

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Optimal Policy Structures of Stochastic Supply Chains with Outsourced Logistics Agreements

Permalink

<https://escholarship.org/uc/item/9hf4w628>

Author

Alper, Osman Engin

Publication Date

2010

Peer reviewed|Thesis/dissertation

**Optimal Policy Structures of Stochastic Supply Chains with Outsourced
Logistics Agreements**

by

Osman Engin Alper

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering – Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Philip M. Kaminsky, Chair
Professor Hyun-Soo Ahn
Professor Andrew Lim
Professor David Aldous

Fall 2010

**Optimal Policy Structures of Stochastic Supply Chains with Outsourced
Logistics Agreements**

Copyright 2010
by
Osman Engin Alper

Abstract

Optimal Policy Structures of Stochastic Supply Chains with Outsourced Logistics Agreements

by

Osman Engin Alper

Doctor of Philosophy in Engineering – Industrial Engineering and Operations Research

University of California, Berkeley

Professor Philip M. Kaminsky, Chair

As evidenced by the rapid growth of the third-party logistics industry, more and more firms are electing to outsource logistics in order to cut costs and to focus on core competencies. One of the key decisions faced by firms when engaging third party logistics providers involves the nature of the agreement with the provider. Agreements range from requirements for service on demand to more structured agreements, in which the timing and size of shipment quantities are specified in advance. By agreeing to more structured arrangements, firms can decrease the uncertainty faced by the logistics provider and thus the logistics provider's costs, and therefore negotiate better rates. In order to understand the impact of using more structured agreements, however, firms need to understand how to effectively utilize the service provided in these agreements.

In this dissertation, motivated by these observations, we develop stylized models of simple production-distribution systems in order to explore the efficient use of structured logistics agreements. First, we introduce the general framework of models that we explore in this manuscript and present our motivation behind it. Second, we present a brief review of the stochastic supply chain literature. In this field, while there has been considerable interest in especially the supply contracts, the emphasis has been mostly on channel coordination and other informational efficiency aspects rather than the operational efficiency that we focus on. Third, we introduce our basic model with a fixed commitment logistics contract in a make-to-order production setting. We mathematically formulate this problem using stochastic dynamic programming and fully characterize the optimal policy structure. We prove some important properties of the optimal policy function that describe its sensitivity to reserved capacity levels and shipment times in the contract. We also provide sufficient conditions for decomposing this problem in time, whereby the complexity in computationally determining the optimal policy parameters can be greatly reduced. Fourth, we extend our model to a make-to-stock production environment and again completely characterize the structure of the optimal policy function. We show monotonicity of the optimal function parameters with

respect to committed capacities in all cases and with respect to time in the periodic shipment case. Fifth, we analyze and extend our results regarding the optimal policy structure to more sophisticated logistics agreements such as option contracts and multi-level option contracts, and introduce additional uncertainties to the system such as stochastic spot market price and stochastic availability of additional capacity. Sixth, we present an initial analysis of the logistics agreements with shipment times chosen dynamically by the contact buyer. Lastly, we provide a computational study illustrating the sensitivity of optimal contract parameters to demand uncertainty and cost parameters, as well as exploring the relative benefits of different logistics agreements under varying operating conditions.

Overall, we investigate simplified models of production-distribution systems with outsourced logistics, our analysis and characterization of optimal policies provide some insight into the practical use of transportation contracts in addition to building a foundation for future investigation of models that incorporate more complicated critical aspects of important real-world problems relating to integrated production and distribution management in the presence of outsourced logistics agreements.

To my family...

Contents

List of Figures	iv
List of Tables	v
1 Introduction	1
1.1 Motivation	2
1.2 Overview of Our Work	3
2 Literature Review	5
2.1 The Clark-Scarf Model and its Extensions	5
2.2 Different Network Structures	11
2.3 Computational Methods	14
2.4 Setup Costs for Transportation	16
2.5 (R, nQ) Policies	18
2.6 Capacity Constraints	22
3 Make-to-Order Production	27
3.1 The Model	27
3.2 One Shipment without Capacity	28
3.3 One Shipment with Capacity	29
3.4 Multiple shipments with capacity	34
4 Make-to-Stock Production	38
4.1 Logistics Agreements with Scheduled Shipments	39
4.2 Logistics Agreements with Options and Stochastic Spot Market Price	45
4.3 Multi-level Option Agreements	48
4.4 Stochastic Availability of Additional Capacity	50
4.5 Logistics Agreements with Flexible Shipments	53
5 Computational Study	55
5.1 Optimal Contract Parameters	55
5.2 Comparison of Contracts Types	63

6 Conclusion	69
6.1 Summary of Results	69
6.2 Future Directions	71
Bibliography	72

List of Figures

3.1	The shifting property of the ship-down-to levels.	33
5.1	The effect of unit discount, $\beta=0.98$	57
5.2	The effect of unit discount, $\beta=0.90$	57
5.3	The effect of additional unit discount, $\gamma=1$	58
5.4	The effect of additional unit discount, $\gamma=0.90$	58
5.5	The effect of demand variance on total cost.	59
5.6	The effect of demand variance on value of a contract.	59
5.7	The effect of demand variance on frequency of shipments.	61
5.8	The effect of waiting cost ratio on optimal contract capacity.	62
5.9	Optimal total expected cost vs reserved capacity level in the fixed commitment contract with $c_r = 7, c_p = 5, c_s = 10, h = 1, p = 5, T = 3, \alpha = 0.95$	63
5.10	Optimal total expected cost vs reserved capacity level in the option contract with $c_r = 2.5, c_e = 5, c_p = 5, c_s = 10, h = 1, p = 5, T = 3, \alpha = 0.95$	64
5.11	Optimal total expected cost vs reserved capacity level in the option contract with $c_r = 1, c_e = 6.5, c_p = 5, c_s = 10, h = 1, p = 5, T = 3, \alpha = 0.95$	64
5.12	Optimal total expected cost vs reserved capacity level in the capacity flexible contract with $c_r = 2.25, c_e = 5, \gamma = 0.15, c_p = 5, c_s = 10, h = 1, p = 5, T = 3, \alpha = 0.95$	65
5.13	Performance comparison of the contracts.	67

List of Tables

5.1	The effect of demand variance on optimal contract parameters.	61
5.2	The effect of waiting cost ratio on optimal contract parameters.	61
5.3	Summary results illustrating how the optimal contract parameters change with respect to costs and demand variability.	62
5.4	Fixed Commitment Contract: The effect of demand variance on optimal contract parameters.	66
5.5	Option Contract: The effect of demand variance on optimal contract parameters.	66
5.6	Capacity Flexible Contract: The effect of demand variance on optimal contract parameters.	66
5.7	The value of contract flexibility for the buyer, $\gamma = 0.15, c_p = 5, c_s = 10, h = 1, p = 5, T = 3, \rho = 0.95$	68

Acknowledgments

It is difficult to overstate my gratitude to my Ph.D. advisors, Professor Philip M. Kaminsky and Professor Hyun-Soo Ahn. Their guidance, from research insights to tips on presenting at research conferences, have been invaluable. Without their steadfast encouragement and generous support, I could not have written this dissertation.

I would also like to express my appreciation to Professor Andrew Lim and Professor David Aldous, my dissertation committee members and teachers, for reviewing this manuscript and giving inspiring lectures in stochastic dynamic programming and applied probability from which I benefited greatly.

Financial support by the Marshall-Oliver-Rosenberger Fellowship is gratefully acknowledged.

I would like to thank Cahit Güngör, Kemal Ozanoğlu, İlker Temizer, Saygın Kahraman, İsmail Ceylan, İsmail Onur Filiz, Şener Aktürk, Can Özcan, Caner Önal, Ankit Jain, Vishnu Bhama Narayanan, Erick Moreno Centeno, and Yannan Shen for their friendship and support. Without all the good times we have had, the graduate life would have been unbearable.

Lastly, I would like to express my deepest love and gratitude to my family, my mother Emine Alper, my father Ali Aykut Alper, and my brother Ersin Alper. Having their unconditional love and support gave me the strength I needed in going through hardships.

Chapter 1

Introduction

The purpose of this dissertation is to investigate the effective utilization of structured logistics agreements in production-distribution systems with stochastic demand in order to understand how optimal policies of such systems are affected by outsourced logistics while gaining managerial insights into practical real-life applications and contributing to the stochastic supply chain literature. To that extent, in this manuscript:

1. We propose simplified models of production-distribution systems with logistics agreements, mathematically formulate them using stochastic dynamic programming, and analyze their optimal policy structures;
2. Through a computational study combined with our analytical results regarding the properties of the optimal policy functions, we explore how relative benefits of the different logistics agreements we consider and their optimal contract parameters change under varying operating conditions.

There has been limited previous work that explicitly consider structured agreements between logistics providers and their clients. As a related area, supply contracts have gathered considerably greater attention, but the emphasis has been mostly on channel coordination and other informational efficiency aspects rather than the operational policies that we focus on. Among the work that have a similar outlook, Alp et al. [1] model a transportation contract design problem; Whittimore [80] shows optimal ordering policies when there are two supply options, one being faster and more expensive than the other; Blumfeld et al. [9] and Gallego and Simchi-Levi [34] look at integrated inventory control and vehicle dispatching problems; Yano and Gerchak [83] and Yano [82] consider expedited shipments in their models, and determine optimal agreement parameters assuming a base-stock ordering policy; Henig et al. [39] investigate optimal ordering policies under a given supply contract.

In our base case model, we consider a firm utilizing a fixed-date, fixed-capacity transportation agreement for shipping to its retail site. In each period, this firm must decide how much to ship, possibly nothing, from a warehouse to a retailer to meet demand at

the retailer, utilizing a combination of shipping capacity already agreed to via a structured logistics agreement, and shipping capacity available on the spot market. Note that since there is a cost to holding inventory at the retailer, it does not always make sense to utilize all of the available contracted capacity, even though this capacity is already paid for. For this setting, we characterize the optimal shipping policy, which we show to be coupled with the production policy in this setting owing to our simplifying assumptions of uncapacitated production with linear cost and no delay. This modeling approach is to some extent similar to [39] in that we study the optimal behavior of a simple system with a single decision variable under a given contract, however we consider both make-to-stock and make-to-order production environments and a general fixed commitment contract instead of a contract with constant reserved capacity in each period. This allows us to capture the time between shipments in the transportation contract realistically as an integer multiple of the review period. Moreover, we extend this base case analysis to more sophisticated contracts such as option contracts and multi-level option contracts, as well as introducing additional uncertainties to the system such as stochastic spot market price and stochastic availability of additional capacity.

In the rest of this chapter, we first briefly give our motivation for choosing this area, and then present a brief overview of our work.

1.1 Motivation

Production and distribution operations are generally critical operational functions in manufacturing supply chains. Indeed, United States industry spends more than \$350 billion on transportation and more than \$250 billion on inventory holding costs annually (Lambert and Stock [52]). Although production and distribution operations can typically be decoupled if there is sufficient inventory between them, this approach leads to overall higher inventory levels and longer lead times in the supply chain than if these functions are more closely linked. Indeed, as supply chain inventory reduction efforts are becoming more and more common, the linkages between production and distribution are becoming tighter and tighter, requiring a significant focus on decision making in which production and distribution operations are optimized in an integrated manner.

On the other hand, over the last decade, more and more firms have elected to outsource their logistics-related responsibilities. Indeed, the last decade has witnessed a rapid growth of the global third-party logistics (3PL) industry, which was estimated at \$390 billion (Quinn [61]) in 2007. Traditionally, the 3PL-client relationship has been an unstructured one, built around long-term strategic relationships in which the 3PL provider commits to providing whatever resources the client demands, whenever the resources are needed. However, Coffey et al. [20] cite dramatically declining margins among 3PL providers, and suggests that this high level of service may be too expensive to continue to offer, particularly in cases when clients do not dramatically benefit from this service. As an example, [20] recounts

the experience of a 3PL provider that reacts to daily shipment requests from a client, even when these requests fluctuate wildly, even though these requests do not accurately reflect the client's customer demand, as the shipped goods may sit for weeks in a warehouse before ultimately being sold. Given these declining margins, and the obvious difficulty faced by 3PL firms trying to take full advantage of the inherent economies of utilizing their resources over a number of clients, [20] suggests that there will be growing use of more structured agreements between 3PL providers and their clients, in which a specific schedule of services is agreed to in advance, so that the 3PL provider can more effectively utilize these resources.

At the same time, increasing cost pressures mean that many firms that previously outsourced logistics are looking to cut the cost of logistics services. Decreasing margins at 3PL providers, however, reduce the likelihood that contract negotiations will result in lower cost for the same service – clearly, these firms will have to agree to a different level of service from their 3PL providers in order to lower the price that 3PL providers charge. It is likely that more structured logistics agreements will increasingly be an avenue for clients to reduce 3PL provider costs. Of course, as logistics agreements become more structured, it is increasingly important for the client firms to develop effective operating strategies that account for these more structured logistics agreements.

1.2 Overview of Our Work

In this dissertation, we analyze effective operating strategies under a variety of different structured operating agreements. By structured operating agreements, we mean arrangements that explicitly limit the timing or capacity of logistics services over a specified time horizon. For example, these agreements might specify the dates and capacity of individual shipments over the next year (see [83]). In some arrangements, these dates and the capacity might be fixed, while in others, these parameters might be flexible. In the base case of our models, we analyze the optimal operating strategy for a firm that already has in place a fixed-date, fixed-capacity logistics agreement with a 3PL provider. In all models we consider, we assume that in addition to the logistics services available to the firm through existing agreements with 3PL providers, the firm can purchase expedited shipping services on an open market, although typically at a higher cost than those available through long term agreements. Indeed, these long term agreements provide the buyer access to lower cost transportation services, and at the same time allow the 3PL provider to plan and use its resources more effectively. However, these agreements restrict the buyer's ability to react to changing market conditions, and thus we explore other types of agreements that allow more flexibility in either timing or capacity of shipments, or provide the opportunity for the buyer to purchase additional shipments at below market prices.

In Chapter 2, we present a brief review of the multi-echelon stochastic supply chain literature. In Chapter 3, we introduce our base case model of a just-in time manufacturer utilizing an outsourced logistic agreement for shipping to a retail site. In this chapter, we

focus our attention on periodic fixed commitment contracts, which specify in advance the volume to be reserved and the frequency to be shipped. In Section 3.2 and Section 3.3, we analyze two special cases of this problem, which facilitate understanding of the fundamental concepts and lay the groundwork for proving results regarding the structure of the optimal policy function, its basic properties and a decomposition result for the general case we analyze in Section 3.4. In Chapter 4, we start analyzing a variation of our base model within a make-to-stock production environment, where the manufacturer can stock up inventory at the retail location instead of reacting to incoming orders alone. Section 4.1 considers fixed commitment contracts, establishes the optimal policy structure policy, and extends two fundamental monotonicity results of the optimal policy parameters from Chapter 3 to the make-or-order production environment. In Section 4.2, we consider option contracts, which give the right to purchase shipping in the future at predetermined price and capacity. This section also introduces additional uncertainty to the model with stochastic spot market rates whereby these option contracts can be used as a hedging mechanism against spot market shipping rates. Section 4.3 analyzes multi-level option agreements (or capacity flexible contracts), in which a predetermined capacity is reserved by the logistics provider and committed by the buyer (either by a fixed commitment or an option agreement), but a certain percentage of this capacity can be supplied in addition to the reserved portion with an additional cost. In Section 4.4, we extend our model to the case where additional capacity can be supplied to the contract buyer at a fixed rate, but the availability of this additional capacity is driven by an independent stochastic process. Section 4.5, on the other hand, models an agreement that provides flexibility in shipping time rather than capacity. In Chapter 5, we present our computational study; the sensitivity of the optimal contract parameters to demand uncertainty and costs is investigated in Section 5.1 and the relative benefits of different logistics agreements to the contract buyer under various operating conditions are explored in Section 5.2. Chapter 6 concludes this dissertation, summarizing our main results and proposing future work on related areas.

Chapter 2

Literature Review

Although the models we consider in this manuscript have contractual elements in them, in essence, they have their roots in classical multi-echelon inventory theory. That is because we focus on the operational aspects of our models such as optimal ordering and shipping decisions that minimize costs instead of channel coordination problems that have been mostly studied in this context. From an operational perspective, classical inventory control problems can be interpreted as special cases of the corresponding inventory control problems with contracts. As a simple example to illustrate this idea, consider the classical finite horizon periodic review newsvendor problem. As Anupindi and Bassok [2] point out, this classical problem can be viewed as analysis of a contract (say, newsvendor supply contract) with a given horizon length, proportional purchase price (often considered fixed), fixed periodicity of ordering (given by the period length), no commitments, and unlimited flexibility (there are usually no limits on the order quantity).

As a consequence of this close relationship, in this chapter, we present a brief review of the multi-echelon inventory models, which are also central to the supply chain management literature. The theory is voluminous with categories by, e.g., demand characteristics (deterministic, stochastic), and control characteristics (periodic review, continuous review). We will focus on stochastic periodic review problems as that is what we adopt in our models. Although, we note that, there is generally a direct correspondence between discrete-time and continuous-time models under certain assumptions. Some well known books that cover the essentials of this field are Porteus [60], Zipkin [86], and Simchi-Levi et al. [74].

2.1 The Clark-Scarf Model and its Extensions

The interest in stochastic inventory problems traces back to Arrow et al. [4], see Arrow [3] for a personal account on the genesis of this paper. Their work studies many aspects of inventory theory including some deterministic models, the now popular single-period newsvendor model, and general dynamic problems. In their analysis of the dynamic problem,

they make the specific assumption that the cost of purchasing stock is composed of two parts: a setup cost K incurred whenever an order is placed, and a unit cost c proportional to the size of the order. This is the case in which the optimal inventory policy was suspected to be an (s, S) policy. The optimality of such a policy was not known when [4] was written. What they did instead was to restrict their attention to policies of this particular form to calculate the discounted expected cost associated with each such policy and to discuss the selection of that pair (s, S) yielding the lowest cost.

Karlin and Scarf [47] extends the inventory models in [4] by adding a time lag, l , between the order and the arrival of the items. This problem is readily representable by extending the dimension of the state space by $l - 1$. However since the decision at any period affects the cost function in l periods, the cost in the next $l - 1$ periods is fixed. Moreover, if we assume the excess demand is backlogged at every period then the order of the extended state vector is irrelevant to the cost function after the order time lag, only the sum of the elements of this vector plays a role. Hence the policy problem can be converted to the original single state dynamic problem very easily with an additional constant term depending on the initial extended state whenever excess demand is backlogged. In other words, the inventory problem with a time lag reduces to one in which essentially no lag exists.

Note that when there is positive order leadtime and excess demand is lost rather than backordered, the total cost is a non-linear function of the vector of outstanding orders and the optimal policy is extremely complex. Because of the complexity of the optimal policy the research on this area is concentrated on finding bounds on the optimal policy and approximating it. See Morton [53] for an approximation to the proportional cost infinite horizon problem and Nahmias [55] for different approximations methods for several leadtime lost sales inventory models including a positive setup cost for ordering, uncertainty in the leadtime, and partial backordering of demand.

Kaplan [45] analyzed an extension of [47] to incorporate leadtime uncertainty. What makes the random leadtime problem difficult is the way one treats order crossing. If orders are placed with one supplier it is unlikely that they would cross in time; that is, an order placed on Monday should arrive before one placed on Tuesday even though the exact arrival times may not be certain. The difficulty is that if orders are not permitted to cross, successive leadtimes are dependent random variables. Kaplan treats this order crossing issue cleverly in his formulation of the problem by defining probabilities, $p_i = Pr\{\text{all orders placed } i \text{ or more periods ago arrive in the current period}\}$. This formulation guarantees that orders do not cross since the arrival of an order i periods ago forces the arrival of orders placed $i + 1, i + 2, \dots, m$ periods ago as well. The likelihood that the leadtime is i periods, q_i , may be computed from p_i . It is the values of q_i one would observe in a real system. In [45] using this formulation and assuming no order setup cost, it is shown that the optimal policy is a critical number policy in every period. Hence the deterministic leadtime results carry over to this model of stochastic leadtimes.

Karlin [48] extends the dynamic inventory model of [4] by allowing the distribution of the stochastic demand function to change from period to period while still maintaining

independence. Hence the demand variables in every period are still independent but not identically distributed. It is shown that the base-stock policy remains to be optimal, although now the critical numbers (order-up-to levels) varies with the period. However, this is not the main objective of this paper. Its real contribution is providing qualitative results describing the variations of the critical numbers over time as a function of the demand distributions in all future periods when the distributions are stochastically ordered. The model in Karlin [49] is a special case of the one in [48] when the demand functions change in a cyclical pattern, which could reflect a situation where the demand undergoes seasonal variations. The structure of the optimal policy does not change as before, only the critical numbers again depend on the period (in particular, that period's relative place in the cycle). More importantly, a computational method for determining explicitly the critical numbers, which characterize the optimal policy is provided. In both papers infinite horizon discounted cost criterion is assumed. The proofs in both papers and the validation of the computational method rely on analysis of the infinite horizon functional equation and its solution. Zipkin [85] considers the same problem, provides a more elementary method of proof which applies to the infinite horizon periodic demand problem under both discounted and average cost criteria, and also extends the results to the case of cyclic cost functions as well as demand distributions. Moreover, qualitative description of the varying optimal critical numbers over time in [48] is extended to the nonstationary costs case, and given a new interpretation based on exponential smoothing technique.

Scarf [68] extends the news-vendor model with setup cost to multiple periods and by naturally devising the now well known K -convexity, shows optimality of the (s, S) policy. Then in their seminal work Clark and Scarf [18] extends this single-installation model to a purely serial supply chain with time lags between each installation. In their model they assume that the demand originates in the system at the lowest installation, and at no other point in the system; and that each installation backlogs excess demand. Clark in his earlier work recursively defines the "echelon stock" as the total stock consisting of the on hand inventory level (for all the installations except the bottom one this is physical on hand stock hence nonnegative, for the bottom installation this is on hand stock minus backlog) at any given installation plus the stock in transit to the next lower installation plus the echelon stock in that installation. Using this concept, they construct a cost structure in which the holding and shortage costs are functions of the echelon stock levels at each installation and the cost of purchasing and shipping an item is linear only exception being the highest installation, where a setup cost is permitted. This cost structure turns out to be a crucial assumption in the decomposition of the total system cost function into a sum of echelon cost functions, which is shown by a simple induction argument on time periods. In numerous extensions of this model in the inventory control literature, this cost structure is implemented as linear holding costs at every installation increasing towards the bottom of the system, which is justified by added value to the item as it progresses down the serial structure. The shortage cost is generally assigned only to the bottom installation, being the zero function at every other installation.

The emerging key idea is that some kind of a penalty, which they call “induced penalty cost”, must be induced on an installation for keeping a quantity of stock on hand which is insufficient to meet the normal requests from the lower installation. This penalty is merely the expected increment in total cost of the lower installation caused by the shortage of items in the higher installation. The result of the decomposition is a series of single-location problems. The lowest installation, which faces customer demand, considers only its own costs; ignoring all others. Under the assumed cost structure a critical-number policy solves this problem. The optimal policy and expected cost function for this installation are then used to define a convex induced penalty function which is added to the one-period cost function of the next installation in the system. This procedure is repeated until the highest installation. The problem of the highest installation, which makes ordering decisions with fixed and proportional cost structure, is solved by an (s, S) policy. This solution constitutes an optimal order policy for the whole system. In this paper, they also extend this model to the case where multiple installations face direct demand and also show by an example that decomposition is not possible in a non-serial supply chain structure. This reference is later cited as one of the 10 most influential papers published in *Management Science* in the 50 years of the journal’s existence. See Clark [17], and Scarf [69] for interesting comments on [18]. Also see Clark [16] for a retrospective on multi-echelon inventory theory, Scarf [70] for his personal reflections on inventory theory, and Veinott [77] for an extensive survey of mathematical inventory theory as of 1966 including both deterministic and stochastic dynamic models. The result of [18] has been reproved using novel approaches; Chen and Zheng [15] use cost allocation lower bounds to show the result in infinite horizon average cost case, and Muharremoglu and Tsitsiklis [54] uses an alternative approach based on item-customer decomposition to prove the result for finite and infinite horizon problems under both average and discounted cost criteria. These two papers will be reviewed shortly later in the document.

While this decomposition and the resulting solution procedure greatly simplifies the original problem, actual computation of an optimal policy still encounters substantial obstacles. First, two sets of recursive functional equations instead of one set of recursive equations on a composed larger dimension state space must be solved numerically. Secondly and in addition to the first, each evolution of the induced penalty cost function itself entails a numerical integration over the optimal cost function for the lower installation’s problem; indeed, with an order leadtime of several periods, the computation requires a double integration.

These computational issues makes the extension of [68] to infinite horizon case especially rewarding. Iglehart makes the first attempt of tackling this task in Iglehart [42] and [43]. In [43], he proves the optimality of (s, S) policy in the infinite horizon discounted cost case by assuming unimodularity of an endogenous value function. Although the complete proof eludes him; he shows that the series $\{s_n\}$ and $\{S_n\}$ are uniformly bounded, and by using this shows the uniform convergence of value functions.

Federgruen and Zipkin [28] show that the qualitative result of [18] extends to the infinite horizon case, first under the criterion of discounted cost, and then by using this under the

long-term average cost. The resulting two single-location problems are computationally much easier to solve than their respective finite-horizon versions as is often the case in dynamic problems. They also show that in the infinite horizon problems the induced penalty cost functions are stationary, do not involve optimal cost functions, and require at most one numerical integration. The case of normally distributed demands requires no explicit numerical integration and the penalty cost can be expressed in terms of univariate and bivariate normal cumulative functions, which are readily available in standard statistical packages. In summary, this important paper does not only fill a theoretical gap but provides considerable simplifications in practical implementation of the results in [18].

The construction of the infinite horizon problems are as in Veinott [75], where the value functions are defined on the class of infinite-horizon measurable policies and the realization of the entire sequence of demands. Firstly, the induced penalty functions are shown to uniformly converge to a limit function. This limit function, as all the induced penalty functions converging to it, is nonnegative and convex. Then using this result, they show that the difference between the cost functions defined with this limit function the original ones uniformly converges to zero as the planning horizon goes to infinity. Hence the defined limiting cost functions represent the stationary versions of the original ones, and the results follow directly from this for both the discounted and average cost criteria. The analysis in this paper draws heavily on [42], [43], uses well-established convergence properties of discounted and average cost infinite horizon stochastic inventory problems in Heyman and Sobel [40] and Bertsekas and Shreve [7]. The reader is further referred to Bertsekas [6] for a general treatment of dynamic programming and to Rudin [67] for a concise treatment of real analysis.

Even though the original proof in [18] can be modified to handle time varying demand distributions, the demand distributions over time would still need to be independent of each other. Chen and Song [13] generalize the results of [18] to the case where the outside demand distribution is modulated by an exogenous Markov process. The Markov-modulated demand process extends the application of the basic model to a wide range of fluctuating demand environments attributable to, e.g., seasonal effects, price changes, market conditions, and demand forecasts. This type of a demand model, where the distribution of demand depends on the state of a modulating Markov chain has been used in many papers in literature considering single-location inventory models. The collective insight of these works is that the optimal policy for a model with fluctuating demand has the same structure as that in its stationary counterpart, except that the policy parameters must be adjusted to reflect the dynamics of the underlying demand environment. [13] shows that this insight can be carried over to multi-echelon settings.

They assume constant leadtime, proportional ordering costs at every stage, and that the demand process is driven by a discrete-time finite state Markov chain, which is time homogeneous and ergodic. These assumptions imply that this Markov chain has a unique stationary distribution. Since they consider the infinite horizon problem under average cost criterion, this setup facilitates the establishment of a lower bound on the long-run average

cost of any feasible policy through a series of demand-state reductions. Then they show that this lower bound is reached by a state-dependent base-stock policy.

They note that the result of this model can be extended to the case where there is a fixed ordering cost at the highest stage in which case the optimal policy at that stage is an echelon (s, S) policy. Since the lower-bounding procedure eventually reduces the multi-stage problem to a single-stage one; this extension is possible by using the results of Beyer and Sethi [8], who considered the single-stage fixed ordering cost model with Markovian demands and showed that the state dependent (s, S) policy is optimal in infinite horizon problem under average cost criterion. To generalize the results to assembly systems with no fixed ordering cost; one can follow Rosling [64], whose work will be summarized in the next section.

Lastly, Muharremoglu and Tsitsiklis [54] extend [13] with non-order crossing stochastic leadtimes, and show the optimality of base-stock policies in such systems considering both finite and infinite horizon problems, and under both average cost and discounted cost criteria. Their stochastic leadtime model extends [45] incorporating the same two important features, i.e., the non-order crossing property and the independence from the current status of other outstanding orders. In their model, just like in [45], an exogenous random variable determines which outstanding orders are going to arrive at a given stage. However, they additionally allow the stochastic leadtimes to depend on the state of a modulating Markov chain, which enable dependencies between the leadtime random variables corresponding to different stages in the system.

They note that what allows them to handle Markov-modulated demand and leadtime model is a new approach to the uncapacitated serial inventory problem. The standard approach is a decomposition into a series of single stage problems. Their approach instead relies on a decomposition into a series of unit-customer pairs. Consider a single unit and a single customer. Assume that the distribution of time until the customer arrives to the system is known and the goal is to move the unit through the system in a way that optimizes the holding versus backorder cost trade-off. Since only a single unit and a single customer are present, this problem is much simpler than the original one. They show that under the assumptions of this paper, the original problem is equivalent to a series of decoupled single-unit single-customer problems.

This approach renders handling of several extensions to the standard model in a simple manner possible. In particular, inductive arguments based on dynamic programming equations, which get in general quite tedious as more and more complexities are added to the model, can be bypassed. On the contrary, they use a simple qualitative argument to establish a monotonicity property of optimal policies for the single-unit single-customer problem. For finite horizon problems, the optimality of echelon base stock policies is an immediate corollary. The same is true for infinite horizon problems, once some required limiting arguments are carried out.

2.2 Different Network Structures

There has been interest in inventory literature to extend the results of [18] to different non-serial and more complex system structures. A divergent or distribution system is defined as one in which each installation has at most one predecessor, where as in an assembly system each installation has at most one successor. A serial system is both a distribution and an assembly system. In practice we also quite often meet general systems, where some installations have multiple predecessors as well as multiple successors. Such systems are very difficult to handle by scientific methods.

In an assembly system, a number of components acquired from outside vendors are assembled, typically in several stages, into subassemblies and then, finally, into a single end product. Assembly networks are therefore trees with the node at the root corresponding with the single end item, the leaves corresponding with the externally acquired components, and all other nodes with intermediate subassemblies. Assembly networks generalize serial systems in that each installation (node) has at most one successor node but may have more than one predecessor node. There has been no exact or approximate solution methods known to us for the direct extension of [18] where there is a fixed cost component for the ordering of the primary components at the leaves of the assembly network.

However, Karmakar [50] for the general system structure (multilocation inventory problem as he calls) with all proportional order, holding, and shortage costs manages to show that the base stock policy is optimal when the starting stock levels in a period are low enough (under optimal base stock level for that period) under the additional restriction that all leadtimes are zero. He takes a significantly different approach, which allows him to get very general results, in formulating the problem than the one taken in [68].

In his model there is absolutely no structure assumption on the network of the installations. Instead he defines an m -dimension activity vector z_t which represents any process that changes the stock levels such as ordering, shipment between locations or disposal. In any period t , the stock levels at all n installations at the beginning of the period are given by the n -vector x_t . The $(n \times m)$ activity coefficient matrix A_t captures the impact of these activities on the stock level. The resulting stock levels are given by $y_t = x_t + A_t z_t; z_t \geq 0$, which effectively defines the feasible target stock vectors y_t in period t . If the multilocation inventory problem involves only exogenous ordering and transshipment between locations, and if it is feasible to raise the stock level at all locations simultaneously, then the activity coefficient matrix A_t can be shown to be Leontief. (An $(m \times n)$ matrix A is Leontief if its columns have at most one positive element and there exists a nonnegative column n -vector x for which Ax is positive.) In this case there is a base stock vector \hat{y} such that if the starting stock levels are lower than \hat{y} , then it is optimal to raise stock levels to \hat{y} . The result is extended to multiperiod case with induction.

[50] has established the characterization of the optimal policies in general multilocation inventory problem. However, computational methods have been limited to the one-period case or to special structures. The difficulty lies in the inseparability of the optimal cost

functions and the stochastic transitions between periods. These factors make it difficult to numerically compute the cost functions and the expectations of the next period value functions. Of course if the state space were very small, discrete dynamic programming could be used, especially for the static, infinite horizon problem. However, in many instances the state space is quite large, since its dimension is equal to the number of locations. Karmakar [51] presents a Lagrangian decomposition of the problem that results in an easily computable lower bound for the problem, and a dual relaxation that gives an upper bound. He also provides some algorithms for computing these bounds and the associated computational test results.

Rosling [64] shows that an assembly system with linear order and assembly costs, linear holding and shortage costs, general leadtime, can be transformed into an equivalent series system, provided the initial stock levels satisfy some simple conditions which he calls long-run balance. He demonstrates that optimal policies (among others) in finitely many periods lead the system into long-run balance and keep it there under reasonable and intuitive restrictions on the values of the holding and shortage cost coefficients. This condition is therefore completely unrestrictive when minimizing long-run average costs over an infinite planning horizon, although it may be an issue in discounted cost criterion particularly when the discount factor is small. All carrying, outside order and assembly costs in the equivalent system remain linear as do the holding costs and the backlog cost for the final product. Hence it is optimal to employ base-stock policy at each location in each period for both finite or infinite horizon problems.

The equivalent series problem is obtained by computing the echelon leadtimes, M_i , which are defined as the total leadtime for item i and all its successors, for all locations. Then all the locations are placed in a series, starting with the location assembling the end product, in descending order of their echelon leadtimes. This ordering guarantees that, if location i is the successor of location j then i is placed before j . Locations are now numbered (renumbered if necessary), starting with the location assembling the end product, according to their placement in this constructed series. Without loss of generality and for convenience, author chooses the units in which items are measured, such that for any location a single unit of an item is required in the assembly of its successor. Then $X_{i,t}^s$ is defined as the echelon inventory position of location i in period t ordered at most s periods ago. Then the system is in long-run balance in period t if and only if for $i = 1, 2, \dots, N - 1$

$$X_{i,t}^s \leq X_{i+1,t}^s \quad \text{for } s = 1, \dots, M_i.$$

Thus, in long-run balance the inventory positions equally close to the end item increase with i , i.e., with total leadtime.

Schmidt and Nahmias [71] considers the simplest of all assembly networks, where two components are purchased from outside vendors, to be assembled into a single end item. Assuming linear order and assembly costs as in [64], they characterize the optimal policy under all possible combinations of initial component's and end product's stock. Under inap-

appropriately matched inventories, the optimal policy may have a tediously complex structure, but only in the initial periods of the planning horizon.

Carlson and Yano [10] consider special cases of assembly systems in which a number of externally purchased components are assembled in a single operation into a single end product. Unlike [64] and [71], they address order and assembly costs with fixed (setup) components, which amounts to a complete generalization of [18] in two echelon case. It is assumed that the periods with assembly runs are predetermined (by some other model) while all the remaining decisions, i.e., assembly quantities in the predetermined assembly periods, and other epochs and quantities for the components, are determined within the model on the basis of heuristics.

In a distribution system, a single node supplies multiple installations which in turn may supply to other installations or to the end customer. Distribution networks are therefore trees with the node at the root corresponding with the externally acquired components, the leaves corresponding with the end items, and all other nodes with intermediate items. Distribution networks generalize serial systems in that each installation (node) has at most one predecessor node but may have more than one successor node.

The basic network structure studied in distribution systems is a two stage (i.e., one warehouse multiple retailer) model. An important within this class of models is that between systems where inventory is carried at the central warehouse and systems without central inventories. The latter applies, e.g., when the warehouse does not represent a physical location at all, but rather a centralized function. The detailed decisions about shipments to ultimate destinations do not need to be made at the time a system-wide order is placed (with an outside supplier), but can be postponed till some time later (assuming positive leadtime for outside orders). Even if the central warehouse does correspond to a physical location it acts a transshipment center rather than a stocking point. Rosenfield and Pendrock [63] refer to systems with centralized stock (i.e., there is an actual warehouse or a transshipment center with positive leadtime for outside orders) as *uncoupled* and to systems without centralized stock (i.e., there is no actual warehouse and no leadtime for outside orders) as *coupled*. The economies of scale in the order costs for outside orders is an advantage of both coupled and uncoupled systems. However the distinct advantage of the uncoupled system the ability to postpone the allocations. This permits one to observe the demand in the intervening periods between the period the outside order is placed and the period it is actually received in the warehouse, and thus to make better informed allocations. Eppen and Schrage [23] coined the phrase *statistical economies of scale* for this effect.

Distribution systems are generally more difficult to analyze than assembly systems. An issue that arises in these systems is the problem of allocating the available inventory at a higher installation to the requesting nodes. Consider for example the simplest distribution network where there are two stages, one warehouse and two identical retailers. Assume that we in a certain period are able to allocate to the retailers' optimal order-up-to levels, $S_1^* = S_2^*$, which are obtained from solving their problems in isolation. After this allocation we get a large period demand at retailer 2 and no demand at all at retailer 1. This means

that we would like to allocate up to S_2^* at retailer 2 in the next period. But this may not be possible due to insufficient supply at the warehouse. In that case we will get unequal inventory positions at the two retailers. However, it is rather obvious that it would be better to distribute the inventory positions equally. This might have been possible if we had saved some more stock at the warehouse in the preceding period, i.e., if we had not allocated $S_1^* = S_2^*$ to both retailers. Due to this ‘balance’ problem the decision rule that was optimal in the serial case is now only approximate. This balance problem and the outlined approximation method is first presented in [18]. The optimal solution of this problem requires heavy computation. The characterization of the optimal policy even for the simplest two-stage distribution system is still an open research question. The difficulty is due to possible stock imbalance among retailers as pointed out in [18] and [23] and given an illustration above. However, Federgruen and Zipkin [26] provided a lower bound on the minimum cost of the system by allowing a free inventory position rebalance among the retailers. Under such a relaxation, the original system reduces to a single-location system whose minimum cost can be easily computed. This minimum cost is a lower bound on the minimum cost of the original system. Also, some heuristic allocation policies (like myopic or cycle allocation policies) and various approximation methods (by relaxation or restriction) have been proposed in literature, see [26] for some approximation methods, Diks and de Kok [21] for a detailed computational study, and Federgruen [24] for a survey on centralized models of multi-echelon inventory systems under uncertainty including a review of various approximation methods for distribution systems.

2.3 Computational Methods

The mathematical assumptions under which (s, S) policies are optimal or close-to-optimal are satisfied by many practical inventory replenishment problems. Moreover, the rules of this type are easy to implement and require no more data than other standard techniques. However, the scientific methods for computing the optimal policy are often considered to be prohibitively expensive in practice. That is why there has been much research on developing optimal and heuristic methods of computing optimal (s, S) policies using as limited an amount of demand data as possible.

Also for most part, for more realistic models with multiple items, non-serial echelons, and nonlinear cost structure; the optimal policy structures have not been characterized and are very complex in nature. Hence even if the optimal strategies for these systems could be computed efficiently, the complexity of the optimal policy structure makes them unattractive for practical purposes. The implementation and proper execution of such a complex policy would be too expensive and difficult. This is another reason for the shift of focus towards the identification of close-to-optimal, but not necessarily fully optimal, policies with relatively simple structure like (s, S) policies, which are easy to compute and to implement. Also, accurate and easily computable approximations of the total system-wide cost have been

developed for use in design and parametric studies.

Most approximation approaches start with an exact formulation of the planning problem as a dynamic program or a Markov decision problem. The large dimension of the associate state and action spaces precludes, in general, exact solution of these dynamic programs. The exact model is therefore replaced by an approximate one through the application of one or more manipulations of the problem, like those used in large scale mathematical programming; relaxations, restrictions, projections, cost approximations.

These distinctions are important, because the properties of an approximation depend on the types and sequence of manipulations applied. If only relaxations are used, for example, then the resulting approximation is a lower bound on the true optimal cost of the problem. (We use the term relaxation in the general sense, that is any approximation of a minimization/maximization model which results in a lower/upper bound, e.g., expansions of the feasible set and/or replacement of the objective function by lower/upper bound functions.) This fact is very helpful in assessing optimality gaps for any heuristic strategy, since the cost of an appropriately constructed feasible strategy provides an upper bound on the optimal cost. (As is the case with most mathematical programming approximation methods, such as Lagrangian relaxation, the heuristic strategy is usually based on the solution of the approximate model.) If the difference between the upper and lower bounds is small, we can conclude both that the approximation is accurate and that the constructed policy is a good one.

Another approach is first to restrict the policy space to a more convenient and qualitatively appealing class. If determination of an optimal strategy within the chosen class is still intractable and the restriction is followed by one or more relaxations, the result is a lower bound, not on the original problem, but on the minimum cost among all policies within the class, so that optimality gaps may be assessed with respect to the chosen class of strategies only.

Optimal (s, S) policies may be computed by either successive approximations using functional equations, by policy iteration methods, or by Markovian methods. The approximate methods include the method of Roberts [62] and modifications of it developed by Wagner et al. [79] and Ehrhardt [22]; and the techniques of Porteus [58] and Freeland and Porteus [31], which are based on the general approach of Norman and White [56]. The optimal methods include Johnson [44], Veinott and Wagner [78], and Federgruen and Zipkin [27], which is the first optimal method that is proven to terminate in a finite number of iterations. There are numerous other approaches and algorithms proposed in the literature, see Porteus [59] for comparison of 17 such methods. He found that several methods seemed to perform very well. As he notes in his conclusion, stationary (s, S) policies are of only limited interest in practice since the distribution of demand is time varying in most real environments. Practitioners seem to favor a continuous review model and simply recompute the lot size and reorder point on a periodic basis as new estimates of the mean and standard deviation of the demand are made. These estimates are generally obtained using a forecasting tool such as simple exponential smoothing.

The algorithm presented in [27] computes an optimal (s, S) policy under standard assumptions; stationary data, well-behaved one-period costs, discrete demand, full backlogging and the long-run-average cost criterion. The overall strategy of the algorithm is policy iteration, modified to exploit an embedding technique, which is related to the renewal-theoretic approach and streamlines many of the computations. The linear systems that need to be solved are all triangular, and hence can be solved by simple substitution. Furthermore, the proposed technique also removes the need for truncation of the state space; so the algorithm is truly exact. Starting with a given (s, S) policy, the algorithm evaluates a sequence of policies, all of this form unlike standard policy iteration, and converges to an optimal one in a finite number of iterations. The policies generated are strictly improving, but not in the usual sense: Average costs decrease, but not necessarily strictly; cycling is precluded by strict improvement in a certain natural lexicographic criterion. In addition, a lower bound on the optimal cost can be computed in every iteration; the algorithm can thus be terminated with a suboptimal policy whose cost achieves any desired level of precision.

2.4 Setup Costs for Transportation

The assumption of linear costs for distribution in [18] is criticized for not considering the economies of scale in shipping costs. Hence Clark and Scarf [19] attempts to incorporate a setup cost associated with the transportation of items between adjacent installations. The main result of [18] was the decomposition of the total system value function into two (or multiple in general) value functions, one for each isolated echelon.

This decomposition enabled a solution procedure which begins by finding the optimal policy for the problem of the first installation without considering the costs of the second installation and neglecting the order capacity induced by second installation's inventory level. Since there was no setup cost for the transportation of items to installation 1, a base stock policy solves this problem, where the echelon inventory position (the on-hand inventory plus orders in transit to installation 1) is increased to S_n if it is below this critical number and no order is placed otherwise. If the echelon inventory level in installation 2 is above S_n then the order is filled completely; if it is below S_n then it is filled as much as possible since the cost function of installation 1 is convex. This specifies part of the system-wide optimal policy.

There remains the problem of determining the appropriate purchasing decisions for echelon 2. If these decisions were to be made independently of their influence on the lower level, insufficient stock would be procured. Hence the 'natural' echelon holding and shortage costs for echelon 2 are augmented by an additional shortage cost function that penalizes echelon 2 for its inability to deliver the required amount of stock to the lower level. This 'induced penalty cost' is the expected increment in total cost of the lower installation caused by the shortage of items in the higher installation, which is convex and depends only on the second stage echelon inventory level (being independent of the echelon inventory position of stage

one given the echelon inventory level of stage two). An (s_n, S_n) type policy, therefore, solves the procurement problem of the second installation which is the remaining part of the system-wide optimal policy.

This intuitive argument summarizes the procedure for the determination of optimal policies when there is no setup cost for the shipment. It is therefore appropriate to ask, still on the intuitive level, for the part played by the assumption of no setup cost in this policy. First of all, the lack of a setup cost was responsible for the simple description of optimal policies at the lower level in terms of a sequence of single critical numbers, S_n . If a setup cost in transportation were included in the problem, the optimal policy would no longer be of this simple form. Instead, the optimal policies would be of the (s_n, S_n) type with a pair of critical numbers relevant for each period.

Is it possible, in this case, to assign an additional shortage cost to echelon 2, as a function of the echelon inventory level of stage 2 alone? This is the crucial point in the simplification described above, and we must see if this simplification is still possible when a setup cost is introduced.

Now let IP_1 be the echelon inventory position at stage 1 and IL_2 be the echelon inventory level at stage 2 at the beginning of period n before ordering decisions are made. Suppose that $IP_1 > s_n$. In this case no ordering at stage 1 is required, and it might seem reasonable to charge no additional shortage cost even if $S_n > IL_2$. On the other hand, if $IP_1 \leq s_n$, the optimal policy would seem to request a shipment of size $S_n - IP_1$ from stage 2. However, if $IL_2 < S_n$, it is impossible to meet this request, and it would seem reasonable to charge an additional shortage cost. Hence we are led to the conclusion that the appropriate shortage cost to be added when $IL_2 < S_n$ seems to depend on whether $IP_1 \leq s_n$ or $IP_1 > s_n$, and is therefore, not a function of IL_2 alone. This conclusion precludes the decomposition of the total system cost function into separate echelon cost functions, and hence, obtaining a simple form for the optimal system policy with a setup in transportation. In fact, the characterization of the optimal policy structure in this case is still an open research question today after more than 40 years of progress in multi-echelon stochastic inventory theory.

Although the conclusion is that there is no simple form optimal policy, [19] makes use of a substantial amount of the preceding argument to propose an approximation method. Instead of attempting to find the correct additional shortage cost, they find upper and lower bound functions of IL_2 alone. The upper bound is obtained by charging an additional shortage cost whenever $IL_2 < S_n$, regardless of the size of IP_1 . On the other hand, the lower bound is obtained by charging an additional cost only when $IL_2 < s_n$. These upper and lower cost functions enables a simple procedure (as in the case where there are no setup costs in transportation), based on the calculations of functions of one variable, for bounding the true optimal cost function. Moreover, there is a natural specific policy associated with the computation of upper bound functions proposed in the paper, which if adopted guarantees the cost incurred will always be less than the upper bound of the true optimal cost function. Therefore, if the upper and lower bounds are close, as the authors claim that the computation of several examples has suggested, then we have not only a good estimate for the true optimal

cost, but also a simple policy whose cost is very close to the true optimal cost.

Hochstädler [41] extends this approach to a two-stage distribution system. He assumes that all the retailers use (s, S) policies that minimize their respective cost functions ignoring the system wide total cost. Under this assumption he provides upper and lower bounds on the total system cost and shows that the difference between these upper and lower bounds is less than a fixed number, which depends on the mean of period demand variables, period holding and shortage costs.

Chen and Zheng [15] obtain lower bounds on the minimum costs of managing certain production-distribution networks with setup costs at all stages and stochastic demands. These networks include serial, assembly, and one-warehouse multi-retailer systems. Cost allocation is one of the well known lower-bounding methodologies (see, e.g., [26]), novel cost allocation schemes are proposed in [15]. For general systems, new lower bounds are generated by combining cost allocation with “physical decomposition”. In particular, imagine that the product consists of a number of fictitious components. Each component is supplied/produced through a subsystem – a part of the original system – and is allocated part of the costs. By assuming that the components can be replenished and sold separately (i.e., physical decomposition), the original system decomposes to a number of independent systems, one for each component. The sum of the minimum costs of these independent systems is a lower bound on the minimum cost of the original system.

2.5 (R, nQ) Policies

Characterizing the optimal policy for the serial multi-echelon stochastic inventory problem with setup costs for transportation turns out to be extremely difficult, as [19] envisioned correctly. It is known that the optimal policy for this problem does not have a simple structure but a very complex one. Thus, an optimal policy, even if it could be identified, would not be easy to implement. In other words the “optimal policy” is no longer optimal or even attractive once the managerial effort of implementation is taken into account.

Prohibitive complexity of the optimal policy in this setting brings about a shift of focus in literature to approximate methods, performance evaluation of simple policies, and the optimal policies when the action space is restricted. The existence of setup costs for distribution in the system suggests that replenishment of inventories should be carried out in batches. If we ignore the setup costs but insist that every stage order in fixed quantities, then we have a new formulation of the problem in a restricted policy space. Although the fixed order quantities may not be a substitute for the setup costs, they can accommodate aspects that are not captured in the setup costs, e.g., the convenience of standardized shipments. Moreover, as we will see in this section, this reformulation makes the problem of finding an optimal policy tractable.

The (R, nQ) policies potentially offer a simple and cost-effective approach to this problem. An (R, nQ) policy operates as follows: whenever the inventory position at a stage is at or

below R , order nQ units where n is the minimum integer required to increase the inventory position above R . Q is called the *base order quantity* and R is called the *reorder point*.

In (R, nQ) multi-echelon models, almost exclusively, it is assumed that the base order quantities at the different stages satisfy an *integer-ratio constraint* where the base order quantity of an upstream stage is always a positive integer multiple of the base order quantity of the succeeding downstream stage. This further restriction of the policy space can be justified by the studies in deterministic counterparts where Roundy [65], [66] show that the so-called power-of-two policies are very close to the optimal solution. Under the power-of-two structure, the reorder quantities at all stages are restricted to be power-of-two multiples of a base quantity. This facilitates the quantity coordination among the different stages. The effectiveness of the power-of-two policies mainly arises from the insensitivity of the cost function in the EOQ (Economic Order Quantity) model. In fact similar insensitivity results hold for single-stage (R, nQ) systems. Zheng [84] provides the following bound:

$$\frac{C(\alpha Q^*)}{C(Q^*)} \leq \frac{1}{2} \left(\alpha + \frac{1}{\alpha} \right)$$

where $C(\cdot)$ represents the optimal cost function and Q^* is the optimal ordering quantity. Notice that this inequality becomes an equality for the EOQ model (see, e.g., [86]), suggesting that the (R, nQ) model is even more robust than the EOQ model with respect to the lot size. This result is very useful for multi-echelon, multi-location systems where quantity coordination is a primary concern.

Two variations of (R, nQ) policy with different informational requirements have been considered in literature. These are *echelon-stock* (R, nQ) policy where each stage uses an (R, nQ) policy to control its echelon stock, and *installation-stock* (R, nQ) policy where each stage uses an (R, nQ) policy to control its installation stock (i.e., its local inventory position). Both types of policies are easy to implement. Installation-stock policies require only local inventory information, while echelon-stock policies require centralized demand information. The relative cost difference between the two policies is called the value of centralized demand information, which is a fundamental issue in supply chain management.

Chen [11] is the first to study this issue under (R, nQ) policies in serial inventory systems. A key result of this paper is that the optimal echelon reorder points can be determined sequentially: first for stage 1, then for stage 2, and so on. This is based on an observation that the steady-state echelon inventory position at each stage can be replicated by the steady-state inventory position of a standard single-location point/reorder quantity (or (R, nQ)) model with a *random reorder point*. These random reorder points at different stages satisfy a simple recursive equation that is also found in Clark-Scarf model [18] with base-stock policies. Hence, the result in a nutshell is that after a proper transformation, the batch-transfer model can be treated as a base-stock model for the purpose of determining the optimal echelon reorder points. However, they could not propose such an exact method for the determination of optimal installation stock reorder points. Instead, they establish easy-to-compute bounds and suggest that optimal reorder points can be found by a search for

the problems without many stages. They also provide a heuristic method for computing installation stock reorder points which they claim gives good results. They also present an extensive numerical study on the value of centralized demand information. In a pool of 1,536 examples, it is found that the value of information has a fairly wide range with the highest value of 9% and a mean of 1.75%. The value of information tends to increase as a result of increases in the number of stages, the leadtimes, or the batch sizes. Interestingly, the higher demand variability decreases the value of information, and extreme levels of customer service (either high or low) tend to increase the value.

Because of its modest informational requirements, installation-stock policies have received more attention particularly in multi-echelon (R, Q) policy models. An (R, nQ) policy reduces to an (R, Q) policy when demand is for a single unit at a time in continuous-review systems. Note that although the initial measurement of a stage's echelon stock requires the inventory information at every downstream stage, its update only requires the demand information at the point of sales, which is readily available for most companies with advanced communication networks. Hence this advantage of installation-stock policy is quickly disappearing as more and more companies are equipped with advanced information technologies.

Axsäter and Rosling [5] compare installation and echelon stock (R, Q) policies in multi-stage inventory systems. They show that in serial systems, installation stock reorder policies are a subset of echelon stock reorder policies and that *nested* echelon stock reorder policies are a subset of installation stock reorder policies. A policy in a serial system is called *nested*, if an order at an upstream stage implies that all the downstream stages have ordered at the same time. Every installation stock policy is nested, but echelon stock policies are not necessarily nested. The result directly implies that in a serial system any given installation stock policy can be replaced by an equivalent (which means given any sequence of demand quantities and any initial state, the sample path of the system state would be same under both policies) echelon stock policy, but not the other way around. Hence, echelon stock policies represent a larger portion of feasible policies. This makes sense intuitively since each stage under an echelon stock policy uses inclusively more information than it would do under an installation stock policy.

Hadley and Whitin [38] showed that the steady state distribution of the inventory position in a single-stage system under (R, nQ) policy is uniform under some mild assumptions on the demand distribution. Let $IP(t)$ be the inventory position at the beginning of period t after order placement and before demand occurrence. Thus $R + 1 \leq IP(t) \leq R + Q$. [38] proves that if the Markov chain $\{IP(t)\}$ is irreducible, then its steady state distribution is uniform over $R + 1, \dots, R + Q$. This Markov chain is, indeed, irreducible for most demand distributions, e.g., any demand distribution where the one-period demand equals one with a positive probability satisfies this. Even this mild assumption can be generalized to the case when this Markov chain is reducible. In that case its steady state distribution is uniform over $r + \Delta, r + 2\Delta, \dots, r + q\Delta$ where r, Δ , and q are integers with $\Delta > 1, R - \Delta < r \leq R$, and q being the largest integer so that $r + q\Delta \leq R + Q$. The value of r is determined by the initial inventory position. See Chen [12] for details.

There are two alternative assumptions on the fixed setup costs in (R, nQ) policy models. Most of the papers that we have seen, although there are exceptions, and that have important results assume that a fixed setup cost is incurred for each Q units ordered. Therefore, for example, the set up cost for an order of $2Q$ units is $2K$. The alternative assumption is that a fixed cost is incurred for each order, independent of its size. Also a common approach in these models, especially in multi-stage models, is to consider infinite horizon problems under average cost criterion because the mathematical analysis is simplified by using steady state arguments in stationary models.

Given $IP(t) = y$, let $G(y)$ represent one-period expected holding and shortage cost, which is a convex function when the holding and shortage costs are convex. If μ is the mean of the demand then the long-run average cost of an (R, nQ) policy can be expressed as

$$C(R, Q) \equiv \frac{\mu K}{Q} + \frac{\sum_{y=R+1}^{R+Q}}{Q}$$

where the first term is the long-run average setup cost, and the second term is the long-run average holding and shortage cost. This expression is jointly convex in R and Q , hence it can be easily minimized by simple algorithmic methods. However, if the setup cost is fixed for each order independent of its size, $C(R, Q)$ is no longer jointly convex in decision variables and more complex algorithms are needed to determine the optimal control parameters. See Wagner et al. [79] for some exact and approximate methods for computing optimal (R, nQ) policies in single-stage models.

Chen and Zheng [14] are the first to adapt the single-location (R, nQ) policy to a multi-echelon system. They provide a recursive procedure to compute the steady state echelon inventory levels, which can be used to evaluate the long-run average holding and backorder costs as well as other performance measures. The procedure is based upon a key simple observation of a relationship between the inventory status of adjacent stages in a serial system under echelon stock (R, nQ) policy. This relationship uniquely determines the echelon inventory position of an adjacent downstream stage given the inventory level of the adjacent upstream stage at any period. Carrying this relationship to steady state, they obtain expressions for system performance measures such as average on-hand inventories all stages, average inventories in transit, and average customer backorders. In their model they assume fixed cost for each order independent of its size. They provide a numerical study which suggests that the optimal echelon stock (R, nQ) policies are close to optimal in most cases.

Veinott [76] shows that the (R, nQ) policy in a single-stage inventory problem is optimal among the set of policies in which the ordering quantities are restricted to be an integer multiple of Q . Chen [12] extends this result to serial systems, i.e., he demonstrates that an echelon-stock (R, nQ) policy is optimal among the policies in which order quantities at each stage satisfy integer-ratio constraints. Both of these optimality results are for the case where a fixed order cost is charged for every Q units. The proofs presented in these papers are similar. Firstly, a lower bound on the one-period cost is established, and then it is shown

that (R, nQ) policy actually achieves this bound. Notice that if we set $Q = 1$ for every stage then an (R, nQ) policy becomes a base-stock policy. Hence these optimality results can be seen as merely extensions of base-stock optimality results. Equivalently, (R, nQ) policies can be seen as generalized base-stock policies, i.e, each stage orders every period to keep its echelon stock within an interval of its base order quantity.

2.6 Capacity Constraints

It is well known that the inventory serves as a hedge against demand variability, and advance demand information and inventory can be interchanged. The basic reason for demand uncertainty is leadtime of production or transportation. If replenishment were instantaneous, there would be no need to hold inventory. But the leadtime itself is generally not modeled in detail; rather, it is commonly represented as a fixed interval or, when it is taken to be stochastic, independent of the rest of the model. In some settings, e.g., when the leadtime is primarily due to transportation delays, a simple model such as this may be a fair representation of reality. But leadtimes can arise not only from external factors like transportation, but also from congestion effects internal to the operation of a system. In particular, when the limits on production capacity is significant, the primary delay in replenishing stock may be due to backlogs in production created by the replenishment orders themselves, rather than to any external mechanism. Hence, for the efficient operation of these systems, explicit consideration of the production capacity is inevitable.

The management of an inventory under finite capacity and the mathematical analysis of the problem are complicated by the same phenomenon: In the uncapacitated case the effect of a large demand in some period can be corrected immediately in the next period. However when there is a capacity constraint, several periods of full production may be required, during which further large demands might occur, requiring still more time to return to a normal stock level. Indeed, it is not obvious that the system is stable enough under any policy to have a finite average cost. A necessary condition is that the production capacity must be larger than the mean of the demand distribution. The possible buildup of backorders, just described, in a single-stage inventory system with production capacity is reminiscent of the behavior of a queue. Indeed, using a standard transformation, this problem can be shown to be equivalent to a certain continuous-time queuing-control model.

Federgruen and Zipkin [29] and [30] consider the most basic single-stage inventory model with limited production capacity for infinite horizon average cost and discounted cost criteria, respectively. [30] also considers the finite horizon discounted problem. They show that, for all the problems considered, a modified base-stock policy, characterized by a single critical number, is optimal: Follow a base-stock policy when possible; when the prescribed production quantity exceeds the capacity, produce to capacity. While the finite production capacity complicates the analysis of the problem in [29], they observe that the set of feasible actions in each state is compact, which permits them to invoke results from Federgruen

et al. [25] for denumerable-state average-cost dynamic programs. This is also the reason why they assume discrete demand in [29] while assuming continuous demand in [30]. They adopt a different approach, based on the limiting behavior of the sequence of finite horizon problems, for the analysis of the problem in [30]. This is a relatively standard approach for uncapacitated problems (see, e.g., [43]). This approach allows them to show also that the base-stock level and optimal cost function are, respectively, the limits of their finite horizon counterparts.

It might be reasonable to expect that if the basic capacity constrained single-stage inventory model in [29] and [30] were to be extended with a fixed setup cost for production, a modified (s, S) policy would be optimal: Follow an (s, S) policy when possible; when the prescribed production quantity exceeds the capacity, produce to capacity. However, first Wijngaard [81] then Shaoxiang and Lambrecht [73] gave examples of such finite horizon problems having a more complex optimal policy. [73] also show that the optimal policy does exhibit a systematic pattern of what they called $X - Y$ band: When the inventory level drops below X , order up to capacity; when the inventory level is above Y , do nothing; if the inventory level is between X and Y , however, the ordering pattern is not defined.

Gallego and Scheller-Wolf [33] further divide the state space between the X and Y values into two and partially characterizes the optimal policy structure in this region: In one of these regions (between X and Y) it is optimal for the decision maker to either order nothing, or to bring the inventory at least up to a specified level, s' . In the other region the parameters of the solution dictate one of the two cases hold. In the first case it is optimal to order, again at least up to a specified level. In the second, the optimal policy is to either order the full capacity or nothing. To facilitate their analysis, they define what they call CK -convexity, which is very closely related to the K -convexity of Scarf [68]. In particular, it is a relaxation of K -convexity on a real interval where K -convexity corresponds to ∞K -convexity and a C_1 -convex function is also C_2 -convex for $\forall C_2 \leq C_1$. They complement their findings with a computational study, which, they believe, suggests that a still further characterization of the optimal policy exists.

Shaoxiang [72] extends [73] to infinite horizon discounted cost criteria. It is proven that the limiting cost function exists, and there exists stationary policies that are optimal in the long-run. The optimal policy is not of the modified (s, S) type in general, but continues to exhibit the $X - Y$ band structure. (C, K) -convexity (which is different from CK -convexity of [33], a restriction of strong CK -convexity in particular) is defined by the author, and then this property is used to show that the length of the $X - Y$ band is not larger than the value of the capacity. By exploring the $X - Y$ band structure, a linear program model is proposed to find the optimal policy in that band. Lastly, a numerical study is presented which indicates that the “best” modified (s, S) policy may perform poorly (more than 11% deviation from the optimal in cost performance).

Kapuściński and Tayur [46] consider the basic single-stage, discrete time production-inventory model where the stochastic demands follow a periodic pattern. For three cases; finite horizon cost, discounted infinite horizon cost, and infinite horizon average cost, they

show that a modified base-stock policy (as in e.g., [29]) with a set of critical numbers (one for each period in the cycle). This extends the results of [49], [85] for uncapacitated, nonstationary model and [29], [30] for capacity constrained, stationary model. They offer a simulation based method using infinitesimal perturbation analysis (IPA) to compute the set of optimal critical numbers, which completely characterize the optimal policy. They also provide an extensive numerical study, indicating that their IPA method is robust and fast, and treating some issues related to managing these systems.

Gallego and Hu [32] analyzed a discrete-time, single-item, single-location, periodic-review production/inventory system with finite production capacity where the demand and supply processes are driven by two independent, discrete-time, finite-state, time-homogeneous Markov chains. We have summarized some important papers (see, [8],[13],[54]) that used Markov modulated demand models to capture the influence on demand by factors such as seasonality, economic conditions, and product age. In this paper, a similar argument is made to model the supply process. In many instances, not only is the production/inventory capacity finite, but the system is also subject to random production yields that are influenced by factors such as breakdowns, repairs, maintenance, learning, and the introduction of new technologies. An example to motivate this assumption given in the paper is the yields in semiconductor industry: Production yields increase as the manufacturing process is fine tuned and drop again when new, more complex products are introduced to replace older products.

Under the above assumptions on the demand and supply processes and some other mild conditions, for both the finite and infinite horizon discounted cost criteria problems, they show that, given the demand state and the yield state, the optimal producing/ordering policy is a modified state-dependent *inflated base-stock policy*, which means that the optimal production/ordering quantity for each period is decreasing with respect to the initial level as well as the optimal order-up-to level. The term inflated base stock policy was coined by Zipkin, (see [86] p.392). They also demonstrate that the finite horizon solution converges to the infinite horizon solution.

There has been very little research on a multi-echelon system with limited capacity at each echelon. Given the difficulty of finding optimal policies for general multi-echelon systems with capacity constraints, it makes sense to restrict attention to a specific class of operating rules. Base-stock policies are attractive because they are simple and are known to be optimal in certain settings.

Glasserman and Tayur [35] analyze the stability of a serial multi-echelon model in which every stage has capacity constraints and follows a modified base-stock policy (modified because of capacity constraints). When capacity limits are introduced, as mentioned earlier, the stability of the system becomes an issue. Speaking loosely, the system is stable if, on average, it can produce finished goods at a greater rate than they are demanded. For the general stationary demands, they show that if the mean demand per period is smaller than the capacity at every stage, then inventories and backlogs are stable, having a unique stationary distribution to which they converge from all initial states. Under independent and

identically distributed demands, they show that the state of the system constitutes a *Harris ergodic* Markov chain, and thus inherits the wide-sense regenerative structure of that class of processes. While Harris recurrence ensures the existence of wide-sense regeneration times, it does not provide a means of identifying these times. Explicit regeneration times are not needed for convergence results, but they are useful in, for example, computing confidence intervals for simulation estimators. Hence, under an additional mild assumption on the demand distribution, they show that the system is regenerative in the classical sense (that is inventories return to their target levels infinitely often, with probability one) and identify explicit regeneration times. Extensions to systems with random leadtimes and periodic demands are also considered.

In their model, the state of the system is represented through *echelon shortfalls*, which is defined as the amount on order that has not yet been produced because of the capacity constraint at an echelon. In this setting, an echelon's shortfall equals to the difference between echelon base stock level and echelon inventory position. The similarities between the capacity constrained inventory problems and some queuing systems were mentioned before. Actually, the shortfall concept is adopted from queuing theory. The shortfalls satisfy a recursive equation a similar of which arises in the study of a D/G/1 queue. The techniques for analyzing this shortfall process are well established in queuing theory. These techniques are used to analyze the steady state distribution of the shortfalls, which directly corresponds to the steady state distribution of inventory levels in the system.

Mathematical models are good for simple situations and to grasp concepts. To compute numbers for real world situations, simulation is preferred. The following study suggests a method to get the best out of two approaches: Firstly, solve a tractable approximation of the real system to obtain a candidate policy; then evaluate the performance of the policy in a simulation of the real system; finally, experiment with the simulation to improve the policy. Glasserman and Tayur [36] consider general multi-echelon inventory models in which each stage has production capacity and operates under a modified echelon base-stock policy. They develop simulation based methodologies for estimating sensitivities of system costs with respect to policy parameters. These sensitivity estimates are, then, used in adjusting optimal parameters approximated by a simplified model to complexities that can be incorporated in a simulation.

They observe that, under a base-stock policy, inventories are continuous functions of base-stock levels, which enables them to use IPA derivative estimates. They note that continuity and a bound on derivatives wherever they exist are the essential conditions for the interchange of derivative and expectation required for IPA to yield unbiased estimates. As a result, they show that these estimates converge to the correct value for finite horizon and infinite horizon discounted and average cost criteria. Their numerical experiments suggest that this convergence is quick. An illustration of the whole method to a real life problem is presented at the end.

Glasserman and Tayur [37] consider the same problem as [36], and develop a simple approximation method for it. Objective is to find echelon base-stock levels that approxi-

mately minimize holding and backorder costs in the system. The key step in their procedure approximates the distribution of echelon inventory by a sum of exponential variables such that the arrival rates of the exponential variables are chosen to match asymptotically exact expressions. The computational requirements of this method are minimal. They also provide a numerical study, where they show that; in a test bed of 72 problems, each with five production stages, the average relative error for the approximate method is 1.9

Parker and Kapuściński [57] consider a capacity constrained two-echelon inventory system, where the constraining capacity (the smallest capacity) is at the lower installation. They also limit the leadtime at the higher installation to one period while permitting general (integer multiple of period length) leadtime leading to the lower installation. Under these assumptions, they show that a simple modification of the echelon base-stock policy is optimal. The policy for lower echelon is unchanged - it orders up to a specific target (subject to the availability from the higher echelon). The policy for the higher echelon is modified. The higher echelon orders up to a specific echelon target, taking care not to exceed a specific installation inventory.

This result is based on a few observations. Firstly, they demonstrate that it will never be optimal for the higher installation to hold more inventory than can be processed in a single period by the lowest stage, which is a bottleneck. Secondly, they call the states restricted by this property “the inventory band”, and argue that given any initial state, the system will get in the inventory band eventually and stay in it afterwards. These are simple and intuitive arguments, which immediately imply that a conventional echelon base-stock policy cannot be optimal. According to the conventional policy, a huge spike of demand would generate the same-size order at the higher installation, which may exceed the capacity constraint, thus generating unnecessary holding costs. Then, using this band argument, they substitute the constraints upon production by a constraint on inventory and decompose the system cost function into echelon cost functions when the system is in the band. The resulting modified base-stock policy differs from the conventional uncapacitated echelon base-stock policy only in that it constrains the system to operate within the inventory band. Lastly, they extend the analysis into Markov modulated demand and infinite horizon case under both average and discounted cost criteria.

Chapter 3

Make-to-Order Production

In this chapter, we introduce our base case model of a just-in time manufacturer utilizing an outsourced logistic agreement for shipping to a retail site. We focus our attention on periodic fixed commitment contracts, which specify in advance the volume to be reserved and the frequency to be shipped. Section 3.1 presents the model in general form. In Section 3.2 and Section 3.3, we analyze two special cases of this problem, which facilitate understanding of the fundamental concepts and lay the groundwork for proving results regarding the structure of the optimal policy function, its basic properties and a decomposition result for the general case we analyze in Section 3.4.

3.1 The Model

In this chapter, we consider a just-in-time manufacturing firm that must ship its product to some remote location to meet stochastic demand, $w_t \geq 0$, through a finite planning horizon. At the beginning of the planning horizon, we assume that the firm is already in possession of a transportation contract. Transportation contracts often specify in advance the frequency and volume to be reserved by the logistics provider (Yano and Gerchak [83]). Thus, a fixed commitment contract in a finite horizon model can be specified by its shipment periods and its reserved capacities for these periods. We assume that the contract in hand has n shipments at periods T_1, \dots, T_n with reserved capacities C_1, \dots, C_n respectively. When the demand is uncertain, the transportation contract alone usually will not provide sufficient service levels. For this reason, in our model we utilize a spot market for shipping that provides expedited service (which is often the case in practice). At each period, an order can be shipped immediately through expedited shipment, which costs c per unit, or it can be delayed to utilize the contracted capacity in which case a waiting cost of p per unit applies every period. The discount factor is denoted by $\alpha \in (0, 1)$. We assume linear production costs with no lead time or capacities, which simplifies the production decision given the shipping decision, as all production will immediately precede shipment in quantities equal to

the shipping quantities. Thus, to simplify subsequent notation, we do not explicitly model production in what follows. The decision problem is to find the optimal level of shipment u_t at each period given the number of pending orders x_t . The demand distributions for periods 0 through $T_n - 1$ are assumed to be mutually independent. At the end of the contract period, T_n , pending orders in excess of the contracted capacity, C_n , are shipped via expedited shipment service. We use the notation $(x)^+$ to represent $\max(0, x)$ throughout this chapter. The dynamic programming equations for this model follows.

$$J_{T_n}(x_{T_n}) = H_{T_n}(x_{T_n}) \tag{P}$$

$$J_t(x_t) = \min_{u_t \in [0, x_t]} \left\{ H_t(u_t) + p(x_t - u_t) + \alpha \mathbb{E}_{w_t} [J_{t+1}(x_t - u_t + w_t)] \right\}, \quad t = 0, \dots, T_n - 1,$$

where $H_t(x) \equiv \begin{cases} c(x - C_t)^+, & \text{if } t \in \{T_1, \dots, T_n\} \\ cx, & \text{otherwise.} \end{cases}$

In the next two sections, we will analyze two special cases of this model, which will facilitate characterizing the optimal policy for the general case at the last section as well as providing insight into the nature of the problem.

3.2 One Shipment without Capacity

In this section, we analyze the case where there is only one shipment period in the contract which has infinite capacity. In other words, this is a special case of the general model, where $n = 1$, $T_1 = T$, and $C_1 = C = \infty$. The dynamic programming equations simplify as follows.

$$J_T(x_T) = 0 \tag{P1}$$

$$J_t(x_t) = \min_{u_t \in [0, x_t]} \left\{ cu_t + p(x_t - u_t) + \alpha \mathbb{E}_{w_t} [J_{t+1}(x_t - u_t + w_t)] \right\} \quad t = 0, \dots, T - 1.$$

The optimal policy is given by the following intuitive result, which says that it is optimal to ship every order immediately until we get sufficiently close to the end of the contract period, and after that it is optimal not to ship at all until the end of the contract period.

Proposition 3.2.1. *(Optimal Policy of (P1)) In (P1), the optimal shipping quantity function, μ_t , is given by:*

$$\mu_t(x) = \begin{cases} x, & \text{if } p \frac{1 - \alpha^{T-t}}{1 - \alpha} > c; \\ 0, & \text{otherwise} \end{cases} \quad t = 0, \dots, T - 1.$$

Proof. We will make this proof by an induction argument. At period $T - 1$ we have

$$J_{T-1}(x) = \min_{u \in [0, x]} \{cu + p(x - u) + \alpha \mathbb{E}[J_T(x - u + w)]\}$$

$$= \min_{u \in [0, x]} \{u(c - p)\} + px.$$

This implies

$$\mu_{T-1}(x) = \begin{cases} x, & \text{if } p > c; \\ 0, & \text{otherwise.} \end{cases}$$

Hence, the asserted result is satisfied for period $T - 1$. Now, suppose that the result holds for $t = k + 1, \dots, T - 1$; where $k \in \{0, \dots, T - 2\}$. Then, at period k we have

$$\begin{aligned} J_k(x) &= \min_{u \in [0, x]} \{cu + p(x - u) + \alpha \mathbb{E}[J_{k+1}(x - u + w)]\} \\ &= \begin{cases} \min_{u \in [0, x]} \{u(c - p) + \alpha \mathbb{E}[c(x - u + w) + \alpha J_{k+2}(w)]\} + px, & \text{if } p \frac{1 - \alpha^{T-k-1}}{1 - \alpha} > c; \\ \min_{u \in [0, x]} \left\{ u(c - p) + \sum_{n=1}^{T-k-1} \alpha^n p(x - u + n \mathbb{E}[w]) \right\} + px, & \text{otherwise} \end{cases} \\ &= \begin{cases} \min_{u \in [0, x]} \{u[(1 - \alpha)c - p]\} + \alpha \mathbb{E}[c(x + w) + \alpha J_{k+2}(w)] + px, & \text{if } p \frac{1 - \alpha^{T-k-1}}{1 - \alpha} > c; \\ \min_{u \in [0, x]} \left\{ u \left(c - p \sum_{n=0}^{T-k-1} \alpha^n \right) \right\} + \sum_{n=0}^{T-k-1} \alpha^n p(x + n \mathbb{E}[w]), & \text{otherwise.} \end{cases} \end{aligned}$$

This implies

$$\mu_k(x) = \begin{cases} x, & \text{if } p \frac{1 - \alpha^{T-k}}{1 - \alpha} > c; \\ 0, & \text{otherwise.} \end{cases}$$

Hence, the asserted result is satisfied for period k as well. This completes the induction and the proof. \square

3.3 One Shipment with Capacity

In this section, we continue analyzing the case with a single shipment period in the contract, but it now has a finite capacity. In other words, this is a special case of the general model, where $n = 1$, $T_1 = T$, and $C_1 = C < \infty$. The dynamic programming equations change as follows.

$$\begin{aligned} J_T(x_T) &= c(x_T - C)^+ & (\text{P2}) \\ J_t(x_t) &= \min_{u_t \in [0, x_t]} \left\{ cu_t + p(x_t - u_t) + \alpha \mathbb{E}_{w_t} [J_{t+1}(x_t - u_t + w_t)] \right\} & t = 0, \dots, T - 1 \end{aligned}$$

Notice that if $p \geq c$ then $\mu_t(x) = x$, or if $(1 - \alpha)c \geq p$ then $\mu_t(x) = 0$, $t = 0, \dots, T - 1$. Thus, from now on we assume that $p < c$ and $(1 - \alpha)c < p$ to avoid trivial cases. We also naturally assume that $x_0 \geq 0$.

Firstly, we present the following well known result (Heyman [40]) and provide a proof of it for completeness. This lemma will be useful in showing the preservation of convexity for the results that follow.

Lemma 3.3.1. (*Convexity Preservation under Minimization*) *If X is a convex set, $Y(x)$ is a nonempty set for every $x \in X$, the set $C = \{(x, y) \mid x \in X, y \in Y(x)\}$ is a convex set, $g(x, y)$ is a convex function on C ,*

$$f(x) = \inf_{y \in Y(x)} g(x, y),$$

and $f(x) > -\infty$ for every $x \in X$, then f is a convex function on X .

Proof. Let x and \bar{x} be arbitrary elements of X . Let $0 \leq \theta \leq 1$, and let $\bar{\theta} = 1 - \theta$. Select arbitrary $\delta > 0$. By the definition of f , there must exist $y \in Y(x)$ and $\bar{y} \in Y(\bar{x})$ such that $g(x, y) \leq f(x) + \delta$ and $g(\bar{x}, \bar{y}) \leq f(\bar{x}) + \delta$. Then,

$$\begin{aligned} \theta f(x) + \bar{\theta} f(\bar{x}) &\geq \theta g(x, y) + \bar{\theta} g(\bar{x}, \bar{y}) - \delta && \text{[properties of } y \text{ and } \bar{y}] \\ &\geq g(\theta x + \bar{\theta} \bar{x}, \theta y + \bar{\theta} \bar{y}) - \delta && \text{[convexity of } g \text{ on } C] \\ &\geq f(\theta x + \bar{\theta} \bar{x}) - \delta. && \text{[(}\theta x + \bar{\theta} \bar{x}, \theta y + \bar{\theta} \bar{y}) \in C] \end{aligned}$$

Because δ is arbitrary, the inequality must hold for $\delta = 0$. (Otherwise, a contradiction can be reached.) \square

The optimal policy is given by the following result.

Proposition 3.3.2. (*Optimal Policy of (P2)*) *In (P2), for a given C , there is a sequence of increasing numbers $\{R_t\}_{t=0}^{T-1}$ between 0 and C such that $\mu_t(x) = (x - R_t)^+$ for $t = 0, \dots, T - 1$.*

Proof. Firstly, notice that $x_0 \geq 0$ implies $x_t \geq 0$ for $t = 0, \dots, T$. Hence, we restrict the domain of the value functions to nonnegative real numbers, that is $J_t: \mathbb{R}^+ \rightarrow \mathbb{R}$, $t = 0, \dots, T$. Let us rewrite the Bellman equation by defining a new variable y , and a new function $G_t: \mathbb{R} \rightarrow \mathbb{R}$.

$$\begin{aligned} J_t(x) &= \min_{u \in [0, x]} \{cu + p(x - u) + \alpha \mathbb{E}_w [J_{t+1}(x - u + w)]\} \\ &= \min_{y \in [0, x]} \{-cy + py + \alpha \mathbb{E}_w [J_{t+1}(y + w)]\} + cx && \text{[where } y \equiv x - u] \\ &= \min_{y \in [0, x]} G_t(y) + cx, \end{aligned}$$

where $G_t(y) \equiv (p - c)y + \alpha \mathbb{E}_w [J_{t+1}(y + w)]$.

Suppose that for some $t = k \in \{0, \dots, T-1\}$ we have (i) $J_{t+1}(x)$ is convex in x ; (ii) $J'_{t+1}(x) \geq c$ if $x > R_{t+1}$, where R_t is the minimizer of G_t for $t = 0, \dots, T-1$, and $R_T \equiv C$.

Now, notice that (i) implies $G_t(y)$ is convex in y , since $J_{t+1}(y+w)$ is convex in y for all w ; on the other hand (ii) implies $G'_t(y) > 0$ when $y > R_{t+1}$, since $p + \alpha c > c$. Then $\exists R_t \in [0, R_{t+1}]$ such that $G_t(R_t) \leq G_t(y)$ for all $y \geq 0$. Since G_t is convex, this implies $\mu_t(x) = (x - R_t)^+$.

Putting $X = [0, \infty)$, $Y(x) = [0, x]$, $g(x, y) = G_t(y) + cx$, and $f(x) = J_t(x)$ in Lemma 3.3.1, we deduce that $J_t(x)$ is convex in x , in other words (i) holds for $t = k-1$. When $x > R_t$, $J_t(x) = G_t(R_t) + cx$, which means (ii) holds for $t = k-1$.

Since (i) and (ii) holds for $k = T-1$, we can repeat this argument sequentially for $k = T-1, T-2, \dots, 0$, which completes the proof. \square

These R_t values can be thought as “maximum levels of pending orders reserved by the future contracted shipments”, and the policy can be interpreted as a “ship-down-to” type policy analogous to a base-stock policy in reverse. Here, the aim is to keep the level of pending orders at or below the reserved levels, which is similar to a base-stock policy keeping the inventory levels at or above the order-up-to levels.

To understand how R_t changes with C , we will make the dependence of R_t to C explicit by defining it as a function of C for the next proposition.

Proposition 3.3.3. *In (P2), $R_t(C)$ is increasing in C , $t = 0, \dots, T-1$.*

Proof. Firstly, let us rewrite the Bellman equation while making the dependencies on C explicit.

$$\begin{aligned} J_t(x, C) &= \min_{u \in [0, x]} \{cu + p(x - u) + \alpha \mathbb{E}_w[J_{t+1}(x - u + w, C)]\} \\ &= \min_{y \in [0, x]} \{-cy + py + \alpha \mathbb{E}_w[J_{t+1}(y + w, C)]\} + cx \quad [\text{where } y \equiv x - u] \\ &= \min_{y \in [0, x]} G_t(y, C) + cx \end{aligned}$$

where $G_t(y, C) \equiv (p - c)y + \alpha \mathbb{E}_w[J_{t+1}(y + w, C)]$.

Next, we define

$$J_t^r(x, C) \equiv \lim_{\epsilon \downarrow 0} \frac{J_t(x + \epsilon, C) - J_t(x, C)}{\epsilon}$$

and similarly,

$$G_t^r(x, C) \equiv \lim_{\epsilon \downarrow 0} \frac{G_t(x + \epsilon, C) - G_t(x, C)}{\epsilon}.$$

Now, assume that $J_{t+1}^r(x, C)$ is decreasing in C for some $t = k \in \{0, \dots, T-1\}$. Then

$G_t^r(x, C)$ is decreasing in C by definition. This implies $R_t(C)$ is increasing in C , since $G_t(x, C)$ is convex in x .

Furthermore,

$$\begin{aligned} J_t^r(x, C) &= \lim_{\epsilon \downarrow 0} \frac{\min_{y \in [0, x+\epsilon]} G_t(y, C) - \min_{y \in [0, x]} G_t(y, C)}{\epsilon} + c \\ &= \begin{cases} c, & \text{if } R_t(C) \leq x \\ c + G_t^r(x, C), & \text{otherwise.} \end{cases} \end{aligned}$$

Hence $J_t^r(x, C)$ is decreasing in C , since $G_t^r(x, C)$ is decreasing in C .

Since $J_T^r(x, C)$ is decreasing in C , we can repeat this argument sequentially for $k = T - 1, T - 2, \dots, 0$, which completes the proof. \square

The next result, which gives a sufficient condition to expedite all orders, directly follows from Proposition 3.2.1 and Proposition 3.3.3.

Corollary 3.3.4. *In (P2), $R_t = 0$ for all t such that $p \frac{1-\alpha^{T-t}}{1-\alpha} > c$ is satisfied.*

Proof. Notice that (P2) \rightarrow (P1) as $C \rightarrow \infty$. Hence $R_t(\infty) = 0$ for all t such that $p \frac{1-\alpha^{T-t}}{1-\alpha} > c$ by Proposition 3.2.1. But by Proposition 3.3.3, $R_t = R_t(C) \leq R_t(\infty)$, $t = 0, \dots, T$. This completes the proof since $R_t \geq 0$ by definition. \square

Proposition 3.3.5. *In (P2), for all $R_t(C) > 0$ and $\delta \geq 0$; $R_t(C + \delta) = R_t(C) + \delta$, $t = 0, \dots, T - 1$.*

Proof. First, we rewrite the Bellman equation while making the dependencies on C explicit.

$$\begin{aligned} J_t(x, C) &= \min_{u \in [0, x]} \{cu + p(x - u) + \alpha \mathbb{E}_w[J_{t+1}(x - u + w, C)]\} \\ &= \min_{y \in [0, x]} \{-cy + py + \alpha \mathbb{E}_w[J_{t+1}(y + w, C)]\} + cx \quad [\text{where } y \equiv x - u] \\ &= \min_{y \in [0, x]} G_t(y, C) + cx, \end{aligned}$$

where $G_t(y, C) \equiv (p - c)y + \alpha \mathbb{E}_w[J_{t+1}(y + w, C)]$.

Second, suppose that for some $t = k \in \{1, \dots, T - 1\}$ we have (i) $G_t(y + \delta, C + \delta) = G_t(y, C) + L_t(\delta)$ for all $\delta \geq 0$, where $L_t: \mathbb{R}^+ \rightarrow \mathbb{R}$ is some real valued function of δ .

Then for any $R_t(C) > 0$ we have

$$R_t(C + \delta) = \operatorname{argmin}_{y \in [0, \infty)} G_t(y, C + \delta) \quad [\text{by definition}]$$




Figure 3.1: The shifting property of the ship-down-to levels.

$$\begin{aligned}
 &= \operatorname{argmin}_{y \in [0, \infty)} G_t(y - \delta, C) \quad [\text{by (i)}] \\
 &= \operatorname{argmin}_{z \in [-\delta, \infty)} G_t(z, C) + \delta \quad [z \equiv y - \delta] \\
 &= R_t(C) + \delta. \quad [\text{since } R_t(C) > 0]
 \end{aligned}$$

Also,

$$\begin{aligned}
 G_{t-1}(y + \delta, C + \delta) &= (p - c)(y + \delta) + \alpha \mathbb{E}_w [J_t(y + \delta + w, C + \delta)] \\
 &= (p - c)y + \alpha \mathbb{E}_w \left[\min_{z \in [0, y + \delta + w]} G_t(z, C + \delta) + c(y + \delta + w) \right] + \delta(p - c) \\
 &= (p - c)y + \delta(p - (1 - \alpha)c) \\
 &\quad + \alpha \mathbb{E}_w \left[\min_{z \in [0, y + \delta + w]} G_t(z - \delta, C) + c(y + w) + L_t(\delta) \right] \\
 &= (p - c)y + \delta(p - (1 - \alpha)c) \\
 &\quad + \alpha \mathbb{E}_w \left[\min_{z \in [0, y + \delta + w]} G_t(z - \delta, C) + c(y + w) \right] + \alpha L_t(\delta) \\
 &= (p - c)y + \alpha \mathbb{E}_w \left[\min_{x \in [0, y + w]} G_t(x, C) + c(y + w) \right] + L_{t-1}(\delta)
 \end{aligned}$$

$$\begin{aligned}
&= (p - c)y + \alpha \mathbb{E}_w[J_t(y + w, C)] + L_{t-1}(\delta) \\
&= G_{t-1}(y, C) + L_{t-1}(\delta),
\end{aligned}$$

where the third equality follows from (i), the fifth follows since $R_t(C) > 0$, and we define $L_{t-1}(\delta) \equiv \delta(p + \alpha c - c) + \alpha L_t(\delta)$. Hence (i) is satisfied for $t = k - 1$.

To see (i) is satisfied for $t = T - 1$, note that

$$\begin{aligned}
G_{T-1}(y + \delta, C + \delta) &= (p - c)(y + \delta) + \alpha \mathbb{E}_w[J_T(y + \delta + w, C + \delta)] \\
&= (p - c)y + \alpha \mathbb{E}_w[(y + \delta + w - C + \delta)^+] + \delta(p - c) \\
&= (p - c)y + \alpha \mathbb{E}_w[(y + w - C)^+] + \delta(p - c) \\
&= (p - c)y + \alpha \mathbb{E}_w[J_T(y + w, C)] + L_{T-1}(\delta) \\
&= G_{T-1}(y, C) + L_{T-1}(\delta).
\end{aligned}$$

Thus we can repeat this argument sequentially for $k = T - 1, T - 2, \dots, 0$, which completes the proof. \square

3.4 Multiple shipments with capacity

In this section, we will analyze the properties of the most general case as described in (P).

The following lemma, which gives the intuitive result that the value functions are increasing in pending orders, will be needed for the next theorem.

Lemma 3.4.1. *In (P), value functions, $J_t(x)$, are increasing in x for all t .*

Proof. Suppose $\underline{x}_k \leq \bar{x}_k$ for some $k \in \{0, \dots, T_n - 1\}$. We will show that $J_k(\underline{x}_k) \leq J_k(\bar{x}_k)$ by a coupling argument. In other words, we will argue that starting at a lower state is less costly than starting at a higher state even if we do not necessarily behave optimally but act like we started at the higher state.

Let $\pi_k = \{\mu_k, \dots, \mu_{T_n-1}\}$ be the optimal policy from period k for a given sample of (P). Consider the policy $\tilde{\pi}_k = \{\tilde{\mu}_k, \dots, \tilde{\mu}_{T_n-1}\}$, where $\tilde{\mu}_t(\underline{x}_t) = \min[\mu_t(\bar{x}_t), \underline{x}_t]$ for $t = k, \dots, T_n - 1$. Under $\tilde{\pi}$ it is obvious that $\underline{x}_t \leq \bar{x}_t$ and $\underline{u}_t \leq \bar{u}_t$ for $t = k, \dots, T_n - 1$, which also means $\underline{x}_T \leq \bar{x}_T$. Note that $g_t(x, u, w_t)$ is increasing in x and u for all t , and $J_{T_n}(x)$ is increasing in x . Then we have,

$$\begin{aligned}
J_k(\underline{x}_k) &= J_k^\pi(\underline{x}_k) \\
&\leq J_k^{\tilde{\pi}}(\underline{x}_k)
\end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E} \left[\sum_{t=k}^{T_n-1} g_t(\underline{x}_t, \underline{u}_t, w_t) + J_{T_n}(\underline{x}_t) \right] \\
 &\leq \mathbb{E} \left[\sum_{t=k}^{T_n-1} g_t(\bar{x}_t, \bar{u}_t, w_t) + J_{T_n}(\bar{x}_t) \right] \\
 &= J_k^\pi(\bar{x}_k) \\
 &= J_k(\bar{x}_k)
 \end{aligned}$$

as desired. \square

The next theorem completely characterizes the optimal policy structure of the general problem (P) introduced earlier in this chapter.

Theorem 3.4.2. (*Optimal Policy of (P)*) In (P), for a given set of C_m and T_m , $m = 1, \dots, n$, there is a sequence of increasing nonnegative numbers $\{\{R_t\}_{t=T_m+1}^{T_m+1}\}_{m=0}^{n-1}$ such that

$$\mu_t(x) = \begin{cases} (x_t - R_t)^+, & \text{if } t \notin \{T_1, \dots, T_n\}; \\ \begin{cases} x, & \text{if } x \leq C_t; \\ C_t + (x - R_t)^+, & \text{otherwise.} \end{cases} & \text{otherwise,} \end{cases}$$

where we define $T_0 \equiv -1$.

Proof. Suppose that for some $t = k \in \{0, \dots, T_n - 1\}$, we have (i) $J_{t+1}(x)$ is convex in x ; (ii) $J'_{t+1}(x) \geq c$ if $x > R_{t+1}$, where $R_{T_n} \equiv C_{T_n}$.

If $t \notin \{T_1, \dots, T_{n-1}\}$, then the claim is satisfied for $t = k$, and moreover (i), (ii) hold for $t = k - 1$ (see the proof of Proposition 3.3.2).

Otherwise we have

$$J_t(x) = \min_{u \in [0, x]} \{c(u - C_t)^+ + p(x - u) + \alpha \mathbb{E}_w[J_{t+1}(x - u + w)]\}.$$

Since $J_{t+1}(x)$ is increasing in x by Lemma 3.4.1, we deduce that $\mu_t(x) = x$ if $x \leq C_t$ and $\mu_t(x) \geq C_t$ if $x > C_t$. Then we can rewrite the term $x - u$ in the expression above as

$$\begin{aligned}
 x - u &= \begin{cases} 0, & \text{if } x \leq C_t \\ (x - C_t) - (u - C_t), & \text{otherwise} \end{cases} \\
 &= (x - C_t)^+ - (u - C_t)^+ \\
 &= (x - C_t)^+ - \tilde{u}
 \end{aligned}$$

where we define $\tilde{u} \equiv (u - C_t)^+$. Putting this back in $J_t(x)$ while making a change of variable to \tilde{u} , we get

$$J_t(x) = \min_{\tilde{u} \in [0, (x - C_t)^+]} \{c\tilde{u} + p((x - C_t)^+ - \tilde{u}) + \alpha \mathbb{E}_w[J_{t+1}((x - C_t)^+ - \tilde{u} + w)]\}$$

$$= \min_{y \in [0, (x - C_t)^+]} \{(p - c)y + \alpha \mathbb{E}_w[J_{t+1}(y + w)]\} + c(x - C_t)^+$$

where we make another change of variable to $y \equiv (x - C_t)^+ - \tilde{u}$,

$$= \min_{y \in [0, (x - C_t)^+]} G_t(y) + c(x - C_t)^+$$

where $G_t(y) \equiv (p - c)y + \alpha \mathbb{E}_w[J_{t+1}(y + w)]$.

Now, notice that (i) implies $G_t(y)$ is convex in y , since $J_{t+1}(y + w)$ is convex for all w ; on the other hand (ii) implies $G'_t(y) > 0$ when $y > R_{t+1}$, since $p + \alpha c > c$. Then $\exists \tilde{R}_t \in [0, R_{t+1}]$ such that $G_t(\tilde{R}_t) \leq G_t(y)$ for all $y \geq 0$. Since G_t is convex, this implies $\mu_t(x) = C_t + (x - C_t - \tilde{R}_t)^+$ if $x > C_t$. Now defining $R_t \equiv C_t + \tilde{R}_t$, we observe that the claim is satisfied for $t = k$.

Putting $X = [0, \infty)$, $Y(x) = [0, (x - C_t)^+]$, $g(x, y) = G_t(y) + c(x - C_t)^+$, and $f(x) = J_t(x)$ in Lemma 3.3.1, we deduce that $J_t(x)$ is convex in x , in other words (i) holds for $t = k - 1$. When $x > R_t$, $J_t(x) = G_t(\tilde{R}_t) + c(x - C_t)$, which means (ii) holds for $t = k - 1$.

We have shown that the claimed result is satisfied at $t = k$, and moreover (i), (ii) hold at $t = k - 1$. Since (i),(ii) hold for $t = T_n - 1$, we can repeat this argument sequentially for $k = T_n - 1, \dots, 0$, which completes the proof. \square

In words, we say that the optimal policy is a “modified ship-down-to” type policy, where it is ship-down-to type in periods with no scheduled shipment, and in periods with scheduled shipment, the standing orders over the capacity of shipment is ship-down-to type.

The next result is the general case analogue of Proposition 3.3.3.

Proposition 3.4.3. *In (P), $R_t(\mathbf{C})$ is increasing in \mathbf{C} for $t = 0, \dots, T_n - 1$, where the vector \mathbf{C} is defined by $\mathbf{C} \equiv (C_1, \dots, C_n)$.*

Proof. Suppose that $J_{t+1}^r(x, \mathbf{C})$ is decreasing in \mathbf{C} for some $t = k \in \{0, \dots, T_n - 1\}$, where $J_t^r(x, \mathbf{C})$ is defined as in the proof of Proposition 3.3.3.

If $t \notin \{T_1, \dots, T_{n-1}\}$, then $R_t(\mathbf{C})$ is increasing in \mathbf{C} , and moreover $J_t^r(x, \mathbf{C})$ is decreasing in \mathbf{C} (see the proof of Proposition 3.3.3).

Otherwise, proceeding as in the proof of Proposition 3.4.2 while making the dependencies on \mathbf{C} explicit,

$$\begin{aligned} J_t(x, \mathbf{C}) &= \min_{u \in [0, x]} \{c(u - C_t)^+ + p(x - u) + \alpha \mathbb{E}_w[J_{t+1}(x - u + w), \mathbf{C}]\} \\ &= \min_{\tilde{u} \in [0, (x - C_t)^+]} \{c\tilde{u} + p((x - C_t)^+ - \tilde{u}) + \alpha \mathbb{E}_w[J_{t+1}((x - C_t)^+ - \tilde{u} + w, \mathbf{C})]\} \\ &= \min_{y \in [0, (x - C_t)^+]} \{(p - c)y + \alpha \mathbb{E}_w[J_{t+1}(y + w, \mathbf{C})]\} + c(x - C_t)^+ \\ &= \min_{y \in [0, (x - C_t)^+]} G_t(y, \mathbf{C}) + c(x - C_t)^+ \end{aligned}$$

where $G_t(y, \mathbf{C}) \equiv (p - c)y + \alpha \mathbb{E}_w[J_{t+1}(y + w, \mathbf{C})]$.

Hence, $G_t^r(x, \mathbf{C})$ is decreasing in \mathbf{C} by definition. This implies $R_t(\mathbf{C})$ is increasing in \mathbf{C} , since $G_t(x, \mathbf{C})$ is convex in x .

Furthermore,

$$\begin{aligned} J_t^r(x, \mathbf{C}) &= \lim_{\epsilon \downarrow 0} \frac{\min_{y \in [0, (x + \epsilon - C_t)^+]} G_t(y, \mathbf{C}) - \min_{y \in [0, (x - C_t)^+]} G_t(y, \mathbf{C})}{\epsilon} \\ &\quad + c \lim_{\epsilon \downarrow 0} \frac{(x + \epsilon - C_t)^+ - (x - C_t)^+}{\epsilon} \\ &= \begin{cases} 0, & \text{if } x < C_t \\ c + G_t^r(x, \mathbf{C}), & \text{if } C_t \leq x < R_t(\mathbf{C}) \\ c, & \text{otherwise.} \end{cases} \end{aligned}$$

where $R_t(\mathbf{C}) \equiv C_t + \tilde{R}_t(\mathbf{C})$, and $\tilde{R}_t(\mathbf{C})$ is the minimizer of $G_t(y, \mathbf{C})$ as before. Hence $J_t^r(x, \mathbf{C})$ is decreasing in \mathbf{C} , since $G_t^r(x, \mathbf{C})$ is decreasing in \mathbf{C} .

We have shown that $R_t(\mathbf{C})$ is increasing in \mathbf{C} , and moreover $J_t^r(x, \mathbf{C})$ is decreasing in \mathbf{C} . Since $J_{T_n}^r(x, \mathbf{C})$ is decreasing in \mathbf{C} , we can repeat this argument sequentially for $k = T_n - 1, \dots, 0$, which completes the proof. \square

The next corollary follows from Proposition 3.2.1 and Proposition 3.4.3, and provides a sufficient condition for the decomposition of the problem in time.

Corollary 3.4.4. *(Decomposition of (P)) In (P), suppose that $p \frac{1 - \alpha^{T_{N(t)} - t}}{1 - \alpha} > c$ is satisfied for some t , where $N(t) = \min_{j \in \{1, \dots, n\}} (j | T_j > t)$. Then $R_t = 0$ and the problem decomposes at period $T_{P(t)}$, where $P(t) = \min_{j \in \{0, \dots, n\}} (j | T_j \leq t)$.*

Proof. Consider problem (P) as all the capacities gets arbitrary large, or as $C_m \rightarrow \infty$ for $m = 1, \dots, n$. Suppose that $p \frac{1 - \alpha^{T(t) - t}}{1 - \alpha} > c$ is satisfied for some t . Let $\pi_t^{T(t) - 1} \equiv \{\mu_t, \dots, \mu_{T(t) - 1}\}$ be the optimal policy from period t to $T(t) - 1$. Obviously, $\pi_t^{T(t) - 1}$ is independent of the capacities before period t , or C_k where $C_k < t$. Also observe that $\pi_t^{T(t) - 1}$ is independent of C_k where $k > T(t)$, since $J_{T(t)}(x) = \alpha J_{T(t)+1}(w)$ is a constant independent of the actions, or $\pi_t^{T(t) - 1}$. Hence finding $\pi_t^{T(t) - 1}$ is very similar to a $T(t) - t$ period problem (P1), where $J_{T(t)}(x)$ is a constant instead of 0, which does not change anything in the proof of Proposition 3.2.1. Thus by Proposition 3.2.1, this means that if $t \notin \{T_1, \dots, T_{n-1}\}$ then $R_t(\infty) = 0$; otherwise $\tilde{R}_t(\infty) = 0$. Now the result follows when $C_m < \infty$ by Proposition 3.4.3. \square

Chapter 4

Make-to-Stock Production

In this chapter, we consider a similar model to the one in Chapter 3. However, this time we let the firm keep inventory at the retail site. Thus, the firm no longer needs to have pending orders to trigger production – it can stock up finished goods in anticipation of future demand. In addition to this modification to the production strategy, we also look at some other possible logistic agreement types including option contracts that provide more flexibility, as well as extending our model with stochastic spot market price and stochastic availability of additional capacity.

In Section 4.1, we introduce our basic periodic review model of a firm utilizing a fixed date/fixed capacity transportation agreement for shipping to its retail site. In this model, in each period a firm must decide how much to ship, possibly nothing, from a warehouse to a retailer to meet demand at the retailer, utilizing a combination shipping capacity already agreed to via a structured logistics agreement, and shipping capacity available on the spot market. Note that since there is a cost to holding inventory at the retailer, it does not always make sense to utilize all of the available contracted capacity, even though this capacity is already paid for. For this setting, we characterize the optimal shipping policy.

In Section 4.2, we consider a similar setting, except that instead of fixed date/fixed capacity agreement, the firm has a logistics agreement in which it has paid up-front for the right (but not the obligation) to purchase shipping later at a previously-agreed-upon price (a so-called option contract), and in Section 4.3 and Section 4.4, for similar settings we consider two ways in which the logistics provider can provide additional capacity to the buyer beyond that agreed to in the fixed date/fixed capacity agreement, and how the buyer should take advantage of this additional capacity. Lastly, in Section 4.5, we model a contract which provides flexible shipping time instead of capacity.

4.1 Logistics Agreements with Scheduled Shipments

As in Chapter 3, we consider a discrete time finite horizon problem faced by a firm that must ship a product from one location (for the purposes of exposition, we identify this as the warehouse) to some destination (the retailer) to meet stochastic demand $w_t \geq 0$ at the destination. We assume that this firm has a scheduled shipment agreement with a logistics provider at the start of the horizon. These types of agreements specify in advance the frequency and volume of shipping capacity to be provided (Yano and Gerchak [83]). Thus, a scheduled shipment agreement in a finite horizon model can be specified by its shipment periods and its supplied capacities in these periods. In our model we let C_t denote the amount of capacity available for shipping at period t ; hence the periods at which $C_t > 0$ represent the scheduled shipment periods. Of course, since demand is uncertain, the scheduled shipments may not provide sufficient capacity to meet the firms desired service level. For this reason, we model a spot market for transportation that provides expedited shipping when it is needed. At the beginning of each period, the decision maker first decides how many units of inventory to acquire (or make) at the warehouse for a unit cost of c_p . We assume that there is no lead time for acquiring this inventory. Next, the decision maker ships inventory to the retailer either via the capacity supplied per the scheduled shipment agreement if capacity is available, or via expedited shipping at a unit cost of c_s . Finally, the demand for that period at the retailer is realized, and fulfilled as much as possible depending on the available inventory at the retailer. Excess demand is backlogged, a unit of which incurs a penalty cost of p per period, and excess inventory on hand at the retailer at the end of the period incurs a holding cost of h per unit per period.

Since for simplicity and analytical tractability we restrict ourselves to linear acquisition cost with no lead time or capacity restrictions, the firm has no incentive to keep inventory at the warehouse at any positive holding cost. Thus, inventory is acquired immediately before it is shipped, and in a quantity equal to the shipping quantity. This observation enables us to reduce the dimensions of both the state and decision spaces for this problem from two to one, where at the start of each period the firm must determine the optimal shipment level u_t from the warehouse to the retailer given the current inventory position x_t at the retailer. In addition, we assume that the demand distributions in all periods are mutually independent, denote the discount factor by $\alpha \in (0, 1)$, and use the notation $(x)^+$ to represent $\max(0, x)$ and $(x)^-$ to represent $\max(0, -x)$ throughout the rest of this chapter.

Then, for this problem, the dynamic programming recursion for $t = 1, \dots, T$ is:

$$J_t(x) = \min_{u \geq 0} \{c_s(u - C_t)^+ + c_p u + \mathbb{E}_w [h(x + u - w)^+ + p(x + u - w)^- + \alpha J_{t+1}(x + u - w)]\}.$$

Throughout the rest of the chapter we denote the right and left derivatives of a function $f(x)$, by superscripting it with letters r and l , in other words $f^r(x) \equiv \lim_{\epsilon \downarrow 0} \frac{f(x+\epsilon) - f(x)}{\epsilon}$ and $f^l(x) \equiv \lim_{\epsilon \uparrow 0} \frac{f(x+\epsilon) - f(x)}{\epsilon}$. We make the following assumptions about cost parameters and

the terminal value function mostly to avoid trivial cases:

$$p > (1 - \alpha)(c_p + c_s) \quad (4.1)$$

$$J_{T+1}(x) \text{ is a convex function} \quad (4.2)$$

$$\lim_{x \rightarrow \infty} J_{T+1}^r(x) > -(h + c_p)/\alpha \quad (4.3)$$

$$\lim_{x \rightarrow -\infty} J_{T+1}^r(x) < -(c_p + c_s - p)/\alpha. \quad (4.4)$$

If assumption (4.1) does not hold then it would be never optimal to produce and ship an item with expedited shipping until possibly the last period. If assumption (4.3) does not hold then it would be optimal to produce infinitely many items at the last period. Conversely, if assumption (4.4) does not hold then it would never be optimal to produce and ship an item with expedited shipping at the last period.

As an example if the terminal value function is piecewise linear with a penalty cost of c per unit unfulfilled order and a salvage revenue of r per unit excess inventory:

$$J_{T+1}(x) = c(x)^- - r(x)^+,$$

then (4.3) implies $\alpha r < h + c_p$ while (4.4) implies $p + \alpha c > c_p + c_s$, both of which are clearly true in all but the trivial problems.

Now, we prove the following theorem, which characterizes the optimal policy for this problem:

Theorem 4.1.1. *(Optimal Policy) The optimal policy at each period is characterized by two critical levels s_t, S_t such that $s_t \leq S_t$ and*

$$\mu_t(x) = \begin{cases} 0, & \text{if } S_t \leq x \\ S_t - x, & \text{if } S_t - C_t \leq x \leq S_t \\ C_t, & \text{if } s_t - C_t \leq x \leq S_t - C_t \\ s_t - x, & \text{if } x \leq s_t - C_t. \end{cases}$$

Proof. Firstly we define $y \equiv x + u$, the order up to quantity, and rewrite the DP recursion.

$$\begin{aligned} J_t(x) &= \min_{y \geq x} \{c_s(y - x - C_t)^+ + c_p y + \mathbb{E}_w[h(y - w)^+ + p(y - w)^- + \alpha J_{t+1}(y - w)]\} - c_p x \\ &= \min_{y \geq x} \{c_s(y - x - C_t)^+ + G_t(y)\} - c_p x, \end{aligned}$$

where we define $G_t(y) \equiv c_p y + \mathbb{E}_w[h(y - w)^+ + p(y - w)^- + \alpha J_{t+1}(y - w)]$.

Let us assume (4.2), (4.3), and (4.4) hold for $J_{t+1}(x)$, then these imply

- (i) $G_t(x)$ is convex (since $\mathbb{E}_w[g(x - w)]$ is convex if $g(x)$ is convex),

- (ii) $\lim_{y \rightarrow \infty} G_t(x) = \infty$ (since $\lim_{y \rightarrow \infty} G_t^r(x) = c_p + h + \alpha J_{t+1}^r(x) > 0$, i.e. the right tail of the integrand is strictly positive),
- (iii) $\lim_{y \rightarrow -\infty} G_t(x) + c_s y = \infty$ (since $\lim_{y \rightarrow \infty} G_t^r(x) = c_p - p + \alpha J_{t+1}^r(x) < c_s$, i.e. the left tail of the integrand is strictly negative),

respectively. Hence, some minimizers s_t and S_t defined as follows:

$$s_t \in \operatorname{argmin}_{y \in \mathbb{R}} \{G_t(y) + c_s y\},$$

$$S_t \in \operatorname{argmin}_{y \in \mathbb{R}} G_t(y)$$

exist. Furthermore $s_t \leq S_t$, since

$$G_t(s_t) + c_s s_t \leq G_t(S_t) + c_s S_t \leq G_t(s_t) + c_s S_t,$$

where the first inequality follows from the definition of s_t , the second inequality follows from the definition of S_t , and the result follows since $c_s > 0$. Now we rewrite the value function as a minimum of two functions and use the convexity of $G_t(y)$ and the definitions of s_t, S_t to get the following:

$$\begin{aligned} J_t(x) &= \min_{y \geq x} \left\{ \min_{x \leq y \leq x+C_t} G_t(y), \min_{y \geq x+C_t} \{c_s y + G_t(y)\} - c_s(x + C_t) \right\} - c_p x \\ &= -c_p x + \begin{cases} \min\{G_t(x), G_t(x + C_t)\}, & \text{if } S_t \leq x \\ \min\{G_t(S_t), G_t(x + C_t)\}, & \text{if } S_t - C_t \leq x \leq S_t \\ \min\{G_t(x + C_t), G_t(x + C_t)\}, & \text{if } s_t - C_t \leq x \leq S_t - C_t \\ \min\{G_t(x + C_t), G_t(s_t) + c_s(s_t - x - C_t)\}, & \text{if } x \leq s_t - C_t \end{cases} \\ &= -c_p x + \begin{cases} G_t(x), & \text{if } S_t \leq x \\ G_t(S_t), & \text{if } S_t - C_t \leq x \leq S_t \\ G_t(x + C_t), & \text{if } s_t - C_t \leq x \leq S_t - C_t \\ G_t(s_t) + c_s(s_t - x - C_t), & \text{if } x \leq s_t - C_t. \end{cases} \end{aligned}$$

This shows that $\mu_t(x)$ is indeed the optimal policy function for period t .

Now to use Lemma 3.3.1, let $X = \mathbb{R}$, $Y(x) = \{y \in \mathbb{R} \mid y \geq x\}$, $C = \{(x, y) \mid x \in X, y \in Y(x)\}$, $g(x, y) = c_s(y - x - C_t)^+ + G_t(y) - c_p x$. Then

$$J_t(x) = \inf_{y \in Y(x)} g(x, y),$$

and also X, C are convex sets; for every $x \in X$, $Y(x)$ is a nonempty set, $J_t(x) > -\infty$; $g(x, y)$ is a convex function on C ; thus $J_t(x)$ is a convex function on \mathbb{R} .

Furthermore,

$$\lim_{x \rightarrow \infty} J_t^r(x) = -c_p + \lim_{x \rightarrow \infty} G_t^r(x), \quad \left[\text{since } \lim_{x \rightarrow \infty} J_t(x) = -c_p x + \lim_{x \rightarrow \infty} G_t(x) \right]$$

$$\begin{aligned}
&= -c_p + c_p + h + \alpha \lim_{x \rightarrow \infty} J_{t+1}^r(x), \quad \left[\text{by definition of } G_t(x) \right] \\
&> h - h - c_p, \quad \left[\text{by (4.3)} \right] \\
&> -(h + c_p)/\alpha,
\end{aligned}$$

and

$$\begin{aligned}
\lim_{x \rightarrow -\infty} J_t^r(x) &= -c_p - c_s, \quad \left[\text{since } \lim_{x \rightarrow -\infty} J_t(x) = -c_p x + G_t(s_t) + c_s(s_t - x - C_t) \right] \\
&< -(c_p + c_s - p)/\alpha, \quad \left[\text{by (4.1)} \right].
\end{aligned}$$

respectively. Hence (4.2), (4.3), and (4.4) hold for $J_t(x)$ as well. Then we can repeat this argument sequentially for $t = T, T - 1, \dots, 1$, which completes the proof. \square

In other words, when inventory level at the retailer is above the higher of the two critical levels, S_t , do nothing. When inventory level at the retailer is below S_t but close enough so that there is sufficient scheduled shipment capacity C_t to raise the inventory level to S_t , purchase and ship enough from the warehouse to do so. When inventory level is low enough that C_t isn't sufficient capacity to raise the inventory level to S_t at the retailer, but there is enough to raise the inventory level to or above the lower critical value, s_t use all C_t units of scheduled capacity. Finally, when the inventory level at the retailer is so low that even using the entire scheduled capacity would not raise the inventory level to s_t , use all of the scheduled capacity plus some expedited capacity to raise the inventory level at the retailer to s_t .

In the following result we show that increasing the scheduled shipment capacity for some period k , while keeping everything else constant, results in a decrease in all the critical levels up to period k .

Proposition 4.1.2. *Both s_t and S_t are decreasing with C_k for any $t < k$.*

Proof. Firstly, we will show that the right derivative of the value function is increasing with the reserved capacity level at that period. Remember that the value functions are continuous and hence both $J_t(x)$ and $G_t(x)$ must have their right derivative defined in their domain. As before we will denote the right derivative of these functions by $J_t^r(x) = \lim_{\epsilon \downarrow 0} \frac{J_t(x+\epsilon) - J_t(x)}{\epsilon}$ and similarly by $G_t^r(x) = \lim_{\epsilon \downarrow 0} \frac{G_t(x+\epsilon) - G_t(x)}{\epsilon}$.

We have the following by Theorem 4.1.1:

$$J_t(x) = -c_p x + \begin{cases} G_t(x), & \text{if } S_t \leq x \\ G_t(S_t), & \text{if } S_t - C_t \leq x \leq S_t \\ G_t(x + C_t), & \text{if } s_t - C_t \leq x \leq S_t - C_t \\ G_t(s_t) + c_s(s_t - x - C_t), & \text{if } x \leq s_t - C_t \end{cases}$$

and hence

$$J_{t-1}^r(x) = -c_p + \begin{cases} G_t^r(x), & \text{if } S_t \leq x \\ 0, & \text{if } S_t - C_t \leq x < S_t \\ G_t^r(x + C_t), & \text{if } s_t - C_t \leq x < S_t - C_t \\ -c_s, & \text{if } x < s_t - C_t. \end{cases}$$

Notice that for some x , an increase in C_t either does not change the partition x belongs to or it shifts x to a higher partition (i.e. a partition to the right of the original). Notice that $G_t^r(x)$ is independent of C_t by definition as of course are the constants c_p , 0 and $-c_s$. $G_t^r(x+C_t)$ on the other hand is increasing with C_t since $G_t(x)$ is convex. Thus we have shown that the partial function $J_t^r(x)$ is increasing with C_t within each partition in its definition. Furthermore, $G_t^r(x) \geq 0$ when $S_t \leq x$ and $-c_s \leq G_t^r(x+C_t) \leq 0$ when $s_t - C_t \leq x < S_t - C_t$, since $G_t(x)$ is convex, S_t is a minimizer of $G_t(x)$, and s_t is a minimizer of $G_t(x) + c_s x$. Thus we have also shown that for some x , if an increase in C_t causes x to shift to another partition, $J_t^r(x)$ would increase in that case as well. Hence we conclude that $J_t^r(x)$ is increasing with C_t .

Secondly, since $J_t^r(x)$ is increasing with C_t , so is $G_{t-1}^r(x)$ by just the definition of $G_{t-1}(x)$. Then this implies that s_{t-1} and S_{t-1} are decreasing with C_t , since they are some minimizers for the convex functions $G_{t-1}(x) + c_s x$ and $G_{t-1}(x)$ respectively.

Lastly, since $G_{t-1}^r(x)$ is increasing with C_t so is $J_{t-1}^r(x)$, since we have again the following from Theorem 4.1.1:

$$J_t^r(x) = -c_p + \begin{cases} G_{t-1}^r(x), & \text{if } S_{t-1} \leq x \\ 0, & \text{if } S_{t-1} - C_{t-1} \leq x < S_{t-1} \\ G_{t-1}^r(x + C_{t-1}), & \text{if } s_{t-1} - C_{t-1} \leq x < S_{t-1} - C_{t-1} \\ -c_s, & \text{if } x < s_{t-1} - C_{t-1}. \end{cases}$$

Then we can clearly repeat this argument recursively for all $k < t$, e.g. since $J_{t-1}^r(x)$ is increasing with C_t , so is $G_{t-2}^r(x)$, which implies that s_{t-2} and S_{t-2} are decreasing with C_t and so on. \square

Intuitively, we can see why this is true by noting that if all the reserved capacity levels in future periods are zero, then we would be primarily concerned with balancing the holding cost of excess inventory versus the penalty cost of shortage. However if there are future periods with positive reserved capacity, then we would like to utilize these contracted shipments as much as possible in addition to balancing the holding and penalty costs, as those reserved capacity levels essentially represent free shipment opportunities. Hence an increase in the reserved shipment capacity level of a future period makes us less willing to produce and ship today, which results in lower inventory levels.

In practice, it is usually the case that the contracted shipments are scheduled periodically, which means that there is a pattern of zero scheduled shipment capacity periods followed by a positive one. In the next result we show that at periods with no scheduled shipment the optimal policy reduces to an order-up-to policy, and we also show that these order-up-to levels are decreasing.

Proposition 4.1.3. *Suppose that $C_k > 0$ and $C_t = 0$ for $k < t \leq l$. Then at periods $k + 1, \dots, l - 1, l$ the optimal policy reduces to an order-up-to policy. Let s_t denote the order-up-to level at period t , where $k < t \leq l$. Then the following holds:*

$$s_l \leq \dots \leq s_{k-1} \leq s_k.$$

Proof. Notice that if $C_t = 0$, then the expression for the value function at period t reduces to the following:

$$J_t(x) = \min_{y \geq x} \{c_s y + G_t(y)\} - (c_s + c_p)x,$$

where $G_t(y) \equiv c_p y + \mathbb{E}_w[h(y - w)^+ + p(y - w)^- + \alpha J_{t+1}(y - w)]$ as before.

We have already shown in the proof of Theorem 4.1.1 that $G_t(y)$ is convex and that there exists a minimizer, s_t , of $c_s y + G_t(y)$. Hence, the optimal policy at period t given $C_t = 0$, is order-up-to s_t , and thus the value function for this period can be written as follows:

$$J_t(x) = -c_p x + \begin{cases} G_t(x), & \text{if } s_t \leq x \\ G_t(s_t) + c_s(s_t - x), & \text{if } x \leq s_t \end{cases}$$

Let us denote the left derivative of $J_t(x)$ and $G_t(x)$ by $J_t^l(x) = \lim_{\epsilon \uparrow 0} \frac{J_t(x+\epsilon) - J_t(x)}{\epsilon}$ and $G_t^l(x) = \lim_{\epsilon \uparrow 0} \frac{G_t(x+\epsilon) - G_t(x)}{\epsilon}$ respectively.

Notice that if $C_t = 0$ then $J_t^l(x) = -c_s - c_p$ for $x \leq s_t$, and also that $J_t^l(x) \geq -c_s - c_p$ regardless of the value of C_t . The latter claim can easily be verified by noting that the left derivative of a convex function is increasing.

Now, for $k \leq t \leq l$, let $\tilde{s}_t \equiv \max\{s : c_s s + G_t(s) \leq c_s y + G_t(y)\}$ and let $F_w(y)$ be the cumulative probability function for the demand distribution. Then

$$\begin{aligned} c_s + G_t^l(\tilde{s}_{t+1}) &= c_s + c_p + (h + p)F_w(\tilde{s}_{t+1}) - p + \alpha \mathbb{E}_w[J_{t+1}^l(\tilde{s}_{t+1} - w)] \\ &\leq c_s + c_p + (h + p)F_w(\tilde{s}_{t+1}) - p + \alpha \mathbb{E}_w[J_{t+2}^l(\tilde{s}_{t+1} - w)] \\ &= c_s + G_{t+1}^l(\tilde{s}_{t+1}) \\ &\leq 0, \end{aligned}$$

where the inequality follows since $J_{t+1}^l(\tilde{s}_{t+1} - w) = -c_s - c_p \leq J_{t+2}^l(\tilde{s}_{t+1} - w)$. This implies that $s_{t+1} \leq \tilde{s}_{t+1} \leq s_t$ since $G_t(y)$ is convex. Then using this inequality recursively for $k \leq t < l$ gives us the result. \square

While Proposition 4.1.2 tells us that an increase in the reserved shipment capacity level of a future period makes us to produce less today, Proposition 4.1.3 tells us that we would also be less willing to produce and ship as we get closer to a period with reserved capacity shipment, since we would again like to utilize that reserved capacity as much as possible.

4.2 Logistics Agreements with Options and Stochastic Spot Market Price

In Section 4.1, we considered scheduled shipment transportation agreements, which do not leave the firm much flexibility. Indeed, these contracts concentrate demand-related shipping risk with the firm, rather than with the logistics provider. One way to mitigate this risk is to incorporate options into the logistics agreement. Options give the firm the right to use the logistics provider's service at a certain predetermined rate if it opts to do so.

To extend our model to include options, we first introduce some notation. Suppose that at any period t , the rate of using the transportation service specified in the option agreement is set to c_o with a reserved capacity level of C_t . Let us adopt the shorthand notation $g_{t+1}(x+u) \equiv \mathbb{E}_w[h(x+u-w)^+ + p(x+u-w)^- + \alpha J_{t+1}(x+u-w)]$ to simplify notational exposition. Then the dynamic programming recursion becomes:

$$\begin{aligned} J_t(x) &= \min_{u \geq 0} \{c_o \min\{u, C_t\} + c_s(u - C_t)^+ + c_p u + g_{t+1}(x+u)\} \\ &= \min_{u \geq 0} \{(c_s - c_o)(u - C_t)^+ + (c_p + c_o)u + g_{t+1}(x+u)\} \end{aligned}$$

Notice that this recursion has exactly the same form with the model we analyzed previously in Section 4.1, where c_s and c_p are replaced with $c_s - c_o$ and $c_p + c_o$ respectively. Consequently, that analysis directly extends to the case with option contracts.

However, this equivalence assumes that the spot market price for expedited shipping is stationary. This was not a point of concern for the scheduled shipments since the marginal cost of using the reserved capacity in that case is zero, which is always cheaper than the spot market price. However, in the case of option agreements, this may not be true. Furthermore, variations in the spot market price may induce a considerable risk for the manufacturer, which may want to utilize option contracts to hedge against this risk.

For these reasons we extend our main theorem to the case where the spot market price is stochastic. More specifically, we assume that the spot market price for expedited shipping, c_t , at each period becomes known to the manufacturer only at the beginning of that period and may depend on all the relevant information gathered up to that point, which we will denote by \mathcal{F}_t . In particular notice that this information set is a filter, $\mathcal{F}_t \subset \mathcal{F}_{t+1}$, and $c_t \in \mathcal{F}_t$. In this case we define:

$$g_{t+1}(x+u) \equiv \mathbb{E}_{w, c_{t+1}} [h(x+u-w)^+ + p(x+u-w)^- + \alpha J_{t+1}(x+u-w, \mathcal{F}_{t+1}) \mid \mathcal{F}_t]$$

Then dynamic programming recursion becomes:

$$\begin{aligned} J_t(x, \mathcal{F}_t) &= \min_{u \geq 0} \{ \min\{c_o, c_t\} \min\{u, C_t\} + c_t(u - C_t)^+ + c_p u + g_{t+1}(x + u) \} \\ &= \min_{u \geq 0} \{ (c_t - \min\{c_o, c_t\})(u - C_t)^+ + (c_p + \min\{c_o, c_t\})u + g_{t+1}(x + u) \} \end{aligned}$$

Once again notice that this recursion has a very similar form to the model analyzed previously in Section 4.1, where c_s and c_p are replaced with $c_t - \min\{c_o, c_t\}$ and $c_p + \min\{c_o, c_t\}$ respectively. The primary substantial difference, however, is that the cost parameters are no longer stationary or deterministic. However a parallel line of reasoning can be used to prove that the structure of the optimal policy is essentially the same, except that the critical levels at any period depend on the realization of the spot market price at that period.

Firstly, we adopt the following shorthand notation to reduce the amount of clutter in the formulations:

$$\begin{aligned} c_{s,t} &\equiv c_t - \min\{c_o, c_t\} \\ c_{p,t} &\equiv c_p + \min\{c_o, c_t\} \end{aligned}$$

Secondly, for notational simplicity, we assume that the spot market price is a Markov process. This assumption also ensures the dimension of the state space does not increase with time. However we note that our results extend directly to the case where \mathcal{F}_t is a general information set as described above.

Thirdly, as before, to avoid trivial cases we make the following assumptions, which correspond to the assumptions (4.1),(4.2),(4.3), and (4.4) of Section 4.1 respectively:

$$p > (1 - \alpha)(c_{p,t} + c_t), \text{ with probability } 1. \quad (4.1')$$

$$J_{T+1}(x, c_{T+1}) \text{ is convex in } x \text{ for any } c_{T+1}. \quad (4.2')$$

$$\lim_{x \rightarrow \infty} J_{T+1}^r(x, c_{T+1}) > -(h + c_{p,t} + c_o)/\alpha, \text{ with probability } 1. \quad (4.3')$$

$$\lim_{x \rightarrow -\infty} J_{T+1}^r(x, c_{T+1}) < -(c_{p,t} + c_{s,t} - p)/\alpha \text{ with probability } 1. \quad (4.4')$$

We, now, prove the following result, which is an extends Theorem 4.1.1 to the case of logistics agreements with options in the presence of stochastic spot market prices:

Theorem 4.2.1. *(Optimal Policy) At each period t , for any realization of the spot market price for expedited shipment, c_t , the optimal policy is characterized by two critical levels $s_t(c_t)$, $S_t(c_t)$ such that $s_t(c_t) \leq S_t(c_t)$ and*

$$\mu_t(x, c_t) = \begin{cases} 0, & \text{if } S_t(c_t) \leq x \\ S_t(c_t) - x, & \text{if } S_t(c_t) - C_t \leq x \leq S_t(c_t) \\ C_t, & \text{if } s_t(c_t) - C_t \leq x \leq S_t(c_t) - C_t \\ s_t(c_t) - x, & \text{if } x \leq s_t(c_t) - C_t. \end{cases}$$

Proof. Firstly we define $y \equiv x + u$, the order up to quantity, and rewrite the DP recursion.

$$\begin{aligned} J_t(x, c_t) &= \min_{y \geq x} \left\{ c_{s,t}(y - x - C_t)^+ + c_{p,t}y \right. \\ &\quad \left. + \mathbb{E}_{w, c_{t+1}} [h(y - w)^+ + p(y - w)^- + \alpha J_{t+1}(y - w, c_{t+1})] \right\} - c_{p,t}x \\ &= \min_{y \geq x} \{ c_{s,t}(y - x - C_t)^+ + G_t(y, c_t) \} - c_{p,t}x, \end{aligned}$$

where we define $G_t(y, c_t) \equiv c_{p,t}y + \mathbb{E}_{w, c_{t+1}} [h(y - w)^+ + p(y - w)^- + \alpha J_{t+1}(y - w, c_{t+1})]$.

Let us assume (4.2'), (4.3'), and (4.4') hold for $J_{t+1}(x, c_{t+1})$, then these imply (i) $G_t(x, c_t)$ is convex in x , (ii) $\lim_{y \rightarrow \infty} G_t(x, c_t) = \infty$, and (iii) $\lim_{y \rightarrow -\infty} G_t(x, c_t) + c_{s,t}y = \infty$ respectively. Hence, the minimizers $s_t(c_t)$ and $S_t(c_t)$ defined as follows:

$$\begin{aligned} s_t(c_t) &\equiv \operatorname{argmin}_{y \in \mathbb{R}} G_t(y, c_t) + c_{s,t}y \\ S_t(c_t) &\equiv \operatorname{argmin}_{y \in \mathbb{R}} G_t(y, c_t) \end{aligned}$$

exist. Furthermore $s_t(c_t) \leq S_t(c_t)$, since

$$G_t(s_t(c_t), c_t) + c_{s,t}s_t(c_t) \leq G_t(S_t(c_t), c_t) + c_s(c_t)S_t(c_t) \leq G_t(s_t(c_t), c_t) + c_s(c_t)S_t(c_t).$$

Now we rewrite the value function as a minimum of two functions and use the convexity of $G_t(y, c_t)$ in y and the definitions of $s_t(c_t), S_t(c_t)$ to get the following:

$$\begin{aligned} J_t(x, c_t) &= \min_{y \geq x} \left\{ \min_{x \leq y \leq x + C_t} G_t(y, c_t), \min_{y \geq x + C_t} \{ c_{s,t}y + G_t(y, c_t) \} - c_{s,t}(x + C_t) \right\} - c_{p,t}x \\ &= -c_{p,t}x \\ &\quad + \begin{cases} \min\{G_t(x, c_t), G_t(x + C_t, c_t)\}, & \text{if } S_t(c_t) \leq x \\ \min\{G_t(S_t(c_t), c_t), G_t(x + C_t, c_t)\}, & \text{if } S_t(c_t) - C_t \leq x \leq S_t(c_t) \\ \min\{G_t(x + C_t, c_t), G_t(x + C_t, c_t)\}, & \text{if } s_t(c_t) - C_t \leq x \leq S_t(c_t) - C_t \\ \min\{G_t(x + C_t, c_t), G_t(s_t(c_t), c_t) + c_s(s_t(c_t) - x - C_t)\}, & \text{if } x \leq s_t(c_t) - C_t \end{cases} \\ &= -c_{p,t}x + \begin{cases} G_t(x, c_t), & \text{if } S_t(c_t) \leq x \\ G_t(S_t(c_t), c_t), & \text{if } S_t(c_t) - C_t \leq x \leq S_t(c_t) \\ G_t(x + C_t, c_t), & \text{if } s_t(c_t) - C_t \leq x \leq S_t(c_t) - C_t \\ G_t(s_t(c_t), c_t) + c_s(s_t(c_t) - x - C_t), & \text{if } x \leq s_t(c_t) - C_t. \end{cases} \end{aligned}$$

This shows that $\mu_t(x, c_t)$ is indeed the optimal policy function for period t .

Setting $X = \mathbb{R}$, $Y(x) = \{y \in \mathbb{R} \mid y \geq x\}$, $g(x, y, c_t) = c_{s,t}(y - x - C_t)^+ + G_t(y, c_t)$ and using Lemma 3.3.1 we deduce that $J_t(x, c_t)$ is convex in x . Furthermore (4.1') implies

$$\lim_{x \rightarrow \infty} J_t^r(x, c_t) = -c_{p,t} + h + c_{p,t} + \alpha \lim_{x \rightarrow \infty} J_{t+1}^r(x, c_{t+1}) > -(h + c_{p,t})/\alpha, \text{ and}$$

$$\lim_{x \rightarrow -\infty} J_t^r(x, c_t) = -c_{p,t} - c_{s,t} < -(c_{p,t} + c_{s,t} - p)/\alpha$$

respectively. Hence (4.2'), (4.3'), and (4.4') hold for $J_t(x, c_t)$ as well. Then we can repeat this argument sequentially for $t = T, T - 1, \dots, 1$, which completes the proof. \square

In other words, this policy is similar to the policy in Theorem 4.1.1, except that all parameters are functions of the spot market price (and policy parameter levels are impacted by the positive marginal cost of utilizing both contracted and spot market shipping).

4.3 Multi-level Option Agreements

More sophisticated option agreements can include several different option levels and rates. That is, in addition to the primary capacity option level and rate, the logistics provider may agree to provide the option to purchase additional capacity at a different (higher) rate.

We model this by extending the model in the previous section by assuming that in any period t , in addition to reserved capacity level C_t available at rate c_o , the logistics provider will provide up to $100\gamma_t\%$ of the reserved capacity at a rate of $c_o + c_e$.

In this case, the new dynamic programming equations are:

$$\begin{aligned} J_t(x) &= \min_{u \geq 0} \{c_o \min\{u, C_t\} + (c_o + c_e) \min\{(u - C_t)^+, \gamma_t C_t\} \\ &\quad + c_s(u - (1 + \gamma_t)C_t)^+ + c_p u + g_{t+1}(x + u)\} \\ &= \min_{u \geq 0} \{(c_s - c_o - c_e)(u - (1 + \gamma_t)C_t)^+ + c_e(u - C_t)^+ + (c_p + c_o)u + g_{t+1}(x + u)\} \end{aligned}$$

Let us denote $(1 + \gamma_t)C_t$ by \bar{C}_t to simplify notation.

To avoid trivial cases we make the following assumptions, which correspond to the assumptions (4.1), (4.2), (4.3), and (4.4) of Section 4.1 respectively:

$$p > (1 - \alpha)(c_p + c_s) \tag{4.1''}$$

$$J_{T+1}(x) \text{ is a convex function} \tag{4.2''}$$

$$\lim_{x \rightarrow \infty} J_{T+1}^r(x) > -(h + c_p + c_o)/\alpha \tag{4.3''}$$

$$\lim_{x \rightarrow -\infty} J_{T+1}^r(x) < -(c_p + c_s - p)/\alpha. \tag{4.4''}$$

Theorem 4.3.1. *(Optimal Policy) The optimal policy at each period is characterized by*

three critical levels $\underline{s}_t, \bar{s}_t, S_t$ such that $\underline{s}_t \leq \bar{s}_t \leq S_t$ and

$$\mu_t(x) = \begin{cases} 0, & \text{if } S_t \leq x \\ S_t - x, & \text{if } S_t - C_t \leq x \leq S_t \\ C_t, & \text{if } \bar{s}_t - C_t \leq x \leq S_t - C_t \\ \bar{s}_t - x, & \text{if } \bar{s}_t - \bar{C}_t \leq x \leq \bar{s}_t - C_t \\ \bar{C}_t, & \text{if } \underline{s}_t - \bar{C}_t \leq x \leq \bar{s}_t - \bar{C}_t \\ \underline{s}_t - x, & \text{if } x \leq \underline{s}_t - \bar{C}_t. \end{cases}$$

Proof. Firstly we define $y \equiv x + u$, the order up to quantity, and rewrite the DP recursion.

$$\begin{aligned} J_t(x) &= \min_{y \geq x} \{ (c_s - c_o - c_e)(y - x - \bar{C}_t)^+ + c_e(y - x - C_t)^+ + (c_p + c_o)y + g_{t+1}(y) \} \\ &\quad - (c_p + c_o)x \\ &= \min_{y \geq x} \{ (c_s - c_o - c_e)(y - x - \bar{C}_t)^+ + c_e(y - x - C_t)^+ + G_t(y) \} - (c_p + c_o)x, \end{aligned}$$

where we define $G_t(y) \equiv (c_p + c_o)y + g_{t+1}(y)$.

Let us assume (4.2''), (4.3''), and (4.4'') hold for $J_{t+1}(x)$, then these imply (i) $G_t(x)$ is convex, (ii) $\lim_{y \rightarrow \infty} G_t(x) = \infty$, and (iii) $\lim_{y \rightarrow -\infty} G_t(x) + (c_s - c_o)y = \infty$ respectively. Hence, the minimizers \underline{s}_t , \bar{s}_t and S_t defined as follows:

$$\begin{aligned} \underline{s}_t &\equiv \operatorname{argmin}_{y \in \mathbb{R}} G_t(y) + (c_s - c_o)y \\ \bar{s}_t &\equiv \operatorname{argmin}_{y \in \mathbb{R}} G_t(y) + c_e y \\ S_t &\equiv \operatorname{argmin}_{y \in \mathbb{R}} G_t(y) \end{aligned}$$

exist. Furthermore $\underline{s}_t \leq \bar{s}_t \leq S_t$, since

$$G_t(\underline{s}_t) + (c_s - c_o)\underline{s}_t \leq G_t(\bar{s}_t) + (c_s - c_o)\bar{s}_t \leq G_t(\underline{s}_t) + c_s \underline{s}_t - c_o \underline{s}_t$$

and

$$G_t(\bar{s}_t) + c_e \bar{s}_t \leq G_t(S_t) + c_e S_t \leq G_t(\bar{s}_t) + c_e S_t.$$

Now we rewrite the value function as a minimum of three functions and use the convexity of $G_t(y)$ and the definitions of $\underline{s}_t, \bar{s}_t, S_t$ to get the following:

$$\begin{aligned} J_t(x) + (c_p + c_o)x &= \min \left\{ \min_{x \leq y \leq x + C_t} G_t(y), \min_{x + C_t \leq y \leq x + \bar{C}_t} \{ G_t(y) + c_e y \} - c_e(x + C_t), \right. \\ &\quad \left. \min_{y \geq x + \bar{C}_t} \{ G_t(y) + (c_s - c_o)y \} - (c_s - c_o)(x + \bar{C}_t) + c_e(\bar{C}_t - C_t) \right\} \end{aligned}$$

$$\begin{aligned}
& \begin{cases} \min\{G_t(x), G_t(x + C_t), G_t(x + \bar{C}_t) + c_e(\bar{C}_t - C_t)\}, \\ \min\{G_t(S_t), G_t(x + C_t), G_t(x + \bar{C}_t) + c_e(\bar{C}_t - C_t)\}, \\ \min\{G_t(x + C_t), G_t(x + C_t), G_t(x + \bar{C}_t) + c_e(\bar{C}_t - C_t)\}, \\ \min\{G_t(x + C_t), G_t(\bar{s}_t) + c_e(\bar{s}_t - x + C_t), G_t(x + \bar{C}_t) + c_e(\bar{C}_t - C_t)\}, \\ \min\{G_t(x + C_t), G_t(x + \bar{C}_t) + c_e(\bar{C}_t - C_t), G_t(x + \bar{C}_t) + c_e(\bar{C}_t - C_t)\}, \\ \min\{G_t(x + C_t), G_t(x + \bar{C}_t) + c_e(\bar{C}_t - C_t), G_t(\underline{s}_t) + (c_s - c_o)\underline{s}_t - c_e(\bar{C}_t - C_t)\}, \end{cases} \\
= & \begin{cases} G_t(x), & \text{if } S_t \leq x \\ G_t(S_t), & \text{if } S_t - C_t \leq x \leq S_t \\ G_t(x + C_t), & \text{if } \bar{s}_t - C_t \leq x \leq S_t - C_t \\ G_t(\bar{s}_t) + c_e(\bar{s}_t - x - C_t), & \text{if } \bar{s}_t - \bar{C}_t \leq x \leq \bar{s}_t - C_t \\ G_t(x + \bar{C}_t) + c_e(\bar{C}_t - C_t), & \text{if } \underline{s}_t - \bar{C}_t \leq x \leq \bar{s}_t - \bar{C}_t \\ G_t(\underline{s}_t) + (c_s - c_o)(\underline{s}_t - x - \bar{C}_t) - c_e(\bar{C}_t - C_t), & \text{if } x \leq \underline{s}_t - \bar{C}_t. \end{cases}
\end{aligned}$$

This shows that $\mu_t(x)$ is indeed the optimal policy function for period t .

Setting $X = \mathbb{R}$, $Y(x) = \{y \in \mathbb{R} \mid y \geq x\}$, $g(x, y) = (c_s - c_o - c_e)(y - x - \bar{C}_t)^+ + c_e(y - x - C_t)^+ + G_t(y)$ and using Lemma 3.3.1 we deduce that $J_t(x)$ is convex. Furthermore (4.1'') implies

$$\begin{aligned}
\lim_{x \rightarrow \infty} J_t^r(x) &= -c_p - c_o + h + c_p + c_o + \alpha \lim_{x \rightarrow \infty} J_{t+1}^r(x) > -(h + c_p + c_o)/\alpha, \text{ and} \\
\lim_{x \rightarrow -\infty} J_t^r(x) &= -c_s + c_o + c_e - c_e - c_p - c_o < -(c_p + c_s - p)/\alpha
\end{aligned}$$

respectively. Hence (4.2''), (4.3''), and (4.4'') hold for $J_t(x)$ as well. Then we can repeat this argument sequentially for $t = T, T - 1, \dots, 1$, which completes the proof. \square

This policy generalizes the policy of the two previous models by adding a third critical value. In other words, the optimal policy is characterized by three time-dependent critical values \bar{s}_t , \underline{s}_t and S_t such that either some or all of the cheaper reserved capacity, or all of the cheaper reserved capacity and some of the more expensive reserved capacity, or all of the reserved capacity, or all of the reserved capacity plus the spot market is used to ship goods to the retailer, depending on the starting inventory at the retailer.

4.4 Stochastic Availability of Additional Capacity

In practice, the firm and its logistics provider may not agree to a formal option agreement, but may instead agree to scheduled shipments as in Section 4.1, with the added provision that if the if the logistics provider is left with some unused capacity in any period, she may promise to provide this additional capacity to the firm at some below-spot-market rate if necessary. This practice helps the firm more easily handle peaks in demand without

committing to excessive capacity levels, and also helps the logistics provider to better utilize capacity while maintaining a good relationship with its customers.

To model this agreement let us assume that at each period t , in addition to the reserved capacity level C_t , the buyer will have access to A_t units of unused capacity by the logistics provider at the rate c_a . Here we assume A_t is a nonnegative independently distributed random variable, the value of which becomes known to the buyer at the beginning of period t .

In this case we define:

$$g_{t+1}(x+u) \equiv \mathbb{E}_{w, A_{t+1}} [h(x+u-w)^+ + p(x+u-w)^- + \alpha J_{t+1}(x+u-w, A_{t+1})]$$

Then dynamic programming recursion becomes:

$$\begin{aligned} J_t(x, A_t) &= \min_{u \geq 0} \{c_a \min\{(u - C_t)^+, A_t\} + c_s(u - C_t - A_t)^+ + c_p u + g_{t+1}(x+u)\} \\ &= \min_{u \geq 0} \{c_a(u - C_t)^+ + (c_s - c_a)(u - C_t - A_t)^+ + c_p u + g_{t+1}(x+u)\} \end{aligned}$$

Notice that the last equality resembles the recursion in Section 4.3. Indeed under similar conditions as in Section 4.3 and if we denote $C_t + A_t$ by \bar{C}_t , we can show that Theorem 4.3.1 applies to this case as well.

More formally if the following holds:

$$p > (1 - \alpha)(c_p + c_s) \quad (4.1''')$$

$$J_{T+1}(x, A_{T+1}) \text{ is convex in } x \text{ with probability } 1 \quad (4.2''')$$

$$\lim_{x \rightarrow \infty} J_{T+1}^r(x, A_{T+1}) > -(h + c_p)/\alpha, \text{ with probability } 1 \quad (4.3''')$$

$$\lim_{x \rightarrow -\infty} J_{T+1}^r(x, A_{T+1}) < -(c_p + c_s - p)/\alpha \text{ with probability } 1 \quad (4.4''')$$

Then the optimal policy function can be described by Theorem 4.3.1, where $\bar{C}_t \equiv C_t + A_t$. For completeness we present the formal statement of the theorem and its proof below.

Theorem 4.4.1. *(Optimal Policy) The optimal policy at each period is characterized by three critical levels $\underline{s}_t, \bar{s}_t, S_t$ such that $\underline{s}_t \leq \bar{s}_t \leq S_t$ and*

$$\mu_t(x) = \begin{cases} 0, & \text{if } S_t \leq x \\ S_t - x, & \text{if } S_t - C_t \leq x \leq S_t \\ C_t, & \text{if } \bar{s}_t - C_t \leq x \leq S_t - C_t \\ \bar{s}_t - x, & \text{if } \bar{s}_t - \bar{C}_t \leq x \leq \bar{s}_t - C_t \\ \bar{C}_t, & \text{if } \underline{s}_t - \bar{C}_t \leq x \leq \bar{s}_t - \bar{C}_t \\ \underline{s}_t - x, & \text{if } x \leq \underline{s}_t - \bar{C}_t. \end{cases}$$

Proof. Firstly we define $y \equiv x + u$, the order up to quantity, and rewrite the DP recursion.

$$\begin{aligned} J_t(x, A_t) &= \min_{y \geq x} \{(c_s - c_a)(y - x - \bar{C}_t)^+ + c_a(y - x - C_t)^+ + c_p y + g_{t+1}(y)\} - c_p x \\ &= \min_{y \geq x} \{(c_s - c_a)(y - x - \bar{C}_t)^+ + c_a(y - x - C_t)^+ + G_t(y)\} - c_p x, \end{aligned}$$

where we define $G_t(y) \equiv c_p y + g_{t+1}(y)$.

Let us assume (4.2'''), (4.3'''), and (4.4''') hold for $J_{t+1}(x)$, then these imply (i) $G_t(x)$ is convex, (ii) $\lim_{y \rightarrow \infty} G_t(x) = \infty$, and (iii) $\lim_{y \rightarrow -\infty} G_t(x) + c_s y = \infty$ respectively. Hence, the minimizers \underline{s}_t , \bar{s}_t and S_t defined as follows:

$$\begin{aligned} \underline{s}_t &\equiv \operatorname{argmin}_{y \in \mathbb{R}} G_t(y) + c_s y \\ \bar{s}_t &\equiv \operatorname{argmin}_{y \in \mathbb{R}} G_t(y) + c_a y \\ S_t &\equiv \operatorname{argmin}_{y \in \mathbb{R}} G_t(y) \end{aligned}$$

exist. Furthermore $\underline{s}_t \leq \bar{s}_t \leq S_t$, since

$$G_t(\underline{s}_t) + c_s \underline{s}_t \leq G_t(\bar{s}_t) + c_s \bar{s}_t \leq G_t(\underline{s}_t) + c_s \underline{s}_t S_t$$

and

$$G_t(\bar{s}_t) + c_a \bar{s}_t \leq G_t(S_t) + c_a S_t \leq G_t(\bar{s}_t) + c_a S_t.$$

Now we rewrite the value function as a minimum of three functions and use the convexity of $G_t(y)$ and the definitions of $\underline{s}_t, \bar{s}_t, S_t$ to get the following:

$$\begin{aligned} J_t(x, A_t) + c_p x &= \min \left\{ \min_{x \leq y \leq x + C_t} G_t(y), \min_{x + C_t \leq y \leq x + \bar{C}_t} \{G_t(y) + c_a y\} - c_a(x + C_t), \right. \\ &\quad \left. \min_{y \geq x + \bar{C}_t} \{G_t(y) + c_s y\} - c_s(x + \bar{C}_t) + c_a(\bar{C}_t - C_t) \right\} \\ &= \begin{cases} \min\{G_t(x), G_t(x + C_t), G_t(x + \bar{C}_t) + c_a(\bar{C}_t - C_t)\}, \\ \min\{G_t(S_t), G_t(x + C_t), G_t(x + \bar{C}_t) + c_a(\bar{C}_t - C_t)\}, \\ \min\{G_t(x + C_t), G_t(x + C_t), G_t(x + \bar{C}_t) + c_a(\bar{C}_t - C_t)\}, \\ \min\{G_t(x + C_t), G_t(\bar{s}_t) + c_a(\bar{s}_t - x + C_t), G_t(x + \bar{C}_t) + c_a(\bar{C}_t - C_t)\}, \\ \min\{G_t(x + C_t), G_t(x + \bar{C}_t) + c_a(\bar{C}_t - C_t), G_t(x + \bar{C}_t) + c_a(\bar{C}_t - C_t)\}, \\ \min\{G_t(x + C_t), G_t(x + \bar{C}_t) + c_a(\bar{C}_t - C_t), G_t(\underline{s}_t) + c_s \underline{s}_t - c_a(\bar{C}_t - C_t)\}, \end{cases} \\ &= \begin{cases} G_t(x), & \text{if } S_t \leq x \\ G_t(S_t), & \text{if } S_t - C_t \leq x \leq S_t \\ G_t(x + C_t), & \text{if } \bar{s}_t - C_t \leq x \leq S_t - C_t \\ G_t(\bar{s}_t) + c_a(\bar{s}_t - x - C_t), & \text{if } \bar{s}_t - \bar{C}_t \leq x \leq \bar{s}_t - C_t \\ G_t(x + \bar{C}_t) + c_a(\bar{C}_t - C_t), & \text{if } \underline{s}_t - \bar{C}_t \leq x \leq \bar{s}_t - \bar{C}_t \\ G_t(\underline{s}_t) + c_s(\underline{s}_t - x - \bar{C}_t) - c_a(\bar{C}_t - C_t), & \text{if } x \leq \underline{s}_t - \bar{C}_t. \end{cases} \end{aligned}$$

This shows that $\mu_t(x)$ is indeed the optimal policy function for period t .

Setting $X = \mathbb{R}$, $Y(x) = \{y \in \mathbb{R} \mid y \geq x\}$, $g(x, y) = (c_s - c_a)(y - x - \bar{C}_t)^+ + c_a(y - x - C_t)^+ + G_t(y)$ and using Lemma 3.3.1 we deduce that $J_t(x, A_t)$ is convex in x . Furthermore (4.1''') implies

$$\begin{aligned} \lim_{x \rightarrow \infty} J_t^r(x, A_t) &= -c_p + h + c_p + \alpha \lim_{x \rightarrow \infty} J_{t+1}^r(x, A_t) > -(h + c_p)/\alpha, \text{ and} \\ \lim_{x \rightarrow -\infty} J_t^r(x, A_t) &= -c_s + c_a - c_a - c_p < -(c_p + c_s - p)/\alpha \end{aligned}$$

Hence (4.2'''), (4.3'''), and (4.4''') hold for $J_t(x)$ as well. Then we can repeat this argument sequentially for $t = T, T - 1, \dots, 1$, which completes the proof. \square

4.5 Logistics Agreements with Flexible Shipments

In this section we consider a different contract which gives more flexibility to the manufacturer. Instead of fixed scheduled dates for shipments as in Section 4.1, this time we assume that only the frequency of shipments, say one shipment of reserved capacity C per T periods, is specified and the manufacturer is allowed to use this reserved capacity whenever she wishes during those T periods. Obviously this type of contract enables the manufacturer to use the contracted capacity more freely while making it more difficult for the contract provider to allocate its resources efficiently.

In this model, at any period we need to keep track of whether we have used the contracted shipment or not. Hence we augment the state space to a vector (x, y) , where x represents the inventory level at the retailer as before and y represents the number of contracted shipments left. Let us adopt the shorthand notation $L(x) \equiv h(x)^+ + p(x)^-$ to simplify notational exposition.

The dynamic programming recursions for this model at $t = 1, \dots, T$ are:

$$\begin{aligned} J_t(x, 1) &= \min \left\{ \min_{u \geq 0} \left\{ (c_s + c_p)u + \mathbb{E}_w[L(x + u - w) + \alpha J_{t+1}(x + u - w, 1)] \right\}, \right. \\ &\quad \left. \min_{u \geq 0} \left\{ c_s(u - C)^+ + c_p u + \mathbb{E}_w[L(x + u - w) + \alpha J_{t+1}(x + u - w, 0)] \right\} \right\} \end{aligned}$$

and

$$J_t(x, 0) = \min_{u \geq 0} \left\{ (c_s + c_p)u + \mathbb{E}_w[L(x + u - w) + \alpha J_{t+1}(x + u - w, 0)] \right\}.$$

Unfortunately, we have been unable to characterize the optimal policy structure for this model so far. However, based on numerical experiments, we conjecture the following for determining the optimal time of the shipment.

Conjecture 4.5.1. For a given C , there is a sequence of critical numbers $\{S_t\}_{t=1}^T$ between 0 and C such that if $x_t < S_t$ then the manufacturer does not call for the contracted shipment at period t , else if $x_t \geq S_t$ then the manufacturer calls for the contracted shipment at period t .

Even if the conjecture above is true, we still need to specify the optimal production and shipping function up to the period in which the shipment is called to completely characterize the optimal policy structure. For this we conjecture the following.

Conjecture 4.5.2. (Optimal Policy) The optimal policy function is given by:

$$\mu_t(x, y) = \begin{cases} (P_t - x)^+, & \text{if } y = 0 \\ \begin{cases} ((R_t - x)^+, 0), & \text{if } K_t < x \\ (S_t - x, 1), & \text{if } S_t - C \leq x \leq K_t \\ (C, 1), & \text{if } s_t - C \leq x \leq S_t - C \\ (s_t - x, 1), & \text{if } x \leq s_t - C \end{cases} & \text{if } y = 1. \end{cases}$$

Furthermore, $s_t \leq S_t$, $K_t \leq S_t$, $R_t \leq P_t$, and s_t, S_t, R_t are decreasing while K_t is increasing in t .

Chapter 5

Computational Study

In this chapter, to gain managerial insights into practical real-life applications of the logistics agreements we considered, we conduct numerical experiments to answer questions such as the following:

- How do key optimal contract parameters for each contract, such as the capacity to be reserved and the frequency of shipments, change with changes in demand variability?
- How does demand variance affect the value of each logistics agreement, which is measured by the difference in optimal expected operational cost of a spot market only system and a system with access to the corresponding logistics agreement?
- In what ways do operational cost parameters such as holding cost and penalty cost affect the value of each logistics agreement?
- How do the performance of different logistics agreements compare under different operating environments?

We focus on analysis of optimal contract parameters with varying operating costs and demand in Section 5.1, while comparing different logistics agreements under varying operating environments in Section 5.2.

5.1 Optimal Contract Parameters

In this section, we explicitly specify the cost structure of a fixed commitment contract and computationally analyze how the optimal contract parameters change with respect to changes in the cost and demand parameters to facilitate gaining managerial insights in the logistics outsourcing problem of the contract buyer.

We use the following contract cost structure in our computational experiments, which captures some interesting discount parameters:

$$F(C) = \beta c_s C^\gamma.$$

In this contract cost form, we call β the unit discount since reserving unit capacity in the contract is $(1 - \beta)100$ percent cheaper than the expedited shipping rate in the spot market. On the other hand, we call γ the additional unit discount since when it is less than one each additional unit reserved in the transportation contract decreases the average cost of reserving capacity in the contract.

Then the optimal total expected cost function for the periodic contract with reserved capacity level C and time between shipments T is:

$$V(x, C, T) = \frac{\alpha^T}{1 - \alpha^T} \beta c_s C^\gamma + J_{\infty,0}(x, C, T).$$

In the rest of this section we will use this total cost function for our numerical experiments.

In Figure 5.1 and Figure 5.2, we plot the expected total cost as a function of reserved capacity level for varying time between shipments to illustrate the effect of the unit discount on the optimal transportation contract. One expects that when the cost of reserving capacity in the transportation contract decreases, the buyer's demand for the contract should increase. In fact when we give more unit discount in the transportation contract by decreasing β from 0.98 in Figure 5.1 to 0.90 in Figure 5.2, we observe that the optimal reserved capacity increases from 1 to 4 while the optimal time between shipments decreases from 3 to 2. Hence, the buyer increases both the reserved capacity and frequency of shipments in increasing its demand for the contract.

Figure 5.3 and Figure 5.4, on the other hand, illustrate the effect of the additional unit discount on the optimal transportation contract. We again plot the expected total cost as a function of reserved capacity level for varying time between shipments and once more expect that when there is more additional unit discount the buyer should increase its demand for the transportation contract. Indeed, when we give more additional unit discount in the transportation contract by decreasing γ from 1 in Figure 5.3 to 0.9 in Figure 5.4, we observe that the optimal reserved capacity increases from 4 to 16, while the optimal time between shipments also increase from 2 to 4. Although the observation that buyer decreases the frequency of shipments when there is more additional unit discount seems unintuitive at first glance, we note that the reserved capacity per period in the contract is increased from 2 to 4. Hence the demand rate for the transportation contract again increases with more additional unit discount, but in this case, unlike in the unit discount case, the buyer increases the time between shipments to take advantage of the additional unit discount without reserving excess capacity.

Another point of interest is to investigate how demand uncertainty effects the optimal contract parameters or the desirability of the transportation contract for the contract buyer.




Figure 5.1: The effect of unit discount, $\beta=0.98$.




Figure 5.2: The effect of unit discount, $\beta=0.90$.




Figure 5.3: The effect of additional unit discount, $\gamma=1$.




Figure 5.4: The effect of additional unit discount, $\gamma=0.90$.




Figure 5.5: The effect of demand variance on total cost.




Figure 5.6: The effect of demand variance on value of a contract.

In Figure 5.5 and Figure 5.6, we respectively plot the expected total cost and value of the contract as functions of reserved capacity level in the contract for a fixed frequency of shipments where the different curves represent different demand variances. We calculate the expected value of a contract by subtracting the expected operational cost of the optimal contract from the optimal expected operational cost of the contract with zero capacity. We observe that the expected total cost decreases while the value of a contract increases with decreasing uncertainty in demand. This observation is actually confirmed by our earlier remark that the fixed commitment transportation contracts gather the risk related to the demand uncertainty on the contract buyer. Hence the decreasing demand variance makes the transportation contract more valuable for the buyer.

In Figure 5.7, we plot the expected total cost with respect to time between shipments for a fixed level of reserved capacity where the different curves correspond to different levels of demand uncertainty as before. As the minimums of the curves shift left with decreasing demand variance we conclude that the buyer increases the frequency of shipments when there is less uncertainty in demand. In Table 5.1, we list the optimal contract parameters and the corresponding optimal cost values for varying demand variances. Unlike the frequency of shipments, we see that the reserved capacity levels do not always increase with decreasing demand variance. This seemingly unintuitive behavior can be easily explained by looking at the corresponding reserved capacity levels per period and noticing that they always increase with decreasing demand, which once again agrees with our earlier observations.

We also look at how the relative magnitudes of the waiting cost and the other operational cost parameters effect the optimal transportation contract. In Figure 5.8, we plot the expected total cost function with respect to reserved capacity level for a fixed frequency of shipments where different curves represent different ratios of the waiting cost to the production cost. Since the nominal values of the cost parameters would be misleading, we look at the ratios instead. We notice that as the relative magnitude of the waiting cost decreases, the curves in Figure 5.8 shift downwards and right, which indicates that the optimal reserved capacity in the contract increases as the expected total cost decreases. In Table 5.2, we list the optimal contract parameters for different ratios of waiting cost to the production cost. As is the case with Figure 5.8, we expect the buyer to increase its demand for the transportation contract as the penalty for making customers wait relatively decreases. Although the reserved capacity levels in the contract always increase with decreasing relative magnitude of the waiting cost as expected, we notice that the time between shipments also increases in the first row. This drop in the frequency of shipments may seem somewhat unintuitive but once more we observe that the reserved capacity in the contract per period always increases. This can be explained by the intuitive argument that as the penalty of making customers wait gets less severe, the contract buyer wants to ship a larger portion of demand via contracted shipment to utilize the discount on the shipping rate compared to the spot market rate.

Lastly Table 5.3 summarizes our numerical results from this section.




Figure 5.7: The effect of demand variance on frequency of shipments.

Demand Variance	Optimal Contract Parameters			Optimal Cost
	C	T	C/T	$V(0, C, T)$
8.25	4	2	2.00	838.17
5.11	5	2	2.50	824.93
2.25	5	2	2.50	809.66
0.49	4	1	4.00	790.12

Table 5.1: The effect of demand variance on optimal contract parameters.

Waiting Cost Ratio	Optimal Contract Parameters		
p/c	C	T	C/T
0.075	8	3	2.67
0.100	4	2	2.00
0.125	3	2	1.50
0.150	2	2	1.00

Table 5.2: The effect of waiting cost ratio on optimal contract parameters.




Figure 5.8: The effect of waiting cost ratio on optimal contract capacity.

Parameter	Optimal Contract		
	C	$1/T$	C/T
$\beta \downarrow$	\uparrow	\uparrow	\uparrow
$\gamma \downarrow$	\uparrow	\downarrow	\uparrow
$\sigma \downarrow$	$\downarrow - \uparrow$	\uparrow	\uparrow
$p/c \downarrow$	\uparrow	\downarrow	\uparrow

Table 5.3: Summary results illustrating how the optimal contract parameters change with respect to costs and demand variability.

5.2 Comparison of Contracts Types

In this section, we present numerical experiments to illustrate how the optimal contract parameters change with varying uncertainty in demand distribution. We will also compare the performances of the contracts we analyzed in previous sections and try to quantify the value of the contract flexibility for the buyer.




Figure 5.9: Optimal total expected cost vs reserved capacity level in the fixed commitment contract with $c_r = 7$, $c_p = 5$, $c_s = 10$, $h = 1$, $p = 5$, $T = 3$, $\alpha = 0.95$.

Figure 5.9 plots the optimal total expected cost as a function of the reserved capacity level in the fixed commitment contract. Different colors represent demand distributions with various levels of uncertainty. Notice that each curve's intercept with the vertical axis corresponds to the total expected cost when only the spot market is used for transportation. Since all the cost curves decrease initially we infer that the fixed commitment contracts are useful. In other words they can be used to lower the costs. Furthermore notice that the cost curves shift downward as the uncertainty in demand decreases. These general observations are intuitive and will hold for other contract types as well.

One last observation specific to fixed commitment contracts is that the minimums of these curves shift right as the demand uncertainty decreases. This means that the fixed commitment contract becomes more desirable for the buyer when the demand uncertainty decreases. This confirms our earlier remark that the fixed commitment contracts gather the risk related to the demand uncertainty on the contract buyer.




Figure 5.10: Optimal total expected cost vs reserved capacity level in the option contract with $c_r = 2.5, c_e = 5, c_p = 5, c_s = 10, h = 1, p = 5, T = 3, \alpha = 0.95$.




Figure 5.11: Optimal total expected cost vs reserved capacity level in the option contract with $c_r = 1, c_e = 6.5, c_p = 5, c_s = 10, h = 1, p = 5, T = 3, \alpha = 0.95$.

Figure 5.10 and Figure 5.11 plot the optimal total expected cost as a function of the reserved capacity level in the option contract. Different colors represent demand distributions with various levels of uncertainty as before. Unlike the fixed commitment contract, which always becomes more attractive for the buyer as the demand uncertainty decreases, the option contracts may turn out either way. When the reservation price is relatively high, as in Figure 5.10, the option contract shows the same behavior as the fixed commitment contracts. However when the reservation price is relatively low, as in Figure 5.11, as the demand uncertainty increases, a spike in demand becomes more and more probable, hence the buyer demands more capacity in the option contract to hedge herself against the high spot market rate.




Figure 5.12: Optimal total expected cost vs reserved capacity level in the capacity flexible contract with $c_r = 2.25$, $c_e = 5$, $\gamma = 0.15$, $c_p = 5$, $c_s = 10$, $h = 1$, $p = 5$, $T = 3$, $\rho = 0.95$.

As explained previously, capacity flexible contracts can be seen as a bundle of fixed commitment and option contracts. Thus their behavior is generally somewhere between them as illustrated by Figure 5.12. In general the optimal reserved capacity levels may increase or decrease because of the same reasons we mentioned previously.

In Table 5.4, Table 5.5 and Table 5.6, we list the optimal contract parameters and the corresponding optimal cost values for varying demand variances for the fixed commitment, option and capacity flexible contracts respectively. In contrast to the frequency of shipments, we see that the reserved capacity levels do not always increase with decreasing demand variance. This seemingly unintuitive behavior can be easily explained by looking at the

corresponding reserved capacity levels per period and noticing that they always increase with decreasing demand, which once again agrees with our earlier observations.

Demand Variance	Optimal Contract Parameters			Optimal Cost
	C	T	C/T	$V(0, C, T)$
8.25	4	2	2.00	1354.61
5.11	5	2	2.50	1240.35
2.25	5	2	2.50	1203.14
0.49	4	1	4.00	1138.52

Table 5.4: Fixed Commitment Contract: The effect of demand variance on optimal contract parameters.

Demand Variance	Optimal Contract Parameters			Optimal Cost
	C	T	C/T	$V(0, C, T)$
8.25	7	3	2.33	1345.11
5.11	6	2	3.00	1223.32
2.25	6	2	3.00	1142.65
0.49	6	1	6.00	1113.34

Table 5.5: Option Contract: The effect of demand variance on optimal contract parameters.

Demand Variance	Optimal Contract Parameters			Optimal Cost
	C	T	C/T	$V(0, C, T)$
8.25	6	3	2.00	1348.87
5.11	5	2	2.50	1235.43
2.25	6	2	3.00	1167.54
0.49	5	1	5.00	1132.96

Table 5.6: Capacity Flexible Contract: The effect of demand variance on optimal contract parameters.

Figure 5.13 shows the optimal expected total cost under different levels of uncertainty in demand, where different colors represent the different contract types. Remarkably, the fixed commitment contract, which gives the least amount of flexibility for the lowest unit cost, is outperformed only when the demand uncertainty is high.

Lastly, we conduct a small sensitivity analysis and come up with the pricing menu in Table 5.7 that would make a buyer approximately indifferent between each contract. Hence

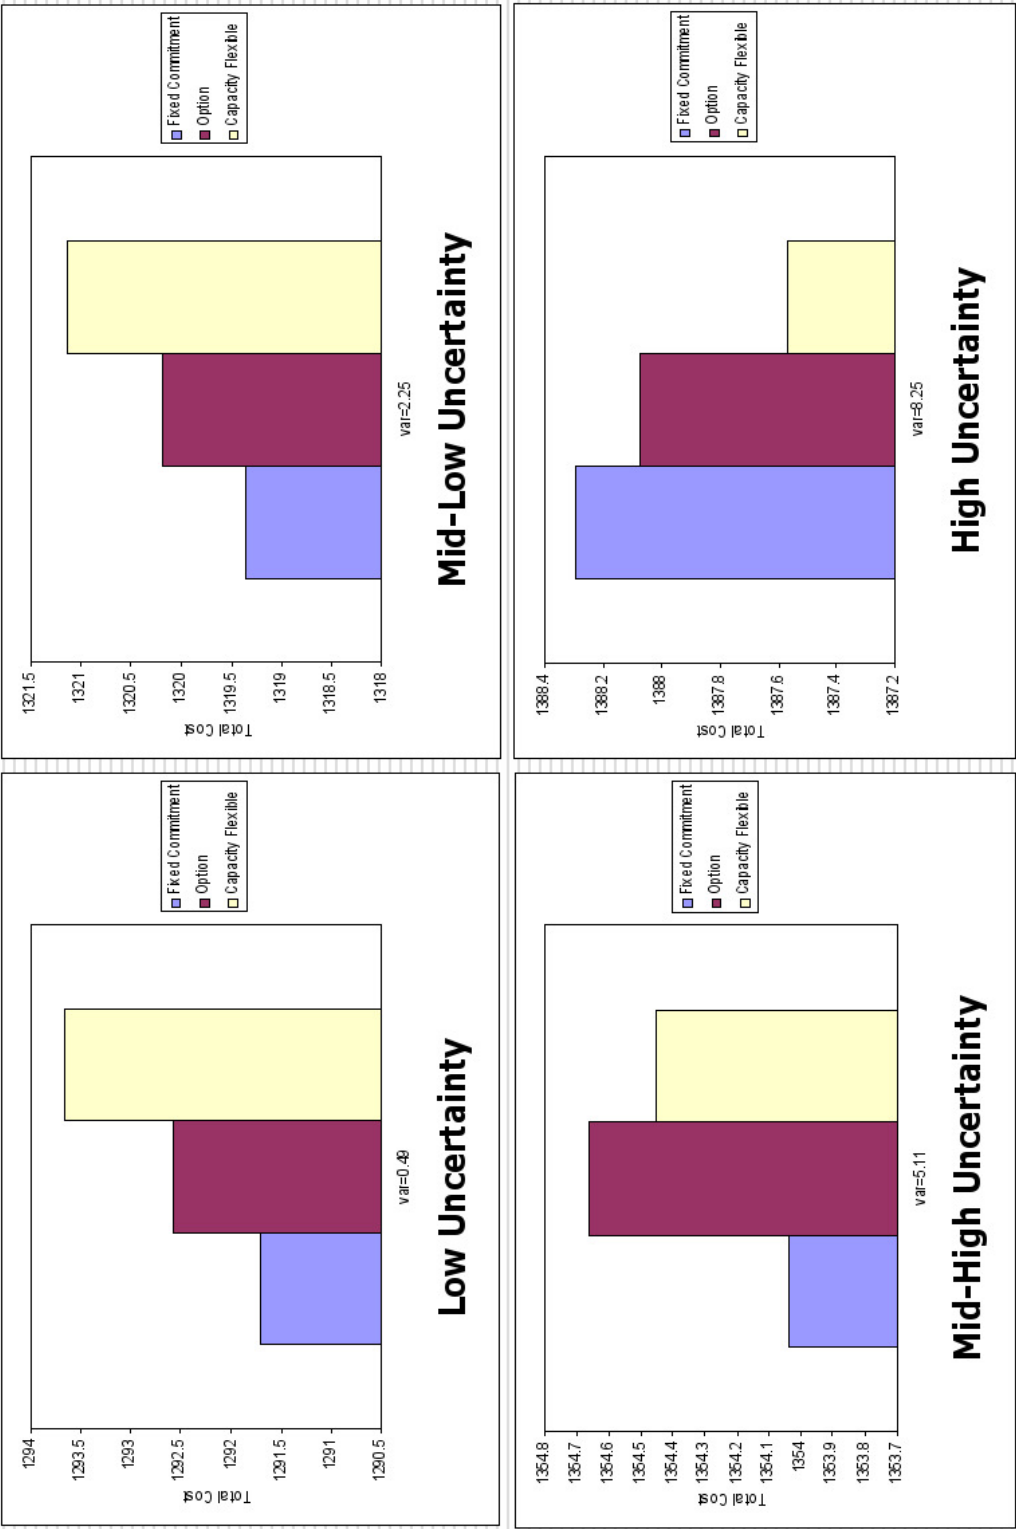


Figure 5.13: Performance comparison of the contracts.

Contract Type	Reservation Price	Exercise Price	Increase in Unit Cost
Fixed Commitment	\$7.00	\$0.00	0.00%
Option	\$2.50	\$5.00	7.14%
Capacity Flexible	\$2.25	\$5.00	3.57%

Table 5.7: The value of contract flexibility for the buyer, $\gamma = 0.15$, $c_p = 5$, $c_s = 10$, $h = 1$, $p = 5$, $T = 3$, $\sigma = 0.95$.

compared to the fixed commitment contract a buyer would be willing to pay approximately \$0.50 (or 7.14%) more per unit for the flexibility of an option contract, and \$0.25 (or 3.57%) more per unit for the flexibility of a capacity flexible contract.

Chapter 6

Conclusion

6.1 Summary of Results

In this dissertation, we set out to investigate the effective utilization of structured logistics agreements in production-distribution systems with stochastic demand in order to understand how optimal policies of such systems are affected by outsourced logistics while gaining managerial insights into practical real-life applications and contributing to the stochastic supply chain literature.

As our base model, we consider a manufacturer that needs to ship its product to a retail site to meet stochastic demand. In addition to a logistics agreement already at hand, the manufacturer has access to a spot market for on demand shipping at a linear rate. The production is assumed to have linear cost with neither delay nor capacity. We explore the optimal operating strategies of this system under varying operating conditions and with different types of agreements.

In Chapter 3, we consider a make-to-order production environment and a fixed-commitment logistic contract with periodic scheduled shipments. This periodic contract allows the manufacturer to ship essentially freely – since the cost of logistic contract is sunk – in periods that have reserved shipping capacity. Theorem 3.4.2 completely characterizes the optimal shipping policy for this problem. This result shows the existence of nonnegative real numbers that essentially represent the ideal level of pending orders, such that it is optimal to ship enough orders to bring the number of pending orders as close to this ideal level as possible. We also show that these ideal levels of pending orders are increasing with the number of time periods to the next scheduled shipment. Proposition 3.4.3 proves that increasing reserved capacity levels in the contract results in an increase in the ideal levels of pending orders, while Proposition 3.3.5 determines that all the ideal levels shift up by exactly the same amount as the increase in reserved capacity level when there is only one scheduled shipment. Section 3.2 analyzes the case of unbounded reserved capacity levels. Proposition 3.2.1 demonstrates that the ideal levels diverge to infinity with unbounded

reserved capacity levels and more importantly proves that any contracted shipment can have effect on a limited and quantifiable number of periods prior to its scheduled shipment. This observation establishes the foundation of a sufficient condition for decomposing the problem in time that we present in Corollary 3.4.4.

In Chapter 4, we extend our model to a make-to-stock production environment and consider more sophisticated logistics agreements as well as additional uncertainties in the system. Section 4.1 extends the analysis of the fixed-commitment contract to a make-to-stock system. The optimal policy structure is fully characterized by Theorem 4.1.1, which shows that there are now two critical levels of inventory. These two critical levels together with the reserved shipment capacity in the contract divide the state space (of inventory at the retailer) into four partitions such that in the lowest partition it is optimal to ship up to the smaller critical level using all the contracted capacity plus spot market shipping as needed, in the second lowest partition it is optimal to fully utilize the contracted capacity, in the second highest partition it is optimal to ship up to the larger critical level using only the contracted shipping, while in the highest partition it is optimal to have no shipment at all. The additional critical level in the make-to-stock model compared to the make-to-order model is clearly a consequence of allowing the manufacturer to stock inventory at the retail site instead of just reacting to pending orders. Proposition 4.1.2 and Proposition 4.1.3 extend the results of Chapter 3 regarding the monotonicity of optimal policy with respect to reserved capacity levels and time to next shipment (in periodic contracts) respectively. Theorem 4.2.1 shows that a similar result holds for optional agreements when the spot market rate is stochastic. Section 4.3 analyzes multi-level option agreement, which allows the buyer to increase the capacity reservation up to a certain percentage with additional cost. Theorem 4.3.1 gives the optimal policy function for this agreement, which has similar characteristics to the fixed-commitment contract but has three critical parameters instead of two. Additional parameter in the optimal policy function stems from the provision in the agreement that provides extra capacity. Theorem 4.4 proves that when the availability of additional capacity is uncertain, a similar result holds under certain conditions.

In Chapter 5, we present a computational study to provide managerial insights into the practical real-life applications of the logistics agreements we consider in this manuscript. Our analysis and numerical experiments show that, for a manufacturing firm that wants to outsource its logistics operations, it is significantly more attractive to buy a transportation contract rather than using only the spot market when pending cost is relatively small compared to other operational costs. However, when pending cost is relatively large, especially in the case of large demand variance, a substantial increase in the frequency of shipments is necessary for the contract to be useful, which could render the contract price too expensive to be profitable under most reasonable contract cost structures. In general, we have found through our computational study that the larger the demand uncertainty the less attractive the transportation contracts and also while unit discounts have the effect of increasing the demand for the transportation contract in both reserved capacity levels and frequency of shipments, additional unit discount increases the reserved capacity levels while decreasing

the frequency of shipments in the optimal contract. Lastly, we observed in our numerical experiments that more structured fixed-commitment contract outperforms more flexible contracts in all cases except when demand variability is very large.

6.2 Future Directions

In all the models we consider in this dissertation, we assume that the production has neither lead time nor capacity and that it has linear cost. These assumptions make the production decision couple with the shipment whereby significantly simplifying the analysis of mathematical models. However, these assumptions may not be realistic in many real-life scenarios. Although extensions of the production process in these directions may render the problem too complicated to solve completely, we envision that efficient policies, which may have influential practical implications, can be devised.

In this dissertation, our work focuses on determining the characteristics of optimal operation in simple production-distribution systems with logistics contracts. In our view, other important areas to work on the overall logistics contracting framework are the optimal contract purchasing problem of manufacturer given a menu of contracts, as well as the corresponding problem of logistics provider for designing and pricing of these contracts.

Bibliography

- [1] O. Alp, N.K. Erkip, and R. Gullu. Outsourcing logistics: Designing transportation contracts between a manufacturer and a transporter. *Transportation Science*, 37(1):23, 2003.
- [2] R. Anupindi, Y. Bassok, S. Tayur, R. Ganeshan, and M. Magazine. Supply Contracts with Quantity Commitments and Stochastic Demand: Quantitative Models for Supply Chain Management, 1999.
- [3] K.J. Arrow. The Genesis of "Optimal Inventory Policy". *Operations Research*, 50(1):1–2, 2002.
- [4] K.J. Arrow, T. Harris, and J. Marschak. Optimal inventory policy. *Econometrica: Journal of the Econometric Society*, 19(3):250–272, 1951.
- [5] S. Axsäter and K. Rosling. Installation vs. echelon stock policies for multilevel inventory control. *Management Science*, 39(10):1274–1280, 1993.
- [6] D.P. Bertsekas. *Dynamic Programming and Optimal Control, vols. 1, 2*. Athena Scientific, 2001.
- [7] D.P. Bertsekas and S.E. Shreve. *Stochastic optimal control: The discrete time case*. Athena Scientific, 1996.
- [8] D. Beyer and S.P. Sethi. Average cost optimality in inventory models with Markovian demands. *Journal of Optimization Theory and Applications*, 92(3):497–526, 1997.
- [9] L.D. Burns, R.W. Hall, D.E. Blumenfeld, and C.F. Daganzo. Distribution strategies that minimize transportation and inventory costs. *Operations Research*, 33(3):469–490, 1985.
- [10] R.C. Carlson and C.A. Yano. Safety stocks in MRP-systems with emergency setups for components. *Management Science*, pages 403–412, 1986.
- [11] F. Chen. Echelon reorder points, installation reorder points, and the value of centralized demand information. *Management Science*, 44(12):221–234, 1998.

- [12] F. Chen. Optimal policies for multi-echelon inventory problems with batch ordering. *Operations research*, pages 376–389, 2000.
- [13] F. Chen and J.S. Song. Optimal policies for multiechelon inventory problems with Markov-modulated demand. *Operations Research*, 49(2):226–234, 2001.
- [14] F. Chen and Y.S. Zheng. Evaluating echelon stock (R, nQ) policies in serial production/inventory systems with stochastic demand. *Management Science*, 40(10):1262–1275, 1994.
- [15] F. Chen and Y.S. Zheng. Lower bounds for multi-echelon stochastic inventory systems. *Management Science*, 40(11):1426–1443, 1994.
- [16] A.J. Clark. Multi-echelon inventory theory—A retrospective. *International Journal of Production Economics*, 35(1-3):271–275, 1994.
- [17] A.J. Clark. Comments on” Optimal Policies for a Multi-Echelon Inventory Problem”. *Management Science*, 50(12):1794–1795, 2004.
- [18] A.J. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management science*, 6(4):475–490, 1960.
- [19] A.J. Clark and H. Scarf. Approximate solutions to a simple multi-echelon inventory problem. *Studies in applied probability and management science*, pages 88–110, 1962.
- [20] D. Coffey, Kesteloo M., and R. Spieker. Third-Party Logistics: The Industry That Needs to Relax, 2007.
- [21] E.B. Diks and A.G. De Kok. Computational results for the control of a divergent N-echelon inventory system. *International Journal of Production Economics*, 59(1-3):327–336, 1999.
- [22] R. Ehrhardt. The power approximation for computing (s, S) inventory policies. *Management Science*, 25(8):777–786, 1979.
- [23] G. Eppen and L. Schrage. Centralized ordering policies in a multi-warehouse system with lead times and random demand. *Multi-level production/inventory control systems: Theory and practice*, 16:51–67, 1981.
- [24] A. Federgruen. Centralized planning models for multi-echelon inventory systems under uncertainty. *Logistics of Production and Inventory*, 4:133–174, 1993.
- [25] A. Federgruen, P.J. Schweitzer, and H.C. Tijms. Denumerable undiscounted semi-Markov decision processes with unbounded rewards. *Mathematics of Operations Research*, pages 298–313, 1983.

- [26] A. Federgruen and P. Zipkin. Allocation policies and cost approximations for multilocation inventory systems. *Naval Research Logistics Quarterly*, 31(1):97–129, 1984.
- [27] A. Federgruen and P. Zipkin. An efficient algorithm for computing optimal (s, S) policies. *Operations research*, 32(6):1268–1285, 1984.
- [28] A. Federgruen and P. Zipkin. Computational issues in an infinite-horizon, multiechelon inventory model. *Operations Research*, 32(4):818–836, 1984.
- [29] A. Federgruen and P. Zipkin. An inventory model with limited production capacity and uncertain demands I. The average-cost criterion. *Mathematics of Operations Research*, 11(2):193–207, 1986.
- [30] A. Federgruen and P. Zipkin. An inventory model with limited production capacity and uncertain demands II. The discounted-cost criterion. *Mathematics of Operations Research*, 11(2):208–215, 1986.
- [31] J.R. Freeland and E.L. Porteus. Evaluating the effectiveness of a new method for computing approximately optimal (s, S) inventory policies. *Operations Research*, 28(2):353–364, 1980.
- [32] G. Gallego and H. Hu. Optimal Policies for Production/Inventory Systems with Finite Capacity and Markov-Modulated Demand and Supply Processes. *Annals of Operations Research*, 126(1):21–41, 2004.
- [33] G. Gallego and A. Scheller-Wolf. Capacitated inventory problems with fixed order costs: Some optimal policy structure. *European Journal of Operational Research*, 126(3):603–613, 2000.
- [34] G. Gallego and D. Simchi-Levi. On the effectiveness of direct shipping strategy for the one-warehouse multi-retailer R-systems. *Management Science*, 36(2):240–243, 1990.
- [35] P. Glasserman and S. Tayur. The stability of a capacitated, multi-echelon production-inventory system under a base-stock policy. *Operations Research*, 42(5):913–925, 1994.
- [36] P. Glasserman and S. Tayur. Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science*, 41(2):263–281, 1995.
- [37] P. Glasserman and S. Tayur. A simple approximation for a multistage capacitated production-inventory system. *Naval Research Logistics*, 43(1):41–58, 1996.
- [38] G. Hadley and T.M. Whitin. A family of inventory models. *Management Science*, 7(4):351–371, 1961.

- [39] M. Henig, Y. Gerchak, R. Ernst, and D.F. Pyke. An inventory model embedded in designing a supply contract. *Management Science*, 43(2):184–189, 1997.
- [40] D.P. Heyman and M.J. Sobel. *Stochastic Models in Operations Research: Stochastic Optimization*. McGraw-Hill College, 1984.
- [41] D. Hochstaedter. An approximation of the cost function for multi-echelon inventory model. *Management Science*, pages 716–727, 1970.
- [42] D. Iglehart. Dynamic Programming and Stationary Analysis of Inventory Problems. Chapter 1 in *Multistage Inventory Models and Techniques*, H. Scarf, D. Gilford, and M. Shelly, 1963.
- [43] D.L. Iglehart. Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Science*, 9(2):259–267, 1963.
- [44] E.L. Johnson. On (s, S) policies. *Management Science*, pages 80–101, 1968.
- [45] R.S. Kaplan. A dynamic inventory model with stochastic lead times. *Management Science*, 16(7):491–507, 1970.
- [46] R. Kapuściński and S. Tayur. A capacitated production-inventory model with periodic demand. *Operations Research*, 46(6):899–911, 1998.
- [47] S. Karlin. Inventory Models of the Arrow-Harris-Marschak Type with Time Lag. Chapter 10 in *Studies in the Mathematical Theory of Inventory and Production* K.J. Arrow, S. Karlin, and H. Scarf, 1958.
- [48] S. Karlin. Dynamic inventory policy with varying stochastic demands. *Management Science*, 6(3):231–258, 1960.
- [49] S. Karlin. Optimal policy for dynamic inventory process with stochastic demands subject to seasonal variations. *Journal of the Society for Industrial and Applied Mathematics*, 8(4):611–629, 1960.
- [50] U.S. Karmarkar. The multiperiod multilocation inventory problem. *Operations Research*, 29(2):215–228, 1981.
- [51] U.S. Karmarkar. The multilocation multiperiod inventory problem: Bounds and approximations. *Management science*, pages 86–94, 1987.
- [52] D.M. Lambert and J.R. Stock. *Strategic logistics management*. Irwin, 1993.
- [53] T.E. Morton. Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM review*, 11(4):572–596, 1969.

- [54] A. Muharremoglu and J.N. Tsitsiklis. A single-unit decomposition approach to multi-echelon inventory systems. *Operations research*, 56(5):1089–1103, 2008.
- [55] S. Nahmias. Simple approximations for a variety of dynamic leadtime lost-sales inventory models. *Operations Research*, 27(5):904–924, 1979.
- [56] J.M. Norman and D.J. White. A method for approximate solutions to stochastic dynamic programming problems using expectations. *Operations Research*, 16(2):296–306, 1968.
- [57] R.P. Parker and R. Kapuscinski. Optimal policies for a capacitated two-echelon inventory system. *Operations Research*, 52(5):739–755, 2004.
- [58] E.L. Porteus. An adjustment to the Norman-White approach to approximating dynamic programs. *Operations Research*, 27(6):1203–1208, 1979.
- [59] E.L. Porteus. Numerical comparisons of inventory policies for periodic review systems. *Operations Research*, pages 134–152, 1985.
- [60] E.L. Porteus. *Foundations of stochastic inventory theory*. Stanford Business Books, 2002.
- [61] J.P. Quinn. Global 3PL Growth Taking Off, 2007.
- [62] D.M. Roberts. Approximations to optimal policies in a dynamic inventory model. Chapter 13 in *Studies in Applied Probability and Management Science*, eds. K. Arrow, S. Karlin, and H. Scarf, 1962.
- [63] D. Rosenfield and M. Pendrock. The effects of warehouse configuration design on inventory levels and holding costs. *Sloan Management Review*, 21(4):21–33, 1980.
- [64] K. Rosling. Optimal inventory policies for assembly systems under random demands. *Operations Research*, 37(4):565–579, 1989.
- [65] R. Roundy. 98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. *Management Science*, 31(11):1416–1430, 1985.
- [66] R. Roundy. A 98%-effective lot-sizing rule for a multi-product, multi-stage production/inventory system. *Mathematics of Operations Research*, 11(4):699–727, 1986.
- [67] W. Rudin. *Real and complex analysis (3rd)*. New York: McGraw-Hill Inc, 1986.
- [68] H. Scarf. The Optimality of (s, S) Policies in Dynamic Inventory Problems. K. Arrow, S. Karlin, P. Suppes, eds., *Mathematical Models in the Social Sciences*, 1960.

- [69] H. Scarf. Comments on "Optimal Policies for a Multi-Echelon Inventory Problem". *Management Science*, 50(12):1791–1793, 2004.
- [70] H.E. Scarf. Inventory theory. *Operations Research*, 50(1):186–191, 2002.
- [71] C.P. Schmidt and S. Nahmias. Optimal policy for a two-stage assembly system under random demand. *Operations Research*, 33(5):1130–1145, 1985.
- [72] C. Shaoxiang. The infinite horizon periodic review problem with setup costs and capacity constraints: A partial characterization of the optimal policy. *Operations Research*, 52(3):409–421, 2004.
- [73] C. Shaoxiang and M. Lambrecht. XY band and modified (s, S) policy. *Operations Research*, 44(6):1013–1019, 1996.
- [74] David. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi. *Designing and managing the supply chain*. Irwin McGraw-Hill, 2003.
- [75] A.F. Veinott Jr. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Science*, 12(3):206–222, 1965.
- [76] A.F. Veinott Jr. The optimal inventory policy for batch ordering. *Operations Research*, 13(3):424–432, 1965.
- [77] A.F. Veinott Jr. The status of mathematical inventory theory. *Management Science*, 12(11):745–777, 1966.
- [78] A.F. Veinott Jr and H.M. Wagner. Computing optimal (s, S) inventory policies. *Management Science*, 11(5):525–552, 1965.
- [79] H.M. Wagner, M. O'Hagan, and B. Lundh. An empirical study of exactly and approximately optimal inventory policies. *Management Science*, 11(7):690–723, 1965.
- [80] AS Whittmore and SC Saunders. Optimal inventory under stochastic demand with two supply options. *SIAM Journal on Applied Mathematics*, 32(2):293–305, 1977.
- [81] J. Wijngaard. An inventory problem with constrained order capacity, 1972.
- [82] C.A. Yano. Optimizing transportation contracts to support just-in-time deliveries: the case of one contracted truck per shipment. *IIE transactions*, 24(2):177–183, 1992.
- [83] C.A. Yano and Y. Gerchak. Transportation contracts and safety stocks for just-in-time deliveries. *Journal of Manufacturing and Operations Management*, 2:314–330, 1989.
- [84] Y.S. Zheng. On properties of stochastic inventory systems. *Management Science*, 38(1):87–103, 1992.

- [85] P. Zipkin. Critical number policies for inventory models with periodic data. *Management Science*, 35(1):71–80, 1989.
- [86] P.H. Zipkin. *Foundations of inventory management*. McGraw-Hill Boston, MA, 2000.