

SCHOOL OF OPERATIONS RESEARCH
AND INDUSTRIAL ENGINEERING
COLLEGE OF ENGINEERING
CORNELL UNIVERSITY
ITHACA, NY 14853-3801

TECHNICAL REPORT NO. 1107

August 1995

**Distribution Inventory
Systems: Lower Bounds
and a Heuristic Policy**

by

John Muckstadt, Juan Pereira
and Robin Roundy

**DISTRIBUTION INVENTORY SYSTEMS:
LOWER BOUNDS AND A HEURISTIC
POLICY**

A Thesis

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Juan Esteban Pereira

August 1997

© Juan Esteban Pereira 1997

ALL RIGHTS RESERVED

ABSTRACT

Large retailers' operations have changed substantially in the last couple of decades, evolving from single stores replenishing their stock directly from suppliers, to large distribution systems where goods are centrally ordered and then allocated and delivered to each individual store. Tremendous progress in information systems' technology has enabled leading retailers to exploit economies of scale to the fullest by implementing centralized control over their distribution systems.

This dissertation studies the problem of replenishing inventory in a centrally controlled, multi-facility distribution system with constant traveling times and serving external, stochastic demand. We formulate an inventory model for this problem which is a natural extension of the classical single-facility inventory models for both continuous and periodic review. Our contribution consists of establishing two bounds for the cost incurred by such system, and also finding a heuristic policy for a two-stage serial system. This policy is 93%-effective within the class of weakly nested policies—a class to be defined later in this work— and is used to extrapolate recommendations relevant to practitioners.

The lower bounds are obtained by decomposing the distribution system into single-facility sub-problems, a strategy based on the use of imputed penalty costs, and which substantially reduces computational complexity and permits the appli-

cation of known single-stage techniques. New extensions of the concepts of 'echelon inventory position' and 'nested policy' are also proposed.

Biographical Sketch

The author was born in Medellín, Colombia in 1966. An early interest in mathematics was further nurtured by his involvement in the International Mathematics Olympiads, a high-school level competition in which he obtained two Third Prizes. He attended two universities simultaneously, obtaining a BS degree in civil engineering from the Engineering School of Antioquia in 1989, and another BS degree in mathematics from the National University of Colombia in 1990. During those years, he also organized the Colombian Mathematical Olympiads in Medellín and was part of the Colombian representation to several other international mathematical competitions.

The author then went to Cornell University, where he earned a PhD in operations research in 1995. He now works as a consultant.

Acknowledgements

I would like to thank the members of my committee for everything I have learned from them. I owe tremendous gratitude to Robin Roundy for going well beyond his duties and working long hours to improve this dissertation. I would also like to thank my advisor Jack Muckstadt for introducing to me the practical situation that motivated this thesis. Guillermo Gallego deserves credit for his suggestions in the methodology.

I also want to thank Katerina for being with me all these years, for all her love and understanding, and for being helpful to the point of proofreading this thesis.

My fondest academic experiences at Cornell were related to teaching. I would like to thank all the people involved with these activities, including many of my students, for making my stay at Cornell worthwhile. I would also like to thank the Operations Research and Industrial Engineering department for fostering such a nurturing environment and for their financial support during five years.

I am grateful with my fellow students for their encouragement and support during these years. In particular, I would like to mention Stuart Carr, Roberto Malamut, Laszlo Ladanyi, and Kathy Caggiano. In this connection I must also thank Edoardo Amaldi.

Finally, I would like to thank my family for providing me with such a wonderful

education, for having kept our house filled with books, and for teaching me the importance of integrity and support for each other. I would like to conclude by thanking María de Losada for providing me with direction and motivation, and for her work with the Mathematical Olympiads.

Table of Contents

1	Introduction	1
2	Literature Review	5
2.1	Single-Facility Problems	6
2.1.1	Single-Period Stochastic Models and Order-up-to Policies . .	7
2.1.2	Multi-Period Models and (s, S) Policies	9
2.2	Multi-Echelon Inventory Systems	13
2.2.1	Serial Systems	13
2.2.2	Distribution Systems under Uncertainty	14
2.2.3	Policies with Guaranteed Good Worst-Case Performance . .	16
2.3	Current Inventory Theory's Limitations	18
3	A New Lower Bound on the Cost of Distribution Systems	20
3.1	Classical Single-Facility Inventory Models	22
3.1.1	Classical Periodic-Review Stochastic Inventory Model	23
3.1.2	Continuous Review Models and (r, Q) Policies	27
3.2	Stochastic Modeling of Distribution Systems	30
3.2.1	Model Formulation	30
3.2.2	Zero Lead-Time Reformulation	39
3.3	The Single-Stage Subproblem	40
3.3.1	Optimal Parameters (s^*, S^*)	41
3.3.2	A Facility as Part of a Larger System	43
3.4	A Lower Bound for the Cost of a Facility	46
3.5	Lower Bounds for Serial Systems	50
3.5.1	Weakly Nested Policies	50
3.5.2	Serial Systems	52
4	Optimality of (s, S) Policies for the Single-Stage Subproblem	59
4.1	Optimality of (s, S) Policies for Continuous-Review Systems	60
4.2	Non-Optimality of (s, S) Policies for Periodic-Review Systems . . .	62
4.2.1	Sufficient Conditions for the Optimality of (s, S) Policies for Periodic Review Systems	64
4.2.2	Concavity Assumption and the One-Period Demand	69

4.2.3	Reducing the Review-Period Length	71
5	93%-Effective Policies for a Two-Stage Inventory System	73
5.1	A Sensitivity Property of the (r, Q) Model	73
5.2	Two-Stage Model Formulation	77
5.3	A Lower Bound for Weakly-Nested Policies	79
5.4	A 93%-Efficient Policy	85
6	Conclusions	98
A	Accounting for Holding and Backordering Costs	100
B	Proof of Lemma 3.1.7.	103
C	Proof of Theorem 3.3.5 for Periodic Review Systems	106
	Bibliography	110

List of Tables

4.1	Data for example 4.2.1.	63
-----	---------------------------------	----

List of Figures

3.1	Activities and Inventory Quantities During a Review Period.	36
-----	---	----

Chapter 1

Introduction

Inventory models found in the literature are not designed for the environment faced by today's leading retailers. One reason is that a typical large retailer's operations have changed substantially in the last couple of decades, evolving from single stores replenishing their stock directly from suppliers, to large distribution systems where goods are centrally ordered and then allocated and delivered to each individual store. Simultaneously with these changes, more value-conscious consumers, saturation, and competition, have reduced retailers' gross margins, turning their logistical capabilities into an important competitive advantage.

To achieve purchasing muscle and economies of scale, major retailers have pursued aggressive growth strategies and have set up their own distribution systems, which are often very large. By the year 1993, America's two largest discount retailers, Wal-Mart and K-Mart, had over 2500 and 4200 stores, respectively. Sears, Roebuck, the largest US department store chain, had over 1800 stores in the same year. Besides stores, a distribution system also includes distribution centers or depots, which receive large amounts of goods from external suppliers and deliver

shipments to individual stores.

A distribution system can be most efficiently operated using centralized information, which includes the amounts sold and the stock held for each product at each store. Centralization of control is considered an important strategic goal by industry leaders and is becoming a trend among large US retailers. During the 1980s, for example, Wal-Mart is reckoned by industry analysts to have spent over \$700m on information technologies for its distribution system, including purchasing its own satellite to transmit its enormous load of information [63]. This strategy has paid off. According to company records, Wal-Mart distribution costs were under 3% of sales in 1992, which compared to 4.5%-5% for the firm's competitors, represent savings of close to \$750m in that year alone [63]. Others are following: 51% of a sample of large US retailers surveyed by Ernst & Young, a management consultancy, reported to be using computer-based merchandise allocation systems in 1994 [27].

This dissertation studies the problem of replenishing inventory in a centrally controlled, multi-facility distribution system, with the purpose of serving external, stochastic demand in a cost-efficient way. We formulate an inventory model for this problem which is a natural extension of classical single-facility inventory models for both continuous and periodic review. Our contribution consists of establishing two bounds for the cost incurred by such systems, and also finding a heuristic policy for a two-stage serial system, which is 93%-effective within the class of weakly nested policies—a class to be defined later in this work. Recommendations relevant to practitioners are extrapolated from this heuristic policy.

We would like to point out two interesting features of this work. First, the

lower bounds are obtained using a decomposition approach, that is, breaking a multi-facility system into several single-facility sub-problems. This decomposition is based on the use of imputed penalty costs, which penalize a facility for the extra costs caused by decisions made by the facility, but incurred elsewhere in the system. Decomposing the system greatly reduces the computational complexity of finding a solution, while at the same time permitting the application of known single-facility techniques to the more complicated multi-facility inventory problem.

Even if centralized information is available, it is often desirable to decentralize decision-making because local managers, often responsible for their facility's profits, will be more motivated if they are given more authority over ordering and shipping decisions at their facilities. Decentralized decision-making may be a source of conflict, however, because decisions that are good for a part of the system may be detrimental for the system as a whole. Hence comes another advantage of the decomposition approach: the induced penalty costs for a distribution center mimic the costs caused on the rest of the system by decisions made by the facility. If such a reward system is implemented at each facility, good local decisions will result in good performance of the system as a whole.

The second interesting feature of this work is the way in which it extends the concepts of 'nestedness' and 'inventory position' to multi-facility environments. Nested policies have been successfully used before for systems with deterministic demand or zero lead times. We introduce the concept of 'weak nestedness,' which is defined for systems with stochastic demands and constant lead times. Finally, this work proposes a re-definition of the concept of 'inventory position' for multi-facility systems. This re-definition, which we call 'shipment-based inventory position', is

more consistent with the multi-echelon nature of such systems.

The rest of this dissertation is organized as follows. Chapter 2 reviews the literature available in inventory theory today and comments on its limitations. Chapter 3 formulates a multi-facility inventory model, introduces the notions of weak nestedness and shipment-based inventory position, and develops two bounds for serial inventory systems which can be extended to distribution systems. Chapter 4 discusses optimal policies for the single-stage sub-problem, an inventory model with non-fixed ordering costs resulting from the decomposition approach used for obtaining the bounds proposed in Chapter 3. Chapter 5 proposes a policy for a two-stage serial inventory system which is 93%-efficient within the class of weakly nested policies. Finally, Chapter 6 presents our conclusions and suggestions for practitioners.

Chapter 2

Literature Review

The purpose of this chapter is to survey the literature available in inventory theory today. This is an ambitious task since inventory management has been an active area of research for over forty years and literally thousands of papers and articles (including many surveys) have been written on this problem. The literature is also extensive in terms of the number of models, the issues that these models intend to address, and the approaches followed. We will concisely summarize those that are relevant to these research efforts, outline their assumptions, state their results and comment on their importance. Details are available in other broader surveys, for which references are provided.

Inventory models are mathematical representations of real inventory systems, formulated with the purpose of providing ordering policies that minimize some cost function. These models can be classified in a number of ways. According to the number of locations considered, they may be single-facility or multi-echelon models. There may also be one or more than one item stocked. The demand may be deterministic or stochastic, stationary or not. Unmet demand may be lost,

backlogged, or partially backlogged. System reviews, the times when inventory is observed and ordering decisions can be made, may be conducted either periodically or continuously. The time horizon over which costs should be minimized may be finite or infinite, or just one time period. A good introduction to these issues, and the classical models which are the basis for later developments, can be found in early surveys, such as Scarf [76] (1963) and Veinott [81] (1966).

Next we provide an overview of the remainder of this chapter, which is divided into three sections. The first one (section 2.1) focuses on single-facility inventory problems. It covers the EOQ formula, the newsvendor problem and its extensions (Section 2.1.1), and a standard periodic review model for both finite and infinite time horizons (Section 2.1.2). More specifically, this last subsection includes results on the optimality of (s, S) policies, algorithms to find optimal policies, an equivalent continuous review formulation, and (r, Q) policies.

The second section (2.2) is concerned with multi-echelon inventory systems. It covers serial systems (Section 2.2.1), distribution systems under stochastic demand (Section 2.2.2), and policies with good worst-case performance, namely, nested power-of-two policies (Section 2.2.3).

2.1 Single-Facility Problems

Managing the stock of a single commodity stored at a single location is the simplest inventory problem that can be studied. Many mathematical models have been proposed for providing ordering policies for such systems. These models normally consider some or all of the following costs: the cost of placing an order, the cost of carrying inventory, the cost of incurring backorders, or the cost of losing a sale.

The oldest and simplest of these models is the famous Economic Order Quantity (EOQ) formula, whose development is credited to Harris [41] (1915). The simplicity and robustness of this result are truly remarkable, so much so that the EOQ formula remains the basis for both practical applications and theoretical studies. This occurs even though virtually no real system satisfies the model's assumptions, which include a deterministic, continuous, and constant rate demand.

A detailed presentation of the properties and extensions of the EOQ model can be found in most inventory theory textbooks (see, for example, Chapter 2 of Hadley and Whitin [39]). The most common extensions include backlogging of orders, integrality of demand, lost sales, finite production rate, capacity constraints and quantity discounts.

A related problem is determining inventory levels and production quantities in multi-product manufacturing/distribution systems, most frequently in non-stationary environments and under capacity constraints. Optimization techniques such as dynamic programming, stochastic programming, and various solution algorithms have been proposed, among others, by Evans [28], Haneveld [40], Peters, Boskma, and Kupper [58], and Sox [79].

2.1.1 Single-Period Stochastic Models and Order-up-to Policies

Virtually all inventory models that consider stochastic demand are built upon a single period model, often referred to as the 'newsvendor model,' which seems to have been first introduced, in its most common form, by Arrow, Harris and Marschack [2]. A good review of this model and some of its extensions can be found in Porteus [61]. The newsvendor formulation is particularly suitable for the

problem of perishable and ‘style’ goods. For a review of the former, see Nahmias [57].

The newsvendor’s problem is to make an ordering decision at the beginning of a discrete time period, during which a random demand for a product will take place. The objective is to minimize the expected cost incurred at the end of the period, which is a function of the decision made and the demand realization. Typically, the resulting optimal policy is of the form “order-up-to- S ”, which means that if the initial stock equals $x \leq S$, one should order $S - x$ units; otherwise, no order should be placed.

The most common classical extensions of the newsvendor problem include continuous and discontinuous demands, partial backlogging and lost sales, linear, convex and non-convex cost functions, minimum and maximum ordering quantities, batch ordering, non-linear production costs, delivery quantities differing from the amount ordered, many products with a resource constraint, competition, and pricing (see Porteus [61]). A recent extension considers the production decision for two markets with non-overlapping selling seasons and exchange rate risk exposure (see Kouvelis and Gutiérrez [50]).

The introduction of a fixed ordering cost is the newsvendor’s most relevant extension in the context of multi-period problems. In such case, and assuming that the cost function is, say, convex, optimal policies turn out to be of the form (s, S) , which means that if the initial stock equals $x \leq s$, an amount of $S - x$ units should be ordered; otherwise, no order should be placed.

2.1.2 Multi-Period Models and (s, S) Policies

The purpose of this section is to review the mathematical models proposed to analyze the multi-period formulation of the simplest and best understood problem in stochastic inventory theory, the single-item, single-facility inventory problem. The two most relevant questions in relation to this problem are: first, determine the form of optimal policies, and second, find the optimal parameters which minimize the cost among all policies of that form. A good review of the fundamental results in this area prior to 1972 can be found in Tijms [80].

The first mathematical models proposed to manage this problem date back to the early 1950's. Early results, for example those by Arrow, Harris and Marschak [2], and Dvoretzky, Kiefer and Wolfowitz [20], [21], were impractical because of the limited scope of their assumptions, which included zero lead times and, sometimes, no ordering costs.

Scarf [74] was the first to find optimal policies for a model of sufficient generality to have wide practical applicability. The main assumptions of this model include constant lead times, fixed ordering costs, periodic review, demand backordering, and a convex inventory cost function; its objective is to minimize the expected costs incurred over an n -period time horizon. By using a dynamic programming argument and introducing the concept of K -convexity, Scarf showed that the optimal policy for this model is of the form (s_k, S_k) for each period k . This proof was extended to the discrete demand case by Zabel [84].

One of the classical results of inventory theory is that (s, S) policies are also optimal for infinite horizon problems, in which the objective is to minimize either the total discounted cost or the undiscounted average cost per period. The results

to be reviewed below show that stationary (s, S) policies can be found which are optimal for infinite horizon models regardless of the initial state. By contrast, optimal policies for finite horizon models in general have different parameters from period to period, even if the data are stationary, due to the end of horizon effect.

Proving the optimality of a policy for an infinite horizon model involves the verification of the so called optimality equation, which, to our knowledge, was first formulated by Bellman, Glicksberg and Gross [10]. One of the terms of this equation is usually referred to as the differential cost function, and must also be found. For a good exposition of dynamic programming and Markovian decision processes, including the optimality equation, see for example, Bertsekas [11] or Puterman [62].

The first proof of the optimality of (s, S) policies for infinite horizon problems was developed by Iglehart [42] in 1963. This proof, which uses Scarf's model with the objective function changed to minimize the discounted expected total cost, shows that both an optimal (s, S) policy and its corresponding differential cost function can be characterized as limits of sequences of finite-horizon optima $\{(s_n, S_n)\}$ and their differential costs $\{f_n(\cdot)\}$. Veinott and Wagner [82] extended this result to the case in which the inventory cost function is quasiconvex instead of convex.

Other authors have proposed proofs of Iglehart's result which bypass the finite horizon considerations. For example Johnson [48] uses a policy improvement argument to prove that an (s, S) policy can be found which satisfies the optimality equation. As pointed out in [51] and [34], this proof is incorrect because it uses finite-state Markovian decision processes results without the appropriate

modifications for infinite state spaces. Zheng [85] proposes a proof for both the discounted and undiscounted cost cases, which includes a closed form expression for the differential cost function and a verification of the optimality equation based on properties of the optimal parameters (s^*, S^*) . This proof is simpler, more illustrative and more insightful than the earlier proofs.

Simultaneously with these developments on the theoretically important issue of the optimality of (s, S) policies, other research efforts were directed at the more practical question of finding optimal parameters (s^*, S^*) . Because of the ill behavior of the cost function, which is known to be non-quasiconvex and which may exhibit more than one local minimum, this used to be considered a computationally expensive task. This is no longer the case, however, since the computational complexity of the fastest algorithm proposed to date is only 2.4 times that of evaluating the cost of a single (s, S) policy.

The oldest algorithm to compute optimal policies, proposed by Veinott and Wagner [82], relied mostly on enumeration aided by bounds on the optimal values. Later improvements of this algorithm are due to Archibald and Silver [1], who expressed the objective function as a convex function in one of the variables, and to Bell [9], who used stopping rules. Heuristic methods have been proposed, among others, by Roberts [64], Porteus [59] [60], Wagner, O'Hagan and Lundh [83], Ehrhardt [22], [23], Freeland and Porteus [35], [36], Sivazlian [78], Schneider [77], and Sahin and Sinha [73]. Each of these heuristics performs well in some instances and poorly in others and it is hard to tell a priori how well a given heuristic will perform.

Federgruen and Zipkin [34], [31], proposed an iterative approach based on the

policy iteration technique from standard dynamic programming theory, and which exploits the structure of the inventory problem. More recently, Zheng [86] developed the algorithm which is computationally least burdensome. This algorithm is based on tighter, monotonically convergent bounds on the optimal parameters, and new properties of the cost function. Its computational complexity is only 2.4 times that of evaluating the cost of a single (s, S) policy.

We conclude this section with two remarks. First, the periodic review model is equivalent to a continuous review formulation in which the demand is generated by a compound counting process and ordering decisions can only be made upon arrival epochs. Such a formulation was first proposed, to our knowledge, by Beckman [8] in 1961. It is also mentioned in [31] and [86]. Federgruen and Schechner [30] develop the cost function in detail and Zipkin [88] presents a good description of the underlying probabilistic aspects and proposes an extension of this model to account for random lead times. This extension only works if the lead times satisfy certain probabilistic conditions.

The second remark refers to the case when the reorder point is always reached exactly, a situation referred to in the literature as ‘no overshooting.’ In this case (s, S) policies are often referred to as (r, Q) policies, which means that whenever the inventory reaches $r(= s)$, $Q(= S - s)$ units are ordered. It is often assumed that the stationary distribution of the inventory position is uniform. In this case the average cost is rather insensitive to the choice of the ordering quantity, see Zheng [87]. Another interesting result provides a procedure to find parameters (r, Q) when only the mean and the variance, but not the form of the distribution, of the lead time demand are known, see Gallego [38] [37]. This procedure was

proposed before by Scarf [75] for the newsvendor problem.

2.2 Multi-Echelon Inventory Systems

2.2.1 Serial Systems

Many of the fundamental concepts for the analysis of multi-facility inventory systems first appeared in a seminal paper by Clark and Scarf [17], which considered a periodic review N -stage serial system in which each stage replenishes its stock from the next up stream facility. This model had been proposed earlier by Clark [16], and its main assumptions include: external demand that is experienced at the lowest stage only, constant lead times at each facility, fixed ordering cost incurred at the highest echelon only, backlogging of unmet demand, a convex inventory cost function at the lowest echelon, and inventory costs dependent only on echelon stock—a concept to be defined shortly.

Clark and Scarf proved that an optimal policy for this model is of the form order-up-to- R for all the facilities except for the highest echelon, which should follow an (s, S) policy. This result was extended to the infinite horizon case by Federgruen and Zipkin [33].

One of the most interesting features of this work lies in its decomposition approach, which not only simplifies the problem by reducing it to N single-stage subproblems, but also permits using existing single-stage techniques to find a solution. Also important is the introduction of induced penalty cost functions, which penalize a facility for the extra costs incurred elsewhere in the system due to the insufficiency of stock at the facility in question to supply orders placed by the next lower stage.

Another important idea introduced by Clark [16] is the concept of a facility's *echelon inventory*, which is defined as all the stock either currently at the facility, or that which has been there previously and is still in the system. Equivalently, a facility's echelon stock is equal to the stock at the facility plus all the stock on-hand at or in-transit to all its descendants. A crucial assumption in the analysis of multi-facility systems, although general enough for most cases, is that inventory costs can be expressed as functions of echelon stocks only. Optimal policies also turn out to be functions of echelon stocks only.

Perhaps the main limitation of the Clark and Scarf serial inventory model is the absence of ordering costs at the lower echelons. The introduction of such costs complicates the analysis to such an extent that no optimal solution is known to date. In view of this difficulty, the same authors [18] proposed an approximation approach for this problem based upon upper and lower bounds for the induced penalty function.

2.2.2 Distribution Systems under Uncertainty

Being able to find an optimal solution for a multi-echelon inventory model is an exception, not the rule. In general, optimal policies for such models are extremely complex, difficult to find and, in most cases, unknown. Classical examples include the serial system just mentioned, and distribution systems, i.e., systems in which each facility has a unique supplier. These are also referred to as arborescence systems.

Most of the literature available in this area focuses on two-echelon distribution systems with one depot and multiple retailers. Policies for such systems may be either decentralized, when management of each facility is based on the observation

of its own status only, or centralized, which allow for system-wide coordination at the expense of higher information requirements. We will address centralized policies first.

Clark and Scarf [17] observed that their optimality result on serial systems could be extended to distribution systems, but the resulting policies would no longer be optimal because of the additional problem of balancing inventory among the retailers. They proposed a relaxation to the problem by assuming that inventory at the retailers always stays in balance. Federgruen and Zipkin [33] showed that this assumption is a satisfactory approximation when the retailers exhibit normally distributed demands with the same coefficient of variation. A related work by the same authors [32] assumes that no inventory is held at the depot and presents a myopic allocation heuristic which is optimal for an approximate model. A review of distribution systems with centralized control can be found in Federgruen [29].

Ehrhardt, Schultz and Wagner [24] present a good heuristic for decentralized policies by assuming that the depot's demand consists of replenishment orders from retail facilities which follow (s, S) policies. A single-facility heuristic is used to find a good policy for the depot and the result is compared to policies obtained by simulation. Other authors have studied this problem assuming that the retailers face Poisson demands and follow a policy of a given form. See, for example, Axsäter [3], Deuermeier and Schwarz [19], and Muckstadt [54]. A review of distribution systems with decentralized policies can be found in Axsäter [3].

An insightful and relevant result concerning the design and operation of cost effective distribution systems is the reduction in uncertainty due to inventory consolidation. Often referred to as 'risk pooling effect,' this principle enables one to

raise service levels while reducing inventory carrying costs. This fact was first reported by Eppen [25], who observed that the holding and penalty costs of a decentralized one-period stochastic system always exceed those of a centralized one, and coined the term ‘statistical economies of scale’ to describe this phenomenon. Related results include Eppen and Schrage [26], who show that the backordering costs of a two-echelon distribution system decrease if the depot delays the allocation of stock to the retailers from the time an order is placed with an external supplier to the time the order arrives to the depot. Jackson and Muckstadt [44] identify further uncertainty reduction benefits when the depot has more than one opportunity to allocate stock to the retailers between depot replenishments.

2.2.3 Policies with Guaranteed Good Worst-Case Performance

Since optimal policies for most multi-echelon inventory systems are either unknown or extremely complex, some authors have sought simple policies with a good worst-case performance. In this connection, the so called nested, power-of-two policies have proven to be highly effective for several multi-stage generalizations of the classical EOQ model. A policy is said to be *nested* if each facility sends out a shipment at every epoch in which it receives stock. A *power-of-two* policy is one in which all the reordering time intervals are of the form $2^k \cdot B$, where B is the length of a base planning period and k is an integer. To ease the presentation of the results below, we also define the *effectiveness* of a policy as 100% times the ratio of the infimum of the costs over all feasible policies to the cost of the policy in question. A good review of the results presented in this section can be found in Muckstadt and Roundy [55].

Early results in nested, power-of-two policies were developed simultaneously and independently by Jackson, Maxwell and Muckstadt [46], who recommended such policies for the joint replenishment problem, and by Roundy [69], who proved that these policies are 98% effective for two-echelon distribution systems. Roundy [68] later extended this 98%-effectiveness result to multi-echelon systems in which the flow of stock can be represented by an acyclic directed graph, under assumptions similar to the classical EOQ model with no backorders. In both cases, the 98% bound rests on the assumption that the length of the base planning period is a variable. If, more realistically, this period is predetermined, the efficiency of nested, power-of-two policies is at least 94%.

General application algorithms to compute 94%- and 98%-effective policies can be found in Maxwell and Muckstadt [52] and Roundy [68]. Faster algorithms tailored to special cases include Roundy [66], for multi-item distribution systems, Muckstadt and Roundy [56], for the joint-replenishment problem; and Jackson and Roundy [47], for tree-structured systems.

Extensions of this model have been proposed by Mitchell [53], who shows that both the effectiveness result and the algorithm for distribution systems can be extended to permit backorders. Additionally, Jackson, Maxwell and Muckstadt [45] and Roundy [67] have introduced capacity constraints into this model.

These results have spurred interest in obtaining extensions to stochastic demand. Most of the research so far has been concerned only with two-stage serial inventory systems. In this dissertation, we present 93%-effective policies for a two-echelon serial inventory system with non-zero lead times at both facilities in such a system. As far as we know, this result is the most general of its kind developed

to date. Previous efforts include Atkins and De [5], who use an approach very similar to ours. Their result, however, is incomplete because their worst performance proof does not hold in all cases. A similar work by Chen [15] presents a valid 94%-efficient policy. However, this work assumes zero lead times at the depot, an assumption much less general and realistic than the weak nestedness condition used in our approach.

2.3 Current Inventory Theory's Limitations

Perhaps the main virtue of most inventory models is that they formulate, in as simple a way as possible, the main trade-offs involved in a stock replenishment decision, namely the cost of carrying inventory, the cost of not having inventory available when it is demanded, and the cost of placing an order. Some inventory theory's shortcomings stem from these models' assumptions. Another important limitation in the theory is that many relevant extensions of these models are not well understood.

The main limitations inherent in the models' assumptions include the difficulty of estimating some of the parameters involved, and the way demand is modeled. It has been written that the main problem in inventory theory is estimating the parameters needed to apply a model. This is particularly true for the backordering cost parameter p , and to a lesser degree for the fixed ordering cost K .

Virtually all inventory models consider the demand to be either deterministic or stochastic and stationary. In many real applications, by contrast, demand falls in between these two extremes. Forecasts for future demand often consist of a deterministic firm-orders component, plus an additional stochastic component that

typically is not stationary. Classical ways of modeling demand may be inadequate for such cases.

Some relevant extensions of the classical inventory models which are not well understood include: non-stationarity, forecasting, capacity constraints, multi-item systems with random demands, multi-facility systems with random demands, or any combination of these. Overcoming some of these difficulties is the main motivation behind this dissertation.

Chapter 3

A New Lower Bound on the Cost of Distribution Systems

This chapter focuses on the problem of making stock-replenishment decisions in a distribution inventory system with the ultimate purpose of serving external, stochastic demand in a cost-effective way.

The availability of system-wide information and the degree of decision-making centralization are two important issues in managing a multi-echelon inventory system. The availability of global information makes system-wide coordination possible at the expense of higher information requirements. Centralized decision-making enables a central planner to minimize total costs in the system but creates a management problem since local managers are often given responsibility for their facilities' profits without authority over ordering and shipping decisions. Delegating these decisions to each manager, even if global information is available, may result in decisions that are good from a facility's viewpoint, but bad for the system as a whole.

This chapter formulates a multi-facility distribution inventory model with centralized control, which generalizes the most widely used single-facility inventory models: Veinott and Wagner’s [82] periodic review model and the continuous demand, continuous review (r, Q) model. Our model allows for constant travel times and fixed shipping costs.

No optimal solution is known for this model. This is not surprising. In fact, as mentioned in Chapter 2, no optimal solution is known for a simpler model in which demand occurs at a deterministic, constant rate. Rather than seeking optimal policies, this chapter focuses on finding lower bounds for the cost of a distribution inventory system governed by an arbitrary centralized policy. These bounds can be used to evaluate the performance of heuristic policies.

Two bounds are presented in this chapter. One uses fixed shipping costs and is valid for any centralized, non-anticipatory policy—a concept to be defined later. The other introduces a non-fixed imputed shipping cost and is valid only for weakly-nested policies—a class of policies proposed for the first time in this dissertation, which extend the concept of nestedness to systems with non-zero lead times and stochastic demands.

These bounds are based on a decomposition approach that reduces a multi-facility model to several single-facility subproblems. This approach is attractive both theoretically and practically. From a theoretical point of view, it has the advantages of reducing the computational effort of finding good policies and permitting the use of existing single-facility techniques. From a practical viewpoint, this approach provides a cost scheme for each facility that allows decentralized decision-making on the basis of global information. This is accomplished by using

induced costs that mimic the effects on the system of a facility's decisions by penalizing each facility for the extra costs its decisions may cause elsewhere in the system.

The rest of this chapter is organized as follows. Section 3.1 reviews some classical results on the two single-facility inventory models to be used in the rest of the chapter, Veinott and Wagner's periodic review model and the (r, Q) continuous review model. Section 3.2 formulates a multi-facility inventory model, presents and comments on its assumptions, and shows how to find an equivalent formulation in which travel times are eliminated using the concept of 'shipment-based inventory position'. Section 3.3 shows how a facility with no successors induces a single-facility model, which we refer to as 'the single-stage subproblem'. Section 3.4 presents the fundamental result of this chapter, a lower bound for the long-term average cost incurred by a facility with no successors. Section 3.5 develops two lower bounds for serial inventory systems. One of these bounds is valid for any centralized non-anticipatory policy. The other bound is tighter, but valid only for weakly nested policies, a new class of policies which extends the concept of nestedness to stochastic systems with non-zero lead times.

3.1 Classical Single-Facility Inventory Models

This section reviews the two most commonly used infinite-horizon, single-item, single facility inventory models, the classical periodic review model (see for example Iglehart [42] and Veinott and Wagner [82]), and the continuous review (r, Q) model (see for example Zheng [87]). In both cases we limit our attention to the minimization of the average cost per time period.

Both of these models share the following assumptions:

Assumption 3.1.1 *The facility replenishes its inventory from an external source whose supply is unlimited, and the amount ordered arrives at the facility with a constant lead time L after the placement of the order. If the review is periodic, L is an integral multiple of the review period. In all cases $L \geq 0$.*

Assumption 3.1.2 *Unmet demand is backordered.*

Assumption 3.1.3 *The cost of placing an order is fixed and equal to K regardless of the size of the order.*

3.1.1 Classical Periodic-Review Stochastic Inventory Model

This is perhaps the most popular infinite-horizon model in stochastic inventory theory, as indicated by the extensive number of articles which use its assumptions. A finite-horizon version of this model was first proposed by Scarf [74] in 1960. Iglehart [42] extended Scarf's results to the infinite horizon case.

This model assumes that the inventory in the system is observed and ordering decisions can be made at equally spaced time intervals referred to as review periods. These review periods are labeled $0, 1, 2, \dots$. The demands $d(0), d(1), \dots$ that occur during these review periods are independent, discrete, nonnegative random variables with a common probability distribution given by $p_k = \Pr\{d(t) = k\}$, $k = 0, 1, 2, \dots$, $t = 1, 2, \dots$.

Definition 3.1.4 *The facility's net inventory is defined as the inventory on-hand minus outstanding backorders. $NI(t)$ is the facility's net inventory immediately after placement and delivery of orders in period t .*

The following sequence of events takes place during each review period t , $t = 1, 2, \dots$:

1. Initially the facility's net inventory is equal to $NI(t-1) - d(t-1)$.
2. A new order for quantity $q(t)$ is placed. This order will arrive at the facility during period $t + L$. Order placement has no effect on the net inventory.
3. Orders due for receipt at period t arrive, causing the net inventory to increase to

$$NI(t) = NI(t-1) - d(t-1) + q(t-L).$$

4. Costs for the review period are incurred which are equal to

$$K\chi_{[q(t)>0]} + g(NI(t)),$$

where $\chi_{[\cdot]}$ is the indicator function and $g(\cdot)$ is called the *inventory cost rate*.

5. The period's demand $d(t)$ occurs and as a result the net inventory decreases to $NI(t) - d(t)$.

The inventory cost rate $g(x)$ is often of the form $E[h(x - d(0))^+ + p(x - d(0))^-]$, where the parameters h and p represent the cost of carrying or backordering one unit of inventory, respectively, during a time period, $x^+ = \max\{x, 0\}$, and $x^- = (-x)^+$. To avoid trivialities, we do not consider policies that incur an infinite backlog or hold an infinite amount of inventory. Thus, the total amount purchased is equal to the total amount demanded in the long run and therefore linear purchasing costs need not be included in the model.

Definition 3.1.5 *The facility's inventory position is defined as the net inventory plus the amount ordered but not yet delivered to the facility. $IP(t)$ is the inventory position immediately after placement and delivery of orders in period t .*

Note that $NI(t)$ and $IP(t)$ are observed at the same point in the review period. Because of Assumption 3.1.1,

$$NI(t + L) = IP(t) - D(t, t + L), \quad (3.1)$$

where $D(t, t + L)$ represents the amount demanded during the interval $[t, t + L)$.

In view of (3.1), Veinott and Wagner [82] point out that an equivalent zero-lead-time formulation can be obtained by charging the inventory cost incurred in period $t + L$ to the inventory position $IP(t)$ in period t . Consistent with this shift in time, the inventory cost function is redefined as

$$G(IP) = E_D[g(IP - D)], \quad (3.2)$$

where the expected value is taken over D , the demand that occurs over a lead time L . Thus $D \sim D(0, L)$.

The renewal function $M(x)$ of the process generated by the sequence $\{X_n, n = 1, 2, \dots\}$, which in this context represents the expected number of periods until at least x units are demanded, is very useful in computing the long-term average cost of an (s, S) policy (see, for example, Veinott and Wagner [82]). The renewal function can be computed by the formula:

$$M(0) = 0; \quad M(i) = \sum_{j=0}^{i-1} m(j), \quad i = 1, 2, \dots; \quad (3.3)$$

where the renewal density, $m(i)$, represents the expected number of review periods

which begin with a total amount demanded of i units. The renewal density can be computed with the expression:

$$m(0) = \frac{1}{1-p_0}; \quad m(i) = \sum_{j=0}^i p_j m(i-j), \quad i = 1, 2, \dots \quad (3.4)$$

Equations (3.3) and (3.4) imply that $M(i)$ exists and is finite for all i provided that $p_j < 1$ for all j .

The renewal function also satisfies the renewal equation

$$M(i) = 1 + \sum_{j=0}^{i-1} p_j M(i-j), \quad i = 1, 2, \dots \quad (3.5)$$

Let $I(s, i)$ denote the expected inventory costs incurred until the inventory position drops to or below s , when the inventory position starts at i and no orders are placed in the meantime. Using renewal theory, we can show that $I(s, i)$ is defined by

$$I(s, i) = \sum_{j=0}^{i-s-1} G(i-j) m(j), \quad i > s, \quad (3.6)$$

and satisfies the renewal equation

$$I(s, i) = G(i) + \sum_{j=0}^{i-s-1} p_j I(s, i-j). \quad (3.7)$$

It is known that optimal policies for this model take the form (s, S) (see Iglehart [42] and Zheng [85]). A standard application of the renewal reward theorem (see for example Ross [65], proposition 7.3.), shows that the average cost per period incurred by an (s, S) policy is equal to

$$c(s, S) = \frac{K + I(s, S)}{M(S-s)}. \quad (3.8)$$

Optimal parameters (s^*, S^*) can be chosen that satisfy the following lemma, which is due to Zheng (see [85], Lemma 1).

Theorem 3.1.6 *Let $G(\cdot)$ be convex, $G(y) \rightarrow \infty$ as $|y| \rightarrow \infty$, and let y^o be a minimum point of $G(y)$. Then there exist s^* and S^* that satisfy*

$$(i) \quad c^* := c(s^*, S^*) = \min_{s < S} c(s, S);$$

$$(ii) \quad s^* < y^o \leq S^*;$$

$$(iii) \quad G(s^*) \geq c^* > G(s^* + 1);$$

$$(iv) \quad G(S^*) \leq c^*.$$

3.1.2 Continuous Review Models and (r, Q) Policies

This continuous review, continuous demand model was first proposed by Zheng [87]. Let $NI(t)$, $IP(t)$ be as defined as in Section 3.1.1. Also, let $D(t)$, $t \geq 0$ denote the cumulative demand that occurs during the interval $[0, t]$.

It is postulated that the demand process $\{D(t) : t \geq 0\}$ satisfies the following axioms:

A1. $D(0) = 0$ and the increments are independent, i.e.

if $0 \leq t_0 \leq t_1 \leq \dots \leq t_k$, then

$$P[D(t_i) - D(t_{i-1}) \in H_i, i \leq k] = \prod_{i \leq k} P[D(t_i) - D(t_{i-1}) \in H_i];$$

A2. The increments over time intervals of the same length are identically distributed, i.e. the distribution of $D_t - D_s$ depends only on $t - s$;

A3. The sample paths are continuous and non-decreasing.

These axioms are limited from a theoretical point of view because the only process that satisfies them is one in which the demand occurs at a constant, continuous, and deterministic rate. To see why, note first that axioms A1, A2 and A3

imply that the demand process is a Lévy process, i.e., a process with stationary independent increments and whose sample paths have at most jump discontinuities of the first kind almost surely. It is a known result that any Lévy process may be represented as the sum of a linear drift, plus a Brownian motion, plus a jump process, where a jump process is understood as a limit of independent superpositions of compound Poisson processes with varying jump sizes. Since the continuity and the monotonicity of the sample paths preclude the existence of the process's jump and Brownian components, respectively, it follows that the demand process consists only of a linear drift. For a good reference on stochastic processes with independent increments, see Ito [43], Section 4. A good summary of these results can be found in Bhattacharya and Waymire [12], p. 349-356.

Despite this limitation, these axioms are approximately satisfied by many real-world systems and remain the basis for insightful and highly effective inventory models.

The demand rate $\lambda := E[D(t)]/t$ is assumed to be finite and strictly positive. As argued in the proof of Lemma 3.1.7 below, the expected time until x units are demanded is equal to

$$M(x) = \frac{x}{\lambda}. \quad (3.9)$$

Inventory costs are assumed to accumulate continuously at a rate $g(NI(t))$, typically equal to $h(NI(t))^+ + p(NI(t))^-$. As in the periodic review, constant lead-time model, (3.1) holds and an equivalent formulation with zero lead times can be obtained by charging the expected inventory cost rate at time $t + L$ to the inventory position at time t . Thus inventory costs can be assumed to accumulate continuously at a rate $G(IP(t))$, where $G(IP) = E_D[g(IP - D)]$ as in (3.2), and

D represents the demand that occurs during a lead time.

The following lemma, new to our knowledge and proven in appendix B, is parallel to equation (3.6).

Lemma 3.1.7 *Assume that $G(\cdot)$ is measurable and nonnegative, that its domain includes the interval $[r, S]$, and that the inventory position at time $t = 0$ is equal to S . Then $I(r, S)$, the expected cost incurred until the inventory position reaches some predetermined level $r < S$, assuming that no orders are placed in the meantime, is equal to*

$$I(r, S) = \frac{1}{\lambda} \int_r^S G(y) dy. \quad (3.10)$$

Proof

The proof of this lemma is included as Appendix B. Note that setting $G(y) \equiv 1$ in (3.10) we obtain (B.5). \square

If the system is governed by an (s, S) policy, the continuity of the sample paths implies that the reorder point s is always reached exactly (a situation often referred to in the literature as no overshooting of the reorder point). When this is the case, (s, S) policies and (r, Q) policies become indistinguishable. In an (r, Q) policy, Q units are ordered whenever the inventory position reaches r .

The cost function, which is parallel to (3.8) and can be derived by a similar argument, is given by

$$c(r, Q) = \frac{\lambda K + \int_r^{r+Q} G(y) dy}{Q}. \quad (3.11)$$

It can be proven that (r, Q) policies are optimal for this problem (see Theorem 4.1.3). Furthermore, optimal parameters (r^*, Q^*) are given by the following theorem, due to Zheng (see [87], Theorem 1).

Theorem 3.1.8 *Assume that $G(\cdot)$ is convex and that $G(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. Then (r^*, Q^*) is optimal for $\min_{r,Q} c(r, Q)$ if and only if*

$$c(r^*, Q^*) = G(r^*) = G(r^* + Q^*). \quad (3.12)$$

3.2 Stochastic Modeling of Distribution Systems

The purpose of this section is to formulate an inventory model which extends the classical single-stage inventory model's assumptions to distribution systems. These assumptions include stochastic demand, constant travel times, non-zero shipping costs, and either periodic or continuous review. This section consists of two subsections; 3.2.1 formulates a distribution system inventory model, and 3.2.2 shows how an equivalent zero-lead time formulation can be obtained.

3.2.1 Model Formulation

This section includes the notation to describe a distribution inventory model, the assumptions for that model, and some comments about these assumptions.

Notation and Assumptions

Definition 3.2.1 Distribution System

Consider the graph representation of a single-item multi-echelon inventory system consisting of N facilities, labeled $1, \dots, N$, in which an arc from m to n means that facility n replenishes its stock from facility m . Let A denote the set of all arcs. If $(m, n) \in A$ we say that m is a predecessor of n and that n is a successor of m . A distribution system is a multi-echelon inventory system that satisfies two additional conditions. First, the facilities can be labeled in such a way that if $(m, n) \in A$,

then $m > n$. Second, each facility $n < N$ has an unique predecessor, and N has no predecessors. Thus N is the only facility that replenishes stock from an external supplier. A descendant of a facility is a location that replenishes its stock directly or indirectly from the facility in question.

Assumption 3.2.2 System Reviews

The distribution system is observed and stock-replenishment decisions can be made only at certain points in time, referred to as review times, which may take place either periodically or continuously. Let T denote the set of review times. If system reviews occur periodically, then $T = \{0, 1, \dots\}$. If system reviews occur continuously, then $T = \{t \geq 0, t \in \mathbb{R}\}$.

Assumption 3.2.3 Demand

$D_n(t)$, $n = 1, \dots, N$, $t \in T$ represents the cumulative demand that occurs at facility n plus all of its descendants from time zero to time t , inclusive. Note that the random variables $D_1(t), D_2(t), \dots, D_n(t)$ are not necessarily independent. The cumulative demand process $D_n(t)$ has independent increments which are identically distributed for time intervals of equal length. External demand occurs only at facilities with no successors. The external demand at each location has non-decreasing sample paths $D_n(t; \omega)$. The sample paths are continuous functions of time if system reviews are continuous. The notation $D_n(s, t)$, $n = 1, \dots, N$, $s, t \in T$ will be used to represent the demand that occurs during the time interval $[s, t) \cap T$. When review is periodic we use $d_n(t) = D_n(t) - D_n(t - 1)$ to represent the one-period demand.

Assumption 3.2.4 Backordering of Demand

Unmet external demand is backordered.

Assumption 3.2.5 Travel Times

Material arrives at facility n a constant lead time L_n after it is shipped by the facility's supplier. If the review is periodic, L_n is an integral multiple of the length of the review period. In all cases $L_n \geq 0$.

Definition 3.2.6 Echelon Net Inventory

A facility's echelon net inventory is defined as all of the stock at or downstream from the facility. This includes the stock on-hand at the facility and at all of its descendants, plus the stock in transit to all of the facility's descendants, minus the amount backordered at the facility and at its descendants. If the system is reviewed periodically, $NI_n(t)$, $n = 1, \dots, N$, $t \in T$ represents facility n 's echelon net inventory immediately after shipments are dispatched and received in review period t but before $d_n(t)$ occurs. If system reviews occur continuously then $NI_n(t)$ is observed continuously and $NI_n(t)$ is a right-continuous function of time for each sample path.

Definition 3.2.7 Shipment-Based Inventory Position

The shipment-based echelon inventory position of facility n is defined as its net echelon inventory, plus the stock already shipped to facility n that has not yet arrived. This definition differs from the conventional use in that it does not include the amount ordered by, but not yet shipped to, the facility. $IP_n(t)$, $n = 1, \dots, N$, $t \in T$ represents facility n 's shipment-based echelon inventory position at time t if system reviews occur continuously, and immediately after shipments are dispatched and received in period t if review is periodic.

Assumption 3.2.8 System Evolution

Let $q_n(t)$, $n = 1, \dots, N$, $t \in T$ represent the amount shipped to facility n at the

beginning of review period t if the system is reviewed periodically, and the amount shipped at time t if the system is reviewed continuously. The following equations hold for periodic review systems and for all $t = 0, 1, \dots$, and $n = 1, \dots, N$:

$$NI_n(t+1) = NI_n(t) - d_n(t) + q_n(t - L_n). \quad (3.13)$$

For continuous review systems, $NI_n(t)$ and $IP_n(t)$ are continuously updated as shipments are dispatched and received, and as demand occurs. They are right continuous as functions of time.

Assumption 3.2.9 Stock Availability Constraint

We must have

$$\sum_{m:(n,m) \in A} IP_m(t) \leq NI_n(t), \quad \forall t \in T, \text{ for all facilities } n. \quad (3.14)$$

This states that inventory cannot be shipped to a facility before it arrives at the supplier of the facility.

Assumption 3.2.10 Shipping Costs

The cost of sending a shipment to facility n ($n = 1, \dots, N$) at time t that raises its shipment-based echelon inventory position to some level $IP_n(t)$, is equal to $K_n + \text{ord}_n(IP_n(t))$, where K_n is a positive constant and $\text{ord}_n(\cdot)$ is a nonnegative function.

Assumption 3.2.11 Inventory-Related Costs

The function $\sum_{n=1}^N g_n(NI_n(t))$ is used to model the inventory-related costs. This is an additive function of the echelon net inventories $NI_n(t)$ ($1 \leq n \leq N$) of the facilities. Each function $g_n(\cdot)$ is convex. For each facility with no successors

$g_n(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. The total inventory-related cost incurred during any interval $[s, t)$, for $s < t$ and $s, t \in T$, is equal to

$$\sum_{n=1}^N \int_s^{t-} g_n(NI(t)) d\mu(t). \quad (3.15)$$

This integral has different interpretations according to whether the system reviews are periodic or continuous. For periodic review, μ is a discrete measure which places a mass of one at each $t \in T$. The integral in (3.15) represents the inventory costs incurred during periods $s, s+1, \dots, t-1$, and is equal to $\sum_{n=1}^N \sum_{k=s}^{t-1} g_n(NI(k))$. If review is continuous, μ is the Lebesgue measure and (3.15) is equal to $\int_s^t g_n(NI(t)) dt$.

Definition 3.2.12 System History

1. $\mathcal{F} = \sigma(D_n(t) : n = 1, \dots, N, t \in T)$ is the sigma-field generated by the demand process.
2. $u(t) = (NI_n(s), q_n(s), D_n(s) : s < t, s \in T, n = 1, \dots, N)$, for all $t \in T$, represents the system history up to time t .

Assumption 3.2.13 Centralized, Non-Anticipatory Policies

We consider centralized, non-anticipatory policies π , which are sets of functions $q_n : (t; u(t)) \rightarrow \mathbb{R}^+$, $n = 1, \dots, N$, $t \in T$, from system histories into feasible shipping decisions. Each function $q_n(t : u(t))$ is measurable \mathcal{F} . $q_n(t : u(t))$ is the amount shipped to facility n at time t . The dependence of q_n on $u(t)$ is often suppressed.

Assumption 3.2.14 Long-Term Average Cost of a Policy

For $t > 0$, let $C(t)$ denote the total shipping and inventory costs incurred during

the time interval $[0, t)$. These costs are defined in Assumptions 3.2.10 and 3.2.11.

The long-term average cost of policy π is defined as

$$\mathcal{C}(\pi) = \liminf_{t \rightarrow \infty} \frac{1}{t} C(t). \quad (3.16)$$

We assume that this limit exists.

If system reviews occur periodically, we will assume that the following sequence of activities takes place during each review period $t > 0$. This sequence of events is consistent with the assumptions above and is illustrated in figure 3.2.1.

At the beginning of period t , facility n has an echelon net inventory of $NI_n(t - 1) - d_n(t - 1)$ and a shipment-based echelon inventory position of $IP_n(t - 1) - d_n(t - 1)$.

A. *Ordering.* Facility N places an order. This does not change either the echelon net inventories or the shipment-based echelon inventory positions.

B. *Shipping and Receipt.* The following activities are performed at facilities N to 1, in that order,

(B1) Shipments sent to facility n at time $t - L_n$ are received. The facility's echelon net inventory increases to $NI_n(t) = NI_n(t - 1) - d_n(t - 1) + q_n(t - L_n)$.

(B2) Quantities $q_m(t)$, for all $(n, m) \in A$, are shipped to each successor m of facility n . The shipment-based inventory position of m increases to $IP_m(t) = IP_m(t - 1) - d_m(t - 1) + q_m(t)$.

Event	Net Inventory and Inventory Position After Event
Start	$\rightarrow NI_n(t-1) - d_n(t-1)$ $IP_n(t-1) - d_n(t-1)$
B1. Receipt $q_n(t - Ln)$	$\rightarrow NI_n(t) = NI_n(t-1) - d_n(t-1) + q_n(t - Ln)$ $IP_n(t-1) - d_n(t-1)$
B2. Dispatch $q_m(t) : (n, m) \in A$	$\rightarrow NI_n(t)$ $IP_m(t) = IP_m(t-1) - d_m(t-1) + q_m(t),$ $\forall (n, m) \in A$
C. Incurring of Costs	$\rightarrow NI_n(t)$ $IP_m(t), \quad \forall (n, m) \in A$
D. Demand $d_n(t)$	$\rightarrow NI_n(t) - d_n(t)$ $IP_n(t) - d_n(t)$

Figure 3.1: Activities and Inventory Quantities During a Review Period.

C. *Incurring of costs.* The costs for the review period are incurred. These costs are equal to

$$\sum_{n=1}^N \chi_{[q_n(t) > 0]} [K_n + \text{ord}_n(IP_n(t))] + \sum_{n=1}^N g_n(NI_n(t)).$$

D. *Demand.* The demands $d_n(t)$, $n = 1, \dots, N$ occur. Both the echelon net inventory and the shipment-based inventory position decrease by $d_n(t)$.

Note that the sequence in which shipments occur is important to ensure that stock is available to be shipped to a facility's successors in the same review period in which it arrives at the facility in question. If review is continuous, these activities take place continuously through time.

We conclude this section with some comments on the model assumptions. Assumptions 3.2.5 and 3.2.10 depart from most inventory models in that they focus on shipments rather than on orders. This difference stems from a fundamental distinction between single-facility and multi-facility systems. In the former, stock replenishment is triggered by orders served by an unlimited supplier. In the latter, stock replenishment is constrained by the availability of inventory and thus shipments and orders do not correspond. Actual inventory levels are determined by shipments, not by orders.

We account for shipping costs rather than ordering costs (Assumption 3.2.10). This implies that if, for example, several orders placed by a facility are consolidated into a single shipment sent from the facility's supplier, the shipping cost is incurred only once. This is realistic. Typically the placement of an order is primarily an information transfer function. The shipment costs is typically much larger. It includes the cost of moving the material, as well as transferring information on what was shipped.

Assumption 3.2.10 also introduces a non-fixed shipping cost term $\text{ord}_n(IP_n(t))$, which is a function of the ship-up-to inventory position at the time of the shipment, independent of the amount shipped. The meaning of this non-fixed shipping cost will be explained in Section 3.4; it will be useful in our method for obtaining lower bounds on costs so that a facility can be eliminated. For now, we only say that this cost represents the penalty cost imputed on a supplier for shipping a quantity so small that the receiver cannot distribute its fixed shipping costs as it had anticipated.

Inventory cost rate functions in Assumption 3.2.11 are typically based on functions of the form

$$\hat{g}_n(x) = h_n x + p_n x^-. \quad (3.17)$$

Parameter h_n represents the *incremental* cost of holding one unit of inventory for one time unit at facility n (also called the echelon holding cost). The parameter p_n is related to the cost of incurring backorders. For facilities n with successors, $p_n = 0$. See Appendix A for a discussion of p_n .

For continuous review systems, a typical choice for the cost rate function is

$$g_n(NI_n(t)) = \hat{g}_n(NI_n(t)), \quad (3.18)$$

and the parameters h_n and p_n are measured in dollars per unit carried or backordered per time unit.

A typical choice for the cost rate for periodic-review systems is function

$$g_n(NI_n(t)) = E_{d_n(t)}[\hat{g}_n(NI_n(t) - d_n(t))], \quad (3.19)$$

and the parameters h_n and p_n are measured in units carried or backordered per review period. This choice of the cost rate models the expected holding and back-

ordering costs that will be incurred at the end of the review period.

The policies defined in Assumption 3.2.13 are non-anticipatory because shipping decisions cannot be based on observations of future demand. They are centralized because shipping decisions are based on centralized control and system-wide information. Compared with their decentralized counterparts, in which shipping decisions are made locally and on the basis of local information, centralized policies have the potential to achieve lower inventory-related costs at the expense of higher information requirements.

3.2.2 Zero Lead-Time Reformulation

The purpose of this section is to describe an equivalent cost accounting scheme. This approach follows the one proposed by Veinott and Wagner [82] for periodic-review, single-facility models and is based on identity (3.1), which holds if the definition of the term $IP(t)$ is appropriately defined for multi-echelon inventory systems.

In contrast to the external unlimited supply case, an internal supplier may—by choice or necessity—delay the shipment of an order, or ship a quantity different from the amount ordered. This results in random lead times and uncertain amounts delivered. Consequently, the conventional definition of inventory position for multi-echelon systems, namely the amounts on-hand plus on order, does not contain enough information to determine the distribution of the net inventory L time units into the future because (3.1) no longer holds.

For this reason we defined the shipment-based inventory position. Using this definition for the notation $IP(t)$, equation (3.1) holds and constant lead times can be introduced into the model—with no additional difficulty—in the same way as

in single-facility models.

Thus we can account for the costs incurred as if shipments arrived immediately. Let $D_n \sim D_n(0, L_n)$ be the external demand served by facility n and which occurs over a time period of length L_n . Define $G_n(\cdot)$ by

$$G_n(IP_n) = E_{D_n}[g(IP_n - D_n)].$$

Instead of incurring cost $g_n(NI_n(t + L_n))$ at time $t + L_n$, we incur a cost of $G_n(IP_n(t))$ at time t . The two formulations are equivalent in expected cost.

We conclude this section by remarking that even though the shipment-based inventory position enables one to compute expected inventory costs without explicit knowledge of net inventories, the observation of such inventories is still necessary. This occurs because the definition of the shipment-based inventory position naturally introduces the constraint (3.14).

3.3 The Single-Stage Subproblem

This section defines and studies the single-stage subproblem, which is related to a facility with no successors in a distribution system and is fundamental for the lower bounds developed in this chapter. The single-stage subproblem is defined as follows:

Definition 3.3.1 *The term single-stage subproblem refers to the special case of the model in section 3.2 when $N = 1$ (i.e., there is only one facility).*

The optimality of (s, S) policies for the single-stage subproblem is a question of theoretical interest (as the basis for the lower bounds presented in this chapter) and practical importance (because (s, S) policies are easy to implement and intuitively

appealing). If the shipping costs are fixed, the results reviewed in Section 3.1 state that an (s, S) policy is optimal for the single-stage subproblem. This is not necessarily the case when shipping costs are not fixed as example 4.1.2 shows. Sufficient conditions for the optimality of (s, S) policies is the subject of Chapter 4. For now, we coin the following term to refer to this optimality.

Definition 3.3.2 *If n is a facility in a distribution system which has no successors, then the single-stage subproblem induced by n is the single-stage subproblem with shipping cost $K_n + \text{ord}_n(\cdot)$, inventory cost function $g_n(NI_n(t))$, demand process $\{D_n(t) : t \in T\}$, and lead time L_n . Note that there are no stock availability constraints.*

Definition 3.3.3 *A facility with no successors in a distribution inventory system is said to satisfy condition Σ if an (s, S) policy is optimal for the single-stage subproblem induced by the facility.*

3.3.1 Optimal Parameters (s^*, S^*)

This section focuses on characterizing optimal parameters (s^*, S^*) , which is a question independent from condition Σ .

Let us start by recalling the terms $I(s, i)$ and $M(x)$. The first term, $I(s, i)$ denotes the expected inventory costs incurred until the inventory position drops to or below s , when the inventory position starts at i and no orders are placed in the meantime. This term can be computed using (3.6) if system reviews occur periodically and (3.10) if they occur continuously. Another term of interest is $M(x)$, the expected time until the cumulative demand reaches or exceeds x units.

This expression can be computed using (3.3) for periodic review and (B.5) for continuous review.

When system reviews occur periodically, the average cost per period incurred by an (s, S) policy for the single-stage subproblem over an infinite time horizon is equal to

$$c(s, S) = \frac{K + \text{ord}(S) + I(s, S)}{M(S - s)}. \quad (3.20)$$

For continuous review, this cost reduces to

$$c(s, S) = \frac{\lambda(K + \text{ord}(S)) + \int_s^S G(y) dy}{S - s}. \quad (3.21)$$

Theorem 3.3.4 *There exist optimal parameters (s^*, S^*) satisfying:*

- (i) $c^* := c(s^*, S^*) = \min\{c(s, S) : s < S\}$;
- (ii) $s^* < y^0$, where y^0 is any point at which $G(\cdot)$ achieves its minimum;
- (iii) For continuous reviews: $G(s^*) = c^*$. If system reviews occur periodically, $G(s^*) \geq c^* > G(s^* + 1)$.

Proof

This proof is based only on the form of the cost functions (3.20) and (3.21) and requires a different argument for periodic and continuous review. The proof for the periodic review case mimics a proof developed by Zheng ([86] Lemma 3 and [85] Lemma 1) assuming fixed shipping costs. The references contain details of the proof. For the continuous review case, the proof is an exercise in calculus.

We start with the periodic review case. Since $G(y) \rightarrow \infty$ as $|y| \rightarrow \infty$ by assumption 3.2.11, there is a finite number of pairs (s, S) for which $c(s, S) \leq G(y_0) + \text{ord}(y_0) \geq c(y_0 - 1, y_0)$. Therefore there exist values s^* and S^* satisfying

(i). Parts (ii) and (iii) follow by contradiction arguments using the convexity of $G(\cdot)$ and the identity

$$c(s-1, S) = \alpha G(s) + (1-\alpha)c(s, S), \quad \text{where } \alpha = \frac{m(S-s)}{M(S-s+1)}.$$

Now we prove the theorem for the continuous review case. Note that $c(s, S)$ is non-negative and continuous for $s < S$, and that $c(s, S) \rightarrow \infty$ as $S-s \rightarrow 0$. Recall that $G(y)$ is nonnegative, convex, and tends to infinity as $|y| \rightarrow \infty$. Thus $c(s, S) \rightarrow \infty$ as $\max(|s|, |S|) \rightarrow \infty$. Consequently $\{(s, S) : c(s, S) \leq c(0, 1)\}$ is compact, and (i) holds. Part (ii) follows by perturbing s^* . Part (iii) follows from the condition $0 = \partial c(s, S^*)/\partial s = (S-s)^{-1}[c(s, S) - G(s)]$. \square

3.3.2 A Facility as Part of a Larger System

This section lays down a probabilistic framework to compare the problem of running a facility based on the observation of local demand only, as occurs in the single-facility subproblem, to the problem of running the same facility under a centralized policy for a larger system, which may use information about events external to the facility in question.

Consider a facility with no successors in a distribution inventory system. In the remainder of this section this facility will be referred to as ‘the facility.’ Let $\{D(t) : t \in T\}$ and $\{\vec{D}(t) : t \in T\} = \{D_n(t) : n = 1, \dots, N, t \in T\}$ denote the demand process generated at the facility and at all facilities in the system, respectively. These demands processes generate sigma-fields $\mathcal{F}^o = \sigma(D(t) : t \in T)$ and $\mathcal{F} = \sigma(\vec{D}(t) : t \in T)$.

For each $t \in T$, we define the sigma-fields $\mathcal{F}_t^o = \sigma(D(s) : s < t, s \in T)$ and $\mathcal{F}_t = \sigma(\vec{D}(s) : s < t, s \in T)$. One can think of \mathcal{F}_t^o as all the information

available from the observation of the demand at the facility up to time t . Similarly, one can think of \mathcal{F}_t as all the information available from the observation of the demand process at all the facilities in the system up to time t . It is clear from these definitions that $\mathcal{F}_t^o \subset \mathcal{F}_t$, for all $t \in T$.

A random variable τ defined on \mathcal{F} and with possible values on T , is called a *stopping time* for \mathcal{F} if

$$\{\tau \leq t\} \in \mathcal{F}_t, \quad \text{for all } t \in T. \quad (3.22)$$

In view of the interpretation of \mathcal{F}_t , this definition means that the event that a stopping time is less than or equal to t is completely decidable by the observation of the system's demand history up to time t . A stopping time τ_o for \mathcal{F}^o is defined in identical fashion as a random variable defined on \mathcal{F}^o and with values on T such that the event $\{\tau_o \leq t\}$ is in \mathcal{F}_t^o .

The following are common examples of stopping times. A constant is a stopping time for either \mathcal{F}^o and \mathcal{F} . The time to reach or exceed x units demanded at the facility is a stopping time for both \mathcal{F}^o and \mathcal{F} . Both the minimum and the maximum of two stopping times for \mathcal{F} are also stopping times for \mathcal{F} . The same applies to \mathcal{F}^o .

The time of sending the next shipment to a given facility in a distribution inventory system governed by a centralized, non-anticipatory policy is another example of a stopping time for \mathcal{F} . In this context, (3.22) precludes prevision on the part of the decision-maker, meaning that he must make the decision to send a shipment at time t based only on the information available to him at that time.

Let τ_1 be a stopping time for \mathcal{F} . Assume that a shipment is sent to the facility at $t = 0$ which raises its shipment-based inventory position to some level

$IP(0) = S$, and that no shipment is sent after time 0 and before time $\tau 1$. The total shipping and inventory cost incurred by the facility during the time interval $[0, \tau 1)$ is a random variable defined on \mathcal{F} , with values in \mathfrak{R}^+ , and defined by

$$R(\tau 1, S; \omega) = K + \text{ord}(S) + \int_0^{\tau 1(\omega)^-} G(S - D(t; \omega)) d\mu(t), \quad \forall \omega \in \mathcal{F} \quad (3.23)$$

where μ depends on whether system reviews take place periodically or continuously (see 3.15).

We are now ready to state the main results of this section, which consist of the following two theorems:

Theorem 3.3.5 *Any stationary policy π measurable \mathcal{F}^o that minimizes the long-term average cost incurred by the facility among all policies measurable \mathcal{F}^o is also optimal among all policies measurable \mathcal{F} .*

Proof

All of the proofs presented in Section 4.1 also hold when \mathcal{F} is used instead of \mathcal{F}^o , which proves the theorem for continuous review systems. The proof for periodic review systems is in Appendix C. \square

Theorem 3.3.6 *Assume that condition Σ is satisfied. Let c^* be defined as in Theorem 3.3.4 and let τ be any stopping time for \mathcal{F} . Then*

$$\liminf_{a \rightarrow \infty} \left[\frac{E_{\omega \in \mathcal{F}}[R(a \wedge \tau, S; \omega)]}{E_{\omega \in \mathcal{F}}[a \wedge \tau(\omega)]} \right] \geq c^* \quad (3.24)$$

for all S . Furthermore, (3.24) holds with equality for a stopping time τ satisfying

$$E_{\omega \in \mathcal{F}}[\tau(\omega)] < \infty. \quad (3.25)$$

Proof

If (3.24) fails then there is a stopping time τ and a finite number $a_1 > 0$ such that $\tau(\omega) = a_1 \wedge \tau(\omega)$ satisfies

$$\frac{E_{\omega \in \mathcal{F}}[R(\tau, S; \omega)]}{E_{\omega \in \mathcal{F}}[\tau(\omega)]} < c^*. \quad (3.26)$$

Note that (3.25) holds for τ . As a result of the strong Markov property, the subproblem renews itself probabilistically if a shipment that raises the inventory position up to S is sent upon observation of τ . Repeating this process an infinite number of times generates a renewal reward process which is equivalent to the inventory policy ‘ship up to S every time τ is observed.’ The cost of this policy is equal to the left-hand side of (3.26) as a result of the renewal reward theorem (see for example Ross [65] proposition 7.3.). Additionally, c^* is equal to the infimum among the long-term average costs of all policies measurable \mathcal{F} because of condition Σ and Theorem 3.3.5. This proves (3.24). Theorem 3.3.4 implies that a stopping time τ exists for which (3.24) holds with equality, and (3.25) follows from (B.5) and (3.3)–(3.4). \square

3.4 A Lower Bound for the Cost of a Facility

The purpose of this Section is to present an alternate cost accounting scheme for a facility with no successors in the distribution inventory model formulated in Section 3.2.1. We will show that for a given policy, the costs incurred by the alternate accounting scheme are lower bounds on the actual costs incurred. The alternate accounting cost scheme accumulates inventory costs at a rate $H(IP)$, a non-increasing function of the facility’s inventory position, and also incurs a non-

fixed shipping cost $J(S)$ when a shipment that brings the inventory position up to S is sent. Both of these cost functions are based on the optimal policy for the single-facility's subproblem.

The bound's inventory cost rate is defined by

$$H(y) = \begin{cases} G(y) & \text{if } y \leq s^*, \\ c^* & \text{otherwise.} \end{cases}$$

This function, proposed by Atkins and De [5] and similar in spirit to the one proposed by Clark and Scarf [17] for serial systems, is based on the following rationale. At all times the cost rate is at least c^* , which is the lowest average cost that the facility can achieve. If the facility's echelon inventory falls below its reorder point it will do so because of a decision made by the supplier. In that case the facility's optimal policy becomes infeasible. The extra costs incurred by the facility are modeled by $H(\cdot)$ and can be charged to the supplier.

The non-fixed shipping cost $J(\cdot)$ rests the observation that an optimal policy achieves cost c^* by distributing the shipping costs over the entire time period when $G(IP(t)) \leq c^*$. If, by shipping a small order quantity, the supplier does not raise the facility's inventory position to a level high enough to distribute all the facility's shipping costs, then the non-distributed portion of these costs is modeled by $J(\cdot)$ and can be charged to the supplier. Formally, the shipping cost function is defined as

$$J(S) = \begin{cases} K + \text{ord}(S) & \text{if } S \leq s^*, \\ K + \text{ord}(S) + I(s^*, S) - c^* \cdot M(S - s^*) & \text{if } s^* < S \leq S^* \wedge \bar{S}, \\ 0 & \text{otherwise;} \end{cases}$$

where

$$\bar{S} := \max\{IP : G(IP) \leq c^*\} \quad (3.27)$$

and the notation $s \wedge t$ means $\min\{s, t\}$.

The following remark is the reason why this lower bound is more tractable than the original cost structure in multi-echelon environments.

Remark 3.4.1 *The two functions defined above satisfy the following properties.*

(i) $H(\cdot)$ is convex and non-increasing;

(ii) $J(S) \geq 0$ for all S .

Part (i) of this remark follows by Theorem 3.3.4 and the convexity of $G(\cdot)$. Part (ii) follows from the optimality of (s^*, S^*) among (s, S) policies. The following theorem is the main result of this section.

Theorem 3.4.2 *Consider a facility with no successors in the distribution inventory system and that condition Σ is satisfied for this facility. Define the following two cost structures for the facility:*

[P] *Inventory costs accumulate at a rate $G(IP)$ and the cost of shipping up to S is equal to $K + \text{ord}(S)$;*

[B] *Inventory costs accumulate at a rate $H(IP)$ and the cost of shipping up to S is equal to $J(S)$.*

Then, for any non-anticipatory policy π , the facility's average cost over an infinite time horizon under [B] does not exceed the same cost under [P].

Proof

Assume that the inventory system is governed by the same policy and costs are accounted by the two cost structures. Hence shipping times and quantities are the same for each realization and we are comparing expected costs. As before, let the

filtrations \mathcal{F}_t and \mathcal{F}_t^o denote the set of events that are measurable on the history of the system and the facility's demand process, respectively, up to time t .

Since the costs incurred over a finite time interval do not affect the long-term average cost per period, let us assume with no loss of generality that the first shipment is sent to the facility at time $t = 0$ and that it raises its inventory position to some level $S = IP(0)$. It is sufficient to show that the expected cost imposed by [B] does not exceed that imposed by [P] between consecutive shipments to the facility. Hence we only consider the interval $[0, \tau 1(\omega))$, where $\tau 1(\omega)$ is the time of the next shipment.

If $S \leq s^*$ there is nothing to prove because [P] and [B] are the same. Thus we assume that $S > s^*$. Note that $\tau 1(\omega)$ is a stopping time on \mathcal{F} . The objective is to prove that

$$\begin{aligned} J(S) + E_{\omega \in \mathcal{F}} \int_0^{\tau 1(\omega)} H(IP(t; \omega)) dt \\ \leq K + \text{ord}(S) + E_{\omega \in \mathcal{F}} \int_0^{\tau 1(\omega)} G(IP(t; \omega)) dt \end{aligned}$$

or equivalently

$$E_{\omega \in \mathcal{F}} \int_0^{\tau 1(\omega)} [H(IP(t; \omega)) - G(IP(t; \omega))] dt \leq K + \text{ord}(S) - J(S). \quad (3.28)$$

If system reviews are periodic, integrals from a to b become summations from $t = a$ to $t = b - 1$ (see (3.15)).

Let $\tau^*(\omega) = \tau(S - s^*; \omega)$ denote the time until at least $S - s^*$ units are demanded, which is a stopping time for both \mathcal{F}^o and \mathcal{F} . The left hand side of (3.28) is equal to

$$E_{\omega \in \mathcal{F}} \int_0^{\tau 1(\omega)} [H(IP(t; \omega)) - G(IP(t; \omega))] dt$$

$$\begin{aligned}
&= E_{\omega \in \mathcal{F}} \int_0^{(\tau 1 \wedge \tau^*)(\omega)} [H(IP(t; \omega)) - G(IP(t; \omega))] dt \\
&\quad + E_{\omega \in \mathcal{F}} \int_{(\tau 1 \wedge \tau^*)(\omega)}^{\tau 1(\omega)} [H(IP(t; \omega)) - G(IP(t; \omega))] dt \\
&= E_{\omega \in \mathcal{F}} \int_0^{(\tau 1 \wedge \tau^*)(\omega)} [c^* - G(IP(t; \omega))] dt.
\end{aligned}$$

The first equation holds because $\tau 1 \wedge \tau^* \leq \tau 1$. The second one follows by the definition of $H(\cdot)$. Substituting into (3.28) we conclude that it is enough to prove

$$E_{\omega \in \mathcal{F}} \int_0^{(\tau 1 \wedge \tau^*)(\omega)} [c^* - G(IP(t; \omega))] dt \leq K + \text{ord}(S) - J(S). \quad (3.29)$$

If $S \geq S^* \wedge \bar{S}$, $J(S) = 0$ and (3.29) follows by Lemma 3.3.6 because $\tau \wedge \tau^*$ is a stopping time on \mathcal{F} . If $S \leq S^* \wedge \bar{S}$, we have

$$\begin{aligned}
E_{\omega \in \mathcal{F}} \int_0^{(\tau 1 \wedge \tau^*)(\omega)} [c^* - G(IP(t; \omega))] dt &\leq E \int_0^{\tau^*} [c^* - G(IP(t; \omega))] dt \\
&= c^* \cdot M(S - s^*) - I(s^*, S) \\
&= K + \text{ord}(S) - J(S),
\end{aligned}$$

where the inequality follows because $\tau 1 \wedge \tau^* \leq \tau^*$ and the integrand is positive for every sample path. The last equality follows from the definition of $J(S)$ in this range. \square

3.5 Lower Bounds for Serial Systems

3.5.1 Weakly Nested Policies

This section presents an extension of the concept of nested policies to stochastic systems with lead times. Nested policies have been used successfully by many authors for deterministic models, mostly in connection with 94%-effectiveness results for deterministic models. Some of these results are reviewed in Section 2.2.3 and

include Jackson, Maxwell and Muckstadt [46], Jackson and Roundy [47], Maxwell and Muckstadt [52], Mitchell [53], and Roundy [69] [68].

Nested policies are not as effective when demand is stochastic. Chen [15] proves that nested policies are optimal for a two-stage serial inventory system with zero lead time at the upper echelon. If $L_0 > 0$, however, nested policies are not optimal, as we argue later in this section.

The following definition provides an extension of nestedness to stochastic models with lead times.

Definition 3.5.1 : Weak Nestedness.

Consider any two facilities m, n in a distribution system such that n (the receiver) replenishes its inventory from m (the supplier). For any system history, let S and R denote the set of shipments made to the supplier and the receiver, respectively. Assume that stock at the supplier is managed on a first in, first out basis and define the function $\phi_{m,n}$ that maps S into R . If $\sigma \in S$, $\rho \in R$, and $\phi_{m,n}(\sigma) = \rho$ then ρ is the first receiver-shipment that contains stock from supplier-shipment σ .

A policy is weakly nested if it induces well defined, one-to-one functions $\phi_{m,n}(\cdot)$ for all system histories and for all arcs (m, n) .

According to this definition, each facility under a weakly nested policy must send part of each shipment it receives to each of its successors (a well defined $\phi_{m,n}(\cdot)$), and must not consolidate shipments (a one-to-one $\phi_{m,n}(\cdot)$). Note that stock that arrives at m in two different shipments, say $\sigma_1, \sigma_2 \in S$, can be sent to n in a single shipment $\rho_2 \in R$, provided that an earlier shipment $\rho_1 \in R$ from m to n contained some of the parts that arrived at m in shipment σ_1 . What is being precluded is to combine all of the units in σ_1 that will eventually reach n , and

some units that reached m in a subsequent shipment, in a single shipment from m to n .

Note that nested policies are weakly nested.

Weakly nested policies are not necessarily optimal. For example, consider a two-facility serial system consisting of a depot and a retailer and suppose that an extremely high demand occurs immediately after the depot places an order, triggering the placement of a second order (by the depot) ϵ time units after the first one. If at the time the first order arrives at the depot, the expected backorder costs to be incurred at the retailer during the next ϵ time units are outweighed by the fixed cost of sending a shipment to the retailer, it will be advantageous to wait for the second order and to consolidate both orders into one shipment to the retailer, hence saving the cost of sending a shipment at the expense of incurring higher backordering costs.

Despite its sub-optimality, weak nestedness appears to be a reasonable assumption in real situations. One reason is that it is common practice to expedite shipments when there is a large amount of backordered demand. Another reason is that the situation depicted in the example above should be highly unlikely, particularly if the demand variance is not extremely high and if the system policy is soundly designed.

3.5.2 Serial Systems

This section applies the results presented earlier in this chapter to find a lower bound for the cost of a serial inventory system.

We consider a serial system consisting of facilities labeled $1, \dots, N$ and satisfying all of the assumptions stated in section 3.2.1. External demand occurs at facility 1

Corollary 3.5.4 *If the definition of $\hat{J}_1(\cdot)$ in [C3] (Definition 3.5.2) is changed to $\hat{J}_1(x) = 0$, for all x , then Theorem 3.5.3 holds for any non-anticipatory policy π , not just for weakly nested policies.*

Proof

All the statements made in the proof of Theorem 3.5.3 still hold for arbitrary non-anticipatory policies, except for the comparison between the shipping costs incurred by [C3] and [C4]. In the corollary $\hat{J}_1(\cdot) \equiv 0$. Thus the shipping cost imposed by both structures is identically equal to 0 and the corollary follows. \square

Algorithm 3.5.5 System Decomposition Algorithm

1. **START.** Start with an N -facility serial inventory system as described in Section 3.2.1 and satisfying the assumptions stated in that section.
2. **FACILITY REDUCTION.** For $n = 1, \dots, N - 1$, do the following:

If condition Σ is satisfied for facility n let (s_n^*, S_n^*) denote the optimal policy for the single-stage subproblem induced by that facility. If condition Σ does not hold then the algorithm fails.

Apply Theorem 3.5.3 to facility n . Obtain functions $H_n(NI_{n+1}(t))$ and $\text{ord}_{n+1}(\cdot)$ as described in [C5]. Let

$$\hat{g}_{n+1}(NI_{n+1}(t)) = g_{n+1}(NI_{n+1}(t)) + H_n(NI_{n+1}(t));$$

and let the cost of sending a shipment to facility $n+1$ that raises its shipment-based echelon inventory position to a level $IP_{n+1}(t)$ be equal to $K_{n+1} + \text{ord}_{n+1}(IP_{n+1}(t))$.

Delete facility n from the system. For facility $n + 1$, define inventory-related costs using $\hat{g}_{n+1}(\cdot)$ in place of $g_{n+1}(\cdot)$.

Note that the intermediate systems resulting from each of the N iterations in these two algorithms satisfy all the assumptions stated in Section 3.2.1, and thus the algorithm terminates successfully as long as condition Σ holds at every iteration.

The following two bounds are the main results of this section. The first of these bounds is tighter than the second one and is valid only within the sub-class of weakly-nested policies. The second bound is not as tight but is valid for any non-anticipatory policy.

Theorem 3.5.6 *If Algorithm 3.5.5 terminates with a single-facility problem and an optimal policy can be found for such problem, then the long-term average cost of this policy is a lower bound for the long-term average cost of any weakly nested policy for the original system.*

Proof

This theorem follows immediately from Theorem 3.5.3. \square

Theorem 3.5.7 *Assume that:*

- (i) $g_n(x) \rightarrow \infty$ as $x \rightarrow \infty$, for all n .
- (ii) $\sum_{m=n}^N g_m(x) \rightarrow \infty$ as $x \rightarrow -\infty$, for all n .
- (iii) *In Step 2 of Algorithm 3.5.5, K_{n+1} is used for the shipment cost at facility $n + 1$, in place of $K_{n+1} + \text{ord}_{n+1}(IP_{n+1}(t))$.*

Then Algorithm 3.5.5 terminates with a single-facility problem and an optimal policy can be found for such problem. The long-term average cost of such policy is

a lower bound for the long-term average cost of any centralized, non-anticipatory policy for the original system.

Proof

The single-facility problem induced by facility n in Step 2 has a fixed shipping cost and the cost rate $\hat{g}_{n+1}(\cdot)$ is convex and satisfies $\hat{g}_{n+1}(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. It follows that an (s, S) policy is optimal for this problem and thus condition Σ is satisfied at each iteration of the algorithm. The rest follows immediately from Theorem 3.5.6. \square

Note that (i) and (ii) for Theorem 3.5.7 are satisfied for inventory cost rates defined according to (3.18) and (3.19), as long as $p > \sum_n h_n$.

Chapter 4

Optimality of (s, S) Policies for the Single-Stage Subproblem

This chapter studies the form of optimal policies for the single-stage subproblem, defined in section 3.3, and which differs from standard single-facility formulations since shipping costs are not necessarily fixed. The optimality of (s, S) policies for this subproblem, also called ‘condition Σ ’ in this dissertation, is fundamental for the lower bounds presented in chapter 3.

This chapter is organized as follows. Section 4.1 shows that (s, S) policies are optimal for the continuous-review version of the single-stage subproblem. Section 4.2 shows by means of an example that (s, S) policies are not necessarily optimal when system reviews occur periodically. Next, in section 4.2.1, sufficient conditions for the optimality of (s, S) policies for periodic-review systems are developed which depend only on the form of the renewal function and are valid for any nonnegative shipping cost function. The chapter concludes with section 4.2.3, which provides evidence that seems to indicate that condition Σ is more likely to hold as system

reviews are performed more frequently.

4.1 Optimality of (s, S) Policies for Continuous-Review Systems

We start the proof that (s, S) policies are optimal for the continuous review version of the single-stage subproblem by recalling some terminology from Section 3.3.2. The sigma-field generated by the demand process at a facility is denoted \mathcal{F}^o and its natural filtration is denoted \mathcal{F}_t^o . Let τ_1 be any stopping time on \mathcal{F}^o . Assume that a shipment is sent at time $t = 0$ which raises the facility's inventory position to some level $IP(0) = S$. Then the cost incurred during the time interval $[0, \tau_1)$ is equal to

$$R(\tau_1, S; \omega) = K + \text{ord}(S) + \int_0^{\tau_1(\omega)^-} G(S - D(t; \omega)) d\mu(t), \quad (4.1)$$

where μ , defined in (3.15), represents the Lebesgue measure when the system is reviewed continuously.

Lemma 4.1.1 *Assume that the inventory cost rate $G(\cdot)$ is convex and that $G(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. Let y^o be any point at which $G(\cdot)$ achieves its minimum and let*

$$c(S) = \inf \left\{ \frac{E_{\omega \in \mathcal{F}^o}[R(\tau, S; \omega)]}{E_{\omega \in \mathcal{F}^o}[\tau(\omega)]} : \tau \text{ a stopping time in } \mathcal{F}^o \right\}. \quad (4.2)$$

Find $s(S) \leq y^o$ satisfying $G(s) = c(S)$. Then the infimum in (4.2) is achieved by $\tau^(\omega) = \tau(S - s(S); \omega)$, the time elapsed until $S - s$ units are demanded.*

Proof

First, note that $\tau^*(\omega)$ is a solution to (4.2) if and only if it achieves the infimum in

$$\inf \{ E_{\omega \in \mathcal{F}^o}[R(\tau, S; \omega) - c(S)\tau(\omega)] : \tau \text{ a stopping time for } \mathcal{F}^o \} = 0. \quad (4.3)$$

Consider a sample path $\omega \in \mathcal{F}^o$ and note that $\tau^*(\omega) \geq \tau(S - y_0; \omega)$ for all ω . This occurs because the sample paths are continuous and the ratio $R(S, \tau; \omega)/\tau$ is decreasing as a function of $\tau(\omega)$ for $\tau(\omega) \leq \tau(S - y_0; \omega)$ for every sample path ω . Because of the convexity of $G(\cdot)$, the accumulation rate in (4.3), $G(S - D(t; \omega)) - c(S)$, increases monotonically and continuously as a function of t as the inventory position descends further from y_0 . Since $\tau^*(\omega)$ must also minimize this difference, the optimal stopping time along any given sample path is the time at which this accumulation rate becomes zero. It follows that $\tau^*(\omega) = \tau(S - s(S); \omega)$. \square

Lemma 4.1.2 *Assume that the inventory cost rate $G(\cdot)$ is convex and that $G(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. Let c^* denote the infimum of the long-term average cost incurred by the facility, taken among all feasible policies. Then*

$$c^* = \inf\{c(S) : S \in \mathfrak{R}\}.$$

Proof

By the renewal reward theorem, $c(S)$ is equal to the cost of the policy ‘ship up to S every time $\tau^*(\omega)$ is observed.’ It follows that $c^* \leq \inf\{c(S) : S \in \mathfrak{R}\}$.

Let π be an arbitrary policy and let c_π denote the long-term average cost incurred by policy π . To prove that $c^* \geq \inf\{c(S) : s \in \mathfrak{R}\}$ it is sufficient to show that $c_\pi \geq \inf\{c(S) : s \in \mathfrak{R}\}$ for all policies π . The last inequality is true because, by definition, the expected average cost per period incurred between any two successive shipments by any policy π is no less than $\inf\{c(S) : S \in \mathfrak{R}\}$. \square

The following theorem is the main result of this section.

Theorem 4.1.3 *If the system is reviewed continuously and the inventory cost rate $G(\cdot)$ is convex and such that $G(y) \rightarrow \infty$ as $|y| \rightarrow \infty$, then an (s^*, S^*) policy is optimal for the single-stage subproblem.*

Proof

Lemmas 4.1.1 and 4.1.2 imply that $c^* = \inf\{c(s, S) : s < S\}$ where $c(s, S)$ denotes the cost of an (s, S) policy, defined in (3.21). Theorem 3.3.4 (i) shows that the infimum is achieved. \square

4.2 Non-Optimality of (s, S) Policies for Periodic-Review Systems

We start this section by showing an instance of a single-stage subproblem with periodic review for which (s, S) policies are not optimal. The subproblem's formulation, cost expression, and optimal parameters are presented in Section 3.3.

Example 4.2.1 *Consider a single-stage subproblem with the following distribution for the one-period demand X_t : $\Pr\{X_t = 5\} = 0.8$ and $\Pr\{X_t = 1\} = 0.2$.*

The cost of sending a shipment that raises the inventory position to some value IP is equal to $1 + \text{ord}(IP)$, where $\text{ord}(IP) = 9$ if $IP \leq 9$ and $\text{ord}(IP) = 0$ otherwise. Stock arrives to the facility immediately upon shipment.

Inventory costs are incurred at the end of each review period, after the corresponding demand has occurred and the facility's stock has been depleted. These costs are incurred at a rate of one dollar per unit carried per period, and one thousand dollars per unit backordered per period.

We will find the best (s, S) policy for this example and then show that a better policy exists which is not an (s, S) policy. Some relevant data are summarized in Table 4.1.

To find the best (s, S) policy, we will start by finding an optimal reorder point s^* first. By Theorem 3.3.4, $s^* < y^o = 5$ (recall that y^o is a point at which $G(\cdot)$

Table 4.1: Data for example 4.2.1.

i	4	5	6	7	8	9	10	11	12
$K + \text{ord}(i)$	10	10	10	10	10	10	1	1	1
$G(i)$	800.6	0.8	1.8	2.8	3.8	4.8	5.8	6.8	7.8
$c(4, i)$		10.80	9.67	10.64	11.57	12.55	4.18	4.59	5.41
$h(i)$	0	-3.33	-3.00	-1.93	-0.72	0	-1.00	0.07	2.13

achieves its minimum). The same theorem states that $G(s^*) \geq c^* \geq G(s^* + 1)$, which implies that s^* equal to 4 because $G(4) = 800.6$, which is much greater than the optimal cost. By inspection over the costs of the policies $(4, S)$ (in the third row of the table) we conclude that the optimal (s, S) policy is $(s^*, S^*) = (4, 10)$, which incurs an average cost per period of $c(s^*, S^*) = 4.18$.

The question remains of whether there exists a better policy. An application of the successive approximation technique from dynamic programming (see for example Bertsekas [11], Section 8.2) results in the following policy: *ship up to 10 if the inventory position is less than or equal to 4 or equal to 9; otherwise do not ship*. The cost of this policy is equal to $c^* = 4.13$, which outperforms the best (s, S) policy. Furthermore, the optimality condition to be discussed later shows that this policy is optimal for this problem (the differential cost function $h(\cdot)$, to be defined in proposition 4.2.3, is included as the last row in table 4.1).

We conclude that condition Σ does not necessarily hold for the periodic-review version of the single-stage subproblem.

4.2.1 Sufficient Conditions for the Optimality of (s, S) Policies for Periodic Review Systems

The main result of this section is that the assumption below is sufficient to ensure that (s, S) policies are optimal for the periodic-review version of the single-stage subproblem (recall that Section 4.1 proves the optimality of (s, S) policies for the continuous review version of the same problem). The application of this assumption to inventory theory was first proposed by Sahin [72], who studied its implications on the behavior of cost functions and (s, S) policies for models with fixed shipping costs.

Assumption 4.2.2 *The renewal function $M(\cdot)$ is concave or, equivalently, the renewal density $m(\cdot)$ is non-increasing.*

In this chapter, concavity is assumed to be non-strict.

Our proof of sufficient conditions for the optimality of (s, S) policies for the single-stage subproblem relies on the so called optimality condition for dynamic programming. This condition is also the basis for all published proofs of the optimality of (s, S) policies for models with fixed shipping costs. Translated into the single-stage subproblem, this condition can be stated as follows (see for example Bertsekas [11] Proposition 1, p. 332).

Proposition 4.2.3 *Assume there exists a bounded function $h : \mathcal{Z} \rightarrow \mathbb{R}$, often referred to as the differential cost function, and a constant c^* , satisfying the optimality equation*

$$h(i) = \inf_{j \geq i} \left[\delta(j - i)[K + \text{Ord}(j)] + G(j) - c^* + \sum_{l=1}^{\infty} p_l h(j - l) \right] \quad \forall i \in \mathcal{Z} \quad (4.4)$$

where the function $\delta(x) = 0$ if $x = 0$ and $\delta(x) = 1$ otherwise.

Then if $\mu^*(i)$ achieves the infimum for all $i \in \mathcal{Z}$, the stationary policy “when at state i , ship up to $\mu^*(i)$ ” is optimal. Furthermore, the average cost per period under policy $\mu^*(\cdot)$ is equal to c^* for all initial states.

The differential cost function for the single-stage subproblem turns out to be identical regardless of whether shipping costs are fixed or non-fixed. An expression of this function in closed form was first proposed by Zheng [85] for the fixed shipping-cost case. Zheng first proves that the unbounded function

$$\hat{h}(i) = \begin{cases} 0, & i \leq s^*, \\ I(s^*, i) - c^*M(i - s^*), & i > s^*, \end{cases}$$

satisfies the optimality equation (4.4), where s^* is an optimal reorder point satisfying Theorem 3.1.6. Zheng also shows how a bounded differential cost function can be obtained by an appropriate relaxation of the problem, a technique to be discussed later.

The concavity Assumption 4.2.2 implies that $\hat{h}(\cdot)$ satisfies the following lemma.

Lemma 4.2.4 *If Assumption 4.2.2 holds, then $\hat{h}(i)$ is convex in the range $i > s^*$.*

Proof Let us start with the definition of $\hat{h}(\cdot)$ in the range $i > s^*$;

$$\begin{aligned} \hat{h}(i) &= I(s^*, i) - c^*M(i - s^*), & i > s^*, \\ &= \sum_{l=0}^{i-s^*-1} m(l)G(i-l) - c^* \sum_{l=0}^{i-s^*-1} m(l), & i > s^*; \end{aligned}$$

where the equality follows from (3.3) and (3.6).

Let $\Delta\hat{h}(i) = \hat{h}(i) - \hat{h}(i-1)$, $\Delta G(i) = G(i) - G(i-1)$, and let y_0 be a point at which $G(\cdot)$ achieves its minimum. We want to prove that $\Delta\hat{h}(\cdot)$ is increasing for $i > s^*$.

$$\Delta\hat{h}(i) = \sum_{l=0}^{i-s^*-2} m(l)\Delta G(i-l) + m(i-s^*-1)[G(s^*+1) - c^*]$$

$$\begin{aligned}
&= \sum_{l=0}^{i-y_0-1} m(l)\Delta G(i-l) + \sum_{l=i-y_0}^{i-s^*-2} m(l)\Delta G(i-l) \\
&\quad + m(i-s^*-1)[G(s^*+1) - c^*].
\end{aligned}$$

Since $G(\cdot)$ is convex, $\Delta G(i)$ is increasing, non-positive for $i \leq y_0$ and nonnegative for $i > y_0$. It follows that the second summation is increasing because it is a sum of non-positive numbers decreasing in absolute value. Similarly, the first summation is also increasing because when the value of i increases by 1, each of the existing (nonnegative) terms increases, and a new nonnegative term is added. Finally, the third term is increasing because $m(\cdot)$ is decreasing and $G(s^*+1) - c^*$ is a negative constant by Theorem 3.3.4, part (iii). \square

We are now ready to state the main result of this section.

Theorem 4.2.5 *If the concavity Assumption 4.2.2 holds, then an (s, S) policy is optimal for the single-stage subproblem.*

Proof

This proof, which closely follows a proof proposed by Zheng [85] for the fixed shipping-cost periodic-review model, consists of the following steps. First, a relaxation of the model is proposed; second, a differential cost function is presented whose associated policy behaves like an (s, S) policy in the long run; third, the optimality condition is verified.

We start by relaxing the model to allow the placement of negative shipments. More precisely, in the relaxed model the inventory position can be lowered from i to $j < i$ at a cost $K + \text{ord}(j)$. This relaxation amounts to deleting the constraint $j \geq i$ from the original model's optimality equation.

The differential cost function for the relaxed model is defined by

$$h(i) = \begin{cases} 0 & \text{if } i \leq s^* \\ I(s^*, i) - c^*M(i - s^*) & \text{if } s^* < i \leq S^o \\ 0 & \text{if } i > S^o \end{cases}$$

where $S^o = \max\{S : I(s^*, S) - c^*M(S - s^*) < 0\}$ and (s^*, S^*) are optimal parameters satisfying Theorem 3.3.4.

The following properties follow immediately from the definition of $h(\cdot)$.

$$K + \text{ord}(i) + h(i) \geq 0 \text{ for all } i, \text{ with equality binding for } i = S^*; \quad (4.5)$$

$$S^* \leq S^o. \quad (4.6)$$

The inequality in (4.5) follows from the minimality of $c(s^*, S^*)$ for i in the interval $s^* < i \leq S^o$, and because $h(\cdot) = 0$ and $\text{ord}(\cdot)$ is nonnegative for i outside that interval. Inequality (4.6) follows because $K + \text{ord}(S^*) + I(s^*, S^*) - c^*M(S^* - s^*) = 0$, which implies that $I(s^*, S^*) - c^*M(S^* - s^*) < 0$ because $\text{ord}(\cdot)$ is nonnegative and K is strictly positive. This last inequality implies that S^o exists and (4.6) holds. Finally, (4.5) is binding for $i = S^*$ because of (4.6) and the definition of $h(\cdot)$.

Intuitively, $-h(S)$ can be interpreted as the fixed shipping cost at which policy (s^*, S) would achieve cost c^* .

The policy corresponding to the differential cost function $h(\cdot)$ calls for shipping up to S^* if the inventory position is at or below s^* ; placing a negative shipment down to S^* if the inventory position is above S^o ; and not shipping otherwise. The average cost per period for this policy is the same as the cost of an (s^*, S^*) policy because a negative shipment is placed at most once (if the initial inventory position is greater than S^o) and the policy becomes indistinguishable from an (s^*, S^*) policy

afterwards. This difference between the two policies over a finite time period does not affect the average cost over an infinite horizon.

We are now in a position to verify the optimality condition. Since $h(\cdot)$ is clearly bounded, the proof reduces to the verification of the optimality equation (4.4), which consists of proving the following two inequalities:

$$h(i) \leq G(i) - c^* + \sum_{l=0}^{\infty} p_l h(i-l) \quad (4.7)$$

with the inequality binding for $s^* < i \leq S^o$, and

$$h(i) \leq \inf_j \{K + \text{ord}(j) + G(j) - c^* + \sum_{l=0}^{\infty} p_l h(j-l)\}, \quad (4.8)$$

with the inequality binding for $i \leq s^*$ and $i > S^o$. The first equation means that not to ship is better than to ship for $s^* < i \leq S^o$. The second equation means that shipping up to S^* is better than not shipping for $i \leq s^*$ and $i > S^o$.

Let us start with (4.7). The following ranges for i will be considered one by one: $i: i \leq s^*$, $s^* < i \leq S^o$, and $i > S^o$. In the first of these ranges, $i \leq s^*$, (4.7) reduces to $0 \leq G(i) - c^*$, which is true by Theorem 3.3.4 parts (ii) and (iii), and by the convexity of $G(\cdot)$.

For the second range, $s^* < i \leq S^o$, (4.7) reduces to

$$h(i) = G(i) - c^* + \sum_{l=0}^{\infty} p_l h(i-l),$$

which can be proved by substituting the renewal equations (3.5) and (3.7) into the definition of $h(\cdot)$ in this range, and noticing that $h(i) = 0$ for $i \leq s^*$.

In the last range, $i > S^o$, the optimality equation reduces to

$$G(i) - c^* + \sum_{l=0}^{\infty} p_l h(i-l) \geq 0, \quad (4.9)$$

whose left-hand side, in view of the renewal equations (3.5) and (3.7), is equal to

$$\begin{aligned}
& I(s^*, i) - c^* M(i - s^*) - \sum_{l=0}^{i-S^o-1} p_l [I(s^*, i-l) - c^* M(i - s^* - l)] \\
& \geq I(s^*, i) - c^* M(i - s^*) - \sum_{l=0}^{i-S^o-1} p_l [I(s^*, i) - c^* M(i - s^*)] \\
& = [I(s^*, i) - c^* M(i - s^*)] (1 - \sum_{l=0}^{i-S^o-1} p_l) \\
& \geq 0.
\end{aligned}$$

The first inequality is true because $I(s^*, i) - c^* M(i - s^*)$ is non-decreasing for $i > S^o$, which follows from the convexity of $I(s^*, i) - c^* M(i - s^*)$ and the fact that $S^o \geq y_0$. The second inequality is true because the term inside the square brackets is positive for $i > S^o$.

We conclude the proof by verifying (4.8). In view of (4.7), it is sufficient to show that

$$h(i) \leq K + \text{ord}(j) + h(j), \quad \text{for all } i, j;$$

which is true because the left hand side is non-positive by definition, and the right hand side is nonnegative by (4.5). Furthermore, the fact that (4.5) is binding at S^* and the definition of $h(\cdot)$ imply that (4.8) is binding for $i \leq s^*$ and $i > S^o$. This concludes the proof. \square

4.2.2 Concavity Assumption and the One-Period Demand

The purpose of this section is to identify two classes of one-period demand distributions that generate concave renewal functions.

Definition 4.2.6 *A nonnegative random variable X is said to have a Decreasing Failure Rate (DFR) distribution if $\Pr\{X = 0\} < 1$ and $\Pr\{X > x + y | X > y\}$ is increasing in $y \geq 0$ for all $x > 0$.*

It is known that a continuous random variable with density $f(\cdot)$ and c.d.f. $F(\cdot)$ has a DFR distribution if and only if its failure rate

$$r(x) = \frac{f(x)}{1 - F(x)}$$

is non-increasing (see for example Barlow and Proschan [7]).

A discrete random variable X with distribution $p_k = \Pr\{X = k\}$ is DFR if and only if its failure rate

$$r(k) = \frac{p_k}{\sum_{j=k}^{\infty} p_j}$$

is non-increasing (see for example Barlow and Proschan [6]).

Proposition 4.2.7 *The following one-period demand distributions generate concave renewal functions:*

- (i) *Any decreasing failure rate (DFR) distribution;*
- (ii) *A Bernoulli distribution.*

Proof

This proposition was first proved by Brown [14] for DFR distributions. Direct substitution shows that the renewal function is linear for Bernoulli distributed demands. \square

Example 4.2.8 *The following are examples of DFR distributions:*

1. *A negative binomial distribution with parameter $\alpha \geq 1$ (see Barlow and Proschan [6] p. 18).*
2. *A geometric distribution, which is a negative binomial with $\alpha = 1$, exhibits a constant failure rate.*

3. An exponential or mixed exponential distribution.
4. Gamma and Weibull distributions with rate parameter α less or equal to one.

4.2.3 Reducing the Review-Period Length

Recall that Theorem 4.2.5 states that (s, S) policies are optimal for the single-stage subproblem with continuous review. In view of this result, one could conjecture that (s, S) policies are more likely to be optimal for periodic review systems as system reviews are performed more frequently. This section presents instances in which reducing the length of the review period results in systems for which (s, S) policies are guaranteed to be optimal.

Definition 4.2.9 *Let P be a non-negative discrete distribution and n a natural number. We say that a distribution Q is the k th root of P if $P = Q^{*k}$, the convolution of Q with itself n times.*

Roots do not exist for all distributions. Common distributions for which roots exist include the binomial, negative binomial, and Poisson distributions.

Definition 4.2.10 *Consider a single-stage subproblem with review period length T_r , one-period demand distribution P and inventory cost rate $G(\cdot)$. Assume that P has an k -th root Q for some natural number k . The k th root of this subproblem is a new subproblem with review period length T_r/k , inventory cost rate $G(\cdot)/k$, one-period demand distribution Q , and otherwise identical to the original subproblem.*

An interpretation of this definition can be given by assuming that there exists an underlying demand process which is observed at intervals of length T_r for the original subproblem and at intervals of length T_r/k for the n th root of the subproblem.

Example 4.2.11 *The one-period demand for the subproblem is binomially distributed.*

A binomial distribution with parameters (n, p) can be expressed as the convolution of n Bernoulli distributions with parameter p . Theorem 4.2.5 and proposition 4.2.7 imply that (s, S) policies are optimal for the n th root of the subproblem.

Example 4.2.12 *The distribution of the one-period demand for the subproblem is a negative binomial.*

It is known that the convolution of two negative binomial distributions with parameters (p, α_1) and (p, α_2) is a negative binomial with parameters $(p, \alpha_1 + \alpha_2)$. Choose any natural number n such that $\alpha/n \leq 1$ and express the one-period demand distribution as the convolution of n negative binomial distributions with parameters $(p, \alpha/n)$. Then an (s, S) policy is optimal for the n th root of the subproblem as a result of Theorem 4.2.5, Proposition 4.2.7, and Example 4.2.8.

Finally, assume that the one-period demand is Poisson distributed. If the underlying demand process is continuous in time and has independent increments, then it follows that the demand process is a Poisson process. We know that as we make the review period smaller and smaller, the probability of two units being demanded during a review period tends to 0 and therefore the demand over a short time interval should look more or less like a Bernoulli trial distribution. It should thus be expected that the renewal function approaches concavity as the mean of the distribution tends to 0. This is in fact the case.

Chapter 5

93%-Effective Policies for a Two-Stage Inventory System

This chapter focuses on two-stage serial inventory systems with continuous review and governed by weakly nested policies. We show that if the penalty cost for incurring backorders is at least six times as large as the incremental holding cost at the lower echelon facility, then a 93%-effective policy can be found.

The chapter is divided into four sections. Section 5.1 discusses the sensitivity of the single-facility (r, Q) model to changes of the shipping quantity when the reorder point is kept at its optimal level. Section 5.2 formulates a two-stage serial inventory model and establishes a lower bound for the long-term average cost incurred by weakly nested policies for this model. Section 5.4 identifies 93%-effective policies.

5.1 A Sensitivity Property of the (r, Q) Model

The purpose of this section is to present a new sensitivity property of the (r, Q) inventory model reviewed in section 3.1.2.

Recall (3.11) states that the cost function for an (r, Q) model is equal to

$$c(r, Q) = \frac{\lambda K + \int_r^{r+Q} G(y) dy}{Q}. \quad (5.1)$$

Assume that inventory costs include holding and backordering costs, which are incurred at a rate proportional to the amount stocked or backordered, respectively.

Thus

$$g(x) = hx^+ + px^-,$$

where the parameters h and p are expressed in dollars per unit held or backordered per day.

Using (3.2), we can write the inventory cost rate as a function of the inventory position as

$$G(y) = E_D[h(y - D)^+ + p(y - D)^-]. \quad (5.2)$$

where D represents the demand that occurs over a lead time. The following remark states some elementary properties of $G(\cdot)$.

Remark 5.1.1 *If D has a density function, the cost rate $G(\cdot)$ satisfies the following properties:*

- (i) $G(\cdot)$ is differentiable almost everywhere, $G'(\cdot)$ is non-decreasing and $G'(y) \rightarrow h$ as $y \rightarrow \infty$;
- (ii) $G(\cdot)$ is convex;
- (iii) $G(y) \rightarrow \infty$ as $|y| \rightarrow \infty$.

Part (iii) is true because $G(y) \geq g(y - E[D])$ for all y , which follows from Jensen's inequality applied to (5.2) (Jensen's inequality can be found, for example, in Billingsley [13], page 283). To prove (i) and (ii) we apply the fundamental

theorem of calculus to $G(y)$, which can be written as $h(y - E[D]) + (h + p)E[(D - y)^+]$, and obtain $G'(y) = h - (h + p)(1 - F_D(y))$.

Remark 5.1.1 implies that $G(\cdot)$ satisfies the conditions in Theorem 3.1.8 and therefore $c(r^*, Q^*) = \min\{c(r, Q) : r \in \mathbb{R}, Q > 0\}$ if and only if

$$c(r^*, Q^*) = G(r^*) = G(r^* + Q^*). \quad (5.3)$$

Theorem 5.1.2 *Let $c^* = c(r^*, Q^*) = \inf\{c(r, Q) : r \in \mathbb{R}, Q > 0\}$. Then for any $\alpha > 0$*

$$\frac{c(r^*, \alpha Q^*) - c^*}{c^*} \leq \frac{1}{2} \frac{(1 - \alpha)^2}{\alpha} \cdot \frac{p + h}{p}. \quad (5.4)$$

Proof

This proof is an exercise in calculus which depends only on the form of the cost function (5.1). The first step of the proof is finding a bound for the numerator of the left hand side of (5.4). The second step is finding a bound for the denominator of the same expression. Dividing these two bounds results in (5.4) and concludes the proof.

We know by remark 5.1.1 (i) that $G'(y)$ is non-decreasing and tends to h as y tends to infinity. This implies that for any numbers a, b , $a \neq b$

$$\frac{G(a) - G(b)}{a - b} \leq h \quad (5.5)$$

regardless of whether a is greater than or less than b .

We now start the proof with the numerator of (5.4).

$$\begin{aligned} & c(r^*, \alpha Q^*) - c^* \\ &= \frac{\lambda K + \int_{r^*}^{r^* + Q^*} G(y) dy + \int_{r^* + Q^*}^{r^* + \alpha Q^*} G(y) dy}{\alpha Q^*} - \frac{\int_{r^*}^{r^* + Q^*} c^* dy + \int_{r^* + Q^*}^{r^* + \alpha Q^*} c^* dy}{\alpha Q^*} \\ &= \frac{\int_{r^* + Q^*}^{r^* + \alpha Q^*} (G(y) - c^*) dy}{\alpha Q^*} \end{aligned} \quad (5.6)$$

where (5.6) holds because $c^*Q^* = \lambda K + \int_{r^*}^{r^*+Q^*} G(y) dy$. Note that the additivity of integrals in the first equation holds regardless of whether $r^* + Q^* > r^* + \alpha Q^*$ or $r^* + Q^* \leq r^* + \alpha Q^*$.

The essential part of the proof is to use (5.5) to bound the integral in (5.6). More explicitly, we now prove that

$$\begin{aligned} 0 \leq \int_{r^*+Q^*}^{r^*+\alpha Q^*} (G(y) - c^*) dy &\leq \int_{r^*+Q^*}^{r^*+\alpha Q^*} h(y - (r^* + Q^*)) dy \\ &= \frac{1}{2}h(1 - \alpha)^2(Q^*)^2. \end{aligned} \quad (5.7)$$

The first inequality is a direct consequence of the convexity of $G(\cdot)$ and of (5.3).

If $\alpha > 1$, the second inequality follows from (5.5) with $b = r^* + Q^*$ and $a = y$.

If $\alpha < 1$, then $-[G(y) - c^*] \leq -[h(y - (r^* + Q^*))]$ for all y between the integration limits by (5.5). Because the lower integration limit $r^* + Q^*$ is greater than the upper integration limit $r^* + \alpha Q^*$, the second inequality holds.

Substitution of (5.7) into (5.6) yields the following bound for the numerator of (5.4):

$$c(r^*, \alpha Q^*) - c^* \leq \frac{1}{2\alpha}h(1 - \alpha)^2(Q^*). \quad (5.8)$$

Next we find a bound for the denominator of (5.4):

$$\begin{aligned} c^* &= G(r^*) = G(r^* + Q^*) \\ &\geq \max\{g(r^* - E[D]), g(r^* + Q^* - E[D])\} \\ &\geq \frac{hp}{h+p}Q^* > 0. \end{aligned} \quad (5.9)$$

The first equality holds by Theorem 3.1.8, the first inequality is an application of Jensen's inequality, and the second inequality follows as a result of known properties of the EOQ model with backorders (see for example, Zheng [87], Lemma 1, p. 6).

The sought result (5.4) is the ratio of (5.8) over (5.9), which concludes the proof. We close by pointing out that the inequality is binding when there are no lead times, in which case the cost function is identical to that of an EOQ model with backorders. \square

Corollary 5.1.3 *If the backordering cost parameter p is at least six times as large as the holding cost parameter h , the extra cost incurred by rounding off the shipping quantity to the nearest power of two, while keeping the reorder point fixed, is at most 7%.*

This corollary follows because rounding off the shipping quantity to the nearest power of two means that $\alpha \in [1/\sqrt{2}, \sqrt{2}]$. The 7% value is a worst-case performance for the bound of Theorem 5.1.2 for α in this range and for $p/h \geq 6$. In expected value the bound is tighter because α most often falls closer to one and because the bound improves for larger values of p/h . Practitioners typically use $p/h \geq 10$. For $p = h$ and $\alpha \in [1/\sqrt{2}, \sqrt{2}]$ the bound is at most 12%.

5.2 Two-Stage Model Formulation

This section considers a continuous review, two-stage serial inventory model with fixed shipping costs. The system consists of two facilities referred to as ‘the retailer’ and ‘the depot’, and denoted by the indices 1 and 2, respectively. External demand occurs at the retailer only. The retailer replenishes its inventory from the depot, which in turn is replenished from an external and unlimited supply.

We assume that the system satisfies the assumptions stated in section 3.2.1. These assumptions include: a demand process with stationary and independent increments and continuous and non-decreasing sample paths, continuous review,

backordering of unmet demand at the retailer, a fixed traveling time L_1 from the depot to the retailer, a fixed lead time L_2 for the shipments sent to the depot, and fixed shipping costs at each facility, which are denoted K_1 and K_2 . The inventory cost rate is an additive function of each facility's echelon net inventory and is the subject of the discussion below.

Notation and Inventory Cost Rate

Inventory in the system can be located either physically at or in transit to any facility. Let $INI_n(t)$ and $TI_n(t)$ ($n = 1, 2$) denote the installation net inventory and the shipments in transit to facility n at time t , respectively. Since backorders occur at the retailer only, all of these quantities are nonnegative except for $INI_1(t)$, which is unrestricted in sign.

A facility's echelon net inventory at time t , defined in section 3.2 and denoted $NI_n(t)$ ($n = 1, 2$), can be calculated by the expressions

$$NI_1(t) = INI_1(t); \quad NI_2(t) = INI_1(t) + IT_1(t) + INI_2(t). \quad (5.10)$$

Similarly, the shipment-based echelon inventory positions at time t , noted $IP_n(t)$ ($n = 1, 2$) and defined in section 3.2.2, can be calculated by the equations

$$IP_n(t) = NI_n(t) + TI_n(t), \quad n = 1, 2.$$

The retailer's shipment-based inventory position before shipping is denoted by

$$IP_1(t^-) = \lim_{\epsilon \rightarrow 0, \epsilon > 0} IP_1(t - \epsilon).$$

Inventory costs are incurred at each facility at a rate linear in the amount carried or backordered at the facility. Thus

$$\tilde{g}(INI_1, IT_1, INI_2, IT_2) = (h_1 + h_2)INI_1^+ + pINI_1^- + h_2IT_1 + h_2INI_2; \quad (5.11)$$

where t has been omitted for convenience. The parameter h_2 represents an *incremental* holding cost. There is only one backordering cost parameter p because demand is backordered only at the retailer. All of these parameters are expressed in dollars per unit carried or backordered per time unit and are assumed to be positive.

An algebraic manipulation of (5.11) and (5.10) shows that the inventory cost rate can be expressed as an additive function of each facility's echelon net inventory by

$$g(NI_1, NI_2) = h_1 NI_1^+ + (p + h_2) NI_1^- + h_2 NI_2 \quad (5.12)$$

where t has been omitted for convenience.

5.3 A Lower Bound for Weakly-Nested Policies

The purpose of this section is to develop a lower bound for the long-term average cost of any weakly nested policy for a two-stage serial inventory system.

Definition 5.3.1 *Let \mathcal{W} denote the class of centralized, non-anticipatory weakly nested policies.*

We start by applying Theorem 3.5.3 to the two-stage system. This results in the following cost structures.

[C0] The inventory cost rate is $g_1(NI_1(t)) = h_1[NI_1(t)]^+ + (p + h_2)[NI_1(t)]^-$.

The cost of sending a shipment is equal to K_1 .

[C1] The inventory cost rate is $G_1(IP_1(t)) = E_{D_1}[g_1(IP_1(t) - D_1)]$. The cost of sending a shipment is equal to K_1 .

[C2] The inventory cost rate is $H_1(IP_1(t))$ and the cost of sending a shipment that raises $IP_1(t)$ to S is equal to the function $J_1(S)$. The functions $H_1(\cdot)$ and $J_1(\cdot)$ are defined in (5.14) and (5.15).

[C3] is the same as [C2] because $J_1(\cdot)$ is non-increasing and therefore $\hat{J}_1(\cdot) = J_1(\cdot)$.

[C4] The inventory cost rate is equal to $H_1(NI_2(t))$ and the cost of receiving a shipment at facility 2 is equal to $J_1(S_2)$, where $S_2 = NI_2(t)$ represents the depot's echelon net inventory just after the time t when the shipment arrives to facility 2.

[C5] The inventory cost rate is equal to $H_1(NI_2(t))$ and the cost of sending a shipment to facility 2 is equal to $\text{ord}_2(S) = E_{D_2}[J_1(S - D_2)]$, where $S = IP_2(t)$ represents the depot's shipment-based echelon inventory position just after the time t when the shipment is sent to facility 2.

Recall that functions $H_1(\cdot)$ and $\text{ord}_2(\cdot)$ are based on the optimal policy (r_1^*, Q_1^*) , for the retailer's single-stage sub-problem. It is shown in Theorem 3.1.8 that the parameters of this policy are characterized by the equation

$$G_1(r_1^*) = G_1(r_1^* + Q_1^*) = \frac{\lambda K_1 + \int_{r_1^*}^{r_1^* + Q_1^*} G_1(y) dy}{Q_1^*}. \quad (5.13)$$

Functions $H_1(\cdot)$ and $\text{ord}_2(\cdot)$ are defined as

$$H_1(IP_1(t)) = \begin{cases} G_1(IP_1(t)) & \text{if } IP_1(t) \leq r_1^*, \\ c_1^* & \text{otherwise.} \end{cases} \quad (5.14)$$

$$J_1(S) = \begin{cases} K_1 & \text{if } S \leq r_1^*, \\ K_1 - \lambda^{-1} \int_{r_1^*}^S (c_1^* - G_1(y)) dy & \text{if } r_1^* < S \leq r_1^* + Q_1^*, \\ 0 & \text{otherwise;} \end{cases} \quad (5.15)$$

We continue our argument by applying Algorithm 3.5.5 to the two- stage serial system.

1. **START.** The two-facility serial inventory system defined in Section 5.2 satisfies the assumptions stated in Section 3.2.1.
2. **FACILITY REDUCTION.** The facility reduction step eliminates the retailer from the system. This technique can be applied because the retailer satisfies condition Σ , a fact proven by the results on the optimality of (r, Q) policies reviewed in Section 3.1.2. The resulting system consists of one facility—the depot—only. The inventory cost rate for this system is equal to

$$\hat{g}_2(NI_2(t)) = h_2NI_2(t) + H_1(NI_2(t)), \quad (5.16)$$

and the cost of sending a shipment that raises the depot's shipment-based echelon inventory position to level $IP_2(t)$ is equal to

$$K_2 + \text{ord}_2(IP_2(t)). \quad (5.17)$$

Definition 5.3.2 For $t > 0$, let $B(t)$ denote the total shipping and inventory costs incurred during the time interval $[0, t)$ under $[C5]$. For any $\pi \in \mathcal{W}$ let

$$\mathcal{B}(\pi) := \liminf_{t \rightarrow \infty} \frac{1}{t} B(t). \quad (5.18)$$

Theorem 3.5.3 implies that for all $\pi \in \mathcal{W}$

$$\mathcal{B}(\pi) \leq \mathcal{P}(\pi). \quad (5.19)$$

Remark 5.3.3 Consider the function

$$C_B(r_2, Q_2) := \frac{\lambda(K_2 + \text{ord}(r_2 + Q_2)) + \int_{r_2}^{r_2 + Q_2} E_{D_2} \hat{g}_2(y - D_2) dy}{Q_2}; \quad (5.20)$$

Then for any $\pi \in \mathcal{W}$ such that under π , the depot follows an (r_2, Q_2) policy to replenish its inventory

$$\mathcal{B}(\pi) = C_B(r_2, Q_2) \quad (5.21)$$

Transferring Shipping Costs

Let $0 \leq \Delta \leq K_1$. Consider a two-stage model identical to the one described in Section 5.2, except that the retailer's shipping cost is equal to $K_1 - \Delta$ instead of K_1 , and the depot's shipping cost is equal to $K_2 + \Delta$ instead of K_2 . Let $\mathcal{C}_\Delta(\pi)$ denote the long-term average cost of policy π for this new shipping cost parameters. Applying Theorem 3.5.3 to the new shipping cost parameters results in a lower bound for $\mathcal{C}(\pi)$, analogous to (5.20), which we will denote $\mathcal{B}_\Delta(\pi)$. Furthermore, if $\pi \in \mathcal{W}$ calls for an (r_2, Q_2) policy for the depot we define

$$C_B(r_2, Q_2; \Delta) := \mathcal{B}_\Delta(\pi). \quad (5.22)$$

Lemma 5.3.4 *For all $\pi \in \mathcal{W}$, $0 \leq \Delta \leq K_1$,*

$$\mathcal{C}_\Delta(\pi) \leq \mathcal{C}(\pi).$$

Proof

Recall from Definition 3.5.1 that for each realization of π , there exist a well defined, one-to-one function $\sigma(\cdot)$ from the set of depot shipments S to the set of retailer shipments R . The lemma is true because, for all realizations of π , the extra cost Δ incurred for each depot shipment σ as a result of the transfer in shipping costs, is offset by a cost saving of equal magnitude for each shipment $\phi^{-1}(\sigma) = \rho$ \square

Theorem 5.3.5 Minimax Property

Function $C_B(r_2, Q_2, \Delta)$, defined in (5.22), satisfies the minimax property. This means that there exist r_2^*, Q_2^*, Δ^* such that

$$\inf_{r_2, Q_2} \sup_{\delta} C_B(r_2, Q_2; \Delta) = \sup_{\delta} \inf_{r_2, Q_2} C_B(r_2, Q_2; \Delta) = C_B(r_2^*, Q_2^*; \Delta^*). \quad (5.23)$$

The value (r_2^*, Q_2^*, Δ^*) is often referred to as a saddlepoint.

Proof

A straightforward argument shows that $C_B(r_2, Q_2; \Delta)$ is continuous because it results of the composition of continuous functions.

For $\delta \in [0, K_1]$ let $(r_2^*(\delta), Q_2^*(\delta))$ be a solution to $\min_{r_2, Q_2} C_B(r_2, Q_2; \delta)$. Such solution exists because of Theorem 3.3.4. Furthermore, the continuity of $C_B(\cdot)$ implies that $Q_2^*(\delta)$ is a continuous function of δ .

For any $r_2 \in \mathfrak{R}$, $Q_2 > 0$, let $\delta^*(r_2, Q_2)$ be a solution to $\max_{0 \leq \delta \leq K_1} C_B(r_2, Q_2; \delta)$. This solution exists and is continuous on r_2, Q_2 because $C_B(\cdot)$ is continuous.

Now consider the function $\bar{\delta}(\delta) := \delta^*(r_2^*(\delta), Q_2^*(\delta))$. The arguments above show that $\bar{\delta}(\delta) : [0, K_1] \rightarrow [0, K_1]$ is continuous. It is a known result that any continuous function from a closed interval onto itself has a fixed point. Thus there exists a point $\delta^f \in [0, K_1]$ such that

$$\delta^f = \delta^*(r_2^*(\delta^f), Q_2^*(\delta^f)). \quad (5.24)$$

To conclude the proof, consider the following inequalities

$$\begin{aligned} \inf_{r_2, Q_2} \sup_{\delta} C_B(r_2, Q_2; \delta) &\leq \sup_{\delta} C_B(r_2^*(\delta^f), Q_2^*(\delta^f); \delta) \\ &= C_B(r_2^*(\delta^f), Q_2^*(\delta^f); \delta^f) \end{aligned} \quad (5.25)$$

$$\begin{aligned} &= \inf_{r_2, Q_2} C_B(r_2, Q_2; \delta^f) \quad (5.26) \\ &\leq \inf_{r_2, Q_2} \sup_{\delta} C_B(r_2, Q_2; \delta); \end{aligned}$$

where (5.25) follows from (5.24) and (5.26) follows from the definition of $r_2^*(\delta)$, $Q_2^*(\delta)$.

The proof is finished because it is known that for any function $f(x, y)$, $\inf_x \sup_y f(x, y) \geq \sup_y \inf_x f(x, y) \square$

Theorem 5.3.6 Lower Bound for Weakly Nested Policies

Let $G_1(IP_1(t)) = E_{D_1}[g_1(IP_1(t) - D_1)]$. For any $Q > 0$ the following quantities are uniquely defined

$$r_1(Q) \quad s.t. \quad G_1(r_1(Q)) = G_1(r_1(Q) + Q), \quad (5.27)$$

$$\Delta(Q) = K_1 - \frac{1}{\lambda} \int_{r_1(Q)}^{r_1(Q)+Q} [G_1(r_1) - G_1(y)] dy. \quad (5.28)$$

Furthermore,

$$c^* := \inf_{\pi \in \mathcal{W}} \mathcal{C}(\pi) \quad (5.29)$$

$$\geq \inf_{r_2, Q_2} C_B(r_2, Q_2; \Delta(Q_2)^+). \quad (5.30)$$

Proof

Quantities $r_1(Q)$ and $\Delta(Q)$ are uniquely defined because $G_1(\cdot)$ is strictly convex and $G(y) \rightarrow \infty$ as $|y| \rightarrow \infty$. We also define $\mathcal{Q} = \{\pi \in \mathcal{W} : \pi \text{ calls for an } (r_2, Q_2) \text{ policy for the retailer, } r_2 \in \mathfrak{R}, Q_2 > 0\}$. It follows that

$$c^* = \inf_{\pi \in \mathcal{W}} \mathcal{C}(\pi) \quad (5.31)$$

$$= \inf_{\pi \in \mathcal{W}} \sup_{0 \leq \delta \leq K_1} \mathcal{C}_\delta(\pi) \quad (5.32)$$

$$\geq \inf_{\pi \in \mathcal{W}} \sup_{0 \leq \delta \leq K_1} \mathcal{B}_\delta(\pi) \quad (5.33)$$

$$= \inf_{\pi \in \mathcal{Q}} \sup_{0 \leq \delta \leq K_1} \mathcal{B}_\delta(\pi) \quad (5.34)$$

$$= \inf_{r_2, Q_2} \sup_{0 \leq \delta \leq K_1} C_B(r_2, Q_2; \delta) \quad (5.35)$$

$$\geq \inf_{r_2, Q_2} \sup_{0 \leq \delta \leq K_1} C_B(r_2, Q_2; \Delta(Q_2)^+) \quad (5.36)$$

The infimum in (5.31) exists because $\mathcal{C}(\pi)$ is non-negative and well defined for all $\pi \in \mathcal{W}$. Expressions (5.32) and (5.33) follow from Lemma 5.3.4 and (5.19), respectively. Theorem 4.1.3 implies that for any $\pi \in \mathcal{W}$, $0 \leq \delta \leq K_1$, there exist $\pi^0 \in \mathcal{Q}$ such that $\mathcal{B}_\delta(\pi^0) \leq \mathcal{B}_\delta(\pi)$; since also $\mathcal{Q} \subset \mathcal{W}$, (5.34) follows. Equation (5.35) is true by (5.21). To conclude, note that $0 \leq \Delta(Q_2)^+ \leq K_1$ by (5.27), (5.28) and $G_1(\cdot)$ is convex; (5.35) follows \square

5.4 A 93%-Efficient Policy

This section shows a class of weakly nested policies, denoted $\mathcal{P}(r_2, Q_1, Q_2)$. We prove that $\mathcal{P}(r_2, Q_1, Q_2)$ always contains a 93%-effective policy.

Features of Policy $\mathcal{P}(r_2, Q_1, Q_2)$

Policy $\mathcal{P}(r_2, Q_1, Q_2)$ is based on the following principles:

1. The depot follows an (r_2, Q_2) policy based on its echelon inventory position.
2. The depot ships stock to the retailer on a first in, first out basis.
3. There is no mixing of stock from different depot-shipments into the same retailer-shipment. This implies that $\mathcal{P}(r_2, Q_1, Q_2)$ is weakly nested.
4. The retailer always places an order when $IP_1(t) = r_1(Q_1)$, defined in (5.27).

The amount shipped to the retailer varies according to the amount of stock available at the depot.

Assumption 5.4.1 *We will assume that $Q_2 \geq Q_1$*

Shipments

A depot-shipment is a shipment from an external supplier to the depot. A retailer-

shipment is a shipment from the depot to the retailer. We will use the following notation for depot-shipments.

1. σ denotes any depot-shipment;
2. $\sigma+$ denotes the next depot-shipment after shipment σ ;
3. $\sigma-$ denotes the last depot-shipment before shipment σ .

Time Epochs and Inventory Quantities Related to a Depot-Shipment

For each depot-shipment σ , we define the following items.

1. $t_a(\sigma)$ is the time at which σ arrives at the depot;
2. We define

$$NI_{2,\sigma}(t) := NI_2(t_a(\sigma)) - D(t_a(\sigma), t), \quad (5.37)$$

where $D(t_2, t_1) := -D(t_1, t_2)$ for $t_1 < t_2$. If $t \geq t_a(\sigma)$ then

$$NI_2(t) - NI_{2,\sigma}(t) \geq 0 \quad (5.38)$$

is the amount of inventory that has arrived at the depot in the time interval $(t_a(\sigma), t]$. Note that $NI_{2,\sigma}(t)$ is continuous and non-increasing, and that

$$NI_{2\sigma}(t) = NI_2(t), \quad \text{if } t_a(\sigma) \leq t < t_a(\sigma+). \quad (5.39)$$

3. $t_u(\sigma)$ is defined by

$$t_u(\sigma+) = \sup\{t : NI_{2,\sigma}(t) > r_1(Q_1)\}. \quad (5.40)$$

Clearly

$$NI_{2,\sigma}(t_u(\sigma+)) = r_1(Q_1). \quad (5.41)$$

$$4. t_c(\sigma) = \max\{t_a(\sigma), t_u(\sigma)\}.$$

Since the retailer places orders when $IP_1(t) = r_1(Q_1)$ and because $IP_1(t) \leq NI_2(t)$, one can think of $t_u(\sigma)$ as the latest time at which σ can arrive at the depot without causing unfilled orders from the retailer.

Policy $\mathcal{P}(r_2, Q_1, Q_2)$ operates in cycles with one cycle per depot-shipment. The cycle corresponding to a depot-shipment σ occurs during the time interval $[\tau_c(\sigma), \tau_c(\sigma+))$. No shipment to facility 1 contains stock from two different depot-shipments. We now present a formal description of Policy $\mathcal{P}(r_2, Q_1, Q_2)$.

Algorithm 5.4.2 : Policy $\mathcal{P}(r_2, Q_1, Q_2)$

Facility 2 Inventory Replenishment. Facility 2 follows an (r_2, Q_2) policy based on its echelon inventory position.

Facility 1 Ordering. Place an order for Q_1 units whenever the facility's shipment-based inventory position $IP_1(t)$ reaches $r_1(Q_1)$. Recall that $r_1(Q_1)$ is defined in (5.27).

Facility 2 Shipping. For each shipment σ incoming at the depot, execute the following steps:

1. *Compute the Ship-up-to Level for the current cycle.* At time $t_c(\sigma)$, calculate the ship-up-to level for the current cycle, denoted $S_c(\sigma)$.

Case 1. If $NI_{2,\sigma}(t_c(\sigma)) \leq r_1(Q_1) + Q_1$, set $k(\sigma) = 0$ and

$$S_c(\sigma) = NI_{2,\sigma}(t).$$

Case 2. If $NI_{2,\sigma}(t_c(\sigma)) > r_1(Q_1) + Q_1$, execute the following steps:

- (a) Compute $AA(\sigma)$, the amount available for shipping after fulfilling the demand backlogged at the beginning of the cycle.

$$AA(\sigma) = NI_{2,\sigma}(t_c(\sigma)) - r_1(Q_1).$$

- (b) Compute $SQ_1(\sigma)$, the shipping quantity for the current cycle, and an integer $k(\sigma) \geq 0$, such that

$$AA(\sigma) = 2^{k(\sigma)} \cdot SQ_1(\sigma); \quad 2^{-1/2} < \frac{Q_1}{SQ_1(\sigma)} \leq 2^{1/2}.$$

- (c) The ship-up-to level for the current cycle is equal to

$$S_c(\sigma) = r_1(Q_1) + SQ_1(\sigma).$$

2. *First Shipment.* Make a shipment to facility 1 at time $t_c(\sigma)$ so that

$$IP_1(t_c(\sigma)) = S_c(\sigma).$$

3. During the rest of the cycle corresponding to σ , ship $SQ_1(\sigma)$ units each time an order from facility 1 comes in. Keep sending these shipments until either the installation inventory at facility 2 is depleted, which will occur after $2^{k(\sigma)} - 1$ more shipments.

If the demand process is assumed to be deterministic and occur at a constant rate, then Policy $\mathcal{P}(r_2, Q_1, Q_2)$ resembles the one proposed by Mitchell [53] for distribution systems under such a demand process. In the random demand case, Policy $\mathcal{P}(r_2, Q_1, Q_2)$ is different from Mitchell's in that shipping quantities to facility 1 change in time to account for randomness in the amount available for shipment at stage 2.

Note that $\mathcal{P}(r_2, Q_1, Q_2)$ is weakly nested. Note also that under $\mathcal{P}(r_2, Q_1, Q_2)$ all retailer-shipments sent during the same cycle raise facility 1's shipment-based inventory position to the same level $S_c(\sigma)$, which may change from cycle to cycle.

Definition 5.4.3 *Note that when $Q_1 = Q_2$, all shipping quantities are equal to $Q_1 = Q_2$. Under such circumstance, retailer-shipments are sent immediately upon the arrival of an incoming depot-shipment or when facility 2 places an order, whichever happens later. We refer to this special case as a lot-by-lot policy.*

A one-to-one correspondence between depot-shipments and retailer-shipments can be established for each realization of a lot-by-lot policy because all depot shipments and retailer-shipments have equal quantities. The following claim follows.

Claim 5.1 *All lot-by-lot policies satisfy $\mathcal{C}(\mathcal{P}(r_2, Q, Q)) = \mathcal{C}_\delta(\mathcal{P}(r_2, Q, Q))$, for all $0 \leq \delta \leq K_1$.*

Lemma 5.4.4 *The following properties are satisfied for all depot-shipments σ :*

- (i) $t_u(\sigma) < t_u(\sigma+)$.
- (ii) $t_c(\sigma) < t_c(\sigma+)$.
- (iii) $NI_{2,\sigma}(t_c(\sigma)) \leq r_1(Q_1)$ if and only if $t_c(\sigma) \geq t_u(\sigma+)$. If $t_c(\sigma) \geq t_u(\sigma+)$ then $t_c(\sigma) = t_a(\sigma)$.
- (iv) Let $t \in [t_c(\sigma), t_c(\sigma+))$.
 - (iv-a) If $t < t_u(\sigma+)$ then $NI_2(t) \geq NI_{2,\sigma}(t) \geq IP_1(t) > r_1(Q_1)$.
 - (iv-b) If $t \geq t_u(\sigma+)$ then $IP_1(t) = NI_2(t) = NI_{2,\sigma}(t) \leq r_1(Q_1)$.
- (v) No shipments to facility 1 are made at time t if $t \in [t_u(\sigma+), t_c(\sigma+))$ and $t > t_c(\sigma)$.

(vi) If $t_u(\sigma+) > t_c(\sigma)$ then $IP_1(t_u(\sigma+)^-) = r_1(Q_1)$.

(vii) Assume $NI_{2,\sigma}(t_c(\sigma)) \leq r_1(Q_1) + Q_1$. Then either $NI_{2,\sigma}(t_a(\sigma)) = IP_1(t_c(\sigma))$

or $NI_{2,\sigma}(t_a(\sigma)) \geq IP_1(t_c(\sigma)) = r_1(Q_1) + Q_1$.

(viii) If $Q_1 = Q_2$ then $NI_{2,\sigma}(t_c(\sigma)) \leq r_1(Q_1) + Q_1$.

Proof

Since $NI_{2,\sigma}(t_a(\sigma)) = NI_{2,\sigma}(t_a(\sigma-)) - D(t_a(\sigma-), t_a(\sigma)) + Q_2$, it follows that

$$NI_{2,\sigma}(t) = NI_{2,\sigma-}(t) + Q_2, \quad \text{for all } t. \quad (5.42)$$

Thus $t_u(\sigma) = \sup\{t : NI_{2,\sigma-}(t) > r_1(Q_1)\} < \sup\{t : NI_{2,\sigma}(t) > r_1(Q_1)\} = t_u(\sigma+)$

by the continuity and monotonicity of $NI_{2,\sigma}(\cdot)$. This proves (i). Part (ii) follows immediately from (i) and $t_a(\sigma) < t_a(\sigma+)$.

The first statement in (iii) follows by (5.40). The second statement follows because, by (i), $t_c(\sigma) \geq t_u(\sigma+) > t_u(\sigma)$.

Note that

$$\begin{aligned} \text{Policy } \mathcal{P}(r_2, Q_1, Q_2) \text{ will not ship any material from depot-} \\ \text{shipment } \sigma+ \text{ to facility 1 before time } t_c(\sigma+). \end{aligned} \quad (5.43)$$

Therefore

$$\begin{aligned} IP_1(t) &\leq NI_{2,\sigma}(t_c(\sigma)) - D(t_c(\sigma), t) \\ &= NI_{2,\sigma}(t), \quad \text{if } t \in [t_c(\sigma), t_c(\sigma+)). \end{aligned} \quad (5.44)$$

Let $t \in [t_c(\sigma), t_c(\sigma+))$. There is enough inventory in shipment σ and in all earlier shipments to keep $IP_1(t)$ above $r_1(Q_1)$ at time t if and only if

$$\begin{aligned} NI_{2,\sigma}(t_c(\sigma)) &> r_1(Q_1) + D(t_c(\sigma), t) \\ \Leftrightarrow r_1(Q_1) &< NI_{2,\sigma}(t) \\ \Leftrightarrow t &< t_u(\sigma+), \end{aligned} \quad (5.45)$$

where the equivalences follow by (5.37) and (5.40). Policy $\mathcal{P}(r_2, Q_1, Q_2)$ is constructed to keep $IP_1(t)$ above $r_1(Q_1)$ as long as possible, i.e.,

$$IP_1(t) > r_1(Q_1) \quad \text{if} \quad t_c(\sigma) \leq t < t_u(\sigma+).$$

This expression, together with (5.38) and (5.44), prove (iv-a).

If $t_u(\sigma+) \leq t_c(\sigma)$ then by (iii) and Policy $\mathcal{P}(r_2, Q_1, Q_2)$, all of the inventory that arrived in shipment σ is dispatched to facility 1 at time $t_c(\sigma)$. If $t_u(\sigma+) > t_c(\sigma)$ then Policy $\mathcal{P}(r_2, Q_1, Q_2)$ ships all the inventory in σ to facility 1 before time $t_u(\sigma+)$. Therefore if $t_u(\sigma+) \vee t_c(\sigma) \leq t \leq t_c(\sigma+)$, then $IP_1(t) \geq NI_{2,\sigma}(t)$. By (5.43), (v) holds. By (5.44) and (5.45), if $t \in [t_c(\sigma) \vee t_u(\sigma+), t_c(\sigma+))$ then $IP_1(t) = NI_{2,\sigma}(t) \leq r_1(Q_1)$. By (5.39), $NI_2(t) = NI_{2,\sigma}(t)$, so (iv-b) holds.

Claim (vi) follows upon taking the limit as $t \uparrow t_u(\sigma+)$ in (iv-a), and noting that $NI_{2,\sigma}(t_u(\sigma+)) = r_1(Q_1)$.

Next we prove (vii). By Policy $\mathcal{P}(r_2, Q_1, Q_2)$ all of the inventory in σ is shipped at time $t_c(\sigma)$. By (5.43),

$$IP_1(t_c(\sigma)) = NI_{2,\sigma}(t_c(\sigma)). \tag{5.46}$$

By (5.42), (5.41) and Assumption (5.4.1),

$$\begin{aligned} NI_{2,\sigma}(t_u(\sigma)) &= NI_{2,\sigma-}(t_u(\sigma)) + Q_2 \\ &\geq r_1(Q_1) + Q_1 \\ &\geq NI_{2,\sigma}(t_c(\sigma)). \end{aligned}$$

Therefore either $NI_{2,\sigma}(t_u(\sigma)) > NI_{2,\sigma}(t_c(\sigma))$ or $NI_{2,\sigma}(t_u(\sigma)) = NI_{2,\sigma}(t_c(\sigma))$. In the former case the monotonicity of $NI_{2,\sigma}(\cdot)$ implies that $t_u(\sigma) < t_c(\sigma) = t_a(\sigma)$. By (5.39) and (5.46),

$$NI_2(t_a(\sigma)) = NI_{2,\sigma}(t_a(\sigma))$$

$$\begin{aligned}
&= NI_{2,\sigma}(t_c(\sigma)) \\
&= IP_1(t_c(\sigma)),
\end{aligned}$$

so (vii) holds. In the latter case $r_1(Q_1) + Q_1 = NI_{2,\sigma}(t_c(\sigma)) = IP_1(t_c(\sigma))$ by (5.46), and $NI_2(t_a(\sigma)) = NI_{2,\sigma}(t_a(\sigma)) \geq NI_{2,\sigma}(t_c(\sigma))$ by (5.39) and the monotonicity of $NI_{2,\sigma}(\cdot)$. Claim (vii) follows.

We conclude by proving (viii). If $Q_1 = Q_2$ then $NI_{2,\sigma}(t_c(\sigma)) = NI_{2,\sigma-}(t_c(\sigma)) + Q_2 \geq NI_{2,\sigma-}(t_u(\sigma)) + Q_1 = r_1(Q_1) + Q_1$, where the first equality follows by (5.42) and the last equality follows by (5.41). \square

The following properties related to the above cost structures are the basis for the effectiveness results to be presented later in this chapter.

Lemma 5.4.5 (i) *The total cost incurred in the interval $[t_c(\sigma) \vee t_u(\sigma+), t_c(\sigma+))$ is the same under [C1] and [C4].*

(ii) *Assume that*

$$t_u(\sigma+) > t_c(\sigma) \tag{5.47}$$

and that Case 1 of Step 1 of Algorithm 5.4.2 holds, that is, that $NI_{2,\sigma}(t_c(\sigma)) \leq r_1(Q_1) + Q_1$. Then the expected total cost incurred in $[t_c(\sigma), t_u(\sigma+))$ is the same under both [C1] and [C4].

(iii) *Assume that $(h_2 + p)/h_1 \geq 6$ and that Case 2 of Step 1 of Algorithm 5.4.2 holds, that is, that $NI_{2,\sigma}(t_c(\sigma)) > r_1(Q_1) + Q_1$. Then the expected costs incurred during the interval $[t_c(\sigma), t_u(\sigma+))$ under [C1] do not exceed the expected costs incurred during the same interval under [C4] by more than 7%.*

Proof

To prove (i), note that the inventory-related costs are equal because of Lemma 5.4.4

(iv-b) and (5.14). If $t_c(\sigma) < t_u(\sigma)$ then no shipping costs are incurred under [C4]. By Lemma 5.4.4 (v), no shipping costs are incurred under [C1]. If $t_c(\sigma) \geq t_u(\sigma+)$ then by Lemma 5.4.4 (v) and Policy $\mathcal{P}(r_2, Q_1, Q_2)$, we incur a cost of K_1 for making a single shipment at time $t_c(\sigma)$ under [C1]. By Lemma 5.4.4 (i), $t_u(\sigma) < t_u(\sigma+) \leq t_c(\sigma)$, so $t_c(\sigma) = t_u(\sigma)$. By Lemma 5.4.4 (iv-b), $r_1(Q_1) \geq NI_2(t_c(\sigma)) = NI_2(t_u(\sigma))$. Under [C4] we incur a shipping cost of $J(NI_2(t_u(\sigma)))$, which is equal to K_1 by (5.15). This proves (i).

To prove (ii), note that Case 1 in Step 2 of Algorithm 5.4.2 applies. By the definition of $\mathcal{P}(r_2, Q_1, Q_2)$ a retailer-shipment is sent at time $t_c(\sigma)$ and no other shipments are sent during the interval $[t_c(\sigma), t_u(\sigma+))$. Thus $IP_1(t)$ is monotonic and continuous for $t \in [t_c(\sigma), t_u(\sigma+))$. By (5.47) and Lemma 5.4.4 (vi), $IP_1(t_u(\sigma+)^-) = r_1(Q_1)$. Lemma 3.1.7 implies that the expected total cost incurred under [C1] during the interval $[t_c(\sigma), t_u(\sigma+))$ is equal to

$$K_1 + \lambda^{-1} \int_{r_1(Q_1)}^{IP_1(t_c(\sigma))} G_1(y) dy. \quad (5.48)$$

Setting $G(y) \equiv c_1^*$ in (3.10), Lemma 3.1.7 also implies that

$$c_1^* \cdot E[t_u(\sigma+) - t_c(\sigma)] = \lambda^{-1} \int_{r_1(Q_1)}^{IP_1(t_c(\sigma))} c_1^* dy. \quad (5.49)$$

By Lemma 5.4.4 (iv-a), $H_1(NI_2(t)) = c_1^*$ for $t \in [t_c(\sigma), t_u(\sigma+))$. By Lemma 5.4.4 (vii), (5.15) and (5.49), the expected total cost incurred under [C4] during the interval $[t_c(\sigma+), t_u(\sigma))$ is equal to

$$J_1(NI_2(t_u(\sigma))) + c_1^* \cdot E[t_u(\sigma+) - t_c(\sigma)] = J_1(IP_1(t_c(\sigma))) + \lambda^{-1} \int_{r_1(Q_1)}^{IP_1(t_c(\sigma))} c_1^* dy. \quad (5.50)$$

Note that by Lemma 5.4.4 (iv-a), $r_1(Q_1) < IP_1(t_c(\sigma)) \leq NI_{2,\sigma}(t_c(\sigma)) \leq r_1(Q_1) + Q_1$. The expected costs (5.48) and (5.50) are equal by (5.15), so (ii) holds.

To prove (iii), note that Case 2 of Step 2 in Algorithm 5.4.2 applies. By (5.39) and the monotonicity of $NI_{2,\sigma}(\cdot)$, $NI_2(t_a(\sigma)) = NI_{2,\sigma}(t_a(\sigma)) \geq NI_{2,\sigma}(t_c(\sigma)) > r_1(Q_1) + Q_1$, so no ordering costs are incurred for the current cycle under [C4].

Let l denote the number of shipments sent to facility 1 during the interval $[t_c(\sigma), t_u(\sigma))$. Policy $\mathcal{P}(r_2, Q_1, Q_2)$ and Lemma 3.1.7 imply that the expected cost incurred under [C1] during this interval is equal to

$$l[K_1 + \lambda^{-1} \int_{r_1(Q_1)}^{r_1(Q_1) + SQ_1} G_1(y) dy].$$

By Lemma 5.4.4 (iv-a), $H_1(NI_2(t)) = c_1^*$ for all $t \in [t_c(\sigma), t_u(\sigma+))$. By (B.5), under [C4] the expected value of the same cost is equal to $l \cdot \lambda^{-1} \cdot c_1^* \cdot SQ_1$. These two expected costs differ by at most 7% as a result of Corollary 5.1.3 and the definition of $SQ_1(\sigma)$ \square

Theorem 5.4.6 *Policy $\mathcal{P}(r_2, Q_1, Q_2)$ satisfies the following properties.*

(i) *If $(h_2 + p)/h_1 \geq 6$ then*

$$\mathcal{C}(\mathcal{P}(r_2, Q_1^*, Q_2)) \leq 1.07 \cdot C_B(r_2, Q_2; 0), \quad \text{for all } r_2, Q_2 > Q_1^*$$

(ii) $\mathcal{C}(\mathcal{P}(r_2, Q_1^*, Q_1^*)) = C_B(r_2, Q_1^*; 0)$, *for all r_2 .*

(iii) $\mathcal{C}(\mathcal{P}(r_2, Q_2, Q_2)) = C_B(r_2, Q_2; \Delta(Q_2))$, *for all $r_2, Q_2 < Q_1^*$.*

Proof

First recall that for all (r_2, Q_1, Q_2) , $\mathcal{C}(\mathcal{P}(r_2, Q_1, Q_2))$ results from accounting costs [C0] to each realization of $\mathcal{P}(r_2, Q_1, Q_2)$ and $\mathcal{B}(\mathcal{P}(r_2, Q_1, Q_2))$ results from accounting costs [C5] to each realization of $\mathcal{P}(r_2, Q_1, Q_2)$. Furthermore, [C0] and [C1] are equivalent in expected cost, and so are [C4] and [C5]. Thus it is sufficient to compare the expected cost incurred by $\mathcal{P}(r_2, Q_1, Q_2)$ under [C4] to the long-term

average cost incurred by $\mathcal{P}(r_2, Q_1, Q_2)$ under [C1] . By Remark 5.4.4 (ii) it is sufficient to consider the expected cost during the interval $[t_c(\sigma), t_c(\sigma+))$ for any σ . For expositional purposes, in this proof we alter [C4] by assuming that the cost $J_1(NI_2(t_a(\sigma)))$ associated with the receipt of shipment σ is incurred at time $t_c(\sigma)$.

First we prove (i). By Lemma 5.4.5(i), the proof of the theorem is complete if $t_u(\sigma+) \leq t_c(\sigma)$. Thus it is sufficient to assume that $t_u(\sigma+) > t_c(\sigma)$ and consider total expected cost incurred during time interval $[t_c(\sigma), t_u(\sigma+))$ under [C1] and [C4]. Consider each of the two cases in Step 1 of Algorithm 5.4.2. If Case 1 holds, the expected costs under [C1] and [C4] are equal by Lemma 5.4.5 (ii). If Case 2 holds, the expected cost under [C1] does not exceed the expected cost under [C4] by more than 7% by Lemma 5.4.5 (iii). This proves (i).

If $Q_1^* = Q_2$, the argument used to prove (i) is still valid, and besides Case 2 never occurs because of Lemma 5.4.4 (viii). This proves (ii).

To prove (iii), note that (5.27) and (5.28) imply that

$$G_1(r_1(Q_2)) = G_1(r_1(Q_2) + Q_2) = \frac{\lambda(K_1 - \Delta(Q_2)) + \int_{r_1(Q_2)}^{r_1(Q_2) + Q_2} G_1(y) dy}{Q_2}; \quad (5.51)$$

note also that $0 < \Delta(Q_2) < K_1$ because $Q_2 < Q_1^*$.

Part (iii) follows because

$$\begin{aligned} \mathcal{C}(\mathcal{P}(r_2, Q_2, Q_2)) &= \mathcal{C}_{\Delta(Q_2)}(\mathcal{P}(r_2, Q_2, Q_2)) \\ &= \mathcal{B}(\mathcal{P}(r_2, Q_2, \Delta(Q_2))) \end{aligned}$$

The first equation is true by Claim 5.1. To prove the second equation, consider the transfer of $\Delta(Q_1)$ shipping cost units from the retailer to the depot. Theorem 3.1.8 and (5.51) imply that Q_2 is the optimal lot size for the single-stage subproblem

induced by the retailer after this transfer of shipping costs. Thus (5.52) follows by

(ii) \square

Claim 5.2 (i) $\Delta(Q_2) > 0 \Leftrightarrow Q_2 < Q_1^*$

(ii) For all r_2 , $Q_2 \leq Q_1^*$, $0 \leq \delta \leq K_1$

$$C_B(r_2, Q_2; \delta) \leq C_B(r_2, Q_2, \Delta(Q_2))$$

Proof

Part (i) follows from definitions (5.27) and (5.28) and because $G_1(\cdot)$ is strictly convex. To prove (ii), note that

$$\begin{aligned} C_B(r_2, Q_2; \delta) &= \mathcal{B}_\delta(\mathcal{P}(r_2, Q_2, Q_2)) \\ &\leq \mathcal{C}_\delta(\mathcal{P}(r_2, Q_2, Q_2)) \\ &\leq \mathcal{C}(\mathcal{P}(r_2, Q_2, Q_2)) \\ &= C_B(r_2, Q_2; \Delta(Q_2)) \end{aligned}$$

The first equation is true by definition. The next two inequalities follow from (5.19) and Lemma 5.3.4, respectively. The last equality holds by Theorem 5.4.6(ii) and (iii). Part (ii) follows \square

Theorem 5.4.7 Optimal Parameters

(i) $C_B(r_2, Q_2; 0)$ achieves its minimum.

(ii) $C_B(r_2, Q_2; \Delta(Q_2)^+)$ achieves its minimum.

(iii) If there exists a solution (r_2^0, Q_2^0) to $\min_{r_2, Q_2} C_B(r_2, Q_2; 0)$ such that $Q_2^0 \geq Q_1^*$ then $(r_2^0, Q_2^0; 0)$ is a solution to $\min_{r_2, Q_2} C_B(r_2, Q_2, \Delta(Q_2)^+)$.

Proof

Theorem 3.3.4 implies (i). Next, note that $C_B(r_2, Q_2, \delta)$ is non-negative and continuous, and that for all (r_2, δ) , $C_B(r_2, Q_2, \delta) \rightarrow \infty$ as $Q_2 \rightarrow 0$. Recall that $\hat{g}_2(NI_2)$ is nonnegative and tends to infinity as $|NI_2| \rightarrow \infty$. Thus for all δ $C_B(r_2, Q_2; \delta) \rightarrow \infty$ as $\max(|r_2|, Q_2) \rightarrow \infty$. Consequently $\{(r_2, Q_2) : C_B(r_2, Q_2, \Delta(Q_2)) \leq C_B(0, 1; 0)\}$ is compact, and (ii) holds.

Next, note that for all r_2, Q_2

$$C_B(r_2^0, Q_2^0; 0) \leq C_B(r_2, Q_2; 0) \leq C_B(r_2, Q_2; \Delta(Q_2)^+). \quad (5.52)$$

The first inequality in (5.52) follows from the optimality of (r_2^0, Q_2^0) . The second inequality follows from Claim 5.2(ii) for $Q_2 \leq Q_1^*$ and from Claim 5.2(i) for $Q_2 > Q_1^*$. Part (iii) follows immediately \square

The following corollary is the main result of this chapter.

Corollary 5.4.8 *If $(h_2 + p)/h_1 \geq 6$, there exist parameters $(r_2^{**}, Q_1^{**}, Q_2^{**})$ such that $\mathcal{P}(r_2^{**}, Q_1^{**}, Q_2^{**})$ is 93%-effective within \mathcal{W} , the class of weakly-nested policies.*

Proof

Theorems 5.4.7(ii) and 5.3.6 imply that there exist parameters r_2^*, Q_2^* such that $\inf_{\pi \in \mathcal{W}} \mathcal{C}(\pi) \geq C_B(r_2^*, Q_2^*; \Delta(Q_2^*)^+)$. If $Q_2^* > Q_1^*$, then $\mathcal{P}(r_2^*, Q_1^*, Q_2^*)$ is 93%-effective by Theorem 5.4.6(i); otherwise $\mathcal{P}(r_2^*, Q_2^*, Q_2^*)$ is optimal by Theorem 5.4.6(ii) and (iii) \square

Chapter 6

Conclusions

The 93%-effective policy for developed in chapter 5 substantiates the following recommendations for managers of distribution inventory systems.

1. Make replenishment decisions based on shipment-based echelon inventories.

This recommendation has the following implications for the management of a depot:

- (a) Shipments to replenish the depot should be triggered by external demand, not by orders placed by the retailer.
- (b) Replenishment of a depot should take into account the stock located downstream from the depot. An empty depot does not necessarily need to be replenished if enough stock has been shipped to the facilities that replenish their stock from the depot.

2. Keep reordering levels fixed and adjust shipping quantities, if necessary, to equally distribute fixed shipping costs.

Extending the validity of these recommendations for more complex systems is an interesting research topic. Policies based on the same principles could be evaluated for any distribution system by comparing their cost, obtained via simulation, to the lower bounds developed in chapter 3.

Some features of this dissertation are of theoretical interest and help to better understand multi-echelon inventory systems. These features include the decomposition approach and the imputed costs for the lower bound developed in Chapter 3. For the 93%-efficiency result developed in Chapter 5, interesting features include the concept of weak nestedness and the use of heuristic policies that keep reordering points fixed while adjusting shipping quantities.

The following problems are not well understood today and would be important areas for future research: the development of algorithms to compute efficient policies, and extensions to multi-item models.

Appendix A

Accounting for Holding and Backordering Costs

The purpose of this appendix is to show the relationship between the parameter p_n used in (3.17), and the basic holding and backordering cost parameters, which are not based in echelon net inventories.

For all facilities n , we define the following:

1. The installation net inventory at facility n is equal to the stock-on hand at the facility minus the amount backordered at the facility. This quantity is restricted to be non-negative for facilities with no successors because backordering of demand does not occur at those facilities. The notation $INI_n(t)$ denotes the installation net inventory of facility n observed after dispatching and delivery of shipments at time t .
2. The inventory in transit to a facility is defined as the stock already dispatched from the facility's supplier that has not yet arrived to the facility. The

notation $IT_n(t)$ denotes the inventory in transit to facility n after dispatching and delivery of shipments at time t .

3. The set consisting of facility n and all its descendants is defined recursively as follows:

$$\sigma^*(n) = \{n\} \cup [\cup_{(n,m) \in A} \sigma^*(m)].$$

Similarly,

$$\pi^*(n) = \{n\} \cup [\cup_{(n,k) \in A} \pi^*(k)],$$

is the set composed of n and all its direct and indirect predecessors.

4. The echelon net inventory for facility n (see Definition 3.2.6) is equal to

$$NI_n(t) = \sum_{m \in \sigma^*(n)} INI_m(t) + \sum_{k \in \sigma^*(n) \setminus n} IT_k(t). \quad (\text{A.1})$$

5. The *conventional holding cost parameter* for facility n is defined as

$$h_n^c = \sum_{m \in \pi^*(n)} h_m. \quad (\text{A.2})$$

We also define $\mathcal{S} = \{m : \sigma(m) \neq \emptyset\}$ to be the set of all facilities with successors, and $\overline{\mathcal{S}} = \{m : \sigma(m) = \emptyset\}$ to be the set of all facilities without successors.

Claim A.1 *The following results follow from these definitions:*

- (i) $n \in \sigma^*(m)$ if and only if $m \in \pi^*(n)$;
- (ii) If $n \in \mathcal{S}$, then $INI_n \geq 0$;
- (iii) If $n \in \overline{\mathcal{S}}$, then $NI_n = INI_n$.

We suppress the dependence of inventory levels on time for notational convenience. Holding and backordering costs are conventionally accounted using the cost rate

$$\begin{aligned}
& \sum_n h_n^c [IN I_n^+ + \sum_{(n,k) \in A} IT_k] + \sum_{n \in \bar{S}} \hat{p}_n IN I_n^- \\
&= \sum_n \sum_{m \in \pi^*(n)} h_m IN I_n + \sum_{m \in \pi^*(n), (n,k) \in A} h_m IT_k + \sum_{n \in \bar{S}} (h_n^c + \hat{p}_n) IN I_n^- \\
&= \sum_m h_m [\sum_{n \in \sigma^*(m)} IN I_n + \sum_{n \in \sigma^*(m), (n,k) \in A} IT_k] + \sum_{n \in \bar{S}} (h_n^c + \hat{p}_n) IN I_n^- \\
&= \sum_m h_m NI_m + \sum_{n \in \bar{S}} (h_n^c + \hat{p}_n) NI_n^-.
\end{aligned}$$

The first equation follows from the definition of h_n^c and (ii). The second one is based on (i). The third one is a result of (A.1).

It follows that the backorder cost in (3.17) is related to the basic holding and backordering cost parameters by the equation:

$$p_n = \hat{p}_n + h_n^c. \quad (\text{A.3})$$

Appendix B

Proof of Lemma 3.1.7.

Let $\tau(x; \omega) = \inf\{t \geq 0 : D(t) \geq x\}$ if $x > 0$, and $\tau(x; \omega) = 0$ otherwise. Then $\tau(x; \omega)$ is a stopping time for $\mathcal{F}^o = \sigma(\{D(t) : t \geq 0\})$ for all x . As in (B.5), for any $x > 0$, let

$$M(x) = E_{\omega \in \mathcal{F}^o}[\tau(x; \omega)],$$

and note that $M(\cdot)$ is non-decreasing in x .

Take any $x, y > 0$. Because of the continuity and the monotonicity of the sample paths, the demand process visits x exactly once before reaching $x + y$. Taking expected values and applying the strong Markov property we conclude that

$$M(x + y) = M(x) + M(y), \quad \text{for all } x, y > 0.$$

It is well known that the only non-decreasing function that satisfies this equation is

$$M(x) = c \cdot x, \tag{B.1}$$

for some c , $0 \leq c \leq \infty$.

The renewal process obtained from observing $D(t)$ only at discrete points in time will be useful in finding the value of c . Let $D'(t) = \{D(t) : t = 0, 1, 2, \dots\}$, $\tau'(x; \omega) = \inf\{t = 0, 1, 2, \dots : D'(t) \geq x\}$, and $M'(x) = E[\tau'(x; \omega)]$. It follows that for all sample paths ω : $\tau'(x; \omega) - 1 < \tau(x; \omega) \leq \tau'(x; \omega)$, which implies that

$$M'(x) - 1 < M(x) \leq M'(x). \quad (\text{B.2})$$

Since $M'(x) < \infty$ by (3.3) and (3.4), it follows that

$$M(x) < \infty \text{ for all } x < \infty. \quad (\text{B.3})$$

Next we prove that $c = 1/\lambda$. This is true because

$$c = \lim_{x \rightarrow \infty} \frac{M(x)}{x} = \lim_{x \rightarrow \infty} \frac{M'(x)}{x} = \frac{1}{\lambda}. \quad (\text{B.4})$$

The first equality holds by (B.1) and (B.3). The second inequality holds by (B.2). The third one is a consequence of the elementary renewal theorem. By (B.1) and (B.4), we conclude that

$$M(x) = \frac{1}{\lambda} x. \quad (\text{B.5})$$

The proof will be complete if we prove that

$$E_{\omega \in \mathcal{F}^0} \left[\int_0^{\tau(S-r; \omega)} G(S - D(t; \omega)) dt \right] = \frac{1}{\lambda} \int_r^S G(y) dy. \quad (\text{B.6})$$

We will prove (B.6) for step functions first, i.e., for functions of the form:

$$G^n(x) = \sum_{k=1}^N c_k \chi_{[x_i \geq x > x_{i+1}]}$$

for some $S = x_0 > x_1 > \dots > x_N = r$, where $\chi_{[\cdot]}$ is the indicator function.

The continuity and monotonicity of the sample paths implies that $\tau(S - x_0; \omega) = 0 < \tau(S - x_1; \omega) < \dots < \tau(S - x_{N-1}; \omega) < \tau(S - x_N; \omega) = \tau(S - r; \omega)$ for all $\omega \in \mathcal{F}^o$. Thus

$$\begin{aligned} E_{\omega \in \mathcal{F}^o} & \left[\int_0^{\tau(S-r; \omega)} G^n(S - D(t; \omega)) dt \right] \\ &= E_{\omega \in \mathcal{F}^o} \left[\sum_{k=1}^N \int_{\tau(S-x_{k-1}; \omega)}^{\tau(S-x_k; \omega)} G^n(S - D(t; \omega)) dt \right] \end{aligned} \quad (\text{B.7})$$

$$\begin{aligned} &= \sum_{k=1}^N E_{\omega \in \mathcal{F}^o} \left[\int_{\tau(S-x_{k-1}; \omega)}^{\tau(S-x_k; \omega)} c_k dt \right] \\ &= \sum_{k=1}^N c_k E_{\omega \in \mathcal{F}^o} [\tau(S - x_k; \omega) - \tau(S - x_{k-1}; \omega)] \\ &= \sum_{k=1}^N c_k \frac{(x_{k-1} - x_k)}{\lambda} \end{aligned} \quad (\text{B.8})$$

$$= \frac{1}{\lambda} \int_r^S G^n(y) dy, \quad (\text{B.9})$$

where (B.7) follows from the definition of $G^n(x)$, and (B.8) holds by (B.5).

Since G is measurable and non-negative, there exists a sequence of step non-negative functions $G_n \leq G$ a.e. which converges to G a.e.; by (B.9), each G^n is measurable. Applying the monotone convergence theorem (as stated in Royden [70], p. 227) we get

$$\begin{aligned} E \int_0^{\tau(r; \omega)} G(S - D(t; \omega)) dt &= E \left[\int_0^{\tau(r; \omega)} \lim G^n(S - D(t; \omega)) dt \right] \\ &= \lim \frac{1}{\lambda} \int_r^S G^n(y) dy \\ &= \frac{1}{\lambda} \int_r^S G(y) dy, \end{aligned}$$

which completes the proof. \square

Appendix C

Proof of Theorem 3.3.5 for Periodic Review Systems

We start by proving that the minimal cost incurred by the facility during any finite time horizon is the same for policies both measurable \mathcal{F}^o and \mathcal{F} . This is accomplished by formulating a dynamic program for each of these cases and proving that any optimal cost for the first dynamic program is also optimal for the second dynamic program. This result is then extended to long-term average costs by taking limits as the time-horizon length tends to infinity.

Dynamic Programming Formulation for Policies Measurable \mathcal{F}^o .

We formulate the single-stage subproblem as a standard dynamic program in which the *time horizon* is equal to the set $\{0, 1, 2, \dots, M\}$.

1. The *state* x_t^o represents the facility's inventory position before any order is placed.
2. The *action* $y_t^o = x_t^o + q_t^o$, where $q_t^o \geq 0$ is equal to the amount shipped to

the facility at the beginning of period t , represents the facility's inventory position after a decision to ship inventory to the facility is made.

3. A *cost* equal to $\bar{g}^o(x_t^o, y_t^o) = \chi_{[y_t^o > x_t^o]}(K + \text{ord}(y_t^o)) + G(y_t^o)$ is incurred immediately after the shipping decision is made. Recall that $\chi_{[\cdot]}$ is the indicator function.
4. A *random demand* d_t^o occurs introducing the system evolution constraint

$$x_{t+1}^o = f^o(x_t^o, y_t^o, d_t^o) = y_t^o - d_t^o.$$

The dynamic programming algorithm states that the cost functional $J_t^o(x_t^o)$, defined as the optimal expected cost for the remaining periods $t, t+1, \dots, M$ when starting at period t with inventory position x_t^o , is uniquely determined by the equations:

$$J_M^o(x_M^o) = \inf_{y^o \geq x_M^o} \{\bar{g}^o(x_M^o, y^o)\} \quad (\text{C.1})$$

$$J_t^o(x_t^o) = \inf_{y^o \geq 0} \{\bar{g}^o(x_t^o, y^o) + E_{d_t^o}[J_{t+1}^o(y^o - d_t^o)]\} \quad (\text{C.2})$$

$$t = 0, 1, \dots, M-1$$

Dynamic Programming Formulation for Policies Measurable \mathcal{F} .

The introduction of information external to the facility results in the following formulation, in which the time horizon is also equal to the set $\{0, 1, 2, \dots, M\}$.

1. The state $x_t = (x_t^o, x_t')$, where x_t' represents the state of the rest of the system before any shipping decision for period t is made.
2. The decision $y_t = (y_t^o, y_t')$, where $y_t^o \geq x_t^o$ and y_t' represents actions taken outside the facility.

3. The one-period cost is given by the function $\bar{g}(x_t, y_t) = \bar{g}^o(x_t^o, y_t^o)$ because only the costs incurred by the facility are considered.
4. The random disturbance $d_t = (d_t^o, d_t')$ introduces the system evolution constraint

$$\begin{aligned} x_{t+1} &= f(x_t, y_t, d_t) \\ &= (y_t^o - d_t^o, f'(x_t, y_t, d_t)). \end{aligned}$$

The dynamic programming algorithm applied to this case states that the cost functional $J_t(x_t)$, defined as the optimal expected cost for the remaining periods $t, t+1, \dots, M$ when starting at period t in state x , is uniquely determined by the equations:

$$J_M(x_M^o, x'_m) = \inf_{y^o \geq x_M^o, y'} \{\bar{g}(x_M, y)\} \quad (\text{C.3})$$

$$\begin{aligned} J_t(x_t^o, x'_t) &= \inf_{y^o \geq 0, y'} \{\bar{g}(x_t, y) + E_{d_t^o, d_t'}[J_{t+1}(f(x_t, y_t, d_t))]\} \\ &\quad t = 0, 1, \dots, M-1 \end{aligned} \quad (\text{C.4})$$

Claim C.1

$$J_t(x_t^o, x'_t) = J_t^o(x_t^o), \quad \text{for all } x_t^o, t = 0, 1, \dots, M.$$

Proof

We prove this claim by induction. The claim is clearly true for $t = M$ by (C.1), (C.3) and $\bar{g}(x_t, y_t) = \bar{g}^o(x_t^o, y_t^o)$.

Now assume that the claim is true for $t+1$ and let us prove that the claim also holds for t .

$$J_t(x_t^o, x'_t) = \inf_{y^o \geq 0, y'} \{\bar{g}(x_t, y) + E_{d_t^o, d_t'}[J_{t+1}(f(x_t, y_t, d_t))]\}$$

$$\begin{aligned}
&= \inf_{y^o \geq x_t^o, y'} \{ \bar{g}^o(x_t^o, y^o) + E_{d_t^o, d_t'} [J_{t+1}(y_t^o - d_t^o)] \} \\
&= \inf_{y^o \geq x_t^o} \{ \bar{g}^o(x_t^o, y^o) + E_{d_t^o} [J_{t+1}(y_t^o - d_t^o)] \} \\
&= J_t^o(x_t^o);
\end{aligned}$$

where the second equation follows from the definition of $\bar{g}(x, y)$ and the induction hypothesis. \square

The theorem we want to prove follows because the problem of minimizing the long-term average cost incurred by the facility starting at state x_0 is equivalent to minimizing the expression

$$\lim_{M \rightarrow \infty} (1/M) E \left\{ \sum_{t=0}^{M-1} \bar{g}(x_t, y_t) \right\},$$

and claim C.1 proves that the minimization for any finite time horizon is equivalent under both \mathcal{F}^o and \mathcal{F} . This concludes the proof. \square

Bibliography

- [1] ARCHIBALD, B., AND E. SILVER, *Policies Under Continuous Review and Discrete compound Poisson Demands*, Management Science **8** (1978), 899–909.
- [2] ARROW, K, T. HARRIS AND J. MARSCHACK, *Optimal Inventory Policy*, Econometrica **19** (1951), 250–272.
- [3] AXSÄTER, S., *Simple Solution Procedures for a Clas of Two-Echelon Inventory Problems*, Operations Research **38** (1990), 64–69.
- [4] AXSÄTER, S., *Continuous Review Policies for Multi-Level Inventory Systems with Stochastic Demand*, in Handbooks in Operations Research and Management Science, Vol. 4, Logistics of Production and Inventory, S. Graves, A. Rinnooy Kan and P. Zipdin (eds.), (1993), 175,197.
- [5] ATKINS, D., AND DE S., *94%-Effective Lot-Sizing for a Two Stage Serial Production/Inventory System with Stochastic Demand*, Working Paper 92–MSC–001, Faculty of Commerce, University of British Columbia, Vancouver, B.C., Canada V6T 1Z2.
- [6] BARLOW, R., AND F. PROSCHAN, *Mathematical Theory of Reliability*, John Wiley & Sons, Inc., New York, 1965.
- [7] BARLOW, R., AND F. PROSCHAN, *Statistical Theory of Reliability and Life Testing*, To Begin With, Silver Spring, MD, 1981.
- [8] BECKMANN, M, *An Inventory Model for Arbitrary Interval and Quantity Discounts of Demand*, Management Science **8** (1961), 35–57.
- [9] BELL, C., *Improved Algorithms for Inventory and Replacemen- Stocking Problems*, SIAM Journal of Applied Mathematics **18** (1970), 558–566.
- [10] BELLMAN, R., I. GLICKSBERG, AND O. GROSS, *On the Optimal Inventory Equation*, Management Science **2** (1955), 83–104.

- [11] BERTSEKAS, D. P., *Dynamic Programming and Stochastic Control*, Academic Press, New York, 1976.
- [12] BHATTACHARYA, R. AND E. WAYMIRE, *Stochastic Processes with Applications*, John Wiley & Sons, Inc., New York, 1990.
- [13] BILLINSLEY P., *Probability and Measure*, John Wiley & Sons, second edition, 1986.
- [14] BROWN, M., *Bounds, Inequalities, and Monotonicity Properties for Some Specialized Renewal Processes*, The Annals of Probability **8** (1980). 227–240.
- [15] CHEN, F., *94%-Effective Policies in a Two-Stage Serial System with Stochastic Demand*, SUBMITTED TO?????
- [16] CLARK, A., *A Dynamic, Single-Item, Multi-Echelon Inventory Model*, RM-2297, Santa Monica, California, The RAND Corporation (December 1958).
- [17] CLARK, A., AND H. SCARF, *Optimal Policies for a Multi-Echelon Inventory Problem*, Management Science **6** (1960), 475–490.
- [18] CLARK, A. AND H. SCARF, *Approximate Solutions to a Simple Multi-Echelon Inventory Model*, in : K. J. Arrow, S. Karlin and H. Scarf (eds), *Studies in Applied Probability and Management Science*, Stanford University Press, Stanford, CA, 88–110.
- [19] DEUERMEYER, B. AND L. SCHWARZ, *A Model for the Analysis of System Service Level in Warehouse-Retailer Distribution Systems: The Identical Retailer Case*, in *Studies in the Management Sciences*, Vol. 16, *Multi-Level Production/Inventory Control Systems: Theory and Practice*, L. Schwarz (ed.) North Holland (1981), 163–193.
- [20] DVORETZKY, A., J. KIEFER AND J. WOLFOWITZ, *The Inventory Problem: I, Case of Known Distributions of Demand; II, Case of Unknown Distributions of Demand*, Econometrica **20** (1952), 187–222 and 450–466.
- [21] DVORETZKY, A., J. KIEFER AND J. WOLFOWITZ, *On the Optimal Character of the (s, S) Policy in Inventory Theory*, Econometrica **21** (1953), 586–596.
- [22] EHRHARDT, R., *The Power Approximation for Computing (s, S) Inventory Policies*, Management Science **25** (1979), 777–786.
- [23] EHRHARDT, R., *(s, S) Policies for a Dynamic Inventory Model With Stochastic Lead Times*, Operations Research **32** (1984), 121–132.

- [24] EHRHARDT, R., C. SCHULTZ AND H. WAGNER, *(s, S) Policies for a Wholesale Inventory System*, in *Multi-Level Production/Inventory Control Systems: Theory and Practice*. L. B. Schwarz (ed.) North Holland, 145–162.
- [25] EPPEN, G., *Effects of Centralization on Expected Costs in a Multi-Location Newsboy Problem*, *Management Science* **25** (1979), 498–501.
- [26] EPPEN G., AND L. SCHRAGE, *Centralized Ordering Policies in a Multi-Warehouse System with Lead Times and Random Demand*, in Schwartz, L. B., Ed., *TIMS Studies in the Management Sciences*, North Holland, Amsterdam (1981), Vol. 16, 51–67.
- [27] ERNST & YOUNG, *Retail information technology : Ernst & Young's survey of expenses and trends*, in *Chain Store Age Executive*, Retail Information Technology Supplement, (1994) p. 6–32.
- [28] EVANS, R., *Inventory Control of a Multi-Product System with a Limited Production Resource*, *Naval Research Logistics Quarterly* **14** (1967), 173–184.
- [29] FEDERGRUEN, A., *Centralized Planning Models for Multi-Echelon Inventory Systems Under Uncertainty*, in *Handbooks in Operations Research and Management Science*, Vol. 4, *Logistics of Production and Inventory*, S. Graves, A. Rinnooy Kan and P. Zipkin (eds.), (1993), 133–173.
- [30] FEDERGRUEN, E. AND Z. SCHECHNER, *Formulas for Continuous Review Inventory Models with Fixed Delivery Lags*, *Operations Research* **31** (1983), 957–965.
- [31] FEDERGRUEN, E. AND P. ZIPKIN, *Computing Optimal (s, S) Policies in Inventory Models with Continuous Demands*, *Advances in Applied Probability* **17** (1985), 424–442.
- [32] FEDERGRUEN, E. AND P. ZIPKIN, *Approximations of Dynamic, Multilocation Production and Inventory Problems*, *Management Science* **30** (1984), 69–84.
- [33] FEDERGRUEN, E. AND P. ZIPKIN, *Computational Issues in an Infinite-Horizon, Multiechelon Inventory Model*, *Operations Research* **32** (1984), 818–836.
- [34] FEDERGRUEN, E. AND P. ZIPKIN, *An Efficient Algorithm for Computing Optimal (s, S) Policies*, *Operations Research* **32** (1984), 1268–1285.
- [35] FREELAND, J. AND E. PORTEUS, *Evaluating the Effectiveness of a New Method for Computing Approximately Optimal (s, S) Inventory Policies* *Operations Research* **28** (1980), 353–364.

- [36] FREELAND, J. AND E. PRORTEUS, *Easily Computed Inventory Policies for Periodic Review Systems: Shortage Cost and Service Level Models*, Working Paper, Graduate School of Business, Stanford University, Stanford, CA.
- [37] GALLEGO, G., *The Distribution Free Newsboy Problem: Review and Extensions*, Journal of Operations Research **44** (1993), 825–834.
- [38] GALLEGO, G., *A Minmax Distribution Free Procedure for the (Q, R) Inventory Model*, Operations Research Letters **11** (1992), 55–60.
- [39] HADLEY, G. AND M. WHITIN, *Analysis of Inventory Systems*, Prentice-Hall, Inc., (1963).
- [40] HANEVELD, W., *A Stochastic Programming Approach to Multiperiod Production Planning*, University of Groningen Institute of Economic Research, Research Memorandum 276 (1988).
- [41] HARRIS, F., *Operations and Cost*, (Factory Management Series), A. W. Shaw Co. (1915), 48–52.
- [42] IGLEHART, D. L., *Optimality of (s, S) Policies in the Infinite Horizon Dynamic Inventory Problem*, Management Science **9** (1963), 259–276.
- [43] ITO, K., *Lectures on Stochastic Processes*, Tata Institute of Fundamental Research, distributed by Springer-Verlag, New York (1984).
- [44] JACKSON, P., AND J. A. MUCKSTADT, *Risk Pooling in a Two-Period, Two-Echelon Inventory Stocking and Allocation Problem*, Naval Research Logistics **36** (1989), 1–26.
- [45] JACKSON, P., W. MAXWELL AND J. MUCKSTADT, *Determining Optimal Reorder Intervals in Capacitated Production-Distribution Systems*, Management Science **34** (1988), 938–958.
- [46] JACKSON, P., W. MAXWELL AND J. MUCKSTADT, *The Joint Replenishment Problem with a Powers-of-Two Restriction*, AIIE Transactions **17** (1985), 25–32.
- [47] JACKSON, P. AND R. ROUNDY, *Minimizing Separable Convex Objectives on Arbitrarily Directed Trees of Variable Upper Bound Constraints*, Mathematics of Operations Research **16** (1991), 504–533.
- [48] JOHNSON, E., *On (s, S) Policies*, Management Science **15** (1968), 80–101.
- [49] KAKUMANU, A., *Relation Between Continuous and Discrete Time Markovian Decision Problems*, Naval Research Logistics Quarterly **24** (1977), 431,440.

- [50] KOUVELIS, P. AND G. GUTIERREZ, *The Newsvendor Problem in a Global Market: Optimal Centralized and Decentralized Control Policies for a Two-Market Stochastic Inventory System*, to appear in *Management Science*.
- [51] KÜENLE, C. AND H. KÜENLE, *Durchschnittsoptimale Strategien in Markovschen Entscheidungsmodellen bei Unbeschränkten Kosten*, *Math. Operationsforsch. Statist., Ser. Optimization* **8** (1977) 549-564.
- [52] MAXWELL, W., AND J. A. MUCKSTADT, *Establishing Consistent and Realistic Reorder Intervals in Production-Distribution Systems*, *Operations Research* **33** (1985), 1316-1341.
- [53] MITCHELL, J., *98%-Effective Lot-Sizing for One-Warehouse, Multi-Retailer Inventory Systems with Backlogging*, *Operations Research* **35** (1987), 399-404.
- [54] MUCKSTADT, J., *On the Probability Distribution for Inventory Position in Two Echelon Continuous Review Systems*, Technical Report No. 336, School of Operations Research and Industrial Engineering, Cornell University (1977).
- [55] MUCKSTADT J., AND R. ROUNDY, *Analysis of Multistage Production Systems*, in *Handbooks in Operations Research and Management Science*, Vol. 4, *Logistics of Production and Inventory*, S. Graves, A. Rinnooy Kan and P. Zipdin (eds.), 59-131.
- [56] MUCKSTADT, J. A., AND R. O. ROUNDY, *Multi-Item, One Warehouse, Multi-Retailer Distribution Systems*, *Management Science* **33** (1987), 1613-1621.
- [57] NAHMIAS, S., *Perishable Inventory Theory: A Review*, *Operations Research* **30** (1982), 680-708.
- [58] PETERS, R., K. BOSKMA AND H. KUPPER, *Stochastic Programming in Production Planning: A Case with Non-Simple Recourse*, *Statistica Neerlandica* **31** (1977), 113-126.
- [59] PORTEUS, E., *Numerical Comparisons of Inventory Policies for Periodic Review Systems*, *Operations Research* **33** (1985), 134-152.
- [60] PORTEUS, E. L., *An Adjustment to the Norman-White Approach to Approximating Dynamic Programs*, *Operations Research* **27**, 1203-1207.
- [61] PORTEUS, E. L., *Stochastic Inventory Theory*, in D. P. Heyman and M. J. Sobel, eds., *Handbooks in OR & MS*, Elsevier Science Publishers B.V., Vol. 2, (1990), Chapter 12.

- [62] PUTERMAN, M. L., *Markov Decision Processes*, in Handbooks in Operations Research and Management Science, Vol. 2., Stochastic Models. D. Heyman and M. Sobel (Eds.), North Holland, Amsterdam (1990), 331–434.
- [63] REID, MICHAEL *Change at the Check-Out*, The Economist, **v334n7904** (1995), ss1–ss18.
- [64] ROBERTS, D., *Approximation to Optimal Policies in a Dynamic Inventory model*, in Studies in Applied Probability and Management Science, K. Arrow, S. Karlin and H. Scarf. Stanford University Press, Standford, Calif., (Chap 13).
- [65] ROSS, S., *Introduction to Probability Models*, fifth ed., San Diego, CA, Academic Press, Inc., 1993.
- [66] ROUNDY, R., *Computing Nested Reorder Invervals for Multi-Item Distribution Systems*, Operations Research **38** (1990), 37–52.
- [67] ROUNDY, R., *Rounding off to Powers of Two in Continuous Relaxations of Capacitated Lot Sizing Problems*, Management Science **35** (1989), 1433–1442.
- [68] ROUNDY, R., *A 98%-Effective Lot-Sizing Rule for a Multi-Product, Multi-Stage Production/Inventory System*, Mathematics of Operations Research **11** (1986), 699–727.
- [69] ROUNDY, R., *98%-Effective Integer-Ratio Lot-Sizing for One-Warehouse Multi-Retailer Systems*, Management Science **31** (1985), 1416–1430.
- [70] ROYDEN H., *Real Analysis*, MacMillan Publishing Co., Inc., New York (1968)
- [71] SAHIN, I., *On the Stationary Analysis of Continuous Review (s, S) Inventory Systems with Constant Lead Times*, Operations Research **27** (1979), 717–729.
- [72] SAHIN, I., *On the Objective Function Behavior in (s, S) Inventory Models*, Operations Research **30** (1982), 709–724.
- [73] SAHIN, I., AND D. SINHA, *On Asymptotic Approximations for (s, S) Policies*, Stochastic Analysis and Applications **5** (1987), 189–212.
- [74] SCARF, H., *The Optimality of (s, S) Policies in the Dynamic Inventory Problem*, Mathematical Methods in the Social Sciences (1960), 196–202.
- [75] SCARF, H., *A Min-Max Solution of an Inventory Problem*, in Studies in The Mathematical Theory of Inventory and Production. (K. Arrow, S. Karlin and H. Scarf, Eds) Stanford University Press, California. 201–209.
- [76] SCARF, H., *A Survey of Analytic Techniques in Inventory Theory*, in Multistage Inventory Models and Techniques, H. Scarf, D. Gilford and M. Shelly (Eds.), Stanford University Press, California (1963), 185–225.

- [77] SCHNEIDER, H., *Methods for Determining the Re-order Point of an (s, S) Ordering Policy when a Service Level is Specified*, Journal of Operations Society **29** (1978), 1181–1194. CHECK THE NAME OF THIS JOURNAL
- [78] SIVAZLIAN, B., *Dimensional and Computational Analysis in (s, S) Inventory Problems with Gamma Distributed Demand*, Management Science **17** (1971), B307–B311.
- [79] SOX, C., *Modeling and Analysis of Quick Response in Production-Inventory Systems*, Doctoral Dissertation, Cornell University, (1992).
- [80] TIJMS, H. C., *Analysis of (s, S) Inventory Models*, Mathematical Centre Tracts, Amsterdam (1972).
- [81] VEINOTT, A., *The Status of Mathematical Inventory Theory*, Management Science **12** (1966) 745–777.
- [82] VEINOTT, A. F. AND H. M. WAGNER, *Computing Optimal (s, S) Inventory Policies*, Management Science **11** (1965), 525–552.
- [83] WAGNER, H., M. O'HAGAN AND B. LUNDH, *An Empirical Study of Exact and Approximately Optimal Inventory Policies*, Management Science **11** (1965), 690–723.
- [84] ZABEL, E., *A Note on the Optimality of (S, s) Policies in Inventory Theory*, Management Science **9** (1962), 123–125.
- [85] ZHENG Y. S., *A Simple Proof for Optimality of (s, S) Policies*, Journal of Applied Probability **28** (1991), pages.
- [86] ZHENG Y. S. AND A. FEDERGRUEN, *Finding Optimal (s, S) Policies is About as Simple as Evaluating a Single Policy*, Operations Research **39** (1991), 654–665.
- [87] ZHENG Y. S., *On Properties of Stochastic Inventory Systems*, Management Science **38** (1992), 87–103.
- [88] ZIPKIN, P., *Stochastic Leadtimes in Continuous-Time Inventory Models*, Naval Research Logistics Quarterly **33** (1986), 763–774.