Chapter 1

INTRODUCTION

Inventory management is one of the most important tasks in business. A business faces inventory problems in its most basic activities. Inventory is held by the selling party to meet the demand made by the buying party. The complexity of inventory problems varies significantly, depending on the situation. While some of the simple inventory problems may be dealt with by common sense, some other inventory problems that arise in complex business processes require sophisticated mathematical tools and advanced computing power to get a reasonably good solution.

The fundamental problem in inventory management can be described by the following two questions: (1) when should an order be placed? and (2) how much should be ordered? There are two basic trade-offs in an inventory problem. One is the trade-off between setup costs and inventory holding costs. By placing orders frequently, the size of each order can be made relatively small. Therefore, the holding costs can be reduced. However, the total setup costs will go up. Conversely, less frequent orders will save on setup costs but incur higher holding costs. The other trade-off is between holding costs and stockout costs. Holding more inventory reduces the likelihood of stockouts, and vice versa. These trade-offs give rise to an optimization problem of finding the optimal ordering policy that minimizes the overall cost.

1.1. Characteristics of Inventory Systems

The primary purpose of inventory control is to manage inventory to effectively meet demand. The effectiveness can be measured in many different ways. The most commonly used methods measure how much the total profit and how little the total cost are. The majority of the

D. Beyer et al., *Markovian Demand Inventory Models*, International Series in Operations Research & Management Science 108,
DOI 10.1007/978-0-387-71604-6_1, © Springer Science+Business Media, LLC 2010

operations research literature on inventory management has used the criterion of cost minimization. In many cases, this criterion is equivalent to that of profit maximization.

There are many factors that should be taken into consideration when solving an inventory problem. These factors are usually the characteristics of, or the assumptions made about, the particular inventory system under consideration. Among them, the most important ones are listed below.

Cost structure. One of the most important prerequisites for solving an inventory problem is an appropriate cost structure. A typical cost structure incorporates the following four types of costs.

- *Purchase or production cost.* This is the cost of buying or producing items. The total purchase/production cost is usually expressed as cost per unit multiplied by the quantity procured or produced. Sometimes a quantity discount applies if a large number of units are purchased at one time.
- *Fixed ordering (or setup) cost.* The fixed ordering cost is associated with ordering a batch of items. The ordering cost does not depend on the number of items in the batch. It includes costs of setting up the machine, costs of issuing the purchase order, transportation costs, receiving costs, etc.
- *Holding (or carrying) cost.* The holding or carrying cost is associated with keeping items in inventory for a period of time. This cost is typically charged as a percentage of dollar value per unit time. It usually consists of the cost of capital, the cost of storage, the costs of obsolescence and deterioration, the costs of breakage and spoilage, etc.
- *Stockout cost.* Stockout cost reflects the economic consequences of unsatisfied demand. In cases when unsatisfied demands are backlogged, there are costs for handling backorders as well as costs associated with loss of customer goodwill on account of negative effects of backlogs on future customer demands. If the unsatisfied demand is lost, i.e., there is no backlogging, then the stockout cost will also include the cost of the foregone profit.

Demand. Over time, demand may be constant or variable. Demand may be known in advance or may be random. Its randomness may depend on some exogenous factors such as the state of the economy, the weather condition, etc. Another important factor often ignored in the inventory literature is that demand can also be influenced directly or indirectly by the decision maker's choice. For example, a promotion decision can have a positive effect on demand.

Leadtime. The leadtime is defined as the amount of time required to deliver an order after the order is placed. The leadtime can be constant (including zero) or random.

Review time. There are two types of review methods. One is called continuous review, where the inventory levels are known at all times. The other is called periodic review, where inventory levels are known only at discrete points in time.

Excess demand. Excess demand occurs when demand cannot be filled fully from the existing inventory. The two common assumptions are that excess demand is either backordered or lost.

Deteriorating inventory. In many cases, the inventory deteriorates over time which affects its utility. For example, food items have limited shelf lives.

Constraints. There are various constraints involved in inventory problems. The typical constraints are supplier constraints – restrictions on order quantities; marketing constraints – minimum customer service levels; and internal constraints – storage space limitations, limited budgets, etc.

1.2. Brief Historical Overview of Inventory Theory

Inventory control has been a main topic in the operations management area for over half a century. Many advanced mathematical methods can be applied to solve inventory control problems. This makes it an interesting topic for researchers from a variety of academic disciplines. The basic question in any inventory problem is to determine the timing and the size of a replenishment order. While the fundamental questions in inventory problems remain the same, the focus of one particular inventory problem may be quite different from another.

It is virtually impossible to summarize the enormous literature on inventory models in this short section. Even with the limitation of focusing on studies aimed at finding optimal policies for dynamic inventory models, it is still a tremendous task to cover all related issues. Different features of inventory systems also lead to different treatments of the resulting problems. Silver (1981) has illustrated such a diversity of inventory problems by introducing a classification scheme consisting of the following characteristics of an inventory system.

Single vs. Multiple Items

Deterministic vs. Probabilistic Demand

Single Period vs. Multiperiod

Stationary vs. Time-Varying Parameters

Nature of the Supply Process

Procurement Cost Structure

Backorders vs. Lost Sales

Shelf Life Considerations

Single vs. Multiple Stocking Points

Among various factors that affect the model characterization, the nature of the demand process is usually the most important one. Assumptions about other parameters of an inventory model also have important implications on the model development. However, these assumptions are more or less standard and well-accepted in practice.

In this book, we will consider finite and infinite horizon inventory models involving a single product, a single stocking point, and a stochastic demand. We will assume that supply is unconstrained and the product has an infinite shelf life. We will allow the system to be with either stationary or time-varying parameters. Both backorders and lost sales models will be treated. We will consider zero leadtime. However, the backlogging models with fixed nonzero leadtimes can be transformed to ones with no leadtime.

The main distinguishing characteristic of the inventory models studied in this book is that the demand will be dependent on a Markov process. The Markov process may be exogenous (i.e., outside our control) or it may be influenced by our inventory decisions.

In what follows, we will briefly summarize the relevant literature based on the assumptions made about the demand process in a single product setting.

1.2.1 Deterministic Demand Models

The most celebrated model in this category is the economic order quantity (EOQ) model, which is the simplest and the most fundamental of all inventory models. It describes the important trade-off between the fixed ordering cost and the holding cost, and is the basis for the analysis of many more complex systems.

The basic EOQ model is based on the following assumptions.

(i) The demand rate is a constant ξ units per unit time.

(ii) Shortages are not permitted.

(iii) The order leadtime is zero.

(iv) The order costs consist of a setup cost K for each ordered batch and a fixed unit cost c for each unit ordered.

(v) The holding cost is h per unit held per unit time.

The famous EOQ formula is given by

$$Q^* = \sqrt{\frac{2K\xi}{h}}.$$

Because of its form, it is also called the *square-root formula; (see also EOQ formula).* The formula is due to Harris (1913); (see Erlenkotter (1989) for a history of the development of this formula).

Some important variations of the EOQ model include planned backorders/lost sales, quantity discounts, a known shelf life, a known replenishment leadtime, or constraints on the order size. A fairly comprehensive discussion on these and related topics can be found in Hadley and Whitin (1963).

Another important work in deterministic demand models is the so-called dynamic lot size problem, where the demand level changes from period to period but in a known fashion. Wagner and Whitin (1958) provide a recursive algorithm for computing the optimal policy. Some interesting properties of the optimal policy are identified by them. The Wagner-Whitin model was republished in a 2004 issue of *Management Science* with a commentary by Wagner (2004).

1.2.2 Stochastic Demand Models

The newsvendor formula for the single period inventory problem is one of the most important results in the stochastic inventory theory; (see Edgeworth (1888), Arrow *et al.*(1951)). Many extensions of the classical newsvendor model have been studied in the literature. These allow for random yields, pricing policies, free distributions, etc. A comprehensive review of the literature on the newsvendor models can be found in Khouja (1999), and references therein.

The base-stock policy has been extensively studied in the literature for single-product periodic review inventory problems with stochastic demand and no fixed setup cost. The early academic discussions on the base-stock policy have focused on proving its optimality under the situation that the demands are independent. Examples can be found in

Gaver (1959,1961), Karlin (1958c), Karlin and Scarf (1958), etc. Bellman *et al.* (1955) and Karlin (1960) are the classical papers for the stationary and nonstationary demand cases, respectively. Optimality of a base-stock policy has been established in these situations.

For single-item models with probabilistic demand and a fixed ordering cost, it has been shown that an (s, S) policy is optimal under a variety of conditions. The policy is to order up to the level S when the inventory level is below s and to not order otherwise. In the case of no fixed ordering cost, the optimal policy becomes a base-stock policy, which is a special case of the (s, S) policy when $s = S$.

Classical papers on the optimality of (s, S) policies in dynamic inventory models with stochastic demands and fixed setup costs are those of Arrow *et al.*(1951), Dvoretzky *et al.* (1953), Karlin (1958c), Scarf (1960), and Veinott (1966). Scarf develops the concept of K-convexity, and uses it to show that if the ordering cost is linear with a fixed setup cost K and if the inventory/backlog cost function is convex, then the optimal ordering policy in any given period for a finite horizon model is characterized by two critical numbers s_n and S_n, with $s_n \leq S_n$, such that if the inventory level x_n at time n is less than s_n, then order $S_n - x_n$; if $x_n \geq s_n$, then do not order. That a stationary (s, S) policy is optimal for the stationary infinite horizon problem is proved by Iglehart (1963a). Optimality of an (s, S) policy in the lost sales case was proved by Shreve (1976). Proof in the lost sales case does not extend to the case when the leadtime is nonzero; (see, e.g., Zipkin (2008b)).

Veinott (1966) presents a model similar to Scarf's but with a different set of conditions, which neither imply nor are implied by the conditions used by Scarf (1960). Instead of using the concept of K-convexity, he proves the optimality of an (s, S) policy by showing that the negative of the expected cost is a unimodal function of the initial inventory level. His model also provides a unified approach for handling both backlogging and lost sales assumptions, as well as the case of perishable products.

1.2.3 Markovian Demand Models

Most classical inventory models assume demand in each period to be a random variable independent of environmental factors other than time. However, many randomly changing environmental factors, such as fluctuating economic conditions and uncertain market conditions in different stages of a product life cycle, can have a major effect on demand. For such situations, the Markov chain approach provides a natural and flexible alternative for modeling the demand process. In such an approach, environmental factors are represented by the demand state or *the state-of-the-world* of a Markov process, and demand in a period is a random variable with a distribution function dependent on the demand state in that period. Furthermore, the demand state can also affect other parameters of the inventory system such as the cost functions.

The effect of a randomly changing environment in inventory models with fixed costs received only limited attention in the early literature. It is conceivable that when the demands in different periods are dependent, the structure of the optimal policy does not change. The main difference is that the parameters of an optimal policy will depend on the demand state in the previous period.

Iglehart and Karlin (1962) consider the problem with dependent demand and no setup cost. They consider an inventory model with the demand process governed by a discrete-time Markov chain. In each period, the current value of the state of the chain decides the demand density for that period. They prove that the optimal policy is a state-dependent base-stock policy.

Karlin and Fabens (1960) introduce a Markovian demand model in which demand depends on the state of an underlying Markov chain. They indicate that given the Markovian demand structure in their model, it appears reasonable, in the presence of a fixed ordering cost, to postulate an inventory policy of (s, S)-type with a different set of critical numbers for each demand state. However, due to the complexity of the analysis, Karlin and Fabens concentrated on optimizing over the class of ordering policies, with each policy characterized by a single pair of critical numbers s and S, irrespective of the demand state.

Song and Zipkin (1993) present a continuous-time formulation with a Markov-modulated Poisson demand and with linear costs of inventory and backlogging. They show the optimality of a state-dependent (s, S) policy for the case when there is a fixed ordering cost. An algorithm for computing optimal policies is also developed using a modified value iteration approach. Sethi and Cheng (1997) provide a discrete-time version of the problem with general demand distributions. They show the optimality of a state-dependent (s, S) policy when the demand is modeled as a Markov-modulated process.

Cheng and Sethi (1999b) extend the results to the case of Markov-modulated demand with the lost sales assumption for unfilled demand. The optimality of a state-dependent (s, S) policy is proved for this case only under the condition of zero supply leadtime.

Another notable development in inventory models with Markovian demand is due to Chen and Song (2001). Their paper considers a multi-stage serial inventory system with a Markov-modulated demand. Random demand arises at Stage 1, Stage 1 orders from Stage 2, etc., and

Stage N orders from an outside supplier with unlimited stock. The demand distribution in each period is determined by the current state of an exogenous Markov chain. Excess demand is backlogged. Linear holding costs are incurred at every stage, and linear backorder costs are incurred at Stage 1. The ordering costs are also linear. The objective is to minimize the long-run average cost in the system. The paper shows that the optimal policy is an echelon base-stock policy with state-dependent order-up-to levels. An efficient algorithm is also provided for determining the optimal base-stock levels. The results can be extended to serial systems in which there is a fixed ordering cost at stage N and to assembly systems with linear ordering costs.

1.2.4 Models with Controllable Demands

When marketing tools such as promotions are used to stimulate consumer demand, the coordination between marketing and inventory management decisions becomes important. While it is clear that an integrated decision-making system should be adopted in order to maximize total benefits and avoid potential conflicts, little theoretical work can be found in the literature to guide such integrated decision making. Most works in the inventory theory do not deal with promotion decisions explicitly. Usually, promotion decisions are taken as given so that the effect of promotions can be determined in advance of inventory replenishment decisions. Thus, the inventory/promotion problem reduces to a standard inventory control problem with demand unaffected by the actions of the decision maker. However, there have been a few papers that address the issue of joint inventory/promotion decision making. Most of these focus on inventory control problems in conjunction with price discount decisions.

Used frequently in the literature, a simplified approach for modeling the demand uncertainty is to assume that the random components of demand are either additive or multiplicative; (see Young (1978)). Based on this assumption, Karlin and Carr (1962), Mills (1962), and Zabel (1970) obtained similar results. They showed that the optimal price is greater than the riskless price in the multiplicative case and is less than the riskless price in the additive case. Furthermore, Thowsen (1975) showed that if the demand uncertainty is additive, the probability density function of the additive uncertainty is PF_2, and if the expected demand function is linear, then the optimal policy is a base-stock list price policy. If the initial inventory level is below the base-stock level, then that stock level is replenished and the list price is charged; if the initial inventory level is above the base-stock level, then nothing is ordered and a price discount is offered.

A dynamic inventory model with partially controlled additive demand is analyzed by Balcer (1983), who presents a joint inventory control and advertising strategy and shows it to be optimal under certain restrictions. However, in the Balcer model, the demand/advertising relationship is deterministic because only the deterministic component of demand is influenced by advertising decisions.

Cheng and Sethi (1999a) extend the result of Markov-modulated demand models by introducing promotional decisions by which demand is affected. For the case of no fixed ordering cost, they show that a state-dependent base-stock policy is optimal. Chen and Simchi-Levi (2004) analyze a single-product, periodic-review model in which pricing and production/inventory decisions are made simultaneously. They show that when the demand model is additive, the profit-to-go functions are K-concave; hence, an (s, S, p) policy is optimal. In such a policy, the inventory is managed using an (s, S) policy, and price is determined based on the inventory position at the beginning of each period.

1.2.5 Other Extensions

Beside different treatments for modeling the demand process, extensions intended to relax other restrictive assumptions on inventory models or to incorporate certain special features into inventory models have also been studied by researchers. A great deal of effort has been made to adopt a more general cost function rather than the linear cost with a setup cost. An interesting case of such a general cost function is that of concave increasing cost. Porteus (1971) shows a generalized (s, S)-type policy to be optimal when the ordering cost is concave and increasing, given an additional assumption that the demand density function is a Pólya frequency function.

Other extensions for single-product inventory models include incorporating more realistic features of the physical inventory systems, such as capacity constraints, service level constraints, and supply constraints. Veinott (1965) considers a model where orders must be placed in multiples of some fixed batch size Q. It is shown that an optimal policy is characterized by Q and a number k as follows. If the initial inventory position is less than k, one orders the smallest multiple of Q, which brings the inventory position to at least k; otherwise, no order is placed.

Among the recent research trends in the analysis of dynamic inventory systems, one particular issue is to model the effect of more flexible procurement options such as multiple delivery modes in connection with the demand forecast updates available over time. Sethi *et al.* (2003) present an inventory model with fixed costs, forecast updates, and two delivery modes, with a forecast-update-dependent (s, S)-type policy shown to be optimal. This and related models are also the subject of the book by Sethi *et al.* (2005a).

Recently, Bensoussan *et al.* (2007a,2008a) have considered Markovian demand models with partially observed demands. In these models, called the censored newsvendor models, only sales are observed when the demand exceeds the inventory level. In other models, the information regarding the inventory is incomplete; (see, e.g., Bensoussan *et al.* (2005a,2007b,2008b)) .

There are several other related works extending the classical inventory models in various ways. We do not intend to provide a complete list of all relevant work in this fairly broad and active research area. Instead, we have chosen to survey only those that are related to the topic of this book either from a historical point of view or because similar modeling issues are also addressed in the book.

1.2.6 Computational Methods

Following the theoretical work on establishing the optimal control policies under various conditions, numerical procedures for computing optimal policies have also been developed. Although the structure of (s, S) policies is fairly simple, the computation of an (s, S) policy turns out to be difficult. Traditionally, two main approaches were explored in computing the actual (s, S) values. The first is based on the dynamic programming formulation of the problem. The second approach, known as stationary analysis, involves the determination of the limiting distribution of the inventory position under an (s, S) policy. Using this distribution, an expected cost function is constructed with s and S as decision variables; (see, e.g., Veinott and Wagner (1965)). Minimization of this function determines the optimal (s, S) policy. A computationally more efficient approach is the one developed by Zheng and Federgruen (1991).

The literature reviewed here is intended to provide the readers with a basic idea about which models are relevant to this book. How models presented in different chapters of this book relate to the literature will be discussed in those chapters.

1.3. Examples of Markovian Demand Models

The classical and best known stochastic inventory models assume the demand in different time periods to be independent and often even identically distributed. Nonetheless, demand forecasters have known all along that the demand over time for most products does not form a sequence of independent random variables. Most forecasting techniques either rely on trends and autocorrelation in the sequence of demands or on the fact that future demands are correlated to some observable early indicator. A specific structure of a Markov-modulated demand process is used in this book in an attempt to formalize these dependencies and allow for them when making inventory decisions. Below are a few examples of such indicators and demand processes consistent with our models.

The influence of weather on the demands for many products is one of the first things that come to mind. The demand for home-heating oil in winter depends heavily on the temperature outside. Likewise, the demand for electricity in summer increases as the temperatures rise and vice versa. Different temperature ranges can be thought of as different states of the world determining the probability distribution of demand. Once the temperature is known, the demand is still uncertain but two things have changed; first, knowing the ambient temperature, demand can likely be predicted more accurately, and therefore the demand distribution will have lower variance. Second, since temperatures in adjacent time periods are correlated, those correlations can be exploited to predict future temperatures and demands. The transition matrices of the temperatures in a geographic area can be estimated from the historical weather data. Also note that these matrices could be nonstationary over time, i.e., they may be different in different months. For more weather related examples, the demand for insecticides in an area may be positively correlated to the amount of total rainfall in the preceding week. A good indicator for gasoline demand at gas stations along stretches of highway I-80 would be the weather forecast for the Lake Tahoe region. Here, the demand could be a high-level categorization of the forecast; for example, pleasant, mediocre, or unpleasant.

When it comes to estimating the demand for consumable materials like ink or toner cartridges for a specific type of printer, it is obvious that the number of such printers currently in use, often referred to as the installed base, is an important piece of information. Therefore, the sales volume combined with data about typical usage patterns of the products requiring the same consumables can be used to provide the state information of the demand. The transition matrix of demand states can also be estimated based on the projected life spans of the relevant products. Similarly, in the case of demand for spare parts, the number of installed units (for example, servers) in a given age range should be a good choice to use as the demand state, since the failure rate of each of the items is usually related to its age.

Demand for replacement units under warranty also exhibits a similar pattern that is directly associated with the number of new units (for example, a specific model of a digital camera or a mainframe computer) sold recently. Many electronic devices show high "infant mortality." A large number of those that fail are either defective on arrival or fail shortly after being put to use. For these reasons, the number of replacement units requested under warranty is positively correlated to the recent sales volume of the product, a reasonable choice for the demand state.

The product life cycle is another important indicator for product demand. Typically, sales volume starts low in the introductory stage, increases significantly in the growth stage, reaches the peak in the mature stage, and then falls in the decline stage. It is a natural and realistic way to model the transitions of a product life cycle from one stage to the next as a Markov chain.

In the setting of competitive market environment, the introduction of a competing product into the market can be modeled as a Markov process with the state variable representing the number of competing products. A product may initially be the market leader when there are no competitive products at the time of its introduction. However, its market share is likely to diminish as competitive products are introduced over time.

In the IT industry, the introduction of new software is often a catalyst for hardware demand. For example, demand for memory upgrades on notebook computers will be higher after the introduction of a new version of software such as a new version of Microsoft Windows operating system. In this case, the introduction of the new Windows operating system can be used as the demand state when a Markovian model is adopted for modeling the demand of an affected hardware product.

Seasonality is a common factor that affects product demand. For example, the demand for many consumer electronics products like notebook computers or inkjet printers exhibits strong seasonal patterns. Demand for these products is the highest during "back to school" season (August, September) and the holiday season (November, December). During the remainder of the year, the monthly demand is lower and more stable. Seasonal demand is also typical in commercial sales. It is well known that IT related purchases and capital investments (servers, large laser printers, storage devices, etc.) follow a so-called "hockey stick" curve. Demand tends to be highly skewed towards the end of a month or a calendar quarter. This is thought to be the result of the spend-it-or-loose-it regime of corporate budgeting, as well as the structure of sales incentives. This seasonality of demand can be captured by a degenerate form of a Markov chain with the transition from one state to the next being deterministic (or almost deterministic). See Section 2.7 of Chapter 2 for further discussion on this type of model.

The examples given above illustrate that Markovian demand models can capture the relationship between the demand and the associated environmental factors. Certainly, these models are more realistic and more flexible representations of the demand process than i.i.d. demand models or nonstationary independent demand models. On the other hand, the inventory models allowing for Markovian demands are more complex than the classical inventory models especially in terms of data and computational requirements. However, inventory decisions have tremendous business consequences in terms of cost as well as customer satisfaction. Because of this, the additional effort required to use these models will easily be justified in many cases by the value of improved decisions resulting from them.

1.4. Contributions

One of the most important developments in the inventory theory has been to show that (s, S) policies (with base-stock policy as a special case) are optimal for a class of dynamic inventory models with random periodic demands and fixed ordering costs.

Dynamic inventory models can be formulated in a continuous-time setting or a discrete-time setting. The models presented in this book are of the latter type.

The main contribution of the book is to generalize a class of inventory models (with fixed ordering costs) to allow for demands that depend on an exogenous Markov process or a controlled Markov process, and to show that the optimal policies are of (s, S)-type. In our models, the distribution of demands in successive periods is dependent on a Markov chain. Special cases of such Markovian demand models include the case of cyclic or seasonal demand. Finite horizon, as well as stationary and nonstationary infinite horizon, problems are considered. Some constraints commonly encountered in practice, namely no-ordering periods, finite storage capacities, and service levels, are also treated. We show that (s, S)-type policies are also optimal for the generalized models as well as their extensions.

We provide rigorous analyses under various modeling assumptions: discounted cost models with backlog or lost sales assumptions, average cost models (also with backlog or lost sales assumptions), models with demand controlled by a Markov decision process (MDP), and models with cost functions of polynomial growth. Required results on the existence of optimal solutions, as well as verification theorems, are established for the cases studied.

1.5. Plan of the Book

This book consists of six parts. Part I (consisting of just Chapter 1) is an introduction to the book. Part II (Chapters 2-4) covers the models with the discounted cost criterion. The models with the average cost criterion are presented in Part III (Chapters 5-7). Part IV (Chapters 8 and 9) collects miscellaneous results that are not covered in the previous chapters. The concluding remarks are discussed in Part V. Appendices are included in Part VI.

Part II starts with Chapter 2, where a discounted cost model with the full backlog assumption is introduced. The model is a generalization of classical inventory models with fixed costs that exhibit (s, S) policies. We model the demand process in a way that the demand distributions in successive periods depend upon the state of an underlying Markov chain. A dynamic programming formulation is used to provide the results on the uniqueness of the solution and the existence of an optimal feedback policy. Also described in the chapter are some new properties of K-convex functions, which provide the technical results needed for the analysis in the rest of the chapter. The optimality of an (s, S)-type policy is first shown in a finite horizon model, with s and S dependent on the demand state of the current time period and on the time remaining. This result is extended to a nonstationary infinite horizon version of the model. Extensions with more realistic modeling assumptions are also presented in this chapter.

Chapter 3 treats the discounted infinite horizon inventory model involving fixed cost in Chapter 2 to allow for unbounded demand and costs with polynomial growth. Finite horizon problems, as well as stationary and nonstationary discounted cost infinite horizon problems, are addressed. Existence of optimal Markov or feedback policies is established with unbounded Markovian demand, ordering costs that are l.s.c., and inventory/backlog (or surplus) costs that are l.s.c. with polynomial growth. Furthermore, optimality of state-dependent (s, S) policies is proved when the ordering cost consists of fixed and proportional cost components and the surplus cost is convex.

Chapter 4 extends the results developed in Chapter 2 to the lost sales case, where demand that is not satisfied on time is lost. The lost sales case is typically more difficult to analyze than its backlog counterpart. Generally it requires a different treatment and sometimes additional assumptions to establish the same type of results as in the backlog case. A new K-convexity result associated with the cost functions in the lost sales case is presented in this chapter. Based on this new result, we are able to show that a state-dependent (s, S) policy is optimal for the Markovian demand inventory models with the lost sales assumption in both finite horizon and infinite horizon formulations. Extensions that incorporate various realistic features and constraints are also developed in the chapter.

Part III, consisting of Chapters 5, 6 and 7, analyzes the average cost models. Chapter 5 deals with a long-run average cost version of the Markovian demand model treated in Chapter 2 with the backlog assumption, fixed ordering cost, and convex inventory cost. We develop a vanishing discount approach to establish the average cost optimality equation and provide the associated verification theorem to show that a state-dependent (s, S) policy is optimal for the average cost model.

Chapter 6 is devoted to studying the long-run average cost version of the model presented in Chapter 3 with unbounded Markovian demands, ordering cost that are l.s.c., and inventory costs that are l.s.c. and of polynomial growth. Finite horizon problems and stationary long-run average cost problems are addressed.

Chapter 7 provides an analysis of the lost sales version of the problem described in Chapter 5. Similar to the approach used in Chapter 5, we examine the asymptotic behavior of the differential discounted value function as the discount rate approaches zero. Then, we establish the average cost optimality equation using the vanishing discount approach. Also included is the proof of the verification theorem as well as the optimality proof for a state-dependent (s, S) policy in this setting.

In Part IV, two additional topics that are not covered by the previous chapters are discussed. In Chapter 8, we present a MDP based model for a joint inventory/promotion decision problem. The state variable of the MDP represents the demand state brought about by changing environmental factors as well as promotion decisions. We show that the optimal joint inventory/promotion decision can be characterized by policy of a simple form under certain conditions, where the promotion decision follows a threshold policy and the inventory decision follows a base-stock policy, with both policies dependent on the demand state.

Chapter 9 revisits the classical papers of Iglehart (1963b) and Veinott and Wagner (1965) devoted to stochastic inventory problems with the criterion of long-run average cost minimization. We indicate some of the assumptions that are implicitly used without verification in their stationary distribution approach to the problems, and provide the missing (nontrivial) verification. In addition to completing their analysis, we examine the relationship between the stationary distribution approach and the dynamic programming approach to the average cost stochastic inventory problems.

We conclude the book with Part V consisting of Chapter 10, where conclusions and open research problems are described.

Part VI consists of Appendices. These Appendices provide some background material as well as the technical results that are used in the book.

The relationships between different chapters are shown in Figure 1.1.

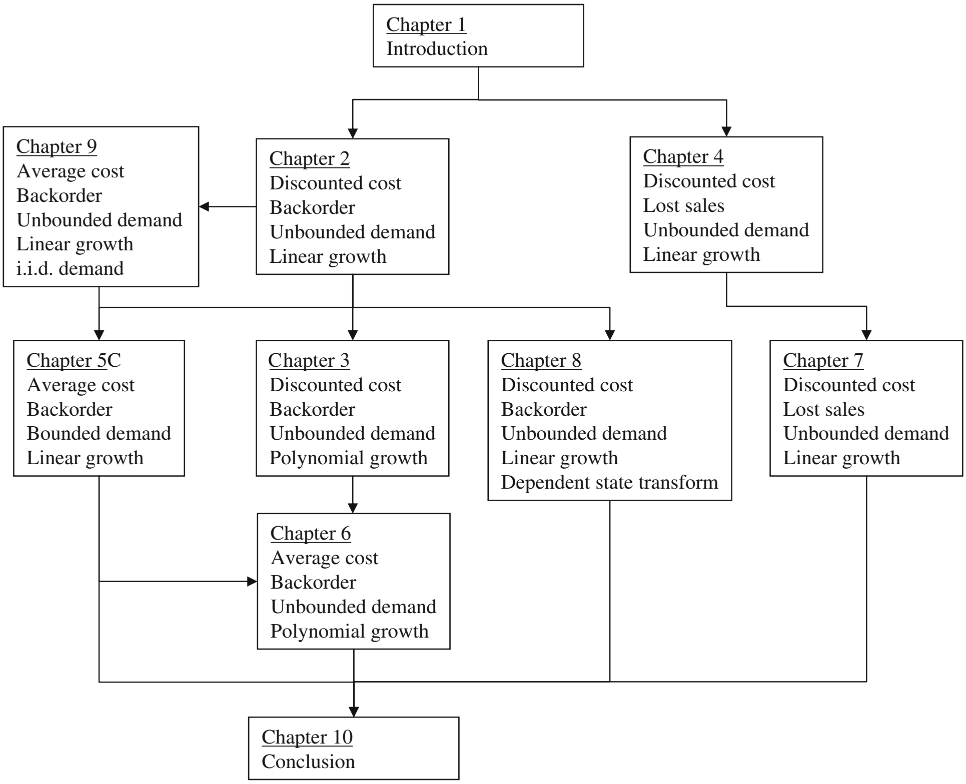

Figure 1.1. Relationships between different chapters.

An extensive bibliography, author index, subject index, and the copyright permissions follow the Appendices.