NORTHWESTERN UNIVERSITY

Product Lifecycle Considerations in Closed-Loop Supply Chain Management

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Industrial Engineering and Management Sciences

By

Kristin Marie Sahyouni

EVANSTON, ILLINOIS

June 2007

ABSTRACT

Product Lifecycle Considerations in Closed-Loop Supply Chain Management

Kristin Marie Sahyouni

This dissertation examines the impact of product returns on effective supply chain management. Within this area of research, known as Closed-Loop Supply Chain Management, we consider both strategic and tactical level reverse logistics and inventory management problems from the perspective of a firm which must efficiently process returned items. More specifically, we explore the effective integration of forward and reverse logistics systems throughout the product lifecycle and the impact of lifetime buys on repair operations for short lifecycle products.

The first part of this dissertation develops three general bidirectional facility location models to deliver products and collect returns in a two-tier supply chain. This research quantifies the value of simultaneously considering forward and reverse product flows when designing an integrated closed-loop logistics network. We develop these models as extensions of the classical uncapacitated fixed charge location problem and present a Lagrangian relaxation-based solution algorithm that is quick and effective. We measure the cost savings opportunities of integrated network design throughout the introductory, maturity, and decline stages of the product lifecycle. In addition, we discuss the resulting network configurations and introduce a new network similarity metric to quantify this analysis.

The second part of this dissertation investigates the impact of lifetime buys on warranty repair operations for short lifecycle products. This work is the first to consider the implications of a single procurement opportunity for repair parts, which is common practice in the electronics and telecommunications industries. Using a deterministic continuous time

model, we show how fixed repair capability costs, variable repair costs, inventory holding costs, and replacement costs affect a firm's optimal repair and replacement decisions for a single product. We extend these models to examine the role of repair capacity constraints for the single product model and the impact of having shared repair facilities that service two products. Our models are applied to an industry case to gain insights for a U.S. mobile device manufacturer.

The third part of this dissertation develops a scenario planning extension of the lifetime buy problem to understand the effects of return rate uncertainty in the inventory planning process. We present two models that address the somewhat competing objectives of minimizing cost and minimizing risk under a discrete set of return rate scenarios. We also explore the tradeoff curve of efficient solutions that are Pareto-optimal in the two objectives. Our analysis empirically shows that solutions to this problem are robust with respect to minimizing cost and risk.

ACKNOWLEDGMENTS

I would like to start by thanking my advisors, Mark Daskin and Canan Savaskan, for their guidance over the last four years. They have helped me to not only develop into a better student and researcher, but have encouraged me to do work that will be meaningful and have impact. Their support, both academically and personally, have made this dissertation possible. Thank you both for your advice over the years and your friendship.

I also thank the other members of my committee, Moritz Fleischmann and Karen Smilowitz. Moritz provided many great ideas and insights that allowed me to see valuable ways to extend my research. He has been very helpful throughout this process and I appreciate his time and kindness. Karen has provided valuable advice and introduced me to a number of professional opportunities. Special thanks to the rest of the IE/MS faculty and staff who have been very gracious to me over the past four years. I'm also extremely grateful to Larry Maye for providing industry insights into my work and for introducing me to the lifetime buy problem. His comments and expertise were invaluable to this research.

I would also like to thank the wonderful faculty at the University of Southern California. Erik Johnson was a great teacher and has been very supportive of me since I first met him. Fernando Ordóñez taught my first graduate optimization class and has since been a source of support and encouragement throughout graduate school. Finally, I would like to express my sincere thanks to Maged Dessouky for being a great mentor to me during my time at USC and Northwestern. He was very influential during my undergraduate years and encouraged me to explore research opportunities and a career in academia. Without his guidance and support, I might not have pursued a PhD in operations research.

Much of the work in this dissertation was supported financially by a number of organizations. I am very thankful to the National Science Foundation Graduate Research Fellowship

Program, the National Highway Institute Eisenhower Transportation Fellowship Program, the Northwestern Transportation Center, and The Graduate School of Northwestern University.

I am also grateful to so many wonderful friends who have encouraged me over the years and helped keep me sane throughout graduate school: Jessica Del Cid has an unshakeable positive spirit and has always been there for me; Katie Dunham and Jason Scott who have always been supportive and make me want to be a better person; Peter Choi who always makes me laugh and reminds me why I work so hard; Chip Murphy has always been encouraging, providing an open ear and helpful advice; and Lesley Perry who has been a good friend, and introduced me to fun Chicago adventures! I'd also like to thank my IE/MS friends: Rob Lien and Jamie Wieland have been great friends since our first year at Northwestern; Peter Francis has been very encouraging and supportive; and finally Frank Curtis who has been a great runch buddy and a true friend. Without Peter and Frank I might not have made it - thank you both for your unending encouragement. I appreciate you all and am grateful that you are a part of my life.

Lastly I'd like to thank my parents Andy and Maureen who have taught me to work hard and provided every opportunity for me to be successful. I'm also forever grateful to my sisters Lori and Katie who keep a smile on my face and laughter in my heart. Thank you all for believing in me and encouraging me throughout this program. Thank you for supporting me in the many challenges I've faced and sharing in my success. Your unconditional love and support have enabled me to become the person I am today.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This dissertation utilizes operations research methods and other quantitative techniques to examine network design and inventory management topics within the context of closed loop supply chain (CLSC) management, a research area that considers the effect of product returns and remanufacturing operations on supply chain activities. Within this domain, this dissertation explores (1) the effective integration of forward and reverse logistics systems throughout the product lifecycle, and (2) the impact of lifetime buys on repair operations for short lifecycle products.

1.1 Closed-Loop Supply Chain Management

Traditional supply chain management research typically deals with the *forward flow* of goods. In this forward flow, raw materials are processed into parts and subassemblies by one or more manufacturers, and undergo numerous operations to produce a finished good. Once these processes are complete, the finished good is moved through a network of distribution centers to arrive at a retail outlet (or similar sales location) where it is sold to a customer. This process is depicted in Figure 1.1. In this forward-oriented supply chain, the product is thought to have left the system once it has been consumed (purchased).

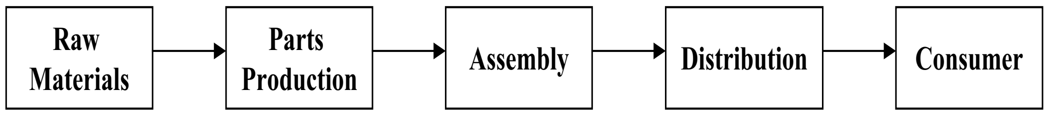

Figure 1.1: A Traditional Forward Supply Chain

This traditional way of thinking about supply chains often ignores the role of product returns and resulting remanufacturing activities on the design and operation of these highly complex systems. In truth, product returns can have a tremendous impact on a firm's supply chain performance. A returned product can enter the system at each tier of the supply chain, depending on its condition and what (if any) processing is required to return the product to a useable state. Throughout this dissertation, we use the term *forward logistics* to refer to the forward flow of products (from raw material to customers) defined above and the term *reverse logistics* to refer to any activities or processes that treat return flows. We use the term *closed loop supply chain* to refer to a flexible system that includes both forward logistics and reverse logistics activites to serve bidirectional flows. Figure 1.2 depicts the major components and activities in a closed-loop supply chain.

Effective closed loop supply chain management is critical to the success of many firms. One recent study shows that the estimated annual value of returned goods in the United States is $60 billion and that firms spend $40 billion a year to manage these returns (Enright, 2003). The US Central Intelligence Agency estimates the 2005 US annual Gross Domestic Product to be roughly $12,370 billion (CIA World FactBook, 2006) in which case returned goods account for approximately 0.5% of the United States economy each year. Another source claims that the value of returned goods exceeds $100 billion annually (Stock, Speh, and Shear, 2002)! In addition to these cost figures, many studies have shown that return rates can reach as high as 35% for catalog retailers, 50% for online retailers, and 40% for textile

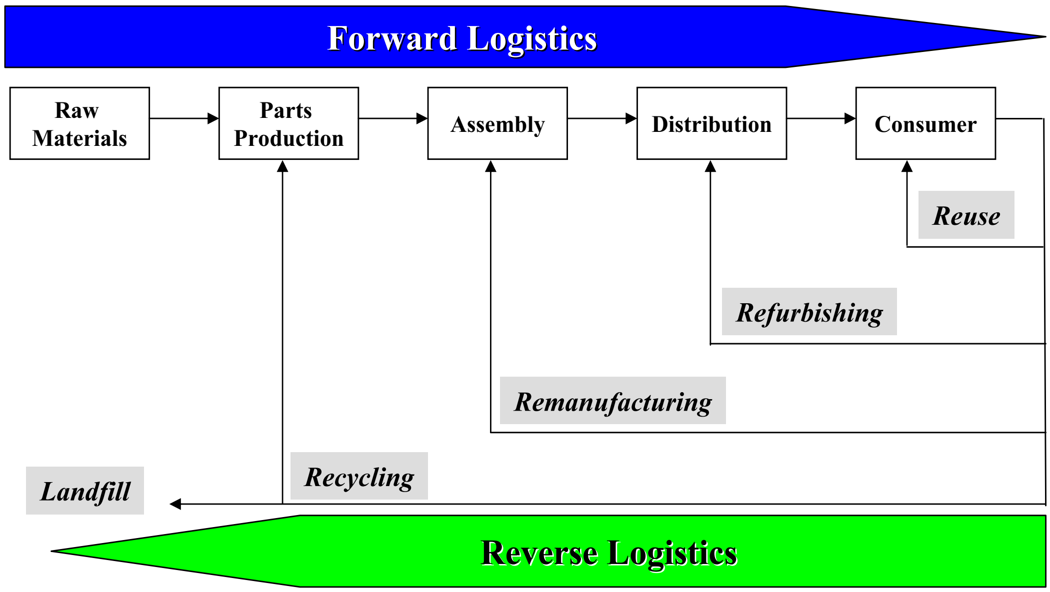

Figure 1.2: A Closed-Loop Supply Chain

manufacturers (Rogers and Tibben-Lembke, 1999). The size, scope, and cost of returns all suggest that the value of effectively managing reverse logistics processes is significant.

Closed loop supply chain and reverse logistics topics have emerged as an important component of supply chain management research due to three current trends: (1) governmental imposition of stricter environmental regulations on product recycling and disposal, (2) shorter product life cycles and subsequent increased supply of used goods with high salvage value, and (3) the explosion of e-commerce activities resulting in a high volume of products returned from consumers for exchange or refund. It is worth noting that the entire business plan of some firms (e.g., Netflix.com) is predicated on the availability of an efficient and effective reverse logistics network. As these trends continue to evolve, effective closed-loop supply chain management practices will become increasingly important for companies. Many companies now recognize that getting their product to the customer is not the end of the story, but the beginning of a new era of challenges that requires the design of effective reverse logistics channels to efficiently handle product returns. Firms that fail to recognize the importance of reverse logistics may be banned from certain markets by legislation and/or left behind in other markets as they fail to recover valuable used products.

The academic community has responded to the need for a better understanding of the needs of effective returns management with the creation of closed loop supply chain research, which explicitly considers the impact of product returns and resulting operations on the entire supply chain - including product design, forward distribution and returns collection, and assembly (or disassembly) operations. Until recently, the subject was solely considered within an environmental or sustainability context, but over the last decade CLSC management has emerged as an important research field and is becoming an increasingly common topic within traditional supply chain management literature. A recent presentation by Guide and Van Wassenhove (2005), gives a brief overview of the history of closed loop supply chain research. The literature on CLSC management includes topics such as: remanufacturing processes and design; reverse logistics and remanufacturing networks; production and inventory control in remanufacturing environments; coordination in reverse supply chains; the time value of commercial returns; marketing issues for new and remanufactured products; and many other topics. We outline the contributions made by this dissertation to the existing literature in Section 1.2.

1.2 Research Contributions

Supply chain management research may be classified according to its scope of impact - e.g. a *strategic-*, *tactical-*, or *operational-* level problem. A *strategic-* level problem considers long term planning scenarios, such as logistic network design or strategic sourcing decisions. A *tactical-* level problem considers midrange planning scenarios, such as inventory management or capacity planning problems. *Operational-* level planning considers short term problems such as production scheduling or distribution dispatching problems. In addition to its scope of impact, it is very important to identify the conditions under which a problem is studied. A scenario in which all parameters are known, also referred to as a *deterministic* problem, is modeled very differently from a *stochastic* scenario in which uncertainty is explicitly considered. A framework for categorizing problems according to these two measures is given in Figure 1.3. Using this taxonomy, we classify the research contributions of this dissertation as *strategic-deterministic, tactical-deterministic,* and *tactical-stochastic* in nature.

Problem level / **Problem type**	**Strategic**	**Tactical**	**Operational**
	Distribution network design	Capacity planning Inventory management	Production planning / scheduling Demand / Return forecasting
Deterministic	**Integrated Bidirectional Network Design**	**Lifetime Buys in Warranty Repair Operations**	
Stochastic		**Scenario Planning for Lifetime Buys in Warranty Repair Operations**	

Figure 1.3: Research Framework

As noted in Section 1.1, there have been many great contributions to the literature on closed loop supply chain management. This dissertation contributes to this area of research by examining (1) the role of product lifecycles in the design of bidirectional CLSC logistics systems and (2) the impact of lifetime buys on warranty repair operations for short lifecycle products. More specifically, this dissertation contributes as follows:

1. We develop a general bidirection facility location model to serve both forward and reverse product flows in a two-tier supply chain. This model is extended to consider the logistics needs of a product during the introductory, maturity, and decline stages

of a product's lifecycle. The embedded forward and reverse logistics networks are determined *sequentially* and *concurrently* to quantify the value of integrated CLSC network design. (Chapter 2)

2. We develop a *network similarity metric* to quantify the spatial similarities and differences among multiple network configurations. Using a modified out-of-kilter flow algorithm, we develop the first known metric to consider the total network distance between competing facility location plans. (Chapter 2)

3. We investigate the impact of lifetime buys on warranty repair operations in the electronics and telecommunications industries. In this work we model a repair operation in which defective warranty items arrive at a manufacturer and must be repaired or replaced with a comparable next generation alternative. There is a single procurement opportunity (lifetime buy) at the beginning of the period when a manufacturer can obtain repair parts from a third party supplier. We develop a deterministic continuous time model that examines how repair capability costs, inventory holding costs, variable repair costs, and replacement costs affect firm decisions in this environment. Using constrained numerical optimization methods, we derive the firms optimal repair/replace policy and order quantity for a single product. (Chapter 3)

4. We extend contribution 3 to consider the impact of repair capacity constraints for a single product as well as the impact of having shared repair facilities to service two products. (Chapter 3)

5. We extend contribution 3 to consider a situation in which the return rate observed by the manufacturer is uncertain but limited to a finite number of scenarios. We consider two planning objectives: (1) minimize the total expected cost, and (2) minimize the maximum regret across all return rate scenarios. We show how the total cost to

the manufacturer and the resulting optimal inventory quantity change as the relative emphasis on each of these two objectives changes. (Chapter 4)

1.3 Outline

The remainder of this dissertation is organized as follows. In Chapter 2 we develop a facility location model for bidirectional flows and present three approaches to modeling closed-loop supply chain systems as extensions of the classic uncapacitated fixed charge facility location model. In Chapter 3 we develop a deterministic continuous-time model to study the impact of lifetime buys on warranty repair operations. We show how a manufacturer should decide whether to repair or replace returned defective items as a function of fixed, variable repair, inventory holding, and replacement costs; we also show how this analysis leads to an optimal inventory order quantity. This work is extended to consider multiple scenarios in Chapter 4. We study the effects of different objective functions based on minimizing cost or maximum regret, and show how the optimal inventory order quantity changes with the objective. Literature related to each of the problems is presented within the appropriate chapter. Finally, conclusions and directions for future research are discussed at the end of each chapter.

Chapter 2

A Facility Location Model For Bidirectional Flows

Recent research on reverse logistics and closed loop supply chains has produced a number of specialized network design models. In this chapter we develop three *generic* facility location models for the *integrated* distribution and collection of products that accommodate a variety of applications and industries. These models quantify the value of integrated decision making in the design of forward and reverse logistics networks throughout different stages of a products life cycle. The formulations extend the uncapacitated fixed charge location model to include the location of used product collection centers and the assignment of product return flows to these centers. In addition, we develop a Lagrangian relaxation based solution algorithm that is both quick and effective. We measure the implications of integrated decision making by comparing the total facility and transportation costs of our joint optimization models to the cost of solutions in which forward and reverse facility location decisions are made *sequentially* or *independently* of one another. In addition, we discuss the implications of integrated decision making on network configuration and introduce a new *network similarity measure* to quantify this analysis. Computational results are also presented.

2.1 Introduction

A number of formulations for designing product collection networks have been proposed in the literature, but they tend to be specific to particular case study applications. In this chapter we formulate three generic facility location models for the *integrated* design of forward and reverse logistics networks. The forward and reverse networks are integrated by locating bidirectional distribution centers in addition to dedicated unidirectional facilities. The generalized models we propose are useful for understanding logistics activities in a variety of industries and can be extended to capture complexities encountered in specific applications. The primary contribution of this work is in developing models to analyze the impact of integrated product distribution and returns collection between the distribution and consumer tiers of the supply chain. Our work addresses the following research questions:

- How can existing facility location models be extended to effectively integrate forward distribution and reverse collection activities and how must solution algorithms be modified to accommodate the presence of return flows?
- What is the value of integrated forward and reverse logistics network design for a company in different stages of a product's life cycle?
- What are the network configuration implications of integrating forward and reverse logistics design decisions?

2.1.1 Sources of Product Returns and Resulting Reverse Logistics Needs

Product returns occur for several reasons. To understand these reasons and determine management needs for an effective response, it is crucial to differentiate among the sources of returns. Empirical research suggests that product returns can be categorized into two primary groups: push returns and pull returns. The first type of product returns are pushed back through the supply chain (from end users towards the origin of supply) as a result of commercial returns, warranty claims, and repairs. The second type of returns are pulled into the reverse logistics network as a result of legislation or the financial benefits of reclaiming high valued items for remanufacturing or recycling. We discuss these two types of returns in detail below.

Reverse logistics has always been a norm of doing business, especially for catalog retailers (Rogers and Tibben-Lembke, 1999). However, the recent growth of business-to-consumer direct distribution channels such as internet sales and telemarketing has led to an explosion of commercial returns for both retailers and manufacturers. These returns can be attributed to relaxed return policies, product quality that does not meet customer expectations, or consumers who are uninformed regarding proper product use. Consequently, the return rates for computer products and for products purchased through catalogs are in the range of 10 - 20% and 18 - 35%, respectively (Rogers and Tibben-Lembke, 1999).

In some cases, companies are forced to set up reverse logistics and processing networks to take back their products after they have been used and disposed of by their consumers. Political concerns for the environment in Europe, in the US, and in Japan have resulted in legislation that requires the recycling and remanufacturing of specific products and materials. Printers, copiers, computers, power tools, and cellular phones are examples of durable products which are collected by manufacturers after use and remanufactured into new products. In non-durables, examples include copy and print cartridges and single use cameras. Marketing activities such as trade-in offers and asset recovery services to induce new product demand can also lead to a stream of used products flowing back to resellers or manufacturers (Savaskan, Coughlan, and Shulman, 2006a; Savaskan, Coughlan, and Shulman, 2006b). The timing and quality of end of life returns differ considerably from commercial returns, and as a result so do their processing needs. Table 2.1 identifies drivers and high frequency time periods for both commercial returns and end of life returns.

Type of Returns	Drivers of Returns	Timing of Returns
Commercial Returns (push returns)	▪ Product is defective ▪ Product does not meet consumer expectations ▪ Consumer does not know how to install or use the product ▪ Wrong product delivered ▪ Relaxed return policies	Higher volumes at the introductory stage of the product life cycle
End of Life Returns Environmental Trade-ins Replacements (pull returns)	▪ Short product life cycles and new product introductions, resulting in trade-ins ▪ Marketing campaigns for product recycling ▪ Legislation ▪ Product stewardship trough asset recovery services	Higher volumes during maturity and decline stages of the product life cycle

Table 2.1: Drivers and Timing of Product Returns

A significant return flow of products can present major obstacles to a firms logistics operations since many of todays supply chain networks were originally designed with only forward product flow in mind (Rogers and Tibben-Lembke, 1999). For this reason, companies now face a considerable challenge in designing a reverse supply chain network that will meet their returns processing needs while complementing their existing forward distribution system.

2.1.2 Three Generic Reverse Network Structures

While there are specific industries that can be characterized as either forward or reverse dominant, many industries experience significant distribution and collection flow changes throughout a products lifecycle. In these cases, network design decisions may change over time.

During the introductory stage there are very few products in the market, which implies even fewer returns. During this time, the firm will focus on forward product distribution and in general there is little need for returns collection. Since the product is newly introduced to the market, the majority of product returns are commercial returns due to unmet customer needs. During this phase, a company may focus on the design of its forward distribution network and collect returns at a subset of forward distribution points. Online retailers with centralized returns handling and decentralized product distribution, such as those found in the textile industry, are an example of **forward dominant** reverse logistics models.

By contrast, a companys primary focus is likely to shift to recovery and collection at the end of a products lifecycle as new product demand declines and used product returns increase over time. During this period a **reverse dominant** reverse logistics network is prevalent. Remanufacturing and recycling networks are examples of reverse dominant network structures, where there is a large number of collection sites and a few centralized processing facilities. Other examples of reverse dominant network design can be found in refillable ink / printer cartridge and single use camera industries.

During the maturity stage, forward and reverse flows are generally stable, so the network tends to favor, but not *require*, the location of bidirectional facilities. During this time neither the forward distribution nor the product collection network dominates the other. The Dell asset recovery program, which collects both end of life and commercial returns is an example of a distribution network whose design is similar to a **co-location** model.

The following graph, Figure 2.1, depicts how product demand, commercial return flows, and end of life return flows change throughout the introductory, maturity, and decline phases of the product lifecycle and the corresponding distribution model.

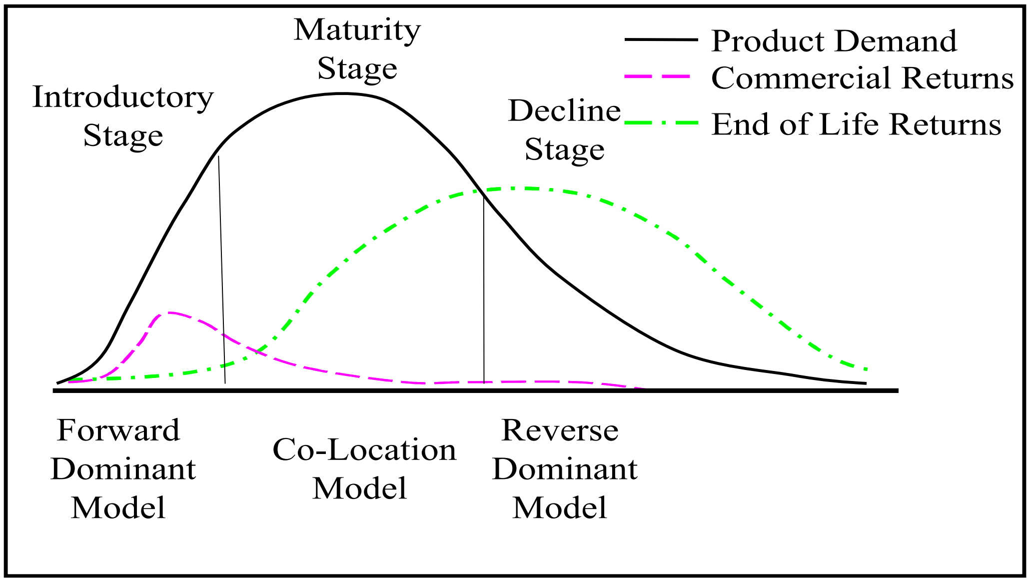

Figure 2.1: Product Flows Throughout the Lifecycle

In practice, firms rarely optimize their forward and reverse network design decisions jointly due to evolving product and market needs or limited management attention or resources. It is important to understand how much value is forgone by not integrating forward and reverse logistics decision making and the conditions under which integration is critical. We measure the value of integrated forward and reverse network design by comparing the total facility and transportation costs of the integrated model to one in which forward and reverse facility decisions are made sequentially over time. We also investigate how much the forward and reverse network configurations change under integrated decision making. We introduce a new network similarity measure which can be used to assess the degree of similarity between any two location plans, providing a quantitative method of comparing solutions not previously found in location analysis literature.

We propose three models that extend the uncapacitated fixed charge location problem to simultaneously locate forward and reverse distribution centers within the frameworks described above, and assign forward and reverse flow between customers and facilities. These three models enable us to study the value and impact of integrated network design throughout the product lifecycle. The models are formulated as mixed integer linear programs and are solved using a specialized Lagrangian relaxation heuristic.

The remainder of this chapter is organized as follows. In Section 2.2 we review some of the current literature on closed loop supply chain management, specifically focusing on papers that integrate forward and reverse logistics design. The uncapacitated fixed charge location problem is extended in Section 2.3 to incorporate the costs of integrating product recovery operations into a traditional distribution network. We provide formulations and solution algorithms for three variations of the integrated distribution network design problem. A network similarity metric is developed in Section 2.4. Computational results and algorithm performance are presented in Section 2.5. Conclusions and directions for future research comprise Section 2.6.

2.2 Literature Review

A considerable amount of empirical work in the realm of reverse logistics and closed loop supply chains has been performed in the last ten years. The bulk of this work, in the form of surveys and case studies, has motivated the advancement of this particular field of research and has proven helpful in quantifying the cost savings opportunities to be realized by optimizing closed loop supply chain processes. This work has been valuable in both defining models for future study and framing the quantitative work that follows.

Rogers and Tibben-Lembke (1999) conducted a survey of manufacturers, wholesalers, retailers, and service providers to elicit information regarding returns policies, reverse logistics activities, and operational challenges. Their work is purely descriptive but is useful in identifying opportunities for future study. Fleischmann et al. (2000) review a number of case studies that detail product recovery network design in various industries. They categorize common activities used in product recovery and remanufacturing processes and contrast these activities with those found in other logistics networks. Krumwiede and Chwen (2002) conduct a field study to examine the various operational and organizational structures required for a company to incorporate reverse logistics practices effectively into their core business and present a decision model to aid companies when considering outsourcing their reverse logistics activities. Finally, De Brito (2003) reviews reverse logistics research to date and outlines a framework for decision making in a returns environment. She reviews a number of case studies and considers returns handling and inventory management issues in depth.

The current modeling literature on closed-loop supply chain management includes topics such as: remanufacturing processes and product design, reverse logistics and remanufacturing networks, remanufacturing production control, inventory control, coordination in reverse supply chains, time value of returns, and marketing issues for remanufactured products. Within this stream of literature, reverse logistics typically deals with three major areas: network design, product design for reuse and recycling, and product recovery strategies and inbound planning. Network design encompasses decisions about how products should be collected and who should collect them (Savaskan and Corbett, 2002; Savaskan, Debo, and Van Wassenhove, 2003; Savaskan, Bhattacharya, and Van Wassenhove, 2004; Savaskan, 2006), where the collection sites should be located (Fleischmann et al., 2001), and whether the collection facilities should be stand-alone operations or should be integrated with existing forward supply chain facilities (Marin and Pelegrin, 1998). Product design deals with the choice of the materials to be used in products, the ease of disassembly of the parts, and part commonality which increases the potential for reuse. Finally, product recovery strategies and inbound planning refers to decisions about whether to reuse, remanufacture, refurbish or recycle the recovered products (Savaskan and Aytekin, 2005). Inventory and capacity management clearly affect and are affected by product recovery decisions.

Much of the network design research in reverse logistics has focused on the location of manufacturing and remanufacturing sites. Bloemhof-Ruwaard et al. (1996) consider a problem in which plants are located to manufacture products to supply given customer demands, and disposal facilities are located to collect waste that is generated during the manufacturing process. Spengler et al. (1997) formulate a multi-staged, multi-product capacitated problem in which processing facilities are located and flows between product sources and recycling facilities are determined. Krikke (1998) proposes a mixed integer linear program to determine optimal facility and processing locations for products to be collected, reprocessed, and sold in a secondary market. Barros et al. (1998) develop a model to minimize total fixed, transportation, and processing costs of locating separate collection and treatment facilities for given waste supply and processed demand quantities and locations. Marin and Pelegrin (1998) propose the Return Plant Location problem in which facilities that process both forward and return flows are located to minimize the total fixed and transportation costs of supplying primary products from a distribution center to customers and returning some proportion of those products back to the DC. Note that all facilities in their model serve both forward and reverse flows. Jayaraman et al. (1999) develop a closed-loop model to determine optimal distribution and remanufacturing facility locations as well as production, stocking, and shipment quantities for used core components and remanufactured products. Fleischmann et al. (2001) introduce a general recovery network design model in which manufacturing plants, forward distribution centers, and separate product collection centers are located to minimize the total fixed location and transportation costs of the system. Jayaraman et al. (2003) propose a reverse logistics model that does not consider forward distribution activities when locating capacitated collection and refurbishing facilities to process used products. Beamon and Fernandes (2004) consider the tradeoff between capacity investment costs and recurring operational costs in a closed-loop supply chain in which collection centers and warehouses are located to accept returned products and remanufacture them for resale. Finally, Fleischman et al. (2004) present a review of papers that specifically consider reverse logistics costs in both deterministic and stochastic demand-based facility location models.

Most, if not all, of the models outlined above have been solved using general purpose optimization solvers such as CPLEX or GAMS. This has tended to limit the size of the problems that can be attacked. By way of contrast, our models can readily be solved using Lagrangian relaxation for instances several times the size of the largest problems reported on earlier. (It is important to note that some of these models include capacity and other complicating constraints which we do not consider and that these may also limit the size of problems that can be studied.) Therefore, our interest is in understanding how facility location models and solution methods change when reverse flows are introduced. This chapter contributes to the existing literature by introducing 3 generic models that consider the impact of integrating forward and reverse network design decisions through the location of bidirectional distribution centers.

2.3 Bidirectional Facility Location Models

In this section, we formulate three extensions to the classical uncapacitated fixed charge location problem (Balinski, 1965) that integrate forward and reverse distribution activities through the location of bidirectional distribution centers in addition to dedicated unidirectional facilities. Numerous empirical studies have shown that most distribution networks are designed with a specific product flow direction, namely forward distribution, in mind (Rogers and Tibben-Lembke, 1999). This emphasis on single direction distribution makes it challenging to effectively combine separate forward and reverse distribution networks into one integrated bidirectional flow network. We consider two methods to accomplish this goal: (1) using financial incentives to induce the location of bidirectional facilities, and (2) imposing constraints that require one type of network to fit within the other. The first method is used to formulate the Co-Location (CL) model while the second is used to formulate the Forward Dominant (FD) and Reverse Dominant (RD) models described in Section 2.2. The CL model locates bidirectional facilities in addition to stand alone forward and reverse facilities. The FD model allows only bidirectional and stand alone forward facilities while the RD model allows only bidirectional and stand alone reverse facilities. Figure 2.2 below illustrates sample network configurations for these three models:

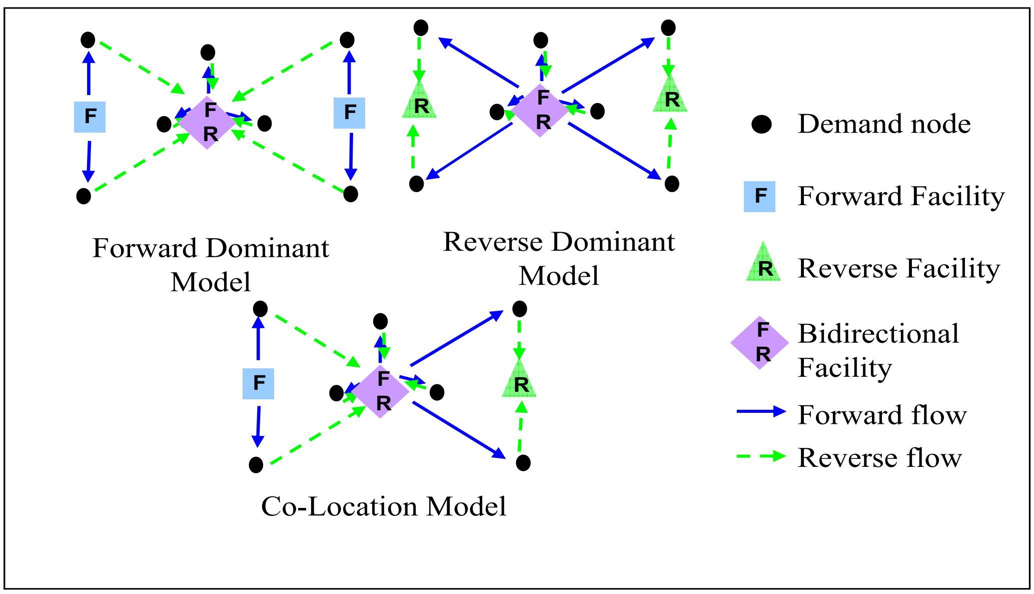

Figure 2.2: Sample Network Configurations for FD, RD, and CL Models

2.3.1 The Co-Location Model

We consider the problem of jointly locating forward and reverse facilities. We assume that the forward demand at some node i is given by h_i and that the unit cost of (forward) shipping from some candidate site j to that demand node is given by c_{ij}. The fixed facility cost of locating a forward distribution center at candidate site j is given by f_j. We further assume that the return rate is given by α_i, so the number of returns per unit time at demand node i is given by $\alpha_i h_i$. Similarly, we assume that the unit cost of shipping a return item is given by $\gamma_{ij} c_{ij}$, where γ_{ij} can be any value greater than 0. In particular, we do not assume that the unit cost of return shipping is either greater than ($\gamma_i j > 1$) or less than ($\gamma_i j < 1$) the forward unit shipping cost. We assume that the fixed facility cost for a reverse site is given by $\beta_j f_j$. Finally, if a forward and a reverse facility are co-located then there is a fixed cost savings of s_j, where $s_j \leq min\{f_j, \beta_j f_j\}$.

The objective of this problem is to identify the locations of forward, reverse, and co-located facilities as well as the assignment of forward and reverse demands to those facilities. (See Figure 2.2 for a sample configuration of a co-location network.) To formulate this problem as a linear program, we define the following decision variables:

$$X_j^F = \begin{cases} 1 & \text{if a forward distribution center is located at site } j, \\ 0 & \text{if not} \end{cases}$$

$$X_j^R = \begin{cases} 1 & \text{if a reverse distribution center is located at site } j, \\ 0 & \text{if not} \end{cases}$$

$$X_j^C = \begin{cases} 1 & \text{if both a forward and reverse distribution center is located at site } j, \\ 0 & \text{if not} \end{cases}$$

$Y_{ij}^F =$ fraction of forward demand at node i that is served by a distribution center at site j

$Y_{ij}^R =$ fraction of returns at node i that is served by a distribution center at site j

With this notation, we can formulate the **Co-Location (CL) model** as follows:

$$\text{Minimize } \sum_{j\in J} f_j X_j^F + \sum_{i\in I}\sum_{j\in J} h_i c_{ij} Y_{ij}^F + \sum_{j\in J} \beta_j f_j X_j^R + \sum_{i\in I}\sum_{j\in J} \alpha_i h_i c_{ij} Y_{ij}^R - \sum_{j\in J} s_j^C X_j^C \tag{2.1}$$

$$\text{Subject to } \sum Y_{ij}^F = 1, \ \forall i \in I \tag{2.2}$$

$$\sum Y_{ij}^R = 1, \ \forall i \in I \tag{2.3}$$

$$Y_{ij}^F \leq X_j^F, \ \forall i \in I \forall j \in J \tag{2.4}$$

$$Y_{ij}^R \leq X_j^R, \ \forall i \in I \forall j \in J \tag{2.5}$$

$$X_j^F = \{0,1\}, \ \forall j \in J \tag{2.6}$$

$$X_j^R = \{0,1\}, \ \forall j \in J \tag{2.7}$$

$$X_j^C = \{0,1\}, \ \forall j \in J \tag{2.8}$$

$$Y_{ij}^F \geq 0, \ \forall i \in I \forall j \in J \tag{2.9}$$

$$Y_{ij}^R \geq 0, \ \forall i \in I \forall j \in J \tag{2.10}$$

$$X_j^F \geq X_j^C, \ \forall j \in J \tag{2.11}$$

$$X_j^R \geq X_j^C, \ \forall j \in J \tag{2.12}$$

The objective function 2.1 minimizes the sum of the fixed forward and reverse facility costs, the forward and reverse transportation costs, and the savings associated with co-location of forward and reverse facilities. Constraints (2.2) and (2.3) require that all forward and reverse demand nodes be assigned to a facility. Constraints (2.4) and (2.5) ensure that a demand node is not assigned to a facility that has not been opened. Constraints (2.6), (2.7), and (2.8), are the necessary binary constraints for locating facilities while constraints (2.9) and (2.10) are standard non-negativity constraints. Constraints (2.11) and (2.12) link the forward and reverse subproblems by allowing the savings for co-location to be realized only if a site has both a forward and a reverse facility. Note that in the absence of constraints (2.11) and (2.12) and the co-location variables, X_j^C, the model decomposes into two uncapacitated fixed charge problems: one for forward sites and one for reverse sites.

2.3.2 Lagrangian Relaxation Solution Algorithm For The Co-Location Model

The Co-Location model can be solved using a standard IP solver such as CPLEX, but the solution times are excessive and hence limit the size and number of problems that can be solved. To address this problem, we develop a Lagrangian relaxation based algorithm to effectively and quickly solve this problem. We relax the assignment constraints (2.2) and (2.2) to obtain the following Lagrangian formulation:

$$\begin{aligned} \min_{X,Y} \max_{\lambda,\mu} & \sum_{j\in J} f_j X_j^F + \sum_{i\in I}\sum_{j\in J} h_i c_{ij} Y_{ij}^F + \sum_{i\in I} \lambda_i \left(1 - \sum_{j\in J} Y_{ij}^F\right) \\ & + \sum_{j\in J} \beta_j f_j X_j^R + \sum_{i\in I}\sum_{j\in J} \alpha_i h_i c_{ij} Y_{ij}^R + \sum_{i\in I} \mu_i \left(1 - \sum_{j\in J} Y_{ij}^R\right) \\ & - \sum_{j\in J} s_j^C X_j^C \end{aligned} \tag{2.13}$$

Subject to Constraints (2.2) $-$ (2.12) above

In addition, we add the following two constraints:

$$\sum_{j\in J} X_j^F \geq 1 \tag{2.14}$$

$$\sum_{j\in J} X_j^R \geq 1 \tag{2.15}$$

These constraints require that the model locate at least one facility in each direction (forward and reverse). Observe that constraints (2.2) - (2.5) in the original CL formulation implicitly enforce these constraints, but the relaxation of constraints (2.2) and (2.3) in the Lagrangian formulation allows a solution that does not locate any facilities which is clearly undesirable.

The objective function can be rewritten as:

$$\begin{aligned} \text{Minimize } & \sum_{j\in J} f_j X_j^F + \sum_{i\in I} \lambda_i + \sum_{i\in I}\sum_{j\in J} \left(h_i c_{ij} - \lambda_i\right) Y_{ij}^F \\ & + \sum_{j\in J} \beta_j f_j X_j^R + + \sum_{i\in I} \mu_i + \sum_{i\in I}\sum_{j\in J} \left(\alpha_i h_i c_{ij} - \mu_i\right) Y_{ij}^R \\ & - \sum_{j\in J} s_j^C X_j^C \end{aligned} \tag{2.16}$$

(Note that λ and μ are the Lagrange variables that correspond to the relaxed forward assignment and reverse assignment constraints, respectively.)

Observe that in the absence of constraints (2.11) and (2.12) the problem decomposes into separable forward and reverse subproblems. We incorporate this constraint into the procedure for solving the problem for fixed λ and μ values using the algorithm described in Figure 2.3

Lower Bound

For given values of the Lagrange multipliers, λ and μ, the formulation above can be easily solved using the following algorithm:

1. Initialize all decision variables X_j^F, X_j^R, Y_{ij}^F, Y_{ij}^R, equal to 0

2. Calculate the "value" of locating a distribution center in each direction for all candidate

```
Lagrangian Relaxation Algorithm:
        Initialize Lagrange multipliers
        While (termination criteria not met)
        {
                Increment iteration counter
                Calculate lower bound solution
                Calculate upper bound solution
                Check optimality gap
                Calculate step size
                Update Lagrange multipliers
        }
```

Figure 2.3: Lagrangian Relaxation Algorithm

sites. The value functions are:

$$V_j^F = \sum_{i \in I} \min\{0, h_i c_{ij} - \lambda_i\} \tag{2.17}$$

$$V_j^R = \sum_{i \in I} \min\{0, \alpha_i h_i c_{ij} - \mu_i\} \tag{2.18}$$

3. Repeat for all candidate sites j:

   (a) If ($V_j^F + f_j < 0$), set $X_j^F = 1$

   (b) If ($V_j^R + \beta_j f_j < 0$), set $X_j^R = 1$

   (Note that this step is executed whether or not a foward site is located in step 3a)

   (c) If ($V_j^F + f_j + V_j^R + \beta_j f_j$ - $s_j^C < 0$) (it is advantageouse to locate a joint facility)

   or, If ($V_j^F + f_j + V_j^R + \beta_j f_j$ - $s_j^C < V_j^R + \beta_j f_j$) (a joint facility is more beneficial than a stand-alone reverse facility)

   or, If ($V_j^F + f_j + V_j^R + \beta_j f_j$ - $s_j^C < V_j^F + f_j$) (a joint facility is more beneficial than a stand-alone forward facility)

Then set $X_j^F = X_j^R = X_j^C = 1$

4. At least one forward and one reverse distribution center must be located in each iteration

   (a) If there are no open forward facilities and no open reverse facilities, then calculate:

      i. m = $\arg\min_{j \in J} (V_j^F + f_j)$ and $\theta = V_m^F + f_m$ (this corresponds to locating a forward facility as a standalone facility)

      ii. n = $\arg\min_{j \in J} (V_j^R + \beta_j f_j)$ and $\phi = V_n^R + \beta_n f_n$ (this corresponds to locating a reverse facility as a standalone facility)

      iii. p = $\arg\min_{j \in J} (V_j^F + f_j + V_j^R + \beta_j f_j - s_j^C)$ and $\omega = V_j^F + f_j + V_j^R + \beta_j f_j - s_j^C$ (this corresponds to locating both a forward and a reverse site at a location)

   If $\omega < \theta + \phi$, then set $X_p^F = X_p^R = X_p^C = 1$.

   Else set $X_m^F = 1$ and $X_p^R = 1$.

   (b) If there is at least one open forward facility, but no open reverse facilities, then let F be the set of all open forward facilities and calculate:

      i. m = $\arg\min_{j \in F} (V_j^R + \beta_j f_j - s_j^C)$ and $\theta = V_m^R + \beta_m f_m - s_m^C$ (this corresponds to locating a reverse facility at a site that already has a forward facility)

      ii. n = $\arg\min_{j \in J \setminus F} (V_j^R + \beta_j f_j)$ and $\phi = V_n^R + \beta_n\ f_n$ this corresponds to locating a stand alone reverse facility at a site that does not have a forward facility)

      iii. p = $\arg\min_{j \in J} (V_j^F + f_j + V_j^R + \beta_j f_j - s_j^C)$ and $\omega = V_j^F + f_j + V_j^R + \beta_j f_j - s_j^C$ (this corresponds to locating both a forward and a reverse site at a location that does not have a forward facility)

If $\omega < \theta$ and $+ \omega < \phi$, then set $X_p^F = X_p^R = X_p^C = 1$.

Else if $\omega < \phi$, set $X_m^R = X_m^C = 1$, else set $X_n^R = 1$.

(c) If there is at least one open reverse facility but no open forward facilities, then locate a forward site using a procedure identical to that of step 4b with the forward and reverse indices exchanged.

5. Determine the assignment of demand nodes to open distribution centers:

   (a) If $X_j^F = 1$ and ($h_i c_{ij}$ - $\lambda_i < 0$), set $Y_{ij}^F = 1$. Else, set $Y_{ij}^F = 0$.

   (b) If $X_j^R = 1$ and ($\alpha_i h_i \gamma_{ij} c_{ij}$ - $\mu_i < 0$), set $Y_{ij}^R = 1$. Else, set $Y_{ij}^R = 0$.

   (c) Repeat steps 5a and 5b for all (i,j) pairs

Using the values of the decision variables obtained through this algorithm, we calculate the lower bound using the objective function of the lagrangian relaxation objective (2.13)

Upper Bound

The upper bound is calculated at each iteration using the following algorithm:

1. All location variables X_j^F and X_j^R retain the same status as determined in the lower bound algorithm

2. All demand nodes are assigned to the single closest open facility in both the forward and the reverse directions.

Using the values of the decision variables obtained through this algorithm, we calculate the upper bound using objective function (2.1)

Lagrange Multiplier Initialization and Updating

The Lagrange multipliers λ and μ are initialized to the following values at iteration 0:

$$\lambda_j^{(0)} = \frac{10 * (\sum_{j \in J} f_j + h_j)}{|J|} \tag{2.19}$$

$$\mu_j^{(0)} = \frac{10 * (\sum_{j \in J} \beta_j f_j + \alpha_j h_j)}{|J|} \tag{2.20}$$

These particular initialization formulae perform well in practice, but others may also be used.

The Lagrange multipliers λ and μ are updated at each iteration using standard subgradient optimization (Fisher, 1981 and Fisher, 1985) with the search direction modified as proposed by Crowder (1976). The following formulae are used at each iteration:

$$\lambda_i^{(n+1)} = \max\{0, \lambda_i^{(n)} + \text{stepsize}^{(n)} \text{direction}_{forward}(i)^{(n)}\} \tag{2.21}$$

$$\mu_i^{(n+1)} = \max\{0, \mu_i^{(n)} + \text{stepsize}^{(n)} \text{direction}_{reverse}(i)^{(n)}\} \tag{2.22}$$

where the superscript (n + 1) indicates that this update is for the $(n+1)^{st}$ iteration.

Iteration Step Size Calculation

To calculate the step size at iteration n, we first define a *direction* for each node $i \in I$:

$$direction_{forward}(i)^{(n)} = \left(1 - \sum_{j \in J} Y_{ij}^F\right) + C * direction_{forward}(i)^{(n-1)} \tag{2.23}$$

$$direction_{reverse}(i)^{(n)} = \left(1 - \sum_{j \in J} Y_{ij}^R\right) + C * direction_{reverse}(i)^{(n-1)} \tag{2.24}$$

where $direction_{forward}(i)^{(0)} = direction_{reverse}(i)^{(0)} = 0 \; \forall i$ and C is a Crowder damping constant (Crowder, 1976). Typically we set C = 0.3.

The step size is calculated at each iteration n using the following formula:

$$\text{stepsize}^{(n)} = \delta \frac{\text{Z}^{*}_{bestupperbound} - \text{Z}^{(n)}_{lowerbound}}{\sum_{i \in I} \{(\text{direction}_{forward}(i)^{(n)})^2 + (\text{direction}_{reverse}(i)^{(n)})^2\}} \tag{2.25}$$

where δ is initialized to 2.0 and is halved if there are k (typically $k = 12$) consecutive iterations in which the lower bound does not increase, and $\text{Z}^{*}_{bestupperbound}$ is the lowest upperbound calculated in iterations 1 through n.

Termination Criteria

The Lagrangian relaxation algorithm utilizes three termination criteria:

- *Maximum Iteration Count* the procedure is stopped once 10,000 iterations have been performed
- *Step Size Multiplier* the procedure is stopped if the step size multiplier δ is less than a defined epsilon (typically when $\delta \leq 0.0001$)
- *Optimality Gap* the procedure is stopped if the optimality gap between the lower and upper bounds is less than a defined tolerance (typically when gap ≤ 0.01

2.3.3 The Forward Dominant Model

In the forward dominant model, the set of reverse distribution centers (product collection centers) is restricted to be a subset of the forward distribution centers. (See Figure 2.2 for a sample configuration of a forward dominant network.) For this model, replace objective function 2.1 with the following objective function:

$$\text{Minimize} \sum_{j \in J} f_j X^F_j + \sum_{i \in I} \sum_{j \in J} h_i c_{ij} Y^F_{ij} + \sum_{j \in J} \beta_j f_j X^R_j + \sum_{i \in I} \sum_{j \in J} \alpha_i h_i c_{ij} Y^R_{ij} \tag{2.26}$$

Thereby eliminating the term related to co-location cost savings. We also remove constraints (2.8), (2.11), and (2.12) and impose the following constraint to ensure that reverse facilities are only located at sites that have an open forward facility:

$$X_j^F \geq X_j^R \qquad \forall j \in J \tag{2.27}$$

The forward dominant model is applicable to products in the early part of their lifecycle when few returns are likely and therefore a firm requires that reverse facilities will be co-located with forward facilities. This model can also be solved using Lagrangian relaxation using a straightforward modification of the algorithm outlined for the co-location model.

2.3.4 The Reverse Dominant Model

The reverse dominant model considers the case in which the set of forward distribution centers is a subset of the product collection facilities. Reverse dominant models are prevalent in recycling networks where there is a large number of collection sites and a few centralized processing facilities. (See Figure 2.2 for a sample configuration of a reverse dominant network.) The reverse dominant model is identical to the forward dominant model except that the sense of inequality (2.27) is reversed, thereby ensuring that forward facilities are located only at sites that house a reverse facility:

$$X_j^F \leq X_j^R \qquad \forall j \in J \tag{2.28}$$

Such a model is applicable to products nearing the end of their useful life, when the number of returns is likely to be on the same order of magnitude, or even greater than, the number of primary sales (forward flow) of the product. This too can readily be solved using Lagrangian relaxation.

2.4 Network Similarity Metric

In addition to considering the financial impact of integrating forward distribution and reverse collection activities within a supply chain, it is also important to measure the extent to which the integrated and sequential (or independent) facility siting plans differ from one another. The current literature simply counts the number of different facility sites and reports this value as a measure of the difference between two location plans. Such a measure can be misleading. Consider three different location plans, each using four facilities on the contiguous 48 states. These plans are identified in Table 2.2 and Figure 2.4 below.

Plan 1	**Plan 2**	**Plan 3**
Chicago, IL	Chicago, IL	Milwaukee, WI
New York, NY	New York, NY	Allentown, PA
Houston, TX	Houston, TX	Pasadena, TX
Los Angeles, CA	Jacksonville, FL	Burbank, CA

Table 2.2: 3 Example Facility Location Plans

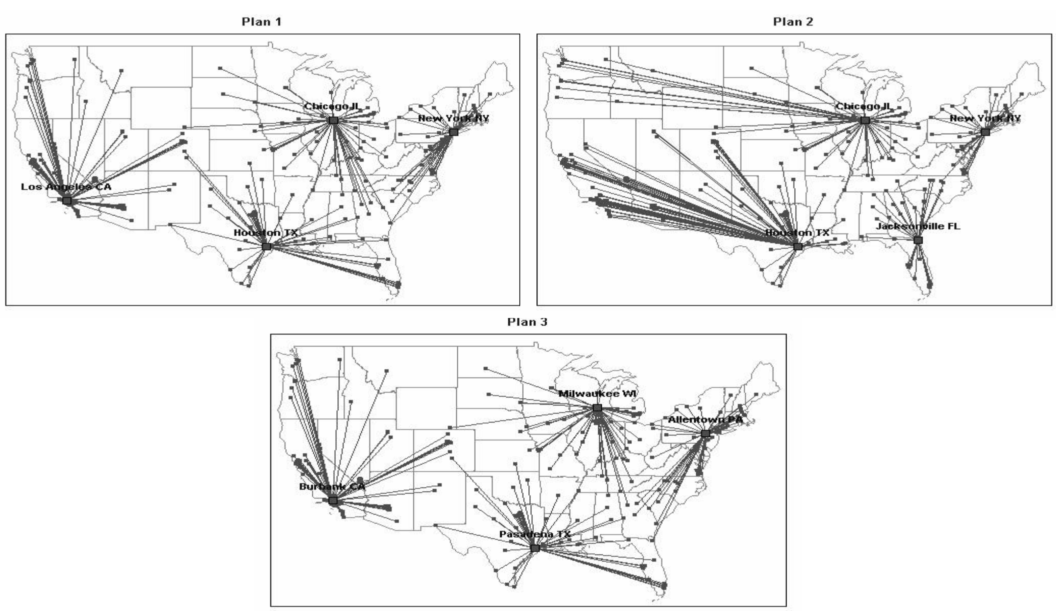

Figure 2.4: 3 Example Facility Location Plans

Plans 1 and 2 share 3 sites in common and differ by only one facility location. However, this single facility has been moved from Los Angeles, CA on the west coast of the U.S. to Jacksonville, FL, on the east coast, some 2153 miles from Los Angeles. A visual examination of the two plans reveals a striking difference. Plans 1 and 3 share no sites in common. However, Milwaukee is only 86 miles from Chicago, Allentown is 81 miles from New York, Pasadena is a mere 16 miles from Houston, and Burbank is only 7 miles from Los Angeles. The total of these four distances is only 190 miles, or less than 10 percent of the distance between Jacksonville and Los Angeles. Visual inspection of plans 1 and 3 indicates that they are indeed quite similar and one could argue that they are in fact more similar than plans 1 and 2.

We introduce a new metric to quantify the differences between any two location plans by capturing the total distance between them. The network similarity measure between two solutions m and n, with known locations $\mathbf{X_m}$ and $\mathbf{X_n}$ respectively, is given by the optimal value of the objective function of the following optimization problem:

$$\text{Minimize} \sum_{i \in \mathbf{X_m}} \sum_{j \in \mathbf{X_n}} c_{ij} Y_{ij} \tag{2.29}$$

$$\text{Subject to} \sum_{i \in \mathbf{X_m}} Y_{ij} \geq 1, \forall j \in \mathbf{X_n} \tag{2.30}$$

$$\sum_{j \in \mathbf{X_n}} Y_{ij} \geq 1, \forall i \in \mathbf{X_m} \tag{2.31}$$

$$Y_{ij} \geq 0, \forall i \in \mathbf{X_m} \quad \forall j \in \mathbf{X_n} \tag{2.32}$$

In this model Y_{ij} is an assignment variable that assigns a facility in $\mathbf{X_m}$ to a facility in $\mathbf{X_n}$ and c_{ij} is the cost of this assignment - which we take as the distance between the two facilities. The problem assigns every facility in solution m to a facility in solution n in such a way that the total assigned distance is minimized. In the event that $|\mathbf{X_m}| \neq |\mathbf{X_n}|$, the model allows multiple facilities in the solution with more locations to be assigned to a single facility in the solution with fewer sites. This is a variant on a simple transportation problem. It is clear that it has an all-integer solution. We solve it using an out-of-kilter flow algorithm (Ahuja et al., 1993).

Since the magnitude of the similarity measure defined above is data set dependent, we report results using a network similarity ratio that standardizes the network impact across data sets. The similarity ratio is the total similarity measure defined above (total distance between all locations) divided by the maximum distance between any two locations in the data set. A similarity ratio of > 0.5 indicates a fair amount of dissimilarity between the two sets being compared, while a similarity ratio < 0.1 indicates that the two sets are fairly similar.

2.5 Computational Results

An extensive computational study was conducted to analyze the effects of key parameters on both the value and impact of integrated bidirectional network design. The results of this study are summarized below.

2.5.1 Numerical Experiments

The solution algorithms were coded in C and run on a Pentium 2.0 GHz PC.

Obtaining Integrated and Sequential Solutions

We solve the models as they appear above to obtain an *integrated* solution. To obtain a *sequential* solution for the forward dominant model we first solve an uncapacitated fixed charge problem (UFC) for forward distribution. Next we solve an independent UFC for reverse collection while restricting the candidate reverse facility sites to be the sites located in the forward UFC solution. A sequential solution for the reverse dominant formulation is obtained similarly. The integrated solution for the co-location model is contrasted with an *independent* solution formed by combining the solutions of two independent forward and reverse UFC instances. If the combined independent solution has a co-located facility at site k, the co-location incentive value, s_k^C , is subtracted from the objective function.

Data Sets

The models were tested on two relatively large data sets. The first data set contains 150 demand and candidate facility site locations and is based on the 150 largest cities in Europe. The second data set contains 263 demand and candidate facility site locations and is based on the 263 largest cities in the contiguous 48 US states. Demand data for these locations is proportional to population. The distribution of demand points is shown below in Figure 2.5.

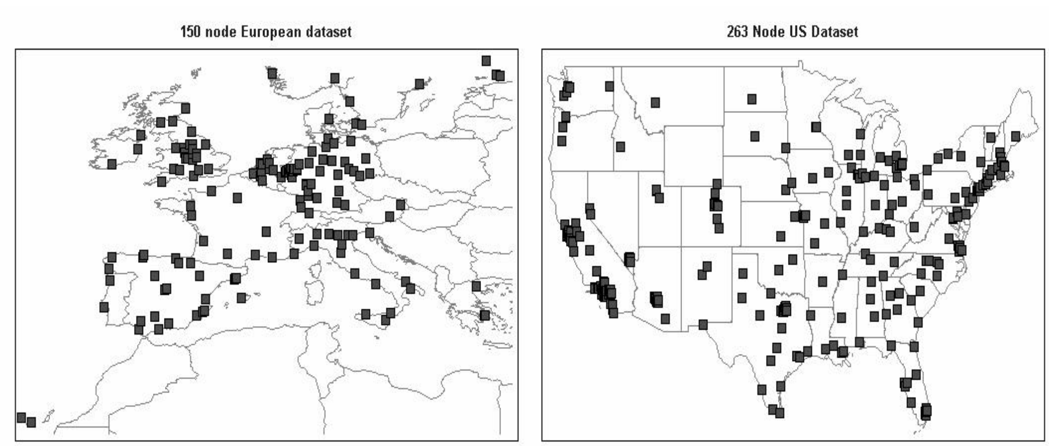

Figure 2.5: Data Sets

Experimental Design

An experiment on return rate (α), reverse fixed cost ratio (β), and reverse transportation ratio (γ) was performed for all three models using the two datasets. Observe that the forward dominant model is appropriate for systems in which return rate $\alpha \leq 1$ while the reverse dominant model is appropriate for systems in which the return rate is relatively high, so the experiments for these two models were conducted using the appropriate return rate values.

Table 2.3 gives the specific values that were tested for each parameter. Each value of a specific parameter was tested for low-low, low-high, high-low, and high-high values of the other two parameters. For example, a return rate value of $\alpha = 0.6$ was tested on four combinations of reverse fixed and transportation cost ratios: ($\beta = 0.25$, $\gamma = 0.50$), ($\beta = 0.25$, $\gamma = 2.0$), ($\beta = 0.75$, $\gamma = 0.50$), and ($\beta = 0.75$, $\gamma = 2.0$). Testing these combinations over both data sets and for both the integrated and the sequential model gives 11,280 problem instances. The co-location savings factor, s_j, was set to s_j=0.25*Min(f_j, $\beta_j f_j$) in all cases below.

Model	α	β	γ	Summary
Co - Location	$0 \leq \alpha \leq 5$ increments of 0.05	$0 \leq \beta \leq 2$ increments of 0.05	$0 \leq \gamma \leq 5$ increments of 0.05	3840 runs
Forward Dominant	$0 \leq \alpha \leq 1$ increments of 0.01	$0 \leq \beta \leq 2$ increments of 0.05	$0 \leq \gamma \leq 5$ increments of 0.05	3840 runs
Reverse Dominant	$0.75 \leq \alpha \leq 5$ increments of 0.05	$0 \leq \beta \leq 2$ increments of 0.05	$0 \leq \gamma \leq 5$ increments of 0.05	3600 runs

Table 2.3: Experimental Design

Performance Metrics

Two performance metrics, (i) the value of integration and (ii) the network impact of integration, were used to evaluate the differences between integrated and sequential design decisions. The network similarity value is calculated as described in Section 2.4. A value of integration is calculated for the total objective cost impact of integrated decision making. This metric captures the total cost penalty of locating facilities sequentially versus locating facilities in an integrated framework. The value of integration (or sequential penalty) is calculated as:

$$\text{Value} = \text{Sequential Penalty} = \frac{Z^*_{sequential} - Z^*_{integrated}}{Z^*_{integrated}} \tag{2.33}$$

where $Z^*_{sequential}$ and $Z^*_{integrated}$ are the optimal total costs of sequentially locating facilities and locating facilities under the integrated formulation, respectively. Table 2.4 shows the average and maximum value of integration and resulting network impact for the three models.

Model	Data Set	Value of Integration		Forward Similarity ratio		Reverse Similarity ratio	
		Average	Maximum	Average	Maximum	Average	Maximum
Co - Location	150 nodes	0.60%	2.48%	0.00	0.00	0.00	0.04
	263 nodes	1.41%	3.87%	0.02	0.15	0.13	0.67
Forward Dominant	150 nodes	2.42%	30.07%	0.27	1.42	0.27	1.42
	263 nodes	6.02%	28.59%	0.15	1.20	0.21	1.20
Reverse Dominant	150 nodes	0.08%	4.64%	0.02	0.47	0.03	0.47
	263 nodes	0.10%	1.69%	0.06	0.24	0.05	0.33

Table 2.4: Value and Impact of Integration Summary

2.5.2 Sensitivity Analysis Results

Co-Location Model Analysis

Our analysis indicates that the reverse fixed cost ratio (β) has the most impact of the three parameters (α, β, γ) on the value of integration in the CL model. The value of integration increases slightly with increasing β values for all α-γ combinations. In the extreme case, the value of integration is approximately 3.5%. Results of our analysis also show that the integrated and independent solution methods produce highly similar forward distribution networks but moderately dissimilar collection networks for all parameter values.

Forward Dominant Model Analysis

Results of our experiments show that the reverse transportation cost ratio (γ) is the critical cost factor in the design of forward dominant distribution systems. The value of integration increases significantly with increasing γ values for all α-β combinations, up to 30% for very high reverse transportation costs. Our analysis also shows that the integrated and sequential formulations produce highly dissimilar networks when all three parameter values are relatively high ($\alpha > 0.6$, $\beta > 0.6$, $\gamma > 2.0$), resulting in similarity ratios close to 1.0 for both the forward and reverse networks. When these values are low to moderate ($\alpha < 0.4$, $\beta < 0.4$, $\gamma < 0.5$), our results show that reverse collection networks can be efficiently incorporated into existing forward distribution networks.

Reverse Dominant Model Analysis

Analysis of the RD experiments shows that the value of integration is minimal with respect to all α-β-γ parameter values and is consistently less than 2% across our entire study. We also see that the integrated and sequential networks are consistently similar for all parameter values tested.

2.5.3 Algorithm Performance

Throughout the experiments, the Lagrangian solution algorithm proved to be both quick and effective. Table 2.5 below shows the speed of this algorithm with average solution times under 1 second for the smaller data set and under 5 seconds for the larger data set with no instance taking longer than 105 seconds. The results in this table also illustrate the effectiveness of this solution procedure, with average optimality gaps of 0.004% and no instance larger than 0.1%.

Model	Data Set	Solution Time (sec)		Optimality Gap	
		Average	Maximum	Average	Maximum
Co - Location	150 nodes	0.93	6.85	0.01%	0.09%
	263 nodes	4.72	69.30	0.00%	0.09%
Forward Dominant	150 nodes	0.84	16.68	0.01%	0.10%
	263 nodes	5.24	104.16	0.00%	0.08%
Reverse Dominant	150 nodes	0.65	2.55	0.00%	0.07%
	263 nodes	6.49	12.05	0.00%	0.05%

Table 2.5: Lagrangian Algorithm Solution Times and Optimality Gaps

In addition to solving these models using the Lagrangian procedure, we used CPLEX to solve a small number of problem instances (27 instances for each model and data set). Using default values, CPLEX was able to solve all 162 instances to optimality but took significantly longer than our solution algorithm. The average CPLEX to Lagrangian solution time ratio for the 150 node data set is 1064, or roughly 17.7 minutes of CPLEX time for every second of Lagrangian solution time. The average solution time ratio for the 263 node data set is 2181, or roughly 36.4 minutes of CPLEX time for every second of Lagrangian time. The worst case ratios are 4955 and 5459 (roughly 1.5 hours of CPLEX time for every second of Lagrangian solution time) for the 150 and 263 node data sets, respectively. Detailed results for these runs, including forward and reverse similarity results, are included in Appendix A in Tables A.1 and A.2. Figure 2.6 below shows the ratio of CPLEX time to Lagrangian solution time for these problems.

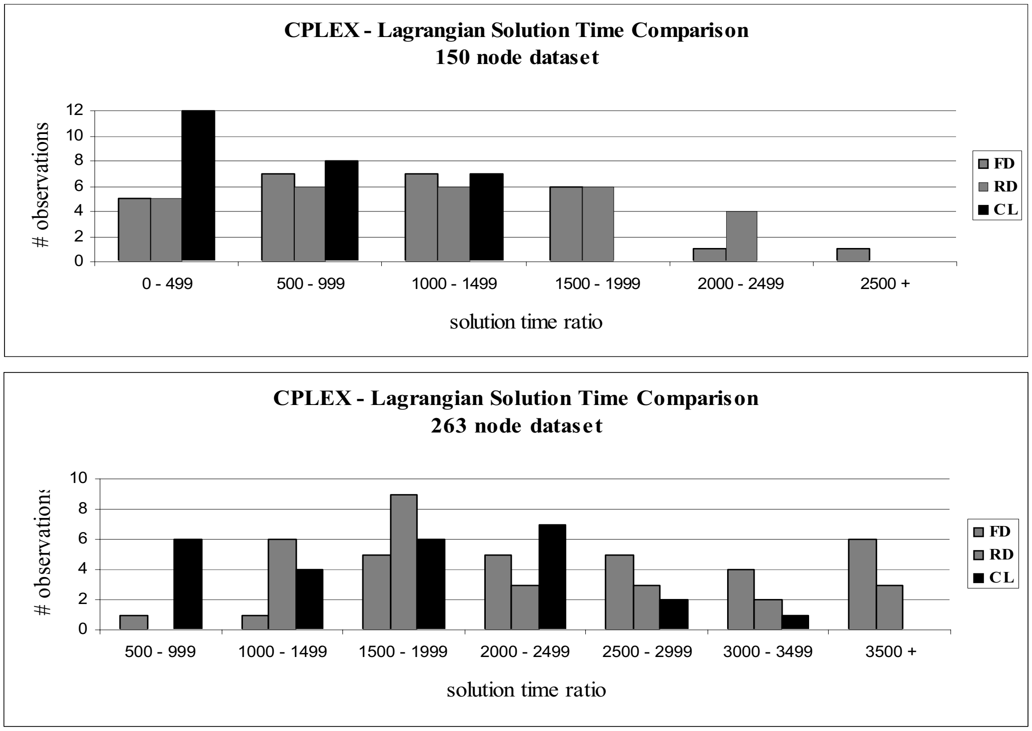

Figure 2.6: CPLEX vs Lagrangian Solution Time Comparison

2.6 Conclusions and Directions for Future Research

In this chapter, we have outlined three extensions of the traditional fixed charge facility location model that account for both forward and reverse flows of product. The forward dominant model assumes that each reverse facility will be co-located at a forward site. Such a model is most applicable to products in the early stages of the product lifecycle when returns are likely to be a small fraction of the total demand. Our computational study suggests that, in such cases, integrated network design is critical if excessive costs are to be avoided as sequential network design can cost up to 30% more than integrated network design. These two approaches also produce highly dissimilar networks as shown in Table 2.4. The reverse dominant model assumes that each forward site is co-located at a reverse facility. This model is most appropriate for products at the tail end of their lifecycle, when return rates are likely to be high relative to the product demand. Our study shows that both the value and impact of integrated network design are minimal for a reverse dominant network structure. Finally, we presented a model with cost incentives for co-location, but which does not assume anything about co-location *a priori*. This model is most appropriate for products in the middle of their lifecycle. Integrated design is moderately valuable in this case but can have a significant impact on the network structure.

We outlined a Lagrangian relaxation approach for solving the models. Computational studies show that the approach is effective (resulting in small optimality gaps) and fast (with solution times under 2 minutes for problems with 263 nodes). We also presented a new measure of the similarity between siting plans that captures the distance between facilities instead of simply counting the number of facilities that differ between the two plans.

At least three directions for future research are suggested by this initial study. First, it would be interesting to develop a time-integrated model that captures the optimal evolution of a closed loop logistics network over the entire span of a product's lifecycle. Second, we would like to extend the models to encompass multiple products at different stages of their lifecycles with differing return rates and processing costs. Third, it would be desirable to make the return rate at each node an endogenous function of the cost of returning the product or of the distance to the nearest return facility. This would enable us to begin to capture the impact of network design in particular of the reverse network design on return rates and possible compliance with mandated return programs.

Chapter 3

The Effect of Lifetime Buys on Warranty Repair Operations

Lifetime buys are a common practice in the electronics and telecommunication industries. Manufacturers often procure their repair parts inventory in one order to support the spare part needs of a product for the duration of its warranty repair period. Lifetime buy decisions are driven by product lifecycle dynamics and the nature of manufacturer-supplier relationships. In this chapter we consider a repair operation in which defective items under warranty are returned to a manufacturer who either repairs these items using its spare parts inventory or replaces each defective unit with a new product. The manufacturer has a single opportunity to procure its spare parts inventory from its supplier. The manufacturer must decide on the number of parts to order and determine if returned products should be repaired using this inventory, or replaced with new products. We show how fixed repair capability costs, variable repair costs, inventory holding costs, and replacement costs affect a firm's optimal repair and replacement decisions. The model is used to gain insights for products from a major mobile device manufacturer in the United States.

3.1 Introduction

Lifetime buys of spare parts are common practice in the electronics and telecommunication industries but have not yet been studied in the operations management literature. A lifetime buy is a process by which a manufacturer places a single order for all repair parts inventory to support its products under warranty. In this chapter, we discuss the optimal management of the lifetime buy process and its implications for maintaining cost effective warranty repair operations. Our work is motivated by current operational challenges in the repair divisions of two major U.S. mobile device (mobile phones, personal digital assistants, and related products) manufacturers. In this industry, lifetime buy decisions are common and are typically driven by two primary factors: product lifecycle dynamics and the nature of manufacturer supplier relationships. We study this problem from the perspective of the electronics and telecommunications product manufacturers but acknowledge that a myriad of other factors may drive lifetime buy decisions in other industries.

Recent technological innovations in the electronics and telecommunications industries have enabled rapid product development, which in turn has led to short lifecycles for these products. In these industries it is not uncommon for a manufacturer to complete production of an item even before it is available in retail outlets. Low cost mobile phones and digital cameras are two examples of manufactured goods for which this is the case. This allows little time for firms to observe demand and subsequent return rates of defective products before procuring the repair parts inventory that must support repair operations during the product's warranty period. In addition, these products span a broad range of functionality, resulting in little component commonality across multiple product lines and/or multiple generations of a single product. This provides manufacturers with limited opportunity to utilize common components to offset the variability in each product's part needs. A short planning horizon coupled with a large number of components can make it very difficult for a manufacturer to manage its spare parts inventory and to plan repair operations for these products. Finally, it is common practice for companies that produce high-end specialty products with considerably longer lifecycles of ten to twenty years (e.g. government communication and/or security systems) to procure all parts for manufacturing and repair at once due to low volumes and high supplier setup costs driven by the specialized nature of these products.

Lifetime buys of spare parts are also affected by the nature of manufacturer-supplier relationships in these industries. The overwhelming majority of components in mobile devices are procured from small Asian suppliers who are either dedicated to a single manufacturer or accept contracts from multiple firms. Dedicated suppliers are central to firms that implement just-in-time manufacturing processes. Such suppliers often incur expensive retooling costs to produce new parts when a manufacturer switches its production to a new product. These high retooling costs make small batch orders cost prohibitive; as a result, suppliers are often unwilling to produce items once primary manufacturing needs have ended. Before the supplier changes its setup, it will often require a lifetime buy on the parts it is currently producing for the manufacturer. Conversely, non-dedicated suppliers may be unwilling to accept long-term contracts or unable to honor future contracts due to limited capacity and competing bids from other firms. In these scenarios, the manufacturer may choose to initiate a lifetime buy for spare parts.

To be competitive in these industries, firms must offer attractive warranty policies to consumers, under which defective items are either repaired resulting in "like-new" condition or replaced by a brand-new product. The role of product warranties from a management strategy perspective is discussed in Murthy and Blischke (2000). Discussions with repair managers at mobile device manufacturers reveal that manufacturers spend anywhere from 1% to 5% of their revenues on repairing defective warranty items. With such a significant amount of money being spent on warranty repair costs, firms are now focusing on controlling their spare parts inventory and reevaluating policies that determine if a claim is satisfied via repair or replacement.

The repair process may be performed by the manufacturer or outsourced to a third party service provider. In most cases, the repair operation is an expensive practice as the firm incurs the cost of maintaining testing equipment, technician training and knowledge management, and all of the costs associated with managing the repair parts inventory. The other option available to a manufacturer is to replace the defective item with a comparable (possibly next-generation) product but this is also an expensive practice. The decision of whether or not to repair or replace a defective item to satisfy a warranty claim is further complicated by the single parts procurement opportunity of a lifetime buy decision.

This work is motivated by our collaboration with two firms from the electronics and mobile device industries that currently use lifetime buys to support their warranty repair operations. One of the firms is a U.S. market leader in converged handheld devices specializing in personal digital assistants and multi-functional mobile devices with telecommunication and organizational capabilities. The other firm is a global leader in the telecommunications industry with expertise in network infrastructure, mobile phones, and wireless communication capabilities. In this chapter, we develop a general modeling framework to analyze the impact of key cost drivers for repair parts inventory lifetime buy decisions. We test our modeling insights using warranty repair data from the motivating firm in Section 3.4.

This chapter contributes to the existing literature on repair operations by considering the following questions within the lifetime buy context described above:

1. Under what conditions should a returned defective item be repaired or replaced, and how does this decision change throughout the warranty period?

2. What is the corresponding optimal order quantity for repair parts?

To answer these questions, we develop a deterministic continuous time model that examines how fixed repair capability costs, variable repair costs, inventory holding costs, and replacement costs affect a firm's optimal repair policy and order quantity for a single repair type. We show how to calculate the true cost of repair during the planning period, and identify a unique "switching time" before which the manufacturer should repair returned items and after which returns should be satisfied with replacement products. We extend this work to consider repair capacity constraints as well as the impact of having shared facilities to support multiple product lines.

The rest of this chapter is organized as follows. In Section 3.2 we discuss the contribution of this chapter to repair parts, warranty costing, and traditional inventory management literatures. The model is developed and extensions are presented in Section 3.3. In Section 3.4 we discuss the industry case that motivated the study. The modeling framework is verified using data from current repair operations at this firm. A final discussion including insights and directions for future research comprises Section 3.5.

3.2 Literature Review

This chapter contributes to three streams of research, each of which we review below. The first is the repairable inventory literature which is concerned with designing inventory systems to serve large scale repair operations. The second stream of literature examines warranty repair operations and focuses on designing product warranty contracts and determining the cost of these agreements. The third stream of literature is classical inventory management research.

The repairable inventory literature has a long, rich history that has yet to explicitly consider the impact of lifetime buys. This area of research typically considers internal repair systems (where the firm repairs items for its own use and not for end consumers) for products that have a relatively long life cycle, such as aircraft, military equipment, and mass transportation equipment. Sherbrookes METRIC model (Sherbrooke, 1968), which determines repair parts inventory levels and their allocation to achieve a minimum service level, was the first model to deal specifically with repair operations. Many extensions to this seminal model have been considered since then, spanning a variety of applications and solution techniques. A comprehensive review by Nahmias (Nahmias, 1981) summarizes a number of studies of repairable inventory systems that consider continuous review, periodic review, and queuing system-based models. A more recent review (Guide and Srivastava, 1997) outlines extensions that have been considered since Nahmias' paper, including multi-echelon models. The models developed in this stream of literature focus on determining stocking quantities at various repair facilities in order to examine the interaction between service level and total cost. While this research provides valuable insights, these models are best suited to applications in longer lifecycle products with large repair networks and multiple parts procurement opportunities. These models are therefore inappropriate for the problem that we consider which serves very short lifecycle products, a single repair facility, and a single parts procurement opportunity.

The second relevant literature stream considers new product warranty policies. Within this stream of literature are a number of approaches for estimating the costs of various warranty policies. One approach defines a Markov process in which multiple components can fail, each with its own cost of repair or replacement (Balachandran, Maschmeyer, and Livingston, 1981). Balcer and Sahin (1986) consider a warranty policy in which failed items are repaired or replaced at a cost that depends on the age of the product. Another approach considers both the remaining length of the warranty period and the state of deterioration of the product in deciding whether to service a failed item by repairing it or replacing it with a new product. Zuo, Liu, and Murthy (2000) consider multiple states of failure and multiple discrete decision periods in showing how to minimize the expected service cost to the manufacturer based on these two parameters when there are multiple parts procurement opportunities. Ja et al. (2001) estimate expected warranty costs over the lifecycle of a complex expensive product when minimal repairs are performed and costs are age-dependent. A review of product warranty literature including warranty cost analysis and the implications of warranty policies for marketing and logistics management is provided by Murthy and Djamaludin (2002). While this literature provides insight into modeling the repair/replace decision, it does not explicitly consider the impact of lifetime buys, which significantly changes both the modeling approach and the characterization of the solution.

Classic inventory management literature represents a third stream of research that is directly related to the models that we develop in Section 3.3. By now (or upon reading Section 3.3), the reader may have observed some similarities between our model and other well-known inventory problems such as the economic order quantity (EOQ) model (Harris, 1913) and the newsvendor model (Wadsworth, 1959). As in to the original EOQ model proposed by Harris, the inventory problem we model considers fixed costs and inventory holding costs incurred over time when determining an optimal order quantity. However, the EOQ model assumes an infinite horizon, constant demand rate, and fixed ordering costs. In contrast, our model assumes a finite horizon while allowing a variable demand rate. Like the newsvendor model, our problem considers the impact of a single procurement opportunity and the penalty associated with not having sufficient inventory to perform repairs. Among the differences between the original newsvendor model and our problem are the fixed cost and inventory holding costs incurred over time as well as demand uncertainty. As one of the oldest streams of literature in operations management research, classic inventory management models such as the EOQ and newsvendor problems have been extended to consider a myriad of additional factors. The purpose of this section is not to review these contributions, but to acknowledge the insight provided by this work and claim that the lifetime buy model presented in Section 3.3 effectively blends the considerations of these two seminal inventory models. The reader is referred to a recent review by Petruzzi and Dada (1999) for extensions of the newsvendor problem. A more comprehensive review of classic inventory management problems is provided by Porteus (1990, 2002).

3.3 Deterministic Models For Repair Inventory Lifetime Buys

In this section, we develop deterministic continuous time models to investigate a manufacturer's optimal lifetime buy decision for a single product, as well as the firm's decision to repair or replace a returned defective item. The basic model we develop in Section 3.3.1 considers the problem of determining the optimal repair period and lifetime buy decision for a single product with uncapacitated repair facilities. This is the simplest form of this model available. In Section 3.3.2, we extend this model to consider the effects of capacity constraints on the repair operation. Finally, in Section 3.3.3 we consider two products with shared repair facilities and discuss how the decision for each product is complicated by the shared resources. Proofs and derivations for Sections 3.3.2 and 3.3.3 are presented in the appendices.

All models developed below make the following assumptions:

1. Defective products arrive to a service center over a known time period [0,T]; The return arrival rate is a known, positive, bounded function of time which is twice differentiable over the planning period.

2. Each returned item is either repaired using the available parts inventory or replaced with a comparable new product.

3. All inventory required to support the repair operation is purchased at the beginning of the planning horizon. For simplicity, we assume a 1 to 1 relationship between returns and repair parts required to service the return.

4. Returns are repaired at the beginning of the planning period.

We begin by identifying key cost components of the manufacturer's problem. We categorize these as fixed costs, variable repair costs, inventory holding costs, and replacement costs. *Fixed* costs are incurred (daily) for the duration of time during which repairs are performed regardless of the number of items repaired during that time period. These costs include facility, employee salary, testing equipment, and overhead costs. *Variable repair* costs are incurred for each item that is repaired. This category includes the unit procurement cost of spare parts as well as the additional cost of cosmetic refurbishing performed to ensure that the repaired items quality is as good as a new unit. *Inventory holding* costs are incurred over the repair planning horizon and represent the carrying cost of repair parts during this time. *Replacement* costs are incurred for each returned item that is replaced with a comparable new product (not repaired). In the cases we consider, replacement products are provided from the manufacturer's current product portfolio on a per-unit cost basis. Note that due to *just in time manufacturing* in these industries, replacement products incur negligible holding costs at repair facilities.

The following notation is used throughout this chapter:

INPUTS

$\lambda(t)$	return arrival rate at time t
T	length of the planning horizon
F	fixed cost per unit time of having repair capability
v	variable cost per item repaired
h	repair parts inventory holding cost per item per unit time
r	unit cost of replacement with a comparable product

DECISION VARIABLES

Q	repair parts order quantity
τ	end of the repair horizon

Qualitatively, we can see the difference in relevant costs between small and large values of τ as shown below in Figure 3.1. Intuitively, a small value of τ is favored over a larger τ when holding costs (h), variable repair costs (v), and fixed costs (F) are high relative to replacement costs (r). Note that for simplicity's sake, in Figure 3.1, we assume a constant rate of inventory depletion (constant return arrival rate).

3.3.1 The Uncapacitated Single Product Model

The deterministic uncapacitated single repair model is the simplest formulation that we develop and provides a foundation for the extensions that follow. When there are no capacity constraints, the repair and replacement horizons are disjoint. This observation gives rise to the concept of a switching time - the time before which returns are repaired and after which they are replaced. In this case, the repair horizon is given by $[0, \tau^*]$ while the replacement horizon is given by $(\tau^*, T]$. This observation follows from a simple marginal analysis argument. Consider scenarios A and B in Figure 3.2. Scenario A shows a repair-

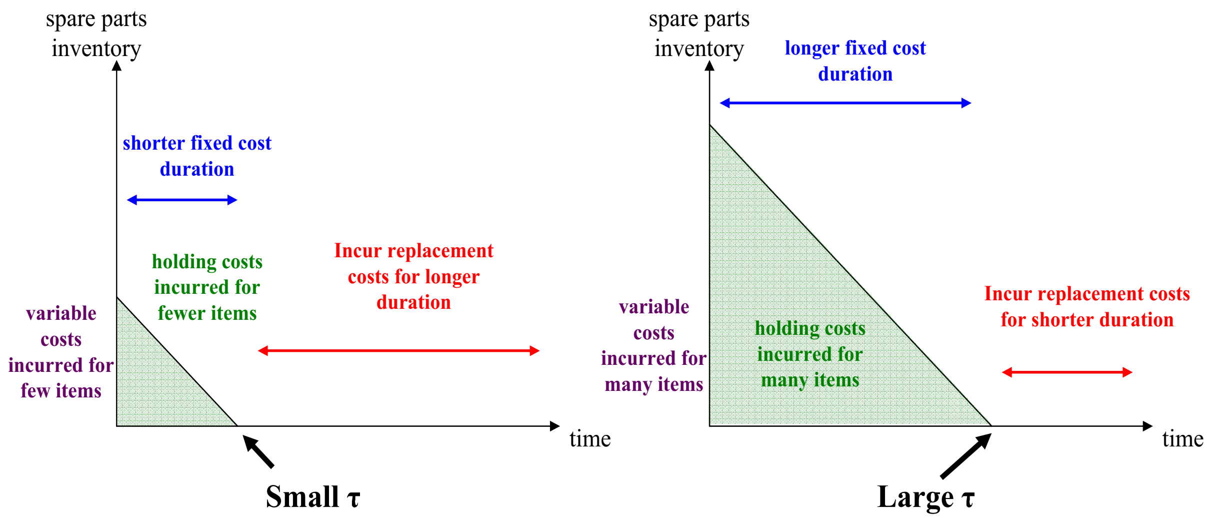

Figure 3.1: Tradeoffs Between Small And Large Values Of τ

replace-repair-replace sample path while scenario B shows a repair-repair-replace-replace sample path. The total number of repaired and replaced items is the same in both scenarios, therefore the total variable and replacement cost incurred in either scenario is equal. Both scenarios model repairs beginning at time 0, but the last repair performed in scenario A occurs at a later time than in scenario B. Therefore the total fixed cost incurred in A is larger than in B. Scenario A requires the manufacturer to hold repair parts inventory for a longer period of time than in scenario B, so the total holding cost in scenario A is greater than in B. The total cost of scenario B is less than the total cost of scenario A, and we see that once the system begins repairing items it will continue to do so as long as there is inventory available. From this argument, we see that the repair and replacement periods are disjoint. Note that a more conservative approach would argue that the fixed cost in scenario A would not be incurred during the first replacement interval and that in fact the fixed cost is the same for both scenarios. If this approach is taken, the total cost for scenario A remains greater than scenario B due to the larger holding cost and the result still holds.

With the repair and replacement horizons defined above, we have the following relation-

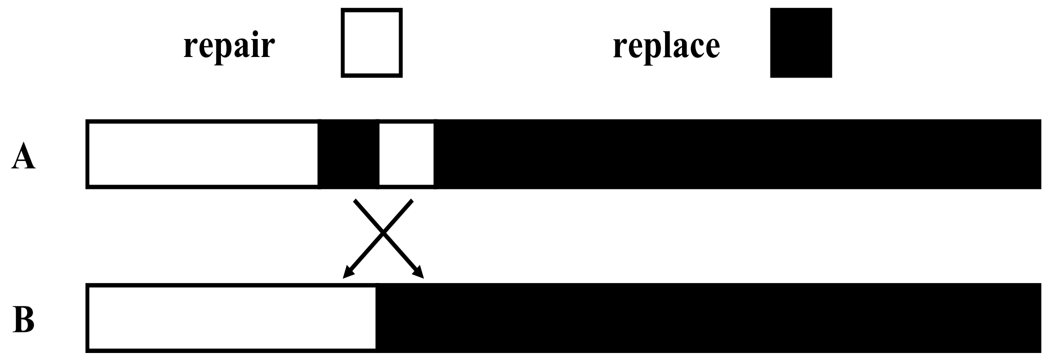

Figure 3.2: Sample Repair and Replacement Scenarios

ship between Q and τ:

$$Q(\tau) = \int_0^{\tau} \lambda(t)dt \qquad \Leftrightarrow \qquad \lambda(\tau) = \frac{dQ}{d\tau} \tag{3.1}$$

Note that for notational simplicity, we refer to $Q(\tau)$ throughout the text as Q but remind the reader that Q is a function of τ.

With this relationship and the notation defined above, we can write expressions for the relevant costs up to time T for a given switching time τ:

$$\text{Fixed Cost} = F\tau \qquad \textit{increasing in } \tau \tag{3.2}$$

$$\text{Variable Repair Cost} = v\int_0^{\tau} \lambda(t)dt \qquad \textit{increasing in } \tau \tag{3.3}$$

$$\text{Inventory Holding Cost} = h\int_0^{\tau} \left[Q - \int_0^{y} \lambda(t)dt\right] dy \qquad \textit{increasing in } \tau \tag{3.4}$$

$$\text{Replacement Cost} = r\int_{\tau}^{T} \lambda(t)dt \qquad \textit{decreasing in } \tau \tag{3.5}$$

Summing these individual cost components, we formulate the manufacturer's optimization problem as follows:

$$\min_{\tau} \qquad F\tau + vQ + hQ\tau - h\int_0^{\tau} \Lambda(t)dt + r\int_{\tau}^{T} \lambda(t)dt \tag{3.6}$$

where $\Lambda(t) = \int_0^t \lambda(u)du$. We observe that the objective function in the unconstrained optimization problem defined by equation (3.6) is continuous and differentiable. The first order optimality condition is given by:

$$\frac{dTC(\tau)}{d\tau} = F + v\lambda(\tau) + h\tau\lambda(\tau) - r\lambda(\tau) = 0 \tag{3.7}$$

Rearranging this equation, a first order critical point is described by:

$$\tau^* = \frac{1}{h}\left(r - v - \frac{F}{\lambda(\tau^*)}\right) \tag{3.8}$$

For a general return rate there may be multiple points that satisfy the first order condition (3.8). If the total cost function (3.6) is non-convex, there may exist multiple local minima. A first order critical point that satisfies the following second order necessary condition is a local minimum:

$$\frac{d^2TC(\tau)}{d\tau^2} = v\lambda'(\tau) + h\tau\lambda'(\tau) + h\lambda(\tau) - r\lambda'(\tau) > 0 \tag{3.9}$$

This condition can be rewritten as:

$$\lambda(\tau) > \frac{\lambda'(\tau)(r - v - h\tau)}{h} \Leftrightarrow \lambda(\tau) > \lambda'(\tau)\left[\frac{r-v}{h} - \tau\right] \tag{3.10}$$

Substituting τ^* from (3.8), we obtain the following condition for local minima:

$$\lambda^2(\tau^*) > \frac{F}{h}\lambda'(\tau^*) \tag{3.11}$$

To obtain the global least cost solution, the total cost is evaluated at each local minimum as well the end points of the planning horizon (e.g., τ = 0 and τ = T).

Confirming our intuition, result (3.8) shows that the optimal switching time, τ^*, decreases as the unit inventory holding cost, fixed cost and variable cost increase. Conversely, τ^* increases as the cost of replacement and the arrival rate increase. By rearranging (3.8), we observe that at the optimal switching time, τ^*, the marginal cost of the last repaired unit is equal to the (marginal) cost of replacement:

$$\frac{F}{\lambda(\tau^*)} + v + h\tau^* = r \tag{3.12}$$

Thus the marginal cost of repair at time τ^* is the total sum of (i) the fixed cost of having repair capabilities at time τ^* divided by the number of repairs performed at that time, (ii) the variable cost of repair, and (iii) the cost of holding the last repair part in inventory for a length of time equal to τ^*. τ^* is then defined as the time at which this marginal cost of repair is exactly equal to the unit replacement cost.

By rearranging the second order condition given by (3.9), we get the following condition which ensures the convexity of the total cost function:

$$\frac{h}{r - ht - v} > \frac{\lambda'(t)}{\lambda(t)} \tag{3.13}$$

If the return arrival rate $\lambda(t)$ satisfies condition (3.13) for all values t $\in$ [0,T], then the total cost function (3.6) is convex. Therefore, if there exists a τ^* that satisfies the first order condition (3.8) in the interval [0,T], then it is a global minimizer. If no such τ^* exists in [0,T], then the minimum cost switching time is at one of the endpoints (e.g., $\tau^* = 0$ or $\tau^* =$ T).

Note that we have assumed that repairs occur during the interval $[0, \tau^*]$ and that products are replaced during $[\tau^*, \mathrm{T}]$. There may be an interval of time early in the planning period during which $\lambda(t)$ is relatively small. If the repair parts purchase can also be delayed or if incurring the fixed costs can be postponed until the beginning of the repair interval, replacements may be cheaper than repairs during this time. Computational results suggest that the cost penalty paid by assuming that repairs begin at time 0 is quite small (less than 1% - 2% in our numerical experiments). Extending the results above to encompass this early replacement interval is trivial and gives result (3.8) as the critical times to begin and end the repair operation. For simplicity we maintain assumptions 3 and 4 throughout this chapter but acknowledge that it may lead to suboptimal solutions for cases of extremely low early return rates.

Comparative Statics For τ^*

From equation (3.8), one can easily investigate how τ^* changes when F, r, v, and h change. We summarize the comparative statics in Table 3.1 and discuss our findings below.

The analysis of the derivatives shows that when the return arrival rate is constant or is decreasing at τ^*, an increase in fixed cost (F), variable unit repair cost (v), and unit inventory holding cost (h) unambiguously results in lower values of τ^*, denoted by a negative sign for the respective derivatives in Table 3.1. On the other hand, an increase in replacement cost (r) leads to a higher value of τ^*, hence the positive derivative sign. Interestingly, we find that when the arrival rate is increasing, the effect of these parameters on the optimal τ^* value is less obvious and can be either increasing or decreasing, depending on the values of

Table 3.1: Comparative Statics For τ^*: Sign of Derivatives

	$\frac{d\tau^*}{dF}$	$\frac{d\tau^*}{dr}$	$\frac{d\tau^*}{dv}$	$\frac{d\tau^*}{dh}$
Expression	$\frac{1}{F\frac{\lambda'(\tau^*)}{\lambda(\tau^*)}-h\lambda(\tau^*)}$	$\frac{1}{h-F\frac{\lambda'(\tau^*)}{\lambda^2(\tau^*)}}$	$\frac{1}{F\frac{\lambda'(\tau^*)}{\lambda^2(\tau^*)}-h}$	$\frac{r-v-\frac{F}{\lambda(\tau^*)}}{h\left(\frac{F\lambda'(\tau^*)}{\lambda^2(\tau^*)}-h\right)}$
$\lambda(t)$ Constant	-	+	-	-
$\lambda(t)$ Increasing	+/-	+/-	+/-	+/-
$\lambda(t)$ Decreasing	-	+	-	-

F, r, v, and h. For instance, when $\lambda'(\tau^*) > 0$, τ^* would be increasing in v and decreasing in r when F is very large. In the following numerical example we further investigate the not so obvious effect of problem parameters on optimal τ^* values.

A Numerical Example

The following numerical example illustrates the sensitivity of the total cost function and the optimal switching time to changes in the relevant instance data. The planning horizon is from t = 0 to T = 100. The return rate pattern follows a truncated bell-shaped curve with 5,000 total returns and maximum return rate of 84 units (see Figure 3.3 below). This example uses a fixed cost of \$2500 per unit time, a replacement cost of \$200 per unit, and unit holding costs for the entire planning period of 15% of the unit variable cost. Figure 3.4 below shows how the total cost function behaves for variable repair costs of \$25, \$65, \$100, \$140. Figure 3.4 shows that the minimum total cost switching time shifts toward the end of the planning period as the variable cost (and therefore holding cost) decreases. For variable repair costs of \$25, \$65, \$100, and \$140, the optimal switching time τ^* is 96, 90, 82, and 65, respectively.

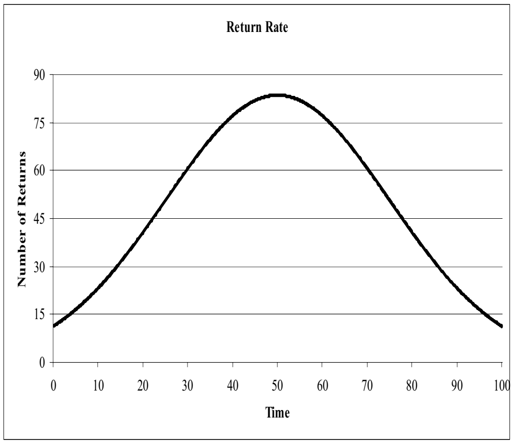

Figure 3.3: Return Rate

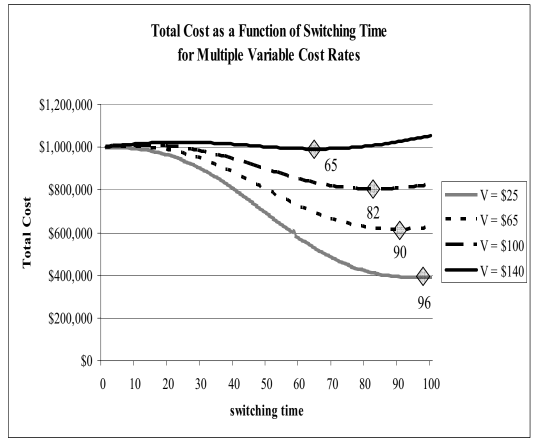

Figure 3.4: Total Cost For Various Switching Times

To obtain the optimal switching times and corresponding optimal order quantities shown in Figures 3.5 and 3.6 below, we vary the fixed cost from $1,000 to $4,000 and we vary the variable repair cost from $0 to $200. As each of the fixed cost curves indicate, the repair horizon contracts as the variable cost increases. For very high variable cost values, the repair option is never cost-effective. There is a variable cost threshold above which it is optimal to replace all returned items, and this threshold decreases as F increases.

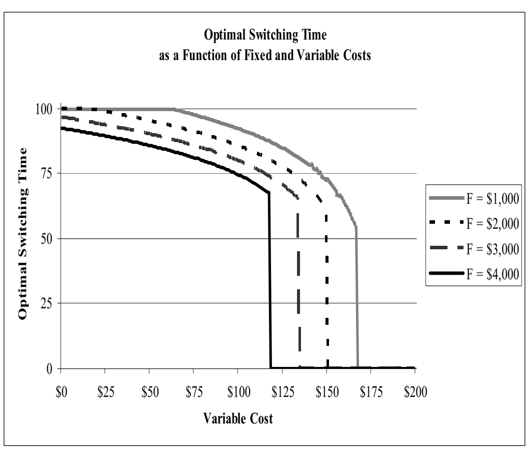

Figure 3.5: Sensitivity of Optimal Switching Time To Fixed And Variable Costs

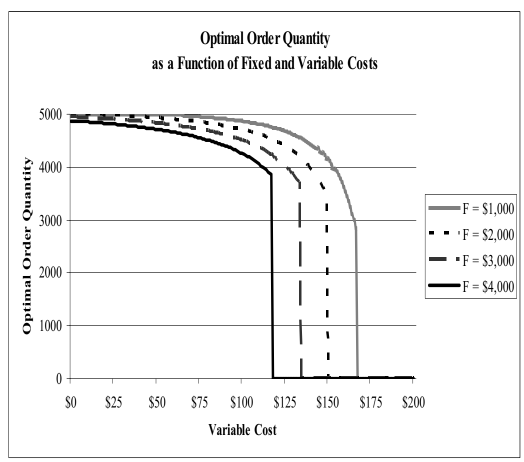

Figure 3.6: Sensitivity of Optimal Order Quantity To Fixed And Variable Costs

The single product uncapacitated model provides a solid foundation for formulating a variety of extensions. In the next section, we extend this basic model to incorporate capacity constraints for the single product model.

3.3.2 The Capacitated Single Product Model

The capacitated single product problem is a natural extension of the basic model in Section 3.3.1. We follow the same notation as in Section 3.3.1 and include one additional parameter, K, which represents the maximum number of repairs per unit time. In addition to the assumptions of the uncapacitated model, we consider return rates such that the time period over which return arrivals exceed capacity is a convex set. This assumption allows us to define 3 disjoint regions as shown below in Figure 3.7.

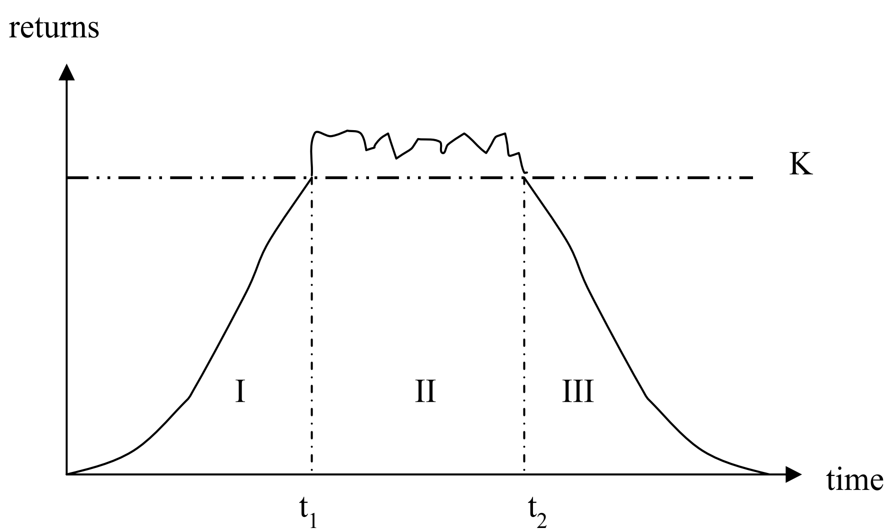

Figure 3.7: Capacity Regions

Since the return rate is known, we can calculate t_1 and t_2, the times at which the arrival rate exceeds capacity for the first time and at which the arrival rate drops below capacity, repectively. This leaves three cases for us to consider:

$$\begin{array}{lll} \text{(i)} & \tau^* \leq t_1 & \text{(region I)} \\ \text{(ii)} & t_1 < \tau^* \leq t_2 & \text{(region II)} \\ \text{(iii)} & t_2 < \tau^* & \text{(region III)} \end{array}$$

Solution Algorithm

The first step in determining the optimal switching time and corresponding order quantity for the capacitated single product model is to define the order quantity expression for each of the three regions:

$$\begin{array}{lll} \text{region I} & Q_I(\tau) & = \int_0^{\tau} \lambda(t)dt \\ \text{region II} & Q_{II}(\tau) & = \int_0^{t_1} \lambda(t)dt + K\tau - Kt_1 \\ \text{region III} & Q_{III}(\tau) & = \int_0^{t_1} \lambda(t)dt + \int_{t_2}^{\tau} \lambda(t)dt + Kt_2 - Kt_1 \end{array}$$

For notational simplicity, we refer to $Q_I(\tau)$, $Q_{II}(\tau)$, and $Q_{III}(\tau)$ throughout the text as Q_I, Q_{II}, and Q_{III} but remind the reader that all order quantities are a function of τ.
Using these order quantities, we obtain the total cost expressions in Table 3.2.

The first order critical points, τ^*, for cases (i) - (iii) given that the first order point is strictly interior to each region are given by:

$$\text{region I} \quad \tau_I^* = \frac{1}{h}\left(r - v - \frac{F}{\lambda(\tau^*)}\right) \tag{3.14}$$

$$\text{region II} \quad \tau_{II}^* = \frac{1}{2}\left(1 + t_1 + \frac{r - v}{h} - \frac{F}{hK}\right) \tag{3.15}$$

$$\text{region III} \quad \tau_{III}^* = \frac{1}{h}\left(r - v - \frac{F}{\lambda(\tau^*)}\right) \tag{3.16}$$

Table 3.2: Total Cost Expressions For The Single Product Capacitated Model

Region	Total Cost Expression
I	$F\tau + vQ_I + hQ_I\tau - h\int_0^\tau \Lambda(t)dt + r\int_\tau^T \lambda(t)dt$
II	$F\tau + vQ_{II} + h\left[\int_0^{t_1}\left[Q_{II} - \int_0^y \lambda(t)dt\right]dy + \int_{t_1}^{\tau}\left[Q_{II} - \Lambda(t_1) - K\right]dy\right] + r\left[\int_{t_1}^{\tau}(\lambda(t) - K)\,dt + \int_\tau^T \lambda(t)dt\right]$
III	$F\tau + vQ_{III} + h\left[\int_0^{t_1}\left[Q_{III} - \int_0^y \lambda(t)dt\right]dy + \int_{t_1}^{t_2}\left[Q_{III} - \Lambda(t_1) - K\right]dy\right] + h\left[\int_{t_2}^{\tau}\left[Q_{III} - \Lambda(t_1) - K(t_2 - t_1) - \int_{t_2}^{y}\lambda(t)dt\right]\right]dy + r\left[\int_{t_1}^{t_2}(\lambda(t) - K)\,dt + \int_\tau^T \lambda(t)dt\right]$
Alternative Expression for III	define $\hat{\lambda}(t) = min(\lambda(t), K)$ $F\tau + vQ_{III} + hQ_{III}\tau - h\int_0^\tau \hat{\Lambda}(t)dt + r\int_\tau^T \lambda(t)dt + r\int_{t_1}^{t_2}\left(\hat{\lambda}(t) - K\right)dt$ where $\Lambda(t) = \int_0^\tau \lambda(t)dt$ and $\hat{\Lambda}(t) = \int_0^\tau \hat{\lambda}(t)dt$

The constrained single product model does not yield a closed-form solution. Therefore, we evaluate the total cost function at the first order citical points defined by 3.14, 3.15, and 3.16 as well as the endpoints (t = 0 and t = T) and the region borders (t = t_1 and t = t_2). We can obtain candidate solutions by solving a constrained optimization problem for each of the three defined regions. Expressions for the first order critical points in each region are given above. The following algorithm can be used to find the global least cost switching time. In all cases, ties are broken arbitrarily.

1. Evaluate the total cost at time t = 0 and t = T using the cost expressions found in

Table 3.2. Let $\hat{\tau}_{endpoint}$ be the switching time which yields the lower of the two values.

2. Find all values $\hat{\tau}$ that satisfy (3.14) and $\hat{\tau} \leq t_1$. Evaluate the total cost expression for region I found in Table 3.2 for each value of $\hat{\tau}$. Let $\hat{\tau}_I$ be the switching time that gives the least total cost from among these candidate switching times.

3. Evaluate the total cost expression for region I found in Table 3.2 using $\tau = t_1$

4. Find all values $\hat{\tau}$ that satisfy (3.15) and $t_1 < \hat{\tau} < t_2$. Evaluate the total cost expression for region II found in Table 3.2 for each value of $\hat{\tau}$. Let $\hat{\tau}_{II}$ be the switching time that gives the least total cost from among these candidate switching times.

5. Evaluate the total cost expression for region II found in Table 3.2 using $\tau = t_2$

6. Find all values $\hat{\tau}$ that satisfy (3.16) and $t_2 < \hat{\tau}$. Evaluate the total cost expression for region III found in Table 3.2 for each value of $\hat{\tau}$. Let $\hat{\tau}_{III}$ be the switching time that gives the least total cost from among these candidate switching times.

7. Select the least total cost value found in steps 1 - 6 above. The corresponding switching time is the global minimum switching time, τ^*.

This algorithm is similar in spirit to the algorithm used for finding the optimal order quantitiy in the EOQ with quantity discounts model (Chopra and Meindl, 2004). We note that the assumption that the period(s) of time during which the return rate exceeds repair capacity form a convex set. This assumption allows us to define the three regions shown in Figure 3.7. For instances in which these periods of time do not form a convex set, a straightforward extension of the formulations and methods presented above can be developed to solve this problem.

The following example illustrates how the solution may fall into regions I, II, III, or on the borders. In this example, returns arrive at a constant rate of 12 returns per time period during [0, 26] and [76,100]; in between, the rate is 15 returns per time period (see Figure 3.8). The capacity is set to 14, the fixed cost is \$10 per unit time, the replacement cost is \$2 per item, the holding cost incurred over the entire planning period is 15% of the variable cost, and the variable cost is changed to show how the optimal switching time responds. For v = \$1.14, $\tau* = 8$ (region I); for v = \$1.10, $\tau* = 59$ (region II); for v = \$1.00, $\tau* = 76$ (on the border between regions II and III); for v = \$0.90, $\tau* = 95$ (region III). For this particular instance and set of parameters, the optimal switching time does not fall on the border between regions I and II. The corresponding total cost curves for these four variable cost values are shown in Figure 3.9.

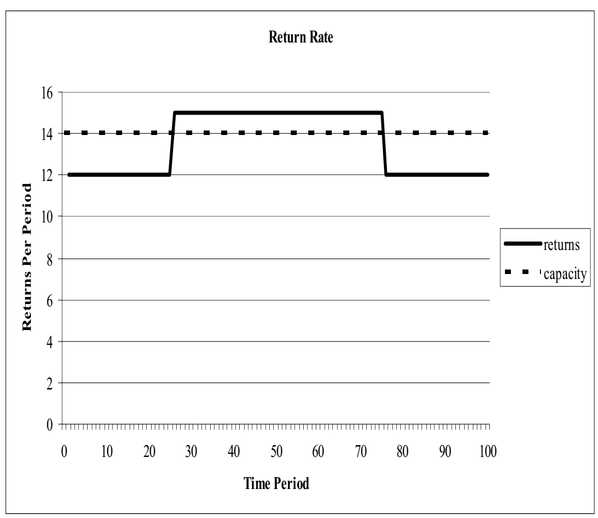

Figure 3.8: Single Product With Capacity Constraints Return Rate

A Numerical Example

We modified the numerical example from Section 3.3.1 to include capacity constraints. The optimal switching time and corresponding optimal total cost are shown below for variable

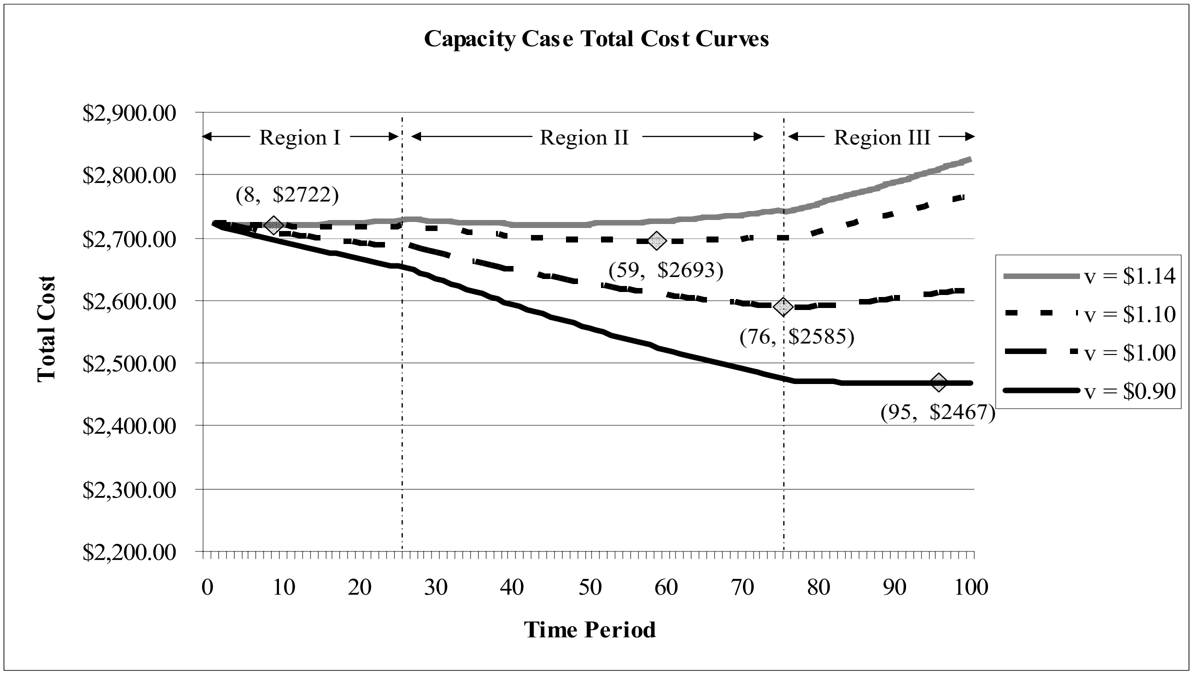

Figure 3.9: Single Product With Capacity Constraints Total Cost Curves

costs ranging from \$0 to \$200 and repair capacity equaling 20, 40, 60, and 80 repairs per time period. Taken together, Figures 3.10 and 3.11 show that while different variable costs may yield the same optimal switching time, the corresponding total cost is significantly higher for lower capacities (due to the need to replace a larger proportion of returns during time periods in which the return rate exceeds capacity). As Section 3.3.1 shows, the repair horizon contracts as the variable cost increases, but there is a higher variable cost threshold as capacity increases. The optimal total cost curves in Figure 3.11 are helpful in making strategic decisions regarding capacity design. For example, if the difference between a small capacity cost curve and a larger capacity cost curve in the region of interest (for a particular variable cost range) is larger than the fixed cost to increase repair capacity to the larger threshold, then the manufacturer should expand its capacity to reduce its overall costs. If the difference between the curves is relatively small, it may be preferable to incur the cost of replacing additional items and not increase repair capacity.

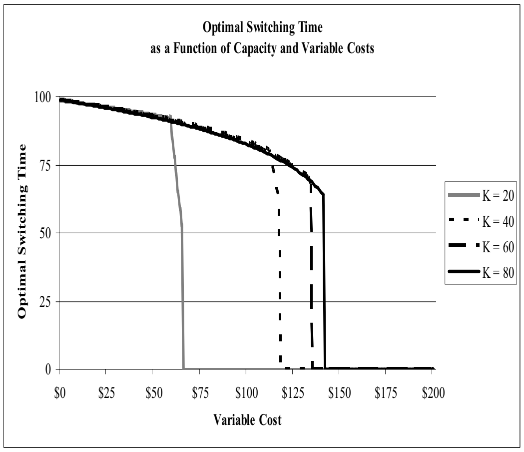

Figure 3.10: Sensitivity of Optimal Switching Time To Capacity And Variable Cost

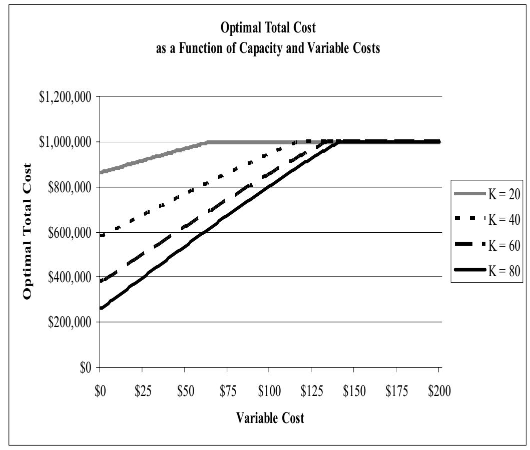

Figure 3.11: Optimal Total Cost As A Function Of Capacity And Variable Cost

The next section examines another extension of the basic model. More specifically, Section 3.3.3 explores the impact of multiple products in a shared repair facility.

3.3.3 The Uncapacitated Two Product Case With Shared Facility Model

The two product with shared facility costs model is a first step in considering the effects of utilizing the same facility to repair multiple products. In this formulation, each repaired product requires a unique inventory but we allow the same fixed cost investment to service both products. Observe that as in the single product model, the fixed cost is incurred if any repairs are performed, so in this case it is incurred for the longer of the two repair horizons. The inventory quantities, Q_1 and Q_2, are calculated for each product as shown in (3.1) and we assume that the warranty periods are identical. We derive the following total cost function where subscripts 1 and 2 denote product 1 and product 2, respectively:

$$\min_{\tau_1,\tau_2} \quad F * max(\tau_1, \tau_2) + v_1 Q_1 + v_2 Q_2 + h_1 Q_1 \tau_1 + h_2 Q_2 \tau_2 \tag{3.17}$$

$$- h_1 \int_0^{\tau_1} \Lambda_1(t)dt - h_2 \int_0^{\tau_2} \Lambda_2(t)dt + r_1 \int_{\tau_1}^{T} \lambda_1(t)dt + r_2 \int_{\tau_2}^{T} \lambda_2(t)dt$$

To find the optimal switching time pair, (τ_1^*, τ_2^*), we can solve two constrained optimization problems - one in which τ_1^* and τ_2^* are equal and one in which they are not. If the two switching times are not equal, then without any loss of generality we can assume that $\tau_1^* < \tau_2^*$. Solving this version of the problem, we get the following first order critical points:

$$\tau_1^* = \frac{1}{h_1}(r_1 - v_1) \qquad \tau_2^* = \frac{1}{h_2}\left(r_2 - v_2 - \frac{F}{\lambda_2(\tau_2^*)}\right) \tag{3.18}$$

This is true when the fixed cost is less than the following threshold, F^*:

$$F^* = \lambda_2(\tau_2^*)\left[r_2 - v_2 - \frac{h_2}{h_1}(r_1 - v_1)\right] \tag{3.19}$$

When $F \geq F^*$, then $\tau_1^* = \tau_2^* = \tau^*$ and the two switching times are equal. The first order point in this case is given by:

$$\tau^* = \frac{\lambda_1(\tau^*)(r_1 - v_1) + \lambda_2(\tau^*)(r_2 - v_2) - F}{h_1\lambda_1(\tau^*) + h_2\lambda_2(\tau^*)} \tag{3.20}$$

Figure 3.12 below qualitatively illustrates the relationship between the two optimal switching times τ_1^* and τ_2^* as the fixed cost of repair capability changes. Case 1 shows the effect of a free ride for product 1 as the fixed cost is driven by τ_2^*. This corresponds to a low fixed cost. As the fixed cost increases up to the critical cost defined in (3.19), τ_2^* decreases as τ_1^* remains unchanged (shown in case 2). Finally, case 3 shows that as the fixed cost increases beyond the critical value F*, τ_1^* and τ_2^* decrease together.

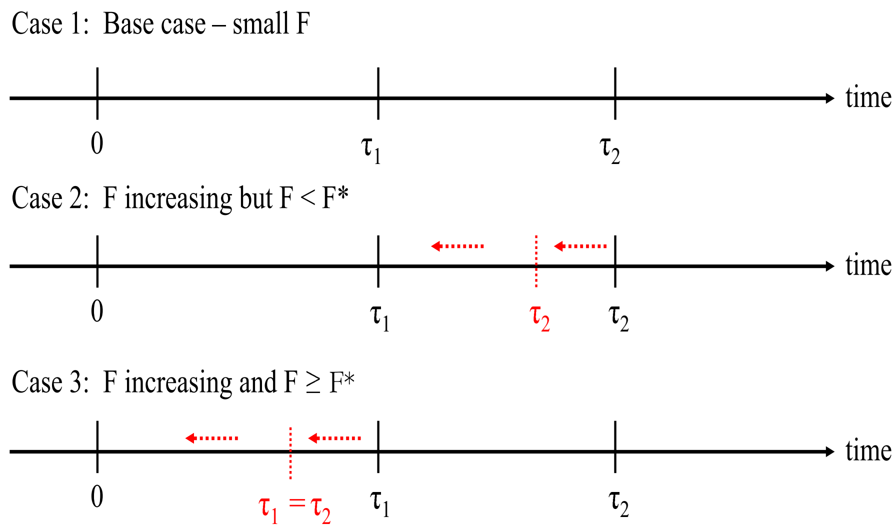

Figure 3.12: Switching Time Relationship As A Function Of Fixed Cost

Restricting our attention to the case $\tau_1^* \leq \tau_2^*$, Figure 3.13 illustrates the solution space for the two product problem and the relationship between the two optimal switching times as a function of the input data. (The interpretation is analogous for the case $\tau_1^* \geq \tau_2^*$.) The quantity $(r_2 - v_2)$ is the "incremental" marginal cost of replacement, so the horizontal axis measures the difference in the incremental marginal cost of replacement adjusted by the holding cost rate of the two products. Note that for the case $\tau_1^* \geq \tau_2^*$, this quantity is negative. In this case, the product indices are reversed. The vertical axis is measured by the fixed cost of having repair capabilities. The slope of the line which separates the region into $\tau_1^* = \tau_2^*$ and $\tau_1^* < \tau_2^*$, is the magnitude of the holding cost on day τ_2^* (since $\lambda_2(\tau_2^*)$ is the number of product 2 returns on day τ_2^*). These values and their relationships result directly from the proof of the two product case values.

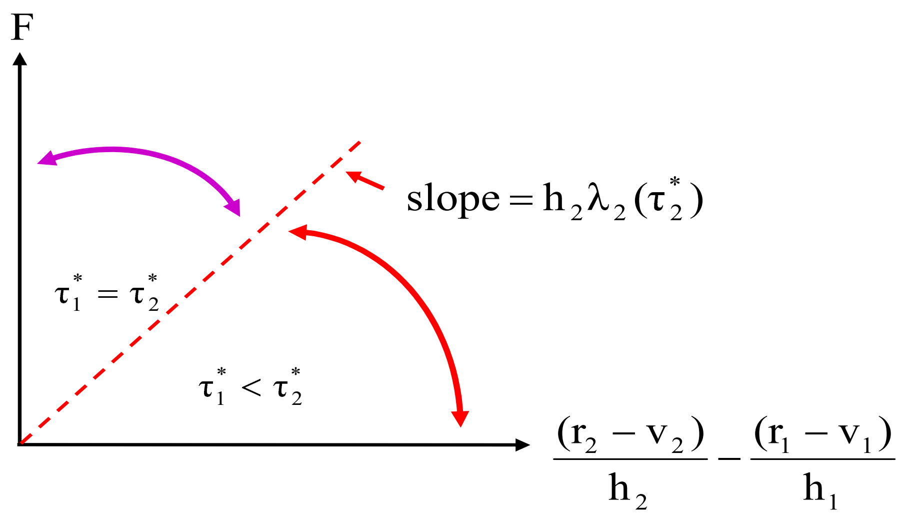

Figure 3.13: Switching Time Relationship As A Function Of Input Data

A Numerical Example

The following numerical example shows how τ_1^* and τ_2^* for a given instance change as the fixed cost and variable costs for the two products change. As in the previous examples, the planning horizon is from t = 0 to t = 100. Product 1 has a maximum return rate of 82 and 3,000 total returns while product 2 has a maximum return rate of 100 and 6,000 total returns (see Figure 3.14).

This example uses a fixed cost of \$15, variable repair costs of \$0.50 (for both products), a holding cost for the entire planning period equaling 15% of the variable cost (for both products), and a replacement cost of \$1 (for both products). This instance yields the optimal switching time pair $(\tau_1^*\ ,\ \tau_2^*)$ = (66, 82) represented by the solid diamond in Figures 3.15 and 3.16 below. These figures show how the optimal switching time pair changes with the instance data, and are consistent with the insights provided by Figures 3.12 and 3.13. Figure 3.15 illustrates the "free ride" effect as τ_1^* remains unchanged for fixed cost values between \$7.50 and \$30.00 (or 0.5x to 2.0x the original fixed cost). As the fixed cost increases beyond this factor of two, the optimal pair $(\tau_1^*\ ,\ \tau_2^*)$ decreases together. Figure 3.16 shows the

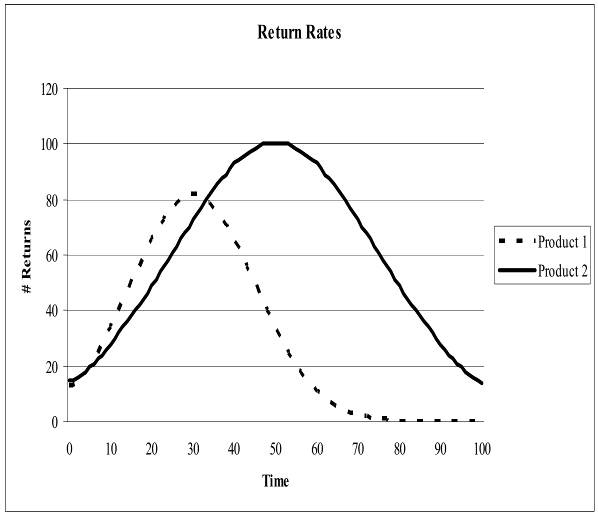

Figure 3.14: Return Rate Curves For Two Product Example

sensitivity of the optimal pair to changes in the variable costs of the two products. Since $\tau_1^* \leq \tau_2^*$, τ_2 is unaffected by moderate (less than a factor of 2) changes in v_1; however when the variable cost of product 2 (v_2) doubles, τ_2 goes to 0 and τ_1 decreases from its original value (for this combination of input data values, we see that $\tau_1^* \geq \tau_2^*$).

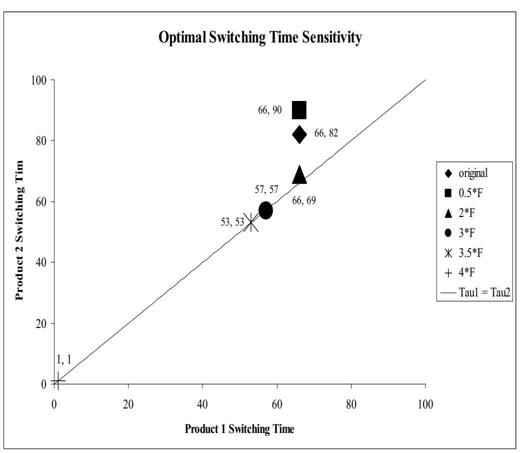

Figure 3.15: Sensitivity of Optimal Switching Time To Changes in Fixed Cost

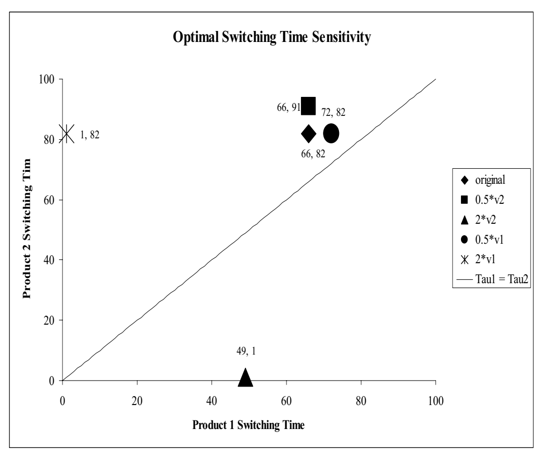

Figure 3.16: Sensitivity of Optimal Switching Time To Changes in Variable Costs

3.4 Industry Case

This research was motivated by our discussions with two U.S. companies that specialize in producing a wide range of personal digital assistants and multi-functional mobile devices. In this section, we test our model on data provided by one of these companies. The firm is currently experiencing high warranty costs that consume between 2 - 3% of its revenues. For most of its products, the firm faces lifetime buy decisions to support warranty repair operations. We tested our models on data provided by the firm and provide results for one of their products below. (To protect the identity of the company, the data have been normalized so that total return volumes and the magnitude of relevant costs are not given; however the relative costs are provided.)

The data below represent a mid- to high-end personal digital assistant product with a standard 1-year warranty period that has been returned due to a defective LCD screen. The weekly return rate over this warranty period is given in Figure 3.17 (rates are given relative to the average) and shows that the observed rates can vary from 10% to over 160% of the average. We normalize our cost data by dividing all of the relevant costs by the cost of providing a replacement unit to the customer. We obtain the following cost values: fixed = \$15.00, variable = \$0.55, holding (per item per week) = \$0.0015, and replacement = \$1.00. Applying the uncapacitated single product model to this data set we obtain the cost curve given in Figure 3.18. The total cost curve over the planning horizon is given relative to the cost at the minimum cost switching time, which is near week 36. For this particular part-product combination, the company should plan to repair returned products during the first 8 - 9 months of the warranty period. After this time, the company should reallocate its repair resources to another product line and satisfy warranty claims with replacement products. Figure 3.18 shows the total cost penalty of not switching from repair to replacement mode at this time. For example, the firm would pay a cost penalty of 16% by not repairing any of its returns; at the other extreme it would pay a cost penalty of 7% by repairing all returns.

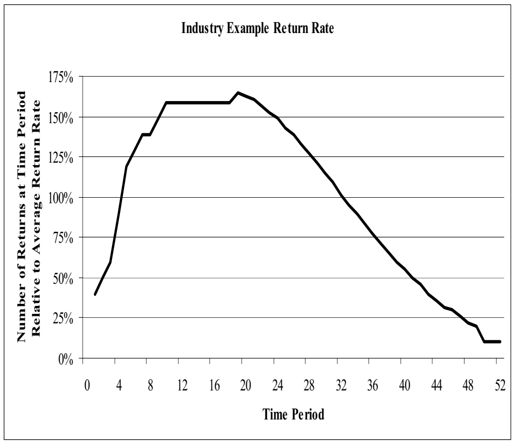

Figure 3.17: Industry Case Return Rate

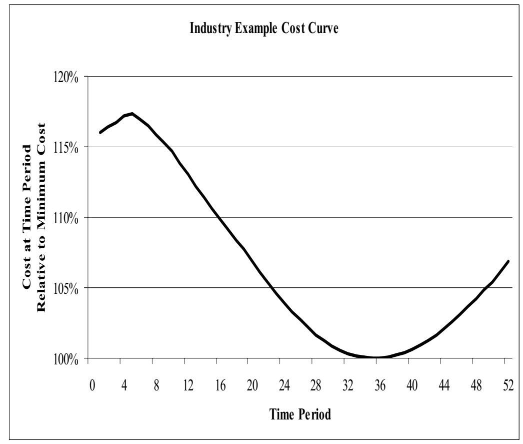

Figure 3.18: Industry Case Cost Curve

3.5 Conclusions and Directions for Future Research

In this chapter, we have developed deterministic models for purchasing repair parts to service warranty returns for short lifecycle products. The first model considers the case of a single product and illustrates the relationship between fixed repair capability costs, inventory holding costs, replacement costs, and variable repair costs in determining an optimal inventory quantity and identifying when a firm should repair or replace returned items in an uncapacitated setting. The second model examines the role of capacity constraints for the single product model. The third model considers the relationship between these decisions for two products that share repair resources, and shows that the decision of when to stop repairing returns and to begin replacing them with alternative products is driven by the fixed cost. Finally, we show how sensitive these decisions are to changes in the relative data.

The deterministic models presented in this chapter provide valuable insights into the interactions between the relevant cost components of this problem. Within the deterministic case, there are other extensions that one would pursue as part of future research. The first of these would examine the role of overlapping product generations and the viability of using procurement opportunities for later generation models to support the repair operations of earlier generation products. A second extension would consider the impact of parts commonality across multiple products. Finally, a third extension of the deterministic case would consider the impact of minimum order quantities and order-size-based pricing for repair parts.

In addition to the deterministic models suggested above, a stochastic version of the basic model is also worth considering. Such an extension would consider the impact of return rate uncertainty and the value of information regarding return patterns throughout the product lifecycle. Supplier contracting options, pricing to reserve supplier capacity, and the value of multiple procurement opportunities throughout the repair horizon are some of the additional aspects one can investigate in a stochastic model. Finally, there may be risk pooling benefits associated with sharing repair resources among product lines if there is sufficient parts commonality.

In summary, this chapter is one of the first studies to investigate lifetime buy decisions and the role of relevant costs in this problem. We believe that the current work establishes a foundation for more detailed modeling efforts in future research.

Chapter 4

Scenario Planning for Warranty Repair Operations with Lifetime Buys

4.1 Introduction

In Chapter 3 we introduced the lifetime buy problem for repair operations which balances repair and replacement costs for satisfying defective product warranty claims. The models developed in Chapter 3 assume that the manufacturer has perfect information about the return rate, including the exact quantity and timing of returns. In practice it is impossible to know this information with complete certainty. Introducing return rate uncertainty complicates the models in Chapter 3 by allowing a *salvage* option in which repair parts inventory is discarded when the manufacturer switches from repair to replacement mode. We include an additional parameter, the *salvage cost*, to consider this consequence in the planning process. In this chapter we develop two models for addressing this problem when the return rate is uncertain.

A first step in considering the impact of return rate uncertainty is to develop scenario planning models in which there is a set of known return rates *(scenarios)*, each with a probability of occurring. The scenario planning approach is a popular method for addressing uncertainty in the facility location literature and proves to be a useful approach for the lifetime buy problem. This method allows the user to include some measure of uncertainty by defining a discrete set of potential outcomes but relieves the user of the burden of having to specify a continuous probability distribution of all possible system realizations. This approach is particularly well-suited for the lifetime buy problem that we consider because it readily accomodates the discrete (unit-based) and continuous (time-based) cost components with relatively straightforward solution methods.

One potential drawback of the scenario planning approach for the lifetime buy problem is a firm's ability (or lack thereof) to define a set of return rate scenarios. However, in the electronics and telecommuncations industries, a firm should have adequate information to readily define these scenarios. Firms spend a considerable amount of time and resources to predict a product's failure rate and have enough knowledge about the quality of the product to be able to identify bounds on the percentage of items that will be returned due to defects. Using this information, a firm should be able to produce at least three estimates of the return rate they will observe: a low rate (best case scenario), expected rate (most likely scenario), and a high rate (worst case scenario). Using these rates and sales data, the firm can determine the expected quantity of returns. The timing of the returns is directly related to the sales data. Taken together, the quantity and timing of returns are sufficient for defining the set of scenarios to be considered.

The scenario planning models in this chapter allow the decision maker to define different objectives for measuring the efficiency of a particular solution, where a solution is the lifetime buy order quantity. In the deterministic models of Chapter 3, the objective is to minimize the total cost of satisfying all returns. It is not immediately clear that any other objective would be more appropriate. By contrast, the existence of multiple return rate scenarios introduces additional objectives which may be appealing to a decision maker, depending on the amount of risk associated with a solution that he is willing to tolerate. This chapter introduces two models which each address a different objective as well as a tradeoff curve to identify solutions that attain both objectives to varying degrees. The first model (expected cost model) uses scenario probabilities to minimize the expected cost over all possible return rate realizations. This approach optimizes the expected performance of the system, but can result in particularly bad solutions if an extreme scenario is realized. By contrast, the second model (minimax regret model) optimizes the worst case performance of the system by minimizing the maximum regret of a solution over all possible scenarios. In this context, the regret for order quantity q and a specific return rate scenario i is defined as the difference between the optimal total cost for scenario i and the total cost of having ordered q before scenario i has been realized. In many applications, the minimax regret objective is overly conservative since the optimal solution is driven by a single scenario which may have a very small probability of being realized. The tradeoff curve of efficient solutions balances the benefits and drawbacks of both modelling approaches.

The rest of this chapter is organized as follows. In Section 4.2 we review literature related to the lifetime buy problem and the scenario planning methodology. The expected cost model is developed in Section 4.3. A numerical example shows why using the expected cost model is preferable to substituting the expected return rate into a single scenario model. In Section 4.4 we present the minimax regret model and results for this approach. Section 4.5 examines the tradeoff between the expected cost model and the minimax regret model and provides a curve of efficent solutions. Conclusions and directions for future research comprise Section 4.6.

4.2 Literature Review

This chapter explores an extension of the single uncapacitated lifetime buy model developed in Chapter 3. Literature related to this problem, including topics such as repairable inventory management, warranty repairs, and classical inventory management remain relevant. The repairable inventory management literature, which has yet to consider the impact of lifetime buys, typically considers internal repair systems where the firm repairs items for its own use and not for end customers. This research generally considers longer lifecycle products with large repair networks and multiple parts procurement opportunities. This area of research provides valuable insight into our work, but is inappropriate for the lifetime buy problem which serves very short lifecycle products, a single repair facility, and a single parts procurement opportunity. (See Sherbrooke, 1968; Nahmias, 1981; Guide and Srivastava, 1997.) The warranty repairs literature estimates the cost of various warranty policies, but does not explicitly consider the impact of lifetime buys, which significantly changes both the modeling approach and the characterization of the solution. (See Balachandran, Maschmeyer, and Livingston, 1981; Balcer and Sahin, 1986; Zuo, Liu, and Murthy, 2000; Ja et al., 2001; Murthy and Djamaludin, 2002.) Finally, the lifetime buy problem blends properties found in classical inventory management problems such as the EOQ and Newsvendor models. (See Harris, 1913; Wadsworth, 1959; Petruzzi and Dada, 1999; Porteus, 1990; Porteus, 2002.) A more detailed review of the related literature found in these topics is given in Section 3.2.

This chapter also draws heavily on the use of scenario planning models which have a rich history in the facility location literature stream. The value of scenario planning from an operations management perspective has long been acknowledged as a way to plan for uncertainty and identify strategies for coping with future realizations (Ringwald, 1998). Scenario planning techniques may be used to plan for future states in which problem parameters may take on a range of potential values (Owen, 1999) or to find *robust* solutions that perform well under expected and worst case realizations (Snyder, 2003). These techniques are used to achieve a variety of objectives, including minimizing the worst case performance of the system when the demand is uncertain (Kouvelis and Yu, 1997; Averbakh and Berman, 1997). Another common objective is maximizing market share (Serra, Ratick, and ReVelle, 1996; Serra and Marianov, 1998; Current, Ratick, and ReVelle, 1997). The *reliability* of a system is another objective in scenario planning where the objective is to design a system that operates reliably at least α-percent of the time (Daskin, Hesse, and ReVelle, 1997; Owen, 1999). Related to this are models in which the expected regret is minimized when considering scenarios that will occur at least α percent of the time (Chen et al., 2006). A review of literature related to scenario planning for facility location problems may be found in Owen and Daskin (1998).

4.3 The Expected Cost Model

The expected cost model minimizes the manufacturer's expected total cost for servicing product returns given a set of possible return rate scenarios and the probability that each will occur. This approach to scenario planning optimizes the expected performance of the system.

Before continuing, we note that the expected cost model of Section 4.3 and the minimax regret model of Section 4.4 make the same assumptions as the lifetime buy model of Chapter 3. Specifically, these models assume:

1. Defective products arrive to a service center over a known time period [0,T]; There is a set of possible return rates, each of which is a known, positive, bounded function of time which is twice differentiable over the planning period.

2. Each returned item is either repaired using the available parts inventory or replaced

with a comparable new product.

3. All inventory required to support the repair operation is purchased at the beginning of the planning horizon. For simplicity, we assume a 1 to 1 relationship between returns and repair parts required to service the return.

4. Returns are repaired at the beginning of the planning period.

These models also assume that repair facilities are uncapacitated. In the absence of a capacity constraint, we note that the repair and replacement periods are disjoint; the repair period is given by $[0, \tau^*]$ and the replacement period is given by $[\tau^*, \text{T}]$. This observation follows from the swapping argument given in Chapter 3 in the formulation of the Uncapacitated Single Product Model. We note that for cases in which the return rate is very low early in the planning period, it may be preferable to replace some units for a short time before switching into repair mode. In the models that follow we assume that the return rate is sufficiently large that this does not happen.

4.3.1 Formulation

The scenario planning models in this chapter use the same notation as in Chapter 3. T represents the length of the planning horizon. The costs associated with repairing items include: F, the fixed cost per unit time of having repair capability; v, the variable cost per item repaired; and h, the repair parts inventory holding cost per item per unit time. The unit cost of replacement is r. The set of possible return rate scenarios is denoted by M. For each scenario $m \in M$, there is a known return rate $\lambda_m(t)$ and a probability p_m that scenario m will be realized. The manufacturer must decide Q, the repair parts order quantity. Given Q and a specific scenario m, an optimal switching time τ_m^* can be calculated. Since the optimal order quantity Q^* is determined by the costs of servicing all possible scenarios, there is likely to be unused repair parts when a low return rate is realized. The salvage cost (purchase price less salvage value) for an unused repair part is given by s. The salvage cost is incurred at the switching time.

In the case of a single scenario, the expected cost problem reduces to the uncapacitated single product model in Chapter 3. The optimal solution will identify the optimal switching time and order quantity as developed in Chapter 3 with zero salvaged repair parts. However, the simultaneous consideration of multiple scenarios is likely to result in optimal order quantities which require salvaging repair parts for lower return rate scenarios. Recognizing this, we develop a different modelling approach wherein the order quantity is given as an input and the optimal solution identifies a switching time and a quantity of parts to be salvaged. For a given return rate $\lambda(t)$ and order quantity Q, we can calculate the maximum feasible repair time $\hat{\tau}$ where $Q = \int_0^{\hat{\tau}} \lambda(t)dt$. We can then calculate the minimum cost switching time τ^* using the following constrained optimization problem:

$$\begin{aligned} \text{Minimize}_\tau \qquad Z(\tau,\lambda,Q) = & F\tau + v\int_0^\tau \lambda(t)dt + h\int_0^\tau \left[Q - \int_0^y \lambda(t)dt\right] dy \qquad (4.1) \\ & + s\left(Q - \int_0^\tau \lambda(t)dt\right) + r\int_\tau^T \lambda(t)dt \\ \text{Subject to} \qquad & \tau \le \hat{\tau}(Q) \qquad (4.2) \end{aligned}$$

Objective function (4.1) is the sum of fixed, variable, inventory holding, inventory salvage, and replacement costs for a single scenario. Constraint (4.2) ensures feasibility of the optimal switching time.

We note that in the discrete planning scenario with M possible return rates, there exist M such problems described above. The manufacturer's problem is to choose Q such that the expected total cost is minimized. With p_m as the probability that return rate m is realized, the manufacturer's problem can be represented by:

$$\text{Minimize}_Q \qquad \sum_{m \in M} p_m Z(\tau_m, \lambda_m, Q) \tag{4.3}$$

$$\text{Subject to} \qquad \tau_m \le \hat{\tau}_m(\lambda_m, Q) \qquad \forall m \in M \tag{4.4}$$

The order quantity is the only linking part of the manufacturer's problem. The scenarios have no interaction with each other. Therefore, for a given Q we solve M independent single scenario problems and recognize that the scenario switching times in the manufacturer's problem are equal to their optimal values when a single scenario is considered. We can then write the manufacturer's problem as:

$$\text{Minimize}_Q \qquad \sum_{m \in M} p_m Z^*(\lambda_m, Q) \tag{4.5}$$

$$\text{where} \qquad Z^*(\lambda_m, Q) = \left\{ \begin{array}{ll} \text{Minimize}_{\tau_m} & Z(\tau_m, \lambda_m, Q) \\ \text{subject to} & \tau_m \le \hat{\tau}_m(\lambda_m, Q) \end{array} \right\} \tag{4.6}$$

Given an order quantity Q, the problems above are straightforward and easy to solve. The optimal solutions identify the least cost switching time and quantity of repair parts to be salvaged for each scenario. From these two values, the least cost solution for a scenario can be easily calculated; in fact, calculating the expected cost over all scenarios becomes trivial. The manufacturer's problem reduces to a search over possible values of Q. To reduce the range of values over which Q must be considered, we observe that:

1. An optimal Q^* value will be integer.

2. An optimal Q^* value will be at least as large as the optimal order quantity found by solving the uncapacitated single product model in Chapter 3 using the smallest return

rate scenario. Let $Q_{lowerbound}$ represent this value.

3. An optimal Q^* value will be no greater than the optimal order quantity found by solving the uncapacitated single product model in Chapter 3 using the largest return rate scenario. Let $Q_{upperbound}$ represent this value.

With these bounds on the set of Q values to be considered, the manufacturer's problem can be solved using the following algorithm:

1. Let $\hat{Q} = Q_{lowerbound}$.

2. For each scenario $m \in M$, solve the single scenario problem defined by the optimization problem in (4.6) to get the minimum total cost, $Z^*(\lambda_m, \hat{Q})$.

3. Calculate the expected cost for $\hat{Q}$, $Z(\hat{Q}) = \sum_{m \in M} p_m Z^*(\lambda_m, \hat{Q})$.

4. $\hat{Q} \leftarrow \hat{Q} + 1$

5. Repeat steps 2 - 4 until $\hat{Q} = Q_{upperbound}$.

6. Q^* is the order quantity which gives the minimum $Z(\hat{Q})$ value found in the search above.

The algorithm above considers every feasible value of Q to find the optimal solution Q^*. Although this method considers a very large set of candidate solutions, empirical testing has shown that the computational time required to do so is quite small (less than 1 minute for the numerical examples in this chapter). Because the objective function is nonlinear, searching over the entire range of values guarantees that a global optimum will be found.

4.3.2 A Numerical Example

The following numerical example demonstrates the impact of multiple scenarios in the expected cost minimization problem. Figure 4.1 shows three possible return rate scenarios over a planning horizon of $T = 500$. The returns in all three scenarios follow a truncated bell-shaped curve with the peak return rate occuring at time $t = 250$. The *low* scenario observes 10,000 returned items while the *medium* and *high* scenarios observe 20,000 and 30,000 returns, repectively.

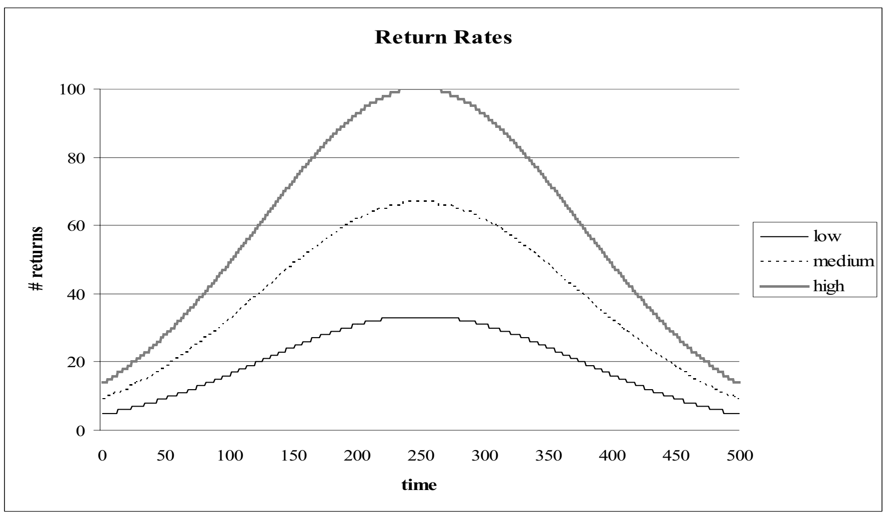

Figure 4.1: Return Rate Patterns For 3 Scenarios

This example uses a fixed cost of \$500 per unit time, a replacement cost of \$100 per unit, variable repair cost of \$75 per unit, unit holding costs of \$0.023 per unit time (this amounts to 15% of the variable cost for the entire planning period), and a net salvage cost of \$30 per unit. To generate the expected cost curve, we use the following probability distribution$(p_{low}, p_{medium}, p_{high}) = (0.333,\ 0.333,\ 0.333)$ to give equal weight to each of the three scenarios.

Figure 4.2 shows how the total cost function changes as the order quantity is varied for the low, medium, and high return rates as well as the expected cost curve. The optimal order quantity for the low scenario, Q^*_{low}, is 0 units, meaning that for this particular combination of parameter values it is cheaper to provide replacement units for all returned items than to repair a single return. The other optimal order quantities are as follows: Q^*_{medium} = 18,100, Q^*_{high} = 28,160, and $Q^*_{E(cost)}$ = 9,420 where $Q^*_{E(cost)}$ is the optimal solution to the problem defined by (4.5) - (4.6). We note that Q^*_{low}, Q^*_{medium}, and Q^*_{high} are optimal when considering each scenario individually and result in no salvaged items.

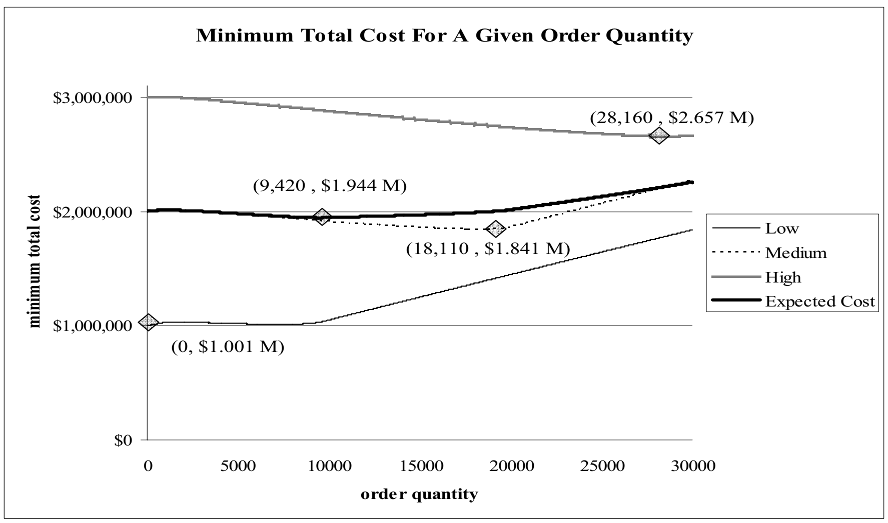

Figure 4.2: Total Cost Curves As A Function Of Order Quantity

Table 4.1 shows the total cost realized in each scenario when the above scenario-optimal order quantities are used. The optimal total cost values for each scenario are shown in the cells along the diagonal. For example, in the medium return rate scenario, the optimal order quantity is Q^*_{medium} = 18,100 and the resulting total cost is approximately $Z(\lambda_{medium}, Q = 18,100) \simeq$ \$1,841,000. This table allows us to see the cost penalty paid in any scenario by ordering the optimal quantity for a different scenario. For example, if the manufacturer is expecting returns to follow the medium return rate scenario and instead the low return rate is realized, he will order Q = 18,100 and will pay approximately \$1.373M. Had he known that the low return rate would occur, he could have ordered 0 units and paid \$1.001M, a difference of 37%.

Scenario	$Q^*_{low} = 0$	$Q^*_{med} = 18{,}110$	$Q^*_{high} = 28{,}610$	$Q^*_{E(cost)} = 9{,}420$
Low	\$1.001 M	\$1.373 M	\$1.785 M	\$1.029 M
Medium	\$2.001 M	\$1.841 M	\$2.215 M	\$1.918 M
High	\$3.001 M	\$2.760 M	\$2.657 M	\$2.886 M
E(cost)	\$2.000 M	\$1.991 M	\$2.219 M	\$1.944 M

Table 4.1: Total Cost Using Each Scenario's Optimal Order Quantity

While Figure 4.2 and Table 4.1 show the cost implications of ordering for a different scenario than what is realized, Figure 4.3 provides a different insight. Figure 4.3 shows what happens to the optimal switching time in each scenario as the order quantity is varied from 0 to 30,000 units. We see that the low return rate scenario has an optimal order quantity of Q^*_{low} = 0 and corresponding optimal switching time of $\tau^* = 0$, illustrating that the manufacturer should immediately begin replacing units. In Figure 4.3 we see that as the order quantity increases beyond a threshold of approximately 900 units, it becomes cost effective to begin repairing returned items in the low scenario. This is because the replacement costs and the added salvage cost of these items (which the manufacturer has already purchased) outweigh the costs of maintaining repair capabilities for the first 100 or so time periods. However, the optimal switching time decreases as the order quantity gets very large - this is mostly due to the increased holding costs of the added salvage items that are incurred. If we were to extend the order quantity axis in Figure 4.3 beyond 30,000 items, we would see the same effect in the medium and high return rate scenarios as well. Taking Figures 4.2 and 4.3 together, we see the relative sensitivity of the solution to changes in Q, specifically the total cost implications and the managerial policy implications (*i.e., switching time changes*).

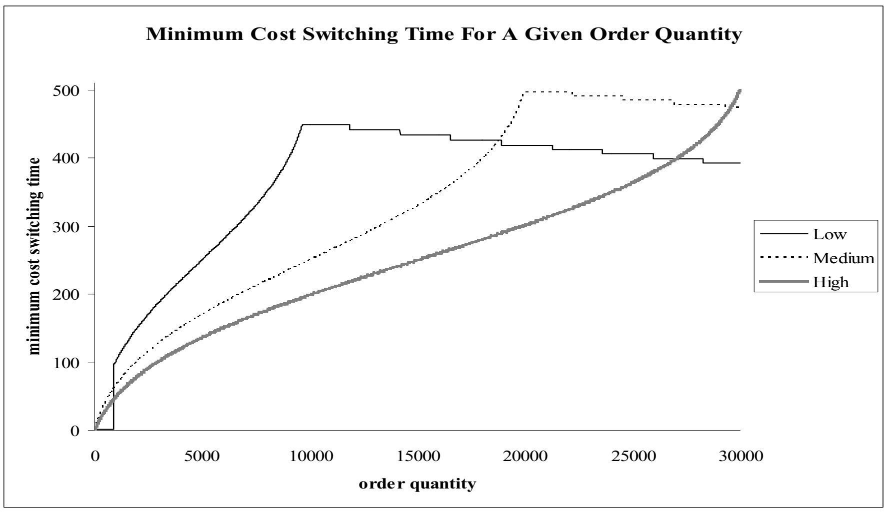

Figure 4.3: Order Quantity Vs. Optimal Switching Time For 3 Scenarios

4.3.3 Performance Of The Expected Cost Model Vs. Expected Return Scenario

Given a range of possible scenarios and the probability that each will occur, it may be natural to believe that solving the single scenario problem by using the expected return rate could be used as a substitute for the expected cost problem defined above. However, the nonlinearity of the objective function can lead to cases in which this approach is a poor substitute for solving the expected cost problem. In the numerical example above, the expected return scenario performs well when the optimal order quantity Q^* is very small or very large, but performs poorly over the most likely range of values for Q^*. Figure 4.4 shows the cost curves for the expected cost model and the single scenario model using a probability distribution of $(p_{low}, p_{medium}, p_{high}) = (0.333,\ 0.333,\ 0.333)$ to generate the expected return rate scenario.

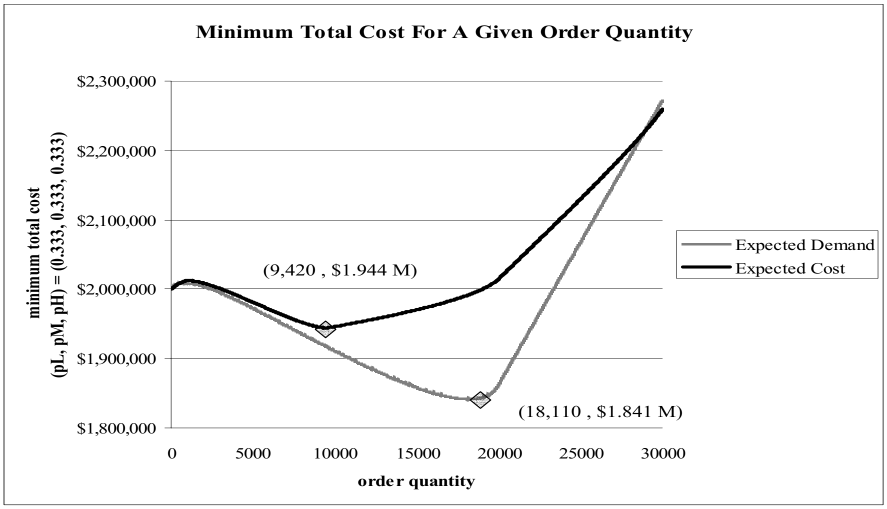

Figure 4.4: Total Cost Curves For The Expected Cost Formulation And The Expected Demand Scenario

From Figure 4.4, we see that using the expected demand in a single scenario as a substitute for the expected cost model would lead the manufacturer to order almost twice as many repair parts as necessary ($Q^*_{E(demand)}$ = 18,100 while $Q^*_{E(cost)}$ = 9,420)! From Table 4.1 we see that the total expected cost of ordering $Q^*_{E(demand)}$ = 18,100 is \$1.991M, which results in a cost penalty of approximately 2.4%. Although the penalty of ordering the expected demand is not very high for this example, it is important to recognize that it may be quite high for others. When possible, the expected cost model should be used since using the expected demand in a single scenario case can lead to drastically different order quantities resulting in severely suboptimal solutions.

4.4 The Minimax Regret Model

The objective of the minimax regret model is to protect the manufacturer against the worst possible outcome. In this context, regret is defined as the difference between the realized objective function value and the optimal objective function value for a particular scenario.

In the liftime buy problem, the manufacturer must decide the repair parts order quantity before the return rate is observed. If the manufacturer knew the return rate before the order was placed, then he would order Q to minimize his total cost and his regret would be zero. We note that there is a regret value for each scenario. By minimizing the maximum regret across all possible scenarios, the model minimizes the worst possible outcome. The objective function value in the minimax regret model is driven by a single scenario and does not consider the probability that this scenario will be realized. As a result, this model may be overly conservative. The solution performs well in the worst case scenario but its expected performance may be poor. The remainder of this section formulates the minimax regret model and presents an algorithm for its solution. We also extend the numerical example from Section 4.3 to examine the minimax regret curve and its sensitivity to changes in Q.

In addition to the notation defined for the expected cost model in Section 4.3, the minimax regret model defines Z_m^* as the optimal objective function value over all possible values of Q in scenario m. We define Y_m to be the regret in scenario m and Y to be the maximum regret over all scenarios. The minimax regret model can be formulated as:

$$\text{Minimize}_{Q,Y,Y_m,\tau_m} \quad Y \tag{4.7}$$

$$\text{Subject to} \quad Y \geq Y_m \qquad \forall m \in M \tag{4.8}$$

$$Y_m \geq Z(\tau_m, \lambda_m, Q) - Z_m^* \qquad \forall m \in M \tag{4.9}$$

$$\tau_m \leq \hat{\tau}_m(\lambda_m, Q) \qquad \forall m \in M \tag{4.10}$$

The objective function (4.7) minimizes the maximum regret over all scenarios. Constraint (4.8) ensures that the overall regret is greater than or equal to the regret in each of the component scenarios. Constraint (4.9) defines the regret in scenario m to be the difference between the cost in scenario m when we order Q and the cost when we order the optimal order quantity for scenario m. Finally, constraint (4.10) ensures that the switching time used in constraint (4.9) is feasible.

A solution algorithm similar to the one presented for the expected cost model in Section 4.3.1 can be used to solve the minimax regret problem. As in the expected cost model, solving the minimax regret problem reduces to a search over possible values of Q. We begin by noting that the optimal Q value will be integer and fall between $Q_{lowerbound}$ and $Q_{upperbound}$ as before (see Section 4.3.1 for definitions of $Q_{lowerbound}$ and $Q_{upperbound}$). With these bounds on the set of Q values to be considered, the minimax regret problem can be solved using the following algorithm:

1. Let $\hat{Q} = Q_{lowerbound}$.

2. For each scenario $m \in M$, solve the single scenario problem defined by the optimization problem in (4.6) to get the minimum total cost, $Z^*(\lambda_m, \hat{Q})$. Store these values in an array indexed ($\hat{Q}$, m)

3. $\hat{Q} \leftarrow \hat{Q} + 1$

4. Repeat steps 2 - 3 until $\hat{Q} = Q_{upperbound}$.

5. For each scenario $m \in M$, find the least cost solution over all values of $\hat{Q}$. Label this solution Z^*_m.

6. For each scenario $m \in M$ and each order quantity $\hat{Q} \in \{Q_{lowerbound}, Q_{upperbound}\}$, let $regret_{\hat{Q},m} = Z^*(\lambda_m, \hat{Q})$ - Z^*_m.

7. For each order quantity $\hat{Q} \in \{Q_{lowerbound}, Q_{upperbound}\}$, let $regret_{\hat{Q}}$ be the maximum value of $regret_{\hat{Q},m}$ over all scenarios $m \in M$.

8. Q^* is the order quantity which gives the minimum $regret_{\hat{Q}}$ value found in the search above.

As in the expected cost model, the algorithm above considers every feasible value of Q to find the optimal solution Q^*. Again, this method takes very little computational effort (less than 1 minute for the numerical examples considered here) and is guaranteed to find the global minimum solution.

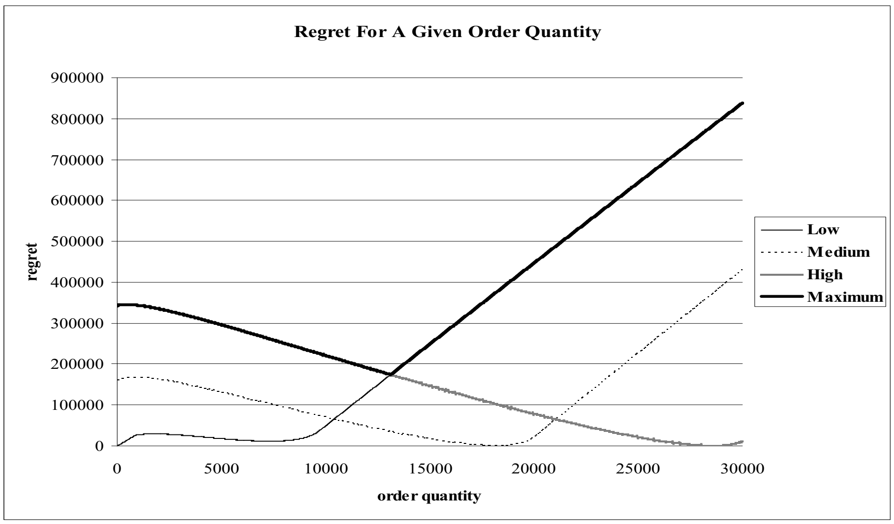

Figure 4.5: Regret In Each Scenario Over All Q Values

Figure 4.5 shows how the regret in each scenario changes as the order quantity Q is varied from 0 to 30,000 units. This figure also shows how the maximum regret curve changes with Q. We can see from Figure 4.5 that the maximum regret is driven by the high return rate scenario for $0 \leq Q \leq 13{,}120$ units. For order quantities larger than this, $13{,}120 \leq Q \leq 30{,}000$, the maximum regret curve is driven by the low return rate scenario. From this figure we see that the order quantity which minimizes the maximum regret is 13,120 which results in a maximum regret of $\simeq$ \$174,000.

4.5 Tradeoffs: Expected Cost Vs. Maximum Regret

In this chapter we present two approaches to the lifetime buy inventory problem when the return rate follows one of a finite number of known scenarios. The first approach protects the manufacturer against the average outcome by minimizing the expected cost over all of the scenarios. The second approach protects the manufacturer against the worst case outcome by minimizing the maximum regret for any single scenario. One could argue that the minimum expected cost approach performs well on average but leaves the manufacturer vulnerable when the realized return rate is extreme. One could just as easily argue that an objective that minimizes the maximum regret is excessively conservative and may produce solutions which do not perform well on average. As with any multi-objective problem there is no one *optimal* solution, but there exist a number of *efficient* solutions that perform well under both objectives. These *efficient* solutions are Pareto optimal - meaning that for a particular expected cost value, there is no other feasible solution which can achieve a lower maximum regret value, and vice versa. The manufacturer then decides which (expected cost, maximum regret) combination is the *best* for him.

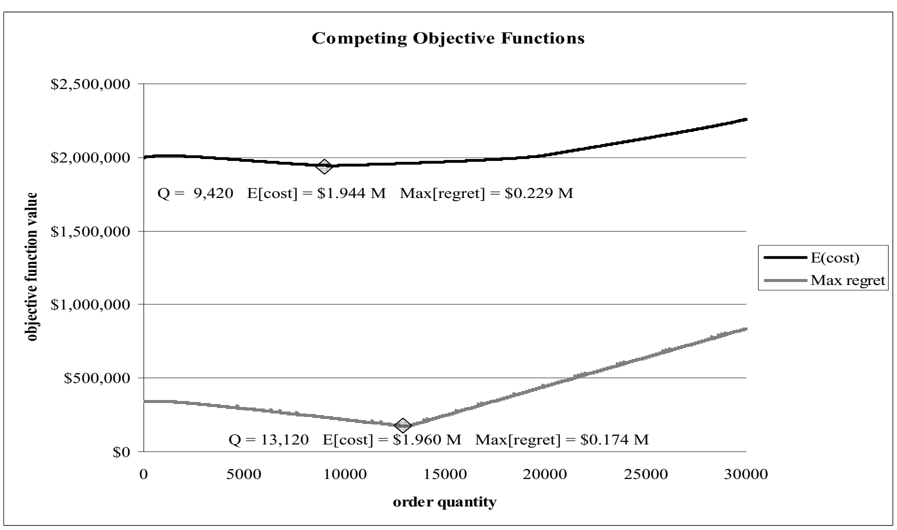

Figure 4.6: Minimum Expected Cost And Minimum Maximum Regret Curves

Using the numerical example from Sections 4.3 and 4.4, we explore the tradeoff between the expected cost and maximum regret models. Figure 4.6 depicts the expected cost curve and the maximum regret curve as the order quantity is varied from 0 to 30,000 units. We see that the optimal solution to the expected cost problem is $Q = 9{,}420$. This solution results in a total expected cost of \$1.944M and a maximum regret of \$0.229M. The optimal solution to the minimax regret problem is $Q = 13{,}120$ units which gives a total expected cost of \$1.960M and maximum regret of \$0.174M. We note that any other order quantity will have an expected cost greater than or equal to \$1.944M and a maximum regret greater than or equal to \$0.174M. We wish to find solutions that are as close as possible to both of these values. To do this, we explore the tradeoff curve of efficient solutions. We generate this tradeoff curve using the following weighting method algorithm (Cohon, 1978 and Solanki, 1991).

1. Solve the minimum expected cost model formulated in Section 4.3 to get order quantity $\hat{Q}_C$ and least cost $\hat{C}_C$. Calculate the maximum regret associated with $\hat{Q}_C$, $\hat{R}_C$. Label this solution ($\hat{C}_C$, $\hat{R}_C$).

2. Solve the minimum maximum regret model formulated in Section 4.4 to get order quantity $\hat{Q}_R$ and minimum maximum regret $\hat{R}_R$. Calculate the expected cost associated with $\hat{Q}_R$, $\hat{C}_R$. Label this solution ($\hat{C}_R$, $\hat{R}_R$).

3. We can draw a line connecting these two endpoints, defined by: $\alpha_1 \hat{C}_C + (1 - \alpha_1)\hat{R}_C = \alpha_1 \hat{C}_R + (1 - \alpha_1)\hat{R}_R$. Solving for α_1, we get $\alpha_1 = \frac{\hat{R}_R - \hat{R}_C}{\hat{C}_C - \hat{C}_R + \hat{R}_R - \hat{R}_C}$. By minimizing a combination of the two competing objective functions, we can obtain the efficient

solutions that define the tradeoff curve. Using α_1 defined above, we minimize:

$$\text{Minimize}_Q \quad \alpha_1 \sum_{m \in M} p_m Z^*(\lambda_m, Q) + (1-\alpha_1)Y \tag{4.11}$$

$$\text{Subject to} \quad Z^*(\lambda_m, Q) = \left\{ \begin{array}{ll} \text{Minimize}_{\tau_m} & Z(\tau_m, \lambda_m, Q) \\ \text{subject to} & \tau_m \leq \hat{\tau}_m(\lambda_m, Q) \end{array} \right\} \tag{4.12}$$

$$Y \geq Y_m \qquad \forall m \in M \tag{4.13}$$

$$Y_m \geq Z(\tau_m, \lambda_m, Q) - Z_m^* \qquad \forall m \in M \tag{4.14}$$

$$\tau_m \leq \hat{\tau}_m(\lambda_m, Q) \qquad \forall m \in M \tag{4.15}$$

This procedure is shown in Figure 4.7.

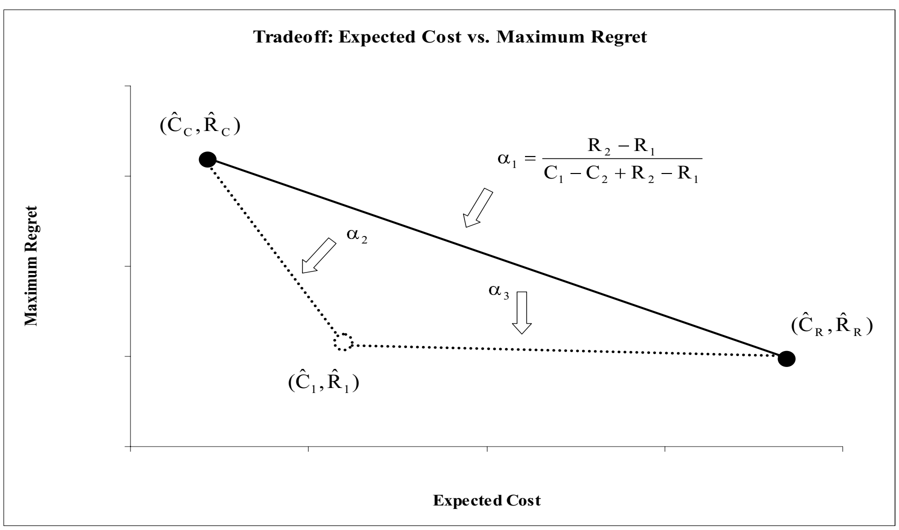

Figure 4.7: The Weighting Method For Generating Tradeoff Curves

4. We repeat Step 3 using solutions $(\hat{C}_C,\ \hat{R}_C)$ and $(\hat{C}_1,\ \hat{R}_1)$ to calculate α_2 and obtain a new solution $(\hat{C}_2,\ \hat{R}_2)$. The same is done for the pair $(\hat{C}_R,\ \hat{R}_R)$ and $(\hat{C}_1,\ \hat{R}_1)$ to calculate α_3 and obtain a new solution $(\hat{C}_3,\ \hat{R}_3)$. This step is repeated for all adjacent pairs of efficient solutions until doing so fails to produce a new solution.

Applying this method to our numerical example, we discover 13 efficient solutions. These order quantities as well as the expected cost and maximum regret associated with each one (rounded to the nearest thousand $) are shown in Table 4.2.

Q	E[*cost*]	Max[*regret*]
9,420	$1.944 M	$0.229 M
9,490	$1.944 M	$0.229 M
9,600	$1.944 M	$0.227 M
9,690	$1.945 M	$0.225 M
10,430	$1.948 M	$0.214 M
10,900	$1.950 M	$0.207 M
11,480	$1.953 M	$0.198 M
11,950	$1.955 M	$0.191 M
12,340	$1.957 M	$0.186 M
12,650	$1.958 M	$0.181 M
12,840	$1.959 M	$0.178 M
13,040	$1.960 M	$0.175 M
13,120	$1.960 M	$0.174 M

Table 4.2: Efficient Solutions

Figure 4.8 depicts the tradeoff curve containing these 13 efficient solutions. We see that aside from the regions immediately surrounding the endpoints, that the tradeoff curve can be approximated by a linear function. Fitting a linear model to the data, we obtain the following equation: Max[Regret] = -3.35 $*$ E[Cost] + $7,000,000 with R^2 value = 0.998 (nearly perfect correlation). From this curve and the solutions in Table 4.2, we see that the expected cost and maximum regret are quite robust with respect to a wide range of variation in the order quantities.

4.6 Conclusions And Directions For Future Research

In this chapter we present a scenario planning approach to the lifetime buy problem. This approach allows a decision maker to account for some measure of uncertainty in the number

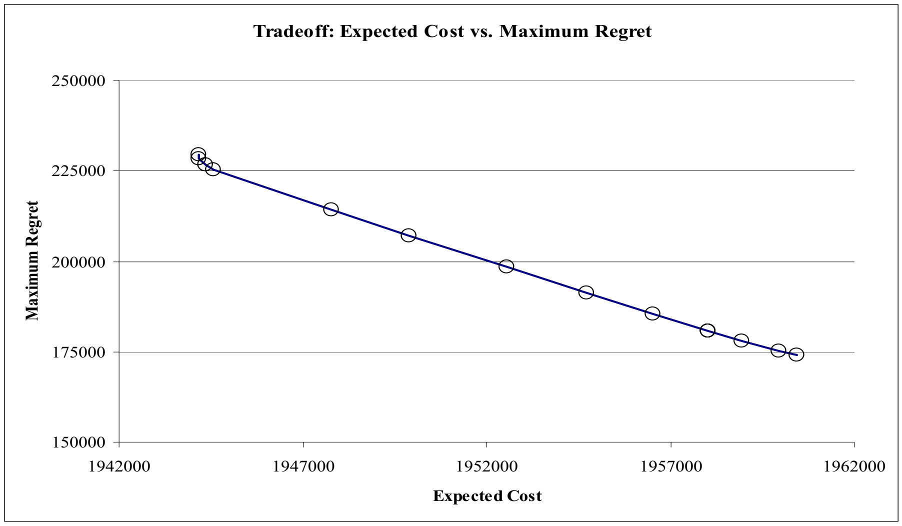

Figure 4.8: The Efficient Frontier Of Minimum Expected Cost And Minimum Maximum Regret Solutions

and timing of returns without having to define a full probability distribution for all possible outcomes by identifying a set of possible return rate scenarios. We develop two models for calculating lifetime buy order quantities to serve these returns based on two competing objectives. The first objective minimizes the expected cost over all scenarios, thus optimizing the average case performance of the system. The second objective minimizes the maximum regret over all scenarios, optimizing the worst case performance. Finally, we explore the tradeoff curve of solutions which are Pareto optimal in this two objectives.

The scenario planning models in this chapter should be extended to include the repair capacity constraints and the shared facilities for two products models considered in Chapter 3. In addition, many of the future research topics identified in Chapter 3 would be beneficial in the scenario planning approach, including examining the role of overlapping product generations and the impact of parts commonality across multiple products. Other extensions such as delaying the timing of the lifetime buy or allowing multiple procurement opportunities would be interesting to pursue in a scenario planning context. Another area of future research concerns the search method for the order quantity. The current solution method for the expected cost and the minimax regret models evalutes a large number of order quantities in find the optimal value Q^*. Although this method executes quickly for the numerical examples in this chapter, a more intelligent search methods for finding Q^* should be explored since this method can become computationally intensive for large problem instances. Finally, it would be interesting to model other scenario planning objectives identified in Section 4.2 to find solutions which are robust to a variety of objective functions.

Chapter 5

Conclusions and Future Research

In this dissertation we present models that explicitly consider the role of the product lifecycle in the effective management of closed-loop supply chains. For our purpose, the term closed-loop refers to supply chains that include both forward and reverse product flows, including traditional supply chain activities such as product manufacturing and distribution as well as product returns, repair operations, and remanufacturing activities. Within this context, we study the impact of product returns on logistics network design in a deterministic setting and the impact of lifetime buys on repair operations in both a deterministic and a stochastic, scenario-planning setting.

The first problem that we consider is a logistics network design problem that effectively serves bidirectional product flows during the introductory, maturity, and decline stages of a product's lifecycle. We examine the role of product returns in the design of distribution networks and the implications of simultaneously considering product delivery and returns collection needs on bidirectional logistics network configurations. We outline three extensions of the traditional uncapacitated fixed charge facility location model and develop Lagrangian relaxation-based solution algorithms which have proven to be quick and effective.

Our computational results suggest that integrated solutions that simultaneously consider bidirectional flows outperform solutions in which forward and reverse network design decisions are made sequentially or independently of each other by as much as 30%. These results further show that these two approaches to bidirectional network design can result in highly dissimilar networks when forward and reverse costs are misaligned. This research indicates that the value and impact of integrated bidirectional network design is significant in a single product, two-tier supply chain system, leading us to identify at least three directions for future study. The first of these would develop a time-integrated model to capture the evolution of a closed-loop logistics network over the entire span of a product's lifecycle. The current models consider each phase of the lifecycle independently, so there would be valuable insight in exploring how these networks evolve over time when all three phases are considered simultaneously. A second extension of this work would develop models that serve multiple products at different stages of their lifecycles. It would be interesting to evaluate how the network structure changes in response to multiple products with unique return rates and processing costs. Finally, it would be valuable to model customer behavior in the network design process. The current models accept return rates as exogenous and independent of the network structure. However, customer behavior drives the return rates in voluntary collection programs. For example, customers may be distance-sensitive. In this case, the return rate would be a function of the network structure. Such behavioral extensions would be valuable to firms that must comply with returns collection legislation.

The second problem that we consider investigates the impact of lifetime buys on warranty repair operations. Chapter 3 studies the lifetime buy problem in a deterministic setting while Chapter 4 examines the impact of return rate uncertainty by developing scenario planning variations of the models in Chapter 3. Among the deterministic models that we develop are a single product uncapacitated model, a single product with repair capacity constraints model, and a two-product with shared repair facilities model. This research develops a framework for identifying the relevant costs for a repair operation facing a lifetime buy decision and allows the manufacturer to identify a *switching time* before which returns are repaired and after which they are replaced. This switching time is used to calculate the optimal lifetime buy order quantity and resulting total cost for the system. The scenario planning models of Chapter 4 extend the single uncapacitated model to consider the impact of return rate uncertainty on the optimal order quantity and resulting switching time. Chapter 4 identifies two suitable objectives for the scenario planning model, one in which the expected total cost of the system is minimized and one in which the maximum regret is minimized. These two somewhat competing objectives optimize the expected and the worst case performance of the system. A tradeoff curve of Pareto-optimal solutions shows that there exists a range of order quantities that are robust in both objectives.

This research is the first study to investigate lifetime buy decisions and the role of relevant costs in this problem. This work establishes a foundation for extensions in both a deterministic and a stochastic return rate environment. One valuable extension of the deterministic models would examine the role of overlapping product generations and the viability of using procurement opportunities for later generation models to support the repair operations of earlier generation products. This would give some insight into the value of a second procurement opportunity during the planning period. This work would lead to another model to explore the impact of parts commonality across multiple products. A third extension of the deterministic case would consider the effect of order-size-based pricing for repair parts. There are many valuable extensions to explore when the return rate is uncertain. Among these are a variety of scenario-planning models based on other objective functions. It would also be valuable to develop stochastic models to explore supplier contracting options, pricing to reserve supplier capacity, and the value of multiple procurment opportunities.

The work in this dissertation is one of the first to explicitly consider the role of the product lifecycle in closed-loop supply chain management. The models presented here establish a foundation for considering extensions in bidirectional logistics network design as well as inventory planning for repair operations for short lifecycle products. It is our hope that these models will provide value and insight to decision makers in these fields.

References

R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows*, Prentice Hall, Upper Saddle River (1993).

I. Averbakh and O. Berman, "Minimax Regret P-Center Location on a Network with Demand Uncertainty," *Location Science* 5, 247 - 254 (1997).

K. R. Balachandran, R. A. Maschmeyer, and J. L. Livingstone, Product Warranty Period: A Markovian Approach to Estimation and Analysis of Repair and Replacement Costs, *The Accounting Review* 56, 115 - 124 (1981).

Y. Balcer and I Sahin, Replacement Costs Under Warranty: Cost Moments and Time Variability, *Operations Research* 34, 554 - 559 (1986).

M.L. Balinski, Integer Programming: Methods, Uses, Computation, *Management Science* 12, 253 - 313 (1965).

A. I. Barros, R. Dekker, and V. Scholten, A Two-Level Network for Recycling Sand: a Case Study, *European Journal of Operational Research* 110, 199 - 214 (1998).

B. M. Beamon and C. Fernandes, Supply Chain Network Configuration for Product Recovery, *Production Planning & Control* 15, 270 - 281 (2004).

J. M. Bloemhof-Ruwaard, M. Salomon, and L. Van Wassenhove, The Capacitated Distribution and Waste Disposal Problem, *European Journal of Operational Research* 88, 490 - 503 (1996).

G. Chen, M. S. Daskin, Z-J. M. Shen, and S. Uryasev, "The α-Reliable Mean-Excess Regret Model for Stochastic Facility Location Modeling," *Naval Research Logistics* 53, 617 - 626 (2006).

S. Chopra and P. Meindl, *Supply Chain Management: Strategy, Planning, and Operation* 2nd edition, Pearson Prentice Hall, Upper Saddle River NJ (2004).

CIA The World Fact Book Rank Order GDP. (n.d.). Retrieved March 13, 2006 from http://www.cia.gov/cia/publications/factbook/rankorder/2001rank.html

J. L. Cohon, *Multiobjective Programming and Planning*, Academic Press, New York (1978).

H. Crowder, Computational Improvements For Subgradient Optimization, *Symposia Mathematica*, Academic Press 19 (1976).

J. Current, S. Ratick, and C. ReVelle, "Dynamic Facility Location when the Number of Facilities is Uncertain: A Decision Analysis Approach," *European Journal of Operational Research* 110, 597 - 609 (1997).

M. S. Daskin, S. M. Hesse, and C. ReVelle, "α-Reliable P-Minimax Regret: A New Model for Strategic Facility Location Modeling," *Location Science* 5, 227 - 246 (1997).

M. De Brito, "Managing Reverse Logistics or Reversing Logistics Management?", Ph.D. thesis, Erasmus University Rotterdam, Rotterdam Netherlands (2003).

T. Enright, Post Holiday Logistics, *trafficWORLD* 20, (2003).

M. L. Fisher, The Lagrangian Relaxation Method For Solving Integer Programming Problems, *Management Science* 32, 1 - 18 (1981).

M. L. Fisher, An Applications Oriented Guide To Lagrangian Relaxation, *Interfaces* 15, 1 - 121 (1985).

M. Fleischmann, H. R. Krikke, R. Dekker, and S. D. P. Flapper, A Characterisation of Logistics Networks for Product Recovery, *Omega* 28, 653 - 666 (2000).

M. Fleischmann, P. Buellens, J. Boemhof-Ruwaard, and L. Van Wassenhove, The Impact of Product Recovery on Logistics Network Design, *Production and Operations Management* 10, 156 - 173 (2001).

M. Fleischmann, J. Bloemhof-Ruwaard, P. Beullens, and R. Dekker, Reverse Logistics Network Design in *Reverse Logistics: Quantitative Models for Closed-Loop Supply Chains*, Springer-Verlag, New York, (2004).

V. D. R. Guide Jr. and R. Srivastava, Repairable Inventory Theory: Models and Applications, *European Journal of Operational Research* 102, 1 - 20 (1997).

V. D. R. Guide Jr., and L. Van Wassenhove, "The Evolution of Closed Loop Supply Chain

Research," *Vanderbilt Closed Loop Supply Chain Conference*, Nashville USA (2005).

F. W. Harris, How Many Parts to Make at Once, *Factory, The Magazine of Management* 10, 135 - 36, 152. (1913).

S. Ja, et al., A Non-Renewable Minimal-Repair Warranty Policy with Time-Dependent Costs, *IEEE Transactions on Reliability* 50, 346 - 352 (2001).

V. Jayaraman, V. D. R. Guide Jr, and R. Srivastava, A Closed-Loop Logistics Model for Remanufacturing, *Journal of the Operational Research Society* 50, 497 - 508 (1999).

V. Jayaraman, R. A. Patterson, and E. Rolland, The Design of Reverse Distribution Networks: Models and Solution Procedures, *European Journal of Operational Research* 150, 128 - 149 (2003).

P. Kouvelis and G. Yu, *Robust Discrete Optimization and Its Applications*, Kluwer Academic Publishers, Amsterdam (1997).

H. Krikke, Recovery Strategies and Reverse Logistic Network Design, Ph.D. thesis, University of Twente, Netherlands (1998).

D. W. Krumwiede and S. Chwen, A Model for Reverse Logistics Entry by Third-Party Providers, *Omega* 30, 325 - 333 (2002).

A. Marin and B. Pelegrin. The Return Plant Location Problem: Modelling and Resolution, *European Journal of Operational Research* 104, 375 - 392 (1998).

D. N. P. Murthy and W. R. Blischke, Strategic Warranty Management: A Life-Cycle Approach, *IEEE Transactions on Engineering Management* 47, 40 - 54 (2000).

D. N. P. Murthy and I. Djamaludin, New Product Warranty: A Literature Review, *International Journal of Production Economics* 79, 231 - 260 (2002).

S. Nahmias, Managing Repairable Item Inventory Systems: A Review in L. B. Schwarz (Ed.), *Multi-Level Inventory Control Systems: Theory and Practice,* North-Holland Publishing Company, New York NY, 253 - 277 (1981).

S. H. Owen and M. S. Daskin, "Strategic Facility Location: A Review," *European Journal of Operational Research* 111, 423 - 447 (1998).

S. H. Owen, "Scenario Planning Approaches to Facility Location: Models and Solution Methods," Ph.D. thesis, Northwestern University, Evanston Illinois (1999).

N. C. Petruzzi and M. Dada, Pricing and the Newsvendor Problem: A Review with Extensions, *Operations Research* 47, 183 - 194 (1999).

E. L. Porteus, Stochastic Inventory Theory, in *Handbooks in Operations Research and Management Science,* Volume 2: Stochastic Models, Elsevier Science Publishers B. V., North Holland, (1990).

E. L. Porteus, *Foundations of Stochastic Inventory Theory,* Stanford University Press, Stanford CA (2002).

G. Ringwald, *Scenario Planning Managing for the Future*, John Wiley & Sons, New York (1998).

D. S. Rogers and R. Tibben-Lembke, *Going Backwards: Reverse Logistics Trends and Practices*, Reverse Logistics Executive Council, Reno, (1999).

R. C. Savaskan and C. Corbett, Contracting and Coordination in Closed-Loop Supply Chains, in *Business Aspects of Closed-Loop Supply Chains, International Management Series: Vol.2*, Carnegie Bosh Institute, (2002).

R. C. Savaskan, L. Debo, and L. N. Van Wassenhove, "Coordination in Closed-Loop Supply Chains, in *Reverse Logistics: Quantitative Models for Closed-Loop Supply Chains*, Springer-Verlag, New York, (2003).

R. C. Savaskan, S. Bhattacharya, and L. N. Van Wassenhove, Closed-Loop Supply Chain Models with Remanufacturing, *Management Science* 50, 239 - 252 (2004).

R. C. Savaskan and U. Aytekin,A Dynamic Life Cycle Pricing Model with Product Remanufacturing under review for *Management Science*, (2005).

R. C. Savaskan, Strategic Decentralization of Reverse Channel in a Closed Loop Supply Chain: The Case of Competing Retailers *Management Science* 52, 1 - 14 (2006).

R. C. Savaskan, A. Coughlan, and J. Shulman, Product Returns Policies and Value of Information, Working Paper, Kellogg School of Management, Northwestern University (2006a).

R. C. Savaskan, A. Coughlan, and J. Shulman, Product Returns Policies under Competition, Working Paper, Kellogg School of Management, Northwestern University (2006b).

D. Serra and V. Marianov, "The P-Median Problem in a Changing Network: The Case of Barcelona," *Location Science* 6, 383 - 394 (1998).

D. Serra, S. Ratick, and C. ReVelle, "The Maximum Capture Problem with Uncertainty," *Environment and Planning B* 23, 49 - 59 (1996).

C. Sherbrooke, METRIC: A Multi-Echelon Technique for Recoverable Item Control, *Operations Research* 16, 122 - 141 (1968).

L. V. Snyder, "Supply Chain Robustness and Reliability: Models and Algorithms," Ph.D. thesis, Northwestern University, Evanston Illinois (2003).

R. Solanki, "Generating the Noninferior Set in Mixed Biojective Linear Programs: An Application to a Location Problem," *Computers and Operations Research* 18, 1 - 15 (1991).

T. Spengler, H. Puchert, T. Penkuhn, and O. Rentz, Environmental Integrated Production and Recycling Management, *European Journal of Operational Research* 97, 308 - 326 (1997).

J. Stock, T. Speh, and H. Shear. "Many Happy (Product) Returns," *Harvard Business Review* July, 16 - 17 (2002).

G. P. Wadsworth, Probability, in *Notes on Operations Research,* 1959, Technology Press, (1959).

M. J. Zuo, B. Liu, and D. N. P. Murthy, Replacement-Repair Policy for Multi-State Deteriorating Products under Warranty, *European Journal of Operational Research* 123, 519 - 530 (2000).

Appendix A: Detailed Results For 150-Node And 263-Node Data Sets

Table A.1: 150 Node Data Set Detailed Results

150 node data set																				
parameter values			Co-Location				Forward Dominant				Reverse Dominant				Forward Similarity			Reverse Similarity		
α	β	γ	lagrangian solution time	cplex solution time	# forward locations	# reverse locations	lagrangian solution time	cplex solution time	# forward locations	# reverse locations	lagrangian solution time	cplex solution time	# forward locations	# reverse locations	CL vs. FD	CL vs. RD	FD vs. RD	CL vs. FD	CL vs. RD	FD vs. RD
0.25	0.25	0.50	0.55	434.51	7	6	0.64	798.88	7	5	0.44	540.54	7	7	0.00	0.00	0.00	0.14	0.17	0.31
0.25	0.25	1.00	0.53	337.31	7	7	0.42	726.41	7	7	0.42	543.26	7	7	0.00	0.00	0.00	0.00	0.00	0.00
0.25	0.25	2.00	0.46	239.35	7	12	0.67	574.54	7	7	0.68	330.21	7	12	0.00	0.00	0.00	0.79	0.00	0.79
0.25	0.75	0.50	3.92	1043.30	7	4	0.61	1278.00	7	2	0.58	1006.74	5	5	0.00	0.32	0.32	0.34	0.20	0.45
0.25	0.75	1.00	1.07	864.24	7	5	2.07	1001.61	7	4	0.58	1054.67	6	6	0.04	0.23	0.19	0.20	0.18	0.28
0.25	0.75	2.00	0.67	635.73	7	7	0.46	812.20	7	6	0.55	797.32	7	7	0.00	0.00	0.00	0.17	0.00	0.17
0.25	2.00	0.50	1.21	1469.05	7	1	0.44	2185.10	7	1	0.65	1554.75	4	4	0.00	0.44	0.44	0.00	0.71	0.71
0.25	2.00	1.00	0.90	1273.83	7	2	1.03	1724.50	7	2	1.06	1623.14	4	4	0.00	0.44	0.44	0.00	0.40	0.40
0.25	2.00	2.00	0.68	899.33	7	4	0.65	1116.92	7	4	0.71	1269.84	4	4	0.00	0.37	0.37	0.00	0.02	0.02
0.75	0.25	0.50	0.51	326.26	7	9	0.42	647.49	7	7	0.43	526.78	7	9	0.00	0.00	0.00	0.16	0.00	0.16
0.75	0.25	1.00	0.59	201.72	7	16	0.59	560.74	9	9	0.60	380.99	7	16	0.16	0.00	0.16	0.95	0.00	0.95
0.75	0.25	2.00	1.13	175.55	7	29	0.85	337.35	12	12	0.48	292.95	7	29	0.84	0.00	0.84	1.20	0.00	1.20
0.75	0.75	0.50	0.58	796.82	7	6	0.59	872.14	7	5	0.43	795.55	7	7	0.04	0.00	0.04	0.14	0.17	0.31
0.75	0.75	1.00	0.55	467.70	7	7	0.46	646.49	7	7	0.39	579.11	7	7	0.00	0.00	0.00	0.00	0.00	0.00
0.75	0.75	2.00	0.85	294.94	8	12	0.54	528.03	9	9	0.46	377.00	7	12	0.06	0.11	0.16	0.63	0.00	0.63
0.75	2.00	0.50	5.82	1172.59	7	4	1.12	1227.60	7	3	0.81	2018.58	4	4	0.03	0.37	0.40	0.17	0.02	0.16
0.75	2.00	1.00	0.56	667.29	7	5	0.59	833.55	7	4	0.56	1239.73	5	5	0.03	0.28	0.31	0.10	0.01	0.10
0.75	2.00	2.00	0.52	700.62	7	7	0.43	678.53	7	7	0.42	866.91	7	7	0.00	0.00	0.00	0.00	0.00	0.00
2.00	0.25	0.50	0.79	186.89	7	20	0.54	384.23	10	10	0.61	309.16	7	20	0.39	0.00	0.39	0.98	0.00	0.98
2.00	0.25	1.00	0.65	203.44	7	30	0.58	264.39	16	16	0.60	242.80	7	30	1.15	0.00	1.15	0.95	0.00	0.95
2.00	0.25	2.00	0.55	148.21	7	47	0.56	183.43	20	20	0.54	203.72	7	47	1.40	0.00	1.40	1.53	0.00	1.53
2.00	0.75	0.50	0.43	381.62	8	9	0.63	595.21	7	7	0.50	480.25	7	9	0.11	0.11	0.00	0.16	0.00	0.16
2.00	0.75	1.00	0.82	281.24	8	16	0.63	363.84	12	12	0.94	358.87	7	16	0.73	0.11	0.83	0.37	0.00	0.37
2.00	0.75	2.00	0.95	175.27	8	27	0.58	230.58	16	16	0.81	308.04	7	27	1.05	0.11	1.16	0.75	0.00	0.75
2.00	2.00	0.50	0.61	714.79	7	6	0.56	832.71	7	5	0.68	1135.16	5	5	0.04	0.32	0.28	0.14	0.15	0.01
2.00	2.00	1.00	0.87	396.08	7	7	0.46	652.08	7	7	0.55	646.94	7	7	0.00	0.00	0.00	0.00	0.00	0.00
2.00	2.00	2.00	0.64	246.41	12	12	0.62	344.18	12	12	0.88	468.93	7	12	0.00	0.84	0.84	0.00	0.00	0.00

α	return rate	CL	Co-location model
β	reverse fixed cost ratio	FD	Forward Dominant model
γ	reverse transportation cost ratio	RD	Reverse Dominant model

Table A.2: 263 Node Data Set Detailed Results

263 node data set																				
parameter values			Co-Location				Forward Dominant				Reverse Dominant				Forward Similarity			Reverse Similarity		
α	β	γ	lagrangian solution time	cplex solution time	# forward locations	# reverse locations	lagrangian solution time	cplex solution time	# forward locations	# reverse locations	lagrangian solution time	cplex solution time	# forward locations	# reverse locations	CL vs. FD	CL vs. RD	FD vs. RD	CL vs. FD	CL vs. RD	FD vs. RD
0.25	0.25	0.50	5.38	10031.00	6	5	2.72	13899.80	6	4	3.02	7741.09	5	5	0.00	0.13	0.13	0.24	0.01	0.25
0.25	0.25	1.00	7.58	7535.81	6	6	5.45	11862.40	6	6	4.78	10665.40	6	6	0.00	0.00	0.00	0.00	0.00	0.00
0.25	0.25	2.00	4.76	6157.52	6	9	4.88	10282.10	6	6	4.39	6657.59	6	9	0.00	0.00	0.00	0.56	0.00	0.56
0.25	0.75	0.50	5.63	10925.70	6	3	5.24	17776.40	6	3	4.04	14558.30	4	4	0.00	0.37	0.37	0.00	0.25	0.25
0.25	0.75	1.00	5.28	11276.50	6	4	9.66	14838.50	6	4	3.60	11325.60	5	5	0.00	0.13	0.13	0.00	0.25	0.25
0.25	0.75	2.00	5.75	8873.41	6	5	5.97	8934.26	6	5	3.70	10943.30	5	5	0.00	0.13	0.13	0.00	0.01	0.01
0.25	2.00	0.50	5.52	11522.90	6	2	6.34	16828.60	6	2	7.78	15546.30	4	4	0.00	0.37	0.37	0.00	0.61	0.61
0.25	2.00	1.00	5.07	10827.40	6	3	6.21	19014.10	6	2	4.34	15523.20	4	4	0.06	0.37	0.35	0.54	0.24	0.58
0.25	2.00	2.00	3.47	10612.30	6	4	5.67	14181.80	6	3	4.98	11018.30	4	4	0.01	0.37	0.37	0.27	0.07	0.25
0.75	0.25	0.50	4.37	6814.25	6	7	5.04	9429.41	6	6	4.76	7953.56	6	7	0.00	0.00	0.00	0.13	0.00	0.13
0.75	0.25	1.00	5.52	4694.08	6	13	2.24	6667.26	7	7	3.39	6768.04	6	13	0.16	0.00	0.16	0.87	0.00	0.87
0.75	0.25	2.00	4.75	3403.68	6	15	2.40	5703.77	9	9	4.73	4794.92	6	15	0.58	0.00	0.58	0.73	0.00	0.73
0.75	0.75	0.50	4.38	10112.10	6	5	4.33	13095.20	6	4	5.15	9227.01	5	5	0.00	0.13	0.13	0.24	0.01	0.25
0.75	0.75	1.00	4.24	8696.25	6	6	2.37	8805.01	6	6	6.87	9564.81	6	6	0.00	0.00	0.00	0.00	0.00	0.00
0.75	0.75	2.00	5.02	6465.67	6	9	2.44	6377.71	7	7	3.41	6556.74	6	9	0.14	0.00	0.14	0.42	0.00	0.42
0.75	2.00	0.50	6.30	13410.60	6	3	6.22	16357.80	6	3	3.63	12961.00	4	4	0.00	0.38	0.38	0.00	0.25	0.25
0.75	2.00	1.00	4.43	13178.80	6	4	2.53	13827.30	6	4	4.68	14198.90	4	4	0.00	0.37	0.37	0.00	0.06	0.06
0.75	2.00	2.00	4.05	10657.70	6	6	6.33	11632.40	5	5	5.38	8436.23	5	5	0.13	0.13	0.00	0.13	0.13	0.00
2.00	0.25	0.50	5.78	5031.64	6	14	1.93	8215.28	7	7	5.12	7059.64	6	14	0.13	0.03	0.16	1.03	0.12	0.98
2.00	0.25	1.00	5.79	3923.01	6	22	2.85	4732.37	10	10	3.64	5372.32	6	22	0.77	0.00	0.77	1.32	0.00	1.32
2.00	0.25	2.00	5.23	1966.16	6	33	2.78	4283.96	14	14	4.11	4166.38	5	33	1.21	0.12	1.43	1.52	0.00	1.52
2.00	0.75	0.50	5.89	8089.92	6	7	4.31	9067.69	6	6	4.22	6634.47	6	7	0.01	0.00	0.01	0.13	0.00	0.13
2.00	0.75	1.00	4.39	5324.77	6	11	2.39	8899.04	7	7	3.00	5500.68	6	11	0.17	0.00	0.17	0.66	0.00	0.66
2.00	0.75	2.00	5.17	4132.92	6	14	5.16	4135.44	13	13	4.70	7820.18	6	14	1.07	0.00	1.07	0.13	0.00	0.13
2.00	2.00	0.50	4.52	10101.30	6	5	2.61	11820.10	6	4	5.47	12862.40	5	5	0.00	0.13	0.13	0.24	0.01	0.25
2.00	2.00	1.00	4.16	7616.88	6	6	2.59	8682.00	6	6	4.13	11720.50	6	6	0.00	0.00	0.00	0.00	0.00	0.00
2.00	2.00	2.00	4.57	7603.73	9	9	3.67	8603.60	7	7	5.44	7084.47	6	9	0.42	0.58	0.14	0.42	0.00	0.42

α	return rate	CL	Co-location model
β	reverse fixed cost ratio	FD	Forward Dominant model
γ	reverse transportation cost ratio	RD	Reverse Dominant model

Appendix B: Second Order Conditions for the Single Product Capacitated Model

We derive the second order conditions for the single product with repair capacity constraints model below. The conditions for each region are presented separately.

Region I

$$\begin{aligned}
&\text{Total Cost} \quad TC = F\tau + vQ_I + hQ_I\tau - h\int_0^\tau \Lambda(t)dt + r\int_\tau^T \lambda(t)dt \\
&\text{F.O.C.} \quad \frac{dTC}{d\tau} = F + v\lambda(\tau) + h\tau\lambda(\tau) - r\lambda(\tau) = 0 \\
&\tau^* = \frac{1}{h}\left(r - v - \frac{F}{\lambda(\tau^*)}\right) \\
&\text{S.O.C.} \quad \frac{d^2TC}{d\tau^2} = v\lambda'(\tau) + h\tau\lambda'(\tau) + h\lambda(\tau) - r\lambda'(\tau) > 0 \\
&\lambda^2(\tau^*) > \frac{F}{h}\lambda'(\tau^*) \qquad \textit{substitute } \tau^* \textit{ from FOC}
\end{aligned}$$

Region II

$$\begin{aligned}
\text{Total Cost}\quad TC &= F\tau + vQ_{II} \\
&+ h\left[\int_0^{t_1}\left[Q_{II} - \int_0^y \lambda(t)dt\right]dy + \int_{t_1}^{\tau}\left[Q_{II} - \Lambda(t_1) - K\right]dy\right] \\
&+ r\left[\int_{t_1}^{\tau}(\lambda(t) - K)\,dt + \int_{\tau}^{T}\lambda(t)dt\right]
\end{aligned}$$

$$\text{F.O.C.}\quad \frac{dTC}{d\tau} = F + vK + 2hK\tau - hKt_1 - hK - rK = 0$$

$$\tau^* = \frac{1}{2}\left(1 + t_1 + \frac{r-v}{h} - \frac{F}{hK}\right)$$

$$\text{S.O.C.}\quad \frac{d^2TC}{d\tau^2} = 2hK > 0$$

Region III

$$\begin{aligned}
\text{Total Cost}\quad TC &= F\tau + vQ_{III} \\
&+ h\left[\int_0^{t_1}\left[Q_{III} - \int_0^y \lambda(t)dt\right]dy + \int_{t_1}^{t_2}\left[Q_{III} - \Lambda(t_1) - K\right]dy\right] \\
&+ h\left[\int_{t_2}^{\tau}\left[Q_{III} - \Lambda(t_1) - K(t_2 - t_1) - \int_{t_2}^{y}\lambda(t)dt\right]\right]dy \\
&+ r\left[\int_{t_1}^{t_2}(\lambda(t) - K)\,dt + \int_{\tau}^{T}\lambda(t)dt\right]
\end{aligned}$$

$$\text{F.O.C.}\quad \frac{dTC}{d\tau} = F + v\lambda(\tau) + h\tau\lambda(\tau) - r\lambda(\tau) = 0$$

$$\tau^* = \frac{1}{h}\left(r - v - \frac{F}{\lambda(\tau^*)}\right)$$

$$\text{S.O.C.}\quad \frac{d^2TC}{d\tau^2} = v\lambda'(\tau) + h\tau\lambda'(\tau) + h\lambda(\tau) - r\lambda'(\tau) > 0$$

$$\lambda^2(\tau^*) > \frac{F}{h}\lambda'(\tau^*) \qquad \textit{substitute } \tau^* \textit{ from FOC}$$

Appendix C: Proof of the First Order Conditions for the Two Product Case Values

To simplify the total cost expression, let:

$$g_1(\tau_1) = v_1Q_1 + h_1Q_1\tau_1 - h_1\int_0^{\tau_1} \Lambda_1(t)dt + r_1\int_{\tau_1}^{T} \lambda_1(t)dt$$

$$g_2(\tau_2) = v_2Q_2 + h_2Q_2\tau_2 - h_2\int_0^{\tau_2} \Lambda_2(t)dt + r_2\int_{\tau_2}^{T} \lambda_2(t)dt$$

$$\Rightarrow TC(\tau_1, \tau_2) = F * max(\tau_1, \tau_2) + g_1(\tau_1) + g_2(\tau_2)$$

We can formulate this problem as a constrained optimization model described by:

$$\begin{aligned} &Minimize && g_1(\tau_1) + g_2(\tau_2) + F\tau_2 \\ &Subject\ to && \tau_2 - \tau_1 \geq 0 \end{aligned}$$

Equivalently, we can write the Lagrangian:

$$Minimize \qquad g_1(\tau_1) + g_2(\tau_2) + F\tau_2 - \mu(\tau_2 - \tau_1)$$

We introduce a Lagrange multiplier, μ, and note that there exist $\tau_1^*, \tau_2^*, \mu^*$ such that the following conditions are satisfied:

$$\frac{dg_1(\tau_1^*)}{d\tau_1^*} + \mu^* = 0 \qquad \frac{dg_2(\tau_2^*)}{d\tau_2^*} + F + \mu^* = 0$$

$$\tau_2^* - \tau_2^* \geq 0 \qquad \mu^* = 0$$

$$\mu^*(\tau_2^* - \tau_2^*) = 0$$

Suppose we have a solution where the constraint $\tau_2^* - \tau_1^* \geq 0$ is satisfied as a strict inequality, that is $\tau_2^* - \tau_1^* > 0$. This implies $\mu^* = 0$. We then have a solution to:

$$\frac{dg_1(\tau_1^*)}{d\tau_1^*} = 0 \qquad \frac{dg_2(\tau_2^*)}{d\tau_2^*} + F = 0$$

Expanding these two equations and solving, we get:

$$\tau_1^* = \frac{1}{h_1}(r_1 - v_1) \qquad \tau_2^* = \frac{1}{h_2}\left(r_2 - v_2 - \frac{F}{\lambda_2(\tau_2^*)}\right)$$

subject to the constraint $\tau_1^* < \tau_2^*$. We observe that this relationship holds for the case:

$$F < \lambda_2(\tau_2^*)\left[r_2 - v_2 - \frac{h_2}{h_1}(r_1 - v_1)\right]$$

This critical value of F, $F^* = \lambda_2(\tau_2^*)\left[r_2 - v_2 - \frac{h_2}{h_1}(r_1 - v_1)\right]$ follows from $\tau_1^* < \tau_2^*$ as given by (3.18). We now consider what happens to τ_1^* and τ_2^* as F changes. Suppose for some value $\hat{F} < F^*$, then we observe that the constraint holds as a strict inequality then we obtain τ_1^* and τ_2^* as shown above. As $\hat{F}$ increases (but remains $< F^*$), we observe that τ_2^* decreases and τ_1^* remains unchanged. This is the result until $\hat{F} = F^*$, in which case the constraint holds at equality, $\tau_1^* = \tau_2^*$ (note that μ^* may be positive or zero). As $\hat{F}$ increases ($\hat{F} > F^*$), then one of two things happens: (i) μ^* increases, forcing $\tau_1^* = \tau_2^*$, implying that τ_1^* and τ_2^* decrease together but remain equal (given by (3.20)), (ii) τ_2^* decreases, forcing τ_1^* to decrease as well (to satisfy the inequality $\tau_1^* < \tau_2^*$). This forces μ^* to change, which in turn will force $\tau_1^* = \tau_2^*$ to satisfy the first order conditions $\mu^*(\tau_1^* = \tau_2^*) = 0$. The equal switching times case $\tau_1^* = \tau_2^* = \tau^*$ is formulated as an unconstrained optimization problem in a single variable τ. The remainder of this analysis is analogous to that presented in Section 3.3.1.

Appendix D: Obtaining the Values in Figure 3.13

From the proof above in Appendix B, $F < \lambda_2(\tau_2^*)\left[r_2 - v_2 - \frac{h_2}{h_1}(r_1 - v_1)\right]$ implies $\tau_1^* < \tau_2^*$. For fixed cost values greater than this, $\tau_1^* = \tau_2^*$. Rearranging the critical F value by factoring out h_2, we get $F < h_2\lambda_2(\tau_2^*)\left[\frac{r_2-v_2}{h_2} - \frac{r_1-v_1}{h_1}\right]$ which is of the form y=ax depicted in Figure 3.13, where $y = F$, $a = h_2\lambda_2(\tau_2^*)$, and $x = \frac{r_2-v_2}{h_2} - \frac{r_1-v_1}{h_1}$.

Appendix E: Second Order Conditions for the Two Product Model

We derive the second order conditions for the two product with shared repair facility costs model below. Conditions for the case of unequal and equal switching times are presented separately.

Case: $\tau_1^* < \tau_2^*$

$$\begin{aligned}
\text{Total Cost} \quad & TC = F * max(\tau_1, \tau_2) + v_1 Q_1 + v_2 Q_2 + h_1 Q_1 \tau_1 + h_2 Q_2 \tau_2 \\
& \quad - h_1 \int_0^{\tau_1} \Lambda_1(t)dt - h_2 \int_0^{\tau_2} \Lambda_2(t)dt + r_1 \int_{\tau_1}^{T} \lambda_1(t)dt + r_2 \int_{\tau_2}^{T} \lambda_2(t)dt \\
\text{F.O.C.} \quad & \frac{dTC}{d\tau_1} = v_1\lambda_1(\tau_1) + h_1\tau_1\lambda_1(\tau_1) - r_1\lambda_1(\tau_1) = 0 \\
& \tau_1^* = \frac{1}{h_1}(r_1 - v_1) \\
& \frac{dTC}{d\tau_2} = F + v_2\lambda_2(\tau_2) + h_2\tau_2\lambda_2(\tau_2) - r_2\lambda_2(\tau_2) = 0 \\
& \tau_2^* = \frac{1}{h_2}\left(r_2 - v_2 - \frac{F}{\lambda_2(\tau_2^*)}\right) \\
\text{S.O.C.} \quad & \begin{bmatrix} 0 & 0 \\ 0 & v_2\lambda_2'(\tau_2) + h_2\tau_2\lambda_2'(\tau_2) + h_2\lambda_2(\tau_2) - r_2\lambda_2'(\tau_2) \end{bmatrix} \succeq 0 \\
& \lambda_2^2(\tau_2^*) > \frac{F}{h_2}\lambda_2'(\tau_2^*) \qquad \textit{substitute } \tau_2^* \textit{ from FOC}
\end{aligned}$$

Case: $\tau_1^* = \tau_2^*$

$$\text{Total Cost} \quad TC = F\tau + v_1 Q_1 + v_2 Q_2 + h_1 Q_1 \tau + h_2 Q_2 \tau - h_1 \int_0^{\tau} \Lambda_1(t)dt - h_2 \int_0^{\tau} \Lambda_2(t)dt + r_1 \int_{\tau}^{T} \lambda_1(t)dt + r_2 \int_{\tau}^{T} \lambda_2(t)dt$$

$$\text{F.O.C.} \quad \frac{dTC}{d\tau} = F + v_1\lambda_1(\tau) + v_2\lambda_1(\tau) + h_1\tau\lambda_1(\tau) + h_2\tau\lambda_2(\tau) - r_1\lambda_1(\tau) - r_2\lambda_2(\tau) = 0$$

$$\tau^* = \frac{\lambda_1(\tau^*)(r_1 - v_1) + \lambda_2(\tau^*)(r_2 - v_2) - F}{h_1\lambda_1(\tau^*) + h_2\lambda_2(\tau^*)}$$

$$\text{S.O.C.} \quad v_1\lambda_1'(\tau) + h_1\tau\lambda_1'(\tau) + h_1\lambda_1(\tau) - r_1\lambda_1'(\tau) + v_2\lambda_2'(\tau) + h_2\tau\lambda_2'(\tau) + h_2\lambda_2(\tau) - r_2\lambda_2'(\tau) > 0$$