

1996

Performance Analysis of Multi-Echelon Inventory Systems.

David Lee Peak

Louisiana State University and Agricultural & Mechanical College

Follow this and additional works at: https://digitalcommons.lsu.edu/gradschool_disstheses

Recommended Citation

Peak, David Lee, "Performance Analysis of Multi-Echelon Inventory Systems." (1996). *LSU Historical Dissertations and Theses*. 6363.

https://digitalcommons.lsu.edu/gradschool_disstheses/6363

This Dissertation is brought to you for free and open access by the Graduate School at LSU Digital Commons. It has been accepted for inclusion in LSU Historical Dissertations and Theses by an authorized administrator of LSU Digital Commons. For more information, please contact gradetd@lsu.edu.

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

UMI

A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor MI 48106-1346 USA
313/761-4700 800/521-0600

**PERFORMANCE ANALYSIS OF
MULTI-ECHELON INVENTORY SYSTEMS**

A Dissertation

**Submitted to the Graduate Faculty of the
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy**

in

Business Administration

by

David Lee Peak

B.S., Louisiana State University at Shreveport, 1989

M.S., Louisiana State University, 1991

December 1996

UMI Number: 9720377

UMI Microform 9720377
Copyright 1997, by UMI Company. All rights reserved.
This microform edition is protected against unauthorized
copying under Title 17, United States Code.

UMI
300 North Zeeb Road
Ann Arbor, MI 48103

ACKNOWLEDGMENTS

I would like to express my gratitude to the members of my dissertation committee, Professors Helmut Schneider, Peter Kelle, Linguo Gong, and particularly, Edward Watson, without whose time, support, comments, suggestions, guidance, and threats, this endeavor may never have been completed. I consider myself particularly fortunate to have been associated with this group of scholars throughout my graduate experience.

Special thanks go to my family. My wife, Gina, and my daughter, Erin, have put up with me (and without me) throughout these graduate school years. They are truly the reason it has all been worth the effort, and I look forward to getting to know them both again.

TABLE OF CONTENTS

	page
ACKNOWLEDGMENTS	ii
LIST OF TABLES	v
LIST OF FIGURES	vi
ABSTRACT	vii
 1. INTRODUCTION	 1
1.1. Contribution of the Research	4
1.2. Organization of the Research	6
 2. REVIEW OF THE LITERATURE	 7
2.1. Classical Multi-Echelon Inventory Models	8
2.2. Single Period Models	10
2.3. Models Where the Warehouse Holds No Stock	12
2.4. Models Where the Warehouse Holds Stock	14
2.5. Systems Using Lot Size/Reorder Point Models	19
2.6. Systems Using Periodic Review Order-up-to-S Models	21
2.7. Systems Using Continuous Review Order-up-to-S Models	23
 3. DEVELOPMENT OF INVENTORY STRATEGIES	 25
3.1. Multi-Store Inventory Model	27
3.2. Inventory Model for Multi-Store System with Central Warehouse	32
3.3. Approximations for the Warehouse/Stores System	37
3.4. Approximate (s,S) Policies Using an α -Service Level	38
3.5. Consideration of Aggregate Measures and Restrictions	43
3.5.1. Aggregate Measures at the Store Level.	44
3.5.1.1. Capital Investment Limit at the Store Level.	45
3.5.1.2. Storage Space Requirements at the Store Level.	48
3.5.2. Aggregate Measures at the Warehouse.	50
3.5.2.1. Capital Investment Limit at the Warehouse.	50
3.5.2.2. Storage Space Requirements at the Warehouse.	54
3.5.2.3. Capacity Measures for Warehouse Deliveries to Stores	 54

4. ANALYSIS OF INVENTORY STRATEGIES	59
4.1. Experimental Design for Analysis of Inventory Policies	59
4.2. Simulation Modeling Assumptions	62
4.3. Comparison of Analytic Approximations to Simulation Results	65
4.3.1. Accuracy of Approximated System Costs	65
4.3.2. Estimations for Average Aggregate On-Hand Inventory	69
4.3.3. Requirement of 'Large' D for (s,S) Order Policies	71
4.3.4. Effect of Parameters on System Gamma Service Levels	73
4.4. Characteristics of Aggregated Orders to the Warehouse	77
4.5. Aggregate Inventories Generated by (s,S) Ordering Policies	80
5. SUMMARY AND CONCLUSIONS	84
BIBLIOGRAPHY	93
APPENDIX A: MATHEMATICAL NOTATION	100
A1. Notation for Store Models	100
A2. Notation for Stores with Warehouse Models	101
A3. Notation for Store Level Restrictions	101
A4. Notation for Warehouse Restrictions	102
A5. Notation for Warehouse Deliveries to Stores Restrictions	102
APPENDIX B: GLOSSARY OF TERMINOLOGY	104
APPENDIX C: EXPERIMENTAL DESIGN MATRIX FOR STORES-WITH- WAREHOUSE FRACTIONAL FACTORIAL DESIGN	106
APPENDIX D: HISTOGRAMS AND DESCRIPTIVE STATISTICS FOR WAREHOUSE ORDERS	107
APPENDIX E: HISTOGRAMS AND NORMAL PROBABILITY PLOTS FOR ON-HAND INVENTORY	139
VITA	201

LIST OF TABLES

Table 4.1. Factor Levels for Experimental Design 61

Table 4.2. Fractional Factorial Effects for Average Total Costs 68

Table 4.3. Total Cost Percent Deviations versus Average D/Mu 72

Table 4.4. Analytic versus Simulated Gamma Service Levels 74

Table 4.5. Means and Standard Deviations for Warehouse Orders 78

Table 4.6. Z-Scores for Equation (42) Confidence Limits 82

Table C1. Stores-with-Warehouse Design Matrix 106

LIST OF FIGURES

Figure 1.1. Two-Echelon Warehouse/Store System	2
Figure 1.2. Inventory Objectives Covered in the Literature	5
Figure 3.1. Multi-Store Inventory/Distribution System	27
Figure 3.2. Multi-Store System with Centralized Warehouse	34
Figure 3.3. (s,S) Inventory Policy Replenishment Cycle	49
Figure 4.1. Relative Deviation Between Analytic and Simulated Holding Cost . .	65
Figure 4.2. Relative Deviation Between Analytic and Simulated Order Cost . . .	66
Figure 4.3. Relative Deviation Between Analytic and Simulated Total Costs . . .	67
Figure 4.4. Simulated Average Inventory Versus Equation 32	69
Figure 4.5. Simulated Average Inventory Versus Equation 38	70

ABSTRACT

A two-echelon inventory and distribution system consisting of a centralized warehouse and N stores is considered in this paper. The inventories of the warehouse as well as the stores are controlled by periodic review (s,S) ordering policies. The expected levels of capital investment, storage space needs, capacity requirements for delivery vehicles, and reliable customer service are issues of great importance to practitioners when considering the introduction of a central warehouse and transportation system.

Helmut Schneider, Dan Rinks, and Peter Kelle have developed a methodology that has been shown to provide approximately optimal (s,S) policies under various demand conditions, and are easy to handle computationally. The approximations of Schneider et al., are used to generate ordering policies for the two-echelon system in order to observe the behavior of the aggregate inventories generated by the (s,S) policies using computer simulation.

The simulation results are used to evaluate the accuracy of the analytic models in predicting the aggregate inventory behavior, and simple computational formulas are proposed to calculate confidence limits for aggregate inventory levels and for shipping volumes and weights.

1. INTRODUCTION

The management of wholesale/retail inventories has been both an important and challenging research area for many years. When inventory performance measures such as number of *inventory turns* per year are examined, the ability for retailers to manage their inventories has lagged far behind other sectors such as manufacturing. Retailers, however, face unique problems in the management of their inventories. Major retailers, due to the nature of their business, must hold stock at many geographically dispersed locations often numbering in the thousands. Retailers are also faced with the problem of managing, in some cases, thousands of *stock keeping units* (SKUs) per location.

The introduction of data processing technology, the increasing use of cash register point-of-sale scanners, and the declining cost of electronic data storage now provide retailers with the ability to manage their inventories much more efficiently. Prior to these developments, most inventory information was reported only in aggregate form, usually considerably delayed and inaccurate (Nahmias and Smith; 1993). As a result, it was nearly impossible to accurately determine the *inventory position* of a SKU at a store, estimate the probability distribution of weekly demand, or calculate the customer service level achieved by individual items. Many major manufacturers and suppliers now use electronic data interchange technology to process purchasing and shipping transactions (Nahmias and Smith; 1993). This in turn reduces paperwork lead time allowing these suppliers to provide more frequent deliveries at reduced cost.

Inventory management models have been an important application area in the literature for many years and the number of papers published in this area is well into the thousands. Only a small percentage of these papers deal specifically with the *two-echelon* retail system proposed in this thesis. However, many of these have the potential for application in this area. A two-echelon inventory structure consisting of a warehouse which supplies N retail stores is shown in Figure 1.1. This structure is common to many wholesale/retail systems today, although there are many variations




Figure 1.1. Two-Echelon Warehouse/Store System

in the methods for their operation. The warehouse receives shipments from suppliers and manufacturers and distributes them to the stores. In some cases, the warehouse itself also maintains inventory. Many retailers consider the role for the warehouse stock to be a strategic issue and they report that their warehouses serve from as few as

fifty to as many as eight hundred stores (Nahmias and Smith; 1994). The warehouse is usually centrally located so that the stores it supports are within one or two days delivery distance by truck.

Figure 1.1 illustrates the type of multiechelon inventory system considered in this thesis. Demand is assumed to originate at the first echelon (stores) only and is transmitted to the second echelon (warehouse). Units are distributed in the opposite direction, from the warehouse to the stores. The warehouse transmits the total demand of the system to its suppliers who ship the required units to the warehouse.

There are several important reasons for introducing a centralized warehouse into a wholesale/retail system. First, it could be less costly to hold inventory at the warehouse than at the retail location. While location is considered one of the most important marketing decisions for a retail outlet, warehousing can often be located in comparatively remote areas where space is less expensive. Also, a warehouse can serve as a hedge against uncertainty. Since demand at the stores for most items is generally uncertain, it is often necessary to redistribute stock amongst the stores. Since *transshipments* (inter-store transfers) between stores are usually uncommon, redistribution can be handled by storing a proportion of the stock at the central warehouse to be shipped to the stores when necessary. Finally, transportation and/or purchasing costs could be reduced when there is a warehouse, since economies of scale result from the warehouse transacting business with suppliers in bulk quantities.

1.1. Contribution of the Research

Multi-item, multi-echelon inventory systems often involve thousands of stock-keeping units and require capital investments in the millions of dollars. Single-location inventory control models which have been modified for use in multi-echelon systems generally perform poorly when compared to methods which are designed to take advantage of the system structure. Taking advantage of system structure is particularly important when there are many items with low demand and high relative cost (Muckstadt and Thomas; 1980). One objective of this research is to investigate the savings potential for a wholesale stores system through introduction of a centralized warehouse.

There is often a serious gap that exists between theory and practice in inventory management in that the most important issues involve aggregate objectives and constraints (Gardner and Dannenbring; 1979). The major goals of inventory management are to minimize inventory investment, maximize customer service, and assure efficient (low cost) operation. It is plain to see that these goals can be inconsistent or even in direct conflict. Marketing managers insist on good customer service, production managers need to take advantage of efficient lot sizes, and purchasing managers wish to maintain low per-unit cost such that all have an inclination toward higher inventories. Conversely, engineers seeking to avoid obsolescence and financial managers insisting on efficient use of capital, prefer to hold inventories at a minimum (Tersine; 1994). Most multi-item, multi-echelon inventory models ignore the fact that aggregate constraints and/or restrictions exist and must be addressed by

practitioners. Figure 1.2 illustrates that problems concerning inventory objectives for a single stocking point both for a single item and for multiple items have been

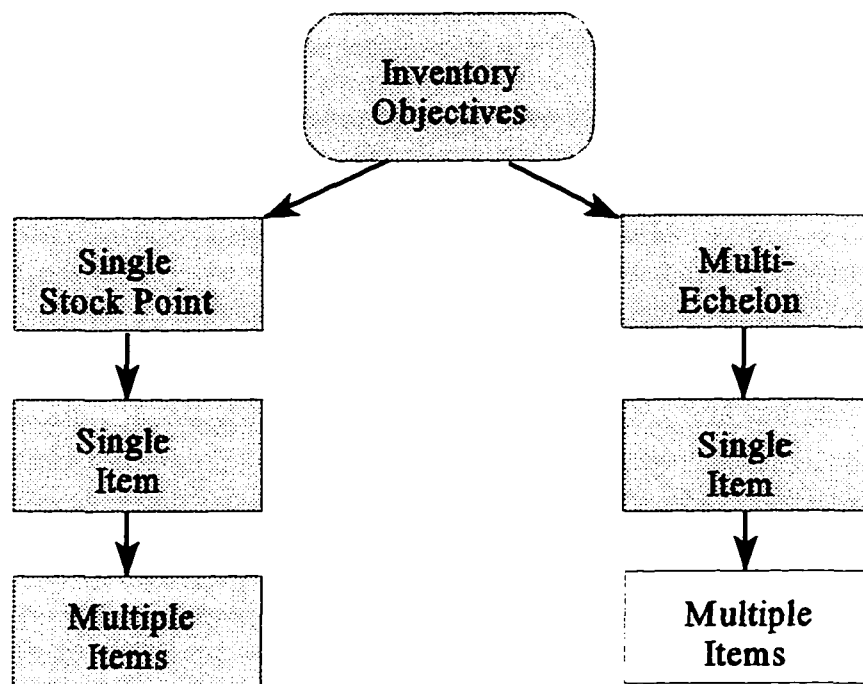


Figure 1.2. Inventory Objectives Covered in the Literature

approached in the literature. Also, in a multi-echelon environment, objectives for a single item have been approached. However, the problem of aggregate restrictions for multiple items in a multi-echelon environment has not been approached in the literature. Accordingly, a second objective of this research is to observe the characteristics of the aggregate inventories in order to determine capital investment and storage space requirements at the stores and at the central warehouse and order quantity or capacity requirements for deliveries from the warehouse to the stores. Functional decision

support can be provided to practitioners through use of the analytic models to supply inventory control and ordering policies for individual parts.

There are two general approaches to inventory policy modeling which are driven by the particular system being studied: (1) models which assume full backordering of excess demand and (2) models in which demands that occur when an item is out of stock are lost sales. Lost sales systems are difficult to analyze mathematically and are often avoided by researchers preferring to approximate lost sales with backorder models. The difficulties encountered in lost sale models will be discussed in detail in Chapter 3. Although the analytical models used in this research will be of the backorder type, the system being studied may operate in an emergency order environment where shortages are backordered, but are filled by special order. Thus, the customer does not wait until the next regular replenishment arrives.

1.2. Organization of the Research

This research is organized into five chapters. In Chapter 2, a review of the relevant background literature is discussed which provides the groundwork for the research. In Chapter 3, ordering policy models for a multi-item, multi-echelon inventory system where aggregate capacity measures are considered will be presented. In Chapter 4, the performance of the models discussed in Chapter 3 is analyzed for various demand parameter assumptions, followed by a summary and conclusions for the research in Chapter 5.

2. REVIEW OF THE LITERATURE

The multi-echelon inventory problem was first motivated by military logistics problems and has played a large role in the materials management of the armed forces such as in Demmy and Presutti (1981). Certain inventory problems faced by retailers today are similar to those faced by the military, and many of the inventory models used today are based on this early work. However, many of the models used for military systems were devoted to so-called *reparable items*, such as aircraft engines and electronic gear.

When a demand cannot be filled from available stock there are two alternatives for the customer to consider: (1) place a *backorder* or (2) cancel the order and obtain it elsewhere . Clearly, a large proportion of the literature on inventory modeling assume full backordering of demand. However, full backordering is not an appropriate assumption for retailers in most situations. Some portion of the customers whose demand cannot be met from available stock will choose to go elsewhere to make their purchase, resulting in *lost sales*. The backorder assumption is dependent upon the type of item being controlled. Items such as clothing and most food products, are typically lost sales items, while durable goods, such as furniture and appliances, are typically backordered items. Other items may be a mixture of backorders and lost sales.

Another characteristic of retail systems that has received little attention in the inventory literature is *substitution*. When an item is out of stock, customers will often purchase another model or color in place of their first choice.

Although there has been considerable literature on the *deterministic demand* problem, the focus of this thesis is *stochastic demand* policies that protect distributors against uncertainty. While deterministic demand has been used to approximate some multi-echelon production systems, it would not be considered a reasonable assumption for wholesale/retail inventory problems.

2.1. Classical Multiechelon Inventory Models

One of the earliest multi-echelon models was developed by Clark and Scarf (1960). They assumed that the system structure consisted of several installations, 1, 2, ..., N, with installation 1 receiving stock from 2, 2 receiving stock from 3, etc., and with demands originating at echelon one only. They also assumed that the cost of purchasing and shipping from any installation to the next is a linear function of the amount shipped without any set-up cost, except at the highest echelon. Finally, positive lead times for shipping between echelons and full backordering were assumed.

The Clark and Scarf study was significant for several reasons. It was the first to depict the form of an optimal policy for a stochastic demand, multi-period, multi-echelon model. It also was important for introducing the concepts of *echelon stock* and implied shortage cost, which formed the basis for the analysis of more complex systems. However, the assumptions of this model make it unlikely that it would be used to manage a real system. For instance, the assumption that all replenishment costs are proportional to the size of the replenishment order is somewhat unrealistic. When fixed ordering costs were applied at all locations, Clark and Scarf (1962) were able to provide only approximately optimal policies. Also, the simple serial system considered, where

items flow through echelons each with a single location, has limited applicability in practice since few actual retail distribution systems have this type of structure.

Bessler and Veinott (1965) extended Clark and Scarf's serial system to a general model using a unidirected tree structure. They assumed that ordering costs were proportional and that each facility could order from an outside supplier. They also assumed full backordering of demand at each facility and allowed redistribution of stock among the outlets at the end of each period. They show that if the initial stock is sufficiently small, it is optimal to order up to a certain base stock level at each facility. Although their results are interesting, their model is not applicable to a system where external demand occurs only at the store level, and where resupply from outside occurs only at the warehouse level. Rosenfield and Pendrock (1980) examine the problems of estimating total inventories and associated costs for alternative distribution warehouse configurations.

Federgruen and Zipkin (1984b), address the computational issues of the Clark and Scarf model. They showed that the optimal policy established by Clark and Scarf for the finite horizon problem can be extended to the infinite horizon versions of the problem under the criterion of discounted cost and for long-term average cost. They also establish simpler computational formulas in the infinite horizon case. Similar simplifications have been used for the single product/single facility problem such as that of Karlin (1958).

Another class of multi-echelon inventory models was developed to deal with the problem of determining suitable stocking levels for repairable items where almost all

of the system items are recycled back into the system after some random period of repair time. The first model of this type was called METRIC (Sherbrooke; 1968). The METRIC model was developed to determine the suitable stocking levels for items valued in the multi-billion dollar range whose availability is crucial for successful military operations. METRIC and several of its successors are the basis of many military inventory control systems. Clark (1972) provides an early survey of these models as well as a discussion of work on related deterministic problems. Nahmias (1981) followed with a comprehensive survey of these models up to 1981. More recent results on the repairable inventory problem can be found in Albright (1989).

2.2. Single Period Models

Many of the early efforts at modeling the warehouse/store system shown in Figure 1.1. focused on a single time period only. These studies can be considered as extensions of the newsboy models (Nahmias; 1989) in which both ordering and redistribution policies are considered. The assumption of a single planning period is important in building an understanding of the nature of the issues faced when managing such a system.

Allen (1958) was the first to consider a one period version of the two-echelon system of Figure 1.1 which is actually a model of a single echelon problem consisting of K locations. The objective of this model was to determine an optimal redistribution of stock among the locations where no reorder decision is involved. Allen assumed that the demands at the locations are random variables with continuous probability density functions. Allen derived an algorithm for finding the optimal redistribution of stock

among the locations. He observed that a shipment from location k to location j will occur only when the shortage probability is higher at j than at k .

Simpson (1959) considered a model similar to Allen's. He assumed that there were a total of Q units to allocate among N outlets. He showed that the optimal allocation is the one that equalized the stockout probabilities at each of the locations. Both Allen's and Simpson's models were valuable to later research because they were the first to present the importance of equal runout probabilities.

Gross (1963) developed a similar single period model but allowed for both replenishment and redistribution at the retail outlets. Similar models were considered by Krishnan and Rao (1965) and Das (1975). In contrast to Gross's model, Das considered only two locations, but allowed for two subperiods and assumed that there would be a redistribution of stock between the subperiods. Both results provide regions of two-space associated with the optimal redistribution policy.

Hoadley and Heyman (1977) also considered a model similar to Gross (1963). Their model is more general in that they allowed for: returns from the store to the warehouse; disposal of stock at that echelon; ordering from the stores to the warehouse; and transshipment of inventories among the retail outlets. Their decision problem is to choose an initial stock level at the warehouse and an initial allocation so as to minimize the initial stock movement costs during the period plus inventory carrying costs and system shortage costs at the end of the period.

Eppen (1979) considered a multi-location newsboy problem with linear holding and penalty cost functions at each location with normal demand. He assumed N

identical retail outlets that order independently according to a simple order-up-to point model obtained by minimizing one period holding and penalty costs, and derived an expression for the expected cost at each facility. The model is used to demonstrate: the expected holding and penalty costs in a decentralized system exceed those in a centralized system; the magnitude of the savings depends on the correlation of demands; and if demands are identical and uncorrelated, the costs increase as the square root of the number of consolidated demands.

While single period models can be used in a dynamic environment, they are often inadequate for handling system dynamics. They usually ignore replenishment lead times, which is the major reason for holding *safety stock*. Also, the problem of redistribution is of less importance than the problem of reordering. In most actual retail systems there is rarely lateral resupply among stores since the retailers effect redistribution by holding inventory at the warehouse level (Nahmias and Smith; 1993).

2.3. Models Where the Warehouse Holds No Stock

Models have been developed for the situation where the warehouse holds no inventory. In these cases, retailers use the warehouse merely as a distribution point through which items flow. The first study of this type system was from Eppen and Schrage (1981). Their model is essentially the warehouse/store system of Figure 1.1. The assumption that the warehouse does not hold stock does not mean that the warehouse serves no purpose. By ordering centrally, more advantageous quantity discounts can be sought. There are also "statistical economies of scale" as observed by Eppen (1979) in which savings are achieved by aggregating orders rather than operating

N individual inventory systems. They show that unless the demand at the stores is perfectly correlated, the coefficient of variation of demand (σ/μ) for the aggregate system is smaller than for the demand at the individual stores.

Eppen and Schrage assume deterministic *lead times* from both the supplier to the warehouse and from the warehouse to the stores. Demand originates only at the store level and is assumed to follow a normal distribution with parameters allowed to differ between stores. This paper shows conditions under which a simple ordering policy is optimal. They call their ordering policy an (m,y) , where every m periods, the inventory position is raised to a base stock of y . The warehouse must have enough stock to resupply the stores so that the probability of a stockout at each store is the same. This is shown likely to hold when the fixed cost of ordering from the warehouse is high and/or the coefficient of variation of demand at the stores is moderate.

Federgruen and Zipkin (1984a) consider the same problem but approach the solution differently. They formulate the problem as a dynamic program with a state space of very large dimension. The dimension of the problem is at least $N+L$ (the number of outlets + the supplier to warehouse lead time), which makes it impractical to solve except for small values of N and L . To avoid the "curse of dimensionality" they show that the model can be systematically approximated by a single location inventory problem.

Federgruen and Zipkin (1984c) consider other approaches to the problem of approximating optimal policies where they assume that the penalty and holding costs are proportional. The results of this later paper deal with the problem of unequal

coefficients of variation of demand at the stores. The approximation techniques and results are similar the those of Federgruen and Zipkin (1984a).

An extension of the Eppen and Schrage (1981) model to the case of correlated demands over time was considered by Erkip, et al. (1990). Johnson and Thompson (1975), Ehrhardt, et al. (1981), and Schneider, et al. (1995) have also considered the impact of correlated demands over time on inventory control strategies.

2.4. Models Where the Warehouse Holds Stock

The models in the previous section assumed that the warehouse holds no stock, meaning that as soon as a shipment arrives from a supplier, the entire shipment is distributed to the stores. Because of the uncertainty of demand at the store level, it is likely that the stores' inventories could become out of balance before the warehouse receives the next shipment. If transshipments between stores are not allowed, it is possible that some stores will run out of stock while others will become overstocked during the *replenishment cycle*. Assuming the warehouse holds no stock defeats a major purpose for operating a warehouse. When safety stock is retained at the warehouse (known as *risk pooling*), balancing of store stocks can be achieved between warehouse replenishments.

Jönsson and Silver (1987a) considered the problem of correcting the imbalance of stocks at the store level. They assumed that the warehouse replenishes the system inventory using a base-stock policy and a predetermined order cycle of H time periods so that the system wide inventory position is returned to a predetermined level based on a given minimum level of customer service. The entire shipment is then allocated

among the stores. They further assumed that there is a cost per unit redistribution of inventory among the stores exactly one period before the start of the next order cycle. They show that, for any desired service level, a lower value of the system order-up-to level can be achieved by the redistribution, but note that these improvements must be weighed against the costs of transshipment.

Jackson (1988) studied a similar problem assuming that the warehouse can initiate a redistribution of stock each period between warehouse replenishments by using a ship-up-to-S policy each period. As long as the warehouse has adequate stock, at the start of each period it will ship to each store an amount equal to the demand of the prior period. This continues until either the warehouse receives a replenishment or it runs out of stock. If the warehouse runs out of stock before replenishment, there will be one period in which there is insufficient stock at the warehouse to raise each store to its desired S level. Both exact and approximate models are formulated to determine the optimal allocation in this runout period which are shown to simplify in the identical store case.

An extension of Jackson's work was presented by Jackson and Muckstadt (1989). They expand on the model of Eppen and Schrage (1981) by assuming zero lead times and two allocation periods within each order cycle where there are two opportunities to ship stock from the warehouse to the stores within each order cycle. Their formulation is very hard to solve unless they assume independent demands and two stores. They conclude that risk pooling by retaining stock at the warehouse appears

to be attractive even when demands are independent and there are only two stores in the system.

Schwarz (1989) considers essentially the same model as Jönsson and Silver (1987a), but focuses his analysis on fixed length lead times rather than on safety stock. He develops two systems. First, each retailer operates independently, receiving goods directly from an outside supplier after a fixed leadtime consisting of the suppliers leadtime and the supplier-to-retailer shipping time. Second, the system order is shipped to the warehouse, arriving after a fixed leadtime. The warehouse then ships the units to the retailers to equalize their inventory positions. However, the system is assumed to have an additional lead time between the warehouse and the store. The primary conclusion of their research is that the value of adding the warehouse depends most importantly on the pipeline inventory costs, since the additional lead time means larger pipeline inventories.

The advantages of risk pooling were also investigated by McGavin et.al. (1990). Between warehouse replenishments an allocation policy is specified by the number of withdrawals from stock, the intervals between successive withdrawals, the quantity of stock withdrawn at each withdrawal, and the partitioning of the withdrawn stock among the stores. This study deviates substantially from previous work. It assumes lost sales which is very common for most low cost retail items. It also assumes that demand for each store is generated by a stochastic process with stationary independent increments. One of the interesting results of this work is that the optimal partitioning scheme is what

they call a "balancing" scheme, where the stock is partitioned to maximize the minimum store inventory.

Kumar et.al. (1995) examine two alternative policies for replenishing and allocating inventories among N stores located along a fixed delivery route. The warehouse is restocked every m periods after a fixed lead time of T periods. The warehouse holds no inventory. Under a "static" policy, allocation are determined for all retailers simultaneously at the moment the delivery vehicle leaves the warehouse. Under the "dynamic" policy, allocations are determined sequentially upon arrival of the delivery vehicle at each store. Simulation models show that the dynamic policies give better results.

Graves (1989) considers a system in which final demand at each retail site is generated by independent Poisson processes with known transshipment times from each supply point. Each site has a single supplier, but may have several successors. This type of structure was also addressed by Bessler and Veinott (1965). Graves assumes that replenishments are made at fixed times at each location, and that the size of a replenishment is equal to the number of units demanded at that location since the last replenishment. Each time a demand occurs, that demand triggers a replenishment request at the facility's supplier which triggers another request to the next supplier, etc.

Graves' approach was also considered by Axsäter (1990a) who assumed order-up-to- S policies. Axsäter assumed Poisson demands and full backordering of demand. Each facility applies a periodic review, order-up-to- S policy. In case of shortages at the warehouse, orders for individual units are filled in the same order as the original

demand at the retailers using a 'virtual allocation' scheme. Axsäter's primary contribution is a recursive approach for computing the holding and shortage costs at each level of the system.

Nahmias and Smith (1994) developed a two level model with partial lost sales and arbitrary mean to variance ratio of the demand at the stores. They assume that the demand process is a Poisson process with a random rate parameter. They also assume that each store is restocked according to a base stock or order-up-to-S policy. A newsboy type model is developed to determine the optimal S level at the stores. Because of partial lost sales, the probability distribution of the total demand on the warehouse is complex. They conclude that the benefit of retaining stock at the warehouse is most significant for low demand, high value items.

A study conducted by Chen and Zheng (1994) establishes lower bounds on the minimum costs of managing certain distribution networks with setup costs at all levels. Transportation lead time at each level is constant and customer demands occur only at the store level and are independent and identically distributed. Unsatisfied customer demand is fully backordered and replenishment decisions are centralized and based on system-wide inventory information under a cost-allocation, physical decomposition framework. By allocating expected holding and backorder cost functions, they derive cost minimizing "induced penalty bounds" for stochastic demand, one-warehouse, multi-echelon systems. Also by allocating cost rates, they generate a new class of bounds, called "parameter-allocation bounds," which are derived for both periodic and continuous review stationary systems operating over an infinite horizon.

2.5. Systems Using Lot Size/Reorder Point Models

Probably the most widely used policy model for inventory management when demands are random is the lot size/reorder point, or (Q,R) , policy. A (Q,R) policy for a single product, single echelon system would be implemented as follows: when the inventory position (quantity on hand plus quantity on order minus quantity on backorder) reaches a level R , an order is placed for Q units. The order arrives following some lead time τ which can either be of fixed length or random. Since the (Q,R) policy is widely applied, it is natural that it be adapted to the warehouse/store system of Figure 1.1.

Deuermeyer and Schwarz (1981) were the first to adapt the (Q,R) to the two-echelon system. They assume that the demands at the stores form independent Poisson processes with respective rates λ_i for $1 \leq i \leq N$. They present an analytic model for estimating the expected service level measured by backorders and fill rates. The fill rate is the proportion of demands that can be met from available stock. They also provide expressions for the time weighted number of backorders at facility j . Their system consists of N identical retailers facing stationary Poisson demand, known lead times, and using stationary (Q,r) ordering policies. Simulation tests showed a close match between observed simulated measures and those computed with the model. This study was an important step in modeling the general warehouse/store problem with lot size/reorder point policies and laid the groundwork for later research.

An extension to include fill rate maximization subject to a constraint on system safety stock, was considered in Schwarz, et.al. (1985). They assumed k identical stores

with a common fill rate F_k . By considering the total dollar investment in safety stock at each store, S_k , and at the warehouse, S_w , the authors determine the proper levels within the total budget available for safety stock S . By varying values of S , they generate solutions where the optimal safety stock policy is shown to be the intersection of a fill-rate policy line and the safety stock budget line. Analysis of the properties of the policy line leads to a simple piecewise linear approximation which gave results within 2% to 3% of the optimum for selected cases.

A similar model was treated by Badinelli and Schwarz (1988) in which the goal of minimizing backorders was considered. Assuming identical stores whose expected backorders are B_k , the authors formulated three optimization problems in which the goal is to choose safety stock at the stores and at the warehouse to minimize system-wide backorders subject to a budget constraint on average system inventory, minimize average system-wide inventory subject to a constraint on expected system backorders, or minimize system-wide backorders subject to a constraint on system-wide safety stock. They introduce a heuristic for minimizing expected backorders with respect to a constraint on average system on-hand inventory. Their heuristic prescribes that a 'near-zero' inventory level at the warehouse is approximately optimal.

Svoronos and Zipkin (1988) consider several refinements of the Deuermeier and Schwarz (1981) model. They obtain an expression for the variance of lead time demand at the warehouse by considering a warehouse-to-store lead time that is composed of two components, one deterministic and one random, deriving the variance of the random component. They also use Poisson distributions to estimate the distribution of

warehouse lead time demand. Using simulation, their results are compared to those of Deuermeier and Schwarz (1981) with a significant improvement in all cases tested.

A similar model was provided by Axsäter (1993a). He develops a different computational method from Svoronos and Zipkin (1988) which he shows to be more efficient for large problems. He develops three approximations which are used to show that the results can be used for the exact or approximate evaluation of more general policies where both the retailers and the warehouse order in batches.

2.6. Systems Using Periodic Review Order-up-to-S Models

Inventory policy models are often presented in which inventory levels are reviewed periodically rather than continuously. Scarf (1960) showed that for a single product system with periodic review, (s,S) policies are optimal. An (s,S) model is implemented as follows: when the starting inventory level x in any period is less than or equal to s , an order is placed for $S-x$ units. When inventory levels are reviewed continuously and demand is smooth, the starting inventory level when an order is placed is exactly s and the order size is $S-s$. By setting $R=s$, and $Q=S-s$, it can be seen that the continuous review (s,S) policy reduces to the (Q,R) policy. We can therefore think of (s,S) policies as a larger class of policies that include (Q,R) policies as a special case. Here, we review the case that inventory levels are reviewed periodically rather than continuously. In this situation, the size of the replenishment order will vary depending on the overshoot of s (the difference between the current inventory level x and the order flag s , or $x-s$).

The first study of (s,S) policies for a two-echelon warehouse/store system came from Ehrhardt et.al. (1981). They assumed that each store follows an (s,S) policy placed on the warehouse for its replenishments and that demands are reviewed periodically. Demands at the stores are assumed to be independent identically distributed random variables with arbitrary probability distributions. Full backordering of excess demand is assumed at all levels of the system, and orders are filled on a first come first served basis.

The total demand on the warehouse in any period is the sum of the demands at the stores. They show that if the stores follow identical (s,S) replenishment rules, the demand pattern at the warehouse will be correlated in time. The authors note that, because of this, the optimal form of the warehouse policy is not an (s,S) policy, but they apply that form because of its wide application in practice. They suggest that the warehouse policy be computed by using the power approximation method of Ehrhardt (1979), which incorporates the period to period correlation.

Simulation was used to search for the nearly optimal warehouse (s,S) replenishment policy to test the effectiveness of the approximation. This study was the first to recognize that the (s,S) policy at the stores induces a correlated demand pattern at the warehouse and that significant errors arise if the correlation is ignored.

Schneider et.al. (1995) extend the Ehrhardt model by focusing on service levels at both the warehouse and store levels. They examine two types of service levels: stockout occasions (α) and time weighted backorders (γ). They determine the distribution of the lead time to the stores as a combination of a fixed delivery time and

a stochastic component which is the result of the warehouse being out of stock. This is accomplished by using the α level of service prescribed for the warehouse. This was analyzed exactly by Svoronos and Zipkin (1988) for the (Q,R) model. The demand distribution for the stores is approximated by a negative binomial. The power approximation of Ehrhardt and Mosier (1984) is used to provide estimates for the (s,S) values at the stores and at the warehouse which requires only knowledge of the first and second moments of the demand distribution. Simulations suggest that their results are good when the service levels at the warehouse are relatively high.

2.7. Systems Using Continuous Review Order-up-to-S Models

Svoronos and Zipkin (1991) consider continuous review (S-1,S) policies in the context of a multi-level inventory system. An (S-1,S) policy is implemented as follows: whenever a demand for a unit occurs, a unit is reordered, this is also called one-for-one replenishment. This policy is optimal for very high value items or for systems where the fixed cost of reordering is negligible and it is often used for expensive repairable items as in Sherbrooke (1968) and Graves (1985) rather than low cost retail items. The multiechelon system which Svoronos and Zipkin consider may have an arbitrary number of layers and only full backordering of excess demand is considered. Demand at the lowest echelon is assumed to follow the Poisson distribution and they do not allow for random transit times between levels. They explore the effect of varying several system parameters on the optimal base stock levels. Their results show that, in contrast to prior multiechelon models, transit-time variances play an important role in system performance.

These results were extended by Zipkin (1991) to the case where demand at the lowest echelon follows a compound Poisson distribution. Their work was significant in that it proves that with proper choice of the compounding distribution, one can model a demand process with an arbitrarily large variance to mean ratio, which is found to be the case for many actual retail items in practice.

Recently, Hausman and Erkip (1994) considered the costs of centralized replenishment policies compared to systems controlled as a network of single-echelon systems for low-demand, high-cost items using an $(S-1, S)$ policy. They assume that lead time for resupplying the stores depends on the availability of stock at the warehouse at the time a store places an order. Using a modification of the single-echelon model of Muckstadt and Thomas (1980), they explore the amount of suboptimization which can occur if multi-echelon systems are managed as independent single-echelon systems. They show that beyond a certain service level of investment in the stores, a saturation point is reached at which investment in the warehouse must be increased to improve the "back-up" service. They conclude that the single-echelon model can be applied at a low penalty cost if the system budget is not excessively constraining and that the managerial simplicity of the single-echelon models may be sufficient inducement for firms to tolerate the suboptimization.

3. DEVELOPMENT OF INVENTORY STRATEGIES

This chapter describes various inventory models which will provide the basis for a comparison of different inventory strategies in Chapter 4. The first is a multi-store model in which all stores are independently supplied by an outside source. The second is a two-echelon inventory model with one centralized warehouse which is supplied by an outside source and several stores which are supplied from the warehouse. The two-echelon model will consider capacity measures for the order frequency and/or order quantities for the stores as well as capacity constraints on the central warehouse.

A system consisting of N geographically dispersed stores selling replacement parts is considered where each store operates independently in terms of inventory control using periodic review (s,S) policies. At the beginning of a review period, if the inventory position (the number of items on-hand, plus on-order, minus back orders) for part j , is less than or equal to s_j , an order is placed for a sufficient number to raise the inventory position to S_j . Orders are placed with the manufacturer which arrive at the beginning of the following period, after the lead time has passed. Arrival time at the store can be estimated fairly accurately and can therefore be considered as deterministic. Parts arriving as regularly scheduled shipments are subject to a purchase discount and reduced freight cost. This system is shown in figure 3.1.

In the case of a *part shortage* (on-hand inventory equal to zero when a customer demand occurs), occasionally, the demanded part is located at another store and is shipped overnight. If the part is not available from another store, an emergency order is placed to the manufacturer which can also be shipped to the store using next-day

freight. Under this scenario, the cost of a stockout is easily obtained assuming that the customer is willing to wait for the part to be shipped. However, in some cases, the customer will not be willing to wait for the part and the sale will be lost. The cost of a lost sale is not so easy to determine. Aside from the obvious lost revenue, other considerations must also be included such as loss of goodwill and its effect on future sales.

Whether or not the customer is willing to wait for an emergency order, this system might be approached using what is known in the literature as a *lost sales model*, although these are different in that the sale does occur with an emergency order. A lost sales problem is very difficult to solve, particularly when there is a positive order lead time. Not only is the cost of the lost sale hard to determine as previously mentioned, but the tendencies of the inventory position when lost sales are considered are extremely difficult to analyze mathematically. Because of the difficulty of dealing with lost sales, few authors have considered it in the context of a two-echelon inventory problem, and only then when the lead time for shipments was small enough to be ignored. Since the lead times for deliveries in the system being considered here are too large to be ignored, and since the introduction of the central warehouse will greatly effect these lead times and therefore the inventory policy selected, the more common assumption of full backordering of demand will be applied.

When full backordering is assumed and a shortage occurs, the customer's order will be filled when the next regular shipment of parts arrives at the store. Since this system supplies replacement parts which are not readily available from nearby

competitors, the backorder assumption is not unreasonable and should result in a close approximation.

3.1. Multi-Store Inventory Model

First, a solution to the (s,S) policy for a single store is considered. Under a periodic review system, the inventory position is checked at the end of the *review period* and an ordering decision is made. The review period is the fixed length of time which passes between ordering decisions and is used as basic time unit for the model. We assume independent demand within review periods at each store i with probability density function $f_i(r)$ having mean μ_i , and variance σ_i^2 , $i=1,2,3,\dots,N$. The cost




Figure 3.1. Multi-Store Inventory/Distribution System

components at each store are ordering cost K_i , holding cost h_i , and stock out cost p_i . Unsatisfied demand is backlogged and a fixed transportation lead time L_i between the

manufacturer and the stores is assumed. An order arrives after its respective lead time at the beginning of a period. The sequence of events in a period is (1) delivery, (2) demand, (3) review, and (4) order. All costs occur at the end of a period after the demand has occurred. For a single echelon periodic review system which operates under a constant ordering cost, linear holding and stock out costs, fixed replenishment lead time and backlogging of unsatisfied demand, an (s,S) policy has been proven to be optimal as shown by Scarf (1960) and Iglehart (1963). The notation used throughout this chapter is defined in Appendix A.

Following the models of Schneider, et al. (1995), let x_i be the inventory position at store i at the beginning of an arbitrary time period t . The conditional expected inventory level $H_i(x_i)$ at the end of period $t+L_i$ is

$$H_i(x_i) = \int_0^{x_i} (x_i - r) f_i(r; L_i + 1) dr \quad x_i \geq 0 \quad (1)$$

where $f_i(r, L_i + 1)$ is the $(L_i + 1)$ -fold convolution to $f_i(r)$. The expected backlog at the end of the period $t+L_i$, $B_i(x_i)$, is given by

$$B_i(x_i) = \begin{cases} \int_0^\infty (r - x_i) f_i(r; L_i + 1) dr & x_i \geq 0 \\ \int_{x_i}^\infty (r - x_i) f_i(r; L_i + 1) dr & x_i \leq 0 \end{cases} \quad (2)$$

Igelhart (1963) gives the steady state distribution of the inventory position for an (s,S) policy as

$$\Psi_i(x) = \frac{m_i(S_i - x)}{1 + M_i(D_i)} \quad s_i \leq x \leq S_i \quad (3)$$

where $D_i = S_i - s_i$, $M_i(D_i)$ is the renewal quantity and $1 + M_i(D_i)$ gives the expected number of periods between placing an order with the manufacturer. The numerator of (3) gives the density function of on-hand inventory minus any backlogs at the beginning of the last period before falling below the order decision point s_i . So, to compute the long-run average cost $C_i(s_i, S_i)$ for store i , we take the sum of ordering, holding, and shortage costs at x and integrate over the steady state distribution of the inventory, x . Let $G_i(x) = h_i H_i(x) + p_i B_i(x)$ be the conditional expected holding and shortage costs, then the long-run average cost at store i is given by

$$C_i(s_i, S_i) = \frac{K + G_i(S_i) + \int_0^{D_i} G_i(S_i - x) m_i(x) dx}{1 + M_i(D_i)}, \quad (4)$$

and since each store operates independently, the resulting long-run average cost for the system of N stores is given by the sum of the costs at the N stores as

$$\sum_{i=1}^N C_i(s_i, S_i) = \sum_{i=1}^N \frac{K_i + G_i(S_i) + \int_0^{D_i} G_i(S_i - x) m_i(x) dx}{1 + M_i(D_i)}. \quad (5)$$

When the cost of a shortage is difficult to estimate, a stockout constraint can be imposed. Such a constraint has been defined by Schneider and Rinqest (1990), and is referred to as the γ -service level:

$$\gamma = 1 - \frac{\text{average cumulative backlog per period}}{\text{average demand per period}}.$$

For this case, the objective is to minimize the long-run average holding and ordering cost per period subject to the constraint that the long-run average backlog per period divided by the average demand per period is equal to $(1-\gamma)$. If $H_i(x_i)$ and $B_i(x_i)$ are averaged with respect to the steady state distribution in (3) then the result is

$$C_i(s_i, S_i) = \frac{K_i + h_i H_i(S_i) + h_i \int_0^{D_i} H_i(S_i - x_i) m_i(x_i) dx_i}{1 + M_i(D_i)} \quad (6)$$

along with the constraints

$$\gamma_i = 1 - \frac{B_i(S_i) + \int_0^{D_i} B_i(S_i - x_i) m_i(x_i) dx_i}{[1 + M_i(D_i)] \mu_i} \quad (7)$$

for $i = 1, \dots, N$. Constraint (7) will always be binding for each store under continuous demand since a higher imposed service level will require higher safety stocks.

Schneider and Rinks (1989) have shown that a good approximation for the total ordering and holding costs at the stores is asymptotically ($D_i \rightarrow \infty$) given by

$$C_i(s_i, S_i) = K_i \left(\frac{\mu_i}{Q_i} \right) + h_i \left[D_i \left(1 - \frac{D_i}{2Q_i} \right) + s_i - \mu_{iL_i+1} \right] + (h_i + p_i)(1 - \gamma_i)\mu_i + o(1) \quad (8)$$

for which an approximately optimal minimum cost is obtained using the power approximations of Ehrhardt and Mosier (1984) to calculate the decision variables

$$D_i = 1.463\mu_i^{0.364} \left(\frac{K_i}{h_i} \right)^{0.498} \sigma_{iL_i+1}^{0.138}, \quad (9)$$

$$Q_i = \left(D_i + \frac{\mu_i^2 + \sigma_i^2}{2\mu_i} \right), \quad (10)$$

and whenever a γ service level at the stores is specified, the shortage cost (p_i) in equation (8) is set=0, and $(1 - \gamma_i)$ as used in (8) is computed using

$$1 - \gamma_i = \frac{\int_{s_i}^{\infty} (r - s_i)^2 f_i(r; L_i + 1) dr}{2\mu_i Q_i}, \quad (11)$$

where we search for the value of s_i that provides the specified γ .

An estimate of the optimal policy parameter s_i can also be obtained when a unit stockout cost (p_i) is known using the power approximation presented by Ehrhardt and Mosier (1984) in the form

$$s_i = \mu_{iL_i+1} + \sigma_{iL_i+1}^{0.832} \left(\frac{\sigma_i^2}{\mu_i} \right)^{0.187} \left(\frac{0.22}{z_i} + 1.142 - 2.866z_i \right) \quad (12)$$

where

$$z_i = \sqrt{\frac{D_i}{\left(\frac{1+p_i}{h_i} \right) \sigma_{L_i+1}}} \quad (13)$$

3.2. Inventory Model for Multi-Store System with Central Warehouse

This proposed inventory/distribution systems consists of the original N stores and includes a warehouse facility which acts as a central stocking point to hold and distribute parts to the entire store system. Although the warehouse will undoubtedly be located in a city where a store currently exists, we will assume that the volume of parts required to be held by the warehouse will be such that no existing store has sufficient capacity to serve as the warehouse. The system will, therefore, incur a fixed cost for expansion of an existing store or rental/purchase of a new facility, and additional operating costs including warehouse-to-store shipping. The warehouse is to be strategically located so that transportation lead time to any store is one day or less.

The stores will operate similar to the multi-store system, using periodic review (s,S) policies with the same weekly review period. However, rather than ordering directly from the manufacturer, the stores will now order from their own central

warehouse. Through minor modification of their inventory information system which is already in place, a store which places an order at the end of a working day will receive it during the next day. In the case of a shortage at a store, the part is again backordered and will be available to the customer upon arrival of the next regular order.

The warehouse is also assumed to operate under a periodic review (s,S) policy where orders for parts are placed to the manufacturer with a weekly review period. The warehouse will, therefore, be operating under an ordering and delivery scheme similar to that of a single store in the current system, however, order quantities are sufficient to satisfy demand for the entire N store system. If a shortage occurs at the warehouse, the store's order for the shortage item is backordered and shipped after it arrives from manufacturing along with the store's next regular replenishment. The schematic of this proposed system is shown in Figure 3.2.

Adding a central warehouse provides several opportunities for cost savings over the multi-store system. Orders placed by the warehouse to the manufacturer will be allowed the same purchase discounts and reduced shipping costs as are available to the individual stores. Additional savings may also be obtained by the warehouse's ability to repackage certain parts for store distribution. For example, a small part may come from the manufacturer in cases of 10 parts when a store may only require an order quantity of two parts. The warehouse could send the smaller order to the store after breaking the case where previously, the store must receive a full case of ten parts.




Figure 3.2. Multi-Store System with Centralized Warehouse

Additionally, a shortage at the warehouse would incur the same cost to the system as a store shortage in the present system.

For the stores, the reduced ordering lead time and resulting reduced demand variance during the protection interval will allow a lower average inventory level to be maintained and therefore lower holding costs. Also, shipping costs from the warehouse will be less than those from the manufacturer in the case of a store stock out while the warehouse will already have received the purchase discount. The savings in total system costs (holding, ordering and shortage) will help to offset the additional fixed costs of warehouse operation and warehouse-to-stores transportation.

The inventory level at the warehouse is reviewed using the same periodic interval as the stores. If the inventory position is at or below s_w , an order is placed which arrives at the warehouse following a fixed leadtime L_w . The distribution of demand from the stores must be determined in order to set up the cost equations at the warehouse. Schultz (1983) derived the distribution of demand $F_w(Q)$ at the warehouse if the stores follow an (s,S) policy.

An order from a store which cannot be immediately filled by the warehouse is backordered and will be shipped upon arrival of a warehouse replenishment. The problem of allocating stock amongst the stores when the warehouse inventory is insufficient to fill all demand at the stores has been discussed by Eppen and Schrage (1981). Although the periodic review case assumes that the orders from the stores arrive simultaneously, a first-come-first-served rule will be used for this model. A stockout occurrence at the warehouse will cause the order of at least one store to be delayed. The result is that leadtime for store replenishment becomes stochastic although the transportation time between warehouse and store is deterministic. Because of this, the stores must increase their safety stocks accordingly to maintain their desired service levels and there will be a trade-off between warehouse and store safety stocks.

Because the impact on leadtime resulting from a warehouse stockout cost is difficult to determine, an α -service level will be imposed at the warehouse. The α -service level is defined as the probability that demand does not exceed supply during an arbitrary period. This type of service level will allow the leadtime distribution to be more easily modeled as a function of α than other service levels.

Assume that the inventory position at the warehouse at the beginning of period t is x . The conditional probability of demand not exceeding supply by the end of the review period $t+L_w$ is $F_w(x; L_w+1) = P(Q \leq x)$, where Q is the demand at the warehouse during leadtime plus review time. If we integrate over the steady state distribution of inventory on hand plus on order at the warehouse for an (s_w, S_w) policy, we obtain the $1-\alpha$ unconditional probability of a stockout during an arbitrary period, where

$$\alpha = \frac{F_w(S_w; L_w+1) + \int_0^{D_w} F_w(S_w - x; L_w+1) m_w(x) dx}{1 + M_w(D_w)} \quad (14)$$

and $D_w = S_w - s_w$.

If the cost components at the warehouse consist of a fixed ordering cost K_w , and an inventory holding cost h_w , the long-run average holding and ordering cost per period for the warehouse using an (s, S) policy is given by

$$C_w(s_w, S_w) = \frac{K_w + h_w H_w(S_w) + h_w \int_0^{D_w} H_w(S_w - x) m_w(x) dx}{1 + M_w(D_w)} \quad (15)$$

where $H_w(\cdot)$ is the conditional expected inventory as given by equation (1).

3.3. Approximations for the Warehouse/Stores System

When a stockout cost is specified at the stores and an α -service level is specified at the warehouse, the long-run average cost per period of the warehouse/stores system is given by

$$C(s_w, S_w, s_1, S_1, \dots, s_N, S_N) = C_w(s_w, S_w) + \sum_{i=1}^N C_i(s_i, S_i) \quad (16)$$

subject to (14).

It can be seen that for the costs $C_i(s_i, S_i)$ in (4), that the demand distributions $f_i(r; L_i + 1)$ can no longer be used with deterministic lead times L_i for the stores as in (1) and (2). Since the leadtime distribution is stochastic due to stockouts at the warehouse, the leadtime demand distribution $f_i(r; L_i + 1)$ must be used. This distribution, $f_i(r; L_i + 1)$, is used in equations (1) through (5) in place of $f_i(r; L_i + 1)$ resulting in the modified cost at the stores (4) denoted by C_i . The distribution of stochastic leadtime demand depending on α is, however, somewhat difficult to obtain. A power approximation will be used that has shown promising accuracy (Ehrhardt; 1979) and requires knowledge of only the first and second moments (μ, σ^2) of the demand distribution.

The optimal set of policies can now be obtained using the Lagrangian method with parameter λ_w for the unit stockout cost model at the stores:

$$\begin{aligned}
\Lambda(s_w, S_w, s_1, S_1, \dots, s_N, S_N, \alpha, \lambda_w) = & C_w(s_w, S_w) + \sum_{i=1}^N C_i(s_i, S_i) \\
& + \lambda_w \left[\alpha - \frac{F_w(S_w; L_w + 1) + \int_0^{D_w} F_w(S_w - x; L_w + 1) m_w(x) dx}{1 + M_w(D_w)} \right] \\
& + \sum_{i=1}^N \lambda_i \left[\gamma_i - 1 + \frac{B_i(S_i) + \int_0^{D_i} B_i(S_i - x) m_i(x) dx}{[1 + M_i(D_i)] \mu_i} \right].
\end{aligned} \tag{17}$$

The optimal policies must satisfy the first order conditions

$$\begin{aligned}
\frac{\partial \Lambda}{\partial s_i} = 0, \quad \frac{\partial \Lambda}{\partial S_i} = 0, \quad \text{and} \quad \frac{\partial \Lambda}{\partial \lambda_i} = 0 \quad \text{for } i=1, \dots, N, \\
\text{and} \quad \frac{\partial \Lambda}{\partial s_w} = 0, \quad \frac{\partial \Lambda}{\partial S_w} = 0, \quad \frac{\partial \Lambda}{\partial \alpha} = 0, \quad \frac{\partial \Lambda}{\partial \lambda_w} = 0.
\end{aligned} \tag{18}$$

Note that these conditions are necessary but not sufficient for optimality.

3.4. Approximate (s,S) Policies Using an α -Service Level

In this section, approximate models are described for the warehouse/stores system where the warehouse policy and the store policies will be related through the α -service level. If there were to be no shortages allowed at the warehouse, each store policy could be determined independent of the warehouse. However, since a warehouse shortage will effect the length of the replenishment leadtime to the store(s), the resulting leadtime distribution must be observed.

Here, approximate ordering policies are obtained for the stores and the warehouse using the power approximations developed by Ehrhardt (1979) and extended by Schneider and Ringuest (1990). Ehrhardt (1979) showed that these approximations, based only on the first and second moments of the demand distribution, are very close to the optimal policies over a wide range of parameters. The power approximation is based on a non-linear regression of the optimal policy (D^*, s^*) on its inventory parameters. Ehrhardt (1979) also showed that the power approximations work well outside the parameter ranges used for the regression.

A revised power approximation was derived by Ehrhardt and Mosier (1984) for use when a stockout cost is given to obtain s_i and D_i when a stockout cost is known for the stores. Let $R_{iL,+1}$ be the demand during the stochastic leadtime plus the review time for store i . The expected value of the demand is shown by Schneider, et al. as

$$\mu_{iL,+1} = E[R_{iL,+1}] = (E[L_i] + 1)\mu_i \quad (19)$$

and the variance is

$$\sigma_{iL,+1}^2 = Var[R_{iL,+1}] = E[L_i + 1]\sigma_i^2 + Var[L_i]\mu_i^2 \quad (20)$$

The power approximation of Ehrhardt and Mosier (1984) for D_i is given by

$$D_i = 1.3\mu_i^{0.494} \left(\frac{K_i}{h_i} \right)^{0.506} \left(\frac{1 + \sigma_{iL,+1}^2}{\mu_i^2} \right)^{0.116} \quad (21)$$

and the power approximation for s_i when the stockout cost (p_i) is known is given by

$$s_i = 0.973\mu_{iL_i+1} + \sigma_{iL_i+1} \left(\frac{0.183}{z_i} + 1.063 - 2.192z_i \right) \quad (22)$$

where z_i is

$$z_i = \sqrt{\frac{D_i}{\left(\frac{\sigma_{iL_i+1} p_i}{h_i} \right)}} \quad (23)$$

In order to obtain a power approximation for the (s,S) policy at the warehouse, the first and second moments of the demand distribution must be derived. Based on Schultz (1983), the expected demand at the warehouse during the leadtime plus review time is

$$\mu_{wL_w+1} = (L_w+1) \sum_{i=1}^N \mu_i \quad (24)$$

and the warehouse demand variance is given by

$$\sigma_{wL_w+1}^2 = L_w \left\{ \sum_{i=1}^N \tau_i^2 - 2 \sum_{i=1}^N \mu_i^2 \right\} + \sum_{i=1}^N \tau_i^2 \quad (25)$$

where

$$\tau_i^2 = \sigma_i^2 + \frac{2D_i^2 \mu_i^2}{2D_i \mu_i + \sigma_i^2 + \mu_i^2}.$$

Again, following Ehrhardt and Mosier (1983), the power approximation for D_w is given by

$$D_w = 1.3 \mu_w^{0.494} \left(\frac{K_w}{h_w} \right)^{0.506} \left(\frac{1 + \sigma_w^2 L_w + 1}{\mu_w^2} \right)^{0.116}. \quad (26)$$

For a fixed α -service level at the warehouse, an approximate can be found for constraint (14) which must be solved because of the condition $\frac{\partial \Lambda}{\partial \alpha} = 0$. Schneider (1978) also provides the following approximation

$$\alpha = 1 - \frac{\int_{s_w}^{\infty} (r - s_w) f_w(r; L_w + 1) dr}{D_w + \frac{\sigma_w^2 + \mu_w^2}{2\mu_w}}, \quad (27)$$

where we search for the value of s_w that provides the appropriate α value.

When the α -service level is specified, the power approximation for s_w has been derived in Schneider, et al. (1995), as

$$s_w = \mu_w(L_w + 1) + p(e) \sigma_w L_w + 1 - \left(\frac{\delta \left(\frac{\sigma_w^2}{\mu_w} - 1 \right) (-1.95269 + 6.39059e)}{1 + 21.17036e} \right) \quad (28)$$

where $\delta(x) = \max(x, 0)$, e is

$$e = \frac{(1-\alpha) \left(D_w + \frac{\mu_w^2 + \sigma_w^2}{2\mu_w} \right)}{\sqrt{\sigma_{wL_w+1}^2}}, \quad (29)$$

and $p(e)$ is approximated using

$$p(e) = \frac{a_0 + a_1 w + a_2 w^2}{b_0 + b_1 w + b_2 w^2 + b_3 w^3 + b_4 w^4}$$

$$\text{where } w = \sqrt{\ln \left(\frac{25}{e^2} \right)}$$

and the coefficients

$a_0 = -5.3925569$	$b_0 = 1.000$
$a_1 = 5.6211054$	$b_1 = -7.2496485 \times 10^{-1}$
$a_2 = -3.8836830$	$b_2 = 5.0732662 \times 10^{-1}$
$a_3 = 1.0897299$	$b_3 = 6.6913686 \times 10^{-2}$
	$b_4 = -3.2912911 \times 10^{-3}$.

An asymptotic expression similar to equation (8) holds for the warehouse which allows us to obtain the following approximate expression for system total cost:

$$C(s_w, s_w, s_1, s_1, \dots, s_N, s_N) \approx K_w \left(\frac{\mu_w}{Q_w} \right) + h_w \left[D_w \left(1 - \frac{D_w}{2Q_w} \right) + s_w - \mu_{wL_w+1} + (1-\gamma_w)\mu_w \right] \\ + \sum_{i=1}^N \left\{ K_i \left(\frac{\mu_i}{Q_i} \right) + h_i \left[D_i \left(1 - \frac{D_i}{2Q_i} \right) + s_i - \mu_{iL_i+1} + (1-\gamma_i)\mu_i \right] + p_i(1-\gamma_i)\mu_i \right\} \quad (30)$$

as $D_i \rightarrow \infty$ for $i = 1, \dots, N$ and $D_w \rightarrow \infty$.

Equation (30) is used to find an approximately optimal α . If (30) is minimized with respect to α , approximately optimal ordering policies can be obtained by using the golden search method of Press, et al. (1988).

Through numerical analysis and simulation, Schneider, et.al. (1995) concluded that the approximations presented here are accurate enough for practical purposes. Also, the effect of the service level at the warehouse proved to be important in improving the ordering policies and decreasing the total cost of the system. They found that the total system cost is very sensitive to the warehouse service level if it is higher than the optimal level. However, if the warehouse service level is lower than the optimal, the resulting increase in stockout costs can be balanced by the adjustment of the stores' ordering policy through the lead time demand parameters. They also note that when the warehouse service level is high, adjustment of the stores' lead time demand parameters is less important since shortages at the warehouse are less likely.

3.5. Consideration of Aggregate Measures and Restrictions.

The models and their approximations discussed in the previous sections optimize the customer service level versus average per-period ordering, holding, and shortage costs under the assumption that sufficient aggregate resources such as investment capital, storage space, and labor capacity are infinite. This is typical of a majority of the inventory models found in the literature, but is rarely the case for practitioners who are often bound by policy and/or resource constraints.

Starr and Miller (1962) derived an 'optimal policy curve' for items under deterministic demand. Schrady and Choe (1977) studied continuous review systems with constraints. Gardner and Dannenbring (1979) extended the approach of Starr and Miller to a stochastic model for a system with continuous review including aggregate objectives and constraints. Alscher and Schneider (1982) considered a multi-item system using (s,S) ordering policies where shortage costs are difficult to estimate. Observing the conflicting objectives of service level, number of replenishments, and average stock on hand, they presented an algorithm for computations which resulted in diagrams and/or tables from which managers could select an appropriate service level and its corresponding ordering policy. Schneider and Rinks (1989) extended the method of Gardner and Dannenbring to a periodic review (s,S) , multi-item system with two objectives: (1) minimizing the cost of a policy and (2) maximizing the service level, where average order handling workload and storage space are limited. They derived approximate solutions based on the asymptotic properties of renewal theory with their results presented in the form of an optimal policy surface showing the tradeoff between policy cost and service level while satisfying the order handling workload and storage space constraints.

3.5.1. Aggregate Measures at the Store Level.

In the literature discussed above, inventory ordering policies have been developed for a single stocking location with one or more constraints. These concepts can be extended to the two-echelon system presented in this thesis. The average capital investment in inventories and available storage space capacity will be considered for

each of the stores. The notation used follows that of the previous formulae for each store i , and will now consider individual products j .

3.5.1.1. Capital Investment Limit at the Store Level.

For any stochastic demand inventory system, the on-hand inventory level at an arbitrary moment in an order cycle is a random variable which by nature is largest just after a replenishment and lowest just before a replenishment. Since this holds for all products in stock, the amount of capital investment in total inventories constantly fluctuates so that a meaningful measure must be based on average inventory levels over all products j .

The inventory of a part at store i is denoted by I_i , and the expected inventory level is $E[I_i]$. The subscript for item j is deleted in the following formulae for simplicity. An exact formulation for the expected inventory level has been expressed by Roberts (1962) as

$$E[I_i] = \frac{\int_0^{S_i} (S_i - r) f_i(r, L_i + 1) dr + \int_0^{D_i} \int_0^{S_i - x} (S_i - x - r) f_i(r, L_i + 1) dr m_i(x) dx}{1 + M_i(D_i)}, \quad (31)$$

where $f_i(r, L_i + 1)$ is the distribution of the demand during the replenishment lead time plus the review time, S_i is the order-up-to level, $D_i = S_i - s_i$, and $M_i(D_i)$ is the renewal quantity. Computational methods have been developed by Veinott and Wagner (1965) for calculating exact optimal policies and their operating characteristics. However, the computational effort required is prohibitive for practical implementation. Also, the

exact computation requires the complete specification of the demand distribution which, in practice, is usually unlikely to be available.

A good approximation for the expected inventory level using power approximations can be expressed as in Schneider and Rinks (1995) as

$$E[I_i] = D_i \left(1 - \frac{D_i}{2Q_i} \right) + (s_i - \mu_{iL_i+1}) + (1 - \gamma) \mu_i \quad (32)$$

where, using the power approximations of Ehrhardt and Mosier (1984)

$$D_i = 1.3 \mu_i^{0.494} \left(\frac{K_i}{h_i} \right)^{0.506} \left(\frac{1 + \sigma_{iL_i+1}^2}{\mu_i^2} \right)^{0.116}, \quad (33)$$

$$Q_i = \left(D_i + \frac{\mu_i^2 + \sigma_i^2}{2\mu_i} \right), \quad (34)$$

and whenever a γ -service level at the stores is specified, $(1 - \gamma_i)$ as used in equation (32) is computed using

$$1 - \gamma_i = \frac{\int_{s_i}^{\infty} (r - s_i)^2 f_i(r; L_i + 1) dr}{2\mu_i Q_i}, \quad (35)$$

and searching for the appropriate s_i .

An estimate of the optimal policy parameter s_i is obtained when the unit stockout cost (p_i) is known using

$$s_i = 0.973\mu_{iL_i+1} + \sigma_{iL_i+1} \left(\frac{0.183}{z_i} + 1.063 - 2.192z_i \right), \quad (36)$$

where z_i is

$$z_i = \sqrt{\frac{D_i}{\left(\frac{\sigma_{iL_i+1} p_i}{h_i} \right)}}. \quad (37)$$

Although the Schneider and Rinks approximation for expected inventory has been shown to be quite accurate, it may still be considered as too complex for practitioners. A somewhat less complex approximation may be applicable for many of the parts in the inventory. Once the operating parameters (s, S) of the inventory policy have been determined, the average inventory level may be determined by estimating the expected maximum inventory and the expected minimum inventory and averaging them. These levels would be observed immediately before and immediately after a replenishment arrives. Since the expected inventory at the time of review is s_i , the expected inventory at the end of the lead time just before a replenishment arrives is ($s_i - \mu_{iL}$), the reorder point minus the expected demand during the lead time. Likewise, immediately after the replenishment arrives, the expected on hand inventory is ($S_i - \mu_{iL}$), the order-up-to level minus the expected demand during the lead time. Averaging

these two expected values should give a reasonable estimate of the expected inventory level, so that

$$E[I_i] = \frac{(s_i - \mu_{iL}) + (S_i - \mu_{iL})}{2} = \left(\frac{s_i + S_i}{2} \right) - \mu_{iL} . \quad (38)$$

The accuracy and applicability of both approximations will be compared in Chapter 4. With a reasonably accurate estimate of the expected inventory level, the expected capital investment in part j at store i , $E[CI_{ij}]$, is given by

$$E[CI_{ij}] = P_j E[I_{ij}] \quad (39)$$

where P_j is the unit purchase price for part j . It follows that the total expected capital investment in inventories at store i , $E[CI_i]$, can be expressed by

$$E[CI_i] = \sum_{j=1}^n E[CI_{ij}] , \quad (40)$$

here there are n stock keeping units in the inventory at store i .

3.5.1.2. Storage Space Requirements at the Store Level.

For storage space limits at a stocking point, a common approach found in the literature uses the expected storage requirement from average stock levels. If capacity decisions such as purchasing or leasing warehouse space are based on these expected values, there would be capacity shortfalls whenever individual shipments or total inventory levels are larger than average. As illustrated in Figure 3.3, at the end of the

review period, time t_1 , the inventory position is reviewed. If the the inventory position is at, or below s , an order is placed, raising the inventory position to S . After passage




Figure 3.3. (s,S) Inventory Policy Replenishment Cycle

of the lead time, L , the order arrives at time t_2 which raises the on-hand inventory level. At this point, inventory position and on-hand are equal. Under an (s,S) inventory policy, the order-up-to level S would be the absolute maximum that the on-hand inventory could reach assuming that lead time demand were zero. Subtracting the mean lead time demand μ_L from S would provide the expected level. A $(1-\delta) \times 100\%$ upper confidence interval for the expected on-hand inventory of part j at store i is given by

$$E[I_{ij,1-\delta}] = \min \left\{ \begin{array}{c} S_{ij} - (\mu_{ijL_i} - z_{1-\delta} \sigma_{ijL_i}) \\ S_{ij} \end{array} \right\}. \quad (41)$$

It follows that the $(1-\delta) \times 100\%$ upper confidence interval for the total storage volume required by the inventories at store i would be expressed using

$$E[V_{i,1-\delta}] = \sum_{j=1}^n V_j(S_{ij} - \mu_{ijL_i}) + z_{1-\delta} \sqrt{\sum_{j=1}^n V_j^2 \sigma_{ijL_i}^2}, \quad (42)$$

where V_j is a volume factor for product j and total on-hand inventory is approximately normally distributed.

3.5.2. Aggregate Measures at the Warehouse.

The addition of a central warehouse to an inventory distribution system will require an investment in additional plant and equipment that in most cases is substantial. Warehouse space and delivery vehicles must be leased or purchased along with material handling equipment, office equipment, and labor. Additions of these types are usually acquired in discrete quantities, i.e., a volume of warehouse is leased or not, a truck of a certain size is purchased or not, an employee is hired or not. At any given point in time, storage capacity is limited by some number of cubic units as are deliveries to the stores limited by truck volume and weight capacities. Additionally, the amount of capital invested in inventories at the warehouse may be restricted by some standing company policy or the short term need to free up working capital.

3.5.2.1. Capital Investment Limit at the Warehouse.

For the warehouse, the average capital investment in inventory may be an important management issue. Like the stores, inventory levels at the warehouse will

fluctuate during the review period so that a useful measure of inventory investment must be based on average levels. The inventory level for a part at the warehouse is given by I_w , and the expected warehouse inventory level is $E(I_w)$. The asymptotic value ($D_w \rightarrow \infty$) for the expected inventory level at the warehouse has been given by the following approximation of Schneider and Rinks (1989) as

$$E[I_w] = D_w \left(1 - \frac{D_w}{2Q_w} \right) + (s_w - \mu_w) + (1 - \gamma_w) \mu_w. \quad (43)$$

In order to obtain the power approximation for the (s,S) policy at the warehouse, the first and second moments for the demand distribution during the review time at the warehouse are needed. Based on Schultz (1983), the expected demand at the warehouse during the review time, μ_w , is shown as the sum of the expected demands at the stores during the review time

$$\mu_w = \sum_{i=1}^N \mu_i, \quad (44)$$

and the variance of demand at the warehouse is

$$\sigma_w^2 = \sum_{i=1}^N \tau_i^2, \quad (45)$$

where

$$\tau_i^2 = \sigma_i^2 + \frac{2D_i^2\mu_i^2}{2D_i\mu_i + \sigma_i^2 + \mu_i^2}. \quad (46)$$

Using Ehrhardt and Mosier (1984) the power approximation for D_w is

$$D_w = 1.3\mu_w^{0.494} \left(\frac{K_w}{h_w} \right)^{0.506} \left(\frac{1 + \sigma_{wL_w+1}}{\mu_w^2} \right)^{0.116}, \quad (47)$$

Q_w is calculated using

$$Q_w = \left(D_w + \frac{\mu_w^2 + \sigma_w^2}{2\mu_w} \right), \quad (48)$$

and s_w is given by

$$s_w = \mu_w(L_w+1) + p(e)\sigma_{wL_w+1} - \left(\frac{\delta\left(\frac{\sigma_w^2}{\mu_w} - 1\right)(-1.95269 + 6.39059e)}{1 + 21.17036e} \right) \quad (49)$$

where $\delta(x) = \max(x, 0)$, e is

$$e = \frac{(1-\alpha) \left(D_w + \frac{\mu_w^2 + \sigma_w^2}{2\mu_w} \right)}{\sqrt{\sigma_{wL_w+1}^2}}, \quad (50)$$

and $p(e)$ is approximated using

$$p(e) = \frac{a_0 + a_1 w + a_2 w^2}{b_0 + b_1 w + b_2 w^2 + b_3 w^3 + b_4 w^4} \quad (51)$$

$$\text{where } w = \sqrt{\ln\left(\frac{25}{e^2}\right)}$$

and the coefficients

$$\begin{aligned} a_0 &= -5.3925569 & b_0 &= 1.000 \\ a_1 &= 5.6211054 & b_1 &= -7.2496485 \times 10^{-1} \\ a_2 &= -3.8836830 & b_2 &= 5.0732662 \times 10^{-1} \\ a_3 &= 1.0897299 & b_3 &= 6.6913686 \times 10^{-2} \\ & & b_4 &= -3.2912911 \times 10^{-3}. \end{aligned}$$

The expected level of capital investment for part j at the warehouse would then be given by

$$E[CI_{wj}] = P_j E[I_{wj}] , \quad (52)$$

where P_j is the unit purchase price for part j . It follows that the total expected capital investment in inventories at the warehouse, $E[CI_w]$, can be expressed using

$$E[CI_w] = \sum_{j=1}^n E[CI_{wj}] , \quad (53)$$

where there are n stock keeping units in the inventory at the warehouse.

3.5.2.2. Storage Space Requirements at the Warehouse.

The availability of storage space at the warehouse will be limited by discrete volumes at any given time. As discussed for the stores, a measure of warehouse capacity requirements based on expected inventories may give a poor estimate of space needs particularly when ordering variances are large. A more useful estimate based on a $(1-\delta) \times 100\%$ upper confidence interval for the expected maximum on-hand inventory at the warehouse for an item j is given by

$$E[I_{wj,1-\delta}] = S_{wj} - (\mu_{wjL_w} - z_{1-\delta} \sigma_{wjL_w}) , \quad (54)$$

and the $(1-\delta) \times 100\%$ confidence interval of the expected volume required for all inventories at the warehouse is given by

$$E[V_{w,1-\delta}] = \sum_{j=1}^n V_j (S_{wj} - \mu_{wjL_w}) + z_{1-\delta} \sqrt{\sum_{j=1}^n V_j^2 \sigma_{wjL_w}^2} , \quad (55)$$

where V_j is a volume factor for product j .

3.5.2.3. Capacity Measures for Warehouse Deliveries to Stores.

Under the original N-store system, orders are individually placed with the manufacturer which are shipped to each store and arrive at the end of the replenishment lead time. The delivery vehicles are owned by, or are agents for, the manufacturer so that the store management have little or no concern or responsibility for balancing the

volume of their orders with the capacity of the vehicles. Their sole concern is that their orders arrive intact and in a timely fashion. With the introduction of a company operated central warehouse, management assumes the responsibility for the direct store replenishments while continuing to place orders with and receive deliveries from the manufacturer on behalf of the entire store system. Among the many fixed costs associated with adding a central warehouse are an improved or modified information system and the lease, purchase, or contracting of vehicles for handling and transporting store orders.

At the end of the review period, the inventory position of every item at each store is noted. For those items whose inventory position is at or below their order flag level s , an order is placed with the warehouse which will raise the inventory position to its maximum level S . If the stores' inventories are computerized and store inventory data are available to the warehouse, store orders could be virtually automatic since the order policy parameters (s, S) for each item at each store could be integrated into the information system. The intricacies of the information system are, however, potentially quite numerous and are beyond the scope of this thesis. Whether they are electronically transferred or physically carried to the warehouse, the store orders must be processed and filled by the warehouse.

The store orders will be picked and assembled by warehouse staff, loaded on a transport vehicle, and routed under the assumption that all replenishments will arrive at the stores on the same day. It is possible that, because of the warehouse's location or the capacity limits of the transport, more than one route may be required to make all

of the store deliveries. Under the regulations of the U.S. Department of Transportation, the on-duty hours per 24-hour period for commercial drivers are limited. As such, it may be that only a subset of the N stores may be resupplied by each of several transports in order to replenish store inventories.

Transport vehicles are available in many configurations. However, they are similar in that they are each restricted by both volume and weight capacities. It becomes a matter of some importance to have an estimate of both the volumes and weights that can be expected to be required of the warehouse transport vehicles. As with warehouse capacity, considering only the expected volume of shipments would result in understating the capacity needed when individual orders are larger than average. It would be prudent, therefore, to have measures more closely reflecting the maximum capacities required, particularly if these capacity needs will be used in the vehicle selection decision. Using a $(1-\delta) \times 100\%$ upper confidence interval, the maximum expected units of item j shipped per cycle from the warehouse is

$$E(I_{wj,1-\delta}) = \mu_{wj} + z_{1-\delta} \sigma_{wj} \quad (56)$$

where

$$\mu_{wj} = \sum_{i=1}^N \mu_{ij} \quad (57)$$

and using equations (46) and (47)

$$\sigma_{wj}^2 = \sum_{i=1}^N \left[\sigma_{ij}^2 + \frac{2D_{ij}^2 \mu_{ij}^2}{2D_{ij} \mu_{ij} + \sigma_{ij}^2 + \mu_{ij}^2} \right]. \quad (58)$$

The $(1-\delta) \times 100\%$ upper confidence interval measure for the total expected shipping volume per cycle from the warehouse to the stores is then given by

$$E[V_{sw,1-\delta}] = \sum_{j=1}^n V_j \mu_{wj} + z_{1-\delta} \sqrt{\sum_{j=1}^n V_j^2 \sigma_{wj}^2}, \quad (59)$$

where V_j is a volume factor for item j . By simply replacing the volume factor for item j , V_j , with a weight factor for item j , W_j , in equation (58), the $(1-\delta) \times 100\%$ upper confidence interval for shipping weight can be obtained.

For individual store orders, the transport volume required is given by

$$TV_i = \sum_{j=1}^m V_j Q_{ij}, \quad (60)$$

where TV_i is the total volume of the order, m is the number of part types in the order, V_j is the volume factor for item j , and Q_{ij} is the order quantity for item j from store i . It follows that the total shipping volume for a delivery route consisting of N stores, TV_r , can be expressed by

$$TV_r = \sum_{i=1}^N TV_i, \quad (61)$$

where, again, replacing the notation for volume with that for weight would give the total shipping weight for the delivery route.

4. ANALYSIS OF INVENTORY STRATEGIES

In this chapter, an analysis of the inventory strategies described in chapter three is presented. In order to determine the potential success of the models when implemented they must be tested for their ability to approximate reality. The goal here is to verify the accuracy of the analytic approximations by comparing their results with those of simulation models that include the relevant realities of inventory management while operating under the policies generated by the analytic models.

During the simulation runs, the major inventory performance characteristics will be monitored, including on-hand inventory levels, ordering and holding costs, frequency and magnitude of backlogs, and orders per period from the stores to the warehouse and from the warehouse to manufacturing. These measurements will form the basis for comparison to the analytic models. The ordering policies will be generated by following a two-level input parameter scheme and the sensitivity of the performance measures to changes in the input parameters will be examined using a two-level factorial design.

4.1. Experimental Design for Analysis of Inventory Policies

In implementing a general factorial design, it is common practice to select a fixed number of levels or 'versions' for each of a number of variables (factors) and then run experiments using all possible combinations. A factorial design using k variables each set at two levels results in an experiment requiring 2^k experimental runs in order to obtain results from all possible combinations of factor levels. The number of runs required by a full 2^k factorial design increases geometrically as k is increased.

However, when k is not small the desired information can often be obtained by performing only a fraction of the full factorial design since there tends to be a hierarchy, in terms of absolute magnitude, of the size of the effects. Main effects tend to be larger than two-factor effects, which in turn tend to be larger than three-factor interactions, and so on. Also, when a moderately large number of variables is introduced into a design, it often happens that some have no distinguishable effects at all. There tends to be redundancy in a 2^k design whenever k is not small in terms of an excess number of interactions that can be estimated and also an excess number of variables that are studied. To exploit this redundancy, a fraction of the full factorial design can be deleted from the experimental design so that the effects of some of the multi-factor interactions are ignored. A two-level design of this type is known as a 2^{k-p} fractional factorial design where k main factors can be analyzed using $k-p$ experiments.

In analyzing the two-echelon inventory system being studied here, it was determined to study the effects of ten input parameters to the models with each set at two levels. A summary of the ten factors and their settings is shown in Table 4.1. The ten factors are arranged into a 2^{10-5} fractional factorial design requiring 32 experiments. Using 32 experiments provides a design of resolution $R=IV$ which allows us to study the ten main factor effects with no confounding (aliasing) of the main effects with any two-factor interactions. However, the resolution IV design does confound two-factor interactions with other two-factor interactions.

The factor levels for each of the variables were selected with two goals in mind. First, to provide results with enough variation to show a difference between the levels,

and second, to implement the analytic models using parameters that were likely to be encountered in reality. Factor A is the fixed length review period, R_i , for the stores in days and is set at the two levels of one day and five days. Factor B is the fixed length lead time, L_i , for delivery of replenishment orders from the warehouse to the stores. The lead times are also set at the levels of one day and five days. Factor C is the value of the γ -service level required at the stores and is set at the two levels of 0.80 and 0.98. Factor D is the cost to place an order, K , which is set at \$3 and \$6. Factor E is the annual holding cost rate, h_r , which is set at 10% and 40%. The order cost and holding rate are assumed equal for both the store echelon and for the warehouse echelon in a given experiment.

Table 4.1.
Factor Levels for Experimental Design

Factor	Input Parameter	L-	L+
A	Store review period, days	1	5
B	Store lead time, days	1	5
C	Store γ -service level	0.80	0.98
D	Order cost	3	6
E	Annual holding rate	0.10	0.45
F	WH review period, days	3	5
G	WH lead time, days	3	5
H	Warehouse α -service level	0.90	0.95
I	Number of stores	5	15
J	Demand rate	slow	fast

Factor F is the fixed length review period, R_w , for the warehouse which is set at three days and five days and Factor G is the fixed replenishment lead time for orders to arrive at the warehouse from manufacturing, L_w , also set at 3 or 5 days. Factor H is the α -service level that is required at the warehouse where settings of 0.90 and 0.95 are used. Factor I is the number of stores in the system, N , that will be placing orders to the warehouse. Systems with five and fifteen stores are used. Finally, Factor J is the demand rate for the parts in the system. Slow moving parts are defined as those whose mean annual demand is $3 \leq \mu \leq 16$ while fast moving parts are defined by $17 \leq \mu \leq 500$ annual demand. For each experimental run, the total population of parts in the system are either fast moving or slow moving. The complete matrix for the 10 factor, 32 experiment design is presented in Appendix C.

4.2. Simulation Modeling Assumptions

Following the experimental design described in the previous section, the annual demand for a set of test parts and their ordering policies were generated. For each experiment, there are 30 parts in the system, with all of the parts stocked at each store. For each part in an experiment, a mean annual demand was generated from a uniform distribution with limits determined by the ranges for fast or slow moving parts. The demand for each store was then randomly generated from a second uniform distribution with a range of $\pm 30\%$ of the mean previously determined. This allowed us to avoid comparing identical store demands by providing for a 'controlled' randomness between the stores for each part.

Using the annual demands for each part at each store, the (s,S) ordering policies were computed using the approximations described in Chapter 3. Since the models require the first and second moments of the demand distribution, store demands for all parts were assumed to be Poisson distributed where $\mu = \sigma^2$. For this study, using the Poisson for all part demands is a reasonable assumption since even the 'fast' moving parts are relatively slow moving in general inventory terms. Additionally, a 260 day operating year and a common part cost of \$5 are assumed. For the 32 experiments, 10,560 individual part ordering policies were computed and, for each of the 32 experiments, five simulation replications were run.

The simulation for each experiment was conducted as follows. Customer demands which are independent and identically distributed (IID) for each part j occur daily at each store i following a Poisson distribution with mean μ_{ij} . The on-hand inventory level for each part is reduced by the random daily demand amount. If the on-hand level is 0 when a demand occurs, a customer backorder is recorded. Upon reaching the end of the store review period, the inventory position of each part is reviewed and orders are placed to the warehouse if the inventory position is at or below the reorder point. The store orders are arranged in a random sequence at each cycle to eliminate bias in the case of warehouse shortages.

At the review time, if there are backorders at the warehouse from previous store orders, the backordered parts are filled first with existing warehouse stock. If current warehouse stock cannot fully fill the backorder, partial filling of the backordered amount is allowed. Once any backordered amount is filled, the current order is

considered. If the order quantity cannot be completely filled from existing stock, the entire order amount is backordered. The total of backordered and current order parts are shipped to the store, scheduled to arrive following the fixed lead time. When the replenishment arrives at the store, any customer backorders are immediately filled and the balance is placed into the on-hand inventory at the store.

Upon reaching the end of the warehouse review period, the inventory position of each part at the warehouse is reviewed and an order is placed to manufacturing if necessary. Manufacturing is assumed to have an infinite supply of parts so that there are no shortages and all orders are completely filled and shipped so that they arrive at the warehouse following the fixed lead time. Upon arrival at the warehouse, the parts are immediately placed into the warehouse on-hand inventory. As stated above, backorders for store orders are filled at the end of the store review period. Because of this sequence of events, both on-hand inventory and store backorders may be positive simultaneously at the warehouse.

As the simulation progresses, a number of values of interest are collected. Following each day's random demands, the total number of parts on hand at each location is recorded. For each store review period, the volume of orders to the warehouse for each part from all of the stores in the system is recorded. Likewise, for each warehouse review period, the volume of orders to manufacturing for each part is recorded. Finally, at the end of each of the five replications of an experiment, the daily holding cost, ordering cost, total cost, and gamma service levels for each part at

each location is computed and recorded. These statistics provide the basis for the analysis of the system performance that follows.

4.3. Comparison of Analytic Approximations to Simulation Results

In this section, the results of the analytic approximations are compared to the results given by the simulation experiments. Comparisons are presented for holding, ordering, and total costs, the estimates for average aggregate on-hand inventory, and the gamma service levels.

4.3.1. Accuracy of Approximated System Costs

The costs from each of the 10,560 ordering policies were compared individually for their deviation from simulated costs. A relative frequency histogram of the deviations for holding costs is shown in Figure 4.1.




Figure 4.1. Relative Deviation Between Analytic and Simulated Holding Cost

These deviations were computed such that a negative deviation implies that the analytic approximation value is lower than the simulated value. The average holding cost deviation is -0.054 or -5.4%, so that over the entire set of order policies computed, on the average, the estimated holding costs are below the simulated holding costs. The maximum observed positive holding cost deviation was 1.261 and the maximum negative deviation was -0.944. 65% of the deviations are between ± 0.10 and 90% are between ± 0.20 .

A histogram of the daily average ordering cost deviations is shown in Figure 4.2. Negative ordering cost deviations reflect observations where the approximated




Figure 4.2. Relative Deviation Between Analytic and Simulated Order Cost

costs are lower than the simulated costs. The average of these deviations was 0.040 or 4.0% so that, on the average, the approximated order costs were higher than those

observed in the simulation. The maximum negative order cost deviation observed was -0.667 while the maximum positive deviation measured was 2.000. 61 % of all order cost deviations were within ± 0.10 and 80% were within ± 0.20 .

Finally, the histogram of relative deviations for total costs is presented in Figure 4.3. The total costs are the sums of the ordering and holding costs for each of the ordering policies in the study. Here again, negative deviations reflect values for the analytic approximations that are below the simulated results. For total cost, the




Figure 4.3. Relative Deviations Between Analytic and Simulated Total Costs

deviations are slightly better grouped about 0 than the deviations for holding and ordering costs individually. The average deviation for total costs was -0.035 or -3.5% so that the average total costs for the approximations were below those of the

simulation. 75% of the total cost deviations were within ± 0.10 and 93% were within ± 0.20 .

A fractional factorial design analysis was performed to estimate the effect of the various model factors (see Table 4.1) on total system costs from each of the 32 experiments. The results of the 2^{10-5} fractional factorial fit are shown in Table 4.2. The resolution IV design allows us to examine main effects that are not confounded with any other factors. But, since some two-factor interactions will be confounded with each other, the two-factors are excluded from the fit.

Table 4.2.
Fractional Factorial Effects for Average
Total Costs

Factor	Variable	Effect	t-value	P
Constant		52.698	25.120	0.000
A	Ri	-3.701	-0.880	0.379
B	Li	0.646	0.150	0.878
C	Gamma	4.517	1.080	0.238
D	K	19.201	4.580	0.000
E	hr	40.882	9.740	0.000
F	Rw	2.874	0.680	0.494
G	Lw	0.836	0.200	0.482
H	Alpha	3.336	0.800	0.428
I	N	47.750	11.380	0.000
J	DRate	66.554	15.860	0.000

Of the ten main effects, four are shown to be active. Changes in the demand rate (slow vs. fast moving parts) have the largest effect on total costs followed by the number of stores, the holding cost rate, and the order cost in order of absolute magnitude. All of the active factors have positive effect, implying that increasing the values of the parameters increases the effect on the total cost response. The analysis

of variance for the experiment shows a fairly large residual error, implying that the cumulative of the interaction effects is also significant. A tactic for future study of this system would be to redesign the experiment to include the active factors found here in a design that would allow high enough resolution to include the two-factor interaction effects with no confounding.

4.3.2. Estimations for Average Aggregate On-Hand Inventory

The holding costs for the inventory policies are computed using the values for the expected on-hand inventories as in equation (32) from Chapter 3. The deviations from the simulated holding costs in Figure 4.1 are, therefore, directly related to the computation of the expected inventory levels. A histogram of the deviations for the expected aggregate inventory for the analytic approximation versus the simulation is shown in Figure 4.4. The values compared here are the total expected inventory for all




Figure 4.4. Simulated Average Inventory Versus Equation 32

parts at each location for the 32 experiments. As can be seen in the histogram, the relative deviations for the analytic on-hand approximations are for the most part lower than the simulated values as the deviations are largely negative. The average deviation for this comparison is -0.098 or a -9.8% difference with standard deviation of 0.177. The largest negative deviation observed was -1.829 and the largest positive deviation was 0.02.

A similar comparison was made using the less complex approximation for expected on-hand stock of equation (38) with the histogram shown in Figure 4.5. The deviations for the expected on-hand approximations of equation (38), although also largely negative, are somewhat better distributed about 0. The average deviation for this comparison was -0.039 or -3.9% with a standard deviation of 0.167. The maximum negative deviation was -1.27 and the maximum positive deviation was 0.099.




Figure 4.5. Simulated Average Inventory Versus Equation 38

Clearly, both of the approximations for the expected on-hand inventories were, on the average, less than the simulated on-hand results and show similar variance. However, the approximations of equation (38) are somewhat closer to the simulated expected on-hand amounts.

4.3.3. Requirement of 'Large' D for (s,S) Order Policies

An essential factor for the application of the asymptotic ordering policy models of Chapter 3 is the requirement 'large' D as computed by equations (9) and (26). Schneider, et al. (1995) advise that a minimal value for D is provided whenever $\frac{D_i}{\mu_i} \geq 1.5$ so that the difference between the order-up-to level and the reorder point divided by the average review period demand exceeds 1.5. If this ratio is not met or exceeded, the application of the (s,S) policy is not appropriate since the item being managed by the policy would tend to be ordered in every review cycle. The values for D/μ were computed for each policy in the study and averaged across all parts in each experiment. The results are shown in Table 4.3 along with the total system cost deviations for each experiment. The computations for the order policies at the store level are independent for each item. However, the warehouse ordering policy for each part results from the aggregation of the demands at all stores in the system. The averaged warehouse D/μ values are therefore presented separately. Our intent here is to see whether the various magnitudes of the average D/μ ratios show any relationship with the percent deviations for total cost.

A fractional factorial design analysis for the average D/μ ratios at the stores shows that changes in the store review period and the demand rate have the largest, and

nearly equal, effects on the D/μ ratios. As each of these is increased, the D/μ ratio decreases. As the holding rate increases, the D/μ ratios also decrease with an effect

Table 4.3.
Total Cost Percent Deviations versus Average D/Mu

Exp	Total Cost % Deviations	Average D/Mu	
		Stores	Warehouse
1	-3.93%	123.49	4.89
2	-8.61%	11.60	11.52
3	-11.08%	700.30	37.68
4	-3.22%	104.76	43.56
5	-6.91%	792.78	83.92
6	-3.46%	90.54	19.53
7	-6.74%	96.48	7.43
8	-5.71%	14.00	6.64
9	-7.57%	911.71	107.58
10	-3.57%	131.13	28.59
11	-0.85%	98.69	8.16
12	-4.92%	18.02	9.16
13	-12.16%	88.31	3.92
14	-3.72%	15.16	16.36
15	-13.02%	1041.81	59.69
16	0.07%	143.25	60.74
17	-6.69%	323.09	17.63
18	-9.79%	47.87	17.12
19	-7.72%	46.18	5.99
20	-25.50%	7.49	1.73
21	-7.57%	35.45	2.47
22	-17.49%	5.66	2.76
23	-20.17%	331.43	8.97
24	-7.28%	41.53	30.18
25	-4.42%	52.69	3.78
26	-15.40%	8.37	4.31
27	-15.62%	478.30	13.69
28	-6.95%	67.59	48.40
29	-5.11%	437.30	25.64
30	-4.44%	68.81	24.72
31	-3.98%	65.77	8.64
32	-16.26%	9.12	2.26

approximately half that of the review period and demand rate. All other factor effects are negligible.

For the warehouse, the largest factor effect on the D/μ ratios is the demand rate followed by the holding rate effect, the warehouse review period, and the number of stores in the system. The former three effects are nearly equal and approximately half the magnitude of the demand rate effect. Increases in all of these parameter levels result in decreases in the ratios.

It is interesting to note, that the experiments with the largest total cost deviations occur when the D/μ ratios are smallest as in experiments 20 and 32. However, there is little or no correlation (less than ± 0.23) between the D/μ ratios for either the stores or the warehouse and the cost deviations.

4.3.4. Effect of Parameters on System Gamma Service Levels

The gamma-service levels for all locations in the system were input parameters used by the analytic approximations. For the stores, the gamma levels were explicitly determined at the levels of 0.80 or 0.98. For the warehouse, the gamma service level is determined by first setting the alpha-service at the level of either 0.90 or 0.95, computing the ordering policy, and calculating the resulting gamma-service level using equation (11). The anticipated gamma-service level for each warehouse item is unique depending on the reorder point and the aggregate demand distribution. The service levels from the approximations, the simulated service levels, and their percent deviations for the stores and the warehouse are shown in Table 4.4. For each of the 32

experiments, the results are the averages of the individual part measures across all locations in the system.

Table 4.4. Analytic versus Simulated Gamma Service Levels

Exp	Stores			Warehouse		
	Input Gamma	Simulation Average	Relative Deviation	Analytic Average	Simulation Average	Relative Deviation
1	0.8	0.9751	0.2189	1.0000	0.8222	-0.1778
2	0.8	0.7535	-0.0581	0.9681	-0.2673	-1.2761
3	0.8	0.9677	0.2097	0.9395	-0.1425	-1.1517
4	0.8	0.8965	0.1206	0.9597	-1.8670	-2.9455
5	0.98	0.9943	0.0146	0.9584	-0.2991	-1.3121
6	0.98	0.9525	-0.0280	0.9092	-1.3571	-2.4926
7	0.98	0.9773	-0.0027	0.9872	0.2213	-0.7759
8	0.98	0.9524	-0.0282	0.9997	0.4820	-0.5178
9	0.8	0.9967	0.2459	0.9554	-0.1292	-1.1353
10	0.8	0.9699	0.2124	0.9777	-0.3320	-1.3395
11	0.8	0.9407	0.1759	1.0000	0.4555	-0.5445
12	0.8	0.7741	-0.0323	0.9881	-0.1422	-1.1439
13	0.98	0.9897	0.0099	0.9977	0.6890	-0.3094
14	0.98	0.9453	-0.0354	0.9994	-0.4739	-1.4741
15	0.98	0.9777	-0.0024	0.9611	-0.1312	-1.1366
16	0.98	0.9271	-0.0540	0.9684	-0.6630	-1.6847
17	0.8	0.9894	0.2367	0.9107	-1.0823	-2.1885
18	0.8	0.9309	0.1637	0.9918	-0.0838	-1.0845
19	0.8	0.9188	0.1486	1.0000	0.8423	-0.1577
20	0.8	0.8544	0.0679	0.9980	0.8829	-0.1152
21	0.98	0.9905	0.0107	0.9997	0.7987	-0.2010
22	0.98	0.9664	-0.0139	1.0000	0.8273	-0.1726
23	0.98	0.9427	-0.0381	0.9989	0.6386	-0.3607
24	0.98	0.7933	-0.1905	0.8404	-1.5340	-2.8254
25	0.8	0.9580	0.1975	1.0000	0.8907	-0.1093
26	0.8	0.8204	0.0255	0.9952	-0.0370	-1.0371
27	0.8	0.9499	0.1874	0.9445	-0.8539	-1.9042
28	0.8	0.8320	0.0400	0.9473	-0.7738	-1.8168
29	0.98	0.9894	0.0096	0.9912	-0.7912	-1.7982
30	0.98	0.9437	-0.0371	0.9014	-0.6486	-1.7195
31	0.98	0.9714	-0.0088	0.9824	0.2577	-0.7377
32	0.98	0.9605	-0.0199	1.0000	0.7707	-0.2292

For the stores, the 'input gamma' column of Figure 4.4. shows the levels of the gamma factor and is the target service level for all store parts. For the percent deviations, a positive value indicates that the simulated service levels were greater than the input values. Across all experiments, the average deviation was 0.054 so that, in general, the simulated gammas were slightly higher than the input values. The factor analysis of the percent deviations shows that, as might be expected, the input gamma value has the largest effect on the deviations. Experiments with the higher setting of 0.98 tend to have the smallest deviations from the input levels. Additionally, the level of the store review time is also an active factor. Experiments with a daily store review tend to have smaller deviations than those with a five day review period. The review period effect is approximately half the magnitude of the input gamma effect while all other factors are insignificant.

For the warehouse, the averages of the analytic average gamma-service levels, the simulation averages, and their relative deviations are shown. The anticipated gammas for nearly all experiments is quite high. In only two experiments, is the average gamma below 0.93. The simulated gammas are quite different as evidenced by the relative deviations for the warehouse where the expected gamma was not met or exceeded in any of the 32 experiments. The gamma-service level is defined as $1 - (\text{average cumulative backlog} / \text{average demand per period})$, and values between 0 and 1 would be expected. As shown, many of the average gammas for the warehouse are negative so that, for these cases, the average accumulated backlog exceeds the average period demand.

A fractional factorial design analysis for the warehouse gamma service levels shows that, by far, the demand rate is the most significant factor in the percent deviations. Experiments where the parts are slow-moving result in the largest errors between the analytic and simulated service levels. Cases where the store review period is five days and cases where the number of stores in the system is small, also show factor effects that are significant, although both of these effects are less than half the magnitude of the demand rate effect. For slow moving parts, the average per-period (daily) demand is likely to be only a small fraction of one part. Since parts enter and leave the inventory one unit at a time, the difference of one backordered unit at the warehouse can result in the gamma level being significantly changed. The average period demand for the warehouse is determined by the demand across all stores so that a lesser number of stores in the system reduces the average warehouse demand, amplifying the effect. It is also assumed that order sizes will be large enough to bring the stock level back to a positive level. When this does not occur, backorders will be 'on the books' for a longer period of time, increasing the average backlog and adversely affecting the service level. Overall, the warehouse service levels are not as large a concern as the store service levels since the warehouse 'customers' are the stores. The real focus of service levels for the system are the store customers. As long as the stores' customer service levels are maintained, the warehouse service levels will not cause great concern to practitioners. Also, there is an insignificant (0.23) correlation between the warehouse gammas and the store gamma-service levels so that the relatively poor service at the warehouse does not appear to directly translate poor service to

customers due to safety stock levels at the stores. However, increasingly extreme warehouse shortages will eventually cause excessive store shortages.

4.4. Characteristics of Aggregated Orders to the Warehouse

The characteristics of aggregated store orders to the warehouse and the resulting requirements for the transportation system were also considered. For each item, a store's orders will tend to be autocorrelated. There will be no order in several periods until, eventually, a large order will occur. The frequency of the orders is determined by the order sizes and the average demand and is reflected by the D/μ ratio. Orders for a single item will exhibit a 'lumpy' pattern over time. However, the aggregation of orders for many items will tend to smooth the pattern of total orders at the warehouse.

The mean and variance of demand at the warehouse for the review period are estimated using equations (57) and (58). During the simulations, the volume of total orders to the warehouse at each review period were recorded and the means and standard deviations of the orders computed. A comparison of the analytic computations and the simulated results is shown in Table 4.5.

The analytic and simulated results were very nearly equal across all of the experiments as shown by the percent deviations. The average relative deviation for the mean warehouse orders is 0.002 and the average deviation for the standard deviations is 0.006.

In order to use the resulting values for the mean and variance of orders to

Table 4.5.
Means and Standard Deviations for Warehouse Orders

Exp	Mean			Standard Deviation		
	Analytic	Simulated	Deviation	Analytic	Simulated	Deviation
1	315.69	315.01	-0.002	142.94	143.03	0.001
2	740.47	739.10	-0.002	193.00	187.41	-0.029
3	16.54	16.45	-0.005	19.66	20.03	0.019
4	25.29	25.62	0.013	21.63	20.41	-0.057
5	4.88	5.14	0.053	10.63	10.67	0.004
6	87.32	86.43	-0.010	43.39	38.19	-0.120
7	142.69	142.26	-0.003	102.13	104.00	0.018
8	1903.28	1893.60	-0.005	305.70	298.91	-0.022
9	6.00	5.90	-0.017	14.29	14.37	0.005
10	86.97	86.10	-0.010	47.05	45.24	-0.038
11	178.67	178.23	-0.002	137.08	138.59	0.011
12	2271.31	2277.00	0.003	370.50	407.31	0.099
13	563.83	562.55	-0.002	239.13	239.32	0.001
14	765.84	768.30	0.003	213.40	228.32	0.070
15	16.29	16.62	0.020	24.27	23.66	-0.025
16	25.30	25.46	0.006	22.50	24.07	0.070
17	5.66	5.60	-0.011	8.06	8.09	0.004
18	76.32	77.28	0.013	24.32	24.44	0.005
19	133.97	134.39	0.003	67.92	67.03	-0.013
20	1841.72	1848.10	0.003	208.80	198.89	-0.047
21	487.34	487.01	-0.001	129.05	128.97	-0.001
22	715.41	713.04	-0.003	126.25	121.54	-0.037
23	16.63	16.83	0.012	14.22	13.94	-0.020
24	29.80	30.32	0.017	16.88	15.63	-0.074
25	488.38	488.82	0.001	145.31	149.36	0.028
26	667.07	668.90	0.003	146.70	142.30	-0.030
27	16.61	16.48	-0.008	16.29	16.49	0.013
28	25.78	26.03	0.010	16.58	16.91	0.020
29	5.92	5.95	0.005	9.92	9.83	-0.009
30	74.99	74.46	-0.007	31.63	28.52	-0.098
31	135.85	135.26	-0.004	79.64	80.68	0.013
32	2301.00	2299.50	-0.001	259.50	274.49	0.058

determine a confidence limit for distribution system requirements, as proposed in

equation (59), we want to have some assurance that the use of the standard normal deviate, $z_{1-\alpha}$, is appropriate.

For each experiment, a histogram of the total units ordered from the warehouse per cycle was constructed using the simulated order sizes. The histograms and descriptive statistics for each are presented in Appendix D. Each simulation replication was run for 1000 days. Experiments with daily store review have 1000 data points for their histograms while those with review every five days have only 200 data points. This does not appear to affect the results. However larger sample sizes would produce better graphs. As can be seen from the histograms, a majority of the order patterns exhibit fairly normal distributions with the best of these being shown by experiments 1, 13, and 25. Conversely, some experiments -- such as 5, 16, and 17 -- show total units ordered patterns that are poorly behaved. In these histograms, high frequencies for zero order sizes is seen. These occur in experiments where both the demand rate of the parts is low and the number of stores in the system is small.

In the most extreme cases of low demand rate, orders for some parts may occur only once or twice over the length of the simulation. With 30 slow moving parts in the experiment, there are many order cycles where no parts are ordered. An exception to the low demand pattern is seen in the histograms for experiments 6 and 10. Although the parts in these experiments are also slow moving, the number of stores is large so that the distributions of orders are fairly normal. Overall, as the number of stores in the system increases, the assumption of normally distributed total orders to the warehouse appears to be reasonable and we can use equation (59) to compute

confidence limits for the total per-period shipping volumes from the warehouse to the stores.

4.5. Aggregate Inventories Generated by (s,S) Ordering Policies

In this section, an analysis of the characteristics of the aggregate on-hand inventories resulting from the application of the (s,S) ordering policies is presented. At the end of each day of the simulation, the total units of on-hand inventory for each location was recorded. Histograms and normal probability plots for each the 352 simulated locations (stores and warehouses) were constructed. Since the graphs for the daily on-hand amounts exhibited similar characteristics between experiments with like parameter settings, a representative set consisting of seven of the experiments (1, 5, 9, 14, 16, and 23) is provided in Appendix E.

The histograms for the aggregate on-hand inventories in the majority of the experiments show fairly well-behaved normal distributions. The best behaved distributions for the stores occur when the demand rate of the parts is high as in experiments 1, 14, and 22. The least well-behaved on-hand distributions for the stores is seen when the demand rates are low as seen in the histograms for experiments 5, 9, 16, and 23. For the warehouse, the tendency for normal behavior of the aggregate on-hand is largely determined by the number of stores in the system while the demand rate at the stores tends to be less of a factor. Experiments 1 and 23 have 15 stores each in the system. However, the demand rate of experiment 1 is high while that of 23 is low. As can be seen in the plots of the warehouse on-hand for these experiments, both exhibit strong normal behavior. The histograms for experiment 22 are included in the

appendix mainly because of the unusual behavior of the warehouse on-hand. There are five stores in this experiment and the demand rates are high. The histograms for the stores are as consistently normally distributed as any in the study but, at the warehouse, a bimodal distribution for on-hand is seen. The review periods for both the stores and the warehouse are set at five days for this experiment. The average orders from the stores to the warehouse and the average orders from the warehouse to manufacturing are identical at 713 units. The difference between the modes of the warehouse on-hand is also approximately 700 units so that flows into, and out of, the warehouse at equal spaced time periods may cause the total on hand to shift between the modes. However, this behavior is not seen for the warehouse in other experiments where the review periods are equal.

The distribution of the aggregate on-hand inventories is important to practitioners so that estimates for the required capital investment and storage space can be determined. Although the average on-hand amounts can be computed accurately with analytic models, the variance of inventories generated by (s,S) ordering policies has proven to be difficult except under very strict modeling assumptions. A formula for computing an upper confidence interval for maximum on-hand inventory was proposed using equation (42) for the stores and (55) for the warehouse.

The 95% upper confidence limit was computed for the aggregate inventory at each of the 352 locations in the experiment and was compared to the simulated distributions for on-hand discussed above. To illustrate these comparisons, the computations for one randomly selected location from each of the experiments are

shown in Table 4.6. The average inventory, $E(I)$, and the upper confidence limit, 95% CI, were computed using the analytic models of Chapter 3, while the standard

Table 4.6.
Z-Scores for Equation (42) Confidence Limits

Exp	Loc	Analytic		Simulated Std Dev	z-score for CI
		$E(I)$	95% CI		
1	8	766.18	1577.93	90.74	8.95
2	4	583.70	1332.77	87.78	8.53
3	8	341.62	685.48	25.83	13.31
4	5	226.26	461.83	22.26	10.58
5	3	328.56	653.91	40.64	8.01
6	9	242.96	486.33	25.37	9.59
7	2	1036.03	2062.42	104.30	9.84
8	3	686.20	1441.69	77.28	9.78
9	2	492.32	989.24	36.27	13.70
10	12	332.14	672.72	47.15	7.22
11	3	1462.36	3008.05	195.73	7.90
12	9	918.99	2026.05	119.56	9.26
13	14	1460.02	2946.36	149.40	9.95
14	1	976.68	2048.11	108.80	9.85
15	5	500.78	997.14	37.93	13.09
16	4	329.32	656.98	31.71	10.33
17	3	163.11	332.53	13.06	12.98
18	8	109.72	229.67	12.08	9.93
19	5	442.25	909.72	53.68	8.71
20	3	297.73	675.02	52.66	7.16
21	3	503.37	995.63	54.70	9.00
22	4	342.77	756.76	53.54	7.73
23	5	182.53	351.10	18.76	8.99
24	4	134.97	262.15	12.41	10.25
25	3	582.61	1215.51	73.48	8.61
26	5	406.07	921.48	65.89	7.82
27	15	234.21	469.82	31.62	7.45
28	3	156.25	318.51	20.72	7.83
29	3	240.51	478.49	21.14	11.26
30	11	161.74	325.67	17.43	9.41
31	4	674.14	1326.47	69.76	9.35
32	14	516.00	1105.20	69.66	8.46

deviations are those observed from the simulation. If the assumption of normality for the on-hand inventories is considered to be reasonable, standardized deviations between the confidence limit and $E(I)$ with values of approximately $2 \leq z \leq 3$ would be valuable in estimating the upper tail of the on-hand distribution and therefore the storage space requirements.

Unfortunately, as evidenced by the z-scores in Table 4.6 which were computed by dividing the difference between the confidence limit and the expected aggregate inventory by the simulated standard deviation of the aggregate on-hand, the upper confidence limits are in the range of 7 to 13 standard deviations and are of little use to us. The average z-score for all 352 confidence limits was 9.53. Investigation of the confidence limits suggests that for individual items, equation (41) is applicable for estimating the upper tail of the on-hand distribution for each item. However, by aggregating across all parts, we are assuming that the orders for every part in the inventory arrive simultaneously. Since this does not occur, and orders for various parts arrive at different points in time, the resulting variance of the total on-hand levels is much smaller than expected in expression (42)

5. SUMMARY AND CONCLUSIONS

A two-echelon inventory and distribution system consisting of N stores and a centralized warehouse has been considered in this study. The problem of multi-echelon inventory management has been a challenging area of research for many years, and as an increasing number of companies in today's economy expand, particularly in terms of vertical integration, efficient control of inventories has become an important strategic issue.

The system studied here was suggested by a wholesale distributor of replacement parts for heavy industrial and agricultural equipment operating several geographically dispersed outlets in the same state. Each distributor (store) places individual periodic replenishment orders to the manufacturer in another state, independent of the other outlets. Having experienced problems with overstocking, obsolete stock, and shortages, they are considering the operation of a central warehouse to transact business with the manufacturer on a corporate scale, and from which to restock and redistribute inventories amongst the stores.

The goal of this research was to study the effectiveness of existing inventory ordering policy models under the range of demand conditions expected to be experienced by the distributor. Our particular interest was in the characteristics of the aggregate inventory behavior in terms of estimated capital investment and storage space requirements, the demand on the transportation system in terms of expected maximum vehicle volumes and weights, and measures of customer service at both echelons of the system.

The two-level inventory system that we consider here operates under stochastic demand with periodic review and fixed lead time (s,S) ordering policies at the stores as well as at the warehouse. (s,S) ordering policy models, which were proven to provide optimal results for a single item at a single location by Scarf (1960), have been published often in the literature under a variety of demand conditions. However, the first application of (s,S) policies to a multi-echelon system was provided by Ehrhardt, et.al. (1981). The models used in this study were the product of Schneider, et al. (1995), which are based on the power approximation models of Ehrhardt (1979) and Ehrhardt and Mosier (1984) and asymptotic renewal theory. The Schneider, et al. (1995) approximations are important in that they consider shortages at the warehouse as a stochastic component in the lead time of deliveries to the stores. They were shown to perform reasonably well under a variety of demand assumptions, and are simple to handle computationally. The demand parameters of the items considered in this study are, in some instances, outside those used by Schneider, et al. (1995) particularly in that their mean per-period demands are very low. Additionally, simple computational formulas for estimating the expected on-hand inventory levels along with formulas for confidence limits for the expected maximum aggregate inventory and the expected maximum per-period shipping volumes were proposed.

In order to study the performance of the analytic models under a variety of demand conditions, a simulation model was designed to test 32 combinations of 10 input parameters set at two levels. The input parameters for the stores consisted of the length of the review period, delivery lead time from the warehouse, the required γ -service

level, and the number of stores in the system. Parameters for the warehouse included the review period, delivery lead time from the manufacturer, and the required α -service level. A common cost to place an order and the annual holding cost rate were applied at both echelons and 'slow' moving parts were tested versus 'fast' movers.

Ordering policies were generated for 10,560 stock keeping units using the analytic models. During each of the simulation experiments, statistics were generated to compare against the estimations from the analytic models for holding costs, ordering costs, total system costs, aggregate on-hand inventories, and shortages as reflected by the γ -service levels. Additionally, the simulated aggregate on-hand inventories and the volume of total orders to the warehouse per-period were recorded in order to observe their distributions.

The holding costs from both the analytic approximations and the simulations were compared for their relative deviations. The holding costs for the approximations were found to be, on the average, 5.4% lower than the simulated costs. The relative deviations for holding cost ranged from -0.944 to 1.261 and 90% were within $\pm 20\%$.

The ordering costs were also analyzed for their percent deviations between the analytic and simulated values. The approximated ordering costs were found to be an average 4% above the simulated costs with a range of -0.667 to 2.00 which resulted in more frequent orders and lower average inventory than the simulation on the average. The ordering cost deviations were more variable than those of the holding cost with only 80% of the deviations within $\pm 20\%$.

The total costs (holding plus ordering costs) were then examined for their relative deviations. The average deviation for total costs was -0.035 with a range of -0.932 to 1.186 and 93% of the deviations were within $\pm 20\%$. The positive deviations for order costs and the negative deviations for holding costs did not tend balance out because holding costs were observed to be much greater in magnitude than ordering costs across nearly all ordering policies. A fractional factorial analysis was performed using total system cost as the response variable to estimate the effects of the input parameters on total costs. In order of magnitude, the demand rate, number of stores in the system, holding cost rate, and order cost were found to be statistically significant in increasing the total system cost above the grand mean. A similar factorial analysis was performed on the total system cost deviations to determine if changes in parameters affected the difference in the analytic and simulated results. Only the holding cost rate was found to be marginally significant such that an increase in the holding rate causes larger negative deviations.

A comparison was made between the accuracy of the asymptotic estimate of average inventory and the simpler computation. Again, the relative deviation from the simulated results was used for the comparison. The less complex estimate was found to be closer, on the average, to the simulated values with an average -3.9% deviation versus the average -9.8% deviation of the asymptotic estimates. Although the simpler estimate appears to be more accurate, both measures give estimates for average on-hand that are below the simulated values.

As Schneider, et al. (1995) point out, the application of their asymptotic (s, S) models is only appropriate when the difference $S-s=D$ approximately exceeds 1.5μ . Since the demand parameters used in this study generated some policies with small D/μ ratios, they were computed and averaged over each experiment to determine if there was any significant effect on the accuracy of the approximations resulting from them. We found that the magnitudes of the D/μ ratios for store policies were most affected by the length of the review period and the demand rate. Increases in either of these parameters tended to decrease the ratios. For the warehouse, which has its ordering policies determined by the aggregated demands at the stores, the D/μ ratios decrease most with increases in the demand rate. It was also found that there was little or no correlation between the magnitude of the D/μ ratios and the relative accuracy of the approximated costs.

Since required levels of customer service were input parameters to the analytic approximations, the simulated service levels at both the warehouse and the stores resulting from the ordering policies were analyzed to determine if they had been accurately predicted. For the stores, it was found that the simulation percent deviations averaged 5.4% above the input gamma service levels. Fractional factorial analysis of the service level deviation showed that the input level had the greatest effect on the deviations. Experiments set at the higher level of 0.98 produced the smallest deviations from the simulated results. Experiments with a shorter review period at the stores also tended to provide more accurate gammas.

For the warehouse, while the predicted service levels were, in general, quite high (close to 100%), the simulated service levels were quite 'poor' for many of the experiments. In over half of the experiments, warehouse γ -service was negative, meaning that the average accumulated backlog at the warehouse exceeded the average per-period demand. Experimental design analysis of the magnitude of the service level deviations showed that the largest deviations were overwhelmingly caused by low demand rate parts. Longer store review periods and a lesser number of stores in the system also resulted in larger deviations in the service levels.

The aggregated per-period orders from the stores to the warehouse were also considered in order to determine volume and weight requirements for the delivery vehicles. We found that the analytic models produced highly accurate estimates of both the means and standard deviations for the per-period orders compared to the simulation, within 0.2% and 0.6% respectively. Observation of the simulated per-period order sizes shows strong support that the order volumes can be approximated using a normal distribution whenever the number of stores in the system is large or the demand rates are high. Low demand rate experiments generally showed order size distributions that were not well defined, mainly due to the high number of review periods when there were no orders. The assumption of normality allows us to apply the standard normal deviate to the means and standard deviations produced by the analytic models for calculating upper confidence limits for the volume and weight of shipments to the stores.

The characteristics of the aggregate on-hand inventories resulting from the application of the (s,S) ordering policies were also investigated. The expected inventory levels can be accurately predicted as discussed earlier. However, calculating the variance of on-hand when (s,S) policies are used has proven difficult and no robust method has as yet been proposed. However, the on-hand inventories were monitored for each of the 352 simulated stocking locations. Histograms and normal probability plots were prepared for each location. Here, again, there is strong evidence that the aggregate on-hand inventory levels exhibit normally distributed behavior. A simple computational formula for estimating an upper confidence limit for the maximum on-hand inventory level was proposed. The variance estimate proved to be inaccurate for measuring the upper tail of the distribution of the aggregate inventory unless the orders for all parts are near-perfectly correlated.

Overall, the ordering policy models used here perform reasonably well when applied to the relatively low demand items of this study. However, they did not perform as consistently as they were shown under the demand assumptions of their authors, especially in estimating total costs. The simulated customer service at the store level proved satisfactory when compared against the preselected input levels, although the service levels of the warehouse are suspect. However, with very slow moving parts, order sizes may not be large enough to bring the on-hand inventory back to a positive position, leaving backlogs 'on the books' for longer periods of time. In these cases, an 'emergency' order might need to be initiated to return the stock to a positive

level. These emergency orders were not used in the simulations and may account for some of the differences between the analytical and the simulated results.

The models proved to be highly accurate in estimating the means and variances of aggregate per-period shipments and should provide good estimates for capacity requirements of the transportation system. In studying the D/μ ratios for the ordering policies, it needs mentioning that in cases of very slow moving parts, a management override should be incorporated. Very slow movers often had values for D that were large enough for several years' supply of the part. When obsolescence is a factor, especially in the case of high cost items, it would be advisable to stock zero of these parts at the store level and keep only a minimal supply at the warehouse.

In continuing the study of two-level inventory and distribution systems, clearly there is a need to develop an accurate technique for determining the variance of the aggregate on-hand inventories generated by the application of (s,S) control policies. The performance of the models used in this study must also be examined for various demand distributions such as the gamma and the compound Poisson. There are many open research problems concerning more complex material flow patterns than are considered with full backordering. Another area of interest for future research will be to study the effects of emergency or 'special' orders on system costs and service levels when shortages occur. With an emergency order policy, shortages at the stores are filled with expedited shipments from the warehouse, further reducing the levels of safety stocks required to be held at the store level. Another interesting problem is the

sensitivity of costs and service levels to situations where the constraints for aggregate capital investment, storage space and/or shipping volumes are binding.

Although there is opportunity for further research in the area, it is clear that the simple approximations studied here provide reasonable results that could be applied in practice and would probably give significant savings when compared to the inventory control methods that are used by many companies today.

BIBLIOGRAPHY

- Albright, S. C. (1989). "An Approximation to the Stationary Distribution of a Multiechelon Repairable Item with Finite Sources and Repair Channels," *Naval Research Logistics*, Vol. 36, pp. 179-195.
- Allen, S. G. (1958). "Redistribution of Total Stock over Several User Locations," *Naval Research Logistics Quarterly*, Vol. 5, pp. 337-345.
- Alscher, J., and Schneider, H. (1982). "Resolving a Multi-item Inventory Problem with Unknown Costs," *Engineering Costs and Production Economics*, Vol. 6, pp. 9-15.
- Axsäter, S. (1990a). "Optimization of Order-Up-To-S Policies in Two-Echelon Inventory Systems with Periodic Review," Unpublished Manuscript, Lulea University of Technology, Lulea, Sweden, September.
- Axsäter, S. (1990b). "Simple Solution Procedures for a Class of Two-Echelon Inventory Problems," *Operations Research*, Vol. 38, No. 1, pp. 64-69.
- Axsäter, S. (1990c). "Modelling Emergency Lateral Transshipments in Inventory Systems," *Management Science*, Vol. 36, No. 11, pp. 1329-1338.
- Axsäter, S. (1993a). "Exact and Approximate Evaluation of Batch-Ordering Policies for Two-Level Inventory Systems," *Operations Research*, Vol. 41, No. 4, pp. 777-785.
- Axsäter, S. (1993b). "Continuous Review Policies for Multi-Level Inventory Systems with Stochastic Demand", Logistics of Production and Inventory, Handbooks in Operations Research and Management Science, (Graves, S. C. et al., eds.), Vol. 4, Ch. 4, (North-Holland, Amsterdam).
- Badinelli, R. D., and Schwarz, L. B. (1988). "Backorders Optimization in a One-Warehouse N-Identical retailer Distribution System," *Naval Research Logistics*, Vol. 35, pp. 427-440.
- Bessler S. A., and Veinott, A. F. (1965). "Optimal Policy for a Dynamic Multiechelon Inventory Model," *Naval Research Logistics Quarterly*, Vol. 13, pp. 355-389.
- Caplin, A. S. (1985). "The Variability of Aggregate Demand with (S,s) Inventory Policies," *Econometrica*, Vol. 53, No. 6, pp. 1395-1409.
- Chen, F., and Zheng, Y.S. (1994). "Lower Bounds for Multi-echelon Stochastic Inventory Systems," *Management Science*, Vol.40, No. 11, pp. 1426-1443.

- Clark, A. J. (1972). "An Informal Survey of Multi-echelon Inventory Theory," *Naval Research Logistics Quarterly*, Vol. 19, pp. 621-650.
- Clark, A. J., and Scarf, H.E. (1960). "Optimal Policies for a Multi-echelon Inventory Problem," *Management Science*, Vol. 6, pp. 475-490.
- Clark, A. J., and Scarf, H.E. (1962). "Approximate Solutions to a Simple Multi-echelon Inventory Problem," Studies in Applied Probability and Management Science, (Arrow, Karlin, and Scarf, eds.), Chapter 5, (Stanford University Press, Stanford, CA.).
- Das, C. (1975). "Supply and Redistribution Rules for Two-Location Inventory Systems: One Period Analysis," *Management Science*, Vol. 21, No. 7, pp. 765-776.
- Demmy, W. S., and Presutti, V. J. (1981). "Multi-Echelon Inventory Theory in the Air Force Logistics Command," Multi-Level Production/Inventory Control Systems: Theory and Practice, TIMS Studies in the Management Sciences, (Schwarz, L., B., ed.), Vol. 16, (North Holland, Amsterdam).
- Deuermeyer, B. L., and Schwarz, L. B. (1981). "A Model for the Analysis of System Service Level in Warehouse-Retailer Distribution Systems: The Identical Retailer Case," Multi-Level Production/Inventory Control Systems: Theory and Practice, TIMS Studies in the Management Sciences, (Schwarz, L., B., ed.), Vol. 16, (North Holland, Amsterdam).
- Ehrhardt, R. A. (1979). "The Power Approximation for Computing (s,S) Inventory Policies," *Management Science*, Vol. 25, No. 8, pp. 777-786.
- Ehrhardt, R. A. (1984). "Analytic Approximations for (s,S) Inventory Policy Operating Characteristics," *Naval Research Logistics Quarterly*, pp. 255-266.
- Ehrhardt, R. A., and Mosier, C. (1984). "A Revision of the Power Approximation for Computing (s,S) Policies," *Management Science*, Vol. 30, No. 5, pp. 619-622.
- Ehrhardt, R. A., Schultz, C. R., and Wagner, H. M. (1981). "(s,S) Policies for a Wholesale Inventory System," Multi-Level Production/Inventory Control Systems: Theory and Practice, TIMS Studies in the Management Sciences, (Schwarz, L. B., ed.), Vol. 16, (North Holland, Amsterdam).
- Eppen, G. D. (1979). "Effects of Centralization on Expected Costs in a Multi-location Newsboy Problem," *Management Science*, Vol. 25, No. 5, pp. 498-501.

- Eppen, G. D. and Schrage, L. E. (1981). "Centralized Ordering Policies in a Multi-Warehouse System with Lead Times and Random Demand," Multi-Level Production/Inventory Control Systems: Theory and Practice, TIMS Studies in the Management Sciences, (Schwarz, L., B., ed.), Vol. 16, (North Holland, Amsterdam).
- Erkip, N., Hausman, W. and Nahmias, S. (1990). "Optimal Centralized Ordering Policies in Multi-Echelon Inventory Systems with Correlated Demands," *Management Science*, Vol. 36, No. 3, pp. 381-392.
- Federgruen, A. (1993). "Centralized Planning Models for Multi-Echelon Inventory Systems Under Uncertainty," Logistics of Production and Inventory, Handbooks in Operations Research and Management Science, (Graves, S. C. et al., eds.), Vol. 4, Chapter 3, (North-Holland, Amsterdam).
- Federgruen, A. and Zipkin, P. (1984a). "Approximations of Dynamic Multilocation Production and Inventory Problems," *Management Science*, Vol. 30, No.1, pp. 69-84.
- Federgruen, A. and Zipkin, P. (1984b). "Computational Issues in an Infinite Horizon Multiechelon Inventory Model," *Operations Research*, Vol. 32, No. 4, pp. 818-836.
- Federgruen, A. and Zipkin, P. (1984c). "Allocation Policies and Cost Approximations for Multilocation Inventory Systems," *Naval Research Logistics Quarterly*, Vol. 31, pp. 97-129.
- Gardner, E. S., and Dannenbring, D. G. (1979). "Using Optimal Policy Surfaces to Analyze Aggregate inventory Tradeoffs," *Management Science*, Vol. 25, No. 8, pp. 709-720.
- Graves, S. C. (1985). "A Multi-Echelon Inventory Model for a Repairable Item with One-For-One Replenishment," *Management Science*, Vol. 31, No. 10, pp. 1247-1256.
- Graves, S. C. (1989). "A Multi-Echelon Inventory Model With Fixed Reorder Intervals," Working Paper #3045-89-MS, Sloan School of Management, Massachusetts Institute of Technology.
- Graves, S. C. (1996). "A Multiechelon Inventory Model with Fixed Replenishment Intervals," *Management Science*, Vol. 42, No. 1, pp. 1-18.

- Gross, D. (1963). "Centralized Inventory Control in Multilocation Supply Systems," Multistage Inventory Models and Techniques, (Scarf, Gilford, and Shelly, eds.), (Stanford University Press, Stanford, CA.).
- Hadley, G., and Whiten, T. M. (1963), "An Inventory-Transportation Model With N Locations," Multistage Inventory Models and Techniques, (Scarf, Gilford, and Shelly, eds.), Chapter 5, (Stanford University Press, Stanford, CA.).
- Hausman, W. H., and Erkip, N. K. (1994). "Multi-echelon vs. Single-echelon Inventory Control Policies for Low-demand Items." *Management Science*, Vol. 40, No. 5, pp. 597-602.
- Hoadley, B., and Heyman, D. P. (1977). "A Two Echelon Inventory Model with Purchases, Dispositions, Shipments, Returns, and Transshipments," *Naval Research Logistics Quarterly*, Vol 24, pp. 1-19.
- Igelhart, D. (1963). "Dynamic Programming and Stationary Analysis of Inventory Problems," Multistage Inventory Models and Techniques, (Scarf, Gilford, and Shelly, eds.), Chapter 1, (Stanford University Press, Stanford, CA.).
- Jackson, P. L. (1988). "Stock Allocation in a Two Echelon Distribution System or 'What to Do Until Your Ship Comes In'," *Management Science*, Vol. 34, No. 7, pp. 880-895.
- Jackson, P. L., and Muckstadt, J. A. (1989). "Risk Pooling in a Two Period, Two Echelon Inventory Stocking and Allocation Problem," *Naval Research Logistics Quarterly*, Vol. 36, pp. 1-26.
- Johnson, G. D., and Thompson, H. E. (1975), "Optimality of Myopic Inventory Policies for Certain Dependent Demand Processes," *Management Science*, Vol. 21, No. 11, pp. 1303-1307.
- Jönsson, G. D., and Silver, E. A. (1987). "Analysis of a Two-Echelon Inventory Control System with Complete Redistribution," *Management Science*, Vol. 33, No. 2, pp. 215-227.
- Karlin, S. (1958a). "Optimal Inventory Policy for the Arrow-Harris-Marchak Dynamic Model," Studies in the Mathematical Theory of Inventory and Production, (Arrow, Karlin, and Scarf eds.), Chapter 9, (Stanford University Press, Stanford, CA.).

- Karlin, S. (1958b). "The Application of Renewal Theory to the Study of Inventory Policies," Studies in the Mathematical Theory of Inventory and Production, (Arrow, Karlin, and Scarf eds.), Chapter 15, (Stanford University Press, Stanford, CA.).
- Krishnan, K. S., and Rao, V. R. K. (1965). "Inventory Control in N Warehouses," *The Journal of Industrial Engineering*, Vol. 16, pp. 212-215.
- Kumar, A. , Schwarz, L. B., and Ward, J. E. (1995). " Risk Pooling Along a Fixed Delivery Route Using a Dynamic Inventory Allocation Policy," *Management Science*, Vol. 41, No. 2, pp. 344-362.
- McGavin, E. J., Schwarz, L. B., and Ward, J. E. (1993). "Two Interval Inventory Allocation Policies in a One-Warehouse N-Identical-Retailer Distribution System." *Management Science*, Vol 39, No. 9, pp. 1092-1107.
- Muckstadt, J. A., and Thomas, L. J. (1980). "Are Multi-Echelon Inventory Methods Worth Implementing in Systems with Low-Demand Rates?" *Management Science*, Vol. 31, No. 5, pp. 483-494.
- Nahmias, S. (1981). "Managing Repairable Item Inventory Systems: A Review," Multi-Level Production/Inventory Control Systems: Theory and Practice, *TIMS Studies in the Management Sciences*, (Schwarz, L. B., ed.), Vol. 16, (North Holland, Amsterdam).
- Nahmias, S. (1989). Production and Operations Analysis, Richard D. Irwin, (Homewood, IL.).
- Nahmias, S., and Smith, S. A. (1994). "Mathematical Models of Retailer Inventory Systems: A Review," Perspectives in Operations Management (R. Sarin, Ed.) Kluwer, Norwell MA, 1993, pp. 249-278.
- Nahmias, S., and Smith, S. A. (1994). "Optimizing Inventory Levels in a Two Echelon System with Partial Lost Sales," *Management Science*, Vol. 40, No. 5, pp. 582-596.
- Press, W., Flannery, B., Teukolsky, S., and Vetterling, W. (1988). Numerical Recipes, (Cambridge University Press, N.Y.).
- Roberts, D. (1962). "Approximations to Optimal Policies in a Dynamic Inventory Model," Studies in Applied Probability and Management Science, (Arrow, Karlin, and Scarf, eds.), (Stanford University Press, Stanford, CA.).

- Rosenfield, D. B., and Pendrock, M. (1980). "The Effects of Warehouse Configuration Design on Inventory Levels and Holding Costs," *Sloan Management Review*, Vol. 21, No. 4, pp. 21-33.
- Roundy, R. (1985). "98% Effective Integer-Ratio Lot Sizing for One Warehouse Multi-Retailer Systems," *Management Science*, Vol. 31, No. 11, pp. 1416-1430.
- Scarf, H. E. (1960). "The Optimality of (s,S) Policies in the Dynamic Inventory Problem," *Mathematical Methods in the Social Sciences*, (Arrow, Karlin, and Suppes, eds.), Chapter 13, (Stanford University Press, Stanford, CA.).
- Schneider, H. (1978). "Methods for Determining the Reorder Point of an (s,S) Ordering Policy when a Service Level is Specified," *Operational Research Quarterly*, Vol. 29, pp. 1181-1193.
- Schneider, H. (1981). "Effect of Service-Levels on Order-Point or Order-Level in Inventory Models," *International Journal of Production Research*, Vol. 19, pp. 615-631.
- Schneider, H., and Ringuest, J. L. (1990). "Power Approximation for Computing (s,S) Policies using Service Level," *Management Science*, Vol. 36, No. 7, pp. 822-834.
- Schneider, H. and Rinks, D. B. (1989). "Optimal Policy Surfaces For a Multi-item Inventory Problem," *European Journal of Operational Research*, Vol. 39, pp. 180-191.
- Schneider, H. and Rinks, D. B. (1991). "Empirical Study of a New Procedure for Allocating Safety Stock in a Wholesale Inventory System," *International Journal of Production Economics*, Vol 24, pp. 181-189.
- Schneider, H., Rinks, D. B., and Kelle, P. (1995). "Power Approximations for a Two-Echelon Inventory System Using Service Levels," *Production and Operations Management*, Vol. 4, No. 3, pp. 1-20.
- Schrady, D. A., and Choe, V. E. (1977). "Models for Multi-Item Continuous Review Inventory Policies Subject to Constraints." *Naval Research Logistics Quarterly*, Vol. 24, pp. 451-463.
- Schultz, C. (1983). "Computing Demand Properties at the Wholesale Warehouse Level," *Naval Research Logistics Quarterly*, Vol. 30, pp. 37-48.

- Schwarz, L. B. (1989). "A Model for Assessing the Value of Warehouse Risk-Pooling: Risk-Pooling Over Outside Supplier Leadtimes," *Management Science*, Vol. 35, No. 7, pp. 828-842.
- Schwarz, L. B., Deuermeyer, B. L., and Badinelli, R. D. (1985). "Fill-Rate Optimization in a One-Warehouse N-Identical Retailer Distribution System," *Management Science*, Vol. 31, No. 4, pp. 488-498.
- Sherbrooke, C. C. (1968). "METRIC: A Multi-echelon Technique of Recoverable Item Control," *Operations Research*, Vol. 16, No. 1, pp. 122-141.
- Simpson, K. E., Jr. (1959). "A Theory of Allocation of Stocks to Warehouses," *Operations Research*, Vol. 7, pp. 797-805.
- Starr, M. K., and Miller, D. W. (1962). Inventory Control: Theory and Practice, Prentice-Hall, (Englewood Cliffs, N.J.).
- Svoronos, A., and Zipkin, P. (1988). "Estimating the Performance of Multi-Level Inventory Systems," *Operations Research*, Vol. 36, No. 1, pp. 57-72.
- Svoronos, A., and Zipkin, P. (1991). "Evaluation of One-For-One Replenishment Policies for Multiechelon Inventory Systems," *Management Science*, Vol. 37, No. 1, pp. 68-83.
- Tersine, R. J. (1994). Principles of Inventory and Materials Management, PTR Prentice-Hall, Inc, (Englewood Cliffs, N.J.).
- Van Houtum, G. J., and Zijm, W. H. M. (1991). "Computational Procedures for Stochastic Multi-echelon Production Systems," *International Journal of Production Economics*, Vol. 23, pp. 223-237.
- Zipkin, P. (1991). "Evaluation of Base Stock Policies in Multiechelon Inventory Systems with Compound-Poisson Demands," *Naval Research Logistics*, Vol. 38, pp. 397-412.

APPENDIX A: MATHEMATICAL NOTATION

A1. Notation for Store Models

μ_i = mean demand at store i during the review period.

μ_{i,L_i-1} = expected demand during lead time plus review time at store i .

σ_i = standard deviation of demand at store i .

σ_{i,L_i-1} = standard deviation of demand during lead time plus review time at store i .

x_i = inventory position at any arbitrary point in time at store i .

N = number of stores.

K_i = ordering cost at store i .

h_i = holding cost at store i .

p_i = stock out cost at store i .

L_i = fixed transportation lead time from manufacturing to store i .

s_i = order decision point for store i .

S_i = order-up-to-level for store i .

D_i = $S_i - s_i$.

Q_i = expected order quantity for store i .

$H_i(x_i)$ = conditional expected inventory at the end of period $t + L_i$ at store i .

$f_i(r)$ = pdf of demand within the review period at store i .

$f_i(r, L_i + 1)$ = pdf of demand during the review period plus the lead time at store i .

$B_i(x_i)$ = expected backlog at the end of period $t + L_i$ at store i .

$\Psi_i(x_i)$ = steady state distribution of the inventory position for an (s, S) policy.

$M_i(D_i)$ = the inventory renewal quantity at store i .

$G_i(x_i)$ = conditional expected holding and shortage cost at store i .

$C_i(s_i, S_i)$ = long run average cost given an (s, S) policy at store i .

$1-\gamma_i$ = per units demanded service level (fill rate) at store i .

τ_i^2 = variance of store i demand given the probability of a warehouse shortage.

A2. Notation for Stores with Warehouse Models

L_i = stochastic leadtime between warehouse and store i .

μ_w = mean demand at the warehouse.

μ_{wL_w+1} = mean demand in leadtime plus review time for N stores at the warehouse.

σ_{wL_w+1} = standard deviation of demand in lead time plus review time at the warehouse.

K_w = ordering cost at the warehouse.

h_w = holding cost at the warehouse.

L_w = fixed replenishment leadtime for the warehouse.

s_w = order decision point for the warehouse.

S_w = order-up-to-level for the warehouse.

α = probability of a shortage during an arbitrary period at the warehouse.

$F_w(x; L_w + 1)$ = conditional probability of demand not exceeding supply at the warehouse during lead time plus review time.

$C_w(s_w, S_w)$ = long run average cost at the warehouse given an (s, S) policy.

$1-\gamma_w$ = per units demanded service level at the warehouse.

A3. Notation for Store Level Restrictions

μ_{iL_i} = mean demand during the lead time at store i .

σ_{i,L_i} = standard deviation of demand during the lead time at store i .

$E[I_i]$ = expected inventory level at store i .

P_j = purchase price per unit for part j .

$E[CI_{ij}]$ = expected capital investment in part j at store i .

$E[CI_i]$ = expected capital investment in total inventories at store i .

$E[I_{ij,1-\delta}] = (1-\delta) \times 100\%$ upper confidence limit for expected inventory.

V_j = square units volume factor for part j .

$E[V_{i,1-\delta}] = (1-\delta) \times 100\%$ upper confidence limit for storage volume at store i .

A4. Notation for Warehouse Restrictions

μ_{w,L_w} = mean demand during the lead time at the warehouse.

σ_{w,L_w} = standard deviation of demand during the lead time at the warehouse.

$E[I_w]$ = expected inventory level at the warehouse.

$E[CI_{wj}]$ = expected capital investment in part j at the warehouse.

$E[CI_w]$ = expected capital investment in total inventories at the warehouse.

$E[I_{wj,1-\delta}] = (1-\delta) \times 100\%$ upper confidence limit for warehouse inventory.

$E[V_{w,1-\delta}] = (1-\delta) \times 100\%$ upper confidence limit for storage volume at the warehouse.

A5. Notation for Warehouse Deliveries to Stores Restrictions

μ_{wj} = mean demand at the warehouse during an order cycle for part j .

σ_{wj}^2 = variance of demand at the warehouse during an order cycle for part j .

$E[I_{wj,1-\delta}] = (1-\delta)$ upper confidence limit of part j shipped per order cycle.

$E[V_{sw,1-\delta}] = (1-\delta)$ upper confidence limit of volume shipped per cycle for part j .

TV_i = delivery vehicle volume required for store i order.

TV_r = delivery vehicle volume required for N store route.

APPENDIX B: GLOSSARY OF TERMINOLOGY

- backorder:** a demand which occurs for an item which is out of stock, is backordered and filled as soon as the next adequately sized replenishment arrives.
- deterministic demand:** a pattern of demand for which the parameters are known or can be calculated with certainty.
- echelon stock:** for echelon j , the number of units that are at, or have passed through, echelon j but have not yet been specifically committed to customers.
- fill rate:** the fraction of demand which is satisfied immediately from inventory.
- inventory:** any stock or store of goods which is held for future use including supplies, raw materials, work-in-process goods, and finished goods.
- inventory turns:** also known as turnover; a measure of the velocity with which materials move through the organization measured by the ratio of the annual cost of goods sold to the average or current inventory investment.
- inventory position:** also called available stock; is defined as the relation: $\text{Inventory Position} = (\text{On hand}) + (\text{On order}) - (\text{Backorders})$.
- multi-echelon system:** an inventory/distribution system having more than one operational level; for example: a manufacturing plant, distribution center, warehouses, and retail outlets if controlled by the same concern would constitute a four-echelon system.
- lead time:** the length of time that passes between the decision to replenish an item and its actual physical addition to stock.
- lost sales:** any demand which occurs for an item that is out of stock is lost; the customer goes elsewhere to satisfy his or her need.
- replenishment cycle:** the length of time that passes between replenishment decisions.
- review period:** see replenishment cycle

repairable item: an item which fails after it is put into service and is recycled, usually through a repair center, and made available for use again.

risk pooling: retaining safety stock at the warehouse in order to balance store stocks between warehouse replenishments and reduce the probability of system stockouts.

safety stock: also called buffer stock; inventory held in reserve when the pattern of demand is uncertain as protection against stockouts.

shortage: a demand occurs for an item that has no items currently in stock.

stock keeping unit: (SKU); a specific item of stock which has been completely defined by function, style, size, color, etc., and for which replenishment decisions must be made.

stockout: see shortage.

stochastic demand: also known as probabilistic demand; a pattern of demand that is not deterministic; some or all of the demand parameters are probabilistic.

substitution: a customer purchases other brands, colors, or sizes of products when the first choice is out of stock.

transshipment: stock items are redistributed (shipped) between stocking locations belonging to the same echelon.

two-echelon structure: an inventory/distribution system having two levels of distribution; specifically, a warehouse (echelon 2) which distributes goods to at least one store (echelon 1).

APPENDIX C: EXPERIMENTAL DESIGN MATRIX FOR STORES-WITH-WAREHOUSE FRACTIONAL FACTORIAL DESIGN

Table C1. Stores-with-Warehouse Design Matrix

Factor	A Ri	B Li	C Gamma	D K	E hr	F Rw	G Lw	H Alpha	I N	J Drate
Exp										
1	1	1	0.8	3	0.1	5	5	0.95	15	2
2	5	1	0.8	3	0.1	3	3	0.9	5	2
3	1	5	0.8	3	0.1	3	3	0.9	15	1
4	5	5	0.8	3	0.1	5	5	0.95	5	1
5	1	1	0.98	3	0.1	3	3	0.95	5	1
6	5	1	0.98	3	0.1	5	5	0.9	15	1
7	1	5	0.98	3	0.1	5	5	0.9	5	2
8	5	5	0.98	3	0.1	3	3	0.95	15	2
9	1	1	0.8	6	0.1	3	5	0.9	5	1
10	5	1	0.8	6	0.1	5	3	0.95	15	1
11	1	5	0.8	6	0.1	5	3	0.95	5	2
12	5	5	0.8	6	0.1	3	5	0.9	15	2
13	1	1	0.98	6	0.1	5	3	0.9	15	2
14	5	1	0.98	6	0.1	3	5	0.95	5	2
15	1	5	0.98	6	0.1	3	5	0.95	15	1
16	5	5	0.98	6	0.1	5	3	0.9	5	1
17	1	1	0.8	3	0.45	5	3	0.9	5	1
18	5	1	0.8	3	0.45	3	5	0.95	15	1
19	1	5	0.8	3	0.45	3	5	0.95	5	2
20	5	5	0.8	3	0.45	5	3	0.9	15	2
21	1	1	0.98	3	0.45	3	5	0.9	15	2
22	5	1	0.98	3	0.45	5	3	0.95	5	2
23	1	5	0.98	3	0.45	5	3	0.95	15	1
24	5	5	0.98	3	0.45	3	5	0.9	5	1
25	1	1	0.8	6	0.45	3	3	0.95	15	2
26	5	1	0.8	6	0.45	5	5	0.9	5	2
27	1	5	0.8	6	0.45	5	5	0.9	15	1
28	5	5	0.8	6	0.45	3	3	0.95	5	1
29	1	1	0.98	6	0.45	5	5	0.95	5	1
30	5	1	0.98	6	0.45	3	3	0.9	15	1
31	1	5	0.98	6	0.45	3	3	0.9	5	2
32	5	5	0.98	6	0.45	5	5	0.95	15	2

**APPENDIX D: HISTOGRAMS AND DESCRIPTIVE STATISTICS
FOR WAREHOUSE ORDERS**




Figure D1. Histogram of Total Orders to the Warehouse for Experiment 1

Descriptive Statistics, Total Orders to the Warehouse for Experiment 1

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 1	1000	315.01	301.50	311.03	142.94	4.52

Variable	Min	Max	Q1	Q3
EXP 1	0.00	952.00	214.00	409.00




Figure D2. Histogram of Total Orders to the Warehouse for Experiment 2

Descriptive Statistics, Total Orders to the Warehouse for Experiment 2

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 2	200	739.1	733.0	735.3	193.0	13.6

Variable	Min	Max	Q1	Q3
EXP 2	280.0	1329.0	614.5	873.3




Figure D3. Histogram of Total Orders to the Warehouse for Experiment 3

Descriptive Statistics, Total Orders to the Warehouse for Experiment 3

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 3	1000	16.446	19.000	14.534	19.657	0.622

Variable	Min	Max	Q1	Q3
EXP 3	0.000	150.000	0.000	26.000




Figure D4. Histogram of Total Orders to the Warehouse for Experiment 4

Descriptive Statistics, Total Orders to the Warehouse for Experiment 4

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 4	200	25.62	19.00	24.14	21.63	1.53

Variable	Min	Max	Q1	Q3
EXP 4	0.00	106.00	13.00	36.00




Figure D5. Histogram of Total Orders to the Warehouse for Experiment 5

Descriptive Statistics, Total Orders to the Warehouse for Experiment 5

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 5	1000	5.137	0.000	3.848	10.629	0.336

Variable	Min	Max	Q1	Q3
EXP 5	0.000	50.000	0.000	0.000




Figure D6. Histogram of Total Orders to the Warehouse for Experiment 6

Descriptive Statistics, Total Orders to the Warehouse for Experiment 6

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 6	200	86.43	82.50	84.53	43.39	3.07

Variable	Min	Max	Q1	Q3
EXP 6	0.00	228.00	54.00	115.50




Figure D7. Histogram of Total Orders to the Warehouse for Experiment 7

Descriptive Statistics, Total Orders to the Warehouse for Experiment 7

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 7	1000	142.26	136.00	136.49	102.13	3.23

Variable	Min	Max	Q1	Q3
EXP 7	0.00	555.00	75.25	205.00




Figure D8. Histogram of Total Orders to the Warehouse for Experiment 8

Descriptive Statistics, Total Orders to the Warehouse for Experiment 8

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 8	200	1893.6	1866.0	1891.6	305.7	21.6

Variable	Min	Max	Q1	Q3
EXP 8	1072.0	2877.0	1670.5	2141.3




Figure D9. Histogram of Total Orders to the Warehouse for Experiment 9

Descriptive Statistics, Total Orders to the Warehouse for Experiment 9

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 9	1000	5.902	0.000	3.976	14.287	0.452

Variable	Min	Max	Q1	Q3
EXP 9	0.000	93.000	0.000	0.000




Figure D10. Histogram of Total Orders to the Warehouse for Experiment 10

Descriptive Statistics, Total Orders to the Warehouse for Experiment 10

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 10	200	86.10	77.50	83.97	47.05	3.33

Variable	Min	Max	Q1	Q3
EXP 10	0.00	236.00	49.25	116.00




Figure D11. Histogram of Total Orders to the Warehouse for Experiment 11

Descriptive Statistics, Total Orders to the Warehouse for Experiment 11

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 11	1000	178.23	126.00	169.31	137.08	4.33

Variable	Min	Max	Q1	Q3
EXP 11	0.00	743.00	96.00	246.00




Figure D12. Histogram of Total Orders to the Warehouse for Experiment 12

Descriptive Statistics, Total Orders to the Warehouse for Experiment 12

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 12	200	2277.0	2233.0	2275.4	370.5	26.2

Variable	Min	Max	Q1	Q3
EXP 12	875.0	3298.0	2048.5	2506.8




Figure D13. Histogram of Total Orders to the Warehouse for Experiment 13

Descriptive Statistics, Total Orders to the Warehouse for Experiment 13

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 13	1000	562.55	540.00	557.74	239.13	7.56

Variable	Min	Max	Q1	Q3
EXP 13	0.00	1658.00	408.00	714.00




Figure D14. Histogram of Total Orders to the Warehouse for Experiment 14

Descriptive Statistics, Total Orders to the Warehouse for Experiment 14

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 14	200	768.3	765.0	763.4	213.4	15.1

Variable	Min	Max	Q1	Q3
EXP 14	280.0	1433.0	627.8	897.0




Figure D15. Histogram of Total Orders to the Warehouse for Experiment 15

Descriptive Statistics, Total Orders to the Warehouse for Experiment 15

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 15	1000	16.623	0.000	13.853	24.267	0.767

Variable	Min	Max	Q1	Q3
EXP 15	0.000	140.000	0.000	35.000




Figure D16. Histogram of Total Orders to the Warehouse for Experiment 16

Descriptive Statistics, Total Orders to the Warehouse for Experiment 16

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 16	200	25.46	23.00	23.98	22.50	1.59

Variable	Min	Max	Q1	Q3
EXP 16	0.00	92.00	0.00	43.00




Figure D17. Histogram of Total Orders to the Warehouse for Experiment 17

Descriptive Statistics, Total Orders to the Warehouse for Experiment 17

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 17	1000	5.604	0.000	4.732	8.061	0.255

Variable	Min	Max	Q1	Q3
EXP 17	0.000	37.000	0.000	12.000




Figure D18. Histogram of Total Orders to the Warehouse for Experiment 18

Descriptive Statistics, Total Orders to the Warehouse for Experiment 18

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 18	200	77.28	76.00	76.54	24.32	1.72

Variable	Min	Max	Q1	Q3
EXP 18	23.00	155.00	59.00	90.75




Figure D19. Histogram of Total Orders to the Warehouse for Experiment 19

Descriptive Statistics, Total Orders to the Warehouse for Experiment 19

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 19	1000	134.39	132.00	132.35	67.92	2.15

Variable	Min	Max	Q1	Q3
EXP 19	0.00	393.00	81.00	179.00




Figure D20. Histogram of Total Orders to the Warehouse for Experiment 20

Descriptive Statistics, Total Orders to the Warehouse for Experiment 20

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 20	200	1848.1	1850.5	1849.0	208.8	14.8

Variable	Min	Max	Q1	Q3
EXP 20	1168.0	2418.0	1710.3	1996.5




Figure D21. Histogram of Total Orders to the Warehouse for Experiment 21

Descriptive Statistics, Total Orders to the Warehouse for Experiment 21

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 21	1000	487.01	485.00	485.48	129.05	4.08

Variable	Min	Max	Q1	Q3
EXP 21	101.00	884.00	399.00	571.00




Figure D22. Histogram of Total Orders to the Warehouse for Experiment 22

Descriptive Statistics, Total Orders to the Warehouse for Experiment 22

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 22	200	713.04	704.00	709.94	126.25	8.93

Variable	Min	Max	Q1	Q3
EXP 22	455.00	1074.00	611.50	803.25




Figure D23. Histogram of Total Orders to the Warehouse for Experiment 23

Descriptive Statistics, Total Orders to the Warehouse for Experiment 23

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 23	1000	16.832	13.000	15.826	14.224	0.450

Variable	Min	Max	Q1	Q3
EXP 23	0.000	67.000	0.000	25.000




Figure D24. Histogram of Total Orders to the Warehouse for Experiment 24

Descriptive Statistics, Total Orders to the Warehouse for Experiment 24

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 24	200	30.32	28.00	29.84	16.88	1.19

Variable	Min	Max	Q1	Q3
EXP 24	0.00	82.00	18.00	38.00




Figure D25. Histogram of Total Orders to the Warehouse for Experiment 25

Descriptive Statistics, Total Orders to the Warehouse for Experiment 25

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 25	1000	488.82	479.00	485.78	145.31	4.60

Variable	Min	Max	Q1	Q3
EXP 25	135.00	971.00	384.00	585.00




Figure D26. Histogram of Total Orders to the Warehouse for Experiment 26

Descriptive Statistics, Total Orders to the Warehouse for Experiment 26

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 26	200	668.9	679.5	668.1	146.7	10.4

Variable	Min	Max	Q1	Q3
EXP 26	312.0	1120.0	549.8	772.5




Figure D27. Histogram of Total Orders to the Warehouse for Experiment 27

Descriptive Statistics, Total Orders to the Warehouse for Experiment 27

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 27	1000	16.480	16.000	15.083	16.293	0.515

Variable	Min	Max	Q1	Q3
EXP 27	0.000	82.000	0.000	28.000




Figure D28. Histogram of Total Orders to the Warehouse for Experiment 28

Descriptive Statistics, Total Orders to the Warehouse for Experiment 28

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 28	200	26.03	23.00	25.33	16.58	1.17

Variable	Min	Max	Q1	Q3
EXP 28	0.00	80.00	12.00	35.00




Figure D29. Histogram of Total Orders to the Warehouse for Experiment 29

Descriptive Statistics, Total Orders to the Warehouse for Experiment 29

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 29	1000	5.950	0.000	4.652	9.921	0.314

Variable	Min	Max	Q1	Q3
EXP 29	0.000	53.000	0.000	15.000




Figure D30. Histogram of Total Orders to the Warehouse for Experiment 30

Descriptive Statistics, Total Orders to the Warehouse for Experiment 30

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 30	200	74.46	71.00	73.41	31.63	2.24

Variable	Min	Max	Q1	Q3
EXP 30	12.00	195.00	53.00	97.50




Figure D31. Histogram of Total Orders to the Warehouse for Experiment 31

Descriptive Statistics, Total Orders to the Warehouse for Experiment 31

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 31	1000	135.26	129.00	132.34	79.64	2.52

Variable	Min	Max	Q1	Q3
EXP 31	0.00	432.00	80.00	187.00




Figure D32. Histogram of Total Orders to the Warehouse for Experiment 32

Descriptive Statistics, Total Orders to the Warehouse for Experiment 32

Variable	N	Mean	Median	TrMean	StDev	SEMean
EXP 32	200	2299.5	2293.5	2295.5	259.5	18.3

Variable	Min	Max	Q1	Q3
EXP 32	1552.0	3118.0	2124.8	2475.8

APPENDIX E: HISTOGRAMS AND NORMAL PROBABILITY PLOTS FOR ON-HAND INVENTORY




Figure E1. Histogram for Total On-Hand, Exp 1, Store 1




Figure E2. Normal Probability Plot for Total On-Hand, Exp 1, Store 1




Figure E3. Histogram for Total On-Hand, Exp. 1, Store 2




Figure E4. Normal Probability Plot for Total On-Hand, Exp. 1, Store 2




Figure E5. Histogram for Total On-Hand, Exp. 1, Store 3




Figure E6. Normal Probability Plot for Total On-Hand, Exp.1, Store 3




Figure E7. Histogram for Total On-Hand, Exp. 1, Store 4




Figure E8. Normal Probability Plot for Total On-Hand, Exp. 1, Store 4




Figure E9. Histogram for Total On-Hand, Exp. 1, Store 5




Figure E10. Normal Probability Plot for Total On-Hand, Exp. 1, Store 5




Figure E11. Histogram for Total On-Hand, Exp. 1, Store 6




Figure E12. Normal Probability Plot for Total On-Hand, Exp. 1, Store 6




Figure E13. Histogram for Total On-Hand, Exp. 1, Store 7




Figure E14. Normal Probability Plot for Total On-Hand, Exp. 1, Store 7




Figure E15. Histogram for Total On-Hand, Exp. 1, Store 8




Figure E16. Normal Probability Plot for Total On-Hand, Exp. 1, Store 8




Figure E17. Histogram for Total On-Hand, Exp. 1, Store 9




Figure E18. Normal Probability Plot for Total On-Hand, Exp 1, Store 9




Figure E19. Histogram for Total On-Hand, Exp. 1, Store 10




Figure E20. Normal Probability Plot for Total On-Hand, Exp. 1, Store 10




Figure E21. Histogram for Total On-Hand, Exp. 1, Store 11




Figure E22. Normal Probability Plot for Total On-Hand, Exp. 1, Store 11




Figure E23. Histogram for Total On-Hand, Exp. 1, Store 12




Figure E24. Normal Probability Plot for Total On-Hand, Exp. 1, Store 12




Figure E25. Histogram for Total On-Hand, Exp. 1, Store 13




Figure E26. Normal Probability Plot for Total On-Hand, Exp 1, Store 13




Figure E27. Histogram for Total On-Hand, Exp. 1, Store 14




Figure E28. Normal Probability Plot for Total On-Hand, Exp. 1, Store 14




Figure E29. Histogram for Total On-Hand, Exp. 1, Store 15




Figure E30. Normal Probability Plot for Total On-Hand, Exp. 1, Store 15




Figure E31. Histogram for Total On-Hand, Exp. 1, Warehouse




Figure E32. Normal Probability Plot for Total On-Hand, Exp. 1, Warehouse




Figure E33. Histogram for Total On-Hand, Exp. 5, Store 1




Figure E34. Normal Probability Plot for Total On-Hand, Exp. 5, Store 1




Figure E35. Histogram for Total On-Hand, Exp. 5, Store 2




Figure E36. Normal Probability Plot for Total On-Hand, Exp. 5, Store 2




Figure E37. Histogram for Total On-Hand, Exp. 5, Store 3




Figure E38. Normal Probability Plot for Total On-Hand, Exp. 5, Store 3




Figure E39. Histogram for Total On-Hand, Exp. 5, Store 4




Figure E40. Normal Probability Plot for Total On-Hand, Exp. 5, Store 4




Figure E41. Histogram for Total On-Hand, Exp. 5, Store 5




Figure E42. Normal Probability Plot for Total On-Hand, Exp. 5, Store 5




Figure E43. Histogram for Total On-Hand, Exp. 5, Warehouse




Figure E44. Normal Probability Plot for Total On-Hand, Exp. 5, Warehouse




Figure E45. . Histogram for Total Daily On-Hand, Exp. 9, Store 1




Figure E46. Normal Probability Plot for Total On-Hand, Exp. 9, Store 1




Figure E47. Histogram for Total Daily On-Hand, Exp. 9, Store 2




Figure E48. Normal Probability Plot for Total On-Hand, Exp. 9, Store 2




Figure E49. Histogram for Total Daily On-Hand, Exp. 9, Store 3




Figure E50. Normal Probability Plot for Total On-Hand, Exp. 9, Store 3




Figure E51. Histogram for Total Daily On-Hand, Exp. 9, Store 4




Figure E52. Normal Probability Plot for Total On-Hand, Exp. 9, Store 4




Figure E53. Histogram for Total Daily On-Hand, Exp. 9, Store 5




Figure E54. Normal Probability Plot for Total On-Hand, Exp. 9, Store 5




Figure E55. Histogram for Total Daily On-Hand, Exp. 9, Warehouse




Figure E56. Normal Probability Plot for Total On-Hand, Exp. 9, Warehouse




Figure E57. Histogram for Total On-Hand, Exp. 14, Store 1




Figure E58. Normal Probability Plot for Total On-Hand, Exp. 14, Store 1




Figure E59. Histogram for Total On-Hand, Exp. 14, Store 2




Figure E60. Normal Probability Plot for Total On-Hand, Exp. 14, Store 2




Figure E61. Histogram for Total On-Hand, Exp. 14, Store 3




Figure E62. Normal Probability Plot for Total On-Hand, Exp. 14, Store 3




Figure E63. Histogram for Total On-Hand, Exp. 14, Store 4




Figure E64. Normal Probability Plot for Total On-Hand, Exp. 14, Store 4




Figure E65. Histogram for Total On-Hand, Exp. 14, Store 5




Figure E66. Normal Probability Plot for Total On-Hand, Exp. 14, Store 5




Figure E67. Histogram for Total On-Hand, Exp. 14, Warehouse




Figure E68. Normal Probability Plot for Total On-Hand, Exp. 14, Warehouse




Figure E69. Histogram for Total On-Hand, Exp. 16, Store 1




Figure E70. Normal Probability Plot for Total On-Hand, Exp. 16, Store 1




Figure E71. Histogram for Total On-Hand, Exp. 16, Store 2




Figure E72. Normal Probability Plot for Total On-Hand, Exp. 16, Store 2




Figure E73. Histogram for Total On-Hand, Exp. 16, Store 3




Figure E74. Normal Probability Plot for Total On-Hand, Exp. 16, Store 3




Figure E75. Histogram for Total On-Hand, Exp. 16, Store 4




Figure E76. Normal Probability Plot for Total On-Hand, Exp. 16, Store 4




Figure E77. Histogram for Total On-Hand, Exp. 16, Store 5




Figure E78. Normal Probability Plot for Total On-Hand, Exp. 16, Store 5




Figure E79. Histogram for Total On-Hand, Exp. 16, Warehouse




Figure E80. Normal Probability Plot for Total On-Hand, Exp. 16, Warehouse




Figure E81. Histogram for Total On-Hand, Exp. 22, Store 1




Figure E82. Normal Probability Plot for Total On-Hand, Exp. 22, Store 1




Figure E83. Histogram for Total On-Hand, Exp. 22, Store 2




Figure E84. Normal Probability Plot for Total On-Hand, Exp. 22, Store 2




Figure E85. Histogram for Total On-Hand, Exp. 22, Store 3




Figure E86. Normal Probability Plot for Total On-Hand, Exp. 22, Store 3




Figure E87. Histogram for Total On-Hand, Exp. 22, Store 4




Figure E88. Normal Probability Plot for Total On-Hand, Exp. 22, Store 4




Figure E89. Histogram for Total On-Hand, Exp. 22, Store 5




Figure E90. Normal Probability Plot for Total On-Hand, Exp. 22, Store 5




Figure E91. Histogram for Total On-Hand, Exp. 22, Warehouse




Figure E92. Normal Probability Plot for Total On-Hand, Exp. 22, Warehouse




Figure E93. Histogram for Total On-Hand, Exp. 23, Store 1




Figure E94. Normal Probability Plot for Total On-Hand, Exp. 23, Store 1




Figure E95. Histogram for Total On-Hand, Exp. 23, Store 2




Figure E96. Normal Probability Plot for Total On-Hand, Exp. 23, Store 2




Figure E97. Histogram for Total On-Hand, Exp. 23, Store 3




Figure E98. Normal Probability Plot for Total On-Hand, Exp. 23, Store 3




Figure E99. Histogram for Total On-Hand, Exp. 23, Store 4




Figure E100. Normal Probability Plot for Total On-Hand, Exp. 23, Store 4




Figure E101. Histogram for Total On-Hand, Exp. 23, Store 5




Figure E102. Normal Probability Plot for Total On-Hand, Exp. 23, Store 5




Figure E103. Histogram for Total On-Hand, Exp. 23, Store 6




Figure E104. Normal Probability Plot for Total On-Hand, Exp. 23, Store 6




Figure E105. Histogram for Total On-Hand, Exp. 23, Store 7




Figure E106. Normal Probability Plot for Total On-Hand, Exp. 23, Store 7




Figure E107. Histogram for Total On-Hand, Exp. 23, Store 8




Figure E108. Normal Probability Plot for Total On-Hand, Exp. 23, Store 8




Figure E109. Histogram for Total On-Hand, Exp. 23, Store 9




Figure E110. Normal Probability Plot for Total On-Hand, Exp. 23, Store 9




Figure E111. Histogram for Total On-Hand, Exp. 23, Store 10




Figure E112. Normal Probability Plot for Total On-Hand, Exp. 23, Store 10




Figure E113. Histogram for Total On-Hand, Exp. 23, Store 11




Figure E114. Normal Probability Plot for Total On-Hand, Exp. 23, Store 11




Figure E115. Histogram for Total On-Hand, Exp. 23, Store 12




Figure E116. Normal Probability Plot for Total On-Hand, Exp. 23, Store 12




Figure E117. Histogram for Total On-Hand, Exp. 23, Store 13




Figure E118. Normal Probability Plot for Total On-Hand, Exp. 23, Store 13




Figure E119. Histogram for Total On-Hand, Exp. 23, Store 14




Figure E120. Normal Probability Plot for Total On-Hand, Exp. 23, Store 14




Figure E121. Histogram for Total On-Hand, Exp. 23, Store 15




Figure E122. Normal Probability Plot for Total On-Hand, Exp. 23, Store 15




Figure E123. Histogram for Total On-Hand, Exp. 23, Warehouse




Figure E124. Normal Probability Plot Total On-Hand, Exp 23, Warehouse

VITA

David Lee Peak studied mechanical engineering at Louisiana Tech University at Ruston from 1973 to 1977. He entered the oil and gas industry in 1978 with The Western Company of North America, where he rose to operations manager of the Bossier City, Louisiana District from 1982 to 1984. He earned a bachelor of science degree in management from Louisiana State University at Shreveport in 1989.

In 1990, Mr. Peak entered the graduate program at Louisiana State University and Agricultural and Mechanical College at Baton Rouge, earning a master of science degree in Quantitative Business Analysis in 1991. Continuing at Louisiana State University, he earned a doctor of philosophy degree in Business Administration in 1996, with concentration in Operations Management. He is currently an Assistant Professor and coordinator for the Senior College program at Louisiana State University, Alexandria, Louisiana.

DOCTORAL EXAMINATION AND DISSERTATION REPORT

Candidate: David Lee Peak

Major Field: Business Administration (Information Systems and
Decision Sciences)

Title of Dissertation: Performance Analysis of Multi-Echelon
Inventory Systems

Approved:

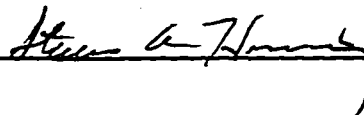


Major Professor and Chairman



Dean of the Graduate School

EXAMINING COMMITTEE:



Date of Examination:

September 27, 1996