A SYSTEMS APPROACH TO OPTIMIZE A MULTI–ECHELON REPAIRABLE ITEM INVENTORY SYSTEM WITH MULTIPLE CLASSES OF SERVICE

by

Paul Lee Ewing Jr.

A thesis submitted to the Faculty and the Board of Trustees of the Colorado School of Mines in partial fulfillment of the requirements for the degree of Doctor of Philosophy (Mineral Economics).

Golden, Colorado

Date ________________

Signed: ________________________
Paul Lee Ewing Jr.

Approved: ________________________
Dr. Alexandra M. Newman
Assistant Professor
Thesis Advisor

Golden, Colorado

Date ________________

____
Dr. Roderick Eggert
Professor and Director
Division of Economics and Business

ABSTRACT

There are tens of thousands of service organizations in the United States and throughout the world that supply spare parts for repair of critical computer systems. Even with their increased reliability, these systems still fail, generating time–sensitive demand for repair parts without which the corporations who depend on these systems become crippled.

To meet the high level of service that customers require, most service parts logistics systems are organized in echelons. The inventory for these logistics systems is held at a large number of warehouses within this multi-echelon framework. Optimization of a multi-echelon repair parts logistics system can result in the reallocation of inventory that reduces the total cost of inventory in the system while increasing the customer level of service.

We develop the Analytical Stocking Model (ASM) to determine inventory levels to balance system level of service with the cost of holding inventory. The ASM is an extension of an existing multi-echelon model; this extension incorporates an additional class of service by allowing replacement parts to be delivered directly to the customer not only from the lower echelon stocking locations but also from the upper echelon. Our research provides solutions to decisions that can save a computer manufacturer millions of dollars annually while adequately meeting customer service requirements. The models, results, and conclusions from this dissertation can, in general, be applied

to many industries with large–scale service parts logistics systems, e.g., commercial industries such as telecommunications, aerospace, automotive, oil field equipment, computer and office equipment manufacturers, and government organizations such as the military.

Our contributions are not only extending an existing multi–echelon model to accommodate the network structure for a particular service parts logistics system, but also implementing a large–scale optimization problem characterized by highly erratic demands, and verifying our results with a simulation model.

TABLE OF CONTENTS

ABSTRACT . . . . **iii**

LIST OF FIGURES . . . . **viii**

LIST OF TABLES . . . . **x**

ACKNOWLEDGMENTS . . . . **xi**

Chapter 1 INTRODUCTION **1**

Chapter 2 BACKGROUND AND LITERATURE REVIEW . . . **4**

2.1 Multi–Echelon Inventory Systems and Parts Flow 4
2.2 Probability Mass Functions and Stochastic Processes 6
2.2.1 The Poisson Probability Mass Function 7
2.2.2 The Negative Binomial Probability Mass Function 9
2.2.3 The Poisson Process 11
2.3 One–for–One Replacement Policy 14
2.4 Measures of Performance 16
2.5 Palm's Theorem and Queuing Theory 19
2.6 Further Literature 21

Chapter 3 MULTI–ECHELON THEORY AND ANALYTICAL STOCKING MODEL **27**

3.1 Multi–Echelon Theory 27
3.2 Analytical Stocking Model Algorithm 35
3.3 Analytical Stocking Model Algorithm Example 36

Chapter 4 SIMULATION METHODOLOGY **46**

4.1 Simulation Model Description 46
4.2 Statistical Output Analysis 50
4.2.1 Initial Transient Problem 50
4.2.2 Output Analysis for Means 52

Chapter 5 RESULTS **54**

5.1 Test Data 54
5.2 Negative Binomial and Poisson Analytical Stocking Model Comparison 59
5.2.1 Using Simulation to Evaluate the Expected Number of Backorders Determined by the Analytical Stocking Models 64
5.3 Comparison of the Analytical Stocking Model and Supply Chain Operation Policies 70
5.3.1 Distribution Resource Planning System Stock Thresholds and Inventory Allocation 71
5.3.2 Optimal Analytical Stocking Model Policies 72
5.3.3 Individual and System Expected Fill Rates 77
5.3.4 Analytical Stocking Model and Supply Chain Operation Base Stock Levels 81
5.3.5 Ratio of Repair Vendor Base Stock to Remote Stocking Location Base Stock 85
5.4 Emergency Lateral Transshipments 88
5.5 Extended Remote Stocking Location Lead–Times 90

Chapter 6 CONCLUSIONS AND FUTURE RESEARCH **94**

6.1 Summary and Conclusions 94
6.2 Future Research 97

REFERENCES CITED **99**

APPENDIX A-CD-ROM INSTRUCTIONS **101**

APPENDIX B-WEEKLY DEMAND DATA **102**

B.1 Remote Stocking Location Demand data 102
B.2 Repair Vendor Demand data 102

APPENDIX C-CALCULATED MEASURES OF PERFORMANCE **103**

C.1 Expected Backorders 103
C.2 Expected Inventory 103
C.3 Average Customer Waiting Time and Stocking Time 103
C.4 Expected Fill Rates 103

APPENDIX D-BASE STOCK LEVELS **104**

D.1 Repair Vendor Base Stock 104

D.2 Remote Stocking Location Base Stock 104

LIST OF FIGURES

2.1 Two–echelon system with two classes of service, one RV, and N RSLs 5

2.2 The Poisson probability mass function for different mean values . . . 8

2.3 The Negative Binomial probability mass functions for a constant mean value of eight 10

2.4 The Negative Binomial probability mass functions for a constant variance to mean ratio of 1.5 12

2.5 Fill Rate –vs.– Backorders 17

2.6 Five RSLs represented as members of two risk pooling groups 24

3.1 Part #1 vs. part #2 sub–problem convexity with non–convex points removed 43

4.1 Simulation flow–chart 47

5.1 Pareto analysis of SCO inventory over a six week period 56

5.2 Forecasted demand at the RSLs for each high–DMV part in the system 57

5.3 Forecasted demand at the RV for each high–DMV part in the system 58

5.4 Optimal ASM exchange curves 61

5.5 Comparison of Poisson and Negative Binomial ASMs 63

5.6 Example output from the Arena Output Analyzer 66

5.7 Percent Error of the Poisson and Negative Binomial ASM 69

5.8 Optimal ASM allocation vs. current SCO allocation 73

5.9 The average time a part spends in inventory versus the average time a customer waits for a part (linear / log scale) 77

5.10 Individual part fill rate comparisons between ASM and SCO policies . 79

5.11 Average total fill rate comparisons between ASM and SCO policies . 80

5.12 Base–stock levels at the RV for the ASM and SCO policies 82

5.13 Base–stock levels at the RSLs for the ASM and SCO policies 84

5.14 Ratio of the total RV demand versus the RSL demand for the ASM and SCO policies 86

5.15 Ratio of RV to RSL aggregate demand 87

5.16 Comparison of ASM using ELT and standard ASM with no ELT . . . 89

5.17 Comparison of ASM exchange curve for RSL lead-times, L_i, of two and nine days 91

5.18 Comparison of ASM using ELT and standard ASM with extended RSL replenishment lead-time of nine days 93

LIST OF TABLES

3.1 Input data for ASM algorithm example 36

3.2 Expected backorder level for part #1 at the RSLs, with RV stock level equal to zero $(* \equiv \beta < 0.0001)$ 37

3.3 Expected backorders for optimal allocation of stock at the RSLs for part #1 given an RV stock of zero 39

3.4 Expected backorders for optimal allocation of stock at the RSLs for part #1 given an RV stock of zero through eight 40

3.5 Expected backorders optimally allocated for part #1 41

3.6 Expected backorders optimally allocated for part #2 42

3.7 Total expected backorders optimally allocated across parts 44

5.1 System expected backorders and total costs for the points shown on Figure 5.8 74

5.2 Percent decrease in expected backorders and total cost between the SCO policy and points **b** and **c** on the optimal ASM curve shown in Figure 5.8 75

5.3 Percent reduction in total RSL base–stock and total cost from standard ASM model and ASM model incorporating ELT 90

ACKNOWLEDGMENTS

I would like to take this opportunity to thank the many people who helped me with this work. First, I would like to thank my dissertation committee for their advice, assistance, and time: Dr. R. E. D. Woolsey, Dr. Jeff Grobman, Dr. Robert Frost, Dr. Irina Khindanova, and especially Dr. Newman who spent endless hours providing guidance for this research and editing this dissertation, and helped me to become a better Operations Research Analyst.

From Sun Microsystems Enterprise Services, I would like to thank the following: Warren White, past director of the Supply Chain Operations for Enterprise Services; Mike White, the manager of the Material Asset Management group, who allowed me to learn his group's operations, and to take the time of the asset planners who provided assistance with the data analysis for this project; Chris Payne and Richard Tyrrell who provided invaluable information concerning service support supply chain operations and assistance with data collection; Blaine Taylor for his insightful and constructive criticism of this dissertation, along with Mark Connolly and the rest of the Global Business Process Reengineering group who provided support and expertise, and most importantly welcomed me as a part of their team.

Finally, special thanks to my wife, Chris, and to my children, Morgan and Joseph, who supported me throughout this endeavor.

Chapter 1

INTRODUCTION

There are tens of thousands of service organizations in the United States and throughout the world that supply spare parts for system repairs. These service organizations provide both low cost consumable and high value repairable spare parts for a multitude of industries, e.g., telecommunications, aerospace, automotive, computer and office equipment manufacturing, and for the military. Though there may be fewer repairable than consumable parts in a typical service parts logistics system, repairable parts usually account for the majority of the costs associated with holding inventory. We focus our efforts on a particular company; nevertheless the models, results, and conclusions from this dissertation can, in general, be applied to many other industries with large scale service parts logistics systems.

The improvement of the operation of a nation–wide enterprise service parts logistics system for Sun Microsystems Inc. motivates our research. Sun Microsystems generates the most UNIX server revenue world wide – about 35% of the market [Sun Microsystems Inc., 2001]. Sun has been able to maintain leadership in the high–end server market by providing innovative, scalable, and reliable products. Sun's continuing success stems not only from its product line, but also from its ability to provide outstanding after–sales customer service.

Specifically, Sun must ensure that the computer systems that provide mission critical services to the corporations who purchase such systems remain functional 24 hours a day, 7 days a week. To provide repair parts to customers in North America, Sun maintains the Sun Virtual Logistics Network (SVLN). New product introductions (NPI) will be more reliable which should result in reduced need for repair parts inventory. However, this increased reliability comes at a greatly increased unit price per part, and as a result, increased inventory holding costs. Future projections for purchases of new inventory for the SVLN suggest that if Sun continues to follow its current inventory policy, SVLN inventory costs could double within the next five years. Thus, even minor reductions in repair parts inventory can save Sun hundreds of millions of dollars.

Sun currently maintains approximately 4300 repair parts, which corresponds to several hundred million dollars of SVLN inventory. Stocking policies developed for the SVLN suggest where and how much of these critical repair parts Sun should hold in inventory. Designing these policies is complex because Sun must balance the costs of holding spare parts with the speed with which customers require these parts.

Storing many and different types of repair parts close to all of Sun's customers guarantees rapid response. However, Sun cannot economically justify such a policy due to the cost of maintaining warehouses in close proximity to all of its customers and of holding a generous number of repairable parts in inventory. In the latter case, holding parts "in storage" causes Sun to incur an opportunity cost, primarily a function of the company's inability to derive revenue–either directly from the sale of the repair part or indirectly through the part's inherent value–while the part sits in a warehouse unused (and potentially becomes obsolete).

We develop the Analytical Stocking Model (ASM) to determine inventory levels to balance system level of service with the cost of holding inventory. The ASM is an extension of an existing multi-echelon model that incorporates two classes of service, i.e., replacement parts may be delivered directly to the customer from either the upper or lower echelon stocking locations. Our research provides solutions to decisions that can save Sun millions of dollars annually while adequately meeting customer service requirements. Our contributions are not only the development of the Analytical Stocking Model to accommodate Sun's network structure, but also implementing a large–scale optimization problem characterized by highly erratic demands, and verifying our results with a simulation model. To our knowledge, such a large–scale, real–world, comprehensive repair inventory analysis does not exist in the literature.

The dissertation is organized as follows: Chapter 2 provides background and literature review of service parts logistics systems. Chapter 3 details the repair part, multi–echelon theory on which our models are based, and Chapter 4 discusses the methodology for the simulation used to validate the Analytical Stocking Model. Chapter 5 contains results for the Analytical Stocking Models and simulation. Conclusions and future research appear in Chapter 6.

Chapter 2

BACKGROUND AND LITERATURE REVIEW

2.1 Multi–Echelon Inventory Systems and Parts Flow

Sun has implemented a *multi–echelon inventory system* to balance the benefits and costs of maintaining inventory in close proximity to its customers. A multi–echelon inventory system has two or more levels of warehouses, e.g., in a two–echelon system, the lower–echelon may contain retail stores that directly service the customer, while the upper–echelon, or distribution center, resupplies the lower echelon. Multi–echelon inventory systems are well suited for situations in which rapid response for spare parts is required, e.g., one hour, and customers are distributed over a large geographical area. We refer to the lower–echelon warehouses, i.e., the responsive warehouses located close to the customer, as Remote Stocking Locations (RSLs). The RSLs hold inventory for emergency repairs only. The Repair Vendor (RV), or the upper echelon warehouse, repairs parts returned from the RSLs and holds inventory for scheduled maintenance or for replenishment of the RSL's inventory.

Multi–echelon systems are characterized as having many RSLs and very few RVs. Those who manage multi–echelon systems must determine the amount of total inventory and the optimal allocation of inventory between the RSLs and the RV. A multi–echelon system may reduce total inventory costs by as much as 50% when

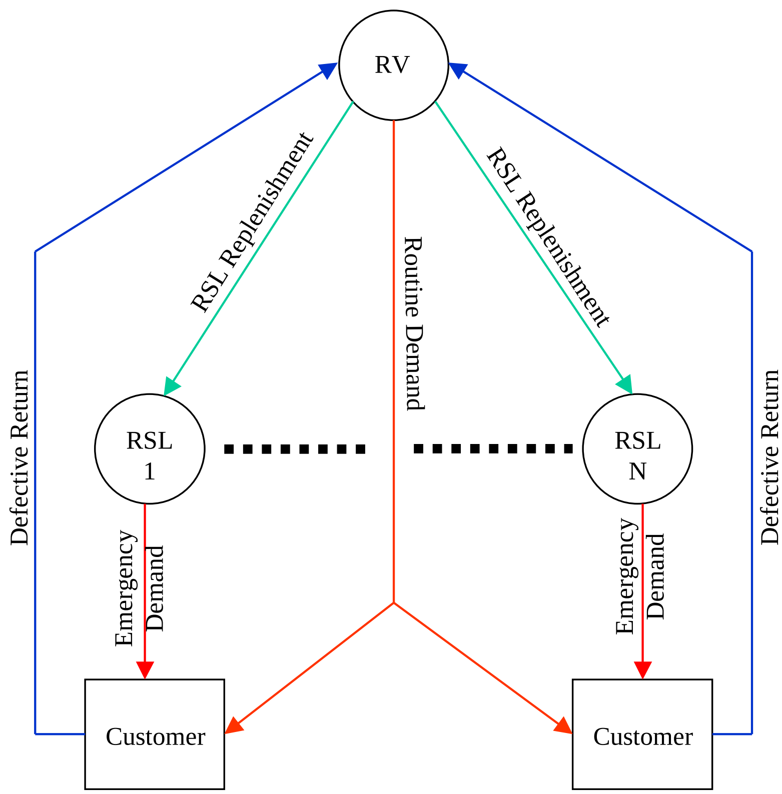

FIG. 2.1. Two–echelon system with two classes of service, one RV, and N RSLs

compared with a single–echelon inventory system [Muckstadt and Thomas, 1980].

Figure 2.1 shows the multi–echelon structure and the parts flow corresponding to Sun's current supply chain operation. In most cases, the RV provides part replenishment for the RSL's emergency, or same–day, demand. However, we refer to the second class of service in which the RV also provides direct customer support for non-time sensitive cases, e.g., scheduled maintenance or part upgrades, as routine demand.

A customer returns a defective part to the RV. If possible, the defective part is repaired at the RV; otherwise it is scrapped. The transportation time from the customer to the RV plus the average time required to repair the part at the RV is referred to as the RV lead–time, e.g., 10 days transportation time + 12 days repair time = 22 days RV lead–time.

When a part fails, the customer also places an order for a new part, thus creating a demand for the part at his local RSL. If the part is in stock at the RSL when the order is received, the customer is immediately issued the new part from the RSL inventory. However, if no parts are available at the RSL, then demand for a replacement part from the RSL inventory exists until serviced, i.e., the part is "backordered". If the RV has the part in its inventory, the RV immediately ships the part to the customer (bypassing the RSL). The time it takes for the RV to ship the replenishment part to the RSL, or to the customer in the case of an RSL stock–out, is referred to as the RSL replenishment lead–time.

We assume that customers arrive at the RSL and the RV according to a Poisson process. We discuss this *counting process* and relevant probability distributions next.

2.2 Probability Mass Functions and Stochastic Processes

This section reviews relevant probability distributions and stochastic processes used throughout this dissertation. We begin by defining the nomenclature.

A *random variable*, X, represents a currently unknown quantity whose value will become known at some future time. For example, X might correspond to demand for a part over the next week. Let us consider a discrete random variable X, with a range of values, S, of a set of integers over some given time interval. The probability that

X will assume a certain value is expressed by the probability mass function (pmf) defined in the next section.

2.2.1 The Poisson Probability Mass Function

The pmf, $g(x)$, is defined for the set of all non–negative integers, S, and is given as:

$$g(x) = Pr\{X = x\}, \quad x \in S. \tag{2.1}$$

The pmf for a Poisson distribution with a single non–negative parameter, ϕ, is:

$$g(x) = \frac{(\phi)^x}{x!} \exp(-\phi), \quad x \geq 0. \tag{2.2}$$

A model that uses a Poisson distribution as the underlying distribution is referred to as a single parameter approximation, because the distribution is completely characterized by its mean, i.e., the mean, $E[X]$, and variance, $Var[X]$, are equal to ϕ.

Figure 2.2 graphs the Poisson pmf for different values of ϕ. Note that $g(x)$ is unimodal, i.e., a global maximum exists, approximately at the value of the mean, ϕ.

For $\phi < 1$ the Poisson pmf is strictly monotonic, but for $\phi \geq 1$ the Poisson pmf is non–monotonic. Monotonicity (or lack thereof) of the Poisson pmf may affect the convexity properties of the corresponding function in which the distribution is used. Possible applications for the Poisson distribution are: (i) the number of items demanded from inventory, (ii) the number of items in a batch of a random size, (iii) and the number of events that occur in an interval of time when the events are

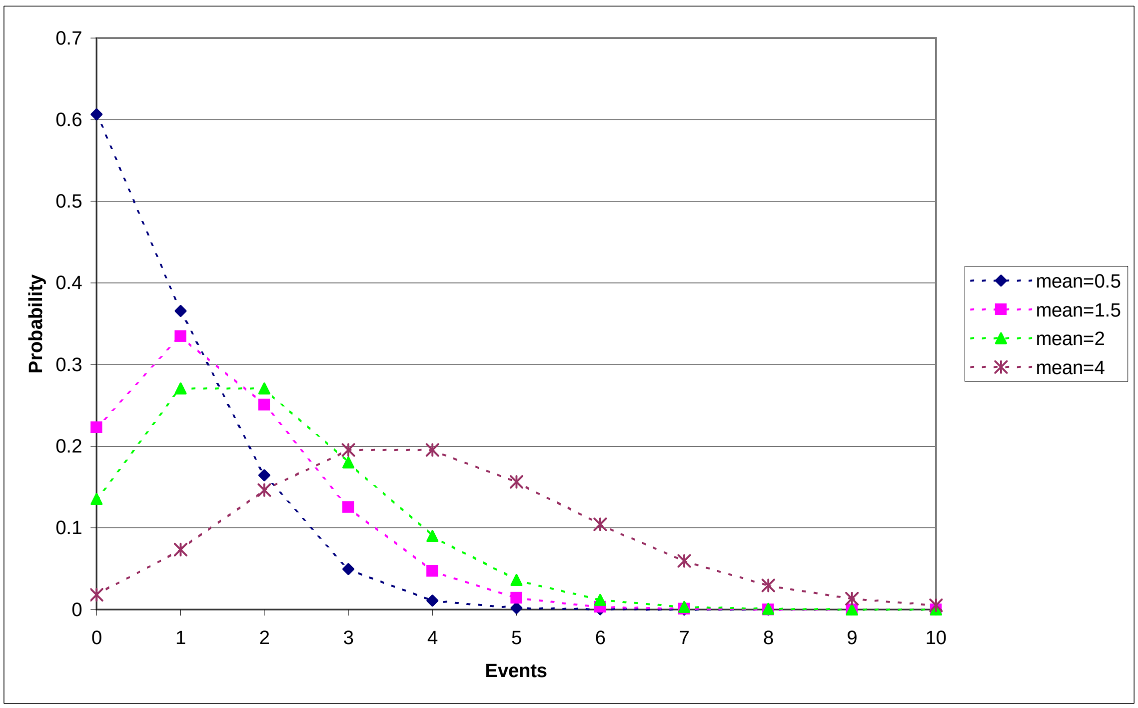

FIG. 2.2. The Poisson probability mass function for different mean values

occurring at a constant rate.

2.2.2 The Negative Binomial Probability Mass Function

The other pmf of interest is a two parameter distribution called the Negative Binomial Distribution. The two parameters are n and p, representing the number of events and probability of the occurrence of an event, respectively, where $0 < p < 1$ and $n > 0$. For example, suppose n is a positive integer and X is the amount of time until n identical machines fail with probability p. The pmf is:

$$g(x) = \binom{n+x-1}{x} p^x (1-p)^n \quad x \in S. \tag{2.3}$$

The mean and variance, $E[X]$ and $Var[X]$, of the time until n identical machines fail with probability p are:

$$E[X] = \frac{np}{1-p} \tag{2.4}$$

$$Var[X] = \frac{np}{(1-p)^2}. \tag{2.5}$$

Figure 2.3 shows several Negative Binomial probability mass functions for a constant mean of eight. The p values are adjusted for given values of n to keep a constant mean. Note that $g(x)$ is unimodal and, like the Poisson pmf, the Negative Binomial pmf is not monotonic for all values of n and p.

We can find the values of n and p in terms of the mean, $E[X]$, and the variance

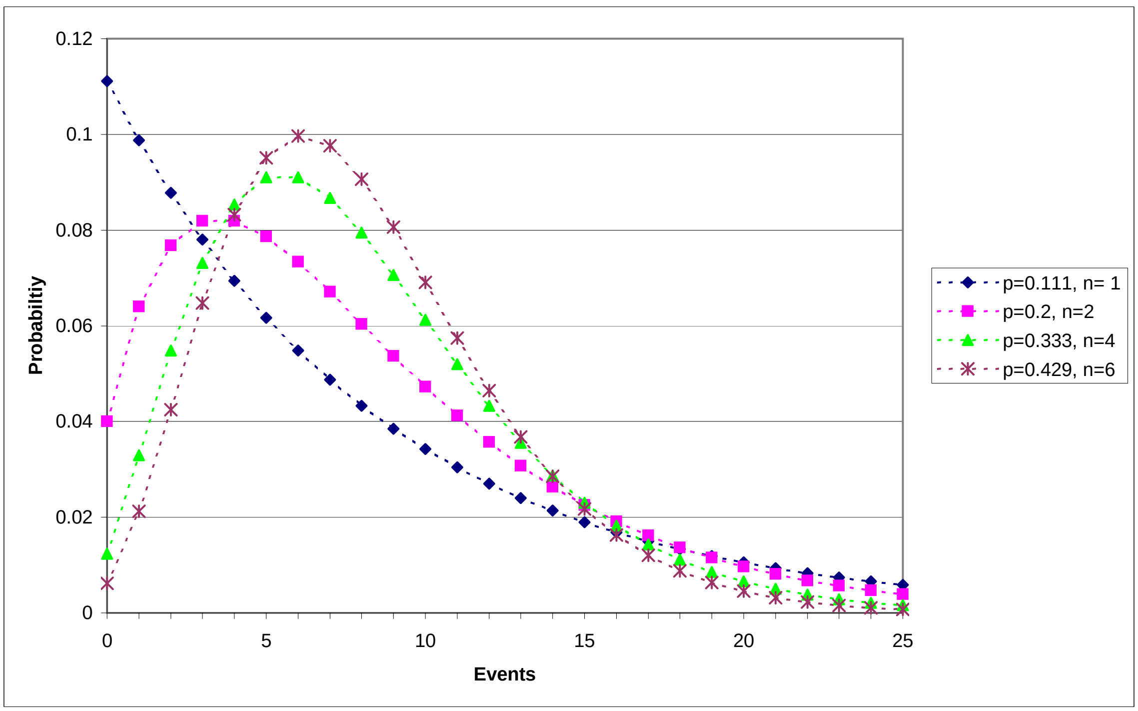

FIG. 2.3. The Negative Binomial probability mass functions for a constant mean value of eight

to mean ratio, V. Solving (2.4) and (2.5) for V, we find:

$$V = \frac{1}{(1-p)}.$$

It follows that:

$$p = \frac{(V-1)}{V}. \tag{2.6}$$

We substitute (2.6) into (2.4) and solve for n:

$$n = \frac{E[X]}{(V-1)} \tag{2.7}$$

We use (2.6) and (2.7) to solve for p and n, respectively, for use in the Analytical Stocking Model. Figure 2.4 shows several Negative Binomial probability mass functions for constant values of the mean to variance ratio, V. Note, as n increases, the mean, $E[X]$, and variance, $Var[X]$, increase.

Possible applications for the Negative binomial distribution are: (i) the number of items demanded from an inventory, (ii) the number of items in a batch of a random size, (iii) and the number of good items inspected before encountering the n^{th} defective item.

2.2.3 The Poisson Process

A stochastic process, $X(t)$, $t \in T$ is a collection of random variables, where t represents time and *X(t)* the state of the process at time *t*. In other words, this process is a set of random variables that describes the evolution through time of some

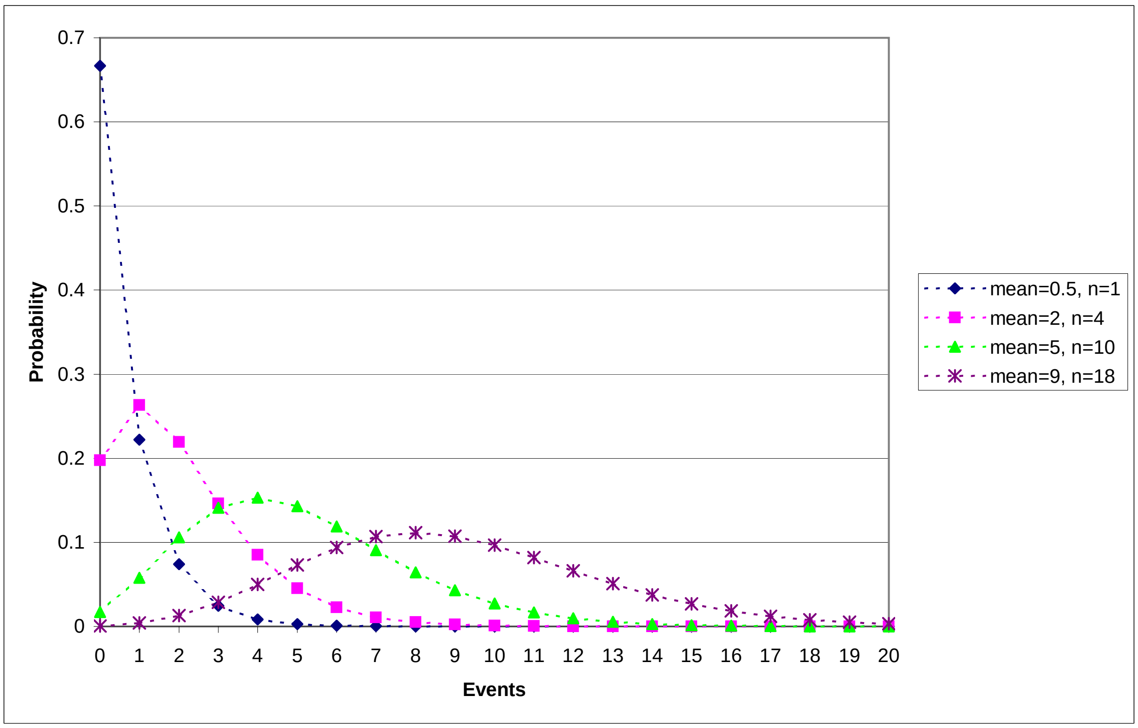

FIG. 2.4. The Negative Binomial probability mass functions for a constant variance to mean ratio of 1.5

"physical" process [Ross, 2000].

A stochastic process, $N(t)$, $t \geq 0$, where $N(t)$ represents the total number of "events" that have occurred up to time t, is said to be a Poisson process with rate, λ, where is $\lambda > 0$, if:

- $N(0) = 0$.
- The process has independent events.
- The number of events in any interval of length t follows a Poisson distribution with mean λ. That is, for s and $t \geq 0$

$$Pr\{N(t+s) - N(s) = x\} = \frac{(\lambda t)^x}{x!} \exp(-\lambda t), \quad x = 0, 1, \cdots. \tag{2.8}$$

A stochastic process has *stationary* increments if the number of events in the interval $[t_1 + s, t_2 + s]$ has the same distribution as the number of events in the interval $[t_1, t_2]$ $\forall$ $t_1 < t_2$, and $s > 0$. It follows from (2.8) that a Poisson process has stationary increments and that the expected number of events that occur by time t is equal to the mean of the Poisson process, λt.

The multi–echelon models we discuss assume that demand at the RSLs follows a Poisson process with a known average customer arrival rate, λ_i, at $RSL_i, \forall\ i = 1, 2, \cdots, N$. The total demand at the RV is Poisson because the sum of Poisson processes is a Poisson process.

2.3 One–for–One Replacement Policy

In Section 2.1, we described an inventory policy where an item is demanded from an RSL, while another is immediately requested from the RV for replenishment. This policy is referred to as a "one–for–one replacement" policy or an $(s-1, s)$ policy where s is the order–up–to point, or the *base–stock* level. The following random variables describe the system:

- OH = on–hand inventory
- BO = number of backorders
- IN = net inventory = OH - BO
- x = number of outstanding orders
- IP = inventory position, defined as

$$IP = IN + x. \tag{2.9}$$

The number of outstanding orders, x, or units in the resupply "pipeline", is the number of parts in transit from the customer back to the RV for repair, in repair at the RV, or enroute to the RSL from the RV.

The objective of an $(s-1, s)$ policy is to maintain an inventory position of s; in other words, we place an order when the net inventory falls to $s-1$ or below. If the initial inventory position of the system is less than s, orders are placed until the inventory position equals s. If the initial inventory position is greater than s, nothing is ordered until customer demand reduces the inventory position to less than s.

When net inventory becomes negative, a backorder exists. We define a backorder as:

- $x - s$ is the quantity backordered if $s - x < 0$
- $s - x$ is the amount of on–hand inventory if $s - x \geq 0$

where $0 \leq x < \infty$.

Most failure processes assume that customer demand for spare parts is a Poisson process. Because each outstanding order corresponds to a unit of demand, the inventory on order has a Poisson distribution with mean λL, where λ is the mean customer demand rate and L is replenishment lead–time.

All the models we review use $(s - 1, s)$ policies and assume fully backlogged demand. Our other assumptions are: (i) Customers are serviced according to first–come–first–served rules, (ii) inventory lead–times and RV repair–times are independent and identically distributed with a mean value L, and these lead–times can follow any general distribution function, and (iii) no repairs occur at the lower echelon, i.e., a defective part is returned to the upper echelon for repair.

Most repairable parts inventory systems are characterized by low usage rates and high holding costs, e.g., demand rates at an RSL of five per month, and high cost items, e.g., thousands of dollars per unit. The $(s - 1, s)$ policy is said to be the *optimal* stocking policy for single–echelon repairable parts inventory systems because it minimizes total inventory holding costs. However, this presumption has not been rigorously proven for multi-echelon systems [Muckstadt and Thomas, 1980, Sherbrooke, 1968, Graves, 1985].

The optimal stocking level, s, is determined by minimizing a system measure of performance subject to a constraint on the total inventory cost. Before we introduce the optimization problem, we must first define the required measures of performance.

2.4 Measures of Performance

In this section, we discuss item measures of performance and how they relate to system measures of performance. The most extensively used item measures of performance are the "fill rate" and "backorder" criteria [Nahmias, 1981]. We define fill rate as the percentage of customer demand that is met at the time the demand is placed. Fill rate measures the number of shortages but does not consider the duration of those shortages. By contrast, the backorder criterion considers the *duration* of shortages. The expected number of backorders represents the amount of unfilled demand that exists over the probability distribution of demand.

We use total expected number of backorders as a system measure of performance to maximize the system availability of service parts. Sherbrooke [1992] shows that minimizing the sum of the expected number of backorders is equivalent to maximizing *supply availability* for a single RSL. We define supply availability as the percentage of the complement of the number computers down due to the lack of spare parts. To use supply availability as a performance measure, the number the number and type of computer servers assigned to each individual RSL, i.e., the *install base*, must be known. Use of the supply availability criteria may not be practical because complete

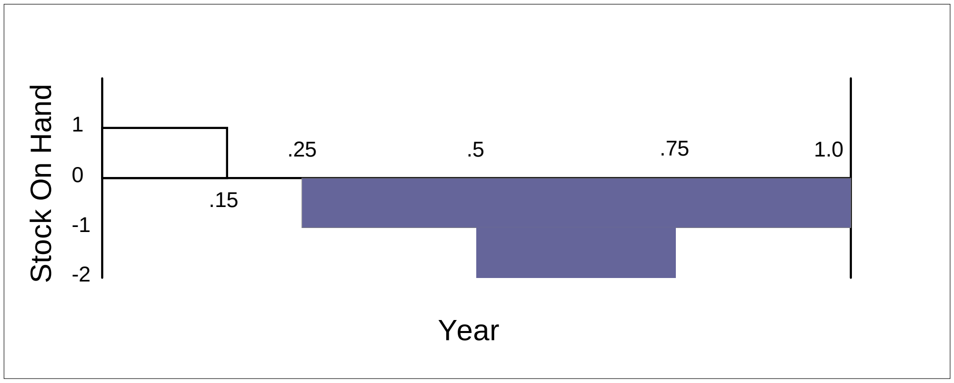

FIG. 2.5. Fill Rate –vs.– Backorders

information about the install base may not be available. Also, the system level availability function is non-convex for multiple RSLs, making the optimization problem difficult to solve.

For the calculations in this paper we use the *expected* number of backorders. However, because the *expected value* of a random variable with a probability distribution is difficult to illustrate, we use the *average* number of backorders in the following example (Figure 2.5) to illustrate the difference between fill rate and average number of backorders [Sherbrooke, 1992]. As the time horizon increases, the average number of backorders becomes a better approximation for the expected number of backorders. The average number of backorders represents the amount of unfilled demand that exists over a predetermined time horizon. The expected number of backorders in this case is said to be a *steady–state* parameter, i.e., for repeated samples, the distributions of the long–run means of the number of backorders are approximately the same.

Suppose, in our example, that the time horizon is one year. An order arrives at

.15 of a year, with one part in stock. Two more orders for the same part arrive at .25 and .5 years before the replenishment of one part arrives from the RV at .75 years. Another replenishment arrives at 1 year, where we stop our measurement. The fill rate for this one–year period is 33%, i.e., of the three demands for the period, only one is met. The average number of backorders is computed as the area of the shaded portion in Figure 2.5, divided by the length of time, i.e., one backorder.

One drawback of the fill rate criterion is that it ignores the length of delay experienced by a customer due to a stock–out at the RSL. Therefore, a consequence of using the fill rate criterion as a system level performance measure is that in order to avoid a shortage of the parts that have not yet suffered a backorder, the backordered part may never be stocked. In addition, Sherbrooke [1968] indicate that maximizing item fill rate for a multi–echelon inventory system causes inventory to migrate to the RSLs. This increase in RSL inventory increases inventory costs, but does not necessarily increase system–level parts availability.

The backorder criterion is not without its faults. It is possible to have high levels of supply availability by minimizing the overall system expected number of backorders. However, the expected fill rates for certain parts may be low for some RSLs [Sherbrooke, 1992], e.g., for the case in which a part has a high cost and very low demand rate at an RSL. Another drawback of the backorder criterion is that this measure of performance is not very intuitive. Specifically, an inventory manager will not necessarily know what an expected number of backorders of “one” means, but he knows how to interpret a fill rate of 33%.

Because our objective is to determine optimal stock levels subject to a given measure of performance, it is important to investigate the characteristics of the expected backorder and fill–rate functions. The expected backorder function, for each part, is a monotonically decreasing and convex function of the base–stock, i.e., as the base–stock increases, the number of backorders decreases. However, the fill–rate function, for each part, is not convex for all values of the base–stock level. Only when the mean, λL, is less than or equal to one is convexity guaranteed for the fill–rate function. Therefore, guaranteeing optimal results using the fill rate criterion would be very difficult, if not impossible, over large values of base–stock and part–location pairs.

2.5 Palm's Theorem and Queuing Theory

Queuing theory is the mathematics of waiting lines, or queues. For example, queuing theory addresses the following questions, among others: (i) What is the expected number of customers present in the queue? (ii) What is the expected time that a customer spends in the queue? and (iii) What is the probability distribution of the number of customers present in the queue?

Scarf [1958] first described the continuous $(s-1, s)$ policy using an $M/G/\infty$ queuing system, i.e., an *infinite channel queuing system* where a part failure is equivalent to a customer entering service. A queuing system is characterized by its arrival process, service process, and by the number of servers. The "M" stands for *Markovian* arrival process; a process that is Markovian is "memoryless", or the arrival process is said to be exponentially distributed. A random variable, which represents the inter–arrival times between demands placed at a stocking location, must be independent and identically distributed. The "G" denotes independent and identically distributed service times that follow some general distribution. The last characteristic assumes an infinite number of servers at the location where the damaged part is returned for immediate repair.

Palm's theorem [Palm, 1938] allows us to determine the expected number of backorders in a system using the probability distribution of the RSL demands and the mean of the repair time distribution. It assumes that the repair times for all failed parts are independent and identically distributed, which is realistic for our application, in which the demand rate for parts is low.

The replenishment order lead–times in the inventory system correspond to the customer service times in the queuing system. The number of customers waiting to be served corresponds to the expected number of backorders. Palm's theorem states that if customers arrive according to a stationary Poisson process and service times are independent and identically distributed random variables with finite mean, then the steady state probability distribution of the number of busy servers per unit time is Poisson, independent of the form of the service distribution. This theorem implies that if customer demands are generated according to a stationary Poisson process, the steady state number of outstanding orders, i.e., parts in the repair pipeline, is a random variable having a Poisson distribution, regardless of the distribution of lead–times, as long as lead–times are independent. This theorem forms the basis of modern $(s-1, s)$ theory for multi–echelon recoverable part inventory systems.

2.6 Further Literature

Scarf was the first to identify the analogy between the $(s-1, s)$ model and the theory of infinite server queues. Applying Palm's theorem, Feeney and Sherbrooke [1966] extend Scarf's results for multi–echelon systems and show that if the demand process is Poisson then the distribution of outstanding orders is also Poisson. Using the results from Feeney and Sherbrooke, the RAND Corporation developed the Multi–Echelon Technique for Recoverable Item Control (METRIC) model for the U.S. Air Force [Sherbrooke, 1968]. The METRIC model is the first successful implementation of a multi–echelon repairable parts system, and determines optimal inventory stocking levels for the RV and RSLs by minimizing the expected number of backorders subject to a system budget constraint. The METRIC model considers the case in which a customer demand arrives according to a Poisson process in batches of demands rather than as single units of demand. However, in order for Sherbrooke to apply Palm's theorem to the METRIC model, he ignores possible RSL lead–time correlation. Correlation between RSL lead–times may exist because a backordered part at several RSLs must be replenished from the same RV on a first–come–first–served–basis. Therefore, the METRIC model is considered an approximation.

Principal METRIC model assumptions are:

- The estimated demand data are stationary and all parts are repairable.
- A certain percentage of repairable parts are repaired with probability r at the RSLs while the RV repairs the balance of the parts.

- There is no condemnation of parts returned to the RV, i.e., all parts can be repaired. This is a reasonable assumption for most recoverable parts inventory systems when typical scrap rates are less than five percent.

- There is an infinite number of servers at the RV. This assumption is not as restrictive as it first appears. Waiting for repair at an RV rarely occurs because the RVs have high repair capacities relative to the low demand requirements for repairable parts. Furthermore, the inaccuracies in the part demand data negate any advantage in modeling the RV repair process exactly.

Even though the METRIC model is the first successful optimization implementation for the multi–echelon repairable parts problem, it does have its shortcomings. The METRIC model tends to underestimate the base–stock levels at the RSLs because the model relies on a single–parameter distribution to approximate the expected number of backorders at the RSLs. Specifically, the METRIC model approximation underestimates required base–stock levels in 11.5% of the cases tested by Graves [1985]. Several extensions to the METRIC model attempt to address this underestimation.

Simon [1971] develops an exact solution approach for solving the $(s-1, s)$ inventory policy, multi–echelon, repairable parts problem. Simon assumes a constant transportation time from the RV to an RSL, that demand at the RSLs follows a Poisson distribution, and that the repair times are deterministic. Shanker [1981] extends Simon's technique to include compound Poisson demands. Simon and Shanker's techniques are computationally intractable for large–scale practical problems, i.e., problems that are characterized by a large number of stocking locations and parts, and highly variable demands. However, their techniques prove accurate for small test problems and the results from these exact solution approaches are often compared to those obtained with approximate models, e.g., the METRIC model.

Graves [1985] and Sherbrooke [1986] improve the METRIC approximation by fitting a Negative Binomial distribution to the RSL expected fill rate and backorder functions, respectively. Graves compares his approximation to exact solutions developed by Simon and Shanker and to METRIC model results for 227 test problems in which the variance is greater than the mean number of (demand) arrivals. Graves' Negative Binomial approximation gives better results than the METRIC model solutions in all but two cases. The Negative Binomial model is accurate in 99.1% of the cases tested, but it tends to overestimate the expected number of backorders.

Cohen et al. [1999] extend a version of METRIC to allow delivery of parts from the upper–echelon directly to the customer. Cohen's results are sub–optimal because he uses a greedy heuristic to maximize the fill rate measure of performance, which, for each part, is non-convex for small values of parts held in inventory, and his model does not use the negative binomial distribution to approximate the expected number of backorders at the RSLs.

Extensions of the METRIC model allow for emergency lateral transshipment (ELT) between the RSLs. If a stockout precludes a customer from receiving an emergency repair part from his assigned RSL, ELT is used to ship a repair part from another RSL within the risk *pooling group.* Figure 2.6 shows two risk pooling groups with ELT. ELTs should not be confused with *inventory fair–sharing*, i.e., the situation

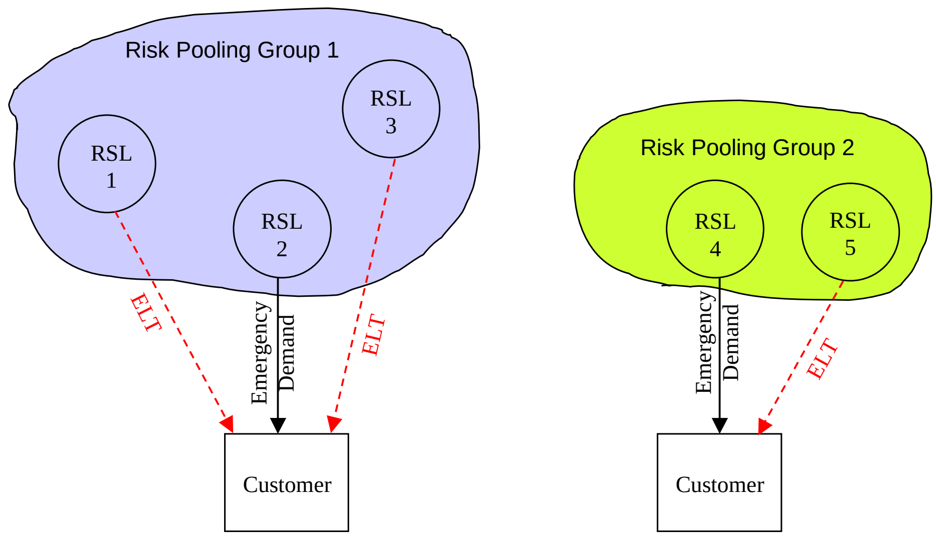

FIG. 2.6. Five RSLs represented as members of two risk pooling groups

in which stock is routinely shifted among RSLs in order to level inventory.

Lee [1987] develops an Emergency Lateral Transshipment model that fits the Negative Binomial distribution to the RV expected backorder function. Lee's ELT model uses a random *stock sourcing rule* to determine how and when RSLs are selected to transship stock to other RSLs within their risk pooling group. His ELT model assumes that all RSLs within the same risk pooling group have identical demand rates and that the demand arrives according to a Poisson process. The transportation time from the RV to each RSL is constant and identical for all RSLs. The remainder of Lee's assumptions is the same as those in the METRIC model.

Axsäter [1990] extends the METRIC model to include ELTs and also uses the random stock sourcing rule. Unlike Lee, Axsäter assumes that the RSLs within a risk pooling group and the transshipment lead–times are not identical. Axsäter finds that for fill rates greater than 70%, his model outperforms Lee's. Both Axsäter's and Lee's models offer analytical solutions, but the model's assumptions appear too restrictive for practical use. Sherbrooke [1992] proposes a different ELT approach combining the METRIC model and linear regression. When compared to the expected number of backorders found with the METRIC model, Sherbrooke estimates that expected backorder reductions of 30% to 50% may be possible by considering inventory policies using ELT.

In contrast to the continuous review, $(s-1, s)$ inventory policy, IBM develops a periodic review model based on the (s, S) inventory policy [Cohen et al., 1990]. When the level of on–hand inventory is less than or equal to s, then an order for the difference between the on–hand inventory, μ, and S is placed. If μ is also the initial inventory in any period, then the (s, S) policy is:

- if $\mu \leq s$, order $S - \mu$

else, do not order.

In general, the IBM model obtains values for s and S by minimizing expected inventory costs subject to product service constraints for each warehouse location within each echelon. Optimal values for s and S are calculated for each echelon and are aggregated to obtain a solution for all echelons combined. The global solution is sub–optimal because the warehouses in each echelon do not function independently. IBM does not use the simpler $(s-1, s)$ inventory policy because "while appropriate for low–demand items, this $(s-1, s)$ policy does not provide adequate cost and service performance for the wide range of demand rates present in IBM's parts environment"

[Cohen et al., 1990].

Though the METRIC model and its extensions have previously been used for systems characterized by high cost and low demand parts, the ASM developed in the next chapter is evaluated for high cost and erratic demand parts.

Chapter 3

MULTI–ECHELON THEORY AND ANALYTICAL STOCKING MODEL

In this chapter, we discuss the equations which describe multi-echelon inventory theory for a specific case, where the backorder functions for the RSLs are approximated by Poisson and Negative Binomial distributions. We also illustrate the ASM algorithm with an example.

3.1 Multi–Echelon Theory

In order to determine the optimal stocking levels for the RV and RSLs, we develop an Analytical Stocking Model (ASM) by extending Sherbrooke's METRIC model for a two–echelon, recoverable parts inventory system to include an additional class of service. We begin by defining the notation:

Parameters

- N: the number of RSLs
- J: the number of parts
- L_{ij}: the transportation time for part j between the RV and RSL_i (days)
- L_{0j}: the replenishment lead–time for the RV, i.e., repair time at the RV plus

defective part j return time from an RSL to the RV (days)(invariant across RSLs)

- L_{RV_j}: the average transportation time for part j between the RV and a customer (days) (invariant across customers)

- λ_{ij} : the demand rate for replacement part j at RSL_i (parts/day)

- $\lambda_{0j} = \sum_{i=1}^{N} \lambda_{ij}$: the total emergency demand rate at the RV for part j (parts/day)

- λ_{RV_j}: the total routine demand rate at the RV for part j (parts/day)

- Φ_{0j}: the average number of part type j in the repair pipeline (# parts)

Decision Variables

- s_{ij}: the base–stock level at RSL_i for part j (# parts)

- s_{0j}: the base–stock level at the RV for part j (# parts)

Random Variables

- I_{ij}: the expected on–hand inventory for part j at RSL_i in steady state (# parts)

- I_{0j}: the expected on–hand inventory for part j at the RV in steady state (# parts)

- β_{ij}: the expected backorder level for part j at RSL_i in steady state (# parts)

- β_{0j}: the expected backorder level for part j at the RV in steady state (# parts)

- W_{ij}: the time delay due to stock–outs for part j at RSL_i in steady state (days)
- W_{0j}: the time delay due to stock–outs for part j at the RV in steady state (days)
- α_{ij}: the fill rate for part j at RSL_i in steady state (%)
- α_{0j}: the fill rate for part j at the RV in steady state (%)
- Θ_{ij}: the expected number of outstanding orders for part j at RSL_i in steady state (# parts)

We first characterize the inventory system using the ASM by determining the expected backorder level, inventory, and average waiting time due to stock–outs for each part. For a single echelon model, an $M/G/\infty$ system describes the steady state number of outstanding orders at a warehouse by assuming a Poisson process with mean λL for demands placed on a warehouse, where λ is the average arrival rate and L is the average service time.

For the two–echelon model, we relax the assumption that L_{0j} is constant. The replenishment lead–times for the RV are assumed to be independent and identically distributed with mean L_{0j}, where L_{0j} can assume any distribution. Using the expected waiting time at the RV and the average RSL lead–time, L_{ij}, we calculate the expected number of backorders at the RSLs for a given RSL base–stock level.

From Section 2.3, we show that inventory on order corresponds directly to a customer demand. Because demand follows a Poisson distribution, the mean number of items in repair at the RV or in transit from the customer to the RV also follows a

Poisson distribution with mean, Φ_{0j}:

$$\Phi_{0j} = \lambda_{0j} L_{0j} + \lambda_{RV_j} L_{RV_j} \tag{3.1}$$

As shown in Section 2.3, when the net inventory is negative, a backorder exists. We can express this expected backorder level at the RV, β_{0j}, at a random point in time, with RV base–stock of s_{0j} as:

$$\beta_{0j} = \sum_{x=s_{0j}+1}^{\infty} (x - s_{0j}) \frac{(\Phi_{0j})^x}{x!} \exp(-\Phi_{0j}) \quad \forall j. \tag{3.2}$$

Applying a well known queuing formula [Little, 1961], we obtain the average waiting time, W_{0j}, due to stock–outs at the RV:

$$W_{0j} = \frac{\beta_{0j}}{(\lambda_{0j} + \lambda_{RV_j})}. \tag{3.3}$$

When the base–stock level at the RV, s_{0j}, exceeds the number of outstanding orders at the RV, x, positive inventory exists, i.e., $s_{0j} - x \geq 0$. We can determine this expected inventory level, I_{0j}, on–hand at the RV for part j as:

$$I_{0j} = \sum_{x=0}^{s_{0j}-1} (s_{0j} - x) \frac{(\Phi_{0j})^x}{x!} \exp(-\Phi_{0j}) \quad \forall j \tag{3.4}$$

and when the number of outstanding orders at the RV, x, exceeds the inventory position at the RV, s_{0j}, a backorder occurs, i.e., $s_{0j} - x < 0$. We compute the expected fill rate for part j, α_0, for the RV as:

$$\alpha_{0j} = \sum_{x=0}^{s_{0j}-1} \frac{(\Phi_{0j})^x}{x!} \exp(-\Phi_{0j}) \quad \forall j. \tag{3.5}$$

For this exposition, the Negative Binomial distribution is used to approximate the RSL backorder function by calculating the mean and variance of the number of outstanding orders at the RSL. However, we must first compute the variance of the expected number of backorders at the RV. The variance, using a Poisson distribution for the expected number of RV backorders, is:

$$Var(\beta_{0j}) = \sum_{x=s_{0j}+1}^{\infty} (x - s_{0j})^2 \frac{(\Phi_{0j})^x}{x!} exp(-\Phi_{0j}) - (\beta_{0j})^2. \tag{3.6}$$

We calculate the expected number of backorders and variance of the expected number of backorders at the RV using Φ_{0j}, the mean number of type j parts in the repair pipeline [Graves, 1985]. For computational efficiency, we use Graves' formula to compute the expected number of backorders at the RV:

$$E\{\beta(s_{0j})\} = E\{\beta(s_{0j} - 1)\} - Pr(x \geq s_{0j}) \tag{3.7}$$

The recursive variance formula is:

$$Var\{\beta(s_{0j})\} = Var\{\beta(s_{0j}-1)\} - [E\{\beta(s_{0j})\} + E\{\beta(s_{0j}-1)\}] [1 - Pr(x \geq s_{0j})] \tag{3.8}$$

We may now calculate the expected lead–time at RSL_i for part j, $\overline{L}_{ij}$, which includes the delay time at RSL_i due to RV backorders. The ASM (like METRIC) is an approximation because it ignores possible correlation between successive RSL replenishment lead–times, $\overline{L}_{ij}$. The RSL replenishment lead–times may not be independent and identically distributed because the RV does not have infinite repair capacity. However, we assume that the RV has sufficient capacity to allow all parts to begin repair immediately upon arrival at the RV. Parts are repaired on a first–come–first–served basis. We ignore the lead–time dependence and again use Palm's theorem, replacing the stochastic lead–time by its mean $\overline{L}_{ij}$. The expected lead–time, $\overline{L}_{ij}$, for replenishment of part j at RSL_i is:

$$\overline{L}_{ij} = L_{ij} + W_{0j}. \tag{3.9}$$

Because the RSL lead–time is known, it is possible to compute the expected number of backorders at the RSL. The probability that a part is in repair is a Poisson random variable with mean, $\overline{L}_{ij}\lambda_{ij}$. We can express the first and second moments of each RSL's number of outstanding orders in terms of the mean and variance, respectively, of the RV's backorder level. Thus, the mean number of outstanding orders at RSL_i is:

$$\Theta_{ij} = \lambda_{ij}\overline{L}_{ij}, \tag{3.10}$$

and the variance of the number of outstanding orders at RSL_i is:

$$Var(\Theta_{ij}) = \left(\frac{\lambda_{ij}}{\lambda}\right)^2 Var(\beta_{0j}) + \left(\frac{\lambda_{ij}}{\lambda}\right)\frac{(\lambda - \lambda_{ij})}{\lambda}\beta_{0j} + \lambda_{ij}L_{ij} \tag{3.11}$$

where $\lambda = \lambda_{0j} + \lambda_{RV_j}$.

In summary, the mean and variance of the RV backorder level are computed using (3.7) and (3.8). We then compute the mean and variance of the expected number of outstanding orders for all values of the RV base–stock levels, s_{0j}, within predetermined bounds using (3.10) and (3.11).

When the number of outstanding orders at RSL_i, x, exceeds the inventory position at RSL_i for part j, s_{ij}, a backorder condition occurs, i.e., $s_{ij} - x < 0$. In general, the expected backorder level, β_{ij}, at RSL_i for part j is:

$$\beta_{ij} = \sum_{x=s_{ij}+1}^{\infty} (x - s_{ij}) f(x|\lambda_{ij}\overline{L}_{ij}). \tag{3.12}$$

$$\text{where} \quad f(x|\lambda_{ij}\overline{L}_{ij}) = \begin{cases} \text{if the pmf for the Negative Binomial ASM, use eq. (2.3)} \\ \text{if the pmf for the Poisson ASM, use eq. (2.2)} \end{cases}$$

Analogous to (3.4), we determine the expected inventory level, I_{ij}, for part j at RSL_i:

$$I_{ij} = \sum_{x=0}^{s_{ij}-1} (s_{ij} - x) f(x|\lambda_{ij}\overline{L}_{ij}). \tag{3.13}$$

The expected fill rate of part j at RSL_i at a random point in time is:

$$\alpha_{ij} = \sum_{x=0}^{s_{ij}-1} f(x|\lambda_{ij}\overline{L}_{ij}). \tag{3.14}$$

Sherbrooke [1968] uses a two–stage heuristic in order to minimize the total system cost subject to a backorder constraint. We follow a similar approach and minimize the expected backorder level subject to a budget constraint on all parts stocked. If we let $\beta(s_{0j}, s_{ij})$ denote the expected number of backorders as a function of an RSL's and RV's base–stock, the ASM optimization problem is written as:

$$min_{(s_{0j},s_{ij})} \sum_{i=1}^{N} \sum_{j=1}^{J} \beta(s_{0j}, s_{ij}) \tag{3.15}$$

subject to:

$$\sum_{j=1}^{J} c_j \sum_{i=0}^{N} s_{ij} \leq B$$

$$s_{ij} \geq 0 \quad \forall\, i, j$$

and where c_j is the unit cost of part j and B is the system budget.

3.2 Analytical Stocking Model Algorithm

We develop a two–phase algorithm, similar to Sherbrooke's two–stage heuristic, to solve the ASM optimization problem. In Phase 1, for a given RSL stock level and part, the expected backorder levels and the corresponding first differences are calculated for each RSL. Then, for each part and RSL stock level, we aggregate backorder levels across all RSLs. In Phase 2, by applying *marginal analysis*, i.e., comparing the first difference of the RSL backorders found in Phase 1, we determine stock levels for each part at the RV and RSLs that minimize backorder levels for a given system cost. A summary of the two–phase algorithm, where Phase 1 consists of Steps 1 and 2 and Phase 2 consists of Step 3, follows:

- Step 1. Start with the RV stock level equal to zero for one part. For each RSL, calculate the expected number of backorders and the corresponding first differences for a predetermined maximum RSL stock level. Increase the RV stock level by one and repeat until the predetermined maximum RV stock level is obtained.

- Step 2a. Again, start with the RV stock level equal to zero. For each RSL stock level, use marginal analysis on the RSL backorders found in Step 1 to determine which RSL should be stocked to minimize backorders for the given part. Increment the RV stock level and repeat until the RV stock level reaches the maximum RV stock level; then, go to Step 2b.

- Step 2b. Find the minimum backorder level for all RV and RSL stock combi-

	Demand Rate (demands/week) RSLs 1 through 5					RV	lead–time (wk) L_0	L_i	Unit Cost
Part #1	1.0	3.0	5.0	10.0	15.0	1.5	0.02	0.01	1.0
Part #2	10.2	10.2	10.2	10.2	10.2	9.0	0.02	0.01	1.0

Table 3.1. Input data for ASM algorithm example

nations for the given part. Repeat Steps 1 and 2 for all parts in the system. Remove any non–convex points before Step 3.

- Step 3. From the results of Step 2b, determine the first difference from the optimal stock allocation for the RVs and RSL. Use marginal analysis to calculate the minimum backorder levels across all parts.

3.3 Analytical Stocking Model Algorithm Example

We illustrate the ASM algorithm by considering a repair parts logistics system with two parts and five RSLs replenished by one RV. The system allows for two–classes of service, i.e., both the RV and the RSLs service customers.

The demands placed on the RSLs follow a Poisson process with the parameters given in columns two through six of Table 3.1. The routine RV demands are shown in column seven, Table 3.1. The RV lead time is shown in column eight and the transportation time for the routine RV demand is equal to the RSL lead–time L_i, shown in column nine. We selected unit costs for both parts of \$1.00 for ease of exposition.

Step 1. Beginning with part #1, the RV stock level is set to zero. For each RSL, calculate the expected backorder levels and the associated first differences, until the predetermined maximum RSL stock level – 15 for this example – is obtained, i.e., for RSL 1 through RSL 5 use (3.2) through (3.12) to calculate the expected backorder level, $\beta_{ij}(s)$, as a function of the base stock level and the corresponding first difference, $\Delta = \beta_{ij}(s-1) - \beta_{ij}(s)$.

For example, the backorder level for RSL 1 is 0.0300 (Table 3.2, first row, second column) with a corresponding first difference of zero (Table 3.2, first row, third column).

RSL	RSL 1		RSL 2		RSL 3		RSL 4		RSL 5	
Stock	β	Δ	β	Δ	β	Δ	β	Δ	β	Δ
0	.0300	0	.0900	0	.1500	0	.3000	0	.4500	0
1	.0004	.0296	.0040	.0860	.0110	.1390	.0410	.2590	.0880	.3620
2	*		.0001	.0039	.0005	.0105	.0040	.0370	.0120	.0760
3	*		*		*		.0003	.0037	.0010	.0110

Table 3.2. Expected backorder level for part #1 at the RSLs, with RV stock level equal to zero $(* \equiv \beta < 0.0001)$

After computing the backorder levels and the first differences for all five RSLs, given an RSL stock level of zero, increment the stock levels at the RSLs by 1 and repeat Step 1. For example, given an RSL stock level of one, RSL 1 now has a backorder level of 0.0004 (Table 3.2, second row, second column) and a first difference of 0.0296 (Table 3.2, second row, third column) obtained by subtracting 0.0004 from 0.0300. We continue the process in Step 1 until we reach the maximum RSL stock level. For ease of exposition, Table 3.2 shows at most the first three RSL stock levels.

Once we determine all five of the RSL's expected backorder levels, for all 15 RSL stock levels, the RV stock level is increased by one. The process repeats, starting at the beginning of Step 1 with an RSL stock level of zero, until the predetermined maximum RV stock level – eight for this example – is met.

At the conclusion of Step 1, nine tables similar to Table 3.2 exist, where each table corresponds to a given RV stock level, e.g., the first table, Table 3.2, to an RV stock level of zero, the second table to an RV stock level of one, $\cdots$, the ninth table to an RV stock level of eight. Each table contains 16 rows, corresponding to RSL stock levels of 0, 1, $\cdots$, 15, and the columns of the tables contain the backorder levels and first differences corresponding to the five RSLs.

Step 2a. We determine the minimum backorder level for a given RV and RSL stock level and for a given part and place the results in a table, where the rows represent the amount of stock held at the RV and the columns contain the total amount of stock held at the RSLs. To calculate the upper bound on part #1's backorder level, we set the RV and RSL stock levels equal to zero, and sum over all RSL backorder levels in the first row of Table 3.2. For this example, the upper bound on the backorder level is: 0.03+0.09+0.15+0.3+0.45=1.02. We put this backorder value, (1.02), in the first position of Table 3.3.

We use a technique called marginal analysis to find the minimum backorder levels as additional parts are added to the RSL inventory. This is done by inspecting the first difference column, i.e., the Δ columns of Table 3.2, for the largest first difference among all RSLs for a single part type. By finding the largest first difference and subtracting it from the current backorder value, a new backorder value is calculated

RV Stock	Total Stock at RSLs (Optimal Allocation)								
	0	1	2	3	4	5	6	7	8
0	1.020	0.658	0.399	0.260	0.174	0.098	0.061	0.031	0.020
Optimal RSL Part Location		5	4	3	2	5	4	1	5

Table 3.3. Expected backorders for optimal allocation of stock at the RSLs for part #1 given an RV stock of zero

that offers the greatest marginal reduction in backorders when another part is added to the inventory.

In our example, the value in the last column, row two, of Table 3.2 (0.362), corresponding to RSL 5, contains the largest first difference value across all RSLs. Therefore, subtracting this first difference from the current backorder level, i.e., 0.362 from 1.02, provides the minimum expected backorder level when a part is stocked at an RSL. The resulting expected backorder level (0.658) is entered in row two, column three of Table 3.3. The number "five", corresponding to RSL 5, is entered in row three, column three of Table 3.3.

Continuing our procedure, we compare all the first differences from Table 3.2, except the first difference value already used, i.e., 0.369, to find the next largest first difference. In our example, this value is in column nine, second row of Table 3.2 (0.259). As before, this value is subtracted from the current total backorder level (0.658) and the result, i.e., the new total backorder level (0.399), is placed in column four of Table 3.3 along with the corresponding RSL number 4. To fill in the rest of the columns of Table 3.3, we continue the above procedure until we reach the

RV	Total Stock at RSLs (Optimal Allocation)								
Stock	0	1	2	3	4	5	6	7	8
0	1.020	0.658	0.399	0.260	0.174	0.098	0.061	0.031	0.020
1	0.533	0.324	0.178	0.103	0.057	0.033	0.018	0.007	0
2	0.381	0.226	0.120	0.066	0.033	0.020	0.009	0	0
3	0.347	0.205	0.108	0.058	0.028	0.018	0	0	0
4	0.342	0.203	0.107	0.057	0.027	0	0	0	0
5	0.341	0.202	0.106	0.056	0	0	0	0	0
6	0.340	0.201	0.105	0	0	0	0	0	0
7	0.340	0.200	0	0	0	0	0	0	0
8	0.340	0	0	0	0	0	0	0	0

Table 3.4. Expected backorders for optimal allocation of stock at the RSLs for part #1 given an RV stock of zero through eight

predetermined maximum base–stock level.

The RV stock level is incremented and we repeat Step 2a in order to determine the optimal allocation of part #1 stock among the RSLs for this RV stock level. Table 3.4 shows the minimal backorder levels that correspond to the optimal allocation of stock at the RSLs for different levels of stock at the RV. The rows in Table 3.4 correspond to a given RV stock level while the columns correspond to the optimal allocation of stock at the RSLs. For this example, we repeat Step 2a and begin by calculating the upper bound for the backorder level when one part is in stock at the RV and zero at the RSL, i.e., 0.533 is placed in row two, column two of Table 3.4. Reapplying Step 2a using marginal analysis, we complete row three of Table 3.4. This process continues until the maximum RV stock level is obtained, i.e., eight for this example.

In Step 2b, we determine the minimum expected backorder level for a given RV

Total Stock level	Minimum Backorder	Backorder Reduction
0	1.020	-
1	0.533	0.487
2	0.324	0.209
3	0.178	0.146
4	0.103	0.075
5	0.057	0.046
6	0.033	0.024
7	0.018	0.015
8	0.007	0.011

Table 3.5. Expected backorders optimally allocated for part #1

and RSL stock allocation. The diagonal entries in Table 3.4 represent the minimum expected backorder levels for a given RV and RSL stock allocation. Therefore, the optimal stock allocation for a given system stock level is obtained by inspecting the diagonal (running from upper–right to lower–left) in Table 3.4 for the lowest backorder level. We use the results from Table 3.4 to determine the minimum backorder level for a given total system stock level, and place these results in Table 3.5. The rows of Table 3.5 correspond to the total system stock level, the second column contains the minimum backorder level and the third column contains the first difference, i.e., backorder reduction, corresponding to the system stock level in column one.

For example, a total system stock level of two corresponds to optimal system backorders of 0.399, 0.324, and 0.381 for part #1. The value 0.324 is placed in row three, column two, of Table 3.5, because it is the minimum backorder level on the diagonal which corresponds to a total system stock level of two, i.e., the RV and RSL receive one part each.

Total Stock level	Minimum Backorder	Backorder Reduction
0	1.530	-
1	0.936	0.594
2	0.649	0.287
3	0.527	**0.122***
4	0.406	**0.121***
5	0.212	0.194
6	0.082	0.130
7	0.040	0.042
8	0.029	0.011

Table 3.6. Expected backorders optimally allocated for part #2

We calculate the backorder reduction values in column three, Table 3.5 the same way we calculated the first differences described previously in Step 2a, e.g., 0.533 - 0.324 for the backorder reduction value of 0.209, entered in row three, column three, of Table 3.5. Steps 1 and 2 are repeated for the remaining parts in the system, i.e., we repeat Steps 1 and 2 for part #2 to obtain the optimal allocation of stock shown in Table 3.6.

As part of Step 2b, any non-convex points points found in Tables 3.5 and 3.6 must be removed prior to Step 3 of the algorithm. The backorder reduction for part #1 is monotonically decreasing. Therefore, the optimal solutions for part #1 all lie on the convex hull and dominate the interior solutions. However, the lack of strict monotonicity for the backorder reduction of part #2 reveals dominated solutions. Because the marginal analysis process for Step 3 requires all points to lie on the convex hull, these interior points shown in bold in Table 3.6 must be removed. The bold values in Table 3.6 are strictly dominated solutions, i.e., they are not optimal.

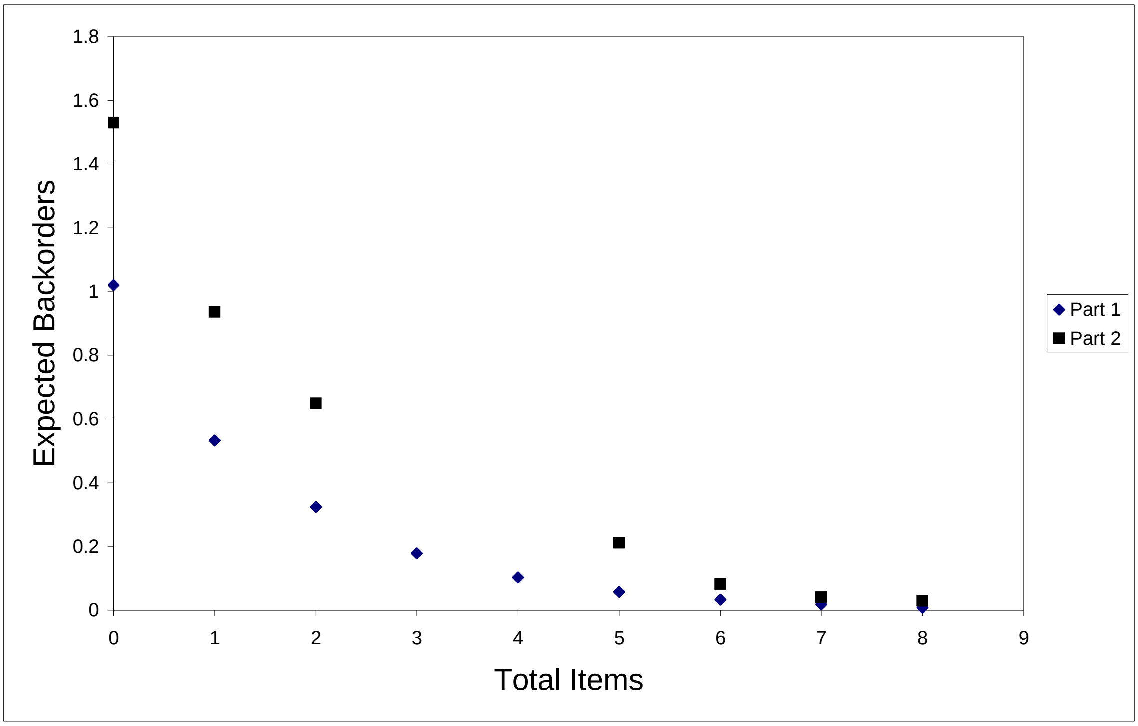

FIG. 3.1. Part #1 vs. part #2 sub–problem convexity with non–convex points removed

Therefore, removal of these points will not affect the optimal solution. Figure 3.1 illustrates the sub–problem convexity from Tables 3.5 and 3.6 with the non–convex points for part #2 removed.

Step 3. We determine the minimum backorder level across parts for a given system stock level. To take the unit costs of parts into account, we divide the backorder reductions from Tables 3.5 and 3.6 by the part unit cost. We use these "adjusted" first differences and again use marginal analysis to calculate the minimum expected number of backorders across parts (Table 3.7).

For each row, column one of Table 3.7 shows the total stock level allocated in

Total	Total		RV Stock		RSL Stock	
Stock Level	Backorder	Cost	Part #1	Part #2	Part #1	Part #2
0	2.55	$0.00	0	0	0	0
1	1.96	$1.00	0	1	0	0
2	1.47	$2.00	1	1	0	0
3	1.18	$3.00	1	2	0	0
4	0.97	$4.00	1	2	5	0
7	0.53	$7.00	1	0	5	1,2,3,4,5
8	0.38	$8.00	1	0	5,4	1,2,3,4,5
9	0.25	$9.00	1	1	5,4	1,2,3,4,5
10	0.18	$10.00	1	1	5,4,3	1,2,3,4,5
11	0.13	$11.00	1	1	5,4,3,2	1,2,3,4,5
12	0.09	$12.00	1	2	5,4,3,2	1,2,3,4,5
13	0.07	$13.00	2	2	5,4,3,2	1,2,3,4,5
14	0.06	$14.00	1	2	5,4,3,2,1,5	1,2,3,4,5
15	0.05	$15.00	1	3	5,4,3,2,1,5	1,2,3,4,5
16	0.04	$16.00	1	3	5,4,3,2,1,5,4	1,2,3,4,5

Table 3.7. Total expected backorders optimally allocated across parts

the system between the RV and the RSLs and column two gives the corresponding optimal total backorder level. To compute the upper bound on system backorders as before, we sum all minimum backorder levels corresponding to zero system stock. For example, we compute the value in row one, column two, in Table 3.7 (2.55), by summing the minimum backorder levels we found for each part when the total stock level is zero, i.e., 1.02 + 1.53 from Tables 3.5 and 3.6.

To find the remaining values in column two of Table 3.7 for different total stock levels in the system, we use marginal analysis as before. For example, 1.96 from Table 3.7 is found by subtracting the larger of the two backorder reduction values – row two, Tables 3.5 and 3.6 – from the current total backorder level when the stock level is one, i.e., 2.55 - max{0.487, 0.594}.

We calculate the total costs in column three of Table 3.7 by summing the number of like parts in the system and multiplying the total number of parts by their respective unit cost.

At the same time that the minimum backorder levels are calculated in Step 3, the system tracks where the parts are stocked using the information stored in Table 3.3. The number of parts held at the RV is displayed in columns four and five for parts #1 and #2, respectively. We display in columns six and seven of Table 3.7, the RSL locations where parts #1 and #2, respectively, are stocked. For example, given a total stock level of 13 parts in the system, part #1 is stocked with two items at the RV and one item each at RSLs five, four, three, and two. Part #2 has two items stocked at the RV and one item stocked at each RSL. In the next chapter we develop a simulation model to evaluate the Poisson and Negative Binomial ASMs.

Chapter 4

SIMULATION METHODOLOGY

The METRIC model and its variants are approximations which, in the past, have been evaluated using exact closed–form solutions. Because of the large number of RSLs and parts contained in the SVLN, obtaining exact closed–form solutions is computationally intractable. Therefore, we employ a discrete–event simulation to compute output parameters to evaluate the ASM's accurateness. We use Arena 5.0 as the simulation software and the Arena Output Analyzer 5.0 to analyze the simulation output [Kelton et al., 2002]. The simulation model is run on a Dell Inspiron 8100 with 384 MB of RAM.

This chapter is divided into two sections. Section 4.1 contains a detailed description of the simulation model and Section 4.2 describes the simulation output analysis methodology and techniques used to generate steady–state output parameters.

4.1 Simulation Model Description

The simulation model consists of four process modules and one counter module (Figure 4.1). A brief description of the simulation modules follows:

- Initial Stock Module: this module sets the initial inventory level at the RV and RSLs to a base–stock level.

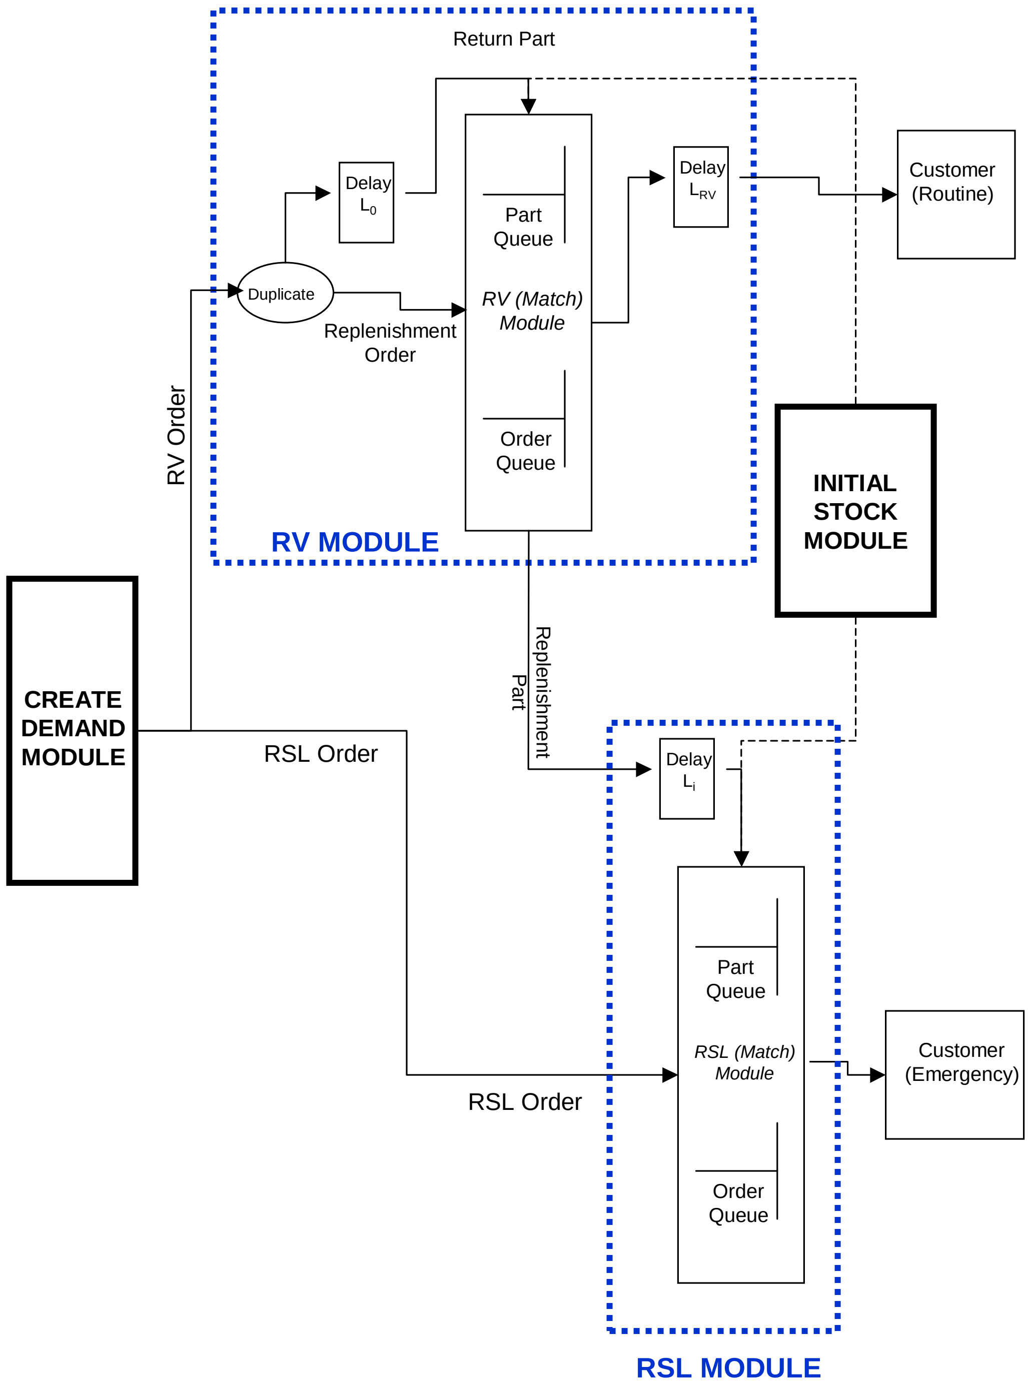

FIG. 4.1. Simulation flow–chart

- Create Demand Module: this module generates demand at the RV and RSLs by sending orders to each respective location.

- RSL Module: this module contains the RSL order and part queues, and generates queue statistics for the RSLs.

- RV Module: this module contains the RV order and part queues, and generates queue statistics for the RV.

- Counter Module: this module collects the backorder statistics from the RV and RSL Modules and stores the result as a *time–persistent* statistic. This module is not shown in Figure 4.1.

We simulate only one part at a time because of the excessive clock time required to run the simulation to steady–state. The input data, which consists of two main components, is read from an Excel spread–sheet. The first segment of the simulation input uses the forecasted mean demand rates obtained from DRP. The second segment of the input data corresponds to the optimal base–stock level, obtained in Section 3, for the simulated part.

When the simulation is initialized, the Initial Stock Module sets the inventory level at the Part Queues in the RV and RSL Modules to the base–stock level. Beginning with the simulation clock time equal to 1.0 days, the Create Demand Module generates orders at the RSLs and RV according to an exponential distribution, where the former represent the customer demands for emergency parts and the latter represent either routine orders or replenishment orders.

When the RV Module receives an order, it is duplicated. The "original order" is immediately stored in the RV Order Queue and the "duplicated order", i.e., the part returned by the customer, is stored in the RV Part Queue after the RV lead–time delay, L_0. Orders are stored in the sequence in which they are received, and are filled on a first–come, first–served basis. If an order arrives at the RV Order Queue, and no parts reside in the RV Part Queue, then a backorder exists. However, if a part is in the RV Part Queue when the order arrives, then the order is filled, i.e., a "match" occurs. If a match occurs, the part is either released immediately to the RSL as a replenishment part, or released after the routine customer lead–time delay, L_{RV}, to the customer as a routine part. The decision of whether a part is transported to an RSL or directly to a customer is dependent upon the order type, i.e., a routine or replenishment order.

Once a replenishment part is released from the RV Part Queue, the replenishment part is placed in the RSL Part Queue after an RSL lead–time delay of L_i. Orders are processed on a first–come, first–served basis. If an order exists in the RSL Order Queue, then the part is issued to the customer and the order is filled. Otherwise, the part is stored in the Part Queue until an order arrives. If an RSL order arrives in the RSL Order Queue, a backorder results if no parts reside in the RSL Part Queue.

Arena calculates the time–average number of orders waiting in the Order Queues of all RV and RSL match modules, i.e., the average number of backorders at the RV and RSLs, respectively. The Counter Module computes the average number of backorders in each RSL Order Queue and stores the result as a time persistent statistic variable. At the end of a simulation run, there is a backorder value for each part and each discrete simulation time event, i.e., one day.

The simulation must be run to steady–state because the ASM is based on steady–state queuing equations. The next section discusses the methodology used to run the simulation in steady–state.

4.2 Statistical Output Analysis

In this section, we address statistical analysis for steady–state parameters, i.e., we analyze the output from a single run of a non–terminating simulation. There are two primary concerns with steady–state simulations: (i) How do we obtain the initial conditions for achieving "steady–state–behavior?" (ii) At what point does a simulation achieve steady–state behavior?" We address the former question first.

4.2.1 Initial Transient Problem

For our purposes, the output statistics from the simulation model are only meaningful in steady–state. In theory, the initial conditions do not matter. Practically speaking, we cannot, however, run the simulation model for an infinite amount of time. Because the simulation output is of finite length, and the observations near the beginning of the simulation may not represent the steady–state behavior of the "true" system, the resulting output may be biased. That is, the output may reflect some arbitrary initial conditions rather than the true long–run average behavior of the system. This problem is referred to as the *initial transient problem.*

Suppose we want to estimate the steady–state mean, v, of the output stochastic process, Y. We define the mean as:

$$v = \lim_{i \to \infty} E(Y_i) \tag{4.1}$$

where $Y_i \; \forall i$ are the random variables that constitute the stochastic process Y.

The expected value of the estimator of Y, $E(\overline{Y})$, may not equal the estimate of the steady–state mean, $\overline{v}$, because of the initial transient problem. In this case, the point estimate of Y is biased. To avoid a biased estimate of the steady–state mean, we run the simulation model for a "sufficient" number of time periods and delete a predetermined number of observations from the beginning of the simulation run. This approach is referred to as "warming up the model" or the "initial–data deletion" approach.

We use Welch's graphical procedure to determine a sufficient amount of warm–up time such that the expected value of the remaining observations of Y is unbiased [Welch, 1983] and [Law and Kelton, 2000, pp. 520–521]. Welch's procedure requires making m replications of a simulation, where $m \geq 5$. We use Welch's procedure only on a few selected parts to find the typical warm–up time because the procedure requires excessive computational effort. We find that a typical warm–up time of 5000 simulation days is more than sufficient to ensure proper initial conditions required for our steady–state simulation. We now address the second question – At what point does a simulation achieve steady–state behavior? We are interested in this question for determining when to terminate the simulation.

4.2.2 Output Analysis for Means

In addition to finding an acceptable warm–up time, specialized techniques for non–terminating simulations are required in order to construct a point estimate and a confidence interval for the steady–state mean, v. In most cases, simulation output data are highly correlated. Therefore, we cannot use classical statistics based on independent and identically distributed observations of a stochastic process for analysis of simulation output. Law and Kelton [2000, pp. 525–537] lists six fundamental approaches so that standard statistical techniques can be used to estimate the steady–state mean of a stochastic process and to construct a confidence interval.

We can implement Law and Kelton's *batch means* method using the Arena Output Analyzer. After a simulation run, we divide the simulation output, i.e., the time persistent statistic, into "batches", $Z_i \; \forall i = 1 \cdots m$, where m is the number of batches of the simulation run, each with k observations. Each batch Z_i consists of $X_j \; \forall j = 1 \cdots k$ observations. We define the point estimator, $\overline{v}$, as:

$$\overline{v} = \sum_{i=1}^{m} Z_i/m \tag{4.2}$$

We choose the batch means method because it does not bias the estimate of the mean. However, this method may bias the estimate for the variance which affects the accuracy of the confidence interval. If the batch size k is large enough, the mean estimates of the batches are approximately uncorrelated [Law and Carson, 1979]. However, if the batch sizes are too small, the mean estimates of the batches become highly correlated which results in a biased estimate for the variance $\overline{var}(\overline{v})$. This can result in confidence intervals which are too narrow, forcing the true mean outside of the confidence interval. To help ensure that the variance $\overline{var}(\overline{v})$ remains unbiased, we increase the batch size until the mean estimates of the batches are uncorrelated.

We simulate only one part at a time, and obtain the steady–state output corresponding to the average number of backordered parts at the RSLs. In the following section, we compare the number of backorders by part for the ASM and the simulation.

Chapter 5

RESULTS

We present results for the ASM, which is written in FORTRAN 95 using the Sun Forte development software, and run on a Sun Blade 1000 equipped with a single 750 MHz Ultra Sparc III processor and one GB of RAM.

This chapter is divided into several sections. Section 5.1 discusses the data set used for the ASM and SCO policy comparisons and for the simulation model. Section 5.2 compares the part backorders from the Negative Binomial and Poisson ASMs against the simulation results. Section 5.3 compares the two ASM policies and the current SCO policy for the data set given in Section 5.1. Section 5.4 compares an ELT case and the ASM results obtained in the previous sections. In the last section, Section 5.5, we increase the RSL lead–time from two days to nine days, and consider two separate cases, a non-ELT case and an ELT case.

5.1 Test Data

We examine the expected number of ASM backorders obtained from a sample data set drawn from over 4,000 parts contained in the SVLN, and compare the ASM with current SCO stocking policies. Inventory management systems typically segment repairable parts inventory by criteria such as location, cost, and part criticality. We choose to segment items using a cost criterion, i.e., demand value (DMV). DMV is the unit cost of the part multiplied by the demand placed for the part during a predetermined time interval, e.g., over a 6 week period. Cohen et al. [1999] indicate that high cost and low demand items constitute approximately 80% or more of the total inventory costs for a replacement parts system. We can control the majority of system inventory costs by concentrating on a small number of high–DMV items. For the parts with low DMV's, we may use a periodic review policy, i.e., a policy that does not require a generous amount of time and resources, to control the remaining repairable and consumable parts.

Pareto analysis identifies the high–DMV parts. Cumulative DMV is given on the **y**–axis and the corresponding proportion of total inventory on the **x**–axis (Figure 5.1). Parts that constitute 80% of the DMV are classified as "A" parts and parts that constitute the next 15% are classified as "B" parts. From Figure 5.1, the Pareto diagram for over 4,000 parts from the SLVN indicates that the "A" and "B" parts constitute approximately 15% of the total stock, which controls 95% of the cumulative DMV in the SVLN inventory. After we remove the end–of–life and "one–time–buy" parts from the "A" and "B" parts, 151 parts remain, which corresponds to a cumulative DMV of 91.4 %. The average unit cost for the 151 parts is $1,433, with ranges from $66 to $12,410.

We obtain weekly demand forecasts from the distribution resource planning (DRP) system for the parts identified by the Pareto analysis. We use the same average forecasted demand rates as inputs to the ASM. This eliminates differences between SCO and ASM stock allocations because of forecast bias.

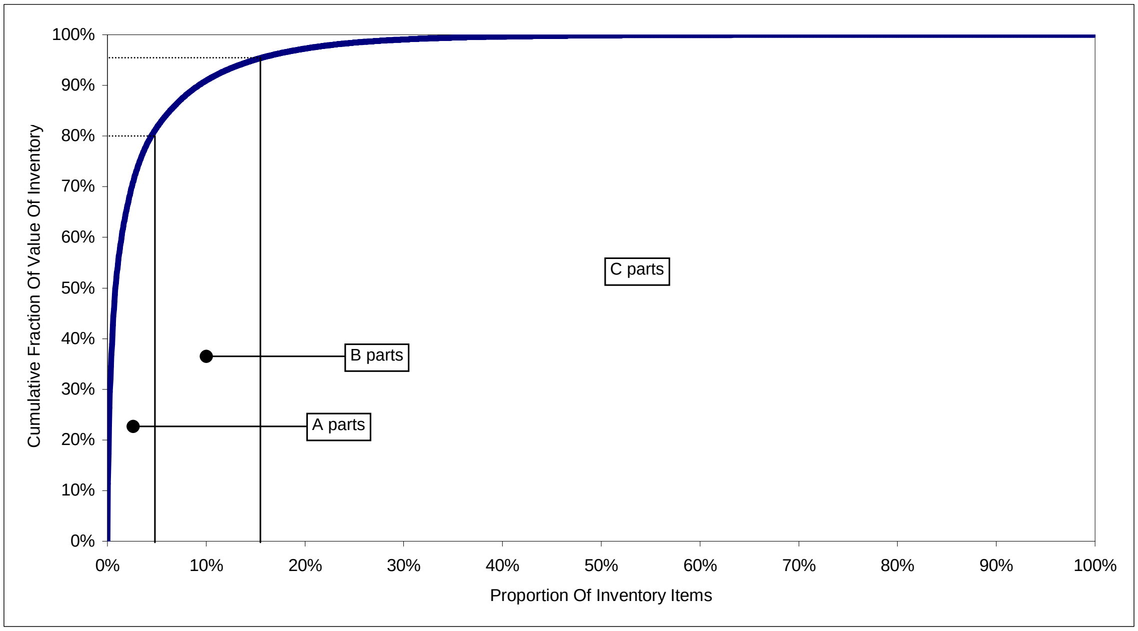

FIG. 5.1. Pareto analysis of SCO inventory over a six week period

The methods used by DRP to determine the average forecasted demand rate are: (i) weighted average, (ii) double smoothing, (iii) moving average, and (iv) single smoothing. Currently, DRP uses the average demand forecasts to establish inventory *thresholds*, i.e., stocking levels, which improves the SCO planner's ability to manage the allocation of inventory throughout the SVLN.

We show the forecasted weekly demand rate at the RSLs in Figure 5.2, with the numbers corresponding to the 151 parts and 100 RSLs along the **x**– and **y**– axes, respectively, and the RSL demand rates along the **z**–axis. From Figure 5.2 we observe high variability of RSL demand across all 151 parts, ranging between zero and 48 parts per week.

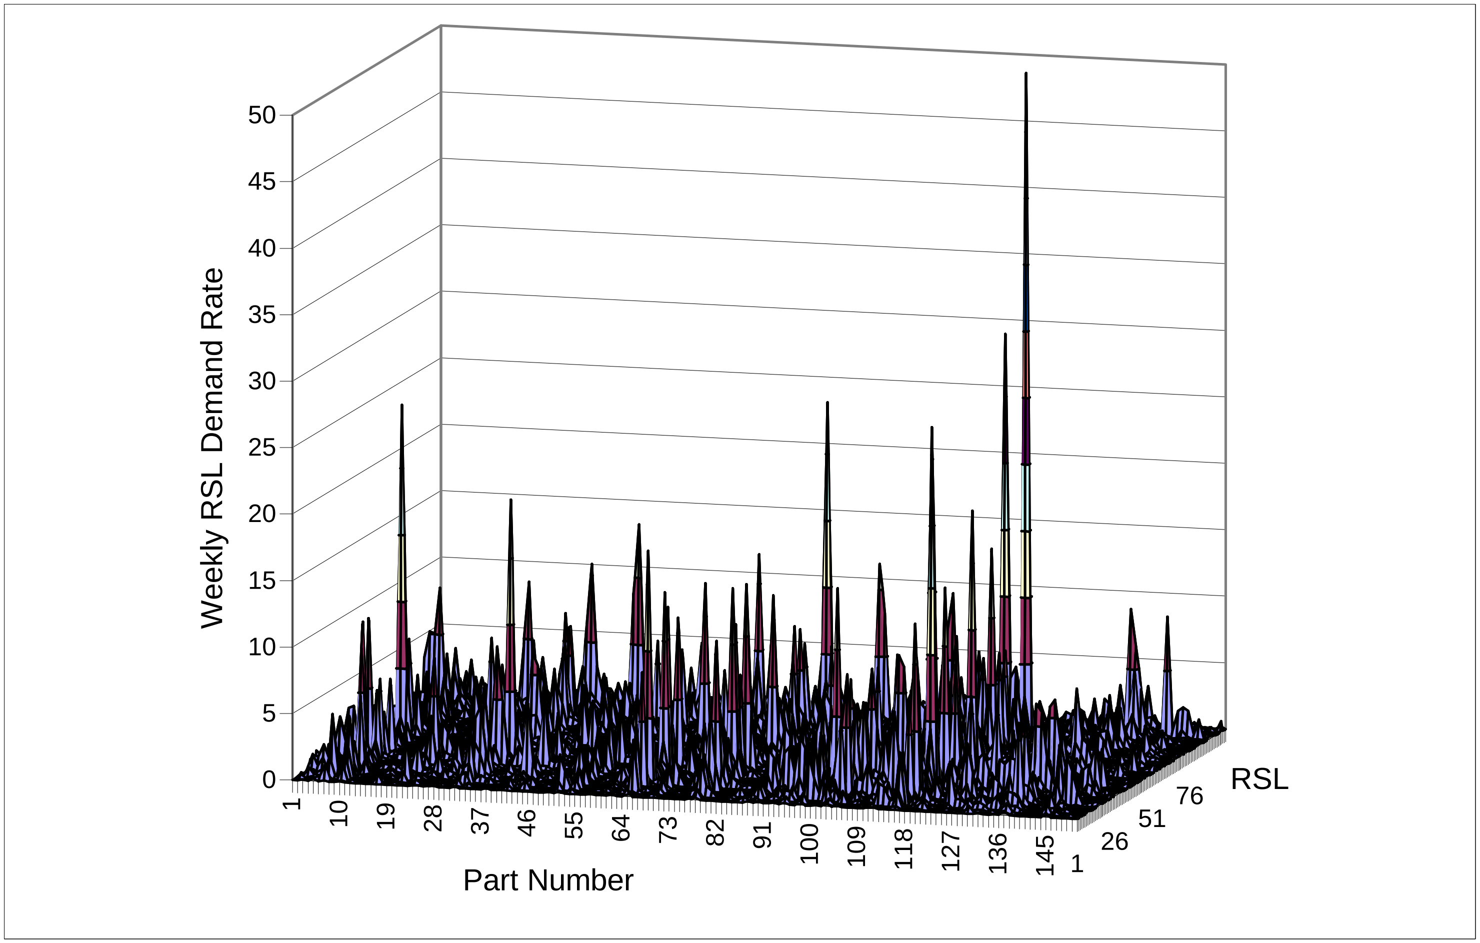

Fig. 5.2. Forecasted demand at the RSLs for each high–DMV part in the system

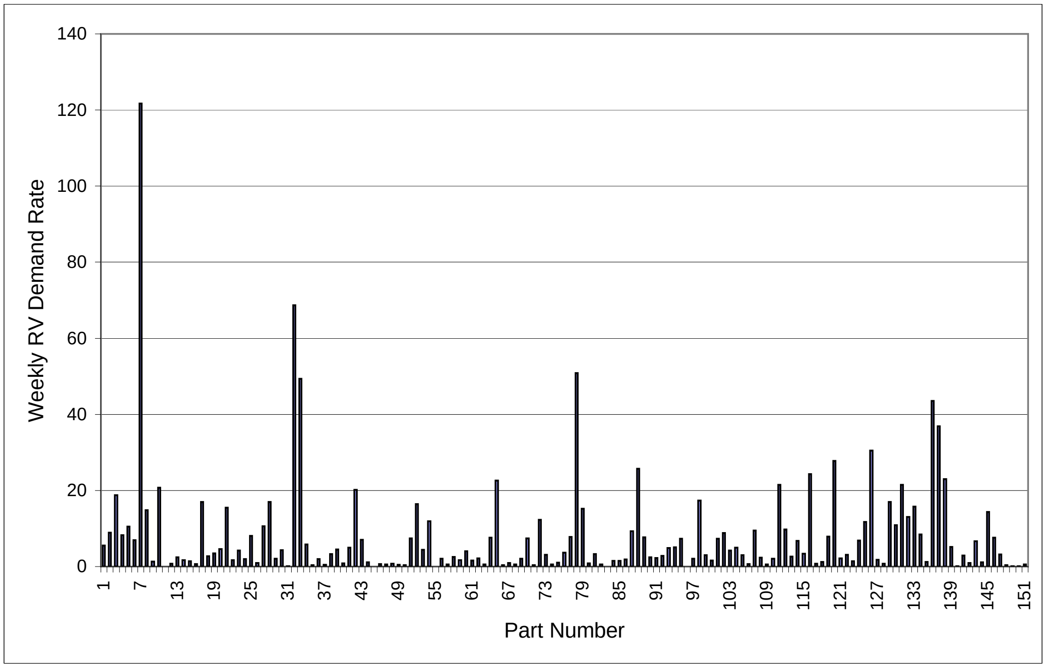

FIG. 5.3. Forecasted demand at the RV for each high–DMV part in the system

Routine RV demand for the 151 high–DMV repairable parts is variable as well. For several parts, RV forecasted demands are zero, although for other parts the forecasted demands exceed 20 parts per week. Figure 5.3 depicts the weekly RV demand rate for each of the 151 parts. Appendices B.1 and B.2 contain the average weekly demand rate data used in this dissertation for the RSLs and RV, respectively.

We develop two versions of the ASM to obtain optimal stock allocation at the RV and the RSLs. We fit the Poisson distribution to the RSL backorder function for one version, and the Negative Binomial distribution to the RSL backorder function for the other version. For both models, we assume an average RV lead–time of 22 days, i.e., 12 days for the customer to return the part to the RV and 10 days for the RV to repair the part. Furthermore, we assume an average RSL lead–time of two days, and a routine delivery lead–time from the RV to the customer of two days. The model may be extended, without loss of generality, to accommodate different RV and RSL lead–times. In Section 5.2 we compare both versions of the ASM, and use simulation results to evaluate both models.

5.2 Negative Binomial and Poisson Analytical Stocking Model Comparison

In this section, we compare the Negative Binomial and Poisson versions of the ASM to determine if significant differences in the expected number of backorders exist. The Negative Binomial and Poisson distributions are used to approximate the backorder function at the RSLs (Section 3). Graves [1985] suggests that a model using the Negative Binomial distribution provides a better approximation of the expected number of backorders than a model using the Poisson distribution, and uses a small number of stocking locations and parts with low demand rates (Section 2.6) to support his claim. By contrast, we apply the ASM to a large scale, multi–class, multi–echelon system with many RSLs, which is characterized by large variations in RV and RSL demand. To our knowledge, this type of analysis does not exist in the literature.

Given the input data set described in Section 5.1, the ASM requires approximately 10 hours to calculate the base–stock levels at all 100 RSLs and at the RV. The ASM output used to form the exchange curves shown in Figure 5.4 illustrates the tradeoff between the expected number of system backorders and total system cost. Each point that comprises the exchange curves corresponds to base–stock levels for all 100 RSLs and at the RV, for each one of the 151 parts. There are over 30,500 discrete points that form each of the exchange curves shown in Figure 5.4. The exchange curves of Figure 5.4 are valuable because they allow management to identify base–stock levels that minimize the expected number of system backorders, given a system budget constraint.

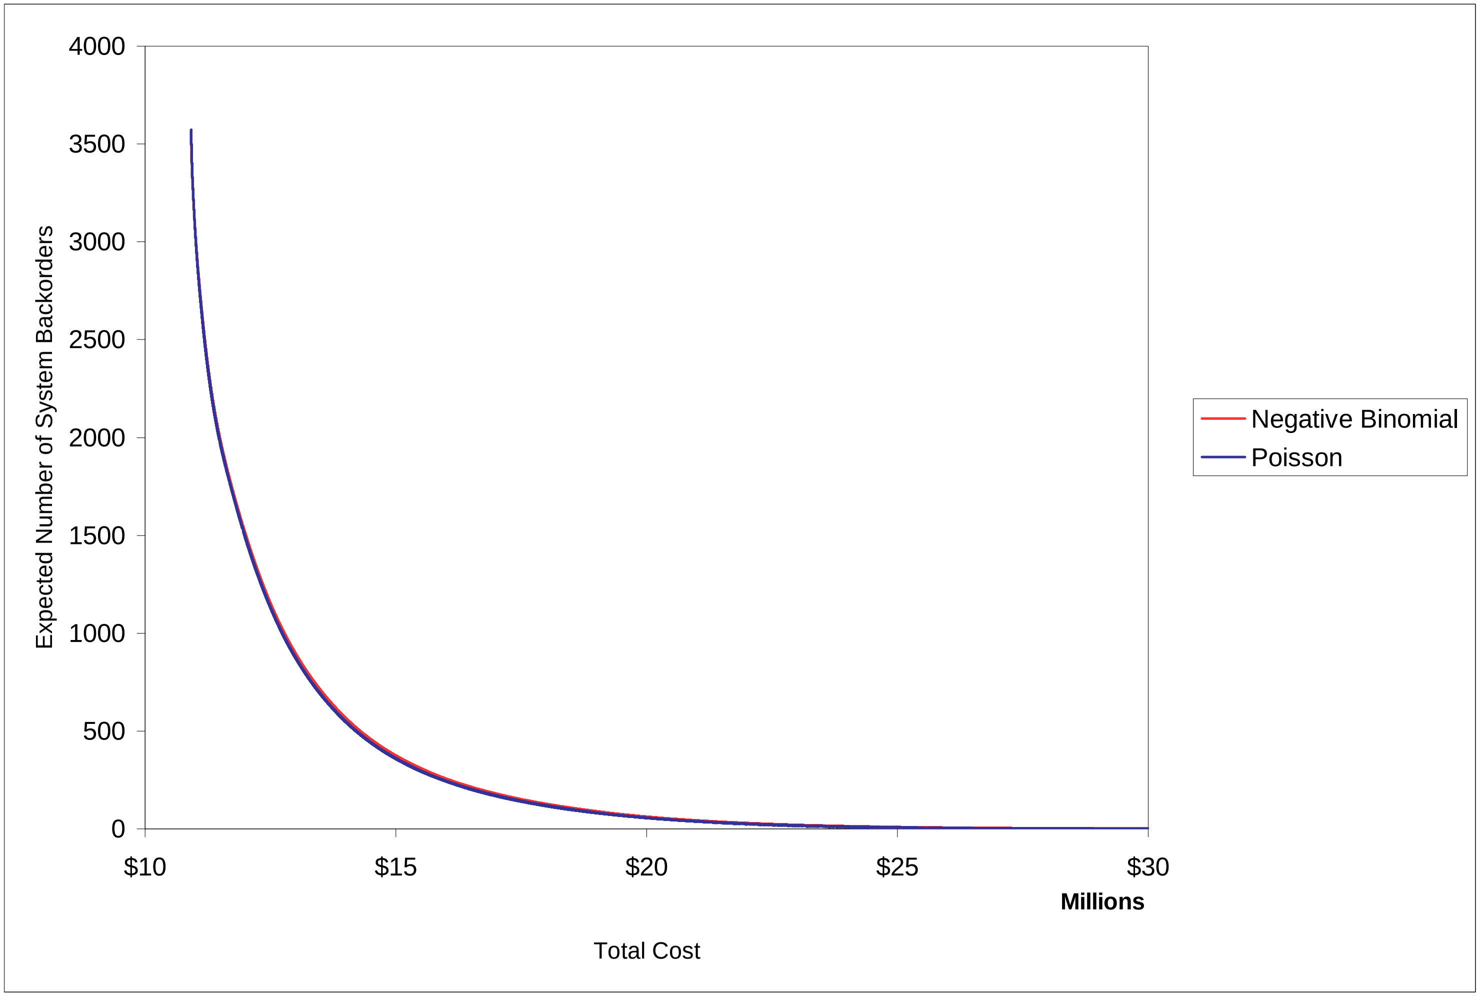

FIG. 5.4. Optimal ASM exchange curves

For a small problem, the cost on the exchange curve would begin at zero, because at that point, there would be no stock in the system. However, for tractability, we must generate the exchange curves in Figure 5.4 with RV stock levels greater than zero. Therefore, the cost on the optimal exchange curves in Figure 5.4 begins at approximately $10,918,842 and the corresponding total backorder level is 3,572 parts.

As expected, the exchange curves in Figure 5.4 are monotonically decreasing because as more stock is added to the system, the total system backorder level decreases. The curves are convex because of the diminishing marginal benefit as more stock is added to the system. The ASM exchange curves are scaled in Figure 5.5 for greater resolution to show the difference between the Negative Binomial and the Poisson ASMs. The Poisson ASM is shifted to the left of the Negative Binomial ASM. This shift shows that for a given backorder level, the Poisson ASM recommends less base–stock than the Negative Binomial ASM. Figure 5.5 indicates that the difference between the system backorder levels is not significant, but slight differences in the expected number of backorders may affect the final allocation of parts throughout the system. In the next section, we evaluate the accuracy of both ASM models by comparing the expected number of backorders to those derived from a simulation model.

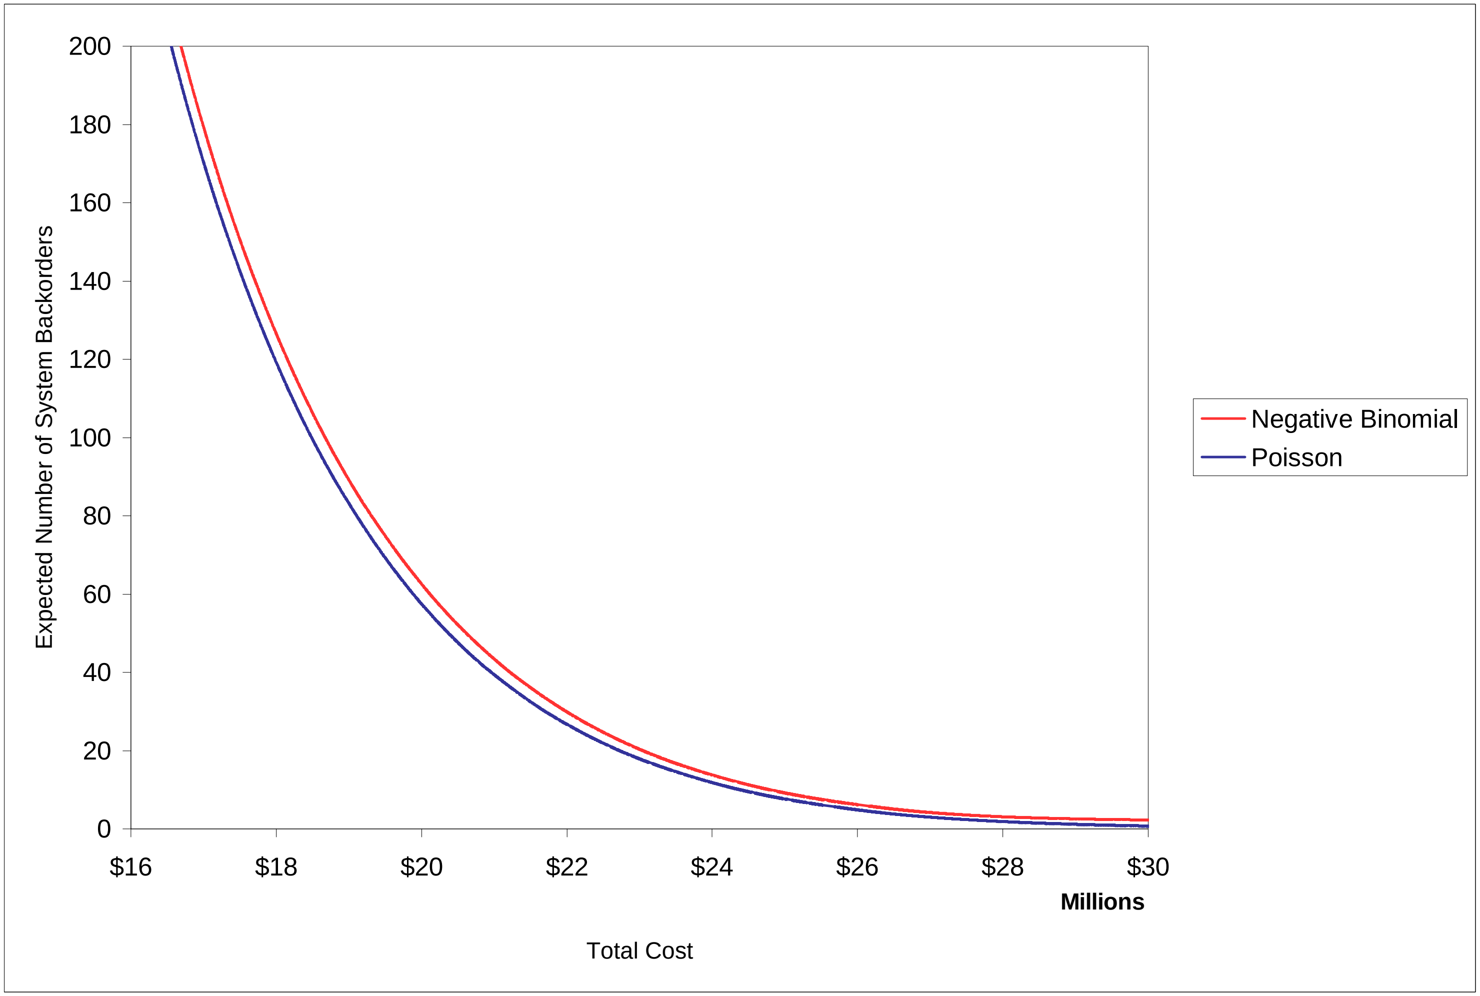

FIG. 5.5. Comparison of Poisson and Negative Binomial ASMs

5.2.1 Using Simulation to Evaluate the Expected Number of Backorders Determined by the Analytical Stocking Models

We evaluate the ASM expected number of backorders for both models using a discrete–event simulation, and compare the ASM expected number of backorders (by part), with the average backorder levels obtained from the simulation.

The simulation uses the same inputs as the ASM, i.e., the ASM optimal stock allocation, s_0 and s_i, average RV and RSL lead–times, L_0 and L_i, respectively, and forecasted demands at the RV and RSLs, λ_0 and λ_i, respectively, to produce average backorder levels within a 95% confidence interval. Each simulation is run to steady–state, typically 50,000 to 75,000 simulated days, each with a 5,000 simulation day warm–up time. The typical simulation wall time is two days for a given part.

Suppose we run a simulation with a warm–up period of 5,000 days and a termination time of 75,000 days. After we start the simulation, the input data is read into the simulation and the warm–up period begins. After 5,000 days, the simulation warm–up period ends, and the simulation then collects the average number of backorders at each RSL and stores the result as a time persistent statistic. The time persistent statistic becomes the input data for the Arena Output Analyzer. We show this output in Figure 5.6. The time persistent statistic is segmented into 37 batches; the batches are 1,890 days long. The correlation between the batches is 0.786, and the average backorder level for this part, the steady–state output parameter, is 1.56 parts. The 95% confidence interval is 1.54 to 1.57 parts. The time series plot shows the simulation time on the **x**–axis, beginning at 5,000 days and ending at 75,000 days, and the average number of part backorders on the **y**–axis.

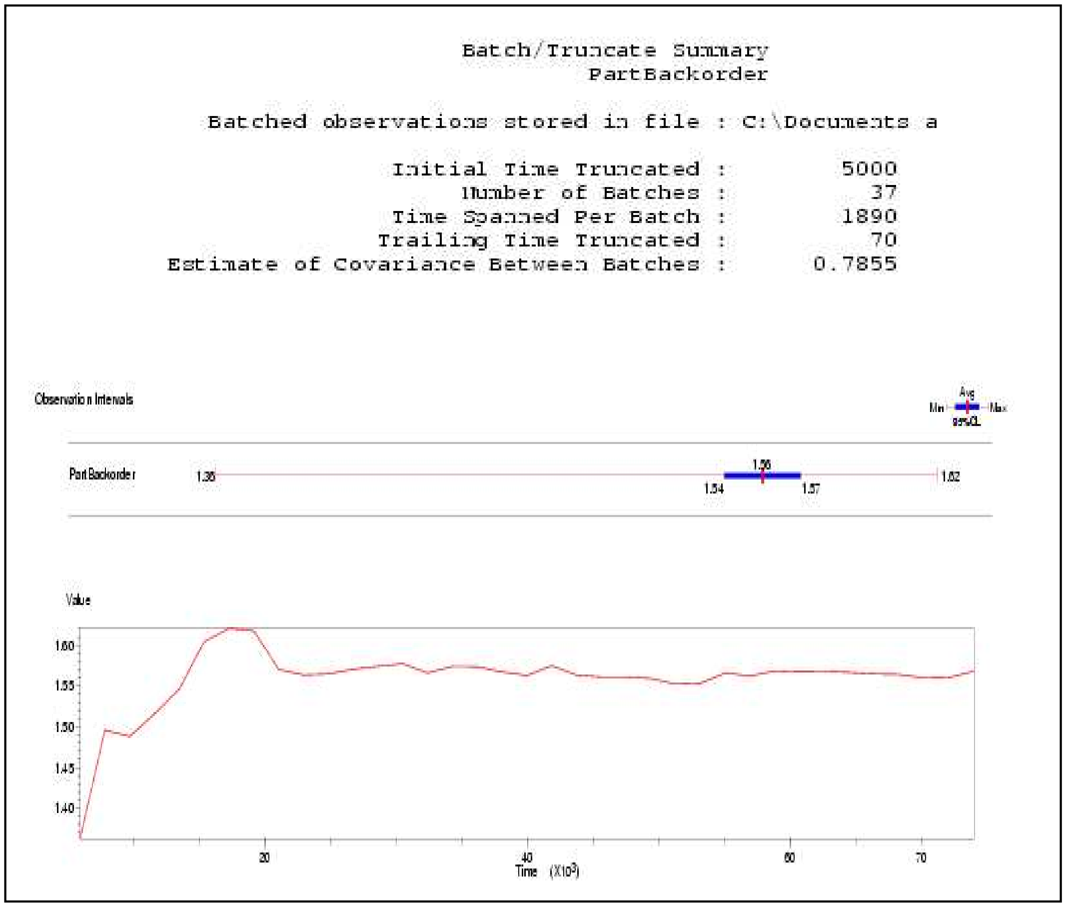

FIG. 5.6. Example output from the Arena Output Analyzer

Using the ASM optimal stock allocation, s_0 and s_i, obtained in the previous section, we generate the expected number of backorders by part at the RV and RSLs using (3.2) and (3.12), respectively, and compare these ASM expected part backorders with the simulation results. If the backorder level obtained from the ASM falls outside the confidence interval, the ASM backorder level is then compared to the closer of the two confidence interval bounds. We determine the percent ASM error for a single run t of the ASM as:

$$\text{percent ASM error} \quad = \frac{A_t - S_t}{S_t} * 100 \tag{5.1}$$

where S_t is the average number of backorders for simulation t and A_t is the expected number of backorders for the ASM.

For example, using the simulation output of Figure 5.6, suppose this is the first simulation and ASM comparison for a given part, i.e., $t = 1$. Then the Negative Binomial ASM backorder level, A_1, is 1.563. The percent ASM error for the Negative Binomial ASM is zero because the ASM backorder level falls within the confidence interval. However, the Poisson ASM backorder level, A_1, is 1.501, and falls outside the lower–bound of the confidence interval. Given the lower–bound of the confidence interval for this run, $S_1 = 1.54$, the percent ASM error for this point is:

$$\text{percent ASM error} \quad = \frac{1.50 - 1.54}{1.54} * 100 = -2.6\% \tag{5.2}$$

To compare the number of backorders for the ASM and the simulation, we first use the simulation to generate the average number of backorders for a selected number of parts. Then we compute the ASM's expected number of backorders for the same selected number of parts. Finally, we calculate the mean percent ASM error using:

$$\text{mean percent ASM error } = \sum_{t=1}^{N} \frac{\frac{A_t - S_t}{S_t}}{N} * 100. \tag{5.3}$$

where N is the total number of observations.

The mean percent ASM errors obtained from (5.3) will equal zero in the absence of systematic over– or under–prediction of part backorders. However, it is possible for the "positive" errors to cancel the "negative" errors. In this case, the resulting mean percent ASM error is negligible, but an individual observation's error may be considerable. To detect this type of problem, we calculate the mean absolute error, (5.4).

The mean absolute ASM error for all N simulations is calculated using:

$$\text{mean absolute ASM error } = \sum_{t=1}^{N} \frac{\left|\frac{A_t - S_t}{S_t}\right|}{N} * 100. \tag{5.4}$$

From the 151 parts in our data set, we select a representative sample of parts based on the variability of demand rates at the RSLs, and on the variance to mean ratio of outstanding orders at the RSLs, which range from 1.01 to 2.1. We use the ASM stock levels, for a sample of 60 parts, from two points selected from Figure 5.4 to generate the $N = 120$ simulation results used to calculate the mean percent and absolute ASM errors shown in Figure 5.7. The **y**–axis represents the percent deviation of the ASM–generated backorders from the simulation results. The **x**–axis gives the corresponding simulation and ASM run number.

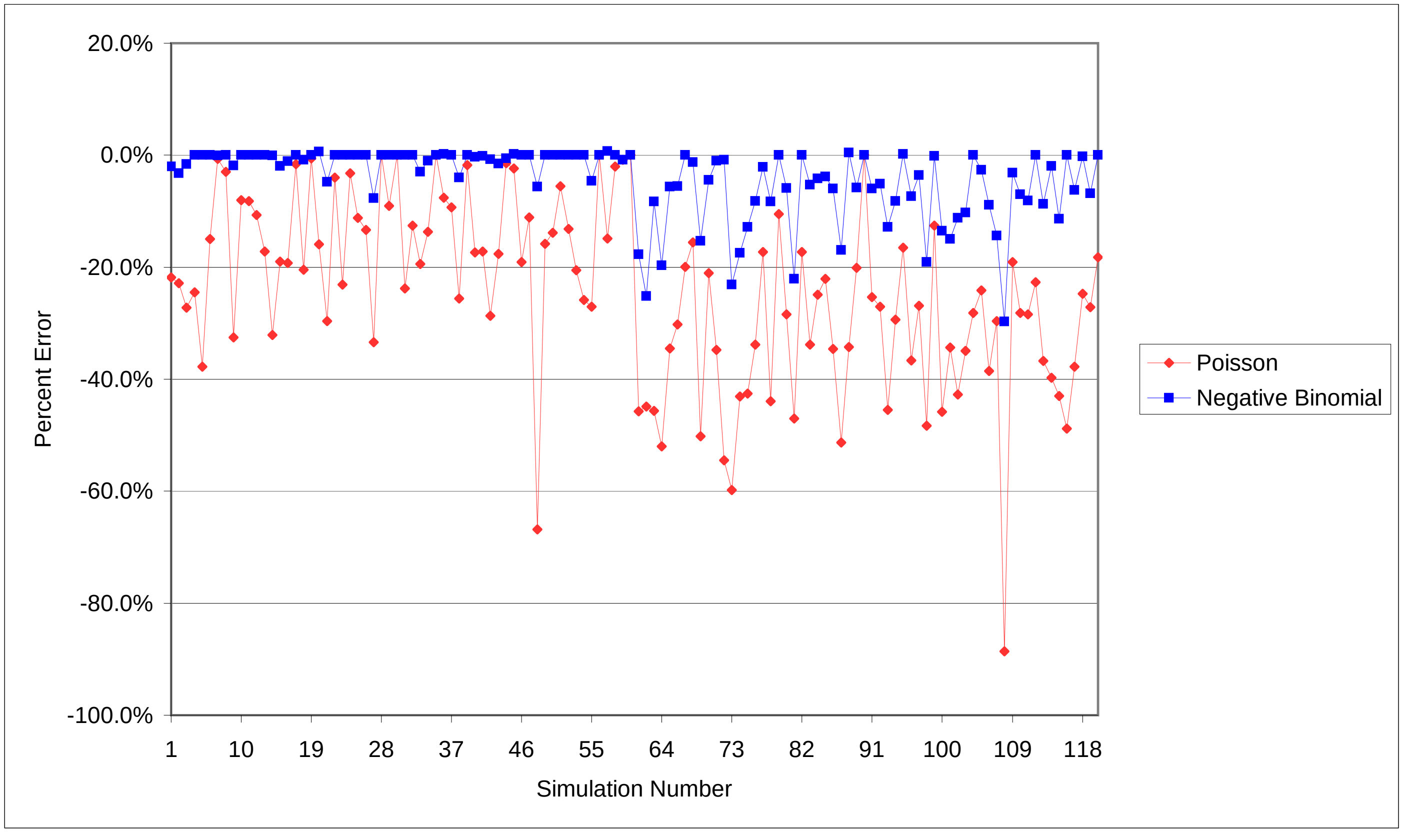

FIG. 5.7. Percent Error of the Poisson and Negative Binomial ASM

Typical mean absolute percent ASM errors may approach 5% to 10% for part backorders that exhibit high variations. The mean percent ASM error for the Poisson model is -24.37%, and we observe that the Poisson ASM exhibits systematic under–prediction of part backorders by 24.37%. The mean absolute ASM error for the Negative Binomial ASM is 4.32%, and the mean percent error for the Negative Binomial ASM is -4.28%. The net result indicates a slight under–prediction of part backorders by 4.28%.

The Poisson ASM does not perform well because the variance to mean ratio for the RSL's expected number of outstanding orders is greater than one for all of the parts we tested. The Negative Binomial ASM uses a two parameter distribution to capture the variability of the outstanding orders, and assumes a variance to mean ratio greater than one. But, the Poisson ASM is unable to capture the variability of the outstanding orders because it uses a single parameter distribution, and it assumes a variance to mean ratio equal to one. For parts 19 and 45 in Figure 5.7, the variance to mean ratio exceeds 2.0 for several RSLs, which is the main contributing factor for the large errors suffered by the Poisson ASM for those parts. We use the Negative Binomial ASM to generate the SCO and ASM comparison results.

5.3 Comparison of the Analytical Stocking Model and Supply Chain Operation Policies

In this section, we compare and evaluate the Negative Binomial ASM and SCO policies using the number of system backorders and total system cost as criteria. We observe the base–stock levels provided by the ASM and SCO policies and compare the policies' corresponding individual fill rates. In Section 5.3.5, we use the ratio of RV to RSL base–stock levels to explain the ASM's superior performance over the SCO policy.

5.3.1 Distribution Resource Planning System Stock Thresholds and Inventory Allocation

Based on forecasted demand rates, DRP calculates the following four stocking thresholds to establish the inventory stock levels in the SLVN: (i) safety stock, (ii) reorder point, (iii) maximum inventory level (MIL), and (iv) MIL plus expected demand over the lead–time, known internally as threshold #4. SCO asset planners and Area Logistics Business Managers (ALBM)s calculate a fifth measure, safety stock (MSS), to manually set stocking levels.

DRP stocks each location using a performance measure called the *piece part fill rate*, i.e., the desired fill rate expressed as a percentage by part and location. DRP does not optimally allocate inventory across the multi–echelon system, and because DRP fails to take the probability of a backorder at the RV into account, delayed replenishment of the RSLs results and the SCO's perception that DRP understocks the RSLs and RV. To compensate, ALBMs and, hence, SCO planners override the DRP – recommended stock thresholds, and manually set MSS levels for many parts and RSLs.

We show in the next several sections that manually setting the MSS levels, regardless of the stocking decisions made for other RSLs and for the RV, results in excess inventory at the RSLs.

5.3.2 Optimal Analytical Stocking Model Policies

We use the expected number of backorders to compare the SCO and ASM policies. Eighty five percent of the SCO's current inventory (including all 151 parts we use) is established by an $(s-1, s)$ policy; the remaining 15% of the SCO inventory (which consists of consumable items) is controlled using a batch ordering policy. The MSS and threshold #4 stock levels calculated by the SCO are used to control the parts using an $(s-1, s)$ policy; thus, they are equivalent to the base–stock levels calculated by the ASM. We use the MSS and threshold #4 stock levels as inputs for (3.2) through (3.12) to calculate the expected number of backorders for the SCO policy.

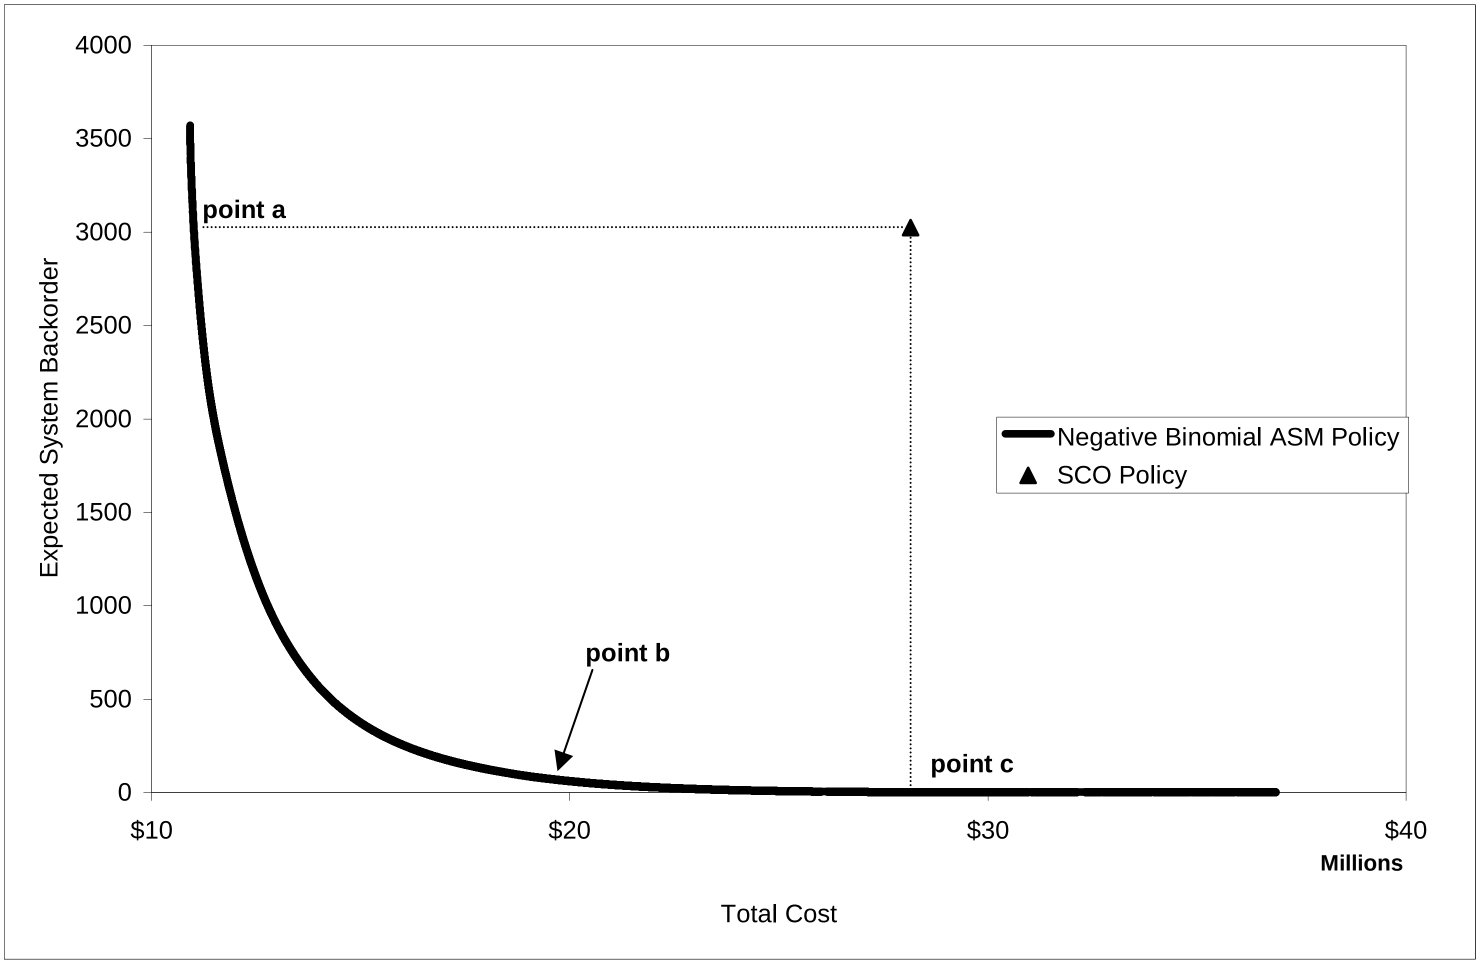

FIG. 5.8. Optimal ASM allocation vs. current SCO allocation

Point on Figure 5.8	Expected Backorder Level (# of parts)	Total Policy Cost
SCO	3022.99	$28,151,971
point a	3022.83	$11,006,190
point b	62.54	$20,003,452
point c	3.05	$28,153,572

Table 5.1. System expected backorders and total costs for the points shown on Figure 5.8

The exchange curve in Figure 5.8 represents the optimal stocking policy for the multi–echelon system, given customer demand and customer service requirements, and the SCO budget. We select three ASM base–stock levels and compare the associated expected number of backorders to the current SCO expected number of backorders. Point **a** of Figure 5.8 is the closest point on the ASM exchange curve that corresponds to the current SCO policy's level of expected backorders. Conversely, point **c** of Figure 5.8 is the closest point on the ASM exchange curve that corresponds to the current SCO policy's total inventory cost. Under ideal conditions, we can move to any point on the optimal exchange curve through reallocation of inventory, e.g., point **b**. Table 5.1 gives the expected backorder levels and total inventory costs for points **a**, **b**, **c**, and the SCO policy of Figure 5.8.

The ASM stocks the RV before the RSLs, i.e., it fills the replenishment pipeline prior to stocking the RSLs. The portion of the curve in Figure 5.8 ranging from an expected backorder level of 3500 to 1500 corresponds to the majority of the stock being placed at the RV, with a small portion of stock placed at the RSLs. The individual RSL expected fill rates for this portion of the curve are, in many cases, less

Point on Figure 5.8	Decrease in Backorder Level	Decrease in Total Cost
point b	97.9%	28.9%
point c	99.9%	0.0%

Table 5.2. Percent decrease in expected backorders and total cost between the SCO policy and points **b** and **c** on the optimal ASM curve shown in Figure 5.8

than 50%. For an expected backorder level below 1500, the majority of the parts are placed at the RSLs.

As the expected system backorders decrease in Figure 5.8, the expected fill rates at the RSLs increase, and the total cost increases. Because each point of the exchange curve corresponds to the optimal base-stock levels at the RSLs and the RV, we should select a point to ensure that the customer service expectations are satisfied while minimizing total cost, i.e., select a point where the marginal cost is greater than or equal to the marginal benefit of stocking more parts at the RSLs. In the following analysis, we compare the current SCO policy and the ASM polices. We do not include ASM point **a** in the analysis because at this point, the expected fill rates at the RSLs are unacceptably low, i.e., less than 70% for most parts.

Stock levels corresponding to points on the curve between points **b** or **c** in Figure 5.8 result in increased fill rates at the RSLs, with less cost than the current SCO policy. Table 5.2 shows the percent decrease in system backorders and the percent decrease in total system cost by moving from the SCO policy to point **b**, or from the SCO policy to point **c** on the optimal exchange curve.

Using the ASM's expected backorder levels, we determine the average time a part spends in inventory as a function of the average time a customer must wait for a part. In general, the expected number of backorders at RSL_i, for part j, β_{ij}, is proportional to the average backorder customer waiting time, W_β, i.e., the average amount of time a customer must wait to receive part j from RSL_i. Zipkin [2000, p. 191–192] shows for general demand and supply processes – where the average backorder level, $\overline{\beta}$, and demand rate, λ, exist – that the average customer waiting time, W_β, can be calculated as follows:

$$W_\beta = \overline{\beta}/\lambda. \tag{5.5}$$

This result is analogous to (3.3). Zipkin states that if the average inventory, $\overline{I}$, exists, then the average time a part spends in inventory, W_I, can be found as follows:

$$W_I = \overline{I}/\lambda. \tag{5.6}$$

Using (3.12) and (3.13) we calculate the expected number of backorders and expected inventory for each part and each RSL (Appendices C.1 and C.2). We then calculate the average customer waiting time and the average time a part spends in inventory with (5.5) and (5.6), respectively, for each part (Appendix C.3). The average time a part spends in inventory and the customer waiting time is plotted using a linear–log scale in Figure 5.9.

Figure 5.9 shows that the average customer waiting times for point **b**, point **c**, and the SCO policy are less than 24 hours. We observe that as the customer waiting time is reduced, the average inventory stock time is increased, and so, as a

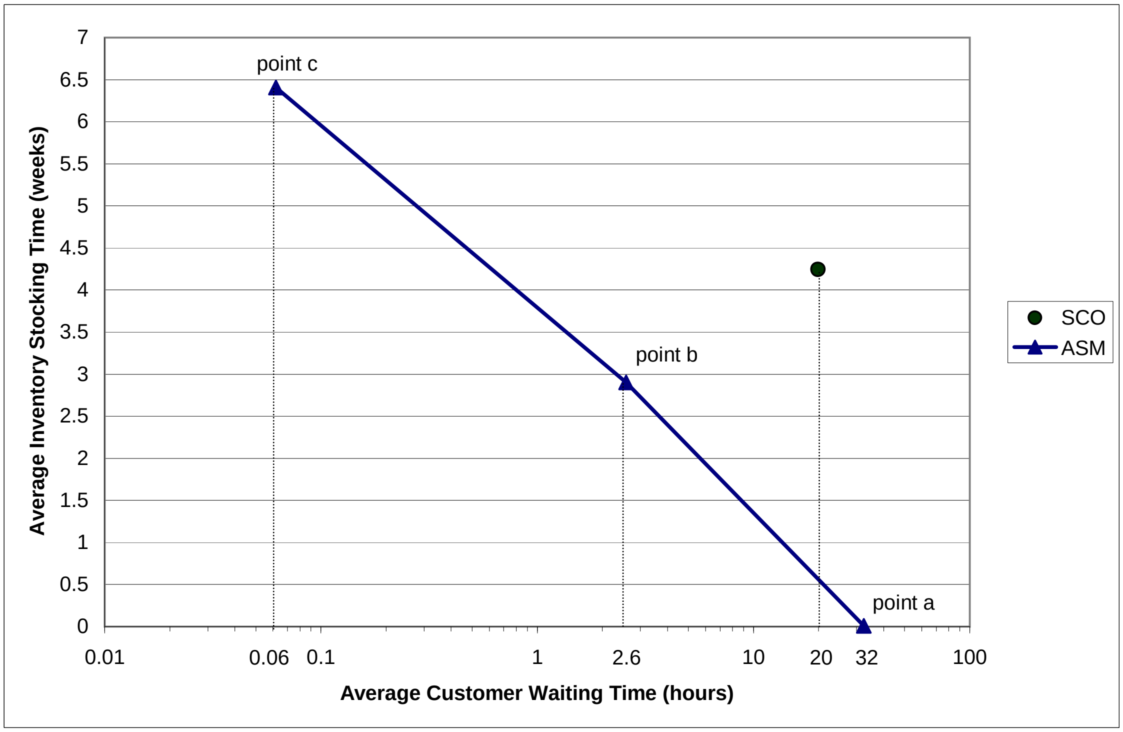

Fig. 5.9. The average time a part spends in inventory versus the average time a customer waits for a part (linear / log scale)

result, are inventory holding costs. For example, point **c** offers the shortest customer waiting time, 0.06 hours, of the three points, and the corresponding average inventory stocking time is 6.4 weeks. At point **b**, the average customer waiting time increases to 2.6 hours, but the average inventory stocking time decreases to less than 3 weeks.

5.3.3 Individual and System Expected Fill Rates

Although we can calculate the expected number of backorders and expected customer waiting times for each stocking location and part, Sun's SCO requires the expected fill rate for each part, and, specifically, the *system fill rate* to be greater than 90%. We calculate the system fill rate by averaging the expected fill rates across all RSLs and parts. To achieve Sun's system fill rate requirement, the target fill rates for each part and location pair are set above 96% for most cases.

We calculate the expected fill rate at the RV and each RSL using (3.5) and (3.14), respectively (Appendix C.4). Figure 5.10 shows the average expected fill rates corresponding to two ASM stock levels, i.e., points **b** and **c**, and the SCO policy across all 100 RSLs, as a function of the 151 parts.

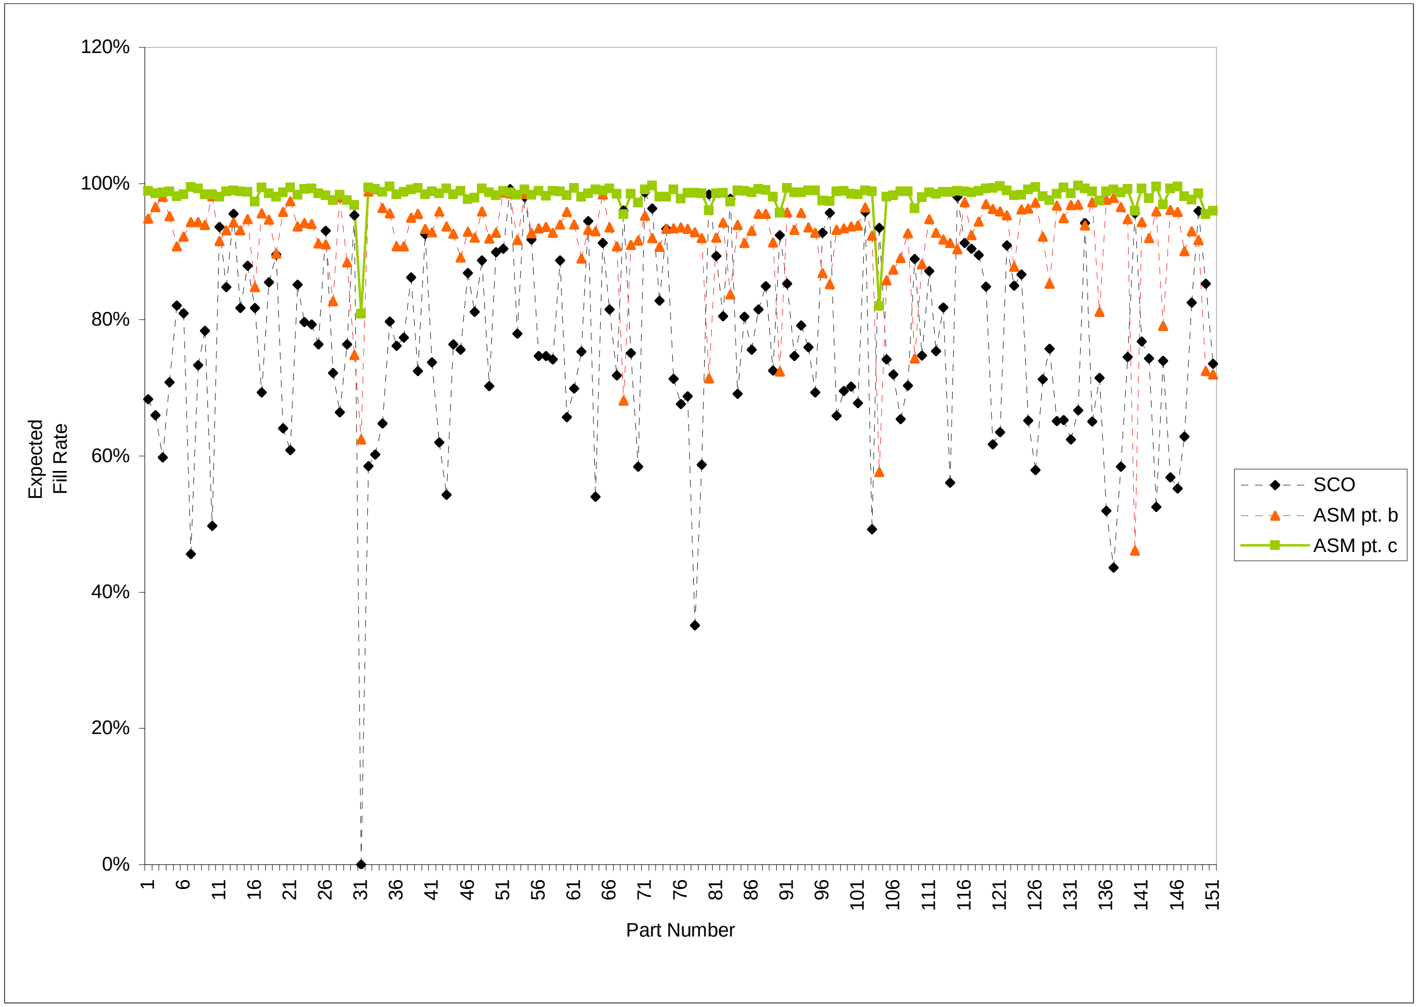

FIG. 5.10. Individual part fill rate comparisons between ASM and SCO policies

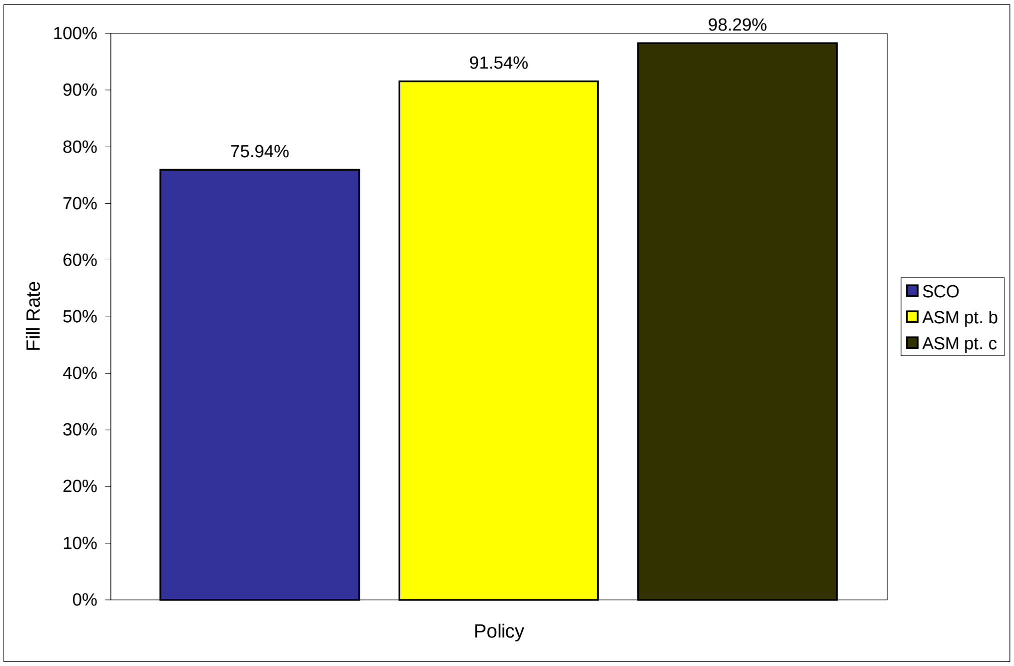

FIG. 5.11. Average total fill rate comparisons between ASM and SCO policies

Figure 5.10 shows that the ASM policy corresponding to point **c** has higher fill rates for all parts than the SCO policy. Furthermore, although the SCO's total inventory costs are 29% greater than those corresponding to the ASM policy at point **b**, the ASM policy fill rates are higher than the fill rates for most parts using the SCO policy.

We aggregate the individual expected fill rates and calculate the system expected fill rate for the ASM and the SCO policies (Figure 5.11). The expected ASM system fill rates corresponding to points **b** and **c** are above 90%. The SCO policy has an system expected fill rate of approximately 76%.

In practice, the typical system fill rate for the SCO policy is greater than 90%.

Sun's SCO is able to achieve higher system fill rates through the use of inventory fair–sharing. Inventory fair–sharing is the transshipment of inventory from one or more RSLs to another. It is used to "level" the on–hand inventory among all of the RSLs within the system, usually on a daily basis. However, inventory fair–sharing policies are not costless. We do not consider the transportation costs and the excess inventory holding costs at the RSLs associated with the SCO's inventory fair–sharing policy. Therefore, we do not compare ASM and SCO policy results that include inventory fair–sharing; the ASM policies achieve high individual fill rates at the RSLs without incurring the additional fair–sharing costs.

5.3.4 Analytical Stocking Model and Supply Chain Operation Base Stock Levels

We now examine the base–stock levels at the RV for the SCO and ASM policies at points **b** and **c**. Figure 5.12 shows the RV base–stock levels for the 151 parts (Appendix D.1). In general, the ASM policies have higher stocking levels than the SCO policy at the RV.

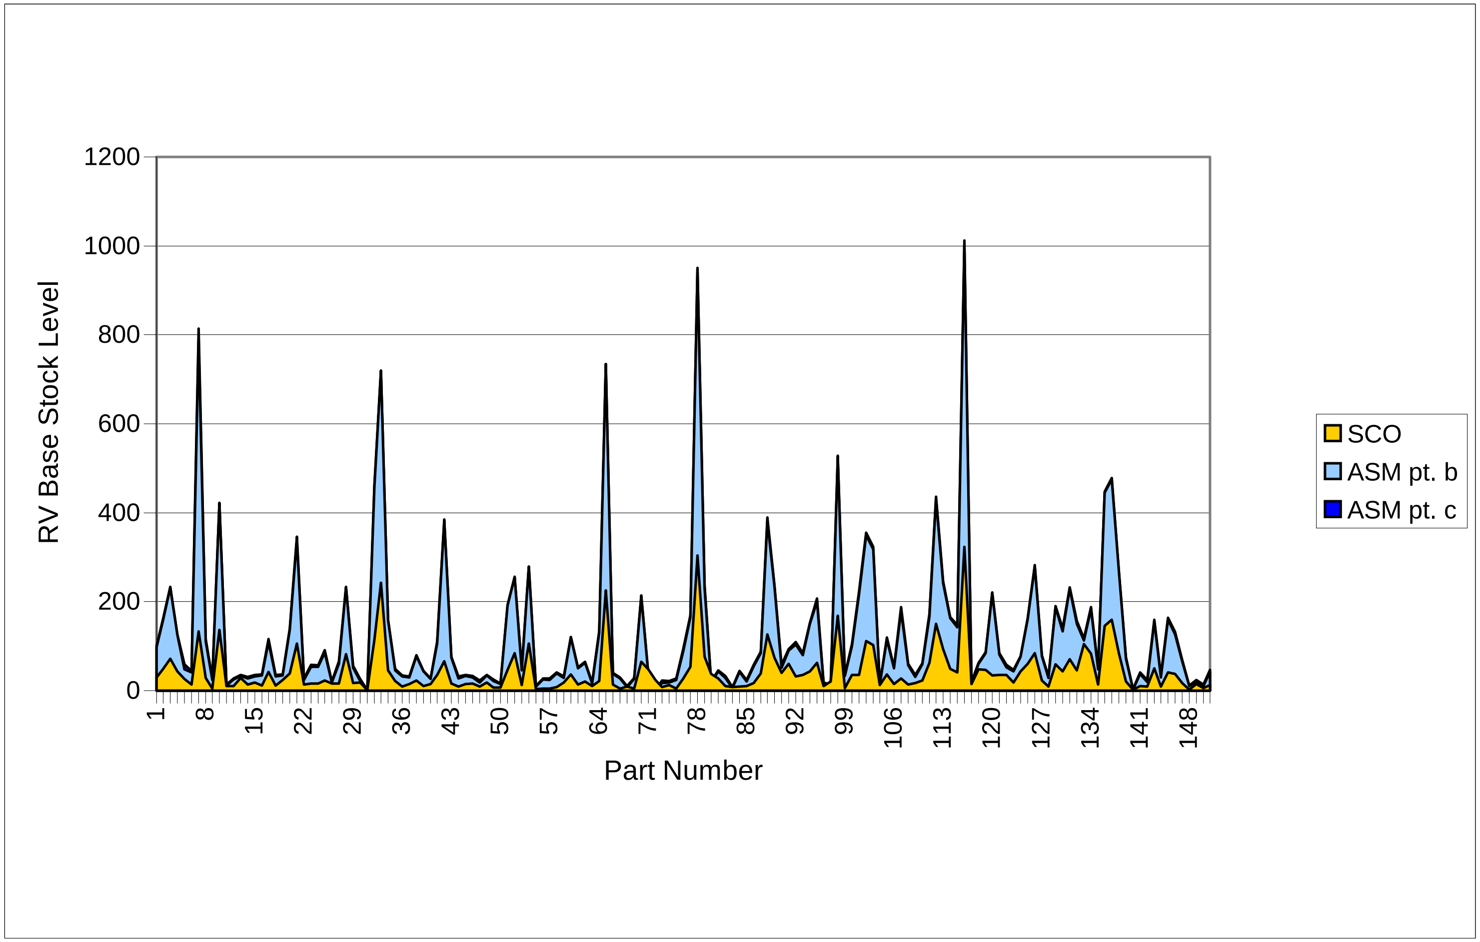

FIG. 5.12. Base–stock levels at the RV for the ASM and SCO policies

Figure 5.13 compares the SCO and ASM base–stock levels at the RSLs for all 151 parts (Appendix D.2). In general, the base–stock levels for ASM at point **c** are slightly lower than most of the SCO's RSL base–stock levels. There are several cases illustrated in Figure 5.13 in which the SCO's base–stock greatly exceeds the ASM's base–stock levels. The ASM does not stock parts for a particular RSL when no forecasted demand exists. Sun allows inventory planners to ignore the DRP–forecasted demand and manually set MSS levels, and in some cases place stock where none exists or is forecasted. Setting MSS levels in this manner will not necessarily increase expected system availability, but it is guaranteed to increase total system inventory costs.

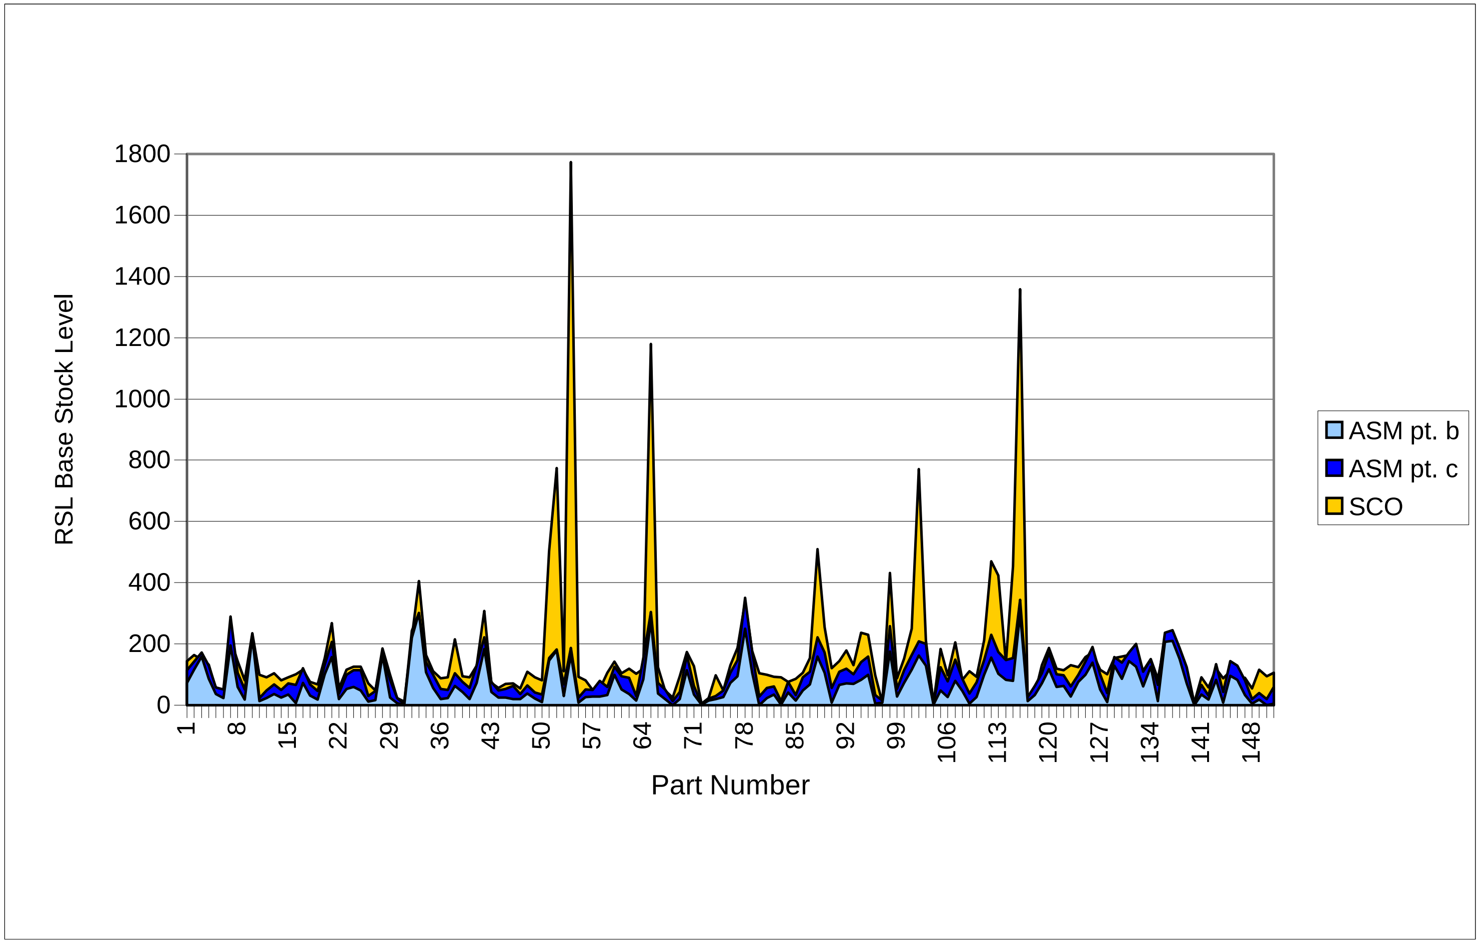

FIG. 5.13. Base–stock levels at the RSLs for the ASM and SCO policies

Figures 5.12 and 5.13 show that SCO and the ASM policies share the same basic trend in stocking levels. Therefore, we conclude that the SCO relies significantly on the forecast obtained from DRP to establish MSS and threshold #4 levels, but the SCO does not optimally allocate the base–stock for this multi–echelon system.

Figure 5.12 illustrates that the ASM places more base–stock at the RV than the SCO policy, while Figure 5.13 shows that the SCO policy places more base–stock at the RSLs than the ASM policy. To investigate the relative placement of stock, we evaluate the ratio of stock at the RV to the stock at all the RSLs.

5.3.5 Ratio of Repair Vendor Base Stock to Remote Stocking Location Base Stock

By comparing the ratio of RV to RSL base–stock, we highlight the primary inventory allocation differences between the SCO and ASM policies. The ratio of RV to RSL base–stock is found by dividing the RV Base Stock for a part by the sum of the base–stock at all of the RSLs for the same part. Figure 5.14 compares, for each part, the RV to RSL base–stock ratio for the SCO policy and the ASM policy corresponding to point **c**. The RV to RSL base–stock ratios are aggregated and shown in Figure 5.15.

The ASM achieves much higher expected service levels, for the same cost, than the SCO policy by accounting for the *risk–pooling* ability of the RV. In other words, the RV "pools" the risk caused by RSL demand uncertainties, and can fill an order from any RSL in addition to providing direct customer support. Therefore, RV stock allows the system to enjoy the advantages of inventory consolidation while still maintaining

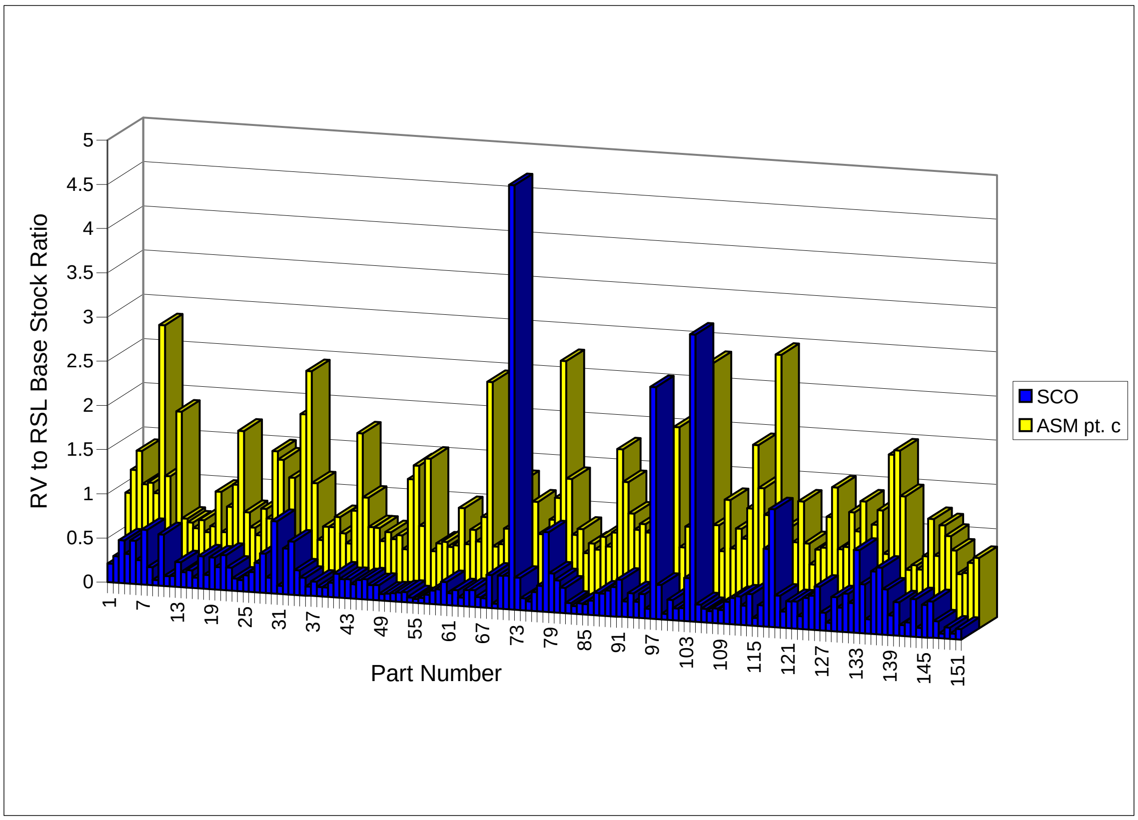

Fig. 5.14. Ratio of the total RV demand versus the RSL demand for the ASM and SCO policies

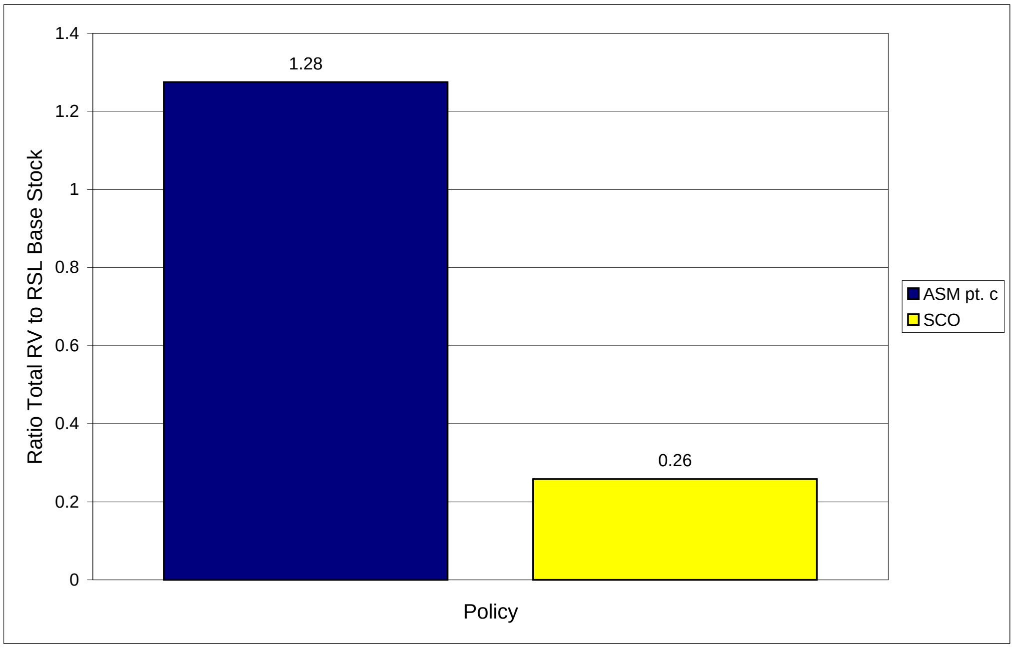

FIG. 5.15. Ratio of RV to RSL aggregate demand

some stock at multiple RSL locations. This RV risk–pooling effect is greater when the RV is very responsive, e.g., two or three days, as is the case with the SLVN in the United States.

The results shown in Figures 5.14 and 5.15 confirm the trends shown in Figures 5.12 and 5.13. The ASM allocates a much greater proportion of its base–stock to the RV than the SCO policy, which results in reduced inventory costs and increased customer level of service. In the next section we observe the effects of risk–pooling at the RSLs by considering an ELT case.

5.4 Emergency Lateral Transshipments

ELT is used to transship repair parts from an RSL to a customer, if the customer is unable to receive a part from his assigned RSL because of a stock–out (Figure 2.1). We mention in Section 2.6 that ELTs are useful in some situations to take advantage of the risk–pooling property of a group of RSLs, which is commonly referred to as a *risk–pooling group*. Lee [1987], Axsäter [1990], and Sherbrooke [1992] suggest that the use of ELT may reduce the amount of base–stock necessary at the RSLs if the average RSL replenishment lead–time is greater than the average transshipment time between RSLs.

Groover et al. [1987] show that multi–echelon models can be used to provide upper and lower bounds on the base–stock levels using ELT. Using Groover's technique, we modify the original data set and use the ASM to produce the ELT lower bound on the base–stock levels. We compare the ASM results obtained in Section 5.3.2 with the ELT results using the same forecasted demand rates at the RV and RSLs.

We select seven RSLs that generate approximately 30% to 40% of the demand on the SVLN. Because of the high demand rates in this area and the close geographical proximity of the RSLs to one another, these RSLs are selected as a "super" RSL. We call this super RSL a *pooling group*, and assume that the transshipment time between RSLs within this pooling group is zero. This assumption is not unreasonable because actual transshipment times between the RSLs range from one to six hours, i.e., much less than the average RSL replenishment lead–times of two or three days. The ASM determines the lower bound of the system backorder levels, i.e, the "best case", by

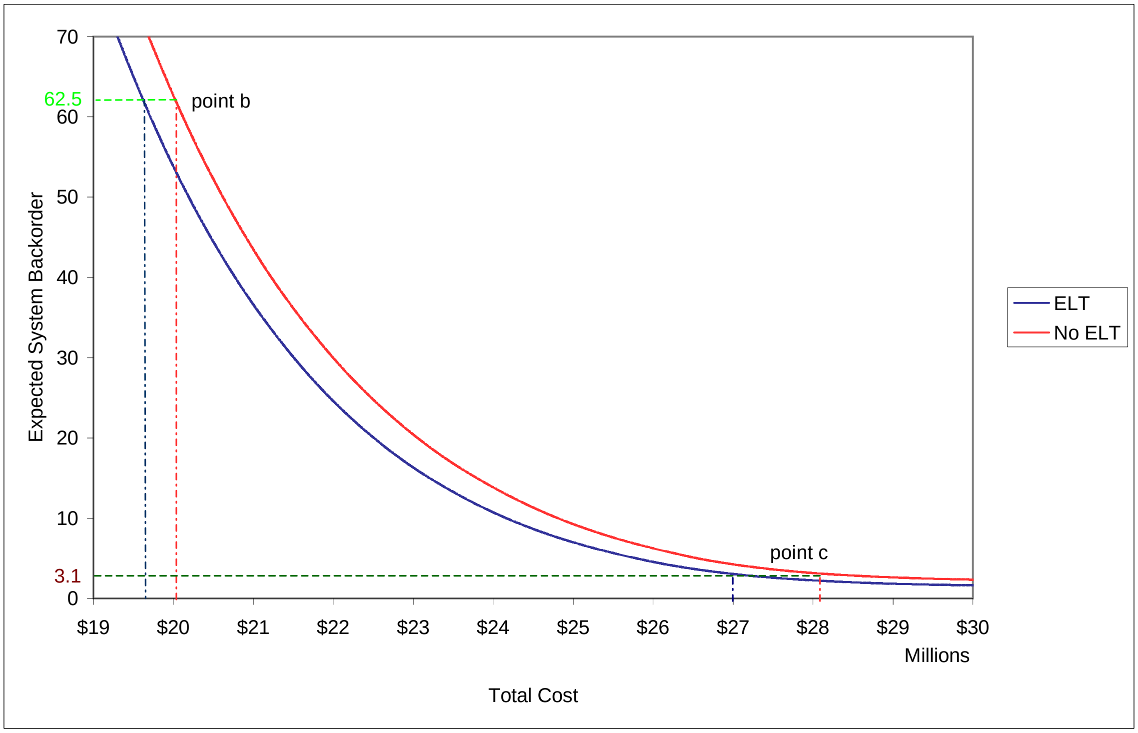

FIG. 5.16. Comparison of ASM using ELT and standard ASM with no ELT

calculating the optimal exchange curve with 94 RSLs – the pooling group and the remaining 93 RSLs in the SVLN.

Figure 5.16 shows two exchange curves representing the ELT case and the standard ASM assuming no ELT. Points **b** and **c** in Figure 5.16 are the same points selected in Section 5.3.2. Table 5.3 tabulates the percent reduction from the standard ASM case in both total cost and total inventory at the RSLs for both points **b** and **c**.

If ELT is employed for the pooling group, the maximum obtainable reduction in base–stock is 5.1% and 4.1% for points **b** and **c**, respectively, corresponding to cost

ASM Policy	Total RSL Base Stock		Total Cost	
	point b	point c	point b	point c
No ELT	10284	15815	\$20,003,452	\$28,153,572
ELT	9703	15003	\$19,602,591	\$26,993,846
% Reduction	5.6%	5.1%	2.0%	4.1%

Table 5.3. Percent reduction in total RSL base–stock and total cost from standard ASM model and ASM model incorporating ELT

reductions of less than 2.0% and 4.1%, respectively.

In practice, more realistic restrictions would be added, e.g., stock sourcing rules (Section 2.6), to the ELT policy. Although, based on geography, three feasible ELT pooling groups may exist in the United States, additional cost reductions of more than 3% to 4% are not likely using ELT, because the combined RSL demand for the three pooling groups is approximately 60% of the total RSL demand. Therefore, it is likely that the costs of implementation and operation of ELTs for the SVLN in the United States will outweigh the benefits. This is due in part to the responsiveness of the RV and to the short transportation time from the RV to the RSLs, i.e., the average RSL replenishment lead–time of two to three days. In the next section, we increase the RSL lead–times and compare non-ELT and ELT cases.

5.5 Extended Remote Stocking Location Lead–Times

We now use the ASM to investigate the case in which the RSL lead–time, L_i, is increased dramatically, e.g., to nine days. This case may exist in Asian or European countries because of increased distances between the distribution centers, or

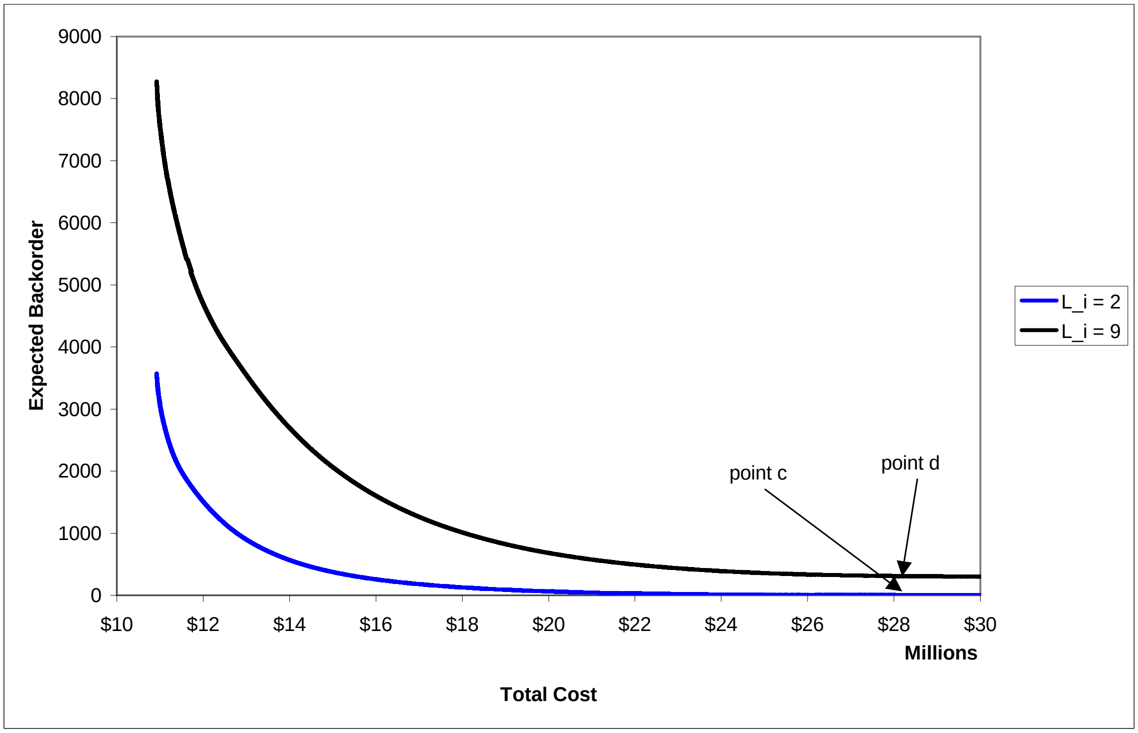

FIG. 5.17. Comparison of ASM exchange curve for RSL lead-times, L_i, of two and nine days

RV, and the RSLs, customs regulations, and lack of express delivery systems in certain geographical regions. We begin by examining the optimal ASM exchange curve corresponding to an RSL lead–time of nine days, and we superimpose the exchange curve (from Figure 5.8) corresponding to an RSL lead–time of two days on Figure 5.17.

Point **d** of Figure 5.17 possesses approximately the same cost as point **c**, i.e., \$28,153,480. The exchange curve corresponding to $L_i = 9$ is shifted up and left from the curve corresponding to $L_i = 2$. Because of the shift, the expected number of system backorders increases from 3.1 (point **c**) to 320 (point **d**). Hariharan and Zipkin [1995], Zipkin [2000, p. 183] provide similar results. As a consequence of the increased RSL replenishment lead–time, the RV risk–pooling effects are mitigated, and the inventory tends to migrate to the RSLs from the RV to reduce the amount of time a customer must wait for a part. To illustrate the shift of stock from the RV to the RSLs, we calculate the RV to RSL stock ratio for point **d**. The stock ratio corresponding to point **d** is 1.06, which is lower than 1.28, the value obtained for point **c** in Section 5.3.5, which indicates (at least for our example) that as the RSL lead–time increases, it is more efficient to increase the inventory held at the RSLs than to rely upon the risk–pooling effect of the RV.

For an ELT with an extended RSL lead–time of nine days, we use the pooling group data from Section 5.4 and generate the exchange curve in Figure 5.18. Reductions in the total RSL base–stock of 5.6% and total cost of 8.0% are possible using ELT with the extended RSL lead–times. These reductions are greater than those observed in Section 5.4 with a two–day replenishment lead-time. As this example illustrates, it may be practical to implement ELT as the RSL lead-times increase.

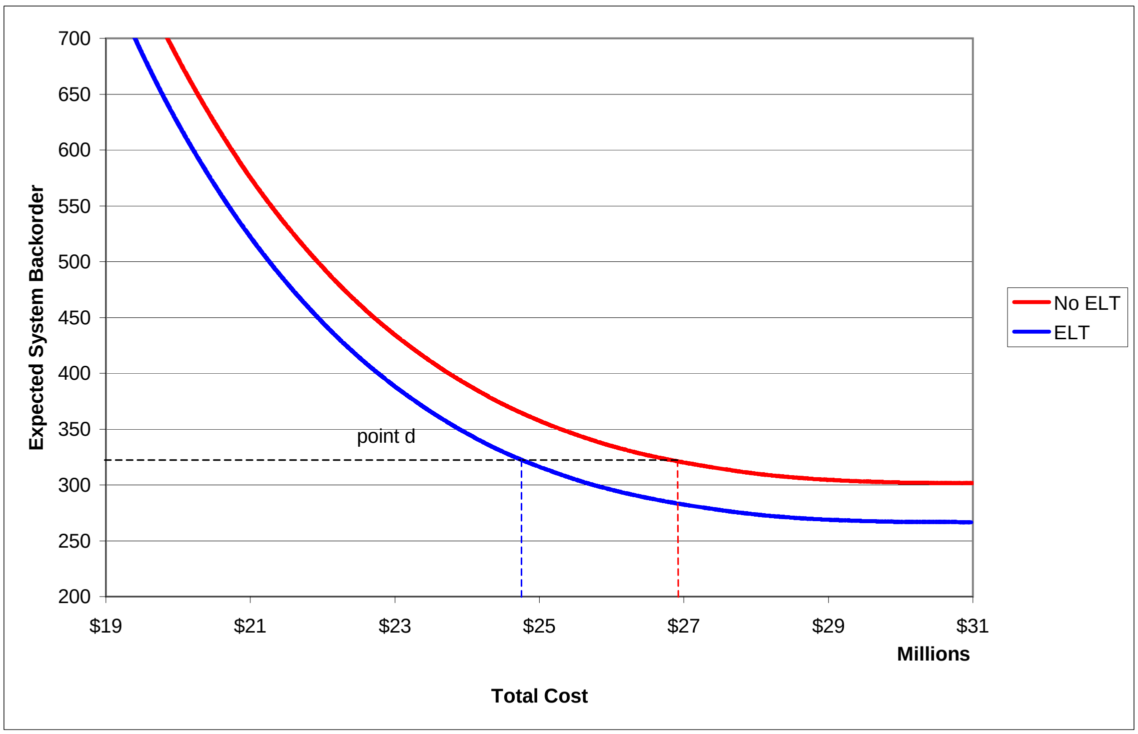

Fig. 5.18. Comparison of ASM using ELT and standard ASM with extended RSL replenishment lead-time of nine days

Chapter 6

CONCLUSIONS AND FUTURE RESEARCH

6.1 Summary and Conclusions

We investigate a large scale repairable parts logistics system with highly variable demands, and determine the optimal allocation of high demand value parts at the RV and the RSLs. Though we can scale our analytical model to allocate stock for all of the SVLN's repairable parts, we use results from a Pareto analysis to focus our efforts on the repairable parts that constitute the majority of the inventory costs. From the Pareto analysis (Figure 5.1), we show that 95% of the cumulative fraction of high demand value parts comprises less than 20% of the parts in the Sun Virtual Logistics Network.

We develop an Analytical Stocking Model (ASM), using two different distributions to fit the backorder function at the RSLs, to characterize the logistics network. As an extension of Graves' negative binomial model that allows for two-classes of service, the ASM determines the optimal allocations of inventory by minimizing expected system backorders subject to a system budget constraint. Our optimization model shows that we can increase system customer level of service while decreasing overall system cost.

The ASM is an approximation using Little's formula. Sherbrooke and Graves show through empirical studies that this approximation is acceptable for parts that are characterized by low customer demand. Our data, however, is characterized by erratic customer demand and high unit part costs. Therefore, we develop a dynamic simulation model to determine how well each version of the ASM model behaves assuming erratic customer demand. The ASM that uses the Negative Binomial distribution to approximate the backorder function outperforms the Poisson ASM. The mean percent error between the simulation and Negative Binomial ASM is -4.28%, while the Poisson ASM error is -24.37%.

Using weekly demand forecasts for the high demand value parts identified by Pareto analysis, the ASM generates an exchange curve depicting system backorders versus total inventory cost (Figure 5.8). From the exchange curve, we determine a reasonable operational range by weighing the marginal benefit of decreasing backorders against increasing total inventory cost. Once we select the optimal cost–benefit ratio from the exchange curve, the associated base–stock levels at the RV and at the RSLs are determined and the expected fill rates and inventory levels at the RV and RSLs are calculated.

We compare the ASM stock allocations for three points selected from the exchange curve and the current SCO policy. An upper limit on cost reductions of 60% is possible for our data set keeping the level of service constant. However, the ASM results show that the expected fill rate for each location and part, calculated by the ASM for this operating range, is too low to be used in practice. Keeping total inventory costs constant, we show that the ASM can reduce system backorders by 97%. A likely operational range – between points **b** and **c** in Figure 5.8 – results in cost reductions of 30% to 40% and an increase in customer service of over 90%.

Figures 5.10 and 5.11 show that the ASM policy results in significant increases in expected fill rates over the SCO policy. We expect this because the ASM uses a global optimization algorithm based on the minimization of system expected backorders and the SCO's inventory allocation satisfies a myopic fill–rate target level. We examine base–stock levels at the RSLs and the RV to help explain why the customer level of service for the ASM policy is, in general, better than the SCO policy.

Figures 5.12, 5.13, and 5.15 suggest that much of ASM's gains in customer service are possible because the current SCO policy understocks many parts at the RV, which diminishes the "risk pooling" effects of centralizing inventory at the higher echelon. This, in turn, causes higher–than–necessary backorders at the RV, resulting in the down–stream effect of greater RSL lead–time variability, and increased expected backorders, or lower expected fill rates. We also consider two additional cases, emergency lateral transshipment (ELT) and an increase in the RSL lead–time.

For the given data set and a single risk–pooling group, we demonstrate that the lower bound for cost savings using ELT is between two and four percent. Because the RVs are very responsive and the RSL lead–times are relatively short for the SLVN, using ELT in this case may not be beneficial. Using the same demand data and risk pooling group, and RSL lead–times extended to nine days, ELT provides a lower bound for cost savings of up to eight percent. Therefore, ELT may reduce the cost of inventory for cases in which RSL lead–times are longer than those of the typical U.S. lead–times, or where multiple risk pooling groups are established.

We conclude that material asset allocation should be performed using a system approach. The practice used widely in industry today is myopic, i.e., parts are allocated by echelon, or allocation is determined on a part and location basis independent of all other parts in the system. Thus, current practice results in sub–optimal performance and excess inventory costs for most multi-echelon systems. The fill–rate service criterion used by Sun may exacerbate this problem by causing inventory migration from the RV to the RSLs, reducing the possible risk pooling benefits of holding inventory at the RVs. Because of the stochastic demands placed on the RV and RSLs, there is a tendency to arbitrarily increase target stock levels without regard to the down–stream effects. Using a model based on the system backorder service criterion, we show that it is possible to increase individual fill rates and decrease system inventory costs through the judicious allocation of inventory between the RV and RSLs.

We have shown how the results, models, and conclusions from this dissertation may be applied to the computer industry. In general, the techniques and results of this dissertation may be applicable to many other industries and organizations with large scale multi-echelon service parts logistics systems.

6.2 Future Research

We implemented a simulation to conduct initial tests of the ASM. The following additional research would provide further insights into the robustness of the ASM:

- Simulate a larger sample of parts.
- Evaluate the effects that changes in RV lead–time and part return delays have on ASM performance.

- Conduct a pilot program to measure the ASM effectiveness under actual conditions.

Because the probability distributions for the number of expected outstanding orders are not calculated explicitly, the ASM may be implemented for many industries that incorporate $(s-1, s)$ policies, regardless of the number of echelons, parts, or stocking locations in the corresponding logistics system.

The enterprise computer service industry, and enterprise services in general, use a standard measure of performance model to evaluate customer levels of service, e.g., fill–rates. However, many customers today are concerned with the availability of their critical systems, not part fill–rates or backorders of spare parts. An interesting extension of this model is the maximization of supply availability subject to a system budget constraint. The computational aspects of a supply availability model would be challenging to implement.

We have shown how to use the ASM to establish bounds on backorders and stocking levels for the ELT case. Another possible extension is to combine the ASM with a simulation model to tighten the bounds for a service parts logistics system that uses inventory fair–sharing, or to implement realistic ELT stock sourcing rules that may be useful in practice.

REFERENCES CITED

S. Axsäter. Modeling Emergency Lateral Transshipments in Inventory Systems. *Management Science*, **36**(11):1329–1338, 1990.

M.A. Cohen, P. Kamesam, P. Kleindorfer, H. Lee, and A. Tekerian. Optimizer: IBM's Multi–Echelon Inventory System for Managing Service Logistics. *Interfaces*, **20**(1):65–82, 1990.

M.A. Cohen, Y. Zheng, and Y. Wang. Identifying Opportunities for Improving Teradyne's Service–Parts Logistics System. *Interfaces*, **29**(4):1–18, 1999.

G.J. Feeney and C.C. Sherbrooke. The $(s-1, s)$ Inventory Policy Under Compound Poisson Demand. *Management Science*, **12**(5):391–411, 1966.

S.C. Graves. A Multi-Echelon Inventory Model for a Repairable Item with One–for–One Replenishment. *Management Science*, **31**(5):1247–1256, 1985.

C.W. Groover, S.J. Culosi, and L.L. Woods. The Logistics Information System Analysis (LISA). In *Analysis of the Stock Control and Distribution System*, volume **2**. Systems Research and Applications Corp., Arlington, V.A, 1987.

R. Hariharan and P. Zipkin. Customer-order Information, Leadtimes, and Inventories. *Management Science*, **41**(10):1599–1607, 1995.

W.D. Kelton, R.P. Sadowski, and D.A. Sadowski. *Simulation with Arena*. McGraw–Hill, Boston, Massachusetts, 2nd edition, 2002.

A.M. Law and J.S. Carson. A Sequential Procedure for Determining the Length of a Steady–State Simulation. *Operations Research*, **27**(5):1011–1025, 1979.

A.M. Law and W.D. Kelton. *Simulation and Modeling and Analysis*. McGraw–Hill, Boston, Massachusetts, 3rd edition, 2000.

H.L. Lee. A Multi–Echelon Inventory Model for Repairable Items with Emergency Lateral Transshipments. *Management Science*, **33**(10):1302–1316, 1987.

J. Little. A Proof of the Queueing Formula $L = \lambda W$. *Operations Research*, **9**(3): 383–387, 1961.

J.A. Muckstadt and L.J. Thomas. Are Multi–Echelon Inventory Methods Worth Implementing in Systems with Low–Demand–Rate Items? *Management Science*, **26**(5):483–494, 1980.

Steven Nahmias. Managing Reparable Item Inventory Systems: A Review. In L.B. Schwarz, editor, *TIMS Studies In The Management Sciences*. 1981.

C. Palm. Analysis of the Erlang Traffic Formula for Busy–Signal Arrangements. *Ericsson Technics*, **5**:39–58, 1938.

Sheldon M. Ross. *Introduction to Probability Models*. Harcourt/Academic Press, San Diego, California, 7th edition, 2000.

H.E. Scarf. Stationary Operating Characteristics of an Inventory Model with Time Lag. In K.J. Arrow, S. Karlin, and H.E. Scarf, editors, *Studies in the Mathematical Theory of Inventory and Production*, chapter 16. Stanford University Press, Stanford California, 1958.

K. Shanker. Exact Analysis of a Two–Echelon Inventory System for Recoverable Items Under Batch Inspection Policy. *Naval Research Logistics Quarterly*, **28**(4): 579–601, 1981.

Craig C. Sherbrooke. METRIC: A Multi–Echelon Technique For Recoverable Item Control. *Operations Research*, **16**(1):122–141, 1968.

Craig C. Sherbrooke. VARI–METRIC: Improved Approximations For Multi–Indenture, Multi–Echelon Availability Models. *Operations Research*, **34**(2):331–319, 1986.

Craig C. Sherbrooke. Optimal Inventory Modeling of Systems. In *Wiley–Interscience*. Wiley, New York, 1992.

R.M. Simon. Stationary Properties of a Two Echelon Inventory Model for Low Demand Items. *Operations Research*, **19**(3):761–777, 1971.

Sun Microsystems Inc. Annual Report, 2001.

P.D. Welch. The Statistical Analysis of Simulation Results. In S.S. Lavenberg, editor, *The Computer Performance Modeling Handbook*. Academic Press, New York, 1983.

P.H. Zipkin. *Foundations of Inventory Management*. McGraw–Hill, Boston, Massachusetts, 2000.

APPENDIX A

-CD-ROM INSTRUCTIONS

The contents of the dissertation and appendices can be found on the CD-ROM labeled: " A SYSTEMS APPROACH TO OPTIMIZE A MULTI–ECHELON REPAIRABLE ITEM INVENTORY SYSTEM WITH MULTIPLE CLASSES OF SERVICE". The contents of this dissertation is in a file named **diss.pdf** and can be read using *Adobe Acrobat Reader*. The appendixes can be read with *Microsoft Excel 2000*.

APPENDIX B

-WEEKLY DEMAND DATA

This appendix contains the data used to obtain the results for this dissertation. The appendix can be found on the Excel file named: **appendixB.xls**.

B.1 Remote Stocking Location Demand data

Appendix B, Tab B.1 found on file **appendixB.xls** on worksheet TabB.1

B.2 Repair Vendor Demand data

Appendix B, Tab B.2 found on file **appendixB.xls** on worksheet TabB.2

APPENDIX C

-CALCULATED MEASURES OF PERFORMANCE

This appendix contains the resulting measures of performance calculated by the ASM. The appendix can be found on the Excel file named: **appendixC.xls**.

C.1 Expected Backorders

Appendix C, Tab C.1 found on file **appendixC.xls** on worksheet TabC.1

C.2 Expected Inventory

Appendix C, Tab C.2 found on file **appendixC.xls** on worksheet TabC.2

C.3 Average Customer Waiting Time and Stocking Time

Appendix C, Tab C.3 found on file **appendixC.xls** on worksheet TabC.3

C.4 Expected Fill Rates

Appendix C, Tab C.4 found on file **appendixC.xls** on worksheet TabC.4

APPENDIX D

-BASE STOCK LEVELS

This appendix contains the Base Stock Levels for all parts and locations. The appendix can be found on the Excel file named: **appendixD.xls**.

D.1 Repair Vendor Base Stock

Appendix D, Tab D.1 found on file **appendixD.xls** on worksheet TabD.1

D.2 Remote Stocking Location Base Stock

Appendix D, Tab D.2 found on file **appendixD.xls** on worksheet TabD.2