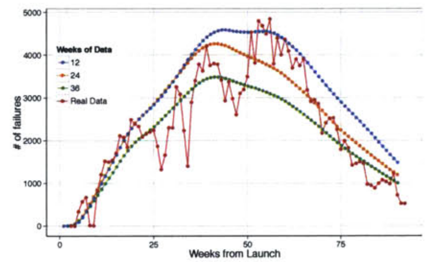

(c) Failure per week - reg. approach

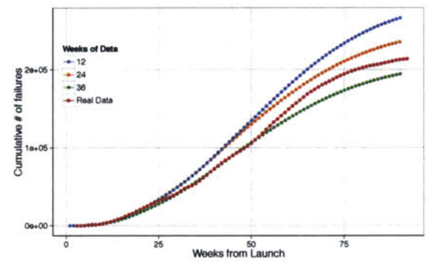

(d) Cumulative failures - reg. approach

Figure 2-8: Simulation results device A using a hazard rate regression approach.

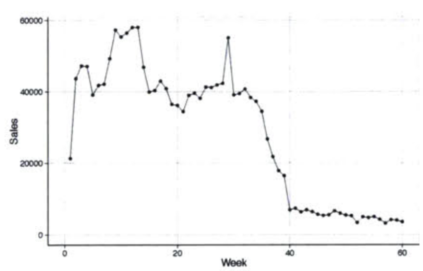

(a) Sales of device A

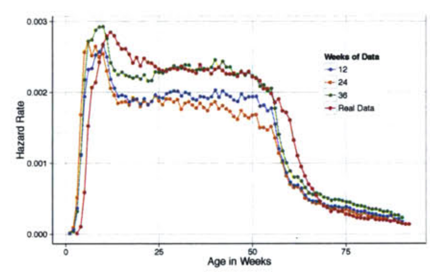

(b) Hazard rate dist. - EM approach.

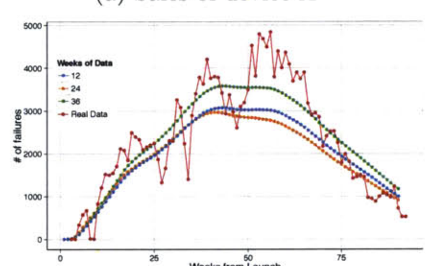

(c) Failure per week - EM approach

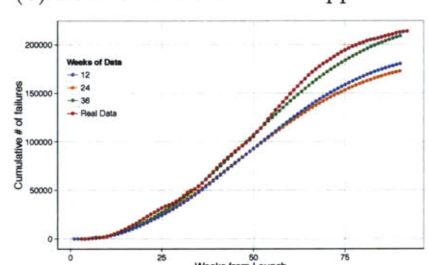

(d) Cumulative failures - EM approach

Figure 2-9: Simulation results device A using EM algorithm.

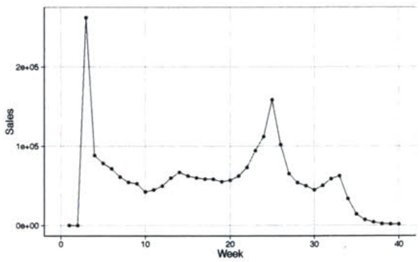

(a) Sales of device E

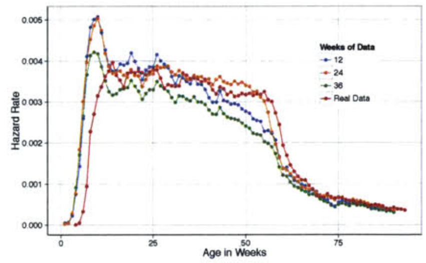

(b) Hazard rate dist. - reg. approach.

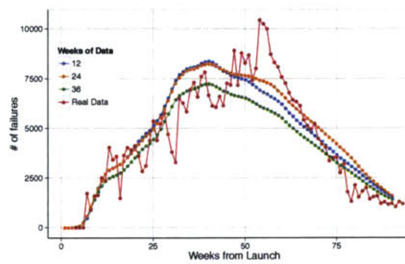

(c) Failure per week - reg. approach

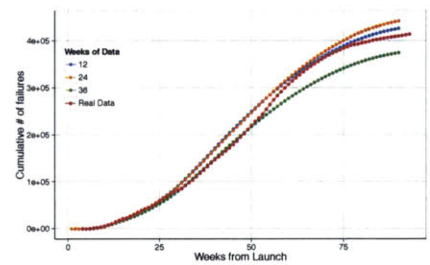

(d) Cumulative failures - reg. approach

Figure 2-10: Simulation results device E using a hazard rate regression approach.

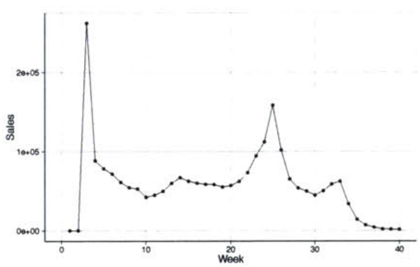

(a) Sales of device E

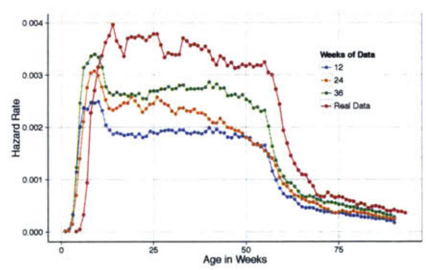

(b) Hazard rate dist. - EM approach.

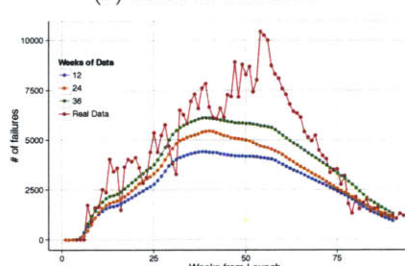

(c) Failure per week - EM approach

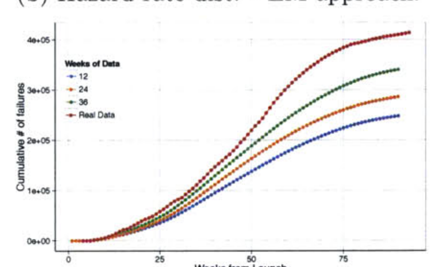

(d) Cumulative failures - EM approach

Figure 2-11: Simulation results device E using the EM algorithm.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Inventory Management in a Closed-Loop Supply Chain

3.1 Introduction

The goal of this chapter is to describe, model, and optimize the inventory at our industrial partner, a Wireless Service Provider (WSP) that is a Fortune 100 company. This WSP sells consumer electronics and offers voice and data plans to its consumers. For this company, as for many other retail businesses, the management of warranty claims and regret returns is a key issue. In fact, the volume of warranty claims for products commercialized by our industrial partner is substantial (in the order of thousands per day), and a significant portion of sold items are returned. Coupled with the short life cycle of their products, usually less than one year, this leads to large levels of inventory of refurbished products that are costly to maintain and expensive to dispose.

Our industrial partner uses a reverse logistics model that is similar to the one adopted by other retailers, especially on-line retailers. In this model, there are two warranty contracts in place: the consumer warranty, and the Original Equipment Manufacturer (OEM) warranty. The consumer warranty protects the consumer against any defects in the purchased product and also provides the consumer a period for "regret returns". In addition, the consumer warranty has strict requirements - when a warranty claim is filed, a new or refurbished item is immediately shipped by our industrial partner to the consumer together with a pre-paid shipping label so that the customer can return his original unit. Thus, a replacement item is sent *before* the original item is received. The OEM warranty, on the other hand, is offered by the OEM to the WSP, covering every device purchased from the OEM. This warranty is slow - a defective product sent to the OEM takes weeks or months to be fixed and a replacement device is not shipped immediately.

Because of the differences between the OEM and consumer warranty contracts, our partner WSP has a reverse logistics facility that is dedicated to processing customer-regret returns and customer warranty claims. This facility can also execute repairs if the defect in the returned product is small or not covered by the OEM warranty. However, if the returned product has a defect that is covered by the OEM warranty, the device is sent to the OEM for repair and refurbishment. Since this reverse logistics facility uses repaired and refurbished devices to satisfy consumer warranty claims, this system fits in the context of closed-loop supply chains (CLSCs).

The remainder of this chapter will address inventory management in the reverse logistics facility. For such, we present a literature review and then proceed to discuss the problem set-up and the different players in the system. We then introduce a discrete-time stochastic model that captures the behavior and the dynamics of this system. We prove the structure of the optimal policy for this problem, starting with a deterministic approximation and proceeding to the stochastic case. We also propose a balancing approximation that allows for a tractable calculation of inventory management policies. This results in a simpler, more tractable newsvendor-type model. In the final part of the chapter the policies we propose will be examined through numerical experiments.

3.2 Literature Review

This chapter builds upon a vast body of literature that addresses closed-loop supply chains (CLSCs) and reverse logistics. Two early examples of papers that consider the management of inventory of repaired/refurbished items are Simpson (1978) and Allen and D'Esopo (1968) that study a finite horizon remanufacturing system. More recently, Guide and Van Wassenhove (2009) and Fleischmann et al. (1997) present an overview of the literature in CLSC and discuss future directions of research. Also, Savaskan et al. (2004) discusses different strategies for remanufacturing products in a CLSC context.

With regard to inventory management for warranty replacements, Huang et al. (2008)

offers a detailed overview of the literature in warranty management, and analyzes inventory management when there is demand for both new and replacement items, and assumes the demand to be stochastic. However, it does not take into account remanufacturing. Khawam et al. (2007) use a simulation approach to obtain inventory management policies for Hitachi. In Toktay et al. (2000), a closed-loop queuing model that captures dynamics of recycling at Kodak is proposed, and a control-policy for the procurement of components is proposed. In Feinberg and Lewis (2005) and in Chen and Simchi-Levi (2009), the authors investigate the case where demand can be positive or negative, being similar to the set-up in this problem, although the authors do not consider time-varying costs.

From a conceptual level, product warranty management is discussed by Murthy and Blischke (1992). Also, the connection between the warranty and logistics literature is discussed by Murthy et al. (2004), as well as the relationship between warranty service and customer satisfaction. The impact of regret returns on inventory management is analyzed in de Brito and van der Laan (2009) where the authors highlight the effect of imperfect information about returns on inventory management. Also, in Geyer et al. (2007), the cost-saving potential of remanufacturing systems is analyzed, although the inventory management problem is not considered.

Our work also draws from the stream of literature that considers inventory management in the presence of incomplete or censored demand information, specifically when ordering quantities or inventory levels impact demand observations. Although this is a setting different than ours, since we consider a reverse logisics system where inventory levels do not influence observations, many in this stream of literature have the same flavor as the results in our work. In Chen and Plambeck (2008), the inventory management of non-perishable goods when there is incomplete demand information is considered, and a policy for managing inventory while learning about demand is studied. The newsvendor problem with censored demand is studied in Godfrey and Powell (2001), and distribution free algorithm for setting the ordering quantity is determined. An asymptotic analysis of inventory planning with censored demand is presented in Huh and Rusmevichientong (2009), and adaptive data-driven inventory control strategies when there is censored information are proposed in Huh et al. (2011).

Finally, our work also fits in the wider field of perishable inventory systems, and a review of the results in this area can be found in Nahmias (2011). As for the theoretical tools that we use in this paper, Bertsekas (2005) presents an overview of Dynamic Programming in the finite and infinite horizon setting, including many examples in inventory management. Balancing policies have also been an active area of study, and Levi et al. (2007) contains an application of this approach in inventory management.

3.3 Problem Description

There are three players 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 players is structured as a traditional supply chain. The customer purchases a mobile device from the WSP and either subscribes to a wireless plan that includes voice and data services provided by the WSP or purchases access to these services through a pre-paid card. The WSP, whom we focus on, is both a retailer of mobile devices and a provider of wireless services to consumers. As mentioned in the introduction, most of the revenue of the WSP comes from the services it provides, and one of their corporate goals is to ensure that the time customers spend "disconnected" from their 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 closed-loop supply chain (CLSC). Before describing the dynamics of this CLSC, we will highlight the two warranty contracts that are in place in this system, namely, the consumer warranty, offered by the Wireless Service Provider (WSP) to the device user, and the Original Equipment Manufacturer (OEM) warranty, offered by the OEM to the WSP.

3.3.1 The Consumer and the OEM warranties

The consumer 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 decides to purchase an extended warranty plan, usually an additional 12 months of coverage. 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 technician is unable to solve the problem, or if a manufacturing defect is verified, 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 so the user can 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 shipped to the consumer is a refurbished device of the same model of the device that failed. More specifically, the replacement device is usually a remanufactured device from some previous warranty claim or regret return. If there are no refurbished devices available, the WSP will send the customer either a new device of the same model or, if no new devices are available, will give the customer an upgrade, and will give the customer a newer model. Giving customers an upgraded device is perceived as a "last resort" for fulfilling warranty claims, since it creates an incentive for customers to file warranty claims as an attempt to obtain a new device for free.

The consumer warranty contract also allows for regret returns, such that the user has a "grace period" of a few weeks after purchasing a new mobile device where the device can be returned to the retail site for a complete refund, net of a stocking fee. These returned devices are shipped from the retail site to the same reverse logistics facility that manages defective devices that originate from warranty claims.

The OEM warranty is designed to protect the WSP from manufacturing defects of devices purchased from the OEM. This warranty usually has a 12 month duration and starts when the WSP purchases the devices, usually in bulk, from the OEM. The warranty asserts that the OEM is responsible for remanufacturing or replacing defective devices that are sold to the WSP. Unlike the consumer warranty, the OEM warranty is not designed to minimize the time that the WSP spends without a working device. However, 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.

The WSP maintains an inventory of refurbished devices in a reverse logistics facility, which is dedicated to processing customer regret returns and faulty devices originated from warranty claims. This facility can also execute repairs if the defect in the returned product 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 a significant number of customers that purchase new devices regret their decision and return their device (usually between 5% and 10%), there are over one million devices in this facility at any given time.

The structure of the consumer and OEM warranties create management issues that ultimately impact the reverse logistics operations of the WSP. An example of this friction 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, both the OEM does not offer a replacement to the WSP, and the WSP might not use this device as a replacement, since it cannot ensure perfect functionality. This generates loss in the system and potential additional costs for the WSP. Another source of friction that can lead to system loss are manufacturing defects that are detected by the WSP and disputed by the OEM. Since there is a large volume of devices going through the reverse logistics facility, there can be a large number of disputed claims, creating additional overhead for the managers at the WSP.

3.3.2 Dynamics of the reverse logistics facility

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 device flow in this reverse supply chain is described below, following the numbers in Figure 1:

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, together with a pre-paid shipping label for the return of the original product held by the customer. Note that the customer receives a replacement *before* returning the original item.

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

3. A second source of products into the facility are regret returns. A customer can return a product within a few weeks of purchase, and receive a full refund, possibly net of a stocking fee. Customers that return an item through this channel do not receive a new product.

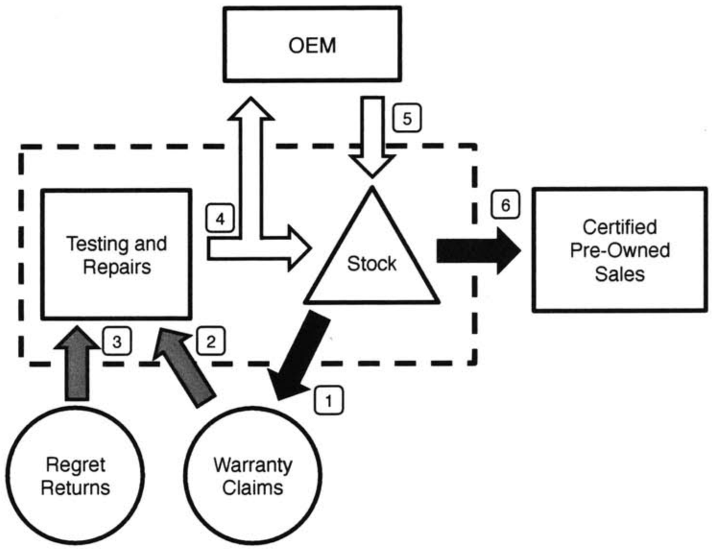

Figure 3-1: Dynamics of the CLSC. The dashed line outlines the reverse logistics facility.

4. Products received from warranty claims and regret returns go through a triage process upon arrival at the facility and different tests are conducted. The triage process generally leads to 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 our industrial partner, and these devices are kept in stock at the reverse logistics facility.

6. The refurbished devices kept in stock 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 products life cycle.

3.3.3 Demand, Costs, and Controls

Demand in this system corresponds to the demand for refurbished devices that will be sent as replacements to users that filed a warranty claim. Refurbished devices enter inventory either from the OEM or directly from the testing and repairs part of the facility. Since this is a closed loop system, there is a correlation between the demand for refurbished devices and the arrivals of refurbished devices in inventory. In addition, the loss in this loop, i.e., the number of devices that cannot be repaired, can be significant and might exceed 20% of all devices that arrive at the facility. Loss can happen for multiple reasons. For example, the consumer warranty might cover a wider range of defects than the OEM warranty or, as mentioned before, the OEM and the WSP do not agree on the nature of a defect.

Due to the consumer 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 these refurbished devices depreciates quickly over time, hence there is a holding cost for units in inventory, corresponding to the opportunity cost of selling inventory through a side channel. There is also a cost associated with not having enough inventory of refurbished products since, in this case, a new device has to be purchased from the OEM and shipped to the consumer or the consumer is given an upgrade. Balancing these two costs in 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 when managing inventory in this system.

Another source of cost in the system are out-of-warranty returns, i.e., defective devices with expired OEM warranty that are returned by consumers who are still covered by the consumer warranty . This issue is particularly acute for devices that have a high failure rate since, in this case, it is not unusual for users to file two or more warranty claims during the period they are covered by the consumer warranty. Defective devices returned with expired

OEM warranty have low salvage value and are costly to repair. Out-of-warranty returns are also accentuated due to the fact that the OEM warranty of a device starts when the WSP purchases it and is not interrupted while it is being remanufactured or when it is in stock in the reverse logistics facility. Thus, devices "age" in inventory, and if an "old" device in stock is sent to a "young" customer, i.e., a customer that is still in the beginning of its consumer warranty contract, this device might fail again while the user is still covered by the consumer warranty but the device has an expired OEM warranty.

3.4 Inventory Model

From a tactical perspective, there are two main decisions involved in managing this system: (i) the decision of how many devices should be kept in the inventory located 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 will uncouple these problems and address them separately. Although, uncoupling these two decisions implies that the policies discussed in this chapter are not optimal in a system-wide sense, it will lead to policies for which we can glean both theoretical and managerial insights. Furthermore, the numerical experiments in the next chapter indicate that making these decisions independently does not lead to a significant additional cost. In this section, we will analyze two versions of the inventory management problem. First, we will discuss a deterministic version, for which hazard rates are fixed fractions of sold devices that fail. We will prove the optimal policy for this case and also present a worst case analysis. The second version is the stochastic problem, for which we prove the structure of the optimal policy and discuss a heuristic for managing this system.

3.4.1 Inventory Dynamics and Costs

The dynamics of the inventory at the reverse logistics facility is depicted schematically in Figure 1. We assume that time is discrete and that, at each time instant t, the demand for refurbished devices, i.e., the number of warranty claims, is described by $d(t)$. The number of refurbished devices that arrive to stock is given by a process $a(t)$, and the inventory level at the beginning of the period is $x(t)$. The number of devices purchased at time t is denoted

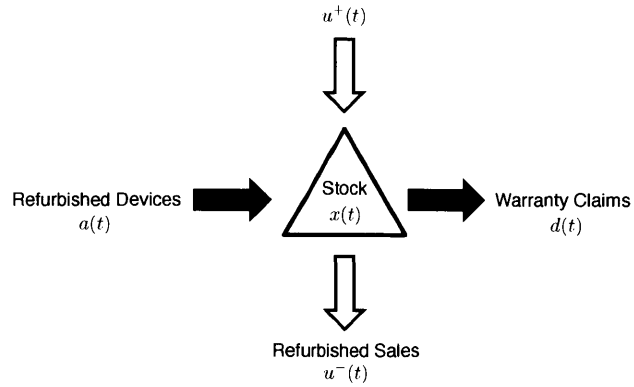

Figure 3-2: Inventory dynamics at the reverse logistics facility.

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, such that $c_t \geq c_{t+1}$, 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, and a new device cannot be purchased and then sold at a later time for a profit.

Thus, the sequence of events in a period t is

1. The inventory level at the beginning of period t is $x(t)$;

2. $a(t)$ refurbished devices arrive in inventory;

3. The demand $d(t)$ for replacement devices is received;

4. At least $\max(d(t) - a(t) - x(t), 0)$ are purchased and demand is satisfied due to the assumption that no backlogging is allowed;

5. Up to $\max(x(t) + a(t) - d(t), 0)$ devices are sold.

6. The inventory at the end of the period is $x(t+1) = x(t) + a(t) - d(t) - u^-(t) + u^+(t)$.

Since this is a CLSC, the demand and arrival processes are correlated, non-stationary, and their distributions might be unknown at the launch of the device. Together with the requirement that no backlogging is allowed, deciding how many devices should be purchased and sold can be challenging. With the notation and assumptions in hand, we are ready to discuss the deterministic version of this problem.

3.5 Deterministic Problem

In this section, we will assume that both the demand and the arrival processes are deterministic and known. For this simplified case, finding the optimal buying and selling quantities in each period is equivalent to solving a Linear Program. More specifically, if we assume that both the demand and the arrival processes are deterministic and known over a finite horizon T, the optimal buying and selling quantities are the solution of the linear program

$$\begin{aligned} \text{maximize} \quad & \sum_{t=0}^{T-1} p_t u^-(t) - c_t u^+(t) \\ \text{s.t.} \quad & x(t+1) = x(t) + a(t) - d(t) - u^-(t) + u^+(t),\ \forall t = 0 \ldots T-1 \\ & x(0) = x_o \\ & x(t), u^-(t), u^+(t) \geq 0,\ \forall t = 0 \ldots T-1 \end{aligned} \tag{3.1}$$

Despite being an optimization problem that can be solved efficiently using a numeric solver, we will prove and discuss the structure of the optimal solution for this problem. There are a three main reasons why we present this analysis: (i) By proving the structure of the optimal policy we can make explicit the relationship between the optimal policy and the dynamic cost structure of the problem; (ii) The results from this analysis will be used to prove the optimal policy of the stochastic version of the problem; (iii) The Certainty-Equivalent approximation of the stochastic problem is exactly the solution to Problem 3.1, since we take the demand and arrival processes to be equal to their expected value.

The first step towards obtaining the optimal solution is proving that there will be an optimal policy where, if we sell items at period t, we will only purchase more items once the cost of a new item falls below p_t. More specifically, if we sell at period t, we will only buy again at or after period $\bar{s}_t$, where $\bar{s}_t = \max\{s | c_s \geq p_t\}$. The definition of $\bar{s}_t$ is depicted in Figure 3-3. This is due to the monotonicity assumptions on c_t and p_t and is stated in the following proposition.

Proposition 1. For some t, let $\bar{s}_t = \max\{s | c_s \geq p_t\}$. Then for the problem in Equation 3.1, there will be an optimal solution $\{u^+_{\text{opt}}(t), u^-_{\text{opt}}(t)\}$ where, for every t where $u^-_{\text{opt}}(t) > 0$, we have $u^+_{\text{opt}}(s) = 0$, $\forall s \in [t, \bar{s}_t]$.

Proof. Assume an optimal solution $\{u^+_{\text{opt}}(t), u^-_{\text{opt}}(t)\}$. Now, let t be the last period where $u^-_{\text{opt}}(t) > 0$, i.e, we are selling devices. If $u^+_{\text{opt}}(s) > 0$ for some $s \in [t, \bar{s}_t]$, construct

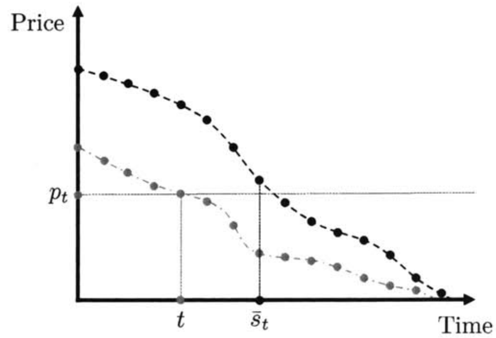

Figure 3-3: Definition of $\bar{s}_t$. The gray line represents $\{p_t\}$ while the black line represents $\{c_t\}$. Thus, as stated in Proposition 2, there will be an optimal solution where if items are sold at time t, no items will be purchased at least until time $\bar{s}(t)$.

a new solution, $\{\hat{u}^+(t), \hat{u}^-(s)\}$, where $\hat{u}^-(t) = \max(u^-_{\text{opt}}(t) - u^+_{\text{opt}}(s), 0)$, and $\hat{u}^+(s) = \max(u^+_{\text{opt}}(s) - u^-_{\text{opt}}(t), 0)$. Note that at most one of $u^-(t)$ and $\hat{u}^+(s)$ is positive. This new solution is still feasible, since less devices are leaving the system and the cost difference between these two solutions is

$$
\begin{aligned}
\left(p_t\hat{u}^-(t) - c_s\hat{u}^+(s)\right) - \left(p_t u^-_{\text{opt}}(t) - c_s u^+_{\text{opt}}(s)\right) &= (p_t - c_s)\max(-u^+_{\text{opt}}(s), -u^-_{\text{opt}}(t)), \\
&= (c_s - p_t)\min(u^+_{\text{opt}}(s), u^-_{\text{opt}}(t)), \\
&\geq 0.
\end{aligned}
$$

Thus, the new solution has at least the same cost as before, and if $c_s > p_t$ there is a strict cost improvement. We can repeat the same procedure for $\{\hat{u}^+(t), \hat{u}^-(s)\}$ and again for each new solution that we obtain. Since this is a finite horizon problem, we will eventually obtain a solution that satisfies the conditions of the proposition and has at least the same cost of the original optimal solution. □

We can determine an optimal purchasing and selling quantity at any time t as a function of the inventory at the beginning of that period. Using the same notation as in the previous proposition, we will denote the optimal buying and selling quantity at time t by $\{u^+_{\text{opt}}(t), u^-_{\text{opt}}(t)\}$. Since the cost of sourcing a device to satisfy a warranty claim is non-increasing, the optimal sourcing strategy will be myopic in the sense that we only buy enough items to satisfy the unmet demand for replacement devices in the current period. Conversely, the optimal selling quantity at some time t, $u^-_{\text{opt}}(t)$, will depend on $x(t)$ and on the *maximum total net demand* in the interval $[t, \bar{s}_t]$. The maximum total net demand in some interval $[t, s]$ is denoted by $v(t, s)$, and is defined as

$$v(t, s) = \max_{r \in [t,s]} \sum_{i=t}^{r} d(i) - a(i). \tag{3.2}$$

More specifically, if inventory $x(t)$ is above $v(t, \bar{s}_t)$, we will sell items down to the level $v(t, \bar{s}_t)$. Thus, the optimal selling quantity at time t does not depend on the demand and arrival after $\bar{s}_t$ since the marginal price of sourcing a device after time $\bar{s}_t$ drops below the marginal revenue of selling a device at time t.

Before we prove the optimal policy, we discuss one more auxiliary proposition. In this proposition, we show that if items are sold at time t, additional items will not be sold before time s_t^*, defined as

$$s_t^* = \min \left\{ s \in [t, \bar{s}_t] \middle| \sum_{i=t}^{s} d(i) - a(i) = v(t, \bar{s}_t) \right\}. \tag{3.3}$$

We will use this proposition to prove the feasibility of the dual solution to the LP in Equation 3.1.

Proposition 2. For some t, let s_t^* be defined as in Equation 3.3. Then, there will be an optimal solution $\{u_{\text{opt}}^+(t), u_{\text{opt}}^-(t)\}$ that satisfies Proposition 1 where, if $u_{\text{opt}}^-(t) > 0$, $u_{\text{opt}}^-(s) = 0$ for $s \in [t+1, s_t^*]$. Furthermore, $x(\tau + 1) = 0$ for some $\tau \in [t+1, \bar{s}_t]$.

Proof. Let t be a period where $u_{\text{opt}}^-(t) > 0$. Then, from Proposition 1, we have that $u_{\text{opt}}^+(s) = 0, \forall s \in [t, \bar{s}_t]$. Thus, we can write the inventory for any $s \in [t, \bar{s}_t]$ as

$$x(s+1) = x(t) - \left(\sum_{r=t}^{s} d(s) - a(s) \right) - \sum_{r=t}^{s} u_{\text{opt}}^-(r).$$

Also, from the definition of s_t^*, we have

$$x(s_t^* + 1) = x(t) - v(t, \bar{s}_t) - \sum_{r=t}^{s_t^*} u_{\text{opt}}^-(r).$$

Since $u_{\text{opt}}^-(t) \geq 0, \forall t$, and from the definition of $v(t, \bar{s}_t)$, note that

$$x(s) > x(s_t^* + 1), \forall s \in [t, s_t^*],$$

and, therefore, $x(s) > 0, \forall s \in [t, s_t^*]$. Now, assume that $u_{\text{opt}}^-(s) > 0$ for some $s \in [t+1, s_t^*]$. Then, we could increase $u_{\text{opt}}^-(t)$ by some $\epsilon > 0$ reduce $u_{\text{opt}}^-(s)$ by ϵ and preserve feasibility, since the inequalities are strict. In addition, since prices are non-increasing, we obtain at least the same profit as before. We can repeat this procedure for every period where items are sold and obtain the result in the proposition.

Finally, assume, for contradiction, that $x(\tau) > 0, \forall \tau \in [t+1, \bar{s}_t + 1]$. Then, in this case, we could sell one extra unit in period t and preserve feasibility in the interval $[t+1, \bar{s}_t]$. Also, we improve the objective even if selling one extra unit in t implies the purchase of one more unit after time $\bar{s}_t$, we still improve the cost sine $c_\tau < p_t, \forall \tau > \bar{s}_t$.

□

In order to prove the optimal solution of this problem, we introduce the concept of an *event*. We define an *event* as a time interval $[t, s]$, $t < s$, such that $x(t) = 0$, $x(s) = 0$, and $x(\tau) > 0, \forall \tau \in (t, s)$. Thus, the interval starts with zero inventory, then has positive inventory in each period, until the end of the interval. We can assume, without loss of generality, that the initial inventory is 0, i.e., $x(0) = 0$. If $x_0 > 0$, then we can re-specify the arrival in the first period as $a(0) = a(0) + x_0$. Note that we can express any optimal solution to the problem in Equation 3.1 as a sequence of events, as $x(T) = 0$ in any optimal solution.

The next proposition states that there will be an optimal solution to the problem where there will be at most a single buy or sale within each of this solution's events.

Proposition 3. Consider the optimal solution to the problem in Equation 3.1. Also, let $[t, s]$ be an event in this solution. Then, there is an optimal solution $\{u_{\text{opt}}^+(s), u_{\text{opt}}^-(s)\}$ where

1. There is at most a single buy within the event $[t, s]$, and if so, it will occur in period $s-1$; i.e., $u_{\text{opt}}^+(\tau) = 0, \tau \in [t, s-2]$, and $u_{\text{opt}}^+(s-1) \geq 0$;
2. There is an optimal solution with at most a single sale within the event $[t, s]$;
3. There is an optimal solution with at most a single transaction (either a buy or a sale) within the event $[t, s]$.

Proof. For part 1 of the proposition, suppose that we have two buying instants, $\tau_1 < \tau_2$ in the interval $[t, s]$. In this case, we can defer the purchase of one unit from period τ_1 to period τ_2 and maintain feasibility since we assume there is positive inventory. In addition, we potentially improve the objective function, since purchasing costs are non-increasing. We can continue until the number of units purchased in period τ_1 goes to zero or until an inventory $x(\tau), \tau \in (\tau_1, \tau_2)$ goes to zero. In both cases we have a contradiction - in the former we have reduced the number of buy epochs by one, while in the latter $[t, s]$ is no longer an event. Hence, there will be a single buy event. Also, since delaying a purchase is always no more expensive, through a similar argument we can show that the single buy should be in period $s - 1$.

For part 2 of the proposition, assume that there are now two selling instants, $\tau_1 < \tau_2$ in the interval $[t, s]$. We can advance the sale of one unit from τ_2 to τ_1 while preserving feasibility since we assume positive inventory. Also, we potentially improve the objective, since prices are non-increasing. We can repeat this procedure until the second sell epoch goes to zero or until an inventory $x(\tau), \tau \in (\tau_1, \tau_2)$ goes to zero. In both cases we have a contradiction - in the former we have reduced the number of selling periods by one, while in the latter $[t, s]$ is no longer an event.

Finally, for part 3 of the proposition, consider an event $[t, s]$. From parts 1 and 2, we know that there will be at most a single buy and a single sell event. Hence, assume that there is a sell event at time τ and a buy event at time $s - 1$. There are two possibilities:

- $p_\tau \leq c_{s-1}$: We can reduce the number of items sold at time τ by one unit, which allows us to reduce the number of items being purchased at time $s - 1$ by one unit. These changes increase the objective by $c_{s-1} - p_\tau \geq 0$. We can repeat this procedure until there is at most one transaction, leading to a contradiction.

- $p_\tau > c_{s-1}$: We can increase the number of items sold at time τ by one unit, which requires that one more unit be bought at time $s - 1$. This increases the objective function by $p_\tau - c_{s-1} \geq 0$. This solution remains feasible as we start with positive inventory in the time interval. We would continue increasing the amount sold at time τ, until this forces the inventory at some point in (t,s) to be zero, at which point [t,s] is not an event, and we obtain a contradiction.

□

Finally, we are ready to state and prove the optimal solution for this problem. We will show that the following procedure yields an optimal solution:

The feasibility of this solution is stated in the following proposition.

Algorithm 1 Procedure for obtaining the optimal solution to the problem in Equation 3.1

for $t = 0, \dots, T-1$ **do**

$u^-_{\text{opt}}(t) = \max\left(x_{\text{opt}}(t) - v(t, \bar{s}_t), 0\right),$

$u^+_{\text{opt}}(t) = \max(d(t) - a(t) - x_{\text{opt}}(t), 0),$

$x_{\text{opt}}(t+1) = x_{\text{opt}}(t) - d(t) + a(t) + u^+_{\text{opt}}(t) - u^-_{\text{opt}}(t).$

end for

Proposition 4. Algorithm 1 generates a feasible solution to the problem described in Equation 3.1. Furthermore, this solution satisfies Propositions 1 to 3.

Proof. In order to prove that $\{u^+_{\text{opt}}(t), u^-_{\text{opt}}(t)\}$ is feasible, note that

$$\begin{aligned} x_{\text{opt}}(t) + a(t) - d(t) + u^+_{\text{opt}}(t) - u^-_{\text{opt}}(t) &= x_{\text{opt}}(t) + a(t) - d(t) + \max(d(t) - a(t) - x_{\text{opt}}(t), 0) \\ &\quad - \max\left(x_{\text{opt}}(t) - v(t, \bar{s}_t), 0\right), \\ &= \max(x_{\text{opt}}(t) + a(t) - d(t), 0) - \max\left(x(t) - v(t, \bar{s}_t), 0\right), \end{aligned}$$

where we combine the first two terms to get the second line. Note that, by definition $v(t, \bar{s}_t) = \max_{s \in [t, \bar{s}_t]} \sum_{i=t}^{s} d(i) - a(i) \geq d(t) - a(t)$. Thus,

$$\begin{aligned} x_{\text{opt}}(t) + a(t) - d(t) + u^+_{\text{opt}}(t) - u^-_{\text{opt}}(t) &= \max(x_{\text{opt}}(t) + a(t) - d(t), 0) - \max\left(x_{\text{opt}}(t) - v(t, \bar{s}_t), 0\right), \\ &\geq \max(x_{\text{opt}}(t) + a(t) - d(t), 0) - \max\left(x_{\text{opt}}(t) - (d(t) - a(t)), 0\right) \\ &= 0. \end{aligned}$$

Thus, $x_{\text{opt}}(t) \geq 0, \forall t$, and the solution is feasible.

Note that this solution satisfies Proposition 2 since if $u^-_{\text{opt}}(t) > 0$, then $x_{\text{opt}}(t) > v(t, \bar{s}_t)$ and, from the definition of $v(t, \bar{s}_t)$, the inventory will never fall below 0 in the interval $[t, \bar{s}_t]$. Also, this solution satisfies Proposition 1 since we sell down to the level $v(t, \bar{s}_t)$, such that the inventory at the beginning of period $s^*_t + 1$ is 0.

The solution also satisfies Proposition 3. Assume that at time s we have $x_{\text{opt}}(s) = 0$. Then, let $\tau \geq s$ be the first period where we buy or sell. If at time τ we sell an item, since the solution satisfies Propositions 1 and 2, we will only sell again after period s^*_t and, $[t, s^*_t + 1]$ will be an event. If we buy items at time τ, since we have, from the definition of the policy, that $x(\tau + 1) = 0$, and $[t, \tau + 1]$ will be an event. Since we only buy or sell in one period during this event, Proposition 3 is satisfied. □

The next proposition addresses the lengths of the events generated by this algorithm.

Proposition 5. If $[t, s]$ is an event generated by Algorithm 1, then $s \leq \bar{s}_t + 1$, where $\bar{s}_t = \max\{s | c_s \geq p_t\}$. Furthermore, if a purchase occurs in this event, it will happen at time t, and if a sale occurs within this event, it will be at time $s - 1$.

Proof. Assume that $x(t) = 0$. Then, if $d(t) - a(t) \geq 0$, we will purchase $d(t) - a(t)$ items, i.e., $u^+_{\text{opt}}(t) = d(t) - a(t)$ and $x(t+1) = 0$; as such we have $s = t + 1$ and $[t, t+1]$ will be an event. If $a(t) - d(t) > 0$, but $v(t, \bar{s}_t) = \max_{r \in [t, \bar{s}_t]} \sum_{i=t}^{r} d(i) - a(i) \geq 0$, then we will not sell any units. In fact, no items will be sold or purchased in the interval $[t, s^*_t)$, since $v(t, \bar{s}_t) \geq 0$. At time s^*_t, we will have that $x(s^*_t) - d(t) + a(t) \leq 0$, i.e., inventory will be 0 or negative, and items might be purchased. In either case, $x(s^*_t + 1) = 0$. Since $s^*_t \leq \bar{s}_t$, the event will have length at most $\bar{s}_t + 1 - t$.

Conversely, if $a(t) - d(t) > 0$ and $v(t, \bar{s}_t) < 0$, then it must be that $a(t) - d(t) \geq -v(t, \bar{s}_t)$, and $-v(t, \bar{s}_t)$ items will be sold. From Proposition 2, we have that $x(s^*_t + 1) = 0$. Once again, the event will have length at most $\bar{s}_t + 1 - t$. If $a(t) - d(t) > 0$ and $v(t, \bar{s}_t) \geq 0$, then no items will be purchased or sold before time s^*_t, and at time s^*_t enough items will be purchased to satisfy on hand demand, such that $[t, s^*_t + 1]$ will be an event □

With these propositions in hand, we are ready to prove that Algorithm 1 leads to an optimal solution to the deterministic inventory management problem. This result is stated and proved in the theorem below.

Theorem 6. *The procedure described in Algorithm 1 generates an optimal solution for the problem in Equation 3.1. Furthermore,this procedure will take* $O(T^2)$ *steps to find an optimal policy.*

Proof. The proof of the theorem will be done by first rewriting the optimization problem (3.1) in a slightly different way and taking its dual. With the dual in hand, we will construct a feasible solution to the dual that satisfies the complementary slackness conditions with the solution $\{u^+_{\text{opt}}(t), u^-_{\text{opt}}(t)\}$, hence proving the optimality of the primal solution. We denote the inventory at time t under this solution by $x_{\text{opt}}(t)$.

The first step is to rewrite the optimization problem. For such we assume, without loss of generality, that $x_0 = 0$, and we let $\delta(t)$ be the cumulative total net supply defined as

$$\delta(t) = \sum_{s=0}^{t} a(s) - d(s), \forall t$$

Then, in Equation 3.1, note that

$$\begin{aligned} x(t+1) &= x(t) + a(t) - d(t) + u_{\text{opt}}^{+}(t) - u_{\text{opt}}^{-}(t), \\ &= \sum_{s=0}^{t} a(s) - d(s) + u_{\text{opt}}^{+}(s) - u_{\text{opt}}^{-}(s), \\ &= \delta(t) + \sum_{s=0}^{t} u_{\text{opt}}^{+}(s) - u_{\text{opt}}^{-}(s). \end{aligned}$$

Thus, this optimization problem can be rewritten as

$$\begin{aligned} \text{maximize} \quad & \sum_{t=0}^{T-1} p_t u^{-}(t) - c_t u^{+}(t) \\ \text{s.t.} \quad & \sum_{s=0}^{t} u^{-}(s) - u^{+}(s) \leq \delta(t), \; \forall t = 0, \ldots, T-1 \\ & u^{-}(t), u^{+}(t) \geq 0, \; \forall t = 0 \ldots T-1. \end{aligned} \tag{P}$$

The dual of this problem is

$$\begin{aligned} \text{minimize} \quad & \sum_{t=0}^{T-1} \delta(t)\, q(t) \\ \text{s.t.} \quad & p_t \leq \sum_{s=t}^{T-1} q(s) \leq c_t, \; \forall t = 0, \ldots, T-1 \\ & q(t) \geq 0, \; \forall t = 0, \ldots, T-1 \end{aligned} \tag{D}$$

We can interpret the dual variable $q(t)$ in the problem above marginal change in profit when increasing or decreasing the total net-demand in one unit.

Consider the sequence of events generated by the algorithm. From the previous proposition we know that, within each event, sales or purchases will happen in only one single period. Also, let $[s, t]$ be a *buying event* if, for some $\tau \in [s, t]$, we have $u_{\text{opt}}^{+}(\tau) > 0$. Conversely, let $[s, t]$ be a *selling event* if, for some $\tau \in [s, t]$, we have $u_{\text{opt}}^{-}(\tau) > 0$. Also, from the nature of the algorithm, we have that inventory at time T will be 0, i.e., no inventory is left over at the end of the horizon.

Now, consider the following candidate dual solution, $\{q^*(t)\}$, defined as follows. Let $[s,T]$ be the last event in the interval $[0,T]$. Then, for $\tau \in [s, T-1]$

$$q^*(\tau) = \begin{cases} c_{T-1}, & \text{if } [s,T] \text{ is a buying event and } \tau = T-1 \\ p_s, & \text{if } [s,T] \text{ is a selling event and } \tau = T-1 \\ 0 & \text{otherwise.} \end{cases}$$

Then, for any two sequential events $[r,s]$ and $[s,t]$, and for every $\tau \in [r,s)$, let $q^*(\tau)$ be

$$q^*(\tau) = \begin{cases} c_{s-1} - q^*(t-1), & \text{if } [r,s] \text{ is a buying event and } \tau = s-1 \\ p_r - q^*(t-1), & \text{if } [r,s] \text{ is a selling event and } \tau = s-1 \\ 0 & \text{otherwise.} \end{cases}$$

Before proving the optimality of the above dual solution, let us briefly go over an intuitive interpretation of its structure. Let us assume that at some time τ contained in the event $[r,s]$ we change the arrivals from the OEM at by ϵ units, and assume that ϵ is small. Then, using the definition of $\delta(k)$, we have that the change in the cost of the dual problem will be

$$\epsilon \sum_{k=\tau}^{T-1} q^*(k) = \begin{cases} \epsilon c_{s-1}, & \text{if } [r,s] \text{ is a buying event,} \\ \epsilon p_r, & \text{if } [r,s] \text{ is a selling event.} \end{cases}$$

Thus, if the marginal impact of changing the demand or arrivals in any time period will either be the cost of a new device at the end of the event to which the time period belongs, if it belongs to a buying event, or the side-sales price of a refurbished at the beginning of the event, if it belongs to a sales event.

In order to prove that this solution is feasible, first note that, from Proposition 5, if $[r,s]$ is an event, then $r \leq \bar{r}_s + 1$, where $\bar{r}_s = \max\{r | c_r \geq p_s\}$, and we have that $p_r \leq c_{s-1}$. Hence, we will have that $p_s \leq \sum_{t=s}^{T-1} q^*(t) \leq c_s, \forall s$.

Also, since the policy satisfies Proposition 1, we know that if $[r,s]$ is a selling event and that $[s,t]$ is a buying event, it must be that $p_r > c_t$. Using this fact, and the fact that $\{c_t\}$

and $\{p_t\}$ are non-increasing, we obtain that for any two consecutive events $[r, s]$ and $[s, t]$,

$$q^*(s-1) \geq p_r - q^*(t-1) \geq p_r - c_{t-1} \geq 0,$$

and $q^*(t) \geq 0, \forall t$. Thus, $\{q^*(t)\}$ is a dual feasible solution.

We are now left to prove that $\{q^*(t)\}$ satisfies the complementary slackness conditions with $\{u^+_{\text{opt}}(t), u^-_{\text{opt}}(t)\}$, i.e., we need to show that

$$q^*(t) \cdot \left(\sum_{s=0}^{t} u^-_{\text{opt}}(s) - u^+_{\text{opt}}(s) - \delta(t) \right) = 0, \quad \forall t \tag{3.4}$$

$$u^-_{\text{opt}}(t) \cdot \left(p_t - \sum_{s=t}^{T-1} q^*(s) \right) = 0, \quad \forall t \tag{3.5}$$

$$u^+_{\text{opt}}(t) \cdot \left(c_t - \sum_{s=t}^{T-1} q^*(s) \right) = 0, \quad \forall t \tag{3.6}$$

Let $x_{\text{opt}}(t)$ be the inventory under the policy $\{u^+_{\text{opt}}(t), u^-_{\text{opt}}(t)\}$. Thus, the first condition can be rewritten as

$$\begin{aligned} q^*(t) \cdot \left(\sum_{s=0}^{t} u^-_{\text{opt}}(s) - u^+_{\text{opt}}(s) - \bar{x}(t) \right) &= q^*(t) \cdot \left(u^-_{\text{opt}}(t) - u^+_{\text{opt}}(t) - \delta(t) + d(t) - a(t) \right), \quad \forall t, \\ &= q^*(t) \cdot (-x(t+1)), \quad \forall t. \end{aligned}$$

Since $q^*(t)$ will only be positive if $t+1$ is the boundary of an event and since, by definition, if $t+1$ is the boundary of an event then $x_{\text{opt}}(t+1) = 0$, we have that this condition is satisfied for all t.

Now, the condition

$$u^-_{\text{opt}}(t) \cdot \left(p_t - \sum_{s=t}^{T-1} q^*(s) \right) = 0, \quad \forall t,$$

will be trivially satisfied if $u^-_{\text{opt}}(t) = 0$. If $u^-_{\text{opt}}(t) > 0$, then, from Proposition 5, we have that t is the beginning of some event $[t, r]$, and that, by construction, $\sum_{s=t}^{T-1} q^*(s) = q^*(r-1) = p_t$. Thus, $p_t - \sum_{s=t}^{T-1} q^*(s) = 0$ and the condition is satisfied.

Finally, the condition $u^+_{\text{opt}}(t) \cdot \left(c_t - \sum_{s=t}^{T-1} q^*(s) \right) = 0$ will always be satisfied since, if $u^+_{\text{opt}}(t) > 0$, then $\sum_{s=t}^{T-1} q^*(s) = c_t$.

Thus, since we have a primal and dual feasible solutions that satisfy complementary slackness, the candidate solution $\{u^+_{\text{opt}}(t), u^-_{\text{opt}}(t)\}$ is an optimal solution to the problem in Equation 3.1.

Since in the procedure described in the statement of the period we must calculate $v(t, \bar{s}_t)$ for every period, which will be a procedure that takes $O(T)$ steps in each period, finding the optimal solution takes $O(T^2)$ periods.

□

From a practical standpoint, if $\bar{s}_t$ is known, the exact values of $p_s, s \in [s, \bar{s}_t]$ and $c_s, s \in [s, \bar{s}_t]$ do not need to be known in order to calculate the optimal sell-down-to level. For example, if we assume that $\{p_t, c_t\}$ depreciate at a factor of α per time period, we have that $\bar{s}_t$ can be found by solving

$$p_t = \alpha^{\bar{s}_t - t} c_t,$$

and we have $\bar{s}_t = \log_\alpha(\frac{p_t}{c_t}) = \log_\alpha(\frac{p_0}{c_0})$.

As another example, if the cost of purchasing a device is constant throughout its life-cycle, then $\bar{s}_t = T - 1, \forall t$ and the optimal solution does not depend on the exact values of prices in the side-sales channel, only on the fact that prices are non-increasing. This particular case is 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 difficult in practice. At our partner WSP, for example, side-sales can occur through auctions, generating price uncertainty. Additionally, estimating the cost of sourcing a new device to be used as a replacement device is also ambiguous: is the cost the cost of purchasing a new device from an OEM, or the opportunity cost of not having that device available in a retail store, where it could be sold to a new customer? A policy that only requires a simple ordering of costs and revenues and not a precise estimation bypasses this issue.

3.5.1 Deterministic demand and arrival processes

As discussed in the previous chapter, we assume that devices fail according to an age dependent hazard rate model that is independent of a device's sales date, where age is defined as the time a device has been with a customer. In addition, the hazard rates are unknown at the launch of the device, although priors for the hazard rates can be constructed using historical data. We denote the estimate of the hazard rate at time t for devices of age s by $h_s(t)$. Thus, denoting all the information available at time t by $\mathcal{F}_t$, the hazard rate for age k is defined as

$$\Pr(\text{device fails in period } s| \text{ device survived up to the beginning of period } s, \mathcal{F}_t) = h_s(t).$$

Also, we denote the number of working devices of age s with customers at time t by $w_s(t)$, such that the vector $w(t) = (w_0(t), \ldots, w_t(t))$ describes the number of "surviving" devices of each age at time t. Thus, the expected number of device failures at time t is

$$E[\# \text{ of failures at time } t] = \sum_{s=0}^{t} w_s(t) \cdot h_s(t).$$

For the deterministic version of this problem, we model the demand process as being the expected number of failures. In the Control Theory literature, this is referred to as a *certainty equivalent* approximation and is discussed in Bertsekas (2005). Thus, we treat the hazard rates as the actual fractions of total devices that fail. If there is information available on the hazard rates up to time $\tilde{t} \leq t$, then the demand for refurbished devices is described by the pair of equations

$$d(t) = \sum_{k=0}^{t} w_s(k) \cdot h_s(\tilde{t}), \forall t,$$

$$w_{s+1}(t+1) = (1 - h_s(\tilde{t})) \cdot w_s(t), \forall t.$$

We can generate the arrival process using a similar model, utilizing remanufacturing rates at the OEM. If $d(t)$ devices were sent to the OEM at time t, we assume that a fraction r_s of these devices will return after s periods. The arrival process then becomes

$$a(t) = \sum_{s=0}^{t} r_{t-s} d(s).$$

In practice, there is loss in this process and not all devices shipped to the OEM can be remanufactured, and thus we would expect that $\sum_s r_s < 1$.

A strategic question that arises when managing the inventory of refurbished devices is: *If the hazard 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, it can be used to bound the operational cost of supporting warranty of a new device, which is useful for planning considerations when releasing a new device into the market.

Since there is considerably less variability in the OEM remanufacturing process compared to the demand process, we assume a fixed lead time of l at the OEM and an efficiency of $0 \leq \alpha \leq 1$, i.e., a fraction of $1-\alpha$ of the devices sent to the OEM cannot be fixed. Thus, the arrival simplifies to

$$a(t) = \alpha \cdot d(t-l).$$

Also, let f_t be the fraction of devices that fail at age t, and let $\overline{w}(t)$ be the number of devices that enter the market at time t. Thus, the demand process can be written as $d(t) = \sum_{s=0}^{t} f_s \overline{w}(t-s)$. Note that we can translate these fractions to deterministic hazard rates, such that the hazard rate at time t is

$$h_t = \frac{f_t}{1 - \sum_{s=0}^{t-1} f_s}.$$

Finally, in order to answer the question, we assume that devices fail only once. The case where multiple failures are allowed significantly complicates the analysis, and we leave it as a possible extension. Thus, if "nature" can pick $\{f_t\}$ adversarialy, and if we assume $\sum_s f_s = \beta \leq 1$, we obtain the following proposition.

Proposition 7. Let f_t be the fraction of devices that fail at age t and $\overline{w}(t)$ the number of devices sold at time t. Furthermore, let $d(t) = \sum_{s=0}^{t} f_s \overline{w}(t-s)$ 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

$$\text{\# of additional units needed to satisfy demand} \leq \beta \max_{t \in [0,T-1]} \sum_{s=0}^{t} \overline{w}(s) - \alpha \overline{w}(s-l).$$

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

Proof. In the Appendix of this chapter. □

As a simple example of how this proposition can be used, assume that sales decay exponentially according to some rate γ, such that $\overline{w}(t+1) = \gamma \overline{w}(t)$. Then, assuming that $\overline{w}(t) = 0, \forall t < 0$, for some $0 \leq \alpha \leq 1$ and lead-time l we have

$$\sum_{s=0}^{t} \overline{w}(s) - \alpha\overline{w}(s-l) = \begin{cases} \frac{1-\gamma^{t+1}}{1-\gamma} \cdot \overline{w}(0), & \text{if } t < l \\ \frac{(1-\alpha)+\gamma^t(\alpha/\gamma^l-1)}{1-\gamma} \cdot \overline{w}(0), & \text{if } t \geq l \end{cases}$$

Thus, this leads to

$$\max_{t\in[0,T-1]} \sum_{s=0}^{t} \overline{w}(s) - \alpha\overline{w}(s-l) \leq \begin{cases} \frac{1-\gamma^{l}}{1-\gamma} \cdot \overline{w}(0), & \text{if } \alpha \geq \gamma^l \\ \frac{(1-\alpha)+\gamma^{T-1}(\alpha/\gamma^l-1)}{1-\gamma} \cdot \overline{w}(0), & \text{if } \alpha < \gamma^l \end{cases}$$

Note that 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 $T-1$, the end of the horizon, because the arrivals from the OEM continue to lag the failures. By making $T \to \infty$, we obtain the bound

$$\max_{t\in[0,T-1]} \sum_{s=0}^{t} \overline{w}(s) - \alpha\overline{w}(s-l) \leq \frac{1-\min(\gamma^l,\alpha)}{1-\gamma} \cdot \overline{w}(0)$$

From the proposition then, the maximum number of devices that will need to be purchased to satisfy warranty claims is $\frac{1-\min(\alpha,\gamma^l)}{1-\gamma} \cdot \beta \cdot \overline{w}(0)$. Since, in this example, the total number of devices that fail is $\frac{1}{1-\gamma}\beta\overline{w}(0)$, a fraction of at most $1-\min(\alpha,\gamma^L)$ of devices that fail will be replaced with new devices. If we assume a 85% efficiency 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 devices needed is bounded by

$$\beta \max_{t\in[0,T-1]} \sum_{s=0}^{t} \overline{w}(s) - \alpha\overline{w}(s-l) \leq \beta\frac{1-\min(\gamma^l,\alpha)}{1-\gamma} \cdot \overline{w}(0) = 0.15\frac{1-\min(0.95^3,0.85)}{1-0.95} \cdot \overline{w}(0)$$

which is about $0.43\overline{w}(0)$.

3.6 Stochastic Problem

Although tractable, the deterministic approximation to the inventory management problem presents a major limitation by not taking into account the variability of the demand and arrival processes, leading to an inventory policy that we expect would oversell refurbished items. In this section, we will analyze the stochastic version of this problem, where replacement requests are generated by the hazard rate process discussed in the previous chapter.

In addition, we will present a simple policy for managing inventory in face of stochastic demand and discuss its connection to the optimal policy.

For an informational state $\mathcal{F}_r$, and stochastic arrival and departure processes $a(t)$ and $d(t)$, the optimal policy for managing inventory will be a solution to the optimization problem

$$\begin{aligned} \text{maximize} \quad & \sum_{t=0}^{T-1} E_{\{d(t),a(t)\}}\left[p_t u^-(t) - c_t u^+(t) \middle| \mathcal{F}_r\right] \\ \text{s.t.} \quad & x(t+1) = x(t) + a(t) - d(t) - u^-(t) + u^+(t),\ \forall t = 0 \ldots T-1 \\ & x(t), u^-(t), u^+(t) \geq 0,\ \forall t = 0 \ldots T, \end{aligned} \tag{3.7}$$

where, once again, the sequence $\{c_t\}$ and $\{p_t\}$ are non-increasing, and $c_t \geq p_t, \forall t$. Here, the state of the system at time t is determined by six variables: (i) the current time t; (ii) the on-hand inventory $x(t)$; (iii) the information available about hazard distribution $\mathcal{F}_r$; and (iv) the realization of the net demand $\Delta(t) = d(t) - a(t)$, (v) the distribution of ages of devices of customers/devices in the market, (vi) the different devices currently with the OEM (on-order inventory) . However, in order to simplify notation, we will denote the *net inventory* at time t by $\overline{x}(t) = x(t) - \Delta(t)$ (which can take a negative value), and the optimal policy will be written as a function $\{u^-_{\text{opt}}(t, \bar{x}), u^+_{\text{opt}}(t, \bar{x})\}$. The structure of the optimal policy $u^*(t, \bar{x})$ is presented in the following proposition, while the proof is in the appendix. The structure of the optimal policy is such that, in every period, we will sell-down to some level $\bar{v}(t)$ that depends on the current informational state of the system.

Proposition 8. For some $\bar{v}(t) \geq 0$, the optimal policy for the stochastic inventory model in (3.7) is

$$\begin{aligned} u^+_{\text{opt}}(t, \overline{x}) &= \max(-\overline{x}, 0), \\ u^-_{\text{opt}}(t, \overline{x}) &= \max(\overline{x} - \bar{v}(t), 0). \end{aligned}$$

Where $\bar{v}(t) \geq 0$ and only depends on the distribution of demand and arrivals, and on the costs of the problem.

Proof. The proof is in the appendix of this chapter. □

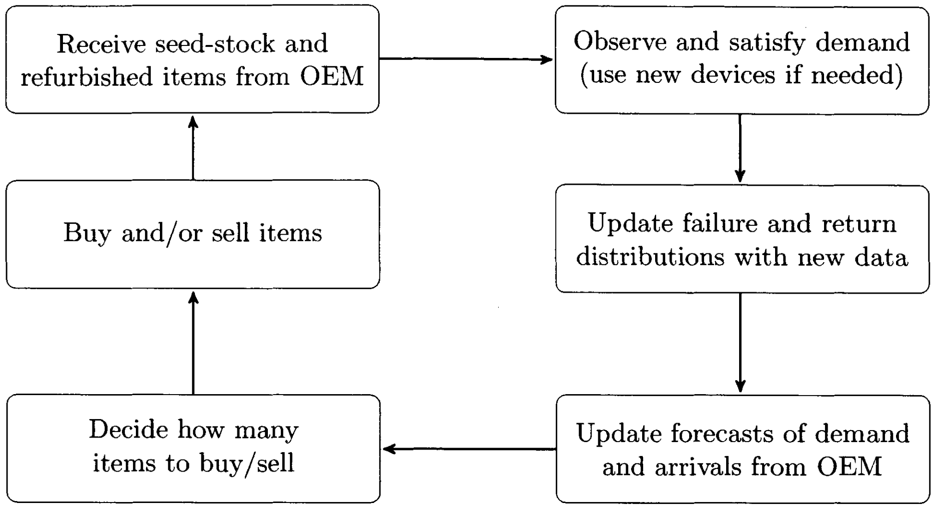

Figure 3-4: Reoptimization scheme for the inventory management problem.

Because the cost of purchasing a new device is non-increasing, an optimal policy only purchases items to satisfy immediate demand needs when $\overline{x} \leq 0$. When there are enough items in inventory to satisfy immediate needs, there will be a sell-down-to level $\bar{v}(t)$ that is time dependent, since the arrival and demand processes are not stationary. Note that the structure of the optimal policy is the same as in the deterministic case and, once again, is an artifact of the cost structure of this problem.

Calculating the sell-down-to levels is potentially difficult depending on the demand and arrival processes and, in general, involves solving the dynamic programming problem. In order to obtain a more tractable stochastic policy that is easy to implement in practice, we consider two heuristics: (i) a certainty equivalent heuristic; and (ii) a cost-balancing heuristic. They are described in the next two subsections. These heuristics can be used in a reoptimization scheme as described in Figure 3-4

3.6.1 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 in 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.

Since the volume of devices in the WSP's reverse logistics system is very large (usually in the order of millions of devices), we expect that the coefficient of variation of replacement requests per day will be low. If this is the case, we can employ the certainty equivalent approximation by taking the demand for replacement devices and the arrival process of devices in inventory as their average values. Then, we can use the results from Section 3.5 to calculate the optimal sell-down-to level.

We construct the certainty-equivalent approximation using the hazard rate model discussed in Chapter 2. In this case, the certainty-equivalent approximation to the optimal policy still purchases items myopically, but uses an approximation to calculate the sell-down-to-level in each time period. More specifically, in order to calculate the sell-down-to-level we assume that, instead of having stochastic demand and arrival processes, the demand and arrival processes are deterministic and equal to their average values.

Since we assume that devices of the same model fail according to the same hazard rate, the certainty equivalent approximation takes hazard rates as fractions of total devices that fail. Following the notation of Section 3.5.1, we have that the expected number of failures at time t will be

$$E[\#\text{ of failures at time } t] = \sum_{s=0}^{t} w_s(t) \cdot h_s(t),$$

where $w_s(t)$ is the number of devices of age s with customers at time t and $h_s(t)$ is the estimate of the hazard rate for devices of age s at time t. With this in hand, along with a forecast of the number of devices that are added to the system in every time-period through sales we can forecast the average number of devices that fail. Furthermore, if the lead time of the customer and at the OEM are also stochastic, if we have an estimate of the repair rate of the OEM, we can use the same procedure to estimate the number of arrivals.

The main limitation of this approach is that it does not take into account the variability of the demand and arrival processes. If there is a large variability in the number of devices that might fail on any given day, the certainty-equivalent approach will lead to a sell-down-to value that is too low. More concretely, let $E[d(t)]$ and $E[a(t)]$ be, respectively, the expected demand for replacement devices and arrivals from the OEM at time t. Then, following the notation from the previous section, the certainty-equivalent sell-down-to level at time

t, denoted by $v^c(t, \bar{s}_t)$ will be

$$v^c(t, \bar{s}_t) = \max_{\tau \in [t, \bar{s}_t]} \sum_{k=t}^{\tau} E[d(k)] - E[a(k)],$$

where $\bar{s}_t = \max\{s | c_s \geq p_s\}$. Note that, from Jensen's Inequality, we have that

$$\max_{\tau \in [t, \bar{s}_t]} \sum_{k=t}^{\tau} E[d(k)] - E[a(k)] \leq E\left[\max_{\tau \in [t, \bar{s}_t]} \sum_{k=t}^{\tau} d(k) - a(k)\right].$$

Thus, the sell-down-to-level generated by the certainty-equivalent policy might be too low. Next, we introduce a policy that takes into account the stochastic variability of the demand and arrival policies.

3.6.2 Cost-Balancing Approximation

The certainty-equivalent approximation assumes that the demand and arrival processes are equal to their average, leading to an underestimate of the sell-down-to-level. We can address this concern by taking the variability of the demand and arrival processes into account. One way to do this is to analyze the distribution of the maximum net-demand. Namely, at period t, we can use the distribution of $v(t, \bar{s}_t)$ and of the time period where the maximum net-demand is achieved, denoted by s_t^* as in Equation 3.3, in order to determine the sell-down-to level. In practice, since the state of the system can be very large[1], and devices can fail multiple times, the distribution of $v(t, \bar{s}_t)$ and of s_t^* can be calculated through numerical simulations.

If the distribution of $v(t, \bar{s}_t)$ and of s_t^* are known, an approximation to the optimal number of items to be sold can be obtained by solving the newsvendor-type problem

$$\min_{0 \leq u^-(t) \leq x(t)} E\left[(c_{s_t^*} - p_t)(v(t, \bar{s}_t) - x(t) + u^-(t))^+ + (p_t - p_{s_t^*})(x(t) - u^-(t) - v(t, \bar{s}_t))^+\right].$$

In this formulation, $c_{s_t^*} - p_t$ represents the marginal underage cost of selling too many items and having to buy additional new units at time s_t^*, while $p_t - p_{s_t^*}$ represents the overage cost, i.e., the opportunity cost of not selling enough items.

We denote the solution to the optimization problem above by $\tilde{u}^-(t)$ and it is obtained

[1]There are anywhere from tens of thousands to millions of customers of different ages that own devices of with varying ages as well.

by differentiating the expression, setting it to 0, and observing if the unconstrained optimal is withing the interval $[0, x(t)]$. Thus, if $\tilde{v}(t)$ satisfies

$$E\left[(c_{s_t^*} - p_{s_t^*})\mathbf{1}(v(t, \bar{s}_t) \leq \tilde{v}(t))\right] = E[c_{s_t^*} - p_t], \tag{3.8}$$

then we have

$$\tilde{u}^-(t) = \max\big(x(t) - \tilde{v}(t), 0\big).$$

From a practical standpoint, the solution to Equation 3.8 can be obtained by simulation. Namely,let $v_i(t, \bar{s}_t)$, $s^{*}_{t,i}{}_{i=1}^{N}$ be N samples of the maximum-net-demand and of the time period where it occurs. Then, the solution to Equation 3.8 can be approximated by finding the smallest $\tilde{v}_N(t)$ that satisfies

$$\frac{1}{N}\sum_{i=1}^{N}(c_{s^*_{t,i}} - p_{s^*_{t,i}})\mathbf{1}(v_i(t, \bar{s}_t) \leq \tilde{v}_N(t)) \geq \frac{1}{N}\sum_{i=1}^{N} c_{s^*_{t,i}} - p_t. \tag{3.9}$$

With the samples in hand, we can find $\tilde{v}_N(t)$ through a simple linear search. This approach is akin to Sample Average Approximation (SAA) methods commonly found in the literature such as in Levi et al. (2007) and in Kleywegt et al. (2002).

Although this method does not ensure optimality, simulating the demand and arrival paths tends to be simpler than using some search or approximate dynamic programming method to find the optimal sell-down-to level in the stochastic setting. Furthermore, if there is no uncertainty in the system, i.e., the demand and arrival processes are deterministic, then $\tilde{v}_N(t)$ will be the same as in the deterministic problem.

We call the policy based on the solving 3.8 the *cost-balancing policy*, since it balances the costs of selling too few or too many items. In the next section we will compare its performance with the certainty equivalent approximation through numerical experiments.

3.7 Numerical Experiments

In this section we will analyze the performance of the policies proposed in the previous section through numerical experiments. We first analyze the *certainty equivalent* policy where we calculate the sell-down-to level in every period by taking hazard rates as fractions of devices that fail and use them to calculate the maximum net demand and, therefore, the sell-down-to level. We simulate the sensitivity of this policy to changes in different parameters of the system such as number of devices, failure distribution, OEM lead time, seed-stock and the loss at the OEM. 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 will use the *total profit over the simulation horizon* as a performance metric in all our simulations. Thus, if the simulation occurs for T time periods and we consider a sequence of costs of new devices $\{c_t\}$ and prices of refurbished devices $\{p_t\}$, then an inventory management policy that purchases a sequence of $\{u^+(t)\}$ new devices and sells $\{u^-(t)\}$ refurbished devices will generate a profit of

$$\sum_{t=0}^{T-1} p_t u^-(t) - c_t u^+(t).$$

Next we compare the certainty equivalent policy with the *cost-balancing* policy for different distributions of device failures. We show that although the cost-balancing policy usually leads to a better performance than the certainty-equivalent approximation, when the number of devices is large (which is the case for our partner WSP) the certainty-equivalent approximation achieves a near-optimal performance and is sufficient for practical applications.

Finally, we analyze the performance of the different policies using real-world sales and failure data of a device sold by the WSP. This simulation incorporates learning, i.e., the methodology discussed in the previous chapter is employed and the hazard rate distribution of the device is updated as new information on failure rates becomes available.

In order to simulate these different scenarios, a large-scale discrete-time simulator was built using the Julia programming language. More details about Julia can be found in Lubin and Dunning (2013). The simulator creates a virtual CLSC where each device and each consumer 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 defined for each customer/device pair. Furthermore, the simulator allows for each individual device to be stored in inventory at the OEM or 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 simulations were designed to be at the finest level of granularity as possible in order to keep the simulation 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.

The variables that can be set by the user and the outputs of the simulator are described in the list below:

- User-defined simulation parameters:
  - Number of periods that the system will be simulated;
  - Number of repetitions of the simulation;
  - Market size, i.e., the total number of customers that purchase devices and will be covered by the customer warranty;
  - Time period in which each customer purchases a device, which corresponds to the activation date of the customer warranty;
  - Seed-stock level as a fraction of new devices sold
  - Failure distribution of each device (alternatively, the user can directly input the hazard rate distribution);
  - Length of the Consumer Warranty;
  - Length of the OEM warranty;
  - OEM and consumer lead time (the lead times are assumed to be deterministic);
  - Probability of a device not returning from the OEM;
  - Sales price p_t of a refurbished device in each time-period;
  - Cost c_t of a new device in each time-period;
  - Failure distribution of refurbished devices (to model multiple device failures);
- User-Defined control parameters
  - Control policy: clairvoyant, certainty-equivalent, or cost-balancing;
  - Number of sample paths for the cost-balancing simulation;
  - Matching Strategy of devices to customers (will be covered in the next chapter);
- The output statistics for each sample path are as follows:

- Number of refurbished devices shipped to customers as a replacement device in each time-period;
- Number of new devices shipped to customers as a replacement device in each time-period;
- Inventory level in each time-period;
- Number of out-of-warranty returns in each time-period: these are customers that are still covered by the customer warranty and that send in defective devices that are out of OEM warranty;
- Seed stock received in each period;
- Sales of new devices sold in each period;
- Volume of side sales in each time period;
- Current forecast of device failures in the remaining time periods;
- Current forecast of device arrivals from the OEM in the remaining time periods;
- Profit in each time period;
- Arrivals from the OEM in each time period;
- Difference between time left in OEM warranty and time left in customer warranty for each replacement device shipped.

Finally, when using pre-defined distributions, we assume that when a device is purchased by a customer, the OEM warranty and the customer warranty start simultaneously at the moment of purchase. However, the simulator does support a mismatch between these warranties at the moment of purchase. Unless specifically stated, one-period in all simulations corresponds to one day.

3.7.1 The certainty-equivalent heuristic

For all deterministic simulations, we assume that one period represents one day and that the customer and a OEM warranties have a 12 month length. Thus, if a customer has a device that fails and he is out of customer warranty, it will not be replaced. We also assume that the demand for new devices is a random variable with the probability distribution in Figure 3-5a, such that each customer samples its purchasing date from this distribution. Since we assume that the customer and device warranties start simultaneously when a customer

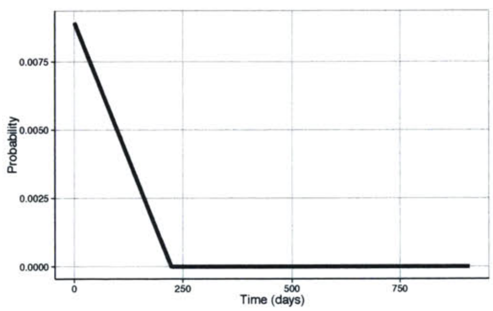

(a) Prob. of a customer purchasing a new device.

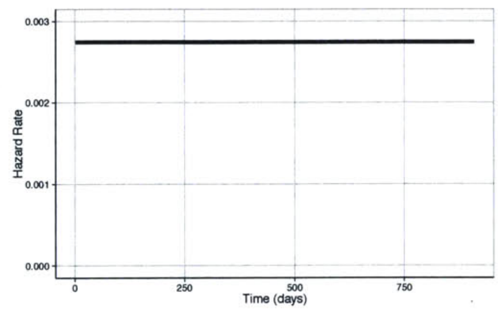

(b) Haz. rate for devices (exponential failure dist.)

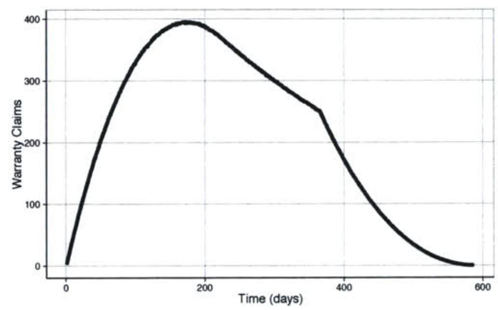

(c) Exp. number of failures for 200,000 devices sold.

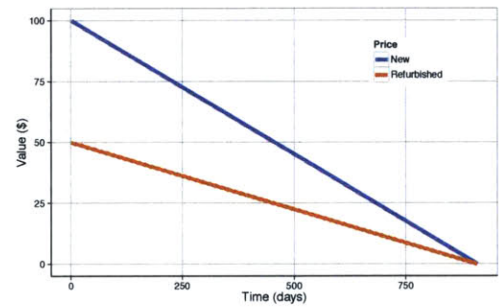

(d) Prices of new devices (c_t) and refurbished devices (p_t)

Figure 3-5: Parameters of the simulation for the deterministic policy. We assume that the failure distribution is exponential with mean 364 days. We also assume that sales only occur during the first 6 months from launch.

acquires a new device, Figure 3-5a also represents the activation date probability 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 3-5a. We 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 3-5d. Finally, we assume that the failure distribution of devices is Exponential with average 12 months. Thus, the hazard rate of devices is constant. Approximately 63% of newly purchased devices will fail under warranty.

One stringent assumption we make for this set of simulations is that the devices fail only once. Albeit unrealistic, this assumption allows us to ignore mismatches between the consumer and OEM warranties, simplifying the simulation and allowing for a larger number of simulated sample paths. In the next chapter, when we analyze strategies aimed at matching devices and customers, we will loosen this assumption.

As a benchmark, we will use the clairvoyant policy, i.e., the policy that knows a priori when all the failures and returns will occur. We compare the profit per/device between the clairvoyant policy and the certainty equivalent policy. We also analyze the ratio between these two profits for each sample path. Note that the clairvoyant policy is the harshest benchmark since there is no uncertainty. Comparing the clairvoyant policy with the optimal policy would be ideal, but a good performance with respect to the clairvoyant policy suffices to guarantee an adequate practical performance of this policy.

In Figure 3-6 we depict the performance of the certainty-equivalent for different numbers of devices sold. We assume that seed-stock corresponds to 1% of sales, that the total lead-time of the customer and the OEM is 15 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. When the number of devices is small (less than 3000), we observe that the performance of the certainty-equivalent policy is, on average, less than 75% of the clairvoyant. This is expected, since there is a small number of devices in the system, and there will be a relatively large coefficient of variation in the system. As the number of devices in the system increases, the coefficient of variation decreases (law of large numbers), and the certainty equivalent approximations improve. In practice, the average number of devices of a given type sold by the WSP is in the order of hundreds of thousands to millions of devices. Thus, if the assumption of stationarity of the hazard rate of a device of a given model holds, the deterministic heuristic will work well in practice.

Figure 3-7a depicts the profit per device sold as a function of the OEM lead-time. Note that, as lead-time increases, the potential profit per device will decrease. This occurs for two reasons. First, at the beginning of the sales-period, if the OEM lead-time is long enough, there will be insufficient seed-stock to cover the incoming warranty requests and new devices will have to be purchased. A second reason is that after the peak of warranty returns there might be an opportunity to sell inventory but, since prices decrease over time, a delay in devices arriving from the OEM leads to these devices being sold at a lower price. Thus, reducing OEM lead-time can be an effective method for reducing costs in a reverse supply chain system. At the limit, if the lead-time were zero and there were no loss at the OEM, no inventory would be needed since devices would be remanufactured instantly.

In Figure 3-7b the profit-per-device as a function of the loss at the OEM is depicted. As the loss at the OEM increases, the number of refurbished devices arriving from the OEM

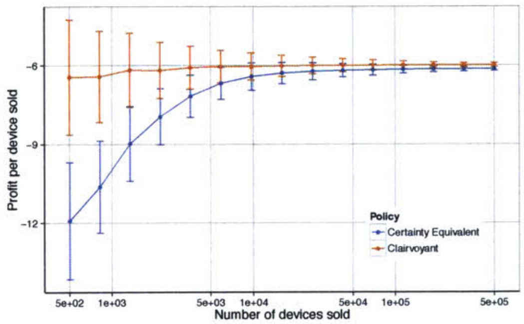

(a) Average profit per-device sold as a function of the number of devices sold.

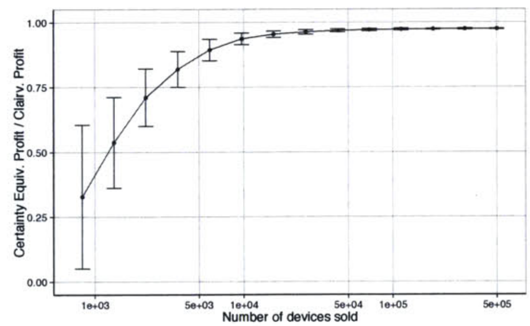

(b) Average ratio between the certainty equivalent and the clairvoyant approximations.

Figure 3-6: Comparison between the certainty equivalent and the clairvoyant policies as a function of the number of devices sold. The error bars represent +/- 2 standard deviations from the average. The average and errors are based on 300 sample paths. We assume that the seed-stock corresponds to 5% of sales, lead-time is 15 days, and loss at OEM is 20%.

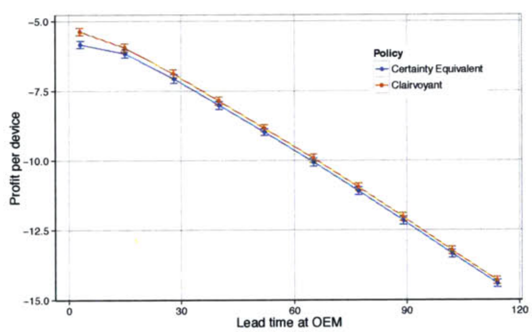

(a) Average profit per-device sold as a function of the OEM lead-time.

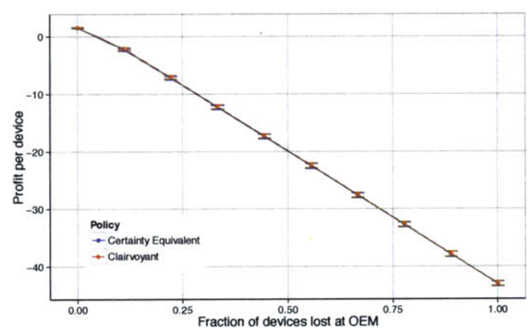

(b) Average profit per-device sold as a function of the loss at the OEM.

Figure 3-7: Comparison between the certainty equivalent and the clairvoyant policies as a function of: (a) the lead-time at the OEM and (b) the loss at the OEM. The error bars represent +/- 2 standard deviations from the average. The average and errors are based on 300 sample paths. We assume that the seed-stock corresponds to 5% of sales, there are 30,000 devices sold, and loss at OEM is 20%.

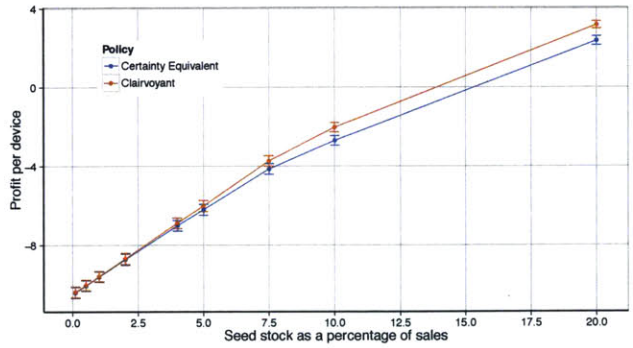

Figure 3-8: Comparison between the certainty equivalent and the clairvoyant policies as a function of seed-stock received from the OEM. The error bars represent +/- 2 standard deviations from the average. The average and errors are based on 300 sample paths. We assume that the OEM lead time is 21 days, there are 30,000 devices sold, and loss at OEM is 20%.

reduces, such that additional new devices might have to be purchased to satisfy demand for replacement devices. Furthermore, there will be potentially less inventory to sell through side-sales, further reducing the profit per device.

Finally, in Figure 3-8 we have the profit-per-device as a function of the OEM seed-stock. An increase in seed-stock counteracts the lead-time at the OEM at the beginning of the horizon, and in the period when warranty claims are increasing in time. It also helps counteract the loss at the OEM, keeping an adequate level of replacement devices.

Simulation using sales and failure data from the WSP

Using data from our partner WSP, we analyze the impact of the deterministic policy compared to the clairvoyant policy when applied to one of their best-selling devices. More specifically, for all devices of this model that the WSP sold, we had data of the sales date of each device as well as the failure dates (if the device failed at all).

In this simulation, each period represents one day, and we assume the price of a new device to be constant and equal to \$100 throughout the life-cycle of the device. Also, we assume that the price of a refurbished device in the side-sales channel to decrease linearly from \$100 to \$10 over a horizon of T = 104 weeks. We assume a standard OEM and consumer warranty length of 1 year. The aggregate lead-time of the customer and the OEM was set to be deterministic and equal to 4 weeks.

Furthermore, we assume that the sales date of each device is "known" by the heuristic, but that the failure dates are not. Thus, we consider that the heuristic does not "know" the hazard rate distribution of this device when it is launched into the market and that it estimates the hazard rate distribution using the regression approach discussed in the last chapter. In particular, we take the hazard rate distributions of five other devices to serve as a basis in our estimation problem. The estimate of the hazard rates and, therefore, of the expected demand and arrivals, is updated in every period as more failures are observed.

The performance of the certainty-equivalent policy is displayed in the table below, for different loss levels at the OEM and different fractions of seed-stock as a fraction of sales.

Table 3.1: % of Clairvoyant Profit Captured by Deterministic Heuristic

	Seed stock as % of sales			
Loss at OEM	0%	1%	2%	3%
10%	92.9%	87.9%	50.5%	35.8%
20%	98%	97.3%	96%	87%
30%	99.3%	99%	98.9%	98.6%

Note that the certainty-equivalent policy performs worse than the clairvoyant case when there is low loss at the OEM and a high seed stock level. This scenario represents a situation where a large volume of items are sold in the side-sales channel, since there is low loss at the OEM and plenty of seed stock. In this scenario, the deterministic policy ends up selling too much inventory, leading to a poor performance. In practice, the WSP received about 1% of devices as seed-stock and the observed loss at the OEM is around 20%.

3.7.2 Cost-Balancing Heuristic

The cost-balancing heuristic uses the strategy described in Section 3.6.2 to approximate the sell-down-to quantity in each period. Unlike the certainty equivalent approximation, it takes into account the uncertainty in the demand and arrival processes. However, this comes at the expense of complexity and, when using the sell-down-to policy, the distribution of the demand and arrival processes must be known or, at least, simulated. In our numerical experiments, we choose to simulate the distributions of device failures and arrivals through a Monte-Carlo simulation. Thus, in every time period of the simulations in this section, the sell-down-to level is calculated using the sample approximation described in Equation 3.9

with 100 samples.

In order to explore the performance of this heuristic, we analyze the distribution of the ratio of the profit generated by the cost-balancing heuristic and the clairvoyant policy for multiple sample paths of different failure distributions. More specifically, we consider 300 sample paths of demands and arrivals. Since this simulation is costly from a computational point of view, we simulated a smaller scale version of the system in the previous section. Namely, instead of considering a time-period as being a day, we assume that it as a week. Thus, we consider a 1 year OEM and customer warranty (52 weeks), and a volume of 5000 devices sold in the system. The sales pattern is generated by the same distribution as the one in Figure 3-5a, and that the price structure is the same as in Figure 3-5d. We also assume an OEM lead-time of 3 weeks, and OEM loss of 20%, and that the seed-stock received from the OEM corresponds to 5% of sales. All these parameters were chosen to resemble parameters actually used by our partner OEM.

The metric we use to contrast the cost-balancing policy and the certainty-equivalent policy is the ratio between the profit obtained by one of the heuristics and the clairvoyant policy. Thus, for each sample path of demand and arrivals, we calculate the difference between the ratio obtained by each heuristic and the clairvoyant profit. The distribution of these ratios suggests the "spread" of the suboptimality of each different heuristic. We visualize these distributions using box plots[2].

If Figure 3-9 a comparison between the different heuristics is depicted for a Lognormal failure distribution. The Lognormal distribution is a "heavy tailed" distribution with a large variance. When the Lognormal has a small mean and variance, as depicted in Figure 3-9a, both heuristics perform very well. This is because there will be a large number of failures and, given the seed-stock level, there will be few opportunities to sell devices in a side-sales channel. When the mean and the variance are increased, as depicted in Figure 3-9b, the cost-balancing policy outperforms the certainty-equivalent policy. This happens precisely because the cost-balancing policy takes into account the distribution of failures and arrivals of devices in inventory, making better selling decisions when failures happen according to a heavy-tail distribution.

Next, as depicted in Figure 3-10 we assume that the failure distribution is Uniform. In Figure 3-10a all devices will fail at some point, leading to a scenario where few refurbished

[2] http://en.wikipedia.org/wiki/Box_plot

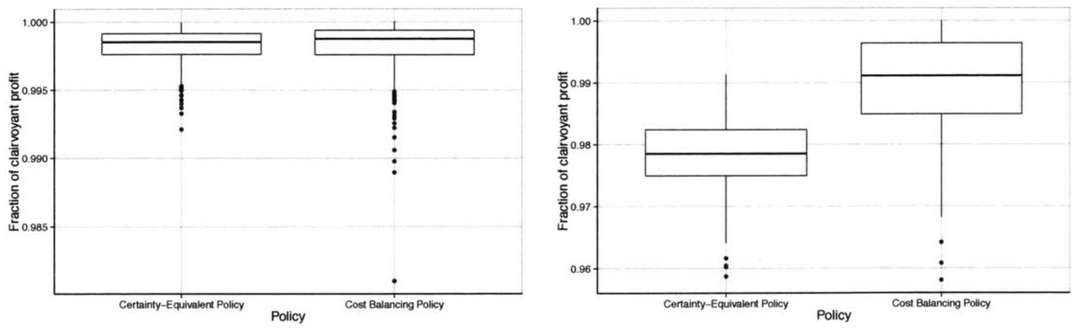

(a) Lognormal parameters: $\mu = \log(13)$, $\sigma^2 = \log(13)$

(b) Lognormal parameters: $\mu = \log(52)$, $\sigma^2 = \log(17.3)$

Figure 3-9: Comparison between the certainty equivalent and the cost-balancing policies policies for a Lognormal failure distribution. Note that the axis on each figure is different. 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.

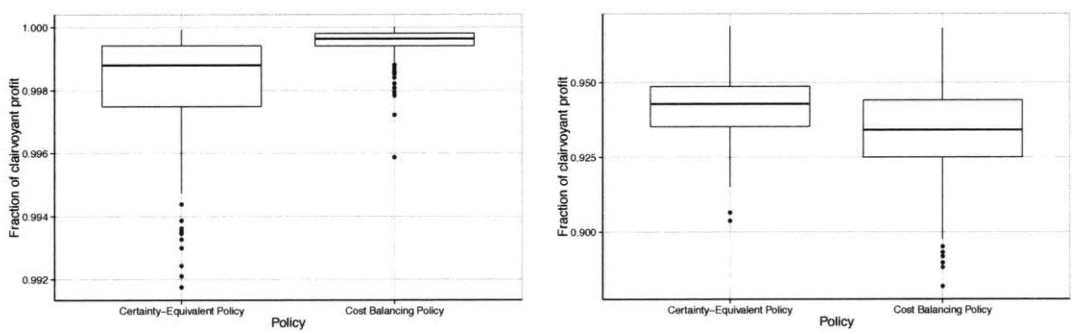

(a) Uniform between $[0, 52]$

(b) Uniform between $[0, 104]$

Figure 3-10: Comparison between the certainty equivalent and the cost-balancing policies policies for a Uniform failure distribution.

devices are sold, and both heuristics perform near-optimal. When only half of devices fail under warranty coverage, as depicted in Figure 3-10b, the certainty-equivalent heuristic and the cost balancing heuristic are not statistically different. In this case, we have a relatively small variability in weekly failures (other than the non-stationary failure average), and the cost-balancing policy does not present an advantage over certainty-equivalent policy.

Finally, we consider the case where failure ages are exponentially distributed, having a constant hazard rate. This is depicted in Figure 3-11. In this case, when the average of the exponential distribution is large, both heuristics are, once again, statistically indistinguishable. Since the hazard rate is constant, the total number of devices that will fail in each

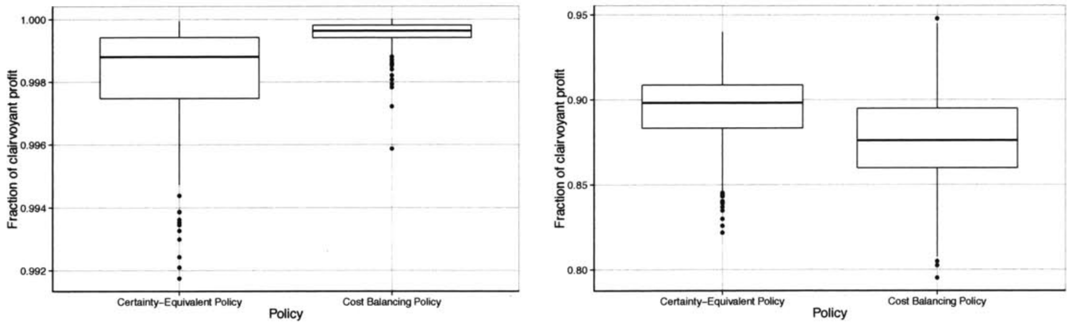

(a) Exponential with mean 52 weeks.

(b) Exponential with mean 104 weeks.

Figure 3-11: Comparison between the certainty equivalent and the cost-balancing policies policies for an Exponential failure distribution.

period will have a small variability and both policies will have similar performances.

In these simulations, we are assuming that devices fail only once in order to simplify our numerical experiments. If devices fail multiple times and if, in addition, we allow for mismatches between the customer warranty and the OEM warranty, estimating the average number of failures and arrivals from the OEM becomes a more challenging task. In practice, the average of the demand and arrival paths can be approximated through a Monte-Carlo simulation, using a strategy similar to the cost-balancing heuristic, and described in Equation 3.9. Thus, in this case, the computational expense of both methods is similar.

Simulation using sales and failure data from the WSP

Using the same data and simulation set-up described in section 3.7.1, we compare the cost-balancing policy with the clairvoyant policy using data from one of their best-selling devices. As before, we assume the price of a new device to be \$100 throughout the life-cycle of the devices and we assume that the price of a device in the side-sales channel decreases from \$100 to \$10 over the horizon of T = 104 weeks. We assume that decisions are made on a daily basis, such that the decision about the number of devices that can be purchased and sold is made every day.

We once again assume that the hazard rate distribution is not known by the heuristic and that it is estimated using the hazard rate regression approach discussed in the previous chapter. With the hazard rate distribution estimate in hand, we use 100 samples to calculate the sell-down-to level every day. The results of the simulation are depicted in the table below:

Table 3.2: % of Clairvoyant Profit Captured by Cost-Balancing Heuristic:

	Seed stock as % of sales			
Loss at OEM	0%	1%	2%	3%
10%	94.3%	90.4%	86.6%	64%
20%	98.2%	97.7%	96.8%	90.4%
30%	99.4%	99.2%	99%	98.7%

Note that the cost-balancing heuristic has a better performance than the deterministic heuristic (shown in the previous section) since it takes into account uncertainty in the arrival and demand processes. However, this difference is, in general, less than 5%. Also, the cost-balancing policy performs better than the certainty equivalent case when there is low loss at the OEM and a high seed stock level. In practice, the WSP usually observes a loss of around 20% at the OEM and receives a red stock of 1% of sales.

3.8 Conclusion

In this chapter, we proposed and analyzed two models for an inventory management problem commonly found in the reverse logistics operations of consumer electronics retailers. We developed this work in the context of a large WSP, which is also one of the largest retailers in its segment. The models that we propose take into account the closed-loop nature of reverse logistics systems in this setting, as well as the short life-cycle of electronic and the fast value depreciation that these devices suffer.

First, we introduced and examined a discrete-time deterministic model for this problem, where the demand for replacement devices and the arrival of refurbished devices in inventory are known. We proved the optimal policy for this case and also presented a worst case analysis. Even though costs, demand, and arrivals change over time, the optimal policy has a simple structure and can be easily implemented in practice. More specifically, the optimal sourcing strategy will be myopic in the sense that we only buy enough items to satisfy the unmet demand for replacement devices in the current period. Conversely, the optimal selling quantity in some time period will depend on the inventory level at the beginning of the period and 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-to 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.

Next we introduced a discrete-time stochastic model, for which we proved the structure of the optimal policy and presented two heuristics. The first heuristic is a *certainty-equivalent* approximation where we obtain a suboptimal inventory control policy by approximating the uncertainty in the problem by its average value. Since the volume of devices in the WSP's reverse logistics system is very large, this heuristic works well in practice. The second heuristic is the *cost-balancing policy* that takes into account the uncertainty of the demand and arrival process. More specifically, we solve a newsvendor-type problem that balances the costs of selling too few or too many items. We find the optimal sell-down-to level through a Sample Average Approximation (SAA), a well studied approach in the Operations Management literature.

Through numerical experiments, we analyzed the performance of these policies and simulated their sensitivity with respect to changes in different parameters of the system. As a benchmark, we compared their performance with the *clairvoyant policy*, the policy that knows ex-ante the sample path of the device failures and the arrivals from the OEM. We observed that although the cost-balancing policy usually leads to a better performance than the certainty-equivalent approximation, when the number of devices is large (which is the case for our partner WSP) the certainty-equivalent approximation achieves a near-optimal performance and is sufficient for practical applications.

Finally, we examined the performance of the different policies using real-world data from a device sold by the WSP. This simulation incorporates learning, i.e., the methodology discussed in the previous chapter is employed and the hazard rate distribution of the device is updated as new information on failure rates becomes available. We observe that in practical settings both policies capture over 90% of the clairvoyant profit.

There are a few open problems that were not addressed in this chapter. First, note that many issues in this supply chain are caused by the fact that the consumer warranty and the OEM warranty have different specifications. Namely, the consumer warranty requires a fast replacement of a failed device, while the OEM warranty allows for a comparatively large lead time for the OEM to refurbish failed devices. Studying how to redesign these contracts by taking into account the incentives and preferences of the consumer, the WSP and the OEM can 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. Although matching customers and devices is the subject of the next chapter, we did not jointly examine the closed-loop inventory management and matching problems. Looking at these two problems together can lead to new policies for managing this system that could have interesting theoretical properties.

3.9 Appendix: Proofs

Proof of Proposition 7

Proposition 7. Let f_t be the fraction of devices that fail at age t and $\overline{w}(t)$ the number of devices sold at time t. Furthermore, let $d(t) = \sum_{s=0}^{t} f_s \overline{w}(t-s)$ 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

$$\# \text{ of additional units needed to satisfy demand} \leq \beta \max_{t \in [0,T-1]} \sum_{s=0}^{t} \overline{w}(s) - \alpha \overline{w}(s-l).$$

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

For failure rates $\{f_t\}$, 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

$$\begin{aligned} v^*(0) &= \max_{k \in [0,T-1]} \sum_{i=0}^{k} d(i) - a(i) \\ &= \max_{k \in [0,T-1]} \sum_{i=0}^{k} d(i) - \alpha d(i-l). \end{aligned}$$

The demand process can be written as

$$d(t) = \sum_{s=0}^{t} f_s \overline{w}(t-s).$$

Thus, we have that

$$
\begin{aligned}
v^*(0) &= \max_{t\in[0,T-1]} \sum_{s=0}^{t}\sum_{r=0}^{s} f_r \overline{w}(s-r) - \alpha \sum_{r=0}^{s-l} f_r \overline{w}(s-l-r), \\
&= \max_{t\in[0,T-1]} \sum_{s=0}^{t} f_s \sum_{r=0}^{t-s} \overline{w}(r) - \alpha \sum_{s=0}^{t-l} f_s \sum_{r=0}^{t-l-s} \overline{w}(r), \\
&= \max_{t\in[0,T-1]} \sum_{s=0}^{t} f_s \left(\sum_{r=0}^{t-s} \overline{w}(r) - \alpha \overline{w}(r-l) \right)
\end{aligned}
$$

In order to find the worst case failure rate, i.e., the failure rate that maximizes $v^*(0)$, consider the problem

$$
\begin{aligned}
\text{maximize} \quad & \sum_{s=0}^{t} f_s \left(\sum_{r=0}^{t-s} \overline{w}(r) - \alpha \overline{w}(r-l) \right) \\
\text{subject to} \quad & \sum_{t} f_t \le \beta, \; f_t \ge 0, \forall i.
\end{aligned}
$$

The optimal cost will be simply $\beta \max_{s\le t} \left(\sum_{r=0}^{t-s} \overline{w}(r) - \alpha \overline{w}(r-l) \right)$ and the optimal failure rate will have only one non-zero component. Thus, we obtain the inequality

$$
\begin{aligned}
v^*(0) &\le \max_{t\in[0,T-1]} \beta \max_{s\le t} \left(\sum_{r=0}^{t-s} \overline{w}(r) - \alpha \overline{w}(r-l) \right) \\
&= \max_{t\in[0,T-1]} \beta \left(\sum_{s=0}^{t} \overline{w}(s) - \alpha \overline{w}(s-l) \right).
\end{aligned}
$$

Which is the bound in the proposition.

Proof of Proposition 8

The proof of the proposition will be done by backwards induction on t, and will share some elements with the proof of the deterministic case. Since no backordering is allowed, we have

$$
J_{T-1}(\overline{x}) = \max_{0\le \overline{x}+u^+-u^-} -c_{T-1}u^+ + p_{T-1}u^-, \tag{3.10}
$$

$$
J_t(\overline{x}) = \max_{0\le \overline{x}+u^+-u^-} -c_t u^+ + p_t u^- + E\left[J_{t+1}(\overline{x}+u^+-u^-) \right], \forall t \in [0, T-1]. \tag{3.11}
$$

Note that the expectation in the second equation is taken over the demand and arrivals in period $t+1$ on-wards. The induction hypothesis are

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

2. The optimal ordering policy is, for some $\bar{v}(t) \geq 0$,

$$u^{+}_{\text{opt}}(\overline{x}, t) = \max(-\overline{x}, 0) \tag{3.12}$$

$$u^{-}_{\text{opt}}(\overline{x}, t) = \max(\overline{x} - \bar{v}(t), 0) \tag{3.13}$$

For $t = T - 1$, we will by or sell all leftover items, such that $\bar{v}(T-1) = 0$ and all the the induction hypothesis holds. Now, assume that the hypotheses hold for $t + 1$. For ease of exposure, let $u(t)$ be defined as the net number of devices purchased and sold, i.e.,

$$u(t) = u^{-}(t) - u^{+}(t)$$

Then,

$$J_t(\overline{x}) = \max_{0 \leq \overline{x} - u} \min(c_t u, p_t u) + E\left[J_{t+1}(\overline{x} - u)\right].$$

which, by defining $y = \overline{x} - u$, can be rewritten as

$$J_t(\overline{x}) = \max_{0 \leq y} \min(c_t(-y + \overline{x}), p_t(-y + \overline{x})) + E\left[J_{t+1}(y)\right].$$

Since the expression

$$\min(c_t(-y + \overline{x}), p_t(-y + \overline{x})) + E\left[J_{t+1}(y)\right]$$

is concave for all $(y, \overline{x})$, then partial maximization over y preserves concavity, and $J_t(\overline{x})$ will be concave, satisfying the first induction hypothesis. Also, the marginal profit variation of adding or removing an item from inventory at time t will always be less than c_t, since removing an item would incur a future cost of at most c_t for any realization of the demand and arrival processes, such that $J_t'(\overline{x}) \leq c_t$.

For the second induction hypothesis, recall Equation (3.11) and note that, if $\overline{x} \leq 0$, we have

$$J_t(\overline{x}) = \max_{0 \leq \overline{x} - u} c_t u + E\left[J_{t+1}(\overline{x} - u)\right].$$

Since, $J_t'(\overline{x}) \leq c_{t+1} \leq c_t$, the expression above is non-decreasing for $u \leq 0$, and $u^*(\overline{x}, t) = \overline{x}$.

Conversely, if $\overline{x} > 0$, we have that $u^*(\overline{x}, t) \geq 0$, since $c_t u + E\left[J_{t+1}(\overline{x} - u)\right]$ is non-decreasing for $u \leq 0$. Therefore, the optimal policy and cost can be found by solving

$$J_t(\overline{x}) = \max_{0 \leq u \leq \overline{x}} p_t u + E\left[J_{t+1}(\overline{x} - u)\right].$$

which, by once again making $y = \overline{x} - u$, we can rewrite as

$$J_t(\overline{x}) = \max_{0 \leq y \leq \overline{x}} p_t \overline{x} - p_t y + E\left[J_{t+1}(y)\right].$$

Let y^u denote the unconstrained minimum of this optimization problem and let $y^*(\overline{x})$ denote the constrained minimum. Then, we have that

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

or, equivalently,

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

If we define $\bar{v}(t) = \max(0, y^u)$, we obtain

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

which satisfies the second induction hypothesis.

Chapter 4

Warranty Matching in a Closed-Loop Supply Chain

4.1 Introduction

In the previous chapter we addressed the issue of managing inventory in a reverse logistics system that supports a warranty contract between a large Wireless Service Provider (WSP) and its customers. The problem of managing inventory at the WSP's reverse logistics facility was studied, and we analyzed various policies for buying and selling refurbished devices. However, once an inventory management policy is fixed, it is still necessary to assign inventory items to customer requests according to the warranties that are in place. We investigate this challenge next.

More specifically, as described in detail in Section 3.1 of the previous chapter, there are two warranties in place: (i) the consumer warranty, offered by the retailer to the consumer, and (ii) the OEM warranty, offered by the OEM to the retailer. Although both warranties have the goal of protecting the players in this supply chain against manufacturing defects, they might have very different characteristics. For example, the consumer warranty might guarantee an immediate replacement of the faulty device for a working (new or refurbished) one, while the OEM warranty might require the vendor to wait for the device to be repaired, such that a replacement device is not immediately sent to the retailer.

Both warranties are valid for a limited period (usually 12 months), and once warranties expire, the coverage to replace or repair a faulty device ends. Namely, a customer does not receive a replacement if he is out of consumer warranty, and the retailer cannot send the device to the OEM for repairs if it is out of OEM warranty. In addition, the OEM warranty is associated to a specific device, while the consumer warranty is specified to the consumer.

The WSP would ideally like to have the two warranties for a device being matched, i.e., the customer would have the same time left in his consumer warranty as the device would have left in the OEM warranty. A mismatch between these warranties can incur costs to the retailer beyond the usual processing costs of warranty requests. Namely, this extra-cost is incurred when a customer still covered by the consumer warranty has a device that fails, and this device is not covered by the OEM warranty. In this case, the WSP will then either pay for the OEM to repair the device, which incurs additional costs to the system, or it will scrap the device and the device leaves the system. If the device leaves the system, it cannot be used in the future as a replacement device and it also cannot be sold through the side-sales channel. At our partner WSP, these out-of-OEM-warranty devices are a significant source of cost for their reverse operations.

Since a device can fail multiple times during its lifecycle, and the failure rate of refurbished devices is about the same as for new devices, the replacement device sent to customers that file warranty requests can lead to out-of-OEM-warranty returns. Also, the OEM warranty does not restart once a device is remanufactured and it is not paused while a device is in stock at the WSP, such that "old" devices, with little OEM warranty left, can potentially be sent to customers as replacements. At the WSP's reverse logistics facility, devices in stock were matched at random to consumers that placed warranty claims. More specifically, refurbished devices received from the OEM were not sorted by time left in OEM warranty, and customer requests were also not sorted according to the time left in their customer warranty. This would lead to "old" devices being sent to "young" customers, creating a scenario where a customer with a few months left in its consumer warranty receives a device with an expired OEM warranty. Conversely, this would also lead to cases where "young" devices were sent to "old" customers, effectively wasting OEM warranty coverage time.

From a practical standpoint, our goal is to propose and analyze assignment strategies that mitigate these out-of-OEM-warranty returns in this system, i.e., minimize the number of replacement requests from customers still covered under the retailer warranty, but whose failed device is not covered by the OEM warranty. For such, we will first discuss the drivers of out-of-OEM-warranty returns and then proceed to formulate the problem of assigning devices in inventory to customers. The final parts of the chapter will be dedicated to discussing and analyzing different strategies for assigning devices to customers.

4.1.1 A note on nomenclature

Throughout this chapter, we use the term *old customers* to refer to customers that have little time left in their consumer warranty. Since the consumer warranty has, in general, about a 1 year duration, an *old* customer can be thought of a customer with only 1 or 2 months left in his warranty. Conversely, *young customers* are customers that still have most of their warranty left. Along these lines, when we refer to the customer's *age*, we are referring to how long he is into the customer warranty.

Similarly, we refer to a device as an *old device* if it has little time left in the OEM warranty. On the other hand *young devices* are new and have most of the OEM warranty left. OEM warranties, at our partner WSP, usually lasted 1 year as well. Thus, if the customer warranty has length T_c and the OEM warranty has length T_{oem}, then a customer of age t has $T_c - t$ left in the customer warranty, and a device of age t has $T_{oem} - t$ left in the OEM warranty.

4.1.2 What drives out-of-warranty returns?

Before formulating mathematically the assignment problem, we will discuss the drivers of out-of-warranty returns from a conceptual level. This will help identify what are the drivers of mismatches and, in a later section, will guide the development of a policy for matching devices to customers. Specifically, assuming that the devices are homogeneous in the sense that they all obey the same failure-rate distribution, we argue that the main drivers of out-of-warranty returns are (i) the length of the sales period, i.e., the amount of time that devices are being sold; (ii) the matching algorithm used to send devices to customers; (iii) random fluctuations in sales and failures.

Sales Period

If all devices were sold on the same day, and assuming that both the OEM and the retailer warranty start on that sales date, out-of-warranty returns in this system would not occur since all devices and customers would have the same age. If the OEM and consumer warranties have the same duration, all devices would have the same amount of time left in the OEM warranty as customers have in their warranty, and a mismatch will never happen.

However, an increase in the length of the sales period leads to devices of different ages failing in the same time period. It also leads to refurbished devices of different ages arriving from the OEM to inventory at the reverse logistics facility in the same time period. Thus, the longer the sales period, the larger the possibility of mismatches being created.

As a simple example of this fact, assume that when a customer purchases a new device, the device and the customer warranty start simultaneously. Also, assume that customers have a deterministic failure rate of α per week, i.e., in each week $\alpha \cdot 100\%$ of customers have devices that fail, independent of the customer's age. Finally, assume that there is a lead time of l weeks for refurbishing devices, i.e., after l weeks a broken device will be available to be used as a replacement. To simplify our analysis, we assume that replacement devices do not fail, i.e., devices can fail at most once. In this setting, consider two scenarios: (i) n devices are sold in one day; (ii) n/T devices are sold each week over a period of T weeks.

In scenario (i), in week t there will be $(1-\alpha)^{t-1}\alpha n$ customers with failed devices that require replacements. All of these customers will have the same age t. The replacement requests from these customers will be satisfied from refurbished devices in inventory (as well as from seed stock), and all devices in inventory will have the same age t as the customers, since they are devices that failed at or before week $t-l$ and were refurbished. Thus, for any assignment strategy there will be no mismatch, since devices in inventory will have the same age as customers requesting replacements.

For scenario (ii), since devices are sold during multiple weeks, failures will be composed of a mix of devices from customers of different ages. More specifically, in week $t \geq T$, there will be $(1-\alpha)^{t-s}\alpha n/T$ failures from customers that purchased devices in week $s \in [1, T]$. This is illustrated in Figure 4-1. There we assume that $T = 6$, $n = 900$, and $\alpha = 0.05$. Note that there will be a mix of devices of different ages in inventory and, depending on the assignment strategy, there could be mismatches between customers and replacement devices. Note that, the longer the sales period, the larger the mix of different devices in inventory. Furthermore, recall that, in this case, the make up of refurbished inventory is a translation of the failures shown in the right figure.

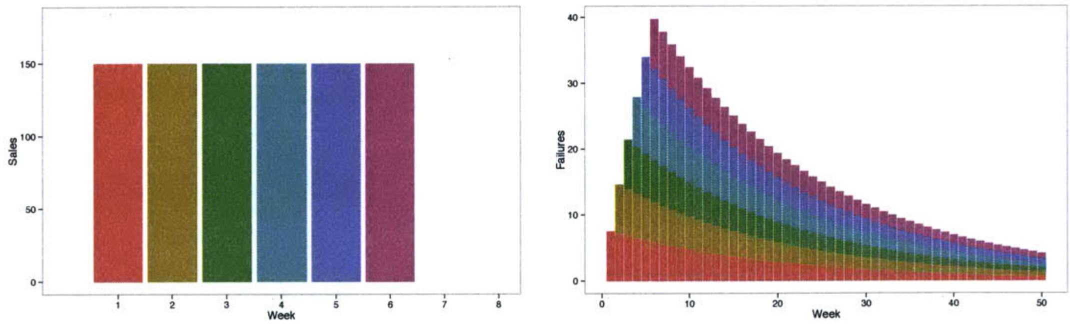

(a) Sales per week

(b) Failures per Week

Figure 4-1: Sales per week and failures per week for the example with deterministic sales and returns. We assume that there are only 6 weeks of sales, 150 items are sold per week, and each week 5% of customers have devices that fail. In (b), each color represents the devices from a sales week. Note that the customers with devices that fail are a mix

Assignment Strategy

The assignment strategy is crucial in mitigating out-of-warranty returns, and simple strategies that do not take into account the closed-loop nature of the system can lead to a large average mismatch. For example, a first-in-first out strategy for assigning replacement devices to customers, not taking into account customer ages, might lead to large mismatches. Similarly, randomly assigning devices to customers, which was the practice used by our partner WSP, might also lead to large mismatches and, consequently, to poorly matched assignments.

In the next few sections of this chapter we will propose and analyze different assignment policies. Namely, we will focus on three strategies:

- The *Random assignment policy*: where devices in inventory are randomly assigned to customers that require replacement devices, ignoring the time remaining in both the customer and device warranties;

- The *Youngest-Out-First (Myopic) assignment policy*: where in every time period devices in inventory are assigned to customers so as to minimize the mismatch in that specific period;

- The *Oldest-Out-First policy*: a policy that always assigns the oldest devices in inventory to the oldest customers that require replacements.

Random fluctuations in sales and failures

As discussed in Chapter 2, we assume that devices of the same type fail according to some hazard rate distribution that is unknown at the beginning of a device's live-cycle. Since devices fail at random times, the mix of the ages of devices that arrive in inventory and the ages of customers requiring replacements depends on the failure distributions of devices.

In addition, there is also uncertainty regarding the volume of new customer-device pairs entering the system in every time-period. Sales of new devices are affected by multiple variables, such as quality of devices, marketing efforts, prices and discounts. This constantly changes the volume of devices entering the market and, therefore, the mix of devices that are failing in every time period.

The rest of the chapter is structured as follows. In Section 4.2, we present a literature review. In Section 4.3, we formulate the matching problem. In Section 4.4, we analyze different matching strategies assuming stationary distributions of failure ages. In section 4.5, we discuss in depth the oldest-out-first assignment strategy. Finally, in Section 4.6, we evaluate the different policies through various numerical experiments.

4.2 Literature Review

Research in matching items and individuals has been very active during the last 50 years. One of the seminal works in this area is on the Hungarian algorithm, and can be found in Kuhn (1955). More recently, the problem of matching under preferences has played a central role in economics and operations research. In particular the stable matching problem, which is discussed in Gusfield and Irving (1989), has been an area of active research. A survey on relationship between market design and matching problems can be found in Roth and Sotomayor (1992), and an application to matching medical students to residency programs is covered in Roth (2003).

In the operations management literature, matching problems have been studied in the context of blood bank inventory management in Jennings (1973) and in Pegels and Jelmert (1970). The role of matching strategies in kidney-exchange programs has also received considerable attention. A framework for studying this problem from a game-theoretic point of view can be found in Ashlagi et al. (2011) and, more recently, in Ashlagi et al. (2013).

Although we build upon these streams of research, our work considers a very different setting. First, we are analyzing a matching problem in the context of a closed-loop supply chain. Furthermore, our set-up is dynamic, and decisions made in one period impact the system's state in the future. Finally, our goal is not to find an allocation that is an equilibrium, but to find a cost minimizing match. Given these characteristics, our work also draws from the literature on on-line allocation. An example of an on-line algorithm for minimum metric bipartite matching is Meyerson et al. (2006). Also, in Bansal et al. (2007), an on-line algorithm for bipartite matching and its competitive ratio are examined.

A review of tools for analyzing on-line algorithms, which encompass the type of policies that we use in the third part of the thesis, can be found in Albers (2003). Finally, the type of matching problem that we consider in this thesis has yet to be covered in the literature, and we believe that our application provides an interesting context for assignment and matching algorithms.

4.3 Problem Formulation

We will formulate the problem of assigning devices to customers as a discrete time problem where, at the beginning of each discrete period, customers of different ages place requests for replacement devices, and devices in inventory are assigned to these customers. In order to maintain the mathematical description of the system as close as possible to real-world reverse-logistics systems, we will assume that each customer and device has a unique numeric identifier. From a practical perspective, the device identification number is its serial number, while the customer identification number is his/her registration number. The set of all device ids in the system at time t will be given by $\mathcal{D}_t$, while the set of all unique customer ids active in the system at time t will be given by $\mathcal{C}_t$. Since customers own devices, we will denote the set of customer-device pairs at time t by $\mathcal{E}_t$. Thus, if customer $i \in \mathcal{C}_t$ owns device $j \in \mathcal{D}_t$ at time t, then $(i, j) \in \mathcal{E}_t$.

The pool of customers, devices, and customer-device assignments changes over time, since customers and devices are constantly entering or leaving the system. In addition, it also depends on the assignment strategy used in assigning customers to new and/or replacement devices. The *activation date*, i.e., the date that the customer warranty starts, will be given by a function $z : \{\mathcal{C}_t\} \cup \{\mathcal{D}_t\} \to \mathbb{N}$, such that if $i \in \mathcal{C}_t$, $z(i)$ will be the start date of i's customer warranty and if $j \in \mathcal{D}_t$, $z(j)$ will be the start date of device j's OEM warranty. We assume that customers have a finite warranty with length given by T_w, such that if, at time t, $t - z(i) = T_w, i \in \mathcal{C}_t$, then the customer will leave the active customer pool at time $t+1$. The set of customer-device pairs that leave the system due to expired customer warranties will be given by $\mathcal{E}_t^{\text{exp}}$. Also, the set of new customer-device pairs that enter the system through sales in the forward chain at time t is denoted by $\mathcal{E}_t^{\text{new}}$.

The set with all devices ids in inventory at time t in the reverse-logistics facility will be denoted by $\mathcal{I}_t$. The set of devices in inventory also changes over time, and depends on the remanufactured devices arriving from the OEM, on devices being sold through side-sales channels, and on devices sent as replacements to customers. We will denote the set of device ids that arrive in inventory from the OEM at time t by $\mathcal{A}_t$ and the ones that leave through side sales channels or due to expired customer warranty (they belong to a customer-device pair in $\mathcal{E}_t^{\text{exp}}$) by $\mathcal{O}_t$. Also, devices that are sent as replacements to customers at time t will be denoted by $\mathcal{R}_t$. Thus, the set of devices in inventory evolves according to the expression

$$\mathcal{I}_{t+1} = \mathcal{I}_t \cup \mathcal{A}_t - \mathcal{O}_t \cup \mathcal{R}_t.$$

The set of customer-device pairs that fail at time t is given by $\mathcal{F}_t$, such that if $(i,j) \in \mathcal{F}_t$ customer i that owns device j has a faulty device and needs a replacement. Also, let the set of customer-device pairs created due to customers receiving replacements be $\mathcal{E}_t^{\text{rep}}$. Hence, if customer-device $(i,j) \in \mathcal{F}_t$ is matched to device $k \in \mathcal{R}_t$, then $(i,k) \in \mathcal{E}_t^{\text{rep}}$. Then, we can write the evolution of $\mathcal{E}_t$ as

$$\mathcal{E}_{t+1} = \mathcal{E}_t \cup \mathcal{E}_t^{\text{new}} \cup \mathcal{E}_t^{\text{rep}} - \mathcal{F}_t \cup \mathcal{E}_t^{\text{exp}}.$$

Let the set of all customers that have devices that fail be given by $\mathcal{C}_t^f$. Thus, $i \in \mathcal{C}_t^f$ if there exists some $j \in \mathcal{D}_t$ such that $(i,j) \in \mathcal{F}_t$. The cost of assigning a device j to a customer i is given by a cost functional, $c : \{\mathcal{C}_t^f\} \times \{\mathcal{I}_t\} \to \mathbb{R}$, defined as

$$c(i,j) = \begin{cases} f(z(i) - z(j)), & i \in \{\mathcal{C}_t^f\}, j \in \{\mathcal{I}_t\}, z(j) \geq z(i), \\ \bar{c}, & z(i) > z(j). \end{cases} \tag{4.1}$$

where f is some non-negative convex non-decreasing function, and $\bar{c} \geq 0$. Compactly,

$$c(i, j) = \max(f(z(i) - z(j)), \bar{c}).$$

Thus, f grows with the difference of the age between the device and the customer and the older the device and the younger the customer, the larger the cost of a mismatch. From a practical standpoint, the function $c(i, j)$ can be thought of as the expected *uncovered warranty cost*, i.e., the expected future cost of sending replacement device j to customer i. In this case, the cost $\bar{c}$ represents the expected future cost when device j has more time left in the OEM warranty than customer i has left in the customer warranty.

The function f captures the cost to the WSP if a customer has a replacement device that fails. If the WSP sends to consumer i a replacement device j with an activation date that is after the consumer's activation date, i.e. $z(j) \geq z(i)$, this means that the device is "younger" than the consumer and, if it eventually fails, the device will be remanufactured through the OEM warranty and the OEM will bare the cost of repair and refurbishment. Conversely, if the activation date of the replacement device is before the consumer's activation date, i.e. $z(j) < z(i)$, the device will be "older" than the consumer and there is a possibility that the device will fail while the consumer is still covered by the consumer warranty but the device is not covered by the OEM warranty. If this is the case, the retailer will still send a replacement device from its reverse logistics center to customer i; but the failed device j has no (or minimal) value. If the retailer wants to include device j in its inventory in the reverse logistics center, this is no longer "costless" and the retailer will have to pay to get device j repaired. If the retailer does not intend to return j to inventory, then it would have liked to have sold the device (after it was repaired by the OEM) into a side channel. However, this option is less attractive as the retailer has to pay to refurbish or repair device j.

As $z(i) - z(j)$ increases, the larger is the probability that, in the future, device j returns with an expired OEM warranty, and the expected cost to the WSP increases. Given estimates of failure rates of a device, the cost function f can represent the expected cost to the system that a device, which is still under consumer warranty but out of OEM warranty, will fail. As an example, let r be the marginal cost of refurbishing a device that is out of OEM warranty, i.e., the additional cost to the WSP of having to deal with a device out of OEM warranty. Then, if the customer warranty and the OEM warranty have the same length, we can set the cost f of sending device j to customer i, if $z(j) < z(i)$, to be

$$f(z(i) - z(j)) = r \cdot \Pr(\text{Device fails and failure time} > z(i) - z(j) \text{ periods }).$$

In practice, for the WSP the value r is around \$100, such that the expected cost of refurbishing a device can be significant. If, for example, the failure age is exponentially distributed, and both the customer warranty and the OEM warranty have length T_w, the cost would be

$$f(z(i) - z(j)) = r \cdot (e^{z(j)-z(i)} - e^{-T_w}).$$

Within this setting, our goal is to minimize the total assignment cost over a finite horizon corresponding to the life-cycle of the device. Our first step will be to formulate this problem as a deterministic optimization, in which all the failures, devices, and customers are deterministic and known for each period. We denote this optimization problem as the *clairvoyant* formulation of the assignment problem since there is no uncertainty and all information is known. We will use the clairvoyant formulation as a benchmark to which other assignment strategies will be compared. More specifically, we will analyze the worst case ratio between the cost obtained by an assignment strategy and the *clairvoyant cost* over all possible sales, failures and arrivals of devices in inventory. Thus, if there is no uncertainty in the device failures, in the new customers that enter the system, and in the arrivals of devices in inventory, the *clairvoyant formulation* will be the assignment problem below, where the binary decision variables are $\{y_{ij}(t)\}$, and $y_{ij}(t)$ will be 1 if device j is assigned to customer i at time t. We call this problem the *Clairvoyant Assignment Problem (CAP)*.

$$
\begin{aligned}
\text{minimize} \quad & \sum_{t=1}^{T} \sum_{i \in \mathcal{C}_t^f} \sum_{j \in \mathcal{I}_t} c(i,j) \cdot y_{ij}(t) \\
\text{s.t.} \quad & \sum_{i \in \mathcal{C}_t^f} y_{ij}(t) \leq 1, \quad \forall j \in \mathcal{I}_t, \forall t \in [1,T] \\
& \sum_{j \in \mathcal{I}_t} y_{ij}(t) = 1, \quad \forall i \in \mathcal{C}_t^f, \forall t \in [1,T] \\
& \mathcal{R}_t = \{(i,j) | y_{ij}(t) = 1\}, \quad \forall t \in [1,T] \\
& \mathcal{I}_{t+1} = \mathcal{I}_t \cup \mathcal{A}_t - \mathcal{O}_t \cup \mathcal{R}_t, \quad \forall t \in [1,T] \\
& \mathcal{E}_{t+1} = \mathcal{E}_t \cup \mathcal{E}_t^{\text{new}} \cup \mathcal{E}_t^{\text{rep}} - \mathcal{F}_t \cup \mathcal{E}_t^{\text{exp}}, \quad \forall t \in [1,T] \\
& y_{ik}(t) \in \{0,1\}, \quad \forall i \in \mathcal{C}_t^f, \forall j \in \mathcal{I}_t, \forall t \in [1,T].
\end{aligned}
\tag{CAP}
$$

Even though this is a large scale optimization problem with potentially millions of decision variables, with no uncertainty in the system this problem can be solved in polynomial time since it is essentially a large assignment problem that has a convex objective function. Algorithms such as the Hungarian Algorithm will provide polynomial time guarantees to this formulation. However, in practice, uncertainty is pervasive in this system. As discussed in the previous chapter, there is considerable uncertainty regarding the failure times of devices. In addition, there is uncertainty related to new customers entering the system, as well as the arrival times of refurbished devices in inventory from the OEM.

Note that we assume an inventory management policy that determines the amount of new devices that enter inventory and of refurbished devices that leave through side-sales channels. Formulating the assignment and inventory management problems as a single (potentially stochastic) optimization problem is challenging. For example, formulating an adequate cost functional would require estimating not only the cost and prices of devices, but also the financial cost of mismatched warranties. However, for the particular case where the mismatch cost function is one-sided, i.e., there is only a cost to the system if the customer that receives a replacement device has more time left in his customer warranty than the replacement device has left in the OEM warranty, any optimal inventory management policy will sell the oldest devices in inventory first. This also come from the fact that we assume that the price of a device in the side-sales market does not depend on the age of the device since once a device leaves the system through side-sales, it is no longer covered by the OEM warranty.

4.4 Assigning Devices to Customers: Stationary Distributions

In this section we will analyze different matching strategies assuming that the activation date of customers that need replacement devices is random. We assume that failure ages of devices are i.i.d. and given by a distribution F with support $[0, T_w]$, where T_w is the length of both the OEM and customer warranties. We also assume that time is discrete and that at least T_w periods have passed since the launch of the device. Thus, if a device that fails at time t, the activation time of this device and its customer will be a random variable $A = t - X$ where X is the failure age and

$$\Pr(A \leq \bar{t}) = \Pr(X \geq t - \bar{t}) = 1 - F(t - \bar{t} + 1).$$

We assume that the total number of device failures is constant in every period and given by n, and that there are $m \geq n$ devices in inventory at the reverse logistics facility in every period. We also assume that the lead time at the OEM is deterministic and equal to l periods, such that devices that fail in period t will be available in period $t + l$. Finally, for an assignment of device j to customer i, we assume that the cost c is linear in the amount of uncovered warranty time and $\bar{c} = 0$. Remembering that $z(i)$ the activation date of customer i and $z(j)$ is the activation date of customer j, we have:

$$c(i,j) = \begin{cases} z(i) - z(j), & i \in \{\mathcal{C}_t^f\}, j \in \{\mathcal{I}_t\}, z(j) \geq z(i), \\ 0, & z(i) > z(j). \end{cases} \tag{4.2}$$

More compactly, $c(i,j) = \max(z(i) - z(j), 0)$. Thus, if X_i is the failure age of the device of customer i, and X_j is the failure age of replacement device j, we have that $z(i) = t - X_i$ and $z(j) = t - X_j - \bar{t} - l$, where $\bar{t}$ is the amount of time device j has been in inventory, and l is the lead time at the OEM. Then, if customer i has a failure at time t and receives device j as a replacement, the expected assignment cost will be

$$E[c(i,j)] = \max(z(i) - z(j), 0) = E[\max(X_j + \bar{t} + l - X_i, 0)]$$

Studying this simple setting will allow us to identify what are the trade-offs between different policies and how they impact the distribution of device ages in inventory. We will look at three policies in the setting: (i) random assignment, which is what is currently being done at the reverse logistics facility; (ii) best assignment, which is a myopic strategy; and (iii) oldest-out-first, where we use first the oldest devices in inventory , matching them to the oldest customers.

4.4.1 Random Assignment

Due to the large volume of incoming refurbished devices and requests, our partner WSP does not sort devices nor customer requests. Although they might prioritize the use of refurbished devices in inventory instead of seed-stock received from the OEM (or vice-versa depending on the situation), the assignment of devices to customers within each class of devices is done in no particular order. In fact, refurbished devices received from the OEM are not sorted by age and are usually mixed and placed in large boxes or containers in inventory. In order to describe this scenario from a conceptual level, we will first look at the case where devices in inventory are randomly assigned to replacement requests. This setting is not entirely realistic, since we do not differentiate between seed-stock and refurbished devices but, as a simplified model, will serve to analyze the potential downsides of random assignment.

Random Assignment when $m = n$

In this section, we assume that the number of items in inventory is the same as the number of request, so that $m = n$. Assume that in some period t the customers that file warranty claims are matched uniformly at random with the devices in inventory. Namely, customers and devices are sampled randomly without replacement and matched. Since $m = n$, the inventory will be renewed in every period, and since the failure ages of all devices and customers are i.i.d., the distribution of device ages in inventory will be the same as incoming requests, albeit shifted by the lead time l. In this case, if X_i and X_j are two random failure ages distributed according to a continuous distribution F, the average cost of the assignment per device will be

$$E[c(i,j)] = \Pr(X_i \leq X_j + l)E[f(X_j + l - X_i)|X_i \leq X_j + l], \forall i \in \mathcal{C}_t^f, j \in \mathcal{I}_t.$$

If $l = 0$, for X and Y distributed according to F, we have that $\Pr(Y \geq X) = \Pr(Y < X) = 1/2$ and

$$\begin{aligned}
E[c(i,j)] &= \frac{1}{2}E[Y - X|Y \geq X] \\
&= \frac{1}{2}\left(\frac{1}{2}E[Y - X|Y \geq X] + \frac{1}{2}E[X - Y|X \geq Y]\right) \\
&= \frac{1}{2}\left(E[\max(Y - X, 0)] + E[\max(X - Y, 0)]\right) \\
&= \frac{1}{2}\left(E[\max(Y, X) - X] + E[\max(-Y, -X)] + X\right) \\
&= \frac{1}{2}\left(E[\max(Y, X)] - E[\min(Y, X)]\right).
\end{aligned}$$

The second equality above comes from the fact that X and Y are i.i.d. and the third equality comes from the fact that $E[\max(Y - X, 0)] = \Pr(Y \geq X)E[Y - X|Y \geq X]$. If, for example, the ages of requests are distributed uniformly between $[t, T + t]$ for some $T \geq 0$, then

$$E[c(i,j)] = \frac{1}{2} \cdot \frac{T}{3} = \frac{T}{6}.$$

If the warranty is 12 months in length, then customers that receive replacements will have, on average, 2 months of warranty coverage time during which the devices are not covered by the OEM warranty.

For a linear cost function as in Equation (4.2), we can bound the expected age difference caused by random matching for any discrete failure distribution with some support $[0, T_w]$. This is shown in the result below.

Proposition 9. For any two discrete i.i.d. random variables X and Y with distribution defined on a support $[0, T_w]$, we have that

$$E[\max(Y + l - X, 0)] \leq \frac{T}{4} + l.$$

Proof. First, note that $E[\max(Y + l - X, 0)] \leq E[\max(Y + l - X, l)] = E[\max(Y - X, 0)] + l$. Now, if $p_x = \Pr(X = x)$ and the joint distribution is $p_{xy} = \Pr(X = x, Y = y) = p_x p_y$, we have

$$E[\max(Y-X,0)] \leq \sum_{y=0}^{T}\sum_{x=0}^{T} p_{xy} \cdot \max(y-x,0).$$

Since $p_{x,y} = p_{y,x}$, the r.h.s. of the expression above can be maximized by solving

$$\begin{aligned}
\text{maximize} \quad & \sum_{y=0}^{T}\sum_{x=0}^{T} p_x p_y \max(y-x,0) \\
\text{s.t.} \quad & \sum_{x} p_x = 1, \\
& p_x \geq 0 \,\forall x, y.
\end{aligned}$$

The solution to this optimization problem is $p_0^* = p_T^* = 1/2$ and $p_x^* = 0$ for all other x leading to a cost of $T/4$.

To prove this, assume a feasible solution to this problem which we denote by $(\tilde{p}_0 \ldots \tilde{p}_T)$. If this solution has only one non-zero component, then the cost will be 0, lower than the proposed solution. Conversely, if $\{\tilde{p}_t\}$ has more than one non-zero component, let $p_{\max}$ be the largest component of $\{\tilde{p}_x\}$, and let $x_{\max}$ be the index of this component. Then,

$$\begin{aligned}
\sum_{y,x\in\mathcal{I}} p_x p_y \max(y-x,0) &\leq p_{\max} \sum_{x<x_{\max}} p_x \max(x_{\max}-x,0) \\
&\leq p_{\max} \sum_{x<x_{\max}} p_x \cdot T \\
&= p_{\max} \cdot (1-p_{\max}) \cdot T.
\end{aligned}$$

The last expression is maximized when $p_{\max} = 1/2$ and, in this case, we obtain

$$\sum_{y,x\in\mathcal{I}} p_x p_y \max(y-x,0) \leq \frac{T}{4}.$$

Since the cost of the solution $p_0^* = p_T^* = 1/2$ and $p_x^* = 0$ for all other x is $T/4$, we are done. □

This bound gives us a limit on how poorly a random assignment strategy can be in a steady-state system. If $T_w = 12$, in this setting, for the piecewise linear cost function, random assignment will lead to at most an average uncovered time of 3 months plus the lead time at the OEM. From a practical perspective, this bound can serve as a guide for practitioners to identify if it is worth investing in the equipment or software to implement more sophisticated matching/sorting strategies.

Random Assignment when $m \geq n$

In the multi-period setting we assume that when a customer purchases a device, the customer warranty and the OEM warranty start at the same time. Also, the failure age of all customers/devices is sampled from the same distribution. As mentioned in the beginning of this section, we assume that there is a constant volume of n requests per day, and that there are $m \geq n$ items in inventory. In addition, we assume the lead-time at the OEM is deterministic and equal to l periods.

In every period n devices with age randomly distributed between 0 and T_w will fail and arrive in inventory. There will also be $m-n$ devices in inventory carried over from previous periods. The age of a device $j \in \mathcal{I}_t$ will be the sum of two random variables: (i) the age of the device when it arrived in inventory; (ii) how long it stayed in inventory. As mentioned in the beginning of this section, we assume that all devices fail according to some c.d.f. F.

To model the time a device stays in inventory when using a random matching strategy, we assume that devices in inventory are chosen uniformly at random and matched to an incoming request. Since we assume that there are n requests per period and m devices in inventory, where $m \geq n$, each device in inventory will be chosen with probability n/m, and $m-n$ devices carried over from one period to another. Thus, on average, a device will stay $m/n-1$ periods in inventory. Thus, we can model the time a device stays in inventory as a Geometric random variable with support $(0, 1, \ldots)$ and parameter n/m.

The age of a device in inventory will then be given by the sum of two random variables Y and Z, where $Y \sim F$ is the age of devices that fail with customers, and Z is a geometric random variable with parameter n/m and support $(0, 1, \ldots)$. With this in hand, we can bound the expected assignment cost in each period, which is presented in the proposition below.

Proposition 10. For any failure time distribution F defined on a finite support $[0, T_w]$, if the age of each customer i that needs a replacement in a period t is a continuous random variable distributed according to F, and the ages of a device j in inventory is given by

$Y+l+Z$, where $Y \sim X$, and Z is is a geometric random variable with parameter n/m and support $(0,1,\ldots)$, then we have

$$E[c(i,j)] = E[\max(Y+l+Z-X,0)] \leq \frac{T_w}{4} + l + \frac{m}{n} - 1.$$

Proof. We have

$$\begin{aligned} E[c(i,j)] &= E[\max(Y+l+Z-X,0)] \\ &= E_Z[E_{X,Y}[\max(Y+l+z-X,0)|Z=z]]. \end{aligned}$$

For the inside expectation, we can use the same proof as in Proposition 9 and we have

$$E[c(i,j)] \leq E_Z[T_w/4 + l + Z] = T_w/4 + l + m/n - 1,$$

completing the proof. By construction, this bound is tight.

□

The term $\frac{m}{n} - 1$ can be interpreted as the safety stock kept in inventory as a fraction of the total incoming warranty claims per period. Also, in the worst case setting, the average mismatch for random assignment grows linearly with the safety stock .

Although a worst-case bound is useful to understand what would be the worst possible mismatch when using a random assignment policy, a lower-bound on the average cost can also help a practitioner decide if investing in a system that tries to improve the matching is worthwhile or not. The next proposition introduces a lower bound using Jensen's inequality.

Proposition 11. For any age distribution F defined on a finite support $[0, T_w]$, if the age of each customer i that needs a replacement in a period t is a random variable distributed according to F, and the ages of a device j in inventory is given by $Y+l+Z$, where $Y \sim X$, and Z is a geometric random variable with parameter n/m and support $(0,1,\ldots)$, then we have

$$E[c(i,j)] \geq l + E[Z] = l + \frac{m}{n} - 1.$$

Furthermore, there is a distribution that achieves this bound.

Proof. Since the function $\max(x, 0)$ is convex, we have

$$\begin{aligned} E[\max(Y + l + Z - X, 0)] &\geq \max(E[Y + l + Z - X], 0) \\ &= l + E[Z]. \end{aligned}$$

Where the last equality comes from the fact that X and Y are i.i.d. The bound is achieved if the distribution F has zero variance, i.e., all the probability mass is concentrated in a single point. □

With these results in hand, we now proceed to analyze the *Youngest-Out-First* assignment policy, where devices in inventory and requests are sorted by age and matched youngest to oldest.

4.4.2 Youngest-Out-First (Myopic Policy)

Another policy that can be used for assigning customers to devices is a *Youngest-Out-First* (YOF) policy. In this policy, we sort all the m devices in inventory and sort all the n customer requests by age, and assign the n youngest devices in inventory to the n customer replacement requests from youngest to oldest. This assignment policy is depicted in Figure 4-2. Assuming that the cost of a match is given by $c(i, j)$ for some device j and customer i, this matching strategy will be optimal if $m = n$. This is shown in the proposition below.

Proposition 12. If $m = n$, the youngest-out-first policy is optimal.

Proof. For contradiction, assume that there is an optimal matching of devices to warranty claims that has a total cost that is strictly less than the one generated by the youngest-out-first policy. Since the matching is different than the one generated by the youngest-out-first policy, there must be at least one "crossing", i.e., two devices in stock, i_1 and i_2, of ages $z(i_1)$ and $z(i_2)$, that are each matched to warranty claims from customers j_1 and j_2 of ages $z(j_1)$ and $z(j_2)$, respectively, and $z(i_1) \leq z(i_2)$ and $z(j_1) \geq z(j_2)$. The term "crossing" is used here because if we arranged all the requests and devices in stock by age, and then looked at the network formed by matching devices in stock to requests, we would have crossing edges in any solution other than the one generated by the policy. Since $c(i, j)$ is non-decreasing and convex in $z(i) - z(j)$, we can match i_1 with j_2 and i_2 with j_1, and potentially improve the cost. To show this, let $g(z(i) - z(j)) = \max(f(z(i) - z(j)), 0)$ and let $\Delta_j = z(j_1) - z(j_2)$.

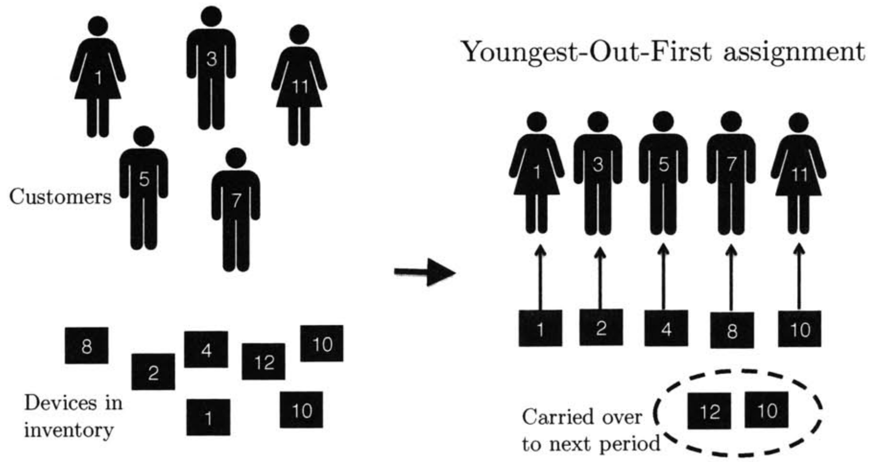

Figure 4-2: Depiction of the Youngest-Out-First assignment strategy. In this example, $m = 7$ and $n = 5$. Customers and devices are sorted by age and matched youngest to oldest. Unused devices in inventory are carried over to the next period and age. Note that the only assignment that occurs a cost is the assignment of the device of age 8 to the customer of age 7.

Then,

$$
\begin{aligned}
&c(i_1, j_1) + c(i_2, j_2) - c(i_1, j_2) - c(i_2, j_1) = \\
&\qquad = g(z(i_1) - z(j_1)) + g(z(i_2) - z(j_2)) - g(z(i_1) - z(j_2)) - g(z(i_2) - z(j_1)) \\
&\qquad = g(z(i_1) - z(j_1)) - g(z(i_1) - z(j_1) + \Delta_j) + g(z(i_2) - z(j_1) + \Delta_j) - g(z(i_2) - z(j_1)) \\
&\qquad \geq g'(z(i_1) - z(j_1))(-\Delta_j) + g'(z(i_2) - z(j_1))\Delta_j \\
&\qquad \geq 0.
\end{aligned}
$$

The first inequality comes from the fact that g is convex (if g were not differentiable, a subgradient can be used), and the last inequality comes from the fact that g is non-decreasing and convex, such that $g'(x) \geq g'(y)$ for any $x \geq y$.

We can repeat this procedure iteratively for all the "crossings" (and crossings that can be potentially added during the procedure), and eventually end in a solution where no crossings exist, which is exactly the one generated by the policy, which is a contradiction. Since inventory is renewed in each period, the proof is complete. □

In the single-period problem of assigning devices to customers, i.e., for the case where $T = 1$, the Youngest-Out-First policy will be optimal. This is proved in the next proposition:

Proposition 13. If $m \geq n$ and $T = 1$. the youngest-out-first policy is optimal.

Proof. First, if $m = n$, the optimality of the YOF policy is asserted from the previous proposition. Next, if $m > n$, i.e., for any subset of size n of the m devices in inventory that we choose, the YOF policy generates the best assignment between the n devices in the subset and the n customers.

Now, we can potentially reduce the assignment cost for any subset of n devices in inventory that are chosen to be matched to the n customer requests using the following procedure: (i) Choose the oldest device in the subset. (ii) If the oldest devices in the subset is *older* than the youngest device outside of this subset, swap the two devices. If it is not older, stop. (iii) With the new subset in hand, use the YOF policy to match devices in the subset to customers. (iv) Return to step (i) and repeat the procedure until the n youngest devices are in the subset. Note that step (ii) will potentially decrease the cost since the cost function is increasing in the age of devices used. Furthermore, step (iv) also potentially decreases the cost since the YOF policy is optimal when $m = n$. Since this procedure is valid for any initial subset of n devices that is chosen, the proof is complete. □

For the multiperiod case and $m > n$, we will present two results regarding the stationary age distribution of devices in inventory when the youngest-out-first policy is used. First, we will prove that as $t \to \infty$ there will be, with probability 1, at least $m - n$ devices in inventory with age greater than T_w, such that they will never be chosen by the youngest-out-first policy since the incoming devices in inventory from the OEM will be younger than the devices carried over from the previous period. This occurs because this policy uses the youngest devices in inventory first, such that the old devices that arrive in inventory tend to not be chosen and will age. Thus, there will be a time period $\bar{t}$, where for any $t > \bar{t}$ the policy will always choose the n devices that have just arrived in inventory to be used as replacements since the devices carried over from the previous period will be older than $T_w + l$, which is the maximum value of the age distribution.

Based on this result, we proceed to prove a distribution-free bound on the tail of the average mismatch between devices in inventory and incoming requests using the youngest-out-first policy. Namely, for any failure distribution, as n increases the average mismatch decays exponentially to zero.

We introduce some notation in order to capture the case where $m > n$. Let $(X_t^{(1)}, \ldots, X_t^{(n)})$

be the random vector representing the sorted ages of customers that file requests at time t. Thus, $X_t^{(i)}$ will be the i-th youngest customer. Also, let $(Y_t^{(1)}, \ldots, Y_t^{(m)})$ be the ages of devices in inventory at time t, after devices arrived from the OEM. We do not assume that $X_t^{(i)}$ and $Y_t^{(i)}$ are equally distributed, although we do assume that if a customer of age X requires a replacement device at time t, then at time $t + l$ there will be a device of age $X + l$ arriving in inventory. Finally, let $(\tilde{X}_t^{(1)} + l, \ldots, \tilde{X}_t^{(n)} + l)$ be the random vector representing the sorted ages of devices arriving from the OEM at time t. Note that the elements of $(\tilde{X}_t^{(1)} + l, \ldots, \tilde{X}_t^{(n)} + l)$ will be contained in the vector $(Y_t^{(1)}, \ldots, Y_t^{(m)})$. Also, note that $\tilde{X}_t^{(i)} = X_{t-l}^{(i)}$, since the lead time is deterministic.

If we match devices from youngest to oldest, and if $m > n$, we have that $\lim_{t \to \infty} Y_t^{(j)} > T_w + l, \forall j > n$ almost surely. This is proved in the proposition below.

Proposition 14. Using a youngest-out-first policy and assuming $m > n$, we have that $\lim_{t \to \infty} Y_t^{(j)} > T_w, \forall j > n$ almost surely. Furthermore, if $\bar{F}(t) = 1 - F(t)$, and for a lead time of l, then for any $t \geq \bar{t}$ we will have, with probability at least $1 - 1/n$ that there will be at least $m - n$ devices in inventory with age greater than T_w, where $\bar{t}$ is defined as

$$\bar{t} = \frac{2}{\bar{F}(t-l)} \cdot \left(\frac{m-n}{n} + \frac{\log n}{n} \right) + l$$

Proof. Since we assume that customers that require replacement devices send to the retailer broken devices of the same age as them, we have that $Y_t^{(n+1)} < T_w$ only if in the interval $[0, t - l]$ the total number of customers of age larger than $T_w - l$ is less than $m - n$. Since in each period the age of customers is independent, and letting $\bar{F}(t - l) = 1 - F(t - l)$ be the probability of an arrival having age larger than $t - l$, we have that, by the Borel-Cantelli lemma, as $t \to \infty$, we will observe an infinite number of devices of age larger than $T_w - l$.

Furthermore, the total number of arrivals of age larger than T_w up to time t is distributed according to a Binomial random variable with $n(t-l)$ trials and success probability $\bar{F}(t-l)$.

Remembering that the Chernoff bound for some binomial random with parameters p and n trials is

$$\Pr\left(\leq \text{k successes in n trials}\right) \leq \exp\left(-\frac{(np-k)^2}{2np}\right),$$

we obtain, in our case, that

$$\Pr(Y_t^{(n+1)} \geq T_w) \geq 1 - \Pr(\leq m-n \text{ successes in } n(t-l) \text{ trials})$$
$$\geq 1 - \exp\left(-\frac{(n\cdot(t-l)\cdot\bar{F}(t-l) - (m-n))^2}{2n(t-l)\bar{F}(t-l)}\right)$$

Thus, the time period $\bar{t}$ where the probability that at least $(m-n)$ devices in inventory are older than T_w is at least $1 - 1/n$ can be obtained by finding the smallest $\bar{t}$ that satisfies

$$\exp\left(-\frac{(n\cdot(t-l)\cdot\bar{F}(t-l) - (m-n))^2}{2n(t-l)\bar{F}(t-l)}\right) \geq \frac{1}{n}.$$

This inequality is satisfied for any

$$t \geq \frac{1}{\bar{F}(t-l)}\cdot\left(\frac{m-n}{n} + \frac{\log n}{n} + \sqrt{\frac{2\log n}{n}\frac{m-n}{n} + \frac{(\log n)^2}{n^2}}\right) + l$$

Using the triangle inequality and the fact that $\log n/n \leq 1/2, \forall n \geq 1$, we have

$$2\cdot\left(\frac{m-n}{n} + \frac{\log n}{n}\right) \geq \frac{m-n}{n} + \frac{\log n}{n} + \sqrt{\frac{2\log n}{n}\frac{m-n}{n} + \frac{(\log n)^2}{n^2}}$$

and we can set $\bar{t}$ as

$$\bar{t} = \frac{2}{\bar{F}(t-l)}\cdot\left(\frac{m-n}{n} + \frac{\log n}{n}\right) + l$$

such that for any $t \geq \bar{t}$, at least $(m-n)$ devices in inventory will have age of at least T_w with probability larger then $1 - 1/n$. □

With this result in hand, we can explicitly describe the limiting distribution of items in inventory under this policy. As $t \to \infty$, we have from the previous proposition, that for some $\tilde{t}$ we will have $Y_t^{(n+1)} \geq T_w$ for all $t \geq \tilde{t}$ and the youngest out first policy will only use the n incoming devices from the OEM as replacements. Also, we defined $\bar{t}$ such that for any $t > \bar{t}$ we will be in this regime, where only incoming refurbished devices are used, with probability at least $1 - 1/n$.

In order to obtain a distribution free bound on the average assignment cost in this case, let $\bar{F}_n^X$ be the empirical CDF for the ages of the n requests arriving in period t. More specifically,

$$F_n^X(t) = \sum_{i=1}^{n} \mathbf{1}(X^i \leq t).$$

where $\mathbf{1}$ is the indicator function. Additionally, let $\tilde{F}_n^{\tilde{X}}$ be the empirical CDF of the age distribution of ages of devices arriving from the OEM at time t, minus the lead time. If we assume that f is linear we can write the random empirical average uncovered time of devices at time t as

$$\frac{1}{n}\sum_{i=1}^{n} \max(\tilde{X}_t^{(i)} + l - X_t^{(i)}, 0).$$

A distribution-free bound for the tail of this cost is given in the proposition below.

Proposition 15. The tail of the distribution of $\frac{1}{n}\sum_{i=1}^{n} \max(\tilde{X}_t^{(i)} + l - X_t^{(i)}, 0)$ will be

$$\Pr\left(\frac{1}{n}\sum_{i=1}^{n} \max(\tilde{X}_t^{(i)} + l - X_t^{(i)}, 0) \geq \epsilon\right) \leq 2\exp\left(-2n\left(\frac{\epsilon - l}{T_w}\right)^2\right).$$

Proof. First, note that,

$$\frac{1}{n}\sum_{i=1}^{n} \max(\tilde{X}_t^{(i)} + l - X_t^{(i)}, 0) \leq \frac{1}{n}\sum_{i=1}^{n} \left|\tilde{X}_t^{(i)} - X_t^{(i)}\right| + l$$

Now, we leverage the results presented by Major (1978) and Levina and Bickel (2001) regarding Mallow's metric. Namely, as noted in Levina and Bickel (2001), we have

$$\frac{1}{n}\sum_{i=1}^{n} \left|\tilde{X}_t^{(i)} - X_t^{(i)}\right| = \frac{1}{n}\min_{(j_1,\ldots,j_n)} \sum_{i=1}^{n} \tilde{X}_t^{(i)} - X_t^{(j_i)}$$

where the minimum is over all permutations of $(1, \ldots, n)$. Also, as posed by Major (1978), we have that this is exactly that Mallow's distance and

$$\frac{1}{n}\min_{(j_1,\ldots,j_n)} \sum_{i=1}^{n} \tilde{X}_t^{(i)} - X_t^{(j_i)} = \int_0^1 |\bar{F}_{Y,n}^{-1}(s) - \bar{F}_{X,n}^{-1}(s)|ds + l.$$

Finally,

$$\begin{aligned}\int_0^1 |\bar{F}_{Y,n}^{-1}(s) - \bar{F}_{X,n}^{-1}(s)|ds + l. &= \int_0^{T_w} |\bar{F}_{Y,n}(s) - \bar{F}_{X,n}(s)|ds + l \\ &\leq T_w \sup_{s\in[0,T_w]} |\bar{F}_{Y,n}(s) - \bar{F}_{X,n}(s)| + l\end{aligned}$$

Thus, from the DvoretzkyKieferWolfowitz inequality, for any $\epsilon > l$, we have

$$\Pr\left(\frac{1}{n}\sum_{i=1}^{n}\max(Y_t^{(i)}+l-X_t^{(i)},0)\geq\epsilon\right)\leq 2\exp\left(-2n\left(\frac{\epsilon-l}{T_w}\right)^2\right).$$

□

In practice, since hundreds of devices of a given type are received per day, sorting all the devices will lead to a mismatch close to l. For example, if we choose $\epsilon = l + T_w/10$, the bound will be $2\exp(-0.02n)$. If we sort 200 devices per day, the probability of the average mismatch being larger than $0.1T_w + l$ is smaller that 4%. Note that the bound is only valid for $\epsilon > l$.

4.4.3 Oldest-Out-First

The Oldest-Out-First (OOF) policy is similar to the youngest-out-first, with the exception that the n oldest devices in inventory are sorted and matched to incoming requests instead of the n youngest devices. The OOF policy is described below.

Algorithm 2 Oldest Out First Policy

for all t **do**
  1) Sort warranty claims by age
  2) Sort devices in stock by age
  3) Match devices in stock to warranty claims from oldest to newest
end for

First, note that this policy will be ineffective if $T_w << m/n$, i.e., the support of the distribution of customer ages is much smaller than the inventory turnover rate. This is because devices age in inventory and since we are using the oldest devices first, we will be unable to clear inventory of old devices. On the other hand, if m is not too much larger than n, we will have a large inventory turn over, and no device will be "stuck" in inventory. The next section will be dedicated to an in-depth analysis of the OOF policy.

4.5 Oldest-Out-First Policy in Depth

In this section we will analyze the Oldest-Out-First (OOF) policy in detail and prove theoretical guarantees on its performance. Note that the random matching policy, the YOF policy, and the OOF policy can be interpreted as on-line algorithms since, at some time t, they only take into account the distribution of devices in inventory and the requests at time t. More specifically, at time t they only take into account the set of customers that require replacements, which we denoted by $\mathcal{C}_t^f$ and the current distribution of devices in inventory, which we denote by $\mathcal{I}_t$. We will determine the performance of the OOF policy by analyzing its *competitive ratio*. The definition of the competitive ratio is given below.

Definition 16. Let $A(\mathcal{C}_t^f, \mathcal{I}_t, \mathcal{A}_t)$ be the cost of an algorithm that, at time t, assigns devices to customers in the Clairvoyant Assignment Problem (CAP). Also, let $C^*(\{\mathcal{C}_t^f, \mathcal{I}_t, \mathcal{A}_t\})$ be the solution of the CAP where the customer requests in each period given by $\{\mathcal{C}_t^f\}$, inventory in the beginning of each period is $\mathcal{I}_t$ and the arrivals from the OEM in each period is $\{\mathcal{A}_t\}$. Then, this algorithm is said to have a competitive ratio α if

$$\max_{\{\mathcal{C}_t^f, \mathcal{I}_t, \mathcal{A}_t\}} \frac{\sum_t A(\mathcal{C}_t^f, \mathcal{I}_t, \mathcal{A}_t)}{C^*(\{\mathcal{C}_t^f, \mathcal{I}_t, \mathcal{A}_t\})} = \alpha.$$

The competitive ratio allows for a comparison between the *clairvoyant policy*, i.e., the policy that "knows" all the information of the system, and an on-line policy. In our case, the clairvoyant policy is the solution to the deterministic optimization problem C^*. Thus, we can interpret the competitive ratio as a game where, given an assignment policy, an adversary chooses a realization of failures and arrivals that maximizes the ratio of the cost of the assignment policy and the clairvoyant cost. Note that if we make $\bar{c} = 0$, it is possible that the competitive ratio is infinity.

We now proceed to formulate the device assignment problem as a transportation problem, and then we propose and analyze the *oldest-out-first* policy, a heuristic for matching devices in stock to customers introduced in the previous section.

4.5.1 The oldest-out-first policy

In order to analyze the OOF policy, we will formulate the assignment problem from the previous section as a closed loop transportation problem, reducing the dimension of the state by allowing the failure ages of customers (and, consequently, the ages of the devices that they hold) to be independent in each period. Note that this assumption is not entirely consistent with the original formulation of the CAP, where the same pool of customers is carried over from one period to another. In our competitive analysis, this assumption gives more "leverage" for the adversary to choose the age of incoming requests and of devices coming into inventory, giving an upper-bound on the competitive ratio for the OOF policy applied to the CAP. However, this assumptions allows for a simpler formulation and analysis of the OOF policy.

Let T_w be the maximum time length of a customer warranty, and let the warranty requests at time t be denoted by a vector $d(t) = (d_1(t), \dots, d_{T_w}(t))$ where $d_j(t)$ represents the number of requests from customers of age j, i.e., customers that are j time-periods into their customer warranty at time t. The warranty claims are satisfied with devices from stock. We will denote the devices in stock by a vector $x(t) = (x_1(t), \dots, x_S(t))$ where S is the age of the oldest devices in inventory, and $x_i(t)$ is the number of devices in stock of age i, i.e., devices that are i months into the OEM warranty. The refurbished items that arrive from the OEM at time t are denoted by the vector $a(t) = (a_1(t), \dots, a_S(t))$. Also, we let $\bar{d}(t) = \sum_k d_k(t)$, $\bar{x}(t) = \sum_k x_k(t)$, and $\bar{a}(t) = \sum_k a_k(t)$.

For this analysis, let l be the OEM lead time, and we assume that the maximum variation of the number of warranty claims from one period to the next is bounded by Δd, i.e.,

$$\left|\bar{d}(t+1) - \bar{d}(t)\right| \leq \Delta d.$$

Also, if we assume a lead time at the OEM of l periods, for any loss level at the OEM we have that

$$\bar{a}(t) - \bar{d}(t) \leq l \cdot \Delta d.$$

As mentioned in the previous section, since a significant fraction of refurbished devices used to fulfill warranty requests fail again, and since faulty devices that are out of the OEM warranty cannot be sent to the OEM for refurbishment, our goal is to match devices in stock to customer warranty requests in order to minimize the *total uncovered warranty cost* over a finite time horizon T. If a device of age j is sent to a customer of age $i < j$, there is a risk that this device will fail when it is out of OEM warranty, but still covered under customer warranty. In addition, as the difference $j - i$ increases, the risk that this device will fail out of OEM warranty increases and, therefore, the expected cost to the vendor also increases. To capture this, we assume a cost function with the same structure as the cost in Equation 4.1, such that the cost of using a device in stock of age j to satisfy a warranty claim from a customer of age i is given by a function $c(i, j)$ that is convex and non-decreasing in the difference $j - i$. As before, the cost $c(i, j)$ can be thought of as the expected future cost for the vendor of sending a device of age j to a customer of age i. The maximum cost will be denoted by c_{max}, while the minimum cost, i.e., when $i \geq j$, will be denoted by $c_{\min}$. Thus, we assume that $c(i,j) = c_{\min}, \forall i > j$.

Customers that file a warranty claim send their broken devices to the vendor and, if the device is still covered under OEM warranty, the vendor sends the faulty device to the OEM for refurbishment. In this section, we assume that there is always enough inventory to satisfy all the warranty claims. More specifically, this means that the WSP is willing to utilize new devices to guarantee that there is no backlog and, therefore, $\bar{x}(t)+\bar{a}(t) \geq \bar{d}(t), \forall t$. Denoting by $y_{i,j}(t)$ the number of devices of age j that are used to satisfy customers of age i, the *clairvoyant allocation*, i.e., the allocation if all the arrivals are known ahead of time is the solution to the following optimization problem:

$$
\begin{aligned}
\text{minimize} \quad & \sum_{t=1}^{T} \sum_{j=1}^{S} \sum_{i=1}^{W} \tilde{c}(i,j) y_{i,j}(t) \\
\text{subject to} \quad & \sum_{j} y_{ij}(t) = d_j(t), \; i = 1, \ldots, W \\
& x_{j+1}(t+1) = x_j(t) + a_j(t) - \sum_{i} y_{ij}(t), \; j = 1, \ldots, S, \;\; t = 1, \ldots, T \\
& y_{ij}(t) \geq 0, \forall i, j, t. \\
& x_j(t) \geq 0, \forall i, t.
\end{aligned}
$$

We will now analyze the performance of the OOF policy as an on-line policy for this problem. The intuition behind this policy is that, since devices in inventory are aging and losing "OEM warranty months", if we use oldest devices first, over time we will not have devices in inventory that are "too old". Additionally, the long-term benefit of using "old" devices first outweighs the short term cost induced by potentially sending "old" devices to "young" customers. If we used the optimal myopic strategy, which is to use the youngest devices first, we could potentially end up with leftover old devices in inventory that continue aging as time passes, eventually expiring their OEM warranty.

For the single period problem, i.e., when $T = 1$, the OOF policy has the following property:

Proposition 17. If T $= 1$ and $\bar{x}(1)+\bar{a}(1) = \bar{d}(1)$, then the oldest out first policy is optimal.

Proof. If $\bar{x}(1) + \bar{a}(1) = \bar{d}(1)$, then the number of incoming requests is the same as the number of devices in inventory, such that there will be no leftover devices at the end of the period. In this case, the OOF policy is the same as the Youngest-Out-First policy and optimality follows directly from Proposition 12. $\square$

The competitive ratio of the OOF policy for a single period is described in the proposition below. Here, we assume that there is enough inventory to satisfy all the warranty claims.

Proposition 18. For the single-period problem, let x, d, and a be the current inventory, the current demand, and the replacements that arrive in the beginning of the period. Also, assume that $\sum_j x_j + \sum_j a_j - \sum_j d_j \leq \Delta x$ and that $\sum_j d_j = n$. Then, if $\mathcal{U} = \{x, a, d | \sum_j x_j + \sum_j a_j - \sum_j d_j \leq \Delta x, \sum_j d_j = n\}$ the oldest out first policy has a competitive ratio of

$$\max_{\{x,a,d\} \in \mathcal{U}} \frac{\text{Cost of OOF given } \{x, a, d\}}{\text{Clairvoyant cost given } \{x, a, d\}} \leq \frac{\alpha c_{\min} + (1 - \alpha) c_{\max}}{c_{\min}},$$

where $\alpha = \frac{(n - \Delta x)^+}{n}$. Furthermore, this bound is tight.

Proof. The proof is in the appendix of this chapter. $\square$

For the multi-period case, the competitive ratio of the OOF policy is given in the next proposition. In order to obtain this ratio, we make two assumptions: (i) that the demand is fairly smooth and does not change by more than Δd between two periods (i.e. a Lipschitz continuity assumption); and (ii) the total number of devices in inventory in a period is within a factor of m of the total demand in that same period. In practice, assuming a hazard rate model for device failures, if sales are smooth (which is often the case), demand for replacements will be fairly smooth. Similarly, the number of devices in inventory will not become arbitrarily large since excess inventory can be sold. This ratio is tight for cases where devices returning from the OEM into stock can have any age, but it is loose if we assume that devices returning from the OEM are warranty requests from previous periods that aged according to the lead time.

Proposition 19. Assume that

- $\sum_i d_i(1) = n$

- $\{d(t)\} \in \mathcal{U} = \{\{d(t)\} \big| |\sum_j d_j(t+1) - \sum_j d_j(t)| \leq \Delta d, \forall t\}$

- For some $m \geq 0$, $\{x(t), d(t)\} \in \mathcal{S} = \{\{x(t), d(t)\} \big| 0 \leq \sum_i x_i(t) \leq m \cdot d_i(t)\}$

Then, the oldest out first policy has a competitive ratio that satisfies the inequality

$$\max_{\{x(t),d(t)\}\mathcal{U}\cap\mathcal{S}} \frac{\text{Cost of OOF given } \{x(t), d(t)\}}{\text{Clairvoyant cost given } \{x(t), d(t)\}} \leq 1 + \frac{c_{\max} - c_{\min}}{c_{\min}} \cdot \left(m + \frac{2\Delta d \cdot l}{n}\right).$$

which does not depend on the horizon T.

Proof. The proof is in the appendix of this chapter. □

As an example of this ratio, let us assume that time is measured in weeks, and that $\Delta d = 0.05n$, $c_{\max} = 3c_{\min}$, $m = 0.2$, and $l = 3$ weeks. Then, the bound would be

$$1 + \frac{c_{\max} - c_{\min}}{c_{\min}} \cdot \left(m + \frac{2\Delta d \cdot l}{n}\right) = 2$$

In addition, as expected, the bound increases with lead time and demand variation ratio $\Delta d/n$. Also, it is worth noting that this bound does not depend on the horizon, so the OOF policy cannot perform arbitrarily poorly. In the next section we will show, through numerical experiments, the performance of the OOF policy. In practice, the OOF policy mitigates significantly the number of uncovered months, especially in comparison to random matching and to the myopic policy.

4.6 Numerical Experiments

In the previous sections we examined three different policies for assigning devices in inventory to customers: (i) random assignment; (ii) the YOF policy and; (iii) the OOF policy. We proved theoretical properties of these policies and also derived bounds for their performance. In this section, we will compare and contrast these three approaches through numerical experiments. In addition, at the end of this section, we will simulate the practical performance of these policies using real-world data from our partner WSP.

Our simulations were done using an extended version of the simulator developed for the numerical experiments in Chapter 3. More specifically, the simulator described in Section 3.7 of Chapter 3 was adapted to include the customer-device assignment strategy as an additional input parameter and to allow for devices to fail multiple times. For the simulations in the first subsection, we used the inventory management policy based on the Certainty-Equivalent approximation described in Section 3.6.1 of the previous chapter. Since we allow devices to fail multiple times, the average demand and arrival paths were estimated in every time period through a Monte-Carlo simulation using 50 samples. For the simulations in the second subsection, we used the clairvoyant policy, where the exact dates of device failures and arrivals from the OEM are assumed to be known.

For the simulations in this section, one discrete period corresponds to one day. We also assume that when a customer purchases a new device, both the OEM and customer warranties start simultaneously. However, once a device fails and is sent to the OEM for repair/refurbishment, the device keeps aging and "consuming" OEM warranty.

We will use two metrics for measuring the performance of the different assignment policies:

- *Average uncovered time per replacement device shipped*: if a refurbished device of age j is sent as a replacement to a customer of age i, the *uncovered time* of the customer will be $\max(j-i,0)$. Since we assume that both the customer warranty and the OEM warranty have the same length, this represents the amount of time that a customer is still covered by the customer warranty while the device he/she owns is not covered by the OEM warranty. Since devices can fail multiple times, this is a measure of exposure of the WSP with regards to out-of-warranty returns;

- *Percentage of failures that are out-of-warranty*: the percentage of all the failures that happened when the customer was covered by the customer warranty but the device that failed was not covered by the OEM warranty. In this case, the device is either scrapped or the WSP has to pay for its repair/refurbishment.

In the remainder of this section we will look at two different sets of numerical experiments. In the first set of experiments we extend the simulator developed in the previous chapter to allow for the assignment policy to be an input to the simulation. We then simulate the life-cycle of devices using different failure distributions and compare the performance of the different policies. The second set of experiments uses real-world data from our partner WSP. We use data from 3 devices to contrast the performance of the different assignment strategies.

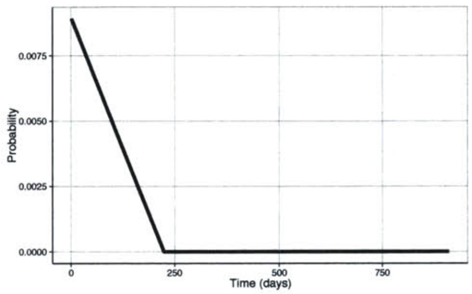

(a) Sales period distribution.

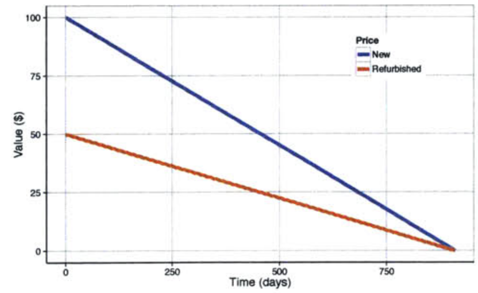

(b) Prices of new and refurbished devices

Figure 4-3: Parameters of the simulation. The failure distribution is exponential with mean 364 days. Sales only occur during the first 6 months from launch.

4.6.1 Comparing different policies

We assume a total simulation horizon of 2.5 years and that both the customer and the OEM warranty have a length of 12 months. Thus, if a customer has a device that fails and he is out of customer warranty, it will not be replaced. We also assume that devices can fail multiple times, such that a replacement device sent to a customer can fail again. In addition, we assume that new devices are sold according to the probability distribution in 4-3a such that each customer samples its purchasing date from this distribution. The prices of new and refurbished devices will, unless explicitly stated, follow the pattern depicted in Figure 4-3b, such that the price of both new and refurbished devices decrease linearly over time. We assume a 3 week total lead time from when a device fails until it returns to inventory (this includes both the customer and OEM lead times), and a 20% probability that a device sent to the OEM cannot be repaired. We set the seed-stock level to be 5% of devices sold.

For our first experiment, we assume that the failure age of devices are exponentially distributed with average 12 months. Thus, the hazard rate of devices is constant and approximately 63% of newly purchased devices will fail under warranty. We also assume that 6000 devices are sold, and that the sales period of each new device sold is sampled from the distribution in 3-5a. Although the volume of devices sold that we use is small compared to real-world volumes (usually one or two orders of magnitude higher), this smaller scale allows us to do a Monte-Carlo simulation of the life-cycle of the device. In this experiment, we simulate the life-cycle of the device 100 times.

The results are depicted in Figure 4-4. The random assignment strategy leads to an

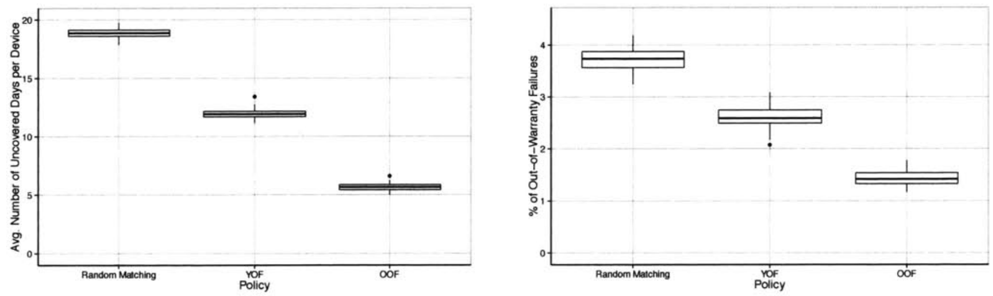

(a) Avg. uncovered time per replacement shipped.

(b) Out-of-warranty failure as a percentage of total failures.

Figure 4-4: Box-plots for the average uncovered days per replacement device shipped and the percentage of failures that are out-of-warranty when assuming that the failure age is an Exponential random variable with average 365 days. The data was generated from 100 simulations of the life-cycle of the device.

average of 18 days of uncovered time per device shipped, the YOF policy generates about 12 days, and the OOF policy only around 6 days, as shown in Figure 4-4a. This happens because, over time, the OOF policy does not let old devices accumulate in inventory and even if the policy is not optimal in a single period, it has a better performance than the myopic policy in this multi-period setting. Furthermore, on average the number of out-of-warranty failures in this experiment was about 1.4% of the total volume of failures, while the YOF and the random matching policies were, respectively, 2.7% and 3.7%. This is illustrated in Figure 4-4b. Since, in practice, the WSP deals with tens-of-thousands of failures per day, the simulation indicates that changing the policy used to assign devices to customers could lead to significant cost savings since, by using the OOF policy, the WSP could refurbish/repair through the OEM warranty more devices than if a random assignment strategy were used. These devices could then either be used as replacement or sold through the side-sales channel.

For the second experiment, we assume that the failure age of devices follow a log-normal distribution. This distribution has a high variance and is "heavy tailed". The results are depicted in Figure 4-5. The OOF policy will still have a superior performance in both metrics when compared to the YOF policy and the random matching policy. In this case, the random assignment policy has an average of 2.4% of devices being out-of-warranty failures, while for the OOF policy this decreases to an average of 1.25%.

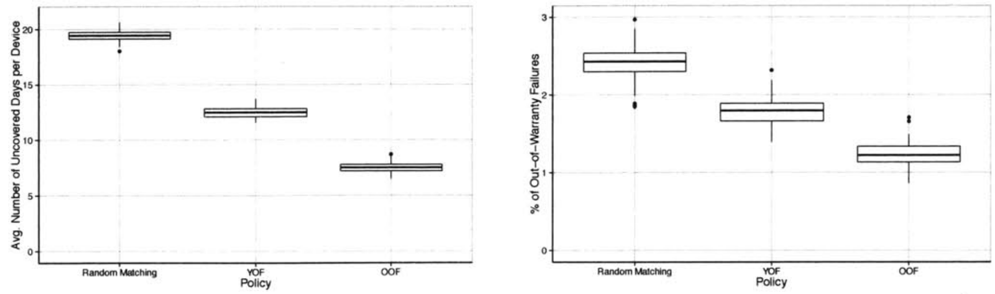

(a) Avg. uncovered time per replacement shipped.

(b) Out-of-warranty failure as a percentage of total failures.

Figure 4-5: Box-plots for the average uncovered days per replacement device shipped and the percentage of failures that are out-of-warranty when assuming that the failure age is a Log-Normal random variable with parameters $\mu = \log(365)$ and $\sigma^2 = \log(120)$. The data was generated from 100 simulations of the life-cycle of the device.

4.6.2 Experiment using real-world data from the WSP

Using data for two different devices sold by the WSP we compare, under a certainty-equivalent inventory management policy, the performance of Random Assignment policy, the YOF policy and the OOF policy. These two devices are from different manufacturers, and use different operating systems. As in the previous subsection, when there are side-sales, we assume that the oldest devices in inventory are sold first.

The data collected by the WSP contains individual sales and failure dates (if the device fails) for every device sold. The sales and the hazard rate distributions for these two devices are depicted in Figure 4-6. We did not have direct access to data from the OEM regarding repair rate and loss. However, the managers at the WSP's reverse logistics measured that the average aggregate customer and OEM lead-time is about 4 weeks. Furthermore, they estimate that 20% of replacement requests received cannot be repaired. We used both of these parameters in our simulation. As before, we assume a customer and OEM warranty lengths of 12 months. We also assume that the cost of sourcing a new device and the price of a refurbished device in a side-sales market behave according to Figure 4-3b. Finally, the contract between the WSP and the OEMs for these devices sets the seed-stock level to be 1% of sales.

Our experiment consisted of a Monte-Carlo simulation where, for each device, we simulated 30 life-cycles of sales and failures, assuming that 15,000 devices were sold per life-cycle. The sales date of each device was sampled from the actual sales dates. The failure dates (if

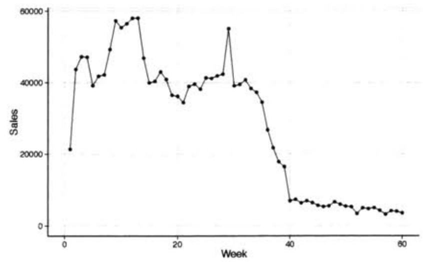

(a) Sales of device A

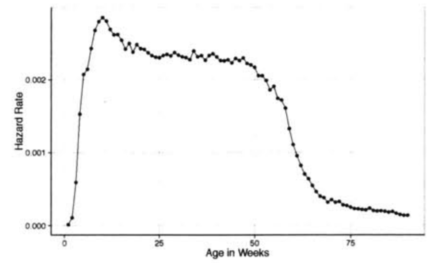

(b) Hazard rate of device A.

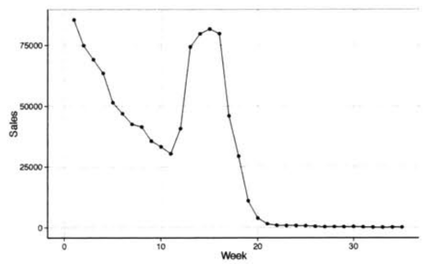

(c) Sales of device B

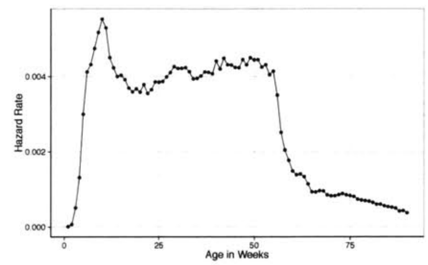

(d) Hazard rate of device B.

Figure 4-6: Sales and hazard rate distributions of two devices sold by the WSP.

a device failed at all) were drawn from the failure-time distribution. As a metric of performance, we assume the *average uncovered time per replacement device shipped.* Finally, we assume that a device can fail more than once throughout its life-cycle, and that refurbished devices have the same failure distribution as new devices.

The outputs of the simulations are summarized in Figure 4-7. Note that the OOF significantly decreases the average number of uncovered weeks with respect to random matching, and that it performs better than myopic matching since "old devices" do not accumulate in stock over time. Also, note that both the OOF and the myopic policy present improvements over random matching due to the *power of sorting* requests and devices. The difference between the performance for both devices can be explained by the different sales patterns of the two devices, and the fact that a larger fraction of devices type B fail than devices of type A.

4.7 Conclusion

In this chapter, we modeled and examined the problem of matching devices to customers in a reverse logistics system when there are two warranties in place: (i) the consumer warranty (offered by a WSP to its consumers), and the (ii) OEM warranty (offered by the OEM to the

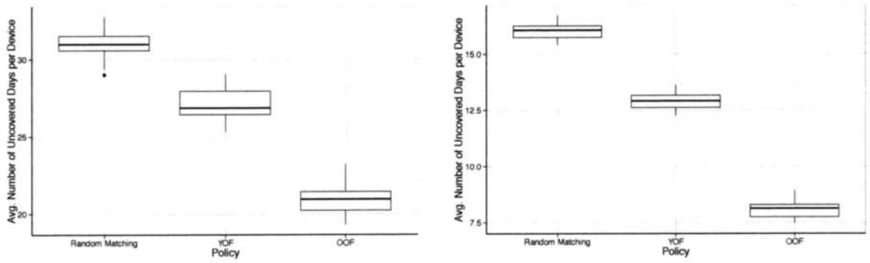

(a) Avg. uncovered days for device A. (b) Avg. uncovered days for device B.

Figure 4-7: Box-plots for the average uncovered days per replacement device shipped using data from the WSP. The data was generated from 30 simulations of the life-cycle of the devices.

WSP). Ideally the two warranties would be matched, i.e., the customer would have the same time left in his consumer warranty as the device would have left in the OEM warranty. A mismatch between these warranties incurs costs to the retailer beyond the usual processing costs of warranty requests. Namely, this extra-cost is incurred when a customer still covered by the consumer warranty has a device that fails, and this device is not covered by the OEM warranty.

Given this setting, we analyzed different assignment strategies and how they impact mismatch costs and out-of-OEM-warranty returns. Namely, the three assignment strategies that we focused on were:

- The *Random assignment policy*: where devices in inventory are randomly assigned to customers that require replacement devices, ignoring the time remaining in both the customer and device warranties. This was the policy used by the WSP at the time that our collaboration began;
- The *Youngest-Out-First (Myopic) policy*: where in every time period devices in inventory are assigned to customers as to minimize the mismatch in that specific period. We prove that this is the optimal single-period assignment strategy in our formulation;
- The *Oldest-Out-First policy*: a policy that always assigns the oldest devices in inventory to the oldest customers that require replacements.

Our first analysis of these policies involved assuming that the activation date of customers that need replacements devices is random and that failure ages of devices are i.i.d. Also, we assumed that the total number of device failures is constant in every period and given by n, and that there are $m \geq n$ devices in inventory at the reverse logistics facility in every period. In this context, we proved distribution-free upper and lower bounds on the expected mismatch cost for the random assignment policy. These bounds have a practical interpretation and can help a plant manager decide if it is worth investing in a matching policy other than random assignment.

We then considered the Youngest-Out-First policy, where customers and devices are sorted by age and matched from youngest to oldest. We proved that, in the long-run and for a lead time of l, the mismatch cost of the YOF policy will be very close to l. In fact, the distribution of the mismatch cost has an exponentially decreasing tail. However, we argued that this policy has a major drawback. If we allow m and n to be random variables that change over time, the YOF policy may lead to an accumulation of "old" devices in the system, since it uses the youngest first, and when n fluctuates, devices that are out of OEM warranty might be sent to customers.

We addressed this through the Oldest-Out-First policy. This policy also sorts devices and customers by age but, instead of matching them from youngest to oldest, matched them from oldest to youngest. The intuition behind this policy is that in the long-run it is not worth allowing devices to "age" in inventory, even though using them immediately is not the optimal short-term thing to do. By assuming certain conditions for the system's behavior, we proved the competitive ratio of this policy.

We also evaluated these policies through numerical experiments that use data from our partner WSP and also through simulations scenarios where failures are chosen from a pre-set distribution. In our experiments, we observe the OOF significantly decreases the average number of uncovered weeks with respect to random matching, and that it performs better than myopic matching since "old devices" do not accumulate in stock over time. We also observe that both the OOF and the myopic policy present vast improvements over random matching due to the *power of sorting* requests and devices.

Finally there are a few promising research directions that we have yet to explore. First, examining the optimal policy for the matching problem could lead to effective on-line matching strategies with provable guarantees. Another unexplored area is the intersection between inventory management and warranty matching. Finally, an analysis of strategies that only require partial sorting of devices and customers may lead to assignment policies that are simple to implement in practice and have a good performance.

4.8 Appendix: Proofs

Proof of Proposition 18

Proposition. For the single-period problem, let x, d, and a be the current inventory, the current demand, and the replacements that arrive in the beginning of the period. Also, assume that $\sum_j x_j + \sum_j a_j - \sum_j d_j \leq \Delta x$ and that $\sum_j d_j = n$. Then, if $\mathcal{U} = \{x, a, d | \sum_j x_j + \sum_j a_j - \sum_j d_j \leq \Delta x, \sum_j d_j = n\}$ the oldest out first policy has a competitive ratio of

$$\max_{\{x,a,d\} \in \mathcal{U}} \frac{\text{Cost of OOF given } \{x, a, d\}}{\text{Clairvoyant cost given } \{x, a, d\}} \leq \frac{\alpha c_{\min} + (1 - \alpha) c_{\max}}{c_{\min}},$$

where $\alpha = \frac{(n - \Delta x)^+}{n}$. Furthermore, this bound is tight.

Proof. The proof will be done in two steps. First, we will explicitly construct an example that achieves the competitive ratio and we will then show by contradiction that this is in fact the largest possible competitive ratio.

Assume we have $\bar{x} + \bar{a} = n + \Delta x$ units in inventory and that $\bar{d} = n$. Also, we assume without loss of generality, that $a_i = 0, \forall i$, i.e., no items arrive (or the items that would arrive are already in inventory).

Assume that the age distribution of items in inventory is $x = (n, 0, \ldots, 0, \Delta x)$, i.e., the only non-zero components of the inventory age distribution are $x_1 = d$ and $x_S = \Delta x$. Also, assume that all the demand for replacement devices is given by $d = (n, 0, \ldots, 0)$. In this case, the optimal single-period policy would allocate all the new devices to satisfy the warranty requests, obtaining a total cost of $c_{\min} n$, while the OOF policy would use the oldest devices first, and have a cost of $c_{\max} \cdot \min(n, \Delta x) + c_{\min}(n - \Delta x)^+$. Note that we can rewrite $\min(n, \Delta x)$ as

$$\min(n, \Delta x) = n - (n - \Delta x)^+.$$

Thus, the cost can be rewritten as

$$c_{\max} \cdot \min(n, \Delta x) + c_{\min}(n - \Delta x)^+ = c_{\max} \cdot \left(n - (n - \Delta x)^+\right) + c_{\min}(n - \Delta x)^+.$$

The ratio between these two costs is

$$\frac{c_{\max} \cdot \left(n - (n - \Delta x)^+\right) + c_{\min}(n - \Delta x)^+}{c_{\min} n} = \frac{c_{\max} \cdot \left(1 - \frac{(n - \Delta x)^+}{n}\right) + c_{\min} \frac{(n - \Delta x)^+}{n}}{c_{\min}}.$$

This achieves the competitive ratio.

Now, for contradiction, assume that there is another combination of customer and device ages that achieves a strictly larger competitive ratio than the distribution stated before. Also, in accordance to the beginning of the chapter, let the set of customer be given $\mathcal{C}^f$ and the set of devices in inventory be given by $\mathcal{I}$. NIt must be that $\text{card}(\mathcal{I}) = \text{card}(\mathcal{C}^f) + \Delta x$, or else we can increase the number of very old devices in stock and potentially increase the cost of the oldest out first policy while potentially decreasing the optimal cost (since there are optimal policy will have more options).

Consider the allocation generated by both the OOF and the optimal policy. With bot allocations in hand, let $\mathcal{I}_{OOF}$ be the set of devices used only by the OOF policy and not by the optimal allocation, let $\mathcal{I}_{OPT}$ be the set of devices devices used only by the optimal allocation and not by the OOF policy, let $\mathcal{I}_{Both}$ bet the devices used by both, and let $\mathcal{I}_{None}$ be the devices that neither of the allocations use. Note that all these sets are disjoints and their union is $\mathcal{I}$. Also, note that $\text{card}(\mathcal{I}_{OOF}) = \text{card}(\mathcal{I}_{OPT})$.

Finally, let $\mathcal{C}^f_{OPT}$ be the set of warranty requests satisfied using devices from $\mathcal{I}_{OPT}$ and let $\overline{\mathcal{C}^f}_{OPT}$ be the complement of this set. The set with all warranty requests is $\mathcal{C}^f = \mathcal{C}^f_{OPT} \cup \overline{\mathcal{C}^f}_{OPT}$.

With these sets in hand, we can increase the competitive ratio through the following procedure:

1. Shift the devices in $\mathcal{I}_{OOF}$ to the oldest age slot, obtaining a competitive ratio that is at least as large as the original stock and warranty claims. This does not change the cost of the optimal allocation;

2. For this new configuration, relocate the units of inventory in $\mathcal{I}_{OPT}$ to the *youngest* time slot. Note that these devices will continue to not be used by the OOF policy, and the optimal cost with this new configuration will potentially be smaller.

3. Reallocate the warranty claims in $\mathcal{C}^f_{OPT}$ to the youngest arrival slot. This will potentially increase the cost of the OOF policy, while not changing the optimal cost, since these devices can still be satisfied with young devices from $\mathcal{I}_{OPT}$. Thus, the competitive ratio of this new allocation will be at least the same as before.

4. If $\mathcal{I}_{Both} = \emptyset$ we are done, since the new allocation has the competitive ratio from the statement of the proposition - in this case it will simply be $\frac{c_{\max}}{c_{\min}}$.

5. In this new setting, if $\mathcal{I}_{Both} \neq \emptyset$, consider a modified policy of the OOF, where the devices in the set $\mathcal{I}_{OOF}$ are sent to the customers in $\mathcal{C}^f_{OPT}$ and the remaining devices are matched to the customers in $\overline{\mathcal{C}^f}_{OPT}$ using the OOF policy (or the youngest-out-first, since the number of remaining devices is the same). Note that this increases overall cost since, if we consider only the devices in $\mathcal{I}_{Both} \cup \mathcal{I}_{OOF}$ to satisfy the n requests, the OOF first assignment will be the same as the youngest-out-first, which is optimal according to Proposition 12. Therefore, using the devices in $\mathcal{I}_{OOF}$ to satisfy the customers in $\mathcal{C}^f_{OPT}$ increases cost.

With the allocation in the last step of the above procedure, note that the allocation of devices of the optimal policy from $\mathcal{I}_{Both}$ to $\overline{\mathcal{C}^f}_{OPT}$ will be the same as the allocation of the rearranged OOF policy, since the number of devices in both sets is the same and will have a total cost that we denote by β. Also, the number of devices in $\mathcal{I}_{OOF}$ and $\mathcal{I}_{OPT}$ will be Δx. Thus, this new competitive ratio can be bounded by:

$$\frac{\text{Cost of OOF policy}}{\text{Optimal Cost}} \leq \frac{\text{Cost of modified allocation policy}}{\text{Optimal Cost}} \leq \frac{c_{\max}\Delta x + \beta}{c_{\min}\Delta x + \beta}.$$

Since $c_{\max} \geq c_{\min}$, the ratio is non-increasing in β. Since the minimum allocation cost for the remaining devices is $(n - \Delta x)c_{\min}$, we obtain

$$\frac{\text{Cost of modified allocation policy}}{\text{Optimal Cost}} \leq \frac{c_{\max}\Delta x + (n - \Delta x)c_{\min}}{c_{\min}\, n},$$

which is exactly the competitive ratio we had before. Hence, we achieve a contradiction and the proof is complete. □

Proof of Proposition 19

Proposition. Assume that

- $\sum_j d_j(1) = n$
- $\{d(t)\} \in \mathcal{U} = \big\{\{d(t)\}\big| |\sum_j d_j(t+1) - \sum_j d_j(t)| \leq \Delta d, \forall t\big\}$
- For some $m \geq 0$, $\{x(t), d(t)\} \in \mathcal{S} = \big\{\{x(t), d(t)\}\big| 0 \leq \sum_i x_j(t) \leq m \cdot d_j(t)\big\}$

Then, the oldest out first policy has a competitive ratio that satisfies the inequality

$$\max_{\{x(t),d(t)\}\mathcal{U}\cap\mathcal{S}} \frac{\text{Cost of OOF given } \{x(t),d(t)\}}{\text{Clairvoyant cost given } \{x(t),d(t)\}} \leq 1 + \frac{c_{\max} - c_{\min}}{c_{\min}} \cdot \left(m + \frac{2\Delta d \cdot l}{n} \right).$$

which does not depend on the horizon T.

Proof. Assume some $\{x(t), d(t)\}$. If we uncouple the ages of devices in stock between periods, and defining the variation of inventory between two periods by $\Delta x(t) = (x(t) + d(t-l) - d(t))^+$, we have from the previous proposition that

$$\begin{aligned}
\max_{\{x(t),d(t)\}\mathcal{U}\cap\mathcal{S}} \frac{\text{Cost of OOF given } \{x(t),d(t)\}}{\text{Clairvoyant cost given } \{x(t),d(t)\}} &\leq \frac{\sum_{t=l}^{T} c_{\max} d(t) - (c_{\max} - c_{\min})(d(t) - \Delta x(t))^+}{\sum_{t=1}^{T} c_{\min} d(t)} \\
&\leq \frac{\sum_{t=l}^{T} c_{\min} d(t) + (c_{\max} - c_{\min})\Delta x(t)}{\sum_{t=1}^{T} c_{\min} d(t)} \\
&= 1 + \frac{c_{\max} - c_{\min}}{c_{\min}} \cdot \frac{\sum_{t=1}^{T} \Delta x(t)}{\sum_{t=1}^{T} d(t)}
\end{aligned}$$

From the assumption that $x(t) \leq m \cdot d(t)$ for some m and since $(x(t) + d(t-l) - d(t))^+ \leq x(t) + (d(t-l) - d(t))^+$, we obtain

$$\max_{\{x(t),d(t)\}\mathcal{U}\cap\mathcal{S}} \frac{\text{Cost of OOF given } \{x(t),d(t)\}}{\text{Clairvoyant cost given } \{x(t),d(t)\}} \leq 1 + \frac{c_{\max} - c_{\min}}{c_{\min}} \left(m + \frac{(d(t-l) - d(t))^+}{\sum_{t=1}^{T} d(t)} \right)$$

Once again using the triangle inequality, note that

$$\frac{\sum_{t=0}^{T} \max(d(t-l) - d(t), 0)}{\sum_{t=1}^{T} d(t)} \leq l \cdot \frac{\sum_{t=1}^{T} \max(d(t+1) - d(t), 0)}{\sum_{t=1}^{T} d(t)}.$$

This expression can be bounded from above by solving

$$\begin{aligned}
\text{maximize} \quad & \frac{\sum_{t=0}^{T-1} \max(d(t) - d(t+1), 0)}{\sum_{t=1}^{T} d(t)} \\
\text{subject to} \quad & |d(t) - d(t+1)| \leq \Delta d, \forall t, \\
& d(0) = N, \\
& d(t) \geq 0, \forall t.
\end{aligned}$$

The solution, $\{d^*(t)\}$ to this optimization problem is

$$\begin{aligned} d^*(t) &= \max(d^*(t-1) - \Delta d, 0), \forall t \\ &= \max(N - \Delta d \cdot t, 0) \end{aligned}$$

We can prove this by contradiction. Assume that we have some other solution with a strictly larger cost than $\{d^*\}$. Then, starting from $d(T)$ and proceeding backwards over t, we can decrease the value of each term as much as possible, increasing the numerator of the objective function while decreasing the denominator. Thus, reduce the overall cost and obtain the solution above, obtaining a contradiction. With the solution in hand, since the $d^*(t) = 0$ for large enough t, we have

$$\sum_{t=1}^{T} d^*(t) \leq \frac{N^2}{2\Delta d},$$

and

$$\sum_{t=0}^{T-1} \max(d(t) - d(t+1), 0) \leq N.$$

Therefore, the bound is

$$\max_{\{d(t)\}\in\bar{\mathcal{U}}, \{x(t)\}\in\mathcal{S}} \frac{\text{Cost of OOF given } (d(t),x(t))}{\text{Clairv. cost given } (d(t),x(t))} = \frac{c_{\max} - c_{\min}}{c_{\min}} \cdot \left(m + \frac{2\Delta d \cdot l}{N}\right) + 1$$

and the proof is complete. □

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

Albers, Susanne. 2003. Online algorithms: a survey. *Mathematical Programming* **97**(1-2) 3–26, **97**(1-2).

Allen, Stephen G., Donato A. D'Esopo. 1968. An ordering policy for repairable stock items. *Operations Research* **16**(3) 669–674, **16**(3).

Ashlagi, Itai, Felix Fischer, Ian A. Kash, Ariel D. Procaccia. 2013. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior Games and Economic Behavior*.

Ashlagi, Itai, Duncan S. Gilchrist, Alvin E. Roth, Michael A. Rees. 2011. Nonsimultaneous chains and dominos in kidney-paired donationrevisited. *American Journal of Transplantation* **11**(5) 984–994, **11**(5).

Bansal, Nikhil, Niv Buchbinder, Anupam Gupta, Joseph Seffi Naor. 2007. An o (log2 k)-competitive algorithm for metric bipartite matching. *AlgorithmsESA 2007*. Springer, 522–533, 522–533.

Bertsekas, Dimitri P. 2005. *Dynamic Programming & Optimal Control, Vol. I*. 3rd ed. Athena Scientific, Athena Scientific.

Chen, Li, Erica L. Plambeck. 2008. Dynamic inventory management with learning about the demand distribution and substitution probability. *Manufacturing & Service Operations Management* **10**(2) 236–256, **10**(2).

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, **23**(04).

Cox, David Roxbee, David Oakes. 1984. *Analysis of survival data*, vol. 21. CRC Press, CRC Press.

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, **194**(1).

Dempster, Arthur P., Nan M. Laird, Donald B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 1–38*Journal of the Royal Statistical Society. Series B (Methodological)*.

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, **19**(01).

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, **103**(1).

Geyer, Roland, Luk N Van Wassenhove, Atalay Atasu. 2007. The economics of remanufacturing under limited component durability and finite product life cycles. *Management Science* **53**(1) 88–100, **53**(1).

Godfrey, Gregory A., Warren B. Powell. 2001. An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Science* **47**(8) 1101–1112, **47**(8).

Goeman, Jelle J. 2010. L1 penalized estimation in the cox proportional hazards model. *Biometrical Journal* **52**(1) 70–84, **52**(1).

Greenwood, Major, others. 1926. A report on the natural duration of cancer. *Reports on Public Health and Medical Subjects. Ministry of Health* (33), (33).

Guide, V. Daniel R, Luk N Van Wassenhove. 2009. OR FORUMthe evolution of closed-loop supply chain research. *Operations Research* **57**(1) 10–18, **57**(1).

Gusfield, Dan, Robert W. Irving. 1989. *The stable marriage problem: structure and algorithms*, vol. 54. MIT press Cambridge, MIT press Cambridge.

Huang, Wei, Vidyadhar Kulkarni, Jayashankar M Swaminathan. 2008. Managing the inventory of an item with a replacement warranty. *Management Science* **54**(8) 1441–1452, **54**(8).

Huh, Woonghee Tim, Retsef Levi, Paat Rusmevichientong, James B. Orlin. 2011. Adaptive data-driven inventory control with censored demand based on kaplan-meier estimator. *Operations Research* **59**(4) 929–941, **59**(4).

Huh, Woonghee Tim, Paat Rusmevichientong. 2009. A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* **34**(1) 103–123, **34**(1).

Jennings, John B. 1973. Blood bank inventory control. *Management Science* **19**(6) 637–645, **19**(6).

Kaplan, Edward L., Paul Meier. 1958. Nonparametric estimation from incomplete observations. *Journal of the American statistical association* **53**(282) 457–481, **53**(282).

Karim, Md Rezaul, Kazuyuki Suzuki. 2005. Analysis of warranty claim data: a literature review. *The International Journal of Quality & Reliability Management* **22**(7) 667–686, **22**(7).

Khawam, John, Warren H Hausman, Dinah W Cheng. 2007. Warranty inventory optimization for hitachi global storage technologies, inc. *Interfaces* **37**(5) 455–471, **37**(5).

Kleywegt, Anton J., Alexander Shapiro, Tito Homem-de Mello. 2002. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* **12**(2) 479–502, **12**(2).

Kuhn, Harold W. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly* **2**(1-2) 83–97, **2**(1-2).

Levi, Retsef, Robin O. Roundy, David B. Shmoys. 2007. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research* **32**(4) 821–839, **32**(4).

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, **56**(5).

Levina, Elizaveta, Peter Bickel. 2001. The earth mover's distance is the mallows distance: some insights from statistics. *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, vol. 2. IEEE, 251–256, 251–256.

Lindgren, Anna. 1997. Quantile regression with censored data using generalized l1 minimization. *Computational Statistics & Data Analysis* **23**(4) 509–524, **23**(4).

Lubin, Miles, Iain Dunning. 2013. Computing in operations research using julia. *arXiv:1312.1431 [cs, math] arXiv:1312.1431 [cs, math]*. arXiv: 1312.1431.

Major, Pter. 1978. On the invariance principle for sums of independent identically distributed random variables. *Journal of Multivariate Analysis* **8**(4) 487–517, **8**(4).

McLachlan, Geoffrey, Thriyambakam Krishnan. 2007. *The EM algorithm and extensions*, vol. 382. John Wiley & Sons, John Wiley & Sons.

Meyerson, Adam, Akash Nanavati, Laura Poplawski. 2006. Randomized online algorithms for minimum metric bipartite matching. *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. 954–959, 954–959.

Murthy, D.N.P., W.R. Blischke. 1992. Product warranty management III: A review of mathematical models. *European Journal of Operational Research* **63**(1) 1–34, **63**(1).

Murthy, D.N.P, O Solem, T Roren. 2004. Product warranty logistics: Issues and challenges. *European Journal of Operational Research* **156**(1) 110–126, **156**(1).

Nahmias, Steven. 2011. *Perishable inventory systems*, vol. 160. Springer, Springer.

Pegels, C. Carl, Andrew E. Jelmert. 1970. An evaluation of blood-inventory policies: a markov chain application. *Operations Research* **18**(6) 1087–1098, **18**(6).

Petersen, Brian J. (Brian Jude). 2013. Reverse supply chain forecasting and decision modeling for improved inventory management. Thesis, Massachusetts Institute of Technology, Massachusetts Institute of Technology.

Roth, Alvin E. 2003. The origins, history, and design of the resident match. *JAMA* **289**(7) 909–912, **289**(7).

Roth, Alvin E., Marilda A. Oliveira Sotomayor. 1992. *Two-sided matching: A study in game-theoretic modeling and analysis*. 18, Cambridge University Press, Cambridge University Press.

Savaskan, R. Canan, Shantanu Bhattacharya, Luk N. van Wassenhove. 2004. Closed-loop supply chain models with product remanufacturing. *Management Science* **50**(2) 239–252, **50**(2). ArticleType: research-article / Full publication date: Feb., 2004 / Copyright 2004 INFORMS.

Simpson, Vincent P. 1978. Optimum solution structure for a repairable inventory problem. *Operations Research* **26**(2) 270–281, **26**(2).

Taylor, Jeremy MG. 1995. Semi-parametric estimation in failure time mixture models. *Biometrics* 899–907*Biometrics*.

Tibshirani, Robert. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* 267–288*Journal of the Royal Statistical Society. Series B (Methodological)*.

Tibshirani, Robert, others. 1997. The lasso method for variable selection in the cox model. *Statistics in medicine* **16**(4) 385–395, **16**(4).

Toktay, L. Beril, Lawrence M. Wein, Stefanos A. Zenios. 2000. Inventory management of remanufacturable products. *Management science* **46**(11) 1412–1426, **46**(11).

Wu, CF Jeff. 1983. On the convergence properties of the EM algorithm. *The Annals of statistics* 95–103*The Annals of statistics*.