Evaluation of Cost Balancing Policies in Multi-Echelon Stochastic Inventory Control Problems

by

Qian Yu

B.Sc, Applied Mathematics, National University of Singapore(2008)

Submitted to the School of Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Computation for Design and Optimization

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Jan 2010

Author . . . .
School of Engineering
Jan, 2010

Certified by . . . .
Restef Levi
Associate Professor of Management
Thesis Supervisor

Accepted by . . . .
Karen Willcox
Associate Professor of Aeronautics and Astronautics
Codirector, MIT CDO Program

Evaluation of Cost Balancing Policies in Multi-Echelon Stochastic Inventory Control Problems

by

Qian Yu

Submitted to the School of Engineering
on Jan, 2010, in partial fulfillment of the
requirements for the degree of
Master of Science in Computation for Design and Optimization

Abstract

We study a periodic-reviewed, infinite horizon serial network inventory control problem. The demands in different periods are independent of each other and follow an identical Poisson distribution. Unsatisfied demands are backlogged until they are satisfied by supply units. In each period, there is a per-unit holding cost is incurred for each unit of supply that stays in the system and a per-unit backorder cost is incurred for each unsatisfied unit of demand. The objective of the inventory control policy is to minimize the long-run expected average cost over an infinite horizon.

The goal of the thesis is to evaluate the empirical performance of the dual-balancing policy and several other variants of cost balancing policies through numerical simulations. The dual-balancing policy is based on two novel ideas: the marginal cost accounting scheme, which assigns to each decision all the costs that are made inevitable after that decision is made; and the cost balancing idea to balance opposing costs. The dual-balancing policy can be modified in several ways to get other cost balancing policies.

It has been proven that the dual-balancing policy has a worst-case guarantee of 2 but this does not indicate the empirical performance. An approximately optimal policy is considered as the benchmark to test the quality of the cost balancing policies. In the computational experiments, the dual-balancing policy shows an average error of 7.74% compared to the approximately optimal policy, much better than the theoretical worst-case guanratee. The three variants of cost balancing policies have made significant improvement on the performance of the dual-balancing policy. The accuracy of the dual-balancing policy is also affected by the system parameters. In addition, with high demand rate and long lead times, we have observed several scenarios when the cost balancing policies dominate the approximately optimal policy.

Thesis Supervisor: Restef Levi
Title: Associate Professor of Management

Acknowledgments

I would like to thank first and foremost my supervisor, Prof. Restef Levi. He has shown incredible support and patience for my research. Even when he was on vacation, he always gave me prompt replies to whatever questions I have asked. Without him pushing me to my very best, this thesis will not be possible in such a short time. I am so glad to have him as my supervisor and learned so much during the process.

I would like to thank my fellow friends in the CDO program. We shared all the joys and pains in the past year. They have made this year much more meaningful to me. I would also like to thank Ms Laura Koller for her support and help on the submission of the thesis. I am grateful that people I met in the United States were so nice and friendly to me and made my life changed after all these. I owe my deepest gratitude to my dear family in China and friends in Singapore for their love and support at my hard times.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis, we study the inventory control problem of a *single-item periodic-reviewed serial network* with stochastic demands and infinite time horizon. This is one of the core problems in inventory management. The design of computationally efficient and provably good inventory control policies for these systems has been a fundamental yet challenging problem in inventory theory and practice. It arises in many domains and has many practical applications in supply chain management and logistics (see examples like Erkip et al. [13] and Lee et al. [14]).

We start with a detailed description of the inventory model. A single commodity moves through a supply chain that consists of an external supplier, distinct warehouses (called *stages or echelons*), to satisfy external demand. Each stage is supplied by the preceding stage in the serial network. The lowest stage (echelon) is facing a sequence of stochastic demands over the time horizon and the highest stage (echelon) is supplied by an external supplier with infinite capacity. There are *lead times* between each two consecutive stages, representing the number of periods it takes to ship a unit of commodity from one stage to the next. At the end of each period, several types of costs are incurred at each of the echelons in the network: a per-unit *ordering cost* for each unit that is first ordered in this period, a per-unit *holding cost*, for each unit of inventory that is on-hand at a stage or in transit between stages, and a per-unit *backordering penalty cost*, for each unit of demand that is not yet satisfied; the latter cost is incurred only at the lowest stage. Unsatisfied demand units are fully backlogged, i.e., wait in the system until they are satisfied. The goal is to find an inventory control policy that minimizes the long-run expected average cost.

Dynamic programming framework has been the most common paradigm to study these problems. It has been very effective in characterizing the structure of the optimal policies. In particular, *echelon base-stock policies* were proven to be optimal for multiple variants of the problem. The first optimality proof is due to Clark and Scarf [7] who study the serial network with independently and identically distributed demands. Subsequently, several researchers have also extended the proof to other models and introduced simpler proofs of optimality of the echelon base-stock policies. Specifically, echelon base-stock policies were proven optimal under more general assumptions on the demand distributions. For example, the optimality of the echelon base-stock policies in serial network has been established for the case where demands follow an exogenous Markov-modulated process [8], Poisson process [9] and compound-Poisson process [10]. Muharremoglu and Tsitsiklis [11] have proposed another simpler approach to prove the optimality of the echelon base-stock policies. Each unit of supply that has been ordered is matched with the corresponding demand unit that is satisfied by the supply unit. The problem is then decomposed into a series of unit supply-demand subproblems, where each subproblem corresponds to a pair of supply and demand unit that are matched together. They have established the optimality of the echelon base-stock policy for the uncapacitated model with Markov-modulated demand and lead times.

The structure of the echelon base-stock policies is rather simple. They are based on the concept of echelon inventory. The echelon inventory at a given stage (echelon) is defined as the total inventory that is at or in transit from that stage to the next stage, and that is at or in transit to any downstream stages in the network. For each stage, there is a target echelon inventory level, defined as *echelon base-stock level*. Whenever the echelon inventory level at the beginning of a period falls below the echelon base-stock level, that stage makes an order to bring up the echelon inventory level, as long as there is enough inventory on hand at the preceding stage; if the echelon inventory level exceeds the echelon base-stock level, no order is placed.

Unfortunately, the simple base-stock form of the optimal policy does not always lead to efficient algorithms for computing the order sizes. The dynamic programming approach can only be tractable in cases where the demands are independent at different periods and do not evolve over time, or in cases where the demands follow an exogenous Markov-modulated process with small-dimension state space [17]. Even in such scenarios, the computations can be tedeous and complex. The difficulty comes from the fact that we need to solve too many subproblems, a phenomenon called *the curse of dimensionality*. Therefore, the straightforward approach to solve the dynamic programs becomes theoretically and often also practically intractable in the presence of correlated and evolving demand distributions.

In recent work, Levi et al. [1] have introduced the *dual-balancing policies* for periodic-reviewed, single-stage inventory models. This is the first computationally efficient algorithm which also admits a worst-case performance guarantee of 2. That is, the expected cost following the dual-balancing policy is guaranteed to be at most twice the expected cost according to the optimal policy. It has been shown in computational experiments in [3] that the policy is computationally efficient and dominates other policies in many common scenarios. The dual-balancing policies were also extended to the multi-echelon stochastic inventory control models with correlated and evolving demands over a finite and infinite horizon (see [2]); this includes the model studied in this thesis. The dual-balancing policies are based on two novel ideas:

I. *Marginal Cost Accounting Scheme* Traditional cost accounting schemes are primarily based on dynamic programming and decompose the cost by stages and time periods. In contrast, the marginal cost accounting scheme decomposes the cost by ordering decisions and assign to each decision all the costs that, after the decision is made, become unaffected by any future decision, even if the costs occur in the future. To be specific, each unit of cost is assigned to a specific decision that caused that cost to be incurred. Similar ideas of marginal cost accounting has been used by Axsäter [4] for inventory control problems in continuous and infinite time horizon with Poisson demands.

There are several cost accounting schemes applied to the serial networks all based on dynamic programming approach. The first cost accounting scheme was established in the classical proof of the optimality of the echelon base-stock policy by Clark and Scarf [7]. The inventory problem is decomposed into a sequence of single-stage subproblems and each subproblem is solved for optimality, assuming that the lower stages have already been optimized. Using the same cost accounting scheme, Zipkin [15] has simplified the proof and later Dong and Lee [16] extended the optimality with time-correlated demand. Another cost accounting scheme was developed by Chen and Zheng for models involving fixed set up costs, that are incurred whenever a positive order is placed [12]. They create a component for each stage in the system and a separate multi-stage inventory model for each component. The cost parameters are then allocated to the different components.

Π. *Cost-Balancing* The simple idea of cost balancing is based on the key observation that any ordering policy incurs two potential opposing costs in the serial network. One of the costs increases with the order size while the other decreases to 0. Hence, it is always possible to order a quantity that will make these two equal.

Next we describe the main ideas underlying the marginal cost accounting scheme developed in [2] for the serial network. The approach is based on the notion of the *critical period*. If the demand arrival time is known in advance, the corresponding supply unit should be ordered in a just-in-time fashion at every stage to minimize the cost incurred. This time period is referred to as critical period. The holding costs are further categorized into three types and assigned to the corresponding decisions that made the cost inevitable. Holding costs incurred when units are in transit between stages are called *pipeline costs*. The pipeline costs are incurred for every unit that is ordered, regardless of the policy followed. Hence, they are not assigned to any decision that has been made. An *early holding cost* is incurred for a supply unit that has been ordered before the critical period. It is then assigned to the stages and periods when the unit is ordered earlier than the corresponding critical period of the stage. Similarly, a *late holding cost* is incurred for a supply unit that is ordered after the critical period. In addition, there are *backorder costs* incurred whenever a demand unit is not satisfied. Details of the assignment rules will be discussed in

Section 3.1.

The ordering decisions at different stages can be done separately. Given a specific ordering size at some stage, we assign all the costs that are made inevitable in the system and categorize the costs into: the conditional expected marginal early holding cost and the conditional expected marginal late holding and backorder cost. The conditional expected marginal early holding cost is an increasing function with the order size while the later cost is a decreasing function. This leads to the dual-balancing policy, which aims to select an order quantity to balance these two opposing cost functions. The policy is computationally efficient, and can be implemented in an on-line manner.

The dual-balancing policy can be modified in two ways to devise policies with better empirical performance. First, we use the interval-constrained bounding technique, which was first proposed by Levi et al [3] for the single-stage inventory control problem. Given the upper and lower bounds on the optimal base-stock level, whenever the after-order inventory level falls out of the bounds, the interval-constrained bounding procedure modifies the ordering decisions. Specifically, when the after-order inventory level exceeds the upper bound, the order size is reduced down so that the after-order inventory level equals the upper bound; if on the other hand it is smaller than the lower bound, then the order size is increased to an appropriate amount to bring up the after-order inventory level to the lower bound. The computational experiments in [3] had indicated that the typical performance of the dual-balancing policy with bounds is better than the dual-balancing policy without bounds. In the serial model with infinite horizon that we study in this thesis, we adopt the upper and lower bounds on the optimal echelon base-stock level developed by Shang and Song [6]. The modified policy works as follows: an order size is computed at each stage following the dual-balancing policy. We then consider the after-order echelon inventory level. If it is smaller than the lower bound of the optimal echelon base-stock level, the order size is increased as long as the on-hand inventory at the preceding stage is positive to bring up the after-order echelon inventory level to the lower bound. On the other hand, if the after-order echelon inventory level exceeds the upper bound, the order size is reduced down by an appropriate quantity. When the after-order echelon inventory level is between the bounds, we make no adjustment to the order size.

Another way to modify the dual-balancing policy is to introduce *parameterized balancing policies*, which were also first developed in [3] for the single-stage model. Here one aims to select an order quantity which balances the two opposing cost functions in a ratio different than 1.

In this thesis, our goal is to evaluate the empirical performance of the dual-balancing policies and several other variants of cost-balancing policies discribed above. It has been proven that the dual-balancing policies are computationally efficient and can be applied in an on-line manner. In [2], it was shown that they have a worst-case guarantee of 2. However, this does not necessarily indicate what is the typical emperical performance. The goal of this thesis is to explore the efficiency and empirical performnace of the dual-balancing policy via computational experiments. Hence, we focus on an inventory model in which the demands in different periods are independent identically distributed following a Poisson distribution. For these models, Shang and Song has developed an approximately optimal policy so one can actually test how far from the optimal the cost balancing policies perform [6]. Moreover, it has been shown that the approximately optimal policy produces the optimal solution in several experiments in [6] and hence it is used as a benchmark of the performance of the cost balancing policies in our experiments.

The performance of the dual-balancing policy is evaluated based on two aspects: accuracy, in terms of long-run expected average cost, and efficiency, in terms of the expected computational time taken for each decision made. First, we examine the error of the dual-balancing policy compared to the approximately optimal policy (based on [6]). We did not observe any case when the dual-balancing policy showed an error larger than 20%. The empirical results showed an average error of 7.74% with the largest error being 19.73%, which is significantly better than the current known worst-case factor of 2. It follows that the worst-case guarantee of 2 does not reflect the empirical performance of the dual-balancing policies in many scenarios. The sensitivity analysis on the system parameters also indicates that the error is monotonically increasing with the per-unit backorder cost and decreasing with the demand rate.

In addition to the dual-balancing policy, we test three variants of cost balancing policies, which are the dual-balancing policy with bounds, the parameterized balancing policy and the parameterized balancing policy with bounds. They all use the same bounds on the optimal echelon base-stock level, obtained from [6]. The details of computation of the balancing ratio of the parameterized balancing policies are given in Chapter 4. The effect of each modification is evaluated based on the difference in error compared to the approximately optimal policy. They all show great improvement over the original dual-balancing policy. Using bounds reduces the error of the dual-balancing policy to an average of 1.15% while the parameterized balancing policy with bounds results in an average error of 1.62%. Our suggestion is hence to take these two modifications to the dual-balancing policy to further improve its performance. Similarly to the dual-balancing policy, the parameterized balancing policy also has an error which increases with the per-unit backorder cost and decreases with the demand rate. However, the modification of system parameters has little impact on the performance of the dual-balancing policy with bounds and the parameterized balancing policy with bounds. In particular, the fluctuations of the errors of these two policies are within 3% when the system parameters change.

Throughout the 52 experiments we have taken, there are several cases when the error of some cost balancing policy becomes negative, that is, the cost balancing policy out performs the approximately optimal policy. Since the magnitude of the error is too small, we further modified the system parameters and found several cases when the cost balancing policy dominates the approximately optimal policy. For example, for a 5-stage system with demand rate $\lambda = 32$, per-unit backorder cost $\pi = 5$, per-unit echelon holding cost of [0.25 0.25 0.25 0.25 2.5] and lead time being [1 1 50 1 1], all the cost balancing policies out perform the approximately optimal policy by at least 2% and especially the dual-balancing policy with bounds results in an improvement of 3.05%.

In terms of efficiency, the approximately optimal policy does not require any computational time at each decision making while the dual-balancing policy takes an average of 0.29 seconds to compute the balancing orders. The computational time increases with the demand rate and the per-unit backorder cost. Since the accuracy of the cost balancing policies is quite high, the computational time for each decision making is still acceptable.

The rest of the thesis is organized as follows. In Chapter 2, we describe the mathematical formulation of the inventory model and specify the sequence of events in each time period. In Sections 3.1 and 3.2, we present the cost decomposition scheme and prove its equivalent with the traditional cost accounting scheme. Then we assign to each desicion the costs that are made inevitable after the decision in Section 3.3. The ordering rules of the dual-balancing policy and the other variant of cost balancing policies are presented in Chapter 4. In Chapter 5, we present the implementation of the numerical simulation. Lastly, we present the details of the experiment design and the numerical results.

Chapter 2

The Serial Inventory Model

In this chapter, we introduce the mathematical formulation of the serial network that will be used throughout the thesis. Consider a periodic-reviewed system with n stages, numbered $1, 2, ..., n$. Stage n orders from an external supplier with infinite capacity and units are shipped downstream through all the stages to stage 1 to meet the external demand. Each stage $k = 1, ..., n-1$ can order only from the on-hand inventory at the preceding stage $k+1$. It takes l_k periods to ship units from stage $k+1$ to stage k. We assume that the lead times l_k are strictly positive; otherwise we can merge two stages without loss of generality. Hence, the cumulative time for transportation from stage $k+1$ to stage 1 is $L_k = \sum_{j=1}^{k} l_j$ periods.

We consider a model with discrete time and infinite horizon. Let D_t denote the demand in period t. The random variables $\{D_t\}_{t=1}^{\infty}$ are i.i.d., following the Poisson distribution with parameter λ. Denote $D_{[1,t]}$ as the cumulative demand arrivals from period 1 (the starting time) to period t. Therefore, $D_{[1,t]} = \sum_{j=1}^{t} D_j$ and hence it is a Poisson random variable with parameter λt. Demands can be observed at all stages but is only satisfied from the on-hand inventory at stage 1. Unsatisfied demands are fully backlogged and accumulate over time until they are satisfied. That is, they wait in the system until they are satisfied.

Two types of costs are incurred at the end of each period. A conventional per-unit holding cost h'_k is incurred for each unit in the on-hand inventory at stage k or in transit to stage k from stage $k+1$, and a per-unit backordering penalty cost π

is incurred for each unit of demand at stage 1 that is not yet satisfied. It is more convenient to use an echelon holding cost accounting scheme, described below.

Denote the number of units on-hand at stage j or in transit from stage $j+1$ to stage j at the end of period τ by $v_{j\tau}$. The *echelon inventory* at stage k at the end of period τ is denoted by $Y_k(\tau)$ and defined as the sum of all $v_{j\tau}$, for $1 \leq j \leq k$ (i.e. $Y_k(\tau) = \sum_{j=1}^{k} v_{j\tau}$). For each unit of the echelon inventory, we charge a per-unit echelon holding cost of h_k where $h_k = h'_k - h'_{k+1}$ (assuming $h'_{n+1} = 0$ and $h'_k \geq h'_{k+1}$, we have $h_k \geq 0$). This implies that the conventional per-unit holding cost h'_k can be expresssed as $\sum_{j=k}^{n} h_j$. Hence, the total holding cost at the end of period τ at all stages is

$$\sum_{k=1}^{n} v_{k\tau} h'_k = \sum_{k=1}^{n} v_{k\tau} \sum_{j=k}^{n} h_j = \sum_{k=1}^{n} h_k \sum_{j=1}^{k} v_{j\tau} = \sum_{k=1}^{n} h_k Y_k(\tau).$$

This shows that the echelon holding cost accounting scheme is equivalent to the conventional holding cost accounting scheme. In addition, denote $B(\tau)$ by the number of backlogged demand unit at the end of period τ.

Denote $\mathcal{P}$ as the set of all feasible policies. For each feasible policy $P \in \mathcal{P}$, the total cost incurred over the first t periods is

$$R_t^P = \sum_{\tau=1}^{t} \pi B^P(\tau) + \sum_{k=1}^{n} h_k Y_k^P(\tau).$$

Define the long-run expected average cost per period as

$$R^P = \lim_{t \to \infty} E(\frac{R_t^P}{t}).$$

The objective of the inventory problem becomes

$$min_{P \in \mathcal{P}} R^P$$

Denote the order placed at stage k at period τ by $Q_k(\tau)$ according to some feasible policy P. Define the *echelon inventory position* at stage k as the echelon inventory of stage k minus current backorders at stage 1. It can be seen that the echelon inventory position goes up when orders are placed and goes down when demand arrives. Let $X_k(\tau)$ denote the echelon inventory position at stage k at the beginning of period k, before the order $Q_k(\tau)$ is placed. The *echelon net inventory* at stage k includes all units that are in the echelon inventory position but not in transit to stage k. We can also observe that the echelon net inventory goes up when past orders arrive and goes down when demand occurs. Let $NI_k(\tau)$ denote the echelon net inventory at stage k after past orders arrive at period τ but before demand arrives. Let $I_k(\tau)$ denote the on-hand inventory at stage k after past orders have arrived at period τ. We can conclude that for any stage k and period τ, $Q_k(\tau) \leq I_{k+1}(\tau)$ and $I_{k+1}(\tau) = NI_{k+1}(\tau) - X_k(\tau)$. Next we specify the sequence of events in each period τ:

1. Period τ begins with echelon inventory position $X_k(\tau)$ and echelon net inventory $NI_k(\tau-1) - D_{\tau-1}$ at each stage k;

2. At each stage k, the orders placed at period $\tau - l_k$ of $Q_k(\tau - l_k)$ units will arrive and the echelon net inventory increases to $NI_k(\tau-1) - D_{\tau-1} + Q_k(\tau - l_k) = NI_k(\tau)$. Consequently, the on-hand inventory at stage k also increase by $Q_k(\tau - l_k)$ units;

3. At each stage k, an order of $Q_k(\tau)$ units is determined by the specific policy. The order size is constrained by the on-hand inventory at the preceding stage $k+1$, i.e. $0 \leq Q_k(\tau) \leq I_{k+1}(\tau)$. Consequently, the echelon inventory position at each stage k increases to $X_k(\tau) + Q_k(\tau)$;

4. At the end of period τ, we observe the demand D_τ and satisfy them from the on-hand inventory at stage 1. The echelon inventory position and echelon net inventory at each stage k will decrease to $X_k(\tau) + Q_k(\tau) - D_\tau$ and $NI_k(\tau) - D_\tau$ respectively;

5. For each demand unit that is not satisfied at stage 1, we charge a backorder cost of π. The total number of backorders is $B(\tau) = (-NI_1(\tau))^+$. For each unit that has been ordered by stage k and is still in the system, we charge a holding cost of h_k. The total number of such units is $Y_k(\tau) = X_k(\tau) + Q_k(\tau) - D_\tau + B(\tau)$.

Hence, the total costs charged in this period is

$$\begin{aligned} & \pi(-NI_1(\tau))^+ + \sum_{k=1}^{n} h_k[X_k(\tau) + Q_k(\tau) - D_\tau + (-NI_1(\tau))^+] \\ = & \quad (h_1' + \pi)(-NI_1(\tau))^+ + \sum_{k=1}^{n} h_k[X_k(\tau) + Q_k(\tau) - D_\tau] \end{aligned}$$

Chapter 3

Marginal Cost Accounting Scheme

In this chapter, we describe the cost accounting scheme introduced by Levi et al. [1] for serial systems. The general approach is to take costs incurred, classify them into categories, and use the categories to assign each unit of cost to the specific decision which caused that cost to be incurred. More rigorously, we assign to a given decision all costs that become inevitable after this specific decision and are not affected by any future decisions. This marginal cost accounting scheme significantly differs from most of the literature on stochastic inventory control problems. In the conventional cost accounting scheme, the costs are accounted in the periods when they are incurred. In the marginal cost accounting scheme, each decision is assigned by all the costs that were made inevitable by that decision, whether they are incurred in the current period or in the future.

To make the discussion more rigorous, we adopt a *distance-numbering scheme* for the units of demand and supply respectively. Let T_D be a half-infinite time line segement $[0, \infty]$ that represents the units of demand realized over the time horizon and T_S be another half-infinite time line segement $[0, \infty]$ that represents the units of supply that can be ordered. Define demand unit i as the unit that is located at a distance of i from the origin on the time line T_D and supply unit i in the same manner respectively on T_S. Without loss of generality, assume that the supply units are consumed by the demand on a first-ordered-first-consumed manner. As a result, we can match each supply unit with the demand unit that has the same distance.

In other words, each unit of demand is satisfied by the unit of supply with the same distance.

The exposition of the accounting scheme is unit-based. That is, we track the sample path of unit i in the system and categorize the costs incurred. Then in Section 3.2, we show that this cost accounting scheme is equivalent to the conventional cost accounting scheme, i.e. it includes all the costs incurred for each unit i. In the last section, we present the assignment of costs to the specific ordering decisions that made the costs inevitable.

3.1 Cost Decomposition

Using the concept of demand-supply unit matching, denote the period when demand unit i arrives by T_i and the period when stage k orders supply unit i by u_{ik}. Suppose we know the value of T_i before hand. It is obvious that to minimize the total cost incurred, we should order the corresponding supply unit in just-in-time manner, defering orders to avoid unnecessary holding cost and ensuring it arrives at stage 1 exactly at period T_i to avoid backorder cost. More precisely, each stage k should order supply unit i at period $T_i - L_k$, which is defined as the *critical* period. However, obviously in general we do not know the value of T_i in advance, making it difficult to anticipate when is critical for ordering. Therefore, we classify the costs incurred by supply unit i into four categories and assign each cost to the ordering decisions that caused that cost to be incurred. That is, after the ordering decision, that unit of cost will incur regardless of any future decision. Let y^+ denote the positive part of y (i.e., $max\{y, 0\}$). For a given sample path, we consider the following costs related to demand-supply unit i:

- *Pipeline holding cost* is incurred when unit i is in transit from one stage to another. It takes L_k periods in transit from the time when it is ordered at some stage k to the time when it reaches stage 1. Hence, for each unit i, we incur a pipeline holding cost of $h_k L_k$ at each stage k. In our infinite model, the pipeline holding costs are inevitable even if one orders the unit according to the critical periods. Thus for any feasible policy, there is a pipeline cost of $\sum_{k=1}^{n} h_k L_k$ incurred for each unit i.

- *Early holding cost* is the holding cost incurred when the supply unit is on-hand inventory at some stage of the system and it is still possible to ship the unit on time to meet the demand unit. Assume that at some stage k, unit i is ordered prior to the respective critical period, i.e. $u_{ik} < T_i - L_k$. In this case, there are $T_i - L_k - u_{ik}$ periods when unit i is in the echelon inventory of stage k and at the same time, it is still possible to ship the supply unit on time to satisfy the demand unit. Hence, a holding cost of $h_k(T_i - L_k - u_{ik})^+$ is incurred in addition to the pipeline costs. We call this *early holding cost* and assign it to the decision of ordering unit i at stage k at period u_{ik}.

- *Late holding cost* is the holding cost incurred when the supply unit is on-hand inventory at some stage of the system but it is no longer possible to ship the unit on time to meet the demand unit. Assume that unit i is on-hand at stage $k+1$ from period τ to period $\tau+1$. However, it is no longer possible to ship the supply unit on time to stage 1, i.e. $\tau > T_i - L_k$. A conventional holding cost of h'_{k+1} is incurred. We call it *late holding cost* and consider it as a result of not ordering that unit at stage k at period τ. Therefore, whenever unit i is stored at stage $k+1$ from period τ to $\tau+1$ with $\tau > T_i - L_k$, then a late holding cost of h'_{k+1} is assigned to the ordering decision at stage k at period τ.

- *Backorder cost.* Consider a period τ such that $\tau > T_i$. If unit i is not delivered to stage 1 at period τ, a backorder cost of π will be accounted. This cost can be avoided if every stage k has ordered the supply unit at the critical period $\tau - L_k$. Let j be the largest stage index at which unit i is not ordered by period $\tau - L_j$. The backorder cost π is considered as a result of the decision not to order that unit at stage j at period $\tau - L_j$. Therefore, for any period $\tau > T_i$ such that unit i has not reached stage 1 by the end of period τ, we assign a backorder cost of π to the ordering decision made at period $\tau - L_j$ at stage j

where j is defined in the manner just described.

Consider the total cost incurred for the i-th demand-supply unit for a given policy P. The period when stage k orders the supply unit is denoted by u_{ik} and the arrival time of the demand unit is denoted by T_i. First, the total pipeline holding cost and early holding cost incurred are $\sum_{k=1}^{n} h_k L_k$ and $\sum_{k=1}^{n} h_k (T_i - L_k - u_{ik})^+$, respectively. A late holding cost of h'_{k+1} is incurred when the supply unit is on-hand at stage $k+1$ from period τ to period $\tau + 1$ with $\tau > T_i - L_k$, for $k = 1, ..., n-1$. According to the ordering policy, the supply unit stays at stage $k+1$ from the time when it arrives at stage $k+1$, i.e. period $u_{i(k+1)} + l_{k+1}$ to the time when it is ordered by stage k, i.e. period u_{ik}. Hence, the late holding cost incurred when the supply unit is on-hand at stage $k+1$ is $h'_{k+1}\left[u_{ik} - max(T_i - L_k, u_{i(k+1)} + l_{k+1})\right]^+$ and the total late holding cost incurred for unit i is

$$\sum_{k=1}^{n-1} h'_{k+1}\left[u_{ik} - max(T_i - L_k, u_{i(k+1)} + l_{k+1})\right]^+$$

The supply unit arrives at stage 1 at period $u_{i1} + l_1$ and thus backorder cost is only incurred if $u_{i1} + l_1 > T_i$. A backorder cost of π is incurred for each period τ such that the supply unit is not delivered to stage 1 at period τ and $\tau > T_i$. Hence, the total backorder cost incurred is $\pi(u_{i1} + l_1 > T_i)^+$. To summarize, the total cost incurred for the demand and supply unit i for policy P is

$$\pi(u_{i1} + l_1 - T_i)^+ + \sum_{k=1}^{n} h_k \left[L_k + (T_i - L_k - u_{ik})^+\right]$$

$$+ \sum_{k=1}^{n-1} h'_{k+1}\left[u_{ik} - max(T_i - L_k, u_{i(k+1)} + l_{k+1})\right]^+$$

3.2 Validity of Marginal Cost Accounting Scheme

In the conventional cost accounting scheme, at the end of period τ, an echelon holding cost h_k is incurred for each stage k such that the supply unit i is within the echelon inventory of stage k and a backorder cost π is incurred if supply unit i has not reached stage 1 but the demand unit has already occurred. To see the validity of the new cost accounting scheme, we need to show that the total cost incurred according to the marginal cost accounting scheme equals the total cost incurred in the conventional cost accounting scheme. Recall that u_{ik} denote the period when the supply unit i is ordered by stage k and T_i denote the period when the demand unit i occurs. The supply unit is included in the echelon inventory of stage k from the time when it is ordered, i.e. u_{ik} to the time it is consumed by the demand unit, i.e. $max(T_i, u_{i1}+L_1)$. Hence, the total echelon holding cost incurred by unit i is

$$\sum_{k=1}^{n} h_k[max(T_i, u_{i1} + L_1) - u_{ik}]$$

The backorder cost is incurred from the time when the demand unit arrives, i.e. T_i to the time when the supply unit actually arrives at stage 1, i.e. $u_{i1} + l_1$. Hence, the total backorder cost is $\pi(u_{i1} + l_1 - T_i)^+$, which is the same as the backorder cost incurred according to the marginal cost accounting scheme. Hence, we only need to prove the equality of the total echelon holding cost incurred according to the two cost accounting schemes, that is, to prove that

$$\sum_{k=1}^{n} h_k \left[L_k + (T_i - L_k - u_{ik})^+\right] + \sum_{k=1}^{n-1} h'_{k+1} \left(u_{ik} - max(T_i - L_k, u_{i(k+1)} + l_{k+1})\right)^+$$

and

$$\sum_{k=1}^{n} h_k[max(T_i, u_{i1} + L_1) - u_{ik}]$$

are equal for any policy P. Next we consider three possible cases of unit i to show that the marginal cost accounting scheme includes all the echelon holding costs.

Case 1 : The supply unit i arrives at stage 1 before the demand unit occurs.

In this case, each stage orders the supply unit before the corresponding critical time, i.e. $T_i - L_k \geq u_{ik}$ for each k. Hence, the total echelon holding cost is

$$\sum_{k=1}^{n} h_k[max(T_i, u_{i1} + L_1) - u_{ik}]^+ = \sum_{k=1}^{n} h_k(T_i - u_{ik})$$

According to the marginal cost accounting scheme, only pipeline holding cost and early holding cost are incurred, which are

$$\sum_{k=1}^{n} h_k L_k \quad \text{and} \quad \sum_{k=1}^{n} h_k (T_i - L_k - u_{ik})^+ = \sum_{k=1}^{n} h_k (T_i - L_k - u_{ik})$$

respectively. Then the sum of holding costs according to the marginal cost accounting scheme is $\sum_{k=1}^{n} h_k(T_i - u_{ik})$, which is the same as the total echelon holding cost according to the conventional cost accounting scheme.

Case 2 : The supply unit i does not reach stage 1 on time to meet the demand. But when stage n orders the supply unit, it is still possible to ship the unit on time. In this case, we have $T_i \leq u_{i1} + L_1$ and $T_i - L_n \geq u_{in}$. Hence, the total echelon holding cost is

$$\sum_{k=1}^{n} h_k[max(T_i, u_{i1} + L_1) - u_{ik}] = \sum_{k=1}^{n} h_k(u_{i1} + L_1 - u_{ik})$$

Denote j as the largest index such that stage j orders the supply unit after the respective critical time, i.e. $T_i - L_j \leq u_{ij}$. Then we have $1 \leq j \leq n-1$. Let's consider the three types of costs in the marginal cost accounting scheme:

- Pipeline holding cost is inevitable, which is $\sum_{k=1}^{n} h_k L_k$;
- Early holding cost is $\sum_{k=1}^{n} h_k (T_i - L_k - u_{ik})^+$. In this case, $T_i - L_k - u_{ik}$ is positive only for stage k such that $j+1 \leq k \leq n$. Hence, the total early holding cost is $\sum_{k=j+1}^{n} h_k (T_i - L_k - u_{ik})$;
- It becomes impossible to ship the unit on time only after stage $j+1$ makes the order. Hence, late holding cost is only incurred when the supply unit i stays at stage 2 up to stage $j+1$. The time period when the supply unit is on-hand inventory at some stage k is $[u_{ik} + l_k, u_{i(k-1)}]$ and it becomes impossible to ship the unit on time after period $T_i - L_{k-1}$. Hence, a late holding cost of $h'_k[u_{i(k-1)} - max(u_{ik} + l_k, T_i - L_{k-1})]^+$ is incurred when the supply unit is on-hand at stage k, for all $2 \leq k \leq j+1$. From the definition

of j, we have $T_i \geq u_{i(j+1)} + L_{j+1}$ and $T_i \leq u_{ik} + L_k$ for $1 \leq k \leq j$. Hence, the total late holding cost becomes

$$\begin{aligned}
& h'_{j+1}(u_{ij} - T_i + L_j) + \sum_{k=2}^{j} h'_k[u_{i(k-1)} - u_{ik} - l_k] \\
= & \sum_{m=j+1}^{n} h_m(u_{ij} - T_i + L_j) + \sum_{k=2}^{j}\sum_{m=k}^{n} h_m[u_{i(k-1)} - u_{ik} - l_k] \\
= & \sum_{m=j+1}^{n} h_m(u_{ij} - T_i + L_j) + \sum_{m=2}^{j-1} h_m \sum_{k=2}^{m}[u_{i(k-1)} - u_{ik} - l_k] \\
& \quad + \sum_{m=j}^{n} h_m \sum_{k=2}^{j}[u_{i(k-1)} - u_{ik} - l_k] \\
= & \sum_{m=j+1}^{n} h_m(u_{ij} - T_i + L_j) + \sum_{m=2}^{j-1} h_m(u_{i1} + L_1 - u_{im} - L_m) \\
& \quad + \sum_{m=j}^{n} h_m(u_{i1} + L_1 - u_{ij} - L_j) \\
= & \sum_{m=j+1}^{n} h_m(u_{ij} - T_i + L_j) + \sum_{m=2}^{j-1} h_m(u_{i1} + L_1 - u_{im} - L_m) \\
& \quad + \sum_{m=j+1}^{n} h_m(u_{i1} + L_1 - u_{ij} - L_j) + h_j(u_{i1} + L_1 - u_{ij} - L_j) \\
= & \sum_{m=j+1}^{n} h_m(u_{i1} + L_1 - T_i) + \sum_{m=2}^{j} h_m(u_{i1} + L_1 - u_{im} - L_m)
\end{aligned}$$

Therefore, the sum of three costs becomes

$$\begin{aligned}
& \sum_{k=1}^{n} h_k L_k + \sum_{k=j+1}^{n} h_k(T_i - L_k - u_{ik}) \\
& \quad + \sum_{m=j+1}^{n} h_m(u_{i1} + L_1 - T_i) + \sum_{m=2}^{j} h_m(u_{i1} + L_1 - u_{im} - L_m) \\
= & \sum_{k=1}^{n} h_k L_k + \sum_{k=j+1}^{n} h_k(u_{i1} + L_1 - L_k - u_{ik}) + \sum_{m=2}^{j} h_m(u_{i1} + L_1 - u_{im} - L_m) \\
= & \sum_{k=1}^{n} h_k L_k + \sum_{k=2}^{n} h_k(u_{i1} + L_1 - u_{ik} - L_k) \\
= & h_1 L_1 + \sum_{k=2}^{n} h_k(u_{i1} + L_1 - u_{ik}) \\
= & \sum_{k=1}^{n} h_k(u_{i1} + L_1 - u_{ik})
\end{aligned}$$

which is equivalent to the total echelon holding cost according to the conventional cost accounting scheme.

Case 3 : The supply unit i does not reach stage 1 on time to meet the demand. Moreover, when stage n orders the supply unit, it is already impossible to ship the unit on time.

In this case, we have $T_i - L_k \leq u_{ik}$ for any stage k. The total echelon holding cost incurred is hence

$$\sum_{k=1}^{n} h_k(u_{i1} + L_1 - u_{ik})$$

In the marginal cost accounting scheme, there is no early holding cost since every stage orders after the corresponding critical time. The pipeline holding cost is the same as in Case 2, i.e. $\sum_{k=1}^{n} h_k L_k$. Late holding cost is incurred when the supply unit is on-hand at any stage except stage 1. Hence, the total late holding cost incurred is

$$\begin{aligned}
& \sum_{k=2}^{n} h'_k[u_{i(k-1)} - u_{ik} - l_k] \\
= \quad & \sum_{k=2}^{n} \sum_{m=k}^{n} h_m[u_{i(k-1)} - u_{ik} - l_k] \\
= \quad & \sum_{m=2}^{n} h_m \sum_{k=2}^{m}[u_{i(k-1)} - u_{ik} - l_k] \\
= \quad & \sum_{m=2}^{n} h_m(u_{i1} + L_1 - u_{im} - L_m)
\end{aligned}$$

The sum of pipeline cost and late holding cost for unit i is

$$\begin{aligned}
& \sum_{k=1}^{n} h_k L_k + \sum_{m=2}^{n} h_m(u_{i1} + L_1 - u_{im} - L_m) \\
& = h_1 L_1 + \sum_{m=2}^{n} h_m(u_{i1} + L_1 - u_{im}) \\
& = \sum_{m=1}^{n} h_m(u_{i1} + L_1 - u_{im})
\end{aligned}$$

Again, it is the same as the total echelon holding cost according to the conventional cost accounting scheme and we proved the equivalence of the two accounting schemes.

3.3 Decision-Based Costs Allocation

Among the four types of costs, the pipeline cost is inevitable and hence independent of the policy. The early holding cost is incurred when we order the units before the critical period. On the other hand, the late holding and backorder costs are incurred when we order the units after the critical period. Therefore, the costs assigned to the ordering decision of size $Q_k(\tau)$ at stage k at period τ include two types of costs, which will be discussed in the following two subsections:

- The early holding costs incurred by the units that have been ordered, defined as *marginal early holding cost*;
- The late holding and backorder costs incurred by the units that have not yet been ordered, defined as *marginal late holding and backorder cost*.

3.3.1 Marginal Early Holding Costs

For any supply unit i that is first ordered by stage k at period τ, if the corresponding demand unit arrives after period $L_k+\tau$, then an early holding cost of $h_k(T_i-L_k-\tau)^+$ is assigned to the ordering decision. In other words, at any period $t \geq L_k+\tau$, if the demand unit has not arrived and the supply unit is within the ordering decision of stage k at period τ, then an early holding cost of h_k is assigned to the ordering decision. To be specific, for each $t > L_k+\tau$, supply unit i contributes an early holding cost of h_k to the ordering decision if and only if it satisfies the following conditions:

1. The demand unit i arrives after period t;

2. The supply unit i has been ordered at the end of period τ at stage k;

3. The supply unit i has not been ordered at the beginning of period τ at stage k.

From the definition of echelon inventory, we can see the number of units that meet condition 2 is $X_k(\tau) + Q_k(\tau)$ and the number of units that fail condition 3 is $X_k(\tau)$. Note that any unit that fails condition 3 will necessarily satisfy condition 2. The number of units that meet both condition 1 and 2 is $(X_k(\tau) + Q_k(\tau) - D_{[\tau,t]})^+$ and the number of units that meet condition 1 but fail condition 3 is $(X_k(\tau) - D_{[\tau,t]})^+$. The total number of units that meet all three conditions is obtained by subtracting the above two,

$$(X_k(\tau) + Q_k(\tau) - D_{[\tau,t]})^+ - (X_k(\tau) - D_{[\tau,t]})^+ = (Q_k(\tau) - (D_{[\tau,t]} - X_k(\tau))^+)^+.$$

Therefore, the expected marginal early holding cost for the ordering decision is $h_k HC_k(\tau)$ where

$$HC_k(\tau) = \sum_{t=L_k+\tau}^{\infty} E[(Q_k(\tau) - (D_{[\tau,t]} - X_k(\tau))^+)^+]$$

3.3.2 Marginal Late Holding and Backorder Costs

Marginal late holding and backorder costs are possible future costs incurred by each unit i that is not ordered at period τ. If the demand unit i arrives after period $L_k+\tau$, at period τ it is still prior to the critical period and hence no late holding or backorder costs caused by the decision. Otherwise if $T_i \leq L_k + \tau$, the decision of not ordering unit i will keep the supply unit at stage $k+1$ from period τ to period $\tau+1$ and the supply unit cannot meet the demand on time. Hence, a late holding cost of h'_{k+1} is assigned to the ordering decision. In addition, the supply unit has not arrived at stage 1 by period $L_k + \tau$ and a backorder cost of π is considered as a result of the ordering decision. Therefore, a total cost of $h'_{k+1} + \pi$ is assigned to the decision if and only if unit i satisfies the following conditions:

4. The supply unit i is on-hand at stage $k+1$ at the beginning of period τ;

5. Stage k is not ordering unit i at period τ;

6. The arrival period of demand unit i satisfies $T_i \leq L_k + \tau$.

Observe that any unit that fails condition 4 will automatically satisfy condition 5. The total number of units satisfying condition 5 and 6 is $(D_{[\tau, L_k+\tau]} - (X_k(\tau) + Q_k(\tau)))^+$. Similarly, the total number of units that meet condition 6 but fail condition 4 is $(D_{[\tau, L_k+\tau]} - NI_{k+1}(\tau))^+$. The total number of units that meet all three conditions is

$$(D_{[\tau, L_k+\tau]} - (X_k(\tau) + Q_k(\tau)))^+ - (D_{[\tau, L_k+\tau]} - NI_{k+1}(\tau))^+$$

Therefore, the expected marginal late holding and backorder cost for the ordering decision is $(h'_{k+1} + \pi)BC_k(\tau)$ where

$$BC_k(\tau) = E[(D_{[\tau, L_k+\tau]} - (X_k(\tau) + Q_k(\tau)))^+ - (D_{[\tau, L_k+\tau]} - NI_{k+1}(\tau))^+]$$

Chapter 4

Cost Balancing Policies

4.1 Dual-Balancing Policies

The dual-balancing policy is based on the marginal cost accounting scheme and aims to balance the effects of ordering too early versus ordering too late. To be specific, the ordering size is determined so that the conditional expected marginal early holding cost caused by the units that have been ordered, and the conditional expected marginal late holding and backorder cost caused by the units that have not yet been ordered can be balanced. In our model, the demands are i.i.d. and the horizon is infinite. Hence, we may assume without loss of generality that current ordering period is 1. For ease of exposition, we omit τ in all denotations. Also, we change the denotation of marginal costs as functions of the order size, the echelon inventory position and the echelon net inventory as follows:

$$HC_k(Q_k, X_k) = \sum_{t=L_k+1}^{\infty} E[(Q_k - (D_{[1,t+1]} - X_k)^+)^+]$$

and

$$BC_k(Q_k, X_k, NI_{k+1}) = E[(D_{[1,L_k+1]} - (X_k + Q_k))^+ - (D_{[1,L_k+1]} - NI_{k+1})^+].$$

In the marginal cost accounting scheme, we assign to each ordering decision all the costs that are made inevitable after the decision is made. These costs may be incurred in current period or in the future, but independent of past orders or the stage at which decision is made. Hence, the ordering decision made at each time period is made separately at each stage. In general, we will discuss the ordering rules at period 1 at stage k for $k = 1, ..., n$.

Suppose the demand of some unit i has arrived but the corresponding supply unit is still on-hand at some stage $k \neq 1$ or in transit. It is clear that at this point the corresponding supply unit should be ordered immediately by each stage k to avoid extra costs. We call this an *immediate order*, and denote the total quantity of immediate orders in stage k by $\bar{Q}_k$, for $k = 1, ..., n$. When the demand unit has not yet arrived by current period, we will then decide whether to order the supply unit according to the specific policy. The order made in stage k after the immediate orders are decided is called a *regular order*, denoted by $\tilde{Q}_k = Q_k - \bar{Q}_k$. Let $\tilde{X}_k = X_k + \bar{Q}_k$ be the echelon inventory position after the immediate orders are placed, but before the regular orders are placed. Then the regular order size will be determined based on the current echelon inventory position $\tilde{X}_k$ and the on-hand inventory at stage $k + 1$ of $NI_{k+1} - \tilde{X}_k$ units.

For the supply units in the immediate orders, the corresponding demand units have already arrived and hence no early holding cost is caused by the ordering decision. Since we order without any delay, there is no late holding cost associated with the ordering decision. The incurred pipeline costs are inevitable and hence no costs assigned to the decision for each unit of immediate orders. The formula for the expected marginal early holding costs still holds for the regular orders. That is, the expected marginal early holding cost assigned to the regular order is $h_k HC_k(\tilde{Q}_k, \tilde{X}_k)$ where

$$HC_k(\tilde{Q}_k, \tilde{X}_k) = \sum_{t=L_k+1}^{\infty} E[(\tilde{Q}_k - (D_{[1,t]} - \tilde{X}_k)^+)^+$$

Also, since $\tilde{Q}_k + \tilde{X}_k = Q_k + X_k$, the marginal late holding and backorder cost assigned to the regular order is $(h'_{k+1}+\pi)BC_k$ where

$$BC_k(\tilde{Q}_k,\tilde{X}_k,NI_{k+1}) = E[(D_{[1,L_k+1]}-(\tilde{X}_k+\tilde{Q}_k))^+ - (D_{[1,L_k+1]}-NI_{k+1})^+]$$

Denote $f(\beta,y), F(\beta,y)$ by the pdf and cdf of Poisson distribution with parameter β evaluated at y, respectively. Hence, the cost functions involve the random variable $D_{[0,t]}$ following a Poisson distribution with parameter λt and the derivations are presented as follows.

$$\begin{aligned}
HC_k(\tilde{Q}_k,\tilde{X}_k) &= \sum_{t=L_k+1}^{\infty} E[(\tilde{Q}_k-(D_{[1,t]}-\tilde{X}_k)^+)^+] \\
&= \sum_{t=L_k+1}^{\infty}\sum_{y=0}^{\infty}(\tilde{Q}_k-(y-\tilde{X}_k)^+)^+ f(\lambda t,y) \\
&= \sum_{t=L_k+1}^{\infty}[\sum_{y=0}^{\tilde{X}_k}\tilde{Q}_k f(\lambda t,y) + \sum_{y=\tilde{X}_k+1}^{\tilde{X}_k+\tilde{Q}_k}(\tilde{Q}_k+\tilde{X}_k-y)f(\lambda t,y)]
\end{aligned}$$

From the definition of Poisson distribution, we have

$$F(\lambda t,y) = \sum_{j=0}^{y} f(\lambda t,j)$$

and

$$yf(\lambda t,y) = ye^{-\lambda t}\frac{(\lambda t)^y}{y!} = \lambda t e^{-\lambda t}\frac{(\lambda t)^{y-1}}{(y-1)!} = \lambda t f(\lambda t,y-1).$$

The expected marginal early holding cost becomes

$$\begin{aligned}
& \sum_{t=L_k+1}^{\infty}[\tilde{Q}_k F(\lambda t,\tilde{X}_k) + (\tilde{Q}_k+\tilde{X}_k)\sum_{y=\tilde{X}_k+1}^{\tilde{X}_k+\tilde{Q}_k} f(\lambda t,y) - \sum_{y=\tilde{X}_k+1}^{\tilde{X}_k+\tilde{Q}_k}\lambda t f(\lambda t,y-1)] \\
= & \sum_{t=L_k+1}^{\infty}[(\tilde{X}_k+\tilde{Q}_k)F(\lambda t,\tilde{X}_k+\tilde{Q}_k) - \tilde{X}_k F(\lambda t,\tilde{X}_k)] \\
- & \sum_{t=L_k+1}^{\infty}\lambda t[F(\lambda t,\tilde{X}_k+\tilde{Q}_k-1) - F(\lambda t,\tilde{X}_k-1)]
\end{aligned}$$

Similarly, the expected marginal late holding and backorder cost is

$$
\begin{aligned}
& BC_k(\tilde{Q}_k, \tilde{X}_k, NI_{k+1}) \\
= \quad & E[(D_{[1,L_k+1]} - (\tilde{X}_k + \tilde{Q}_k))^+ - (D_{[1,L_k+1]} - NI_{k+1})^+] \\
= \quad & \sum_{y=0}^{\infty}[(y - (\tilde{X}_k + \tilde{Q}_k))^+ - (y - NI_{k+1})^+] f(\lambda(L_k+1), y) \\
= \quad & \sum_{y=\tilde{X}_k+\tilde{Q}_k+1}^{NI_{k+1}} (y - \tilde{X}_k - \tilde{Q}_k) f(\lambda(L_k+1), y) \\
& \quad + \sum_{y=NI_{k+1}+1}^{\infty} (NI_{k+1} - \tilde{X}_k - \tilde{Q}_k) f(\lambda(1+L_k), y) \\
= \quad & \sum_{y=\tilde{X}_k+\tilde{Q}_k+1}^{NI_{k+1}} \lambda L_k f(\lambda(L_k+1), y-1) - (\tilde{X}_k + \tilde{Q}_k) \sum_{y=\tilde{X}_k+\tilde{Q}_k+1}^{NI_{k+1}} f(\lambda(L_k+1), y) \\
+ \quad & (NI_{k+1} - \tilde{X}_k - \tilde{Q}_k)[1 - F(\lambda(L_k+1), NI_{k+1})] \\
= \quad & \lambda L_k [F(\lambda(L_k+1), NI_{k+1} - 1) - F(\lambda(L_k+1), \tilde{X}_k + \tilde{Q}_k - 1)] \\
- \quad & (\tilde{X}_k + \tilde{Q}_k)[1 - F(\lambda(L_k+1), \tilde{X}_k + \tilde{Q}_k)] + NI_{k+1}[1 - F(\lambda(L_k+1), NI_{k+1})]
\end{aligned}
$$

Note that the marginal early holding cost is 0 for order size $Q_k = 0$ and increases with the size of order. The marginal late holding and backorder cost is 0 for order size $Q_k = NI_{k+1} - X_k$ and decreases with the size of order. There must exist an order size Q_k^* such that the two opposing functions can be balanced and hence the balancing order size is well-defined. Since demand follows a Poisson distribution, by the matching of demand and supply unit, we may assume that the supply units are ordered in integer units, i.e., the order sizes are all integers. The two cost functions are well-defined with integer value of order sizes and we expand the domain for any value of order sizes by interpolating piecewise linear extensions of the integer values. It is clear that the extended functions preserve the properties of monotonicity. However, the order size of Q_k^* that balances the two functions is not necessarily an integer. Hence, for our infinite horizon model, we adopt a randomized balancing policy with two consecutive integers Q_k^l and $Q_k^u = Q_k^l + 1$ such that $Q_k^l < Q_k^* < Q_k^u$. In particular, $Q_k^* = \delta Q_k^l + (1-\delta) Q_k^u$ for some $0 < \delta < 1$. The randomized balancing policy then makes an order size of Q_k^l with probability δ and Q_k^u with probability $1-\delta$. The exact value of δ is hence determined by balancing the expected marginal early holding cost

$$\delta h_k HC_k(Q_k^l, \tilde{X}_k) + (1-\delta) h_k HC_k(Q_k^u, \tilde{X}_k)$$

and the expected marginal late holding and backorder cost

$$\delta(\pi + h'_{k+1}) BC_k(Q_k^l, \tilde{X}_k, NI_{k+1}) + (1-\delta)(\pi + h'_{k+1}) BC_k(Q_k^u, \tilde{X}_k, NI_{k+1}).$$

4.2 Modified Cost Balancing Policies

Based on the idea of dual-balancing policy discussed in the previous section, we next propose two modified cost balancing policies that perform better empirically. In Chapter 6, the three balancing policies will be tested and compared over a set of computational experiments.

4.2.1 Parameterized Balancing Policy

A *parameterized balancing policy* is defined as a policy which makes an order size of Q_k^* such that $h_k HC_k(Q_k^*, \tilde{X}_k) = \gamma(\pi + h'_{k+1}) BC_k(Q_k^*, \tilde{X}_k, NI_{k+1})$. The constant γ is called the *balancing ratio*. Specifically, for $\gamma = 1$, we get the dual-balancing policy discussed in the previous section. This modification method was originated from [3] in which the parameterized balancing policy was defined in a single-stage system.

Define $C(\gamma)$ as a function of the balancing ratio γ and it takes the value of the long-run expected average cost of γ-balancing policy. Based on empirical study on the function, we assume that there exists a local minimum γ^* in the subdomain of $[1, \infty]$. Consider a 4-stage serial network with demand following a Poisson distribution with parameter $\lambda = 4$. The unit echelon holding cost and unit backorder cost are $h = [0.25 \quad 0.25 \quad 0.25 \quad 0.25]$ and $\pi = 9$, respectively. The transportation time between every two consecutive stages is just one time period, i.e. $l_k = 1$ for $k = 1, 2, 3, 4$. First, we find that $C(11) > C(1)$ and suggest that the local minimu is within the range of $[1, 11]$. To show its existence, we plot the function $C(\gamma)$ numerically on the subdomain of $(1, 11)$. The subdomain is divided into small intervals, each of length 0.1. Hence, we have sample points $\gamma = 1 + 0.1i, i = 0, ..., 100$ and the long-run expected cost is approximated as the average cost of the first 10,000 periods. As illustrated in Figure 4-1, we can see clearly the existence of a local minimum. For all the experiments to

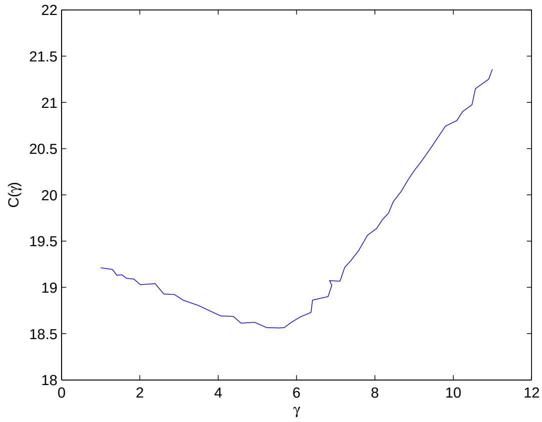

Figure 4-1: The long-run expected cost function $C(\gamma)$ on the balancing ratio γ

be taken in Chapter 6, we first establish the existence of such a local minimum and then find the exact value using bisection search, as described in Section 5.3.

4.2.2 Dual-Balancing Policy with Bounds

We can also modify the dual-balancing policy with the interval-constrained bounding technique, which was first developed in [3] for the single-stage inventory model. In this thesis, we use the same technique in the multi-echelon inventory model with the bounds of the optimal echelon base-stock level: whenever the after-order inventory position falls out of the bounds, the ordering decisions are modified correspondingly. To be specific, denote the lower bound and the upper bound of the optimal echelon base-stock level at stage k by s_k^l and s_k^u, respectively. The ordering decision at any period at stage k is made in the following procedure:

An order size Q_k is determined according to the dual-balancing policy as discussed in Section 4.1. If $Q_k \leq s_k^l - X_k$ and the on-hand inventory I_{k+1} at the preceding stage is positive, we change the order size to $min(s_k^l - X_k, I_{k+1})$; If order size $Q_k \geq s_k^u - X_k$, we reduce the order size to $s_k^u - X_k$. In all other cases, the order size just remains as Q_k.

With the constraint of bounds, we keep the echelon inventory level in the range of $[s_k^l, s_k^u]$ as close as possible. The performance of the dual-balancing policy with bounds can be better or worse than the dual-balancing policy, depending on the value of the bounds. In this thesis, we adopt the bounds on the optimal echelon base-stock level developed by Shang and Song [6].

It is well known that an echelon base-stock policy $s = (s_1, ..., s_n)$, where s_j is the echelon base-stock inventory position for stage $j, j = 1, ..., n$, is optimal for the system we study (for detailed proofs, refer to Chen and Zheng [12]). According to the policy, whenever the echelon inventory position at any stage k falls below s_k, an order will be placed to bring up the position to s_k. Otherwise, we do not make any order. It has been proven in [6] that the optimal base-stock level s can be found through minimizing the following n nested convex functions recursively with $\underline{C}_0(x) = (\pi + h_1')max\left\{0, -x\right\}$ and for $j = 1, ..., n$:

$$
\begin{aligned}
&\hat{C}_j(x) = h_j x + \underline{C}_{j-1}(x),\\
&C_j(y) = E[\hat{C}_j(y - D_{[0,l_k]})],\\
&s_j^* = argmin\left\{C_j(y)\right\},\\
&\underline{C}_j(x) = C_j(min\left\{s_j^*, x\right\}).
\end{aligned}
$$

However, the simple form does not guarantee an easy computational procedure to obtain the optimal base-stock levels and the optimal expected cost. Instead Shang and Song developed an upper bound and a lower bound on the average total echelon cost $C_j(y)$ for each stage j, assuming all downstream stages follow the optimal policy. The bounding functions can be obtained by solving $2n$ seperate newsvendor-type problems. Minimizing these functions, they derived upper and lower bounds for the optiaml echelon base-stock level. Denote the total leadtime demand of stages 1 to k by $\tilde{D}_k = \sum_{j=1}^{k} D_{[0,L_j]}$ and its cdf by $G_k(\cdot)$. The lower and upper bounds of the optimal base-stock level are

$$s_k^l = G_k^{-1}(\frac{\pi + \sum_{j=k+1}^n h_j}{\pi + \sum_{j=1}^n h_j})$$

and

$$s_k^u = G_k^{-1}(\frac{\pi + \sum_{j=k+1}^n h_j}{\pi + \sum_{j=k}^n h_j})$$

respectively. The simple average of the bounds form a heuristic base-stock policy that performs surprisingly well on some instances, with average error of 0.24% and max error of less than 1.5%.

Since the optimal base-stock level is within the bounds, we can modify the dual-balancing orders so that the order size is close to the range of the optimal order size and the resulting average cost may be closer to the optimal cost. As an attempt to improve the performance of the dual-balancing policy, the echelon inventory position after orders are placed is constrained by the upper and lower bounds. Moreover, we also implement the base-stock policy with base-stock level $s_k = \frac{s_k^l + s_k^u}{2}$ for each stage k and compare its performance with all the other balancing policies.

Next we show how to adopt the results of Shang and Song to our setting. In our model (with discrete time) the inbound and outbound shipments occur at the beginning of the period and the costs are accounted at the end of the period. In this case, we need to add one more period in calculating $\tilde{D}_k$. Therefore, function $G_k(\cdot)$, the cdf of $\tilde{D}_k$, is the cdf of Poisson distribution with parameter $\lambda(1 + \sum_{j=1}^k l_j)$. The bounds s_k^l and s_k^u are purely dependent on demand distribution, unit holding cost and unit backorder cost, which is easy to compute. Since the order sizes are bounded above by the on-hand inventory at the preceding stage, it is not always possible to keep the echelon inventory position after ordering in the range of $[s_k^l, s_k^u]$.

As we will show through numerical experiments in later chapters, the bounding behavior improves the performance of the dual-balancing policy we discussed in Section 4.1.

Chapter 5

Numerical Implementation

The purpose of the thesis is to evaluate the performance of the dual-balancing policy and several variants of this policy discussed in Section 4.2. through extensive numerical experiments. In this chapter, we will illustrate the implementation procedure of the dual-balancing policies.

5.1 Computation of Cost Formulas

The decision at stage k at any period with order size Q_k is assigned a marginal early holding cost of $h_k HC_k(\tilde{Q}_k, \tilde{X}_k)$ where

$$\begin{aligned} HC_k(\tilde{Q}_k, \tilde{X}_k) \quad &= \quad \sum_{t=L_k+1}^{\infty} [(\tilde{X}_k + \tilde{Q}_k)F(\lambda t, \tilde{X}_k + \tilde{Q}_k) - \tilde{X}_k F(\lambda t, \tilde{X}_k)] \\ &- \sum_{t=L_k+1}^{\infty} \lambda t [F(\lambda t, \tilde{X}_k + \tilde{Q}_k - 1) - F(\lambda t, \tilde{X}_k - 1)] \end{aligned}$$

Since $\tilde{X}_k, \tilde{Q}_k$ are given, let's define the following function on the integers:

$$\begin{aligned} H(L) \quad &= \quad \sum_{t=L}^{\infty} [(\tilde{X}_k + \tilde{Q}_k)F(\lambda t, \tilde{X}_k + \tilde{Q}_k) - \tilde{X}_k F(\lambda t, \tilde{X}_k)] \\ &- \sum_{t=L}^{\infty} \lambda t [F(\lambda t, \tilde{X}_k + \tilde{Q}_k - 1) - F(\lambda t, \tilde{X}_k - 1)] \end{aligned}$$

Hence, the marginal early holding cost is just $h_k H(L_k + 1)$. However, it is computationally demanding to exactly compute the value of $H(L_k + 1)$ because of the infinite sum. Instead, we manage to find an integer L^* such that $h_k H(L^*) \leq 10^{-3}$. Therefore, $h_k[H(L_k + 1) - H(L^*)]$ becomes a reasonable approximation of the value of $h_k H(L_k + 1)$. It only involves finite sum and is hence computable. Next we will prove the existence of L^* and show how to find its value.

Note that $H(L) \leq (\tilde{X}_k + \tilde{Q}_k) \sum_{t=L}^{\infty} F(\lambda t, \tilde{X}_k + \tilde{Q}_k)$. Define the function of t as $u(t) = F(\lambda t, \tilde{X}_k + \tilde{Q}_k)$ and it is decreasing with t. Therefore, the sum of $u(t)$ for all $t \geq L_k$ can be represented by the area of the shaded region in Figure 5-1, which is less than the area under the function $u(t)$ in the interval $[L-1, \infty]$. This implies that

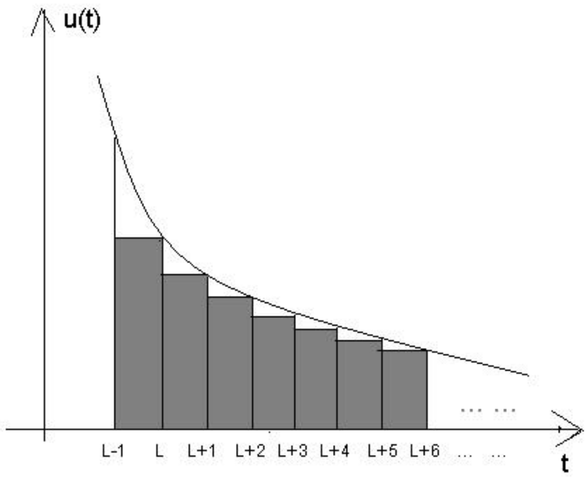

Figure 5-1: Graphical Representation

$$
\begin{aligned}
H(L) &\leq (\tilde{X}_k + \tilde{Q}_k) \int_{L-1}^{\infty} F(\lambda t, \tilde{X}_k + \tilde{Q}_k) dt \\
&= (\tilde{X}_k + \tilde{Q}_k) \sum_{j=0}^{\tilde{X}_k + \tilde{Q}_k} \int_{L-1}^{\infty} f(\lambda t, j) dt
\end{aligned}
$$

where $f(\beta, y)$ is the pdf of Poisson distribution with parameter β evaluated at y.

Define function $v(j) = \int_{L-1}^{\infty} f(\lambda t, j)dt$. Integrating by parts, we get

$$\begin{aligned} v(j) &= \int_{L-1}^{\infty} e^{-\lambda t}\frac{(\lambda t)^j}{j!}dt \\ &= \frac{e^{-\lambda(L-1)}}{\lambda}\frac{(\lambda(L-1))^j}{j!} + \int_{L-1}^{\infty} e^{-\lambda t}\frac{(\lambda t)^{j-1}}{(j-1)!}dt \\ &= \frac{f(\lambda(L-1), j)}{\lambda} + v(j-1) \end{aligned}$$

Using this recursive equation and $v(0) = \int_{L-1}^{\infty} e^{-\lambda t}dt = \frac{f(\lambda(L-1),0)}{\lambda}$, we obtain that $v(j) = \frac{F(\lambda(L-1),j)}{\lambda}$ and hence

$$\begin{aligned} H(L) &\leq (\tilde{X}_k + \tilde{Q}_k)\sum_{j=0}^{\tilde{X}_k+\tilde{Q}_k}\int_{L-1}^{\infty} f(\lambda t, j)dt \\ &= (\tilde{X}_k + \tilde{Q}_k)\sum_{j=0}^{\tilde{X}_k+\tilde{Q}_k}\frac{F(\lambda(L-1), j)}{\lambda} \\ &\leq (\tilde{X}_k + \tilde{Q}_k)(\tilde{X}_k + \tilde{Q}_k + 1)\frac{F(\lambda(L-1), \tilde{X}_k + \tilde{Q}_k)}{\lambda} \end{aligned}$$

Since the value of $F(\lambda(L-1), \tilde{X}_k + \tilde{Q}_k)$ decreases with the value of L and converges to 0 as L tends to infinity, there must exist an integer L^* such that

$$h_k(\tilde{X}_k + \tilde{Q}_k)(\tilde{X}_k + \tilde{Q}_k + 1)\frac{F(\lambda(L^*-1), \tilde{X}_k + \tilde{Q}_k)}{\lambda} \leq 10^{-3}$$

The exact value of L^* can be found using bisection search on the decreasing function $F(\lambda(L-1), \tilde{X}_k + \tilde{Q}_k)$.

As a result, the approximation of the marginal early holding cost becomes $h_k HC_k$ where

$$\begin{aligned} HC_k(\tilde{Q}_k, \tilde{X}_k) \quad &= \quad \sum_{t=L_k+1}^{L^*}[(\tilde{X}_k + \tilde{Q}_k)F(\lambda t, \tilde{X}_k + \tilde{Q}_k) - \tilde{X}_k F(\lambda t, \tilde{X}_k)] \\ &\quad - \sum_{t=L_k+1}^{L^*}\lambda t[F(\lambda t, \tilde{X}_k + \tilde{Q}_k - 1) - F(\lambda t, \tilde{X}_k - 1)] \end{aligned}$$

On the other hand, the marginal late holding and backorder cost can be computed easily as $(h'_{k+1} + \pi)BC_k$ where

$$BC_k(\tilde{Q}_k, \tilde{X}_k, NI_{k+1}) = \lambda L_k[F(\lambda(L_k + 1), NI_{k+1} - 1) - F(\lambda(L_k + 1), \tilde{X}_k + \tilde{Q}_k - 1)]$$

$$-(\tilde{X}_k + \tilde{Q}_k)[1 - F(\lambda(L_k + 1), \tilde{X}_k + \tilde{Q}_k)] + NI_{k+1}[1 - F(\lambda(L_k + 1), NI_{k+1})]$$

5.2 Balancing Order Sizes

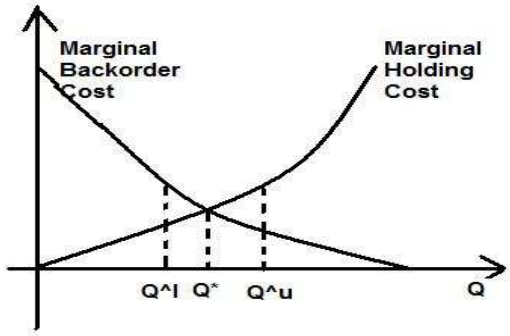

Figure 5-2: Decision of Order Size

The order size Q_k^* that balances the two opposing costs is obtained from two consecutive integers Q_k^l and Q_k^u and a probability value δ such that

$$\delta h_k HC_k(Q_k^l, \tilde{X}_k) + (1 - \delta)h_k HC_k(Q_k^u, \tilde{X}_k)$$

and

$$\delta(\pi + h'_{k+1})BC_k(Q_k^l, \tilde{X}_k, NI_{k+1}) + (1 - \delta)(\pi + h'_{k+1})BC_k(Q_k^u, \tilde{X}_k, NI_{k+1})$$

can be balanced. As illustrated in Figure 5-2, we can find Q_k^u as the first integer such that

$$h_k HC_k(Q_k^u, \tilde{X}_k) \geq \gamma(h'_{k+1} + \pi)BC_k(Q_k^u, \tilde{X}_k, NI_{k+1})$$

and Q_k^l is just $Q_k^u - 1$. This search method is executable if the order size at any stage is bounded above. For any stage $k \neq n$, the order size is bounded above by the on-hand inventory at the preceding stage, i.e. $\tilde{Q}_k \leq NI_{k+1} - X_k$. Stage n orders from an external supplier with infinite capacity but empirically the order size is a finite number. We may choose an upper bound such that if stage k makes an order size of the upper bound, the expected marginal early holding cost incurred is always less than the expected marginal late holding and backorder cost incurred given any echelon inventory. Given the optimal base-stock level s_n, we tried the upper bound of $s_n, 2s_n, 4s_n$, until we found that $4s_n$ can meet the requirement and it is used throughout all the experiments.

5.3 Local Minimum Balancing Ratio

In general, for each of the experiments to be tested, we follow these steps to establish the existence of local minimum and finally find its exact value:

Step 1 Find the smallest integer j such that $C(1+2j) > C(1)$ and the initial guess of the range where the local minimum reside in is $(1, 2j)$;

Step 2 Divide the subdomain into small intervals, each of length 0.1 and compute the average cost of the first 10,000 periods following the γ-balancing policy for each sample point γ;

Step 3 Plot the function $C(\gamma)$ on the subdomain $(1, 2j)$, if a local minimum can be observed, continue to next step; otherwise, extend the subdomain and go back to step 2;

Step 4 Use bisection search to find the exact value of the local minimum γ^* on the subdomain.

Chapter 6

Experimental Results

The goal of this thesis is to evaluate the empirical performance of the cost balancing policies in the periodic-reviewed serial network. Instead of applying the exact optimal base-stock policy, we adopt the approximately optimal policy obtained from Shang and Song [6], denoted by $APPROX$. The policy adopts an echelon base-stock level which equals the average of the lower bound and the upper bound of the corresponding optimal echelon base-stock level. It was shown in [6] that the approximately optimal policy produces the true optimal solution in 12 out of 32 experiments; the average error and the maximum error compared to the optimal base-stock policy are 0.116% and 0.557%, respectively. In our experiments, the approximately optimal policy is only considered as optimal in the base case. In all the other scenarios, the approximately optimal policy is not necessarily optimal but we still use it as a benchmark to evaluate the performance of the cost balancing policies. All policies considered and their short-hand names are listed in Table 6.1. Here the balancing ratio γ^* is computed using the method as described in Section 5.3.

Policy Name	Description
$APPROX$	approximately optimal policy based on [6]
DB	dual-balancing policy
DB_{Bound}	dual-balancing policy with bounds
$B(\gamma^*)$	γ^*-balancing policy
$B_{Bound}(\gamma^*)$	γ^*-balancing policy with bounds

Table 6.1: Policies considered

In all the experiments, we consider a simulated horizon of $T = 10,000$ periods and the long-run expected average cost is approximated by the average cost within this time horizon. Following a specific policy, denote the total cost incurred in period t as C_t^{Policy} and the computational time for the decision at stage k as τ_{kt}^{Policy}. The average total cost, denoted as C^{Policy}, and the average computational time for each decision, denoted as τ^{Policy}, can be computed as

$$C^{Policy} = \frac{1}{10000} \sum_{t=1}^{10000} C_t^{Policy} \quad , \tau^{Policy} = \frac{1}{10000} \sum_{t=1}^{10000} \frac{1}{n} \sum_{k=1}^{n} \tau_k^{Policy} t.$$

We use the approximately optimal policy as a benchmark and the accuracy of all the other policies is based on the error of the given policy against the approximately optimal policy. Specifically, we compute

$$Error^{Policy} = \frac{C^{Policy} - C^{Approx}}{C^{Approx}}.$$

	$n = 4$	$n = 5$
DB	8.55%	9.83%
DB_{Bound}	0.23%	0.80%
$B(\gamma^*)$	4.93%	7.75%
$B_{Bound}(\gamma^*)$	1.17%	0.90%

Table 6.2: Errors of the cost balancing policies in the base cases with n-stages

The Base Case Let's first consider two simple serial networks, one with 4 stages and another with 5 stages. In both systems, the per-unit backorder cost is $\pi = 9$ and the demand rate is $\lambda = 4$. Both systems have the same lead times at all stages. In particular, $l_k = 1$ for $k = 1, 2, 3, 4$ in the 4-stage system and $l_k = 1$ for $k = 1, 2, 3, 4, 5$ in the 5-stage system. The per-unit echelon holding cost at each stage is also set at the same value of 0.25. The dual-balancing policy performs surprisingly good in both cases, with an error of 8.55% in the 4-stage system and 9.83% in the 5-stage system. With the constraint of bounds, the errors are reduced to 0.23% and 0.80% respectively. We can also see that, as shown in Table 6.2, the performance of the other two cost balancing policies is much better than the dual-balancing policy. It is then possible that as we modify the system parameters, the performance of the cost balancing policies can be closer to the approximately optimal policy or even better than that. Hence, in the next section we explore the performance of the policies by variating the system parameters.

6.1 Experiment Design

There are several parameters involved in the system, including the number of stages n, the per-unit echelon holding cost h, the per unit backorder cost π, the leadtime l and the demand rate λ. The experiment design is oriented around the Base Case discussed in the previous section and five sets of scenarios, each of which expands the Base Case in one dimension. In each set of the scenarios we vary specific system parameters. The first two sets of scenarios study the effect of the per-unit holding cost h and the per-unit backorder cost π. The third set of scenarios studies the effect of demand rate λ. Then we consider a set of scenarios when the lead times at different stages vary. Moreover, we study the effect of the per-unit backorder cost with long lead times in the last scenario. Next we describe the manner in which the parameters of the Base Case are varied, in each of the five sets of scenarios.

Holding Cost Scenarios In the base case, we have flat per-unit echelon holding cost at all stages. To study the effect of the per-unit echelon holding cost at some stage on the whole system, we increase the per-unit echelon holding cost to 10 for one stage in one scenario. Table 6.3 shows the variation scheme for the two serial networks.

Backorder Cost Scenarios In the base case, the per-unit backorder cost is set at $\pi = 9$. Our sensitivity analysis takes in a full range of $\pi = 5, 9, 29, 49, 99$ to test the performance of the cost balancing policies when the backordering cost is a lot higher than the echelon holding cost for each unit. Given the same system parameters and echelon inventory, with higher per-unit backorder cost, the balancing order size may increase and yet it is constrained by the on-hand inventory available at the preceding stage.

Scenarios for h in 4-stage system					
$Flat$	0.25	0.25	0.25	0.25	
H_1	2.5	0.25	0.25	0.25	
H_2	0.25	2.5	0.25	0.25	
H_3	0.25	0.25	2.5	0.25	
H_4	0.25	0.25	0.25	2.5	
Scenarios for h in 5-stage system					
$Flat$	0.25	0.25	0.25	0.25	0.25
H_1	2.5	0.25	0.25	0.25	0.25
H_2	0.25	2.5	0.25	0.25	0.25
H_3	0.25	0.25	2.5	0.25	0.25
H_4	0.25	0.25	0.25	2.5	0.25
H_5	0.25	0.25	0.25	0.25	2.5

Table 6.3: Holding Cost Scenarios Scheme

Demand Scenarios The demand at the end of each period follows a Poisson distribution. Its expectation and variation are both set at $\lambda = 4$ in the base case. By varying this parameter, we study the performance of the policies under both extremely small demand and extremely large demand. In particular, experiments are taken on $\lambda = 1, 4, 8, 16, 32$.

Lead Time Scenarios We assume the same lead time at all stages in the base case while it is not practical in the real case. Hence, we take several considerations in the sensitivity analysis, as listed in Table 6.4.

Scenarios for l in 4-stage system					
$Flat$	1	1	1	1	
$Lead_1$	10	1	1	1	
$Lead_2$	1	10	1	1	
$Lead_3$	1	1	10	1	
$Lead_4$	1	1	1	10	
Scenarios for l in 5-stage system					
$Flat$	1	1	1	1	1
$Lead_1$	10	1	1	1	1
$Lead_2$	1	10	1	1	1
$Lead_3$	1	1	10	1	1
$Lead_4$	1	1	1	10	1
$Lead_5$	1	1	1	1	10

Table 6.4: Lead Time Scenarios Scheme

Backorder Cost with Long Lead Time Scenarios In this scenario, we want to study the effect of a combination of long lead time and high backorder cost on the performance of the cost balancing policies. The lead time at all stages are assumed to be 10 instead of 1 and we take the backorder cost as $\pi = 5, 9, 29, 49, 99$.

6.2 Sensitivity Analysis

Holding Cost Effect As we increase the per-unit echelon holding cost in the 4-stage system, the error of the dual-balancing policy has reduced from 8.55% to the least of 2.82% and the most of 4.46%. The increment in per-unit echelon holding cost decreases the balancinng order size and also decreases the long-run expected average cost, which implies that in some periods we may order too much according to the dual-balancing policy. Hence, with the constraint of bounds, we see a significant improvement in the overall performance. The dual-balancing policy with bounds results in less than 1% of error in all the scenarios and only 0.02% in the last scenario, which shows almost the same performance as the approximately optimal policy. On the other hand, when we apply the balancing ratio γ^*, the expected average cost is always closer to the approximately optimal policy than that of the original dual-balancing policy with the balancing ratio of 1 (see Table 6.5). The constraint of bounds also reduces the error of the γ^*-balancing policy on an average of 2.11%. Comparing the different scenarios for the γ^*-balancing policy, the worst performance appears in the $Flat$ scenario when the per-unit echelon holding cost at all stages are the same. Similar trends can also be observed in the 5-stage system, as illustrated in Table 6.6.

Scenarios	$Flat$	H_1	H_2	H_3	H_4
DB	8.55%	4.46%	2.82%	3.30%	2.85%
DB_{Bound}	0.23%	0.65%	0.42%	0.10%	0.02%
$B(\gamma^*)$	4.93%	3.06%	2.23%	3.06%	2.06%
$B_{Bound}(\gamma^*)$	1.17%	0.86%	1.72%	1.06%	0.02%

Table 6.5: Error of policies in holding cost scenarios for 4-stage system

Scenarios	$Flat$	H_1	H_2	H_3	H_4	H_5
DB	9.83%	3.44%	3.66%	4.45%	4.00%	2.30%
DB_{Bound}	0.80%	0.55%	0.45%	0.41%	0.19%	0.23%
$B(\gamma^*)$	7.75%	3.39%	3.19%	3.91%	3.07%	2.05%
$B_{Bound}(\gamma^*)$	0.90%	1.90%	2.11%	1.23%	0.91%	0.90%

Table 6.6: Error of policies in holding cost scenarios for 5-stage system

Backorder Cost Effect As the per-unit backorder cost increases, the balancing orders are increasing, which shows an opposite effect of the per-unit echelon holding cost. The empirical performance again reveals that the errors of the dual-balancing policy and the γ^*-balancing policy increase almost linearly with the per-unit backorder cost in both the 4-stage and the 5-stage system (see Figure 6-1). In the 4-stage system, the errors of the γ-balancing policy is about half of the errors of the dual-balancing policy in all the scenarios. The constraint of bounds on the dual-balancing policy again improves the performance, but is not consistently good in all the scenarios in the 4-stage system, as shown in Table 6.7. With the per-unit backorder cost of 49, the dual-balancing policy with bounds shows an error of 2.29% while in the other scenarios the errors are all less than 1%. However, when the cost increases to 99, the error becomes -1.46%, which shows a better performance than the approximately optimal policy. On the other hand, when there are five stages in the system, the dual-balancing policy with bounds performs consistently good with an average error of 0.60%, as shown in Table 6.8. The γ^*-balancing policy with bounds shows a lot better performance than the dual-balancing policy or the γ^*-balancing policy. However, it has similar performance as the dual-balancing policy with bounds, resulting in an average error of 0.34% and 0.76% in the 4-stage and 5-stage systems, respectively. Hence, we may predict that as the per-unit backorder cost increases, the dual-balancing policy and the γ^*-balancing policy will become much less accurate but with the constraint of bounds the performance will still remain consistently good.

Demand Effect Similarly to the previous sets of scenarios, we did not observe any obvious effect of the demand rate λ on the two cost balancing policies with bounds, as shown in Table 6.9 and 6.10. These two policies perform consistently good with

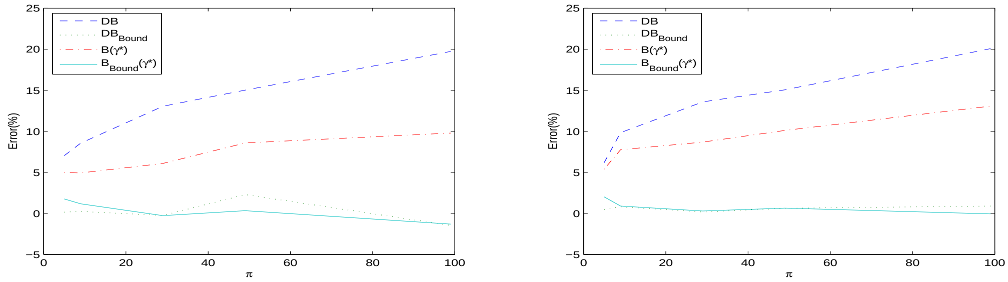

(a) $n = 4$ (b) $n = 5$

Figure 6-1: Error of policies in backorder cost scenarios for n-stage system

Scenarios	$\pi = 5$	$\pi = 9$	$\pi = 29$	$\pi = 49$	$\pi = 99$
DB	7.03%	8.55%	13.05%	15.01%	19.73%
DB_{Bound}	0.16%	0.23%	-0.24%	2.29%	-1.46%
$B(\gamma^*)$	4.98%	4.93%	6.09%	8.58%	9.79%
$B_{Bound}(\gamma^*)$	1.76%	1.17%	-0.27%	0.33%	-1.30%

Table 6.7: Error of policies in backorder cost scenarios in 4-stage system

Scenarios	$\pi = 5$	$\pi = 9$	$\pi = 29$	$\pi = 49$	$\pi = 99$
DB	6.17%	9.83%	13.58%	15.05%	20.07%
DB_{Bound}	0.49%	0.80%	0.17%	0.62%	0.90%
$B(\gamma^*)$	5.38%	7.75%	8.69%	10.11%	13.07%
$B_{Bound}(\gamma^*)$	2.02%	0.90%	0.29%	0.64%	-0.06%

Table 6.8: Error of policies in backorder cost scenarios in 5-stage system

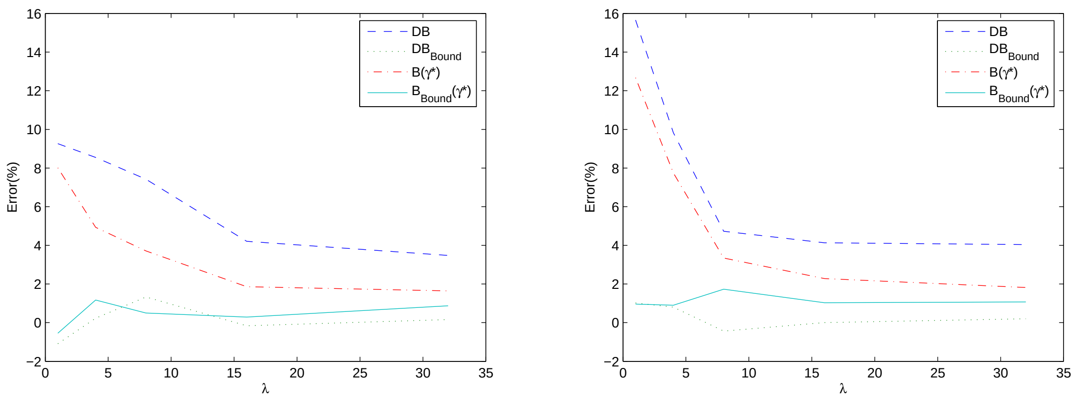

(a) $n = 4$ (b) $n = 5$

Figure 6-2: Error of policies in demand scenarios for n-stage system

Scenarios	$\lambda = 1$	$\lambda = 4$	$\lambda = 8$	$\lambda = 16$	$\lambda = 32$
DB	9.26%	8.55%	7.42%	4.21%	3.48%
DB_{Bound}	-1.09%	0.23%	1.33%	-0.16%	0.16%
$B(\gamma^*)$	8.01%	4.93%	3.71%	1.86%	1.64%
$B_{Bound}(\gamma^*)$	-0.55%	1.17%	0.50%	0.29%	0.87%

Table 6.9: Error of policies in demand scenarios in 4-stage system

Scenarios	$\lambda = 1$	$\lambda = 4$	$\lambda = 8$	$\lambda = 16$	$\lambda = 32$
DB	15.66%	9.83%	4.73%	4.13%	4.04%
DB_{Bound}	1.02%	0.80%	-0.44%	0.00%	0.20%
$B(\gamma^*)$	12.67%	7.75%	3.34%	2.28%	1.82%
$B_{Bound}(\gamma^*)$	0.96%	0.90%	1.73%	1.03%	1.07%

Table 6.10: Error of policies in demand scenarios in 5-stage system

errors less than 2% in all the scenarios and there are several cases when the error becomes negative. Specifically, in the 4-stage system, the dual-balancing policy with bounds has a long-run expected average cost 1.09% less than that of the approximately optimal policy. The γ^*-balancing policy with bounds does not perform any better than the dual-balancing policy with bounds except in the case of 5-stage system with $\lambda = 1$ where the error of the policy is 0.96% and the error of the dual-balancing policy with bounds is 1.02%. Since the constraint of bounds is implemented after we have computed the balancing order size which is not the same for the policy with balancing ratio 1 and γ^*, the overall performance of the two cost balancing policies with bounds is not comparable. Without the constraint of bounds, the dual-balancing policy and the γ^*-balancing policy are significantly affected by the demand rate λ. As shown in Figure 6-2, the errors of these two policies decrease sharply at small λ but remains consistent after some critical points. If we can further increase the demand rate, the computational time may be much longer but it is also possible that all the cost balancing policies will show very small errors or even negative ones.

Lead Time Effect The lead time at all stages are assumed to be at the same value in the base case and the long-run expected average cost of the dual-balancing policy in the 4-stage system is 8.55% larger than that of the approximately optimal policy. As the lead time at the lowest stage increases to ten periods, the error is reduced down to only 4.43%. In the other scenarios, we also observe an average decrease of 0.77% in the error of the dual-balancing policy(see Table 6.11). Hence, the performance of the dual-balancing policy might be better with longer lead times with different values at the stages. The constraint of bounds improves the performance of the dual-balancing policy in all the scenarios and the dual-balancing policy with bounds has an error within 1% in all the scenarios in the 4-stage system. Another modified cost balancing policy, the γ^*-balancing policy provides an average of 2.89% decrease in the error of the dual-balancing policy. But it is no better than the dual-balancing policy with bounds in all the scenarios. Among all the scenarios, the γ^*-balancing policy has the worst performance in the *Flat* scenario when the lead time at all stages are the same. This trend has also been observed for the dual-balancing policy. Hence, for the cost balancing policies without bounds, the variation of the lead time does not only change the ordering sizes at each period but also improves the accuracy.

Scenarios	$Flat$	$Lead_1$	$Lead_2$	$Lead_3$	$Lead_4$
DB	8.55%	4.43%	7.30%	7.92%	6.07%
DB_{Bound}	0.23%	-0.71%	0.59%	0.84%	0.90%
$B(\gamma^*)$	4.93%	2.25%	3.28%	2.27%	3.07%
$B_{Bound}(\gamma^*)$	1.17%	1.14%	0.83%	0.58%	-0.06%

Table 6.11: Error of policies in different scenarios of lead time for 4-stage system

Scenarios	$Flat$	$Lead_1$	$Lead_2$	$Lead_3$	$Lead_4$	$Lead_5$
DB	9.83%	5.47%	3.45%	4.46%	7.42%	6.87%
DB_{Bound}	0.80%	-0.17%	-0.55%	0.14%	1.77%	0.91%
$B(\gamma^*)$	7.75%	3.08%	2.55%	2.68%	3.62%	4.55%
$B_{Bound}(\gamma^*)$	0.90%	1.24%	0.09%	0.83%	1.00%	0.38%

Table 6.12: Error of policies in different scenarios of lead time for 5-stage system

Backorder Cost with Long Lead Time Effect From the previous experiments, we have found that the increase of the per-unit backorder cost will result in worse performance of the dual-balancing policy and the γ^*-balancing policy but has little effect on the cost balancing policies with bounds. As the lead time becomes longer, the two policies without bounds show better performance and the bounded balancing policies remains consistently good in accuracy. In this set of scenarios, we increase the lead time at all stages to 10 and test the performance with high per-unit backorder cost.

In the 4-stage system, with $\pi = 29$, the errors of the cost balancing policies are on average 2.21% lower than those in the base case. When the per-unit backorder cost has increased to 99, the cost balancing policies show different levels of accuracy. In particular, the dual-balancing policy has a decrease of 0.98% in error. In contrast, the error of the γ^*-balancing policy increases from 4.93% to 6.31% and the two cost balancing policies with bounds also results in a larger error (see Table 6.13). Hence, in the 4-stage system, the combined effect of the two system errors does depend on the extend to which each parameter is modified. On the other hand, the experiments taken on 5-stage system shows that all the policies are performing better than they do in the base case.

Scenarios	$\pi = 5$	$\pi = 9$	$\pi = 29$	$\pi = 49$	$\pi = 99$
DB	4.03%	4.94%	5.28%	5.38%	7.57%
DB_{Bound}	0.89%	0.16%	-1.76%	-0.64%	1.43%
$B(\gamma^*)$	3.61%	3.86%	4.31%	5.24%	6.31%
$B_{Bound}(\gamma^*)$	1.16%	-0.32%	-1.79%	0.61%	1.98%

Table 6.13: Errors of Policies in backorder cost with long lead time scenario for 4-stage system

Scenarios	$\pi = 5$	$\pi = 9$	$\pi = 29$	$\pi = 49$	$\pi = 99$
DB	5.32%	5.61%	6.78%	7.56%	8.21%
DB_{Bound}	0.52%	0.45%	-0.17%	-1.81%	-0.42%
$B(\gamma^*)$	4.33%	5.31%	6.20%	6.81%	7.18%
$B_{Bound}(\gamma^*)$	0.37%	0.31%	-0.11%	-1.80%	0.44%

Table 6.14: Errors of Policies in backorder cost with long lead time scenario for 5-stage system

Computational Time Here we consider another aspect of the performance of the cost balancing policies, that is, the computational time taken per decision making. As illustrated in Figure 6-3(a), the computational time increases linearly with the demand rate λ in both systems. The effect of the per-unit backorder cost on the compuational time for each decision is slightly different. In the 4-stage system, the computational time is consistent as the per-unit backorder cost increases. As the number of stages increases to 5, the compuational time remains consistent for small π but becomes extremely large as π has increased to 99. The extra long time may be due to the bisection search for the balancing order size since the large value of π expands the range of the search process.

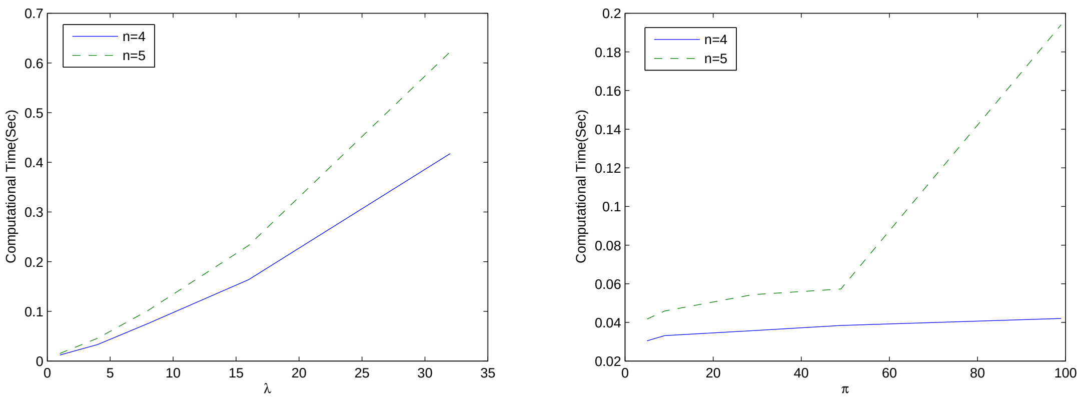

(a) Effect of demand parameter λ (b) Effect of backorder cost π

Figure 6-3: Computational time of dual-balancing orders for n-stage system

6.3 Robust Analysis

In this section, we study the performance of the policies over all the 52 scenarios. First, we compute the times each policy is the best among all the policies. The approximately optimal policy is dominating in 36 cases but there are 11 cases when the dual-balancing policy with bounds over performs and another 5 cases when the γ^*-balancing policy with bounds results in the best. Unfortunately we do not observe any case when the balancing policy or the γ^*-balancing policy show better performance than the approximately optimal policy. In Section 6.4, we will show some further examples in which all the policies considered are more accurate than the approximately optimal policy.

	$APPROX$	DB	DB_{Bound}	$B(\gamma^*)$	$B_{Bound}(\gamma^*)$
$n=4$	16	-	5	-	4
$n=5$	20	-	6	-	1

Table 6.15: Number of times each policy has the least total cost in an n-stage system

Next we compare the improvement of the three types of modification on the dual-balancing policy, which are the constraint of bounds, the local minimum balancing ratio γ^* and the combination of both. The improvement is expressed as the difference between the error of the dual-balancing policy and the modified cost balancing policy. The constraint of bounds shows an average of 7.40% and a maximum of 21.20% improvement on the dual-balancing policy. The dual-balancing policy with bounds has an error greater than 1% in only five out of the 52 cases, which is very close to the approximately optimal policy. Changing the balancing ratio also improves the performance of the dual-balancing policy but the maximum decrease in error is just 9.94%, which is not as effective as the constraint of bounds. The combination of both the bounds and the balancing ratio shows equivalent improvement on the dual-balancing policy as the constraint of bounds only, with an average of 6.93% and a maximum of 21.03% decrease in the error. Hence, there are two practical methods to improve the performance of the dual-balancing policy, namely, the constraint of bounds and the combination of the bounds and the balancing ratio.

	average	min	max
Bound	7.40%	2.08%	21.20%
Ratio	2.77%	0.05%	9.94%
Bound&Ratio	6.93%	1.10%	21.03%

Table 6.16: Improvement after modification on the dual-balancing policy

6.4 Dominating Cases

In the previous sensitivity analysis, we find several cases when the performance of some of the cost balancing policies was better than the approximately optimal policy. For example, the long-run expected average cost according to the dual-balancing policy with bounds is 0.71% less than that according to the approximately optimal policy when we increase the lead time at stage 1 in the base case from 1 to 10 in the 4-stage system. It is then possible that as we modify the system paramters furthermore, the approximately optimal policy may be dominated by some of the cost balancing policies. That is, the cost balancing policies are closer to the optimal policy than the approximately optimal policy.

It has been found that the error of the cost balancing policies without bounds decreases with the demand rate λ and increases with the per-unit backorder cost π. Hence, in this section, we set the parameters at $\lambda = 32$ and $\pi = 5$. Consider the 5-stage system with unit echelon holding cost $h = [0.25 \quad 0.25 \quad 0.25 \quad 0.25 \quad 2.5]$. The lead time at each stage is just one period. We increase the lead time at stage 2 from 1 period up to 70 periods to explore the evolution of the performance of the cost balancing policies.

The errors of the dual-balancing policy and the γ^*-balancing policy are decreasing sharply as the lead time increases. When the lead time at stage two is 10 periods, we observe only one negative error among the cost balancing policies. That is, the dual-balancing policy with bounds shows an improvement of 0.55% over the approximately optimal policy (see Table 6.17). But as the lead time has increased to 30 periods, all the cost balancing policies show a lower long-term expected average cost than the approximately optimal policy, with the maximum improvement of 1.54%. Moreover, in the case when the lead time is 50 periods, the cost balancing policies dominate the approximately optimal policy for at least 2%. Hence, we see a better performance of the cost balancing policies compared to the approximately optimal policy with long lead time at stage 2 as the lead time increases. With extremely long lead times (50 or 70 periods), the cost balancing policies show consistent dominance of 2-3% over the approximately optimal policy. Lastly, we consider the case of $l = [6 \quad 10 \quad 14 \quad 18 \quad 22]$, when the total lead time is 70 periods and the lead time at every stage is different from each other. The dominance over approximately optimal policy is again observed but is on average 0.52% less than those in the case of $l = [1 \quad 70 \quad 1 \quad 1 \quad 1]$.

lead time					Error of Policies			
l_1	l_2	l_3	l_4	l_5	DB	DB_{Bound}	$B(\gamma^*)$	$B_{Bound}(\gamma^*)$
1	1	1	1	1	9.83%	0.80%	7.75%	0.90%
1	10	1	1	1	3.45%	-0.55%	2.55%	0.09%
1	30	1	1	1	-0.11%	-1.54%	-0.31%	-1.33%
1	50	1	1	1	-2.13%	-2.99%	-2.26%	-3.01%
1	70	1	1	1	-2.08%	-3.00%	-2.17%	-3.12%
6	10	14	18	22	-1.54%	-2.39%	-1.77%	-2.60%

Table 6.17: Error of policies, for 5-stage system with $\lambda = 32, \pi = 5$ and $h = [0.25 \quad 0.25 \quad 0.25 \quad 0.25 \quad 2.5]$

Chapter 7

Conclusion

We consider a periodic-reviewed serial network with n stages and an external supplier with infinite capacity. The demands in each period are independently distributed Poisson random variables and unsatisfied demands are backlogged until the supplies arrive to meet the demands. The objective of the inventory problem is to minimize the long-run expected cost including the holding cost for each unit staying in the system and the backorder cost for each backlogged demand unit. The echelon base-stock policies are known to be optimal, but it can be hard to compute the optimal base-stock levels. The dual-balancing policy aims to choose an order quantity which makes these two cost functions equal. It can be further modified with the constraint of additional bounds on the echelon inventory level, different balancing ratio, or a combination of both.

In this thesis, we study the performance of the balancing policies through extensive numerical experiments. The long-run expected cost is approximated by the average cost in the first 10,000 periods. We adopt the upper and lower bounds on the optimal echelon inventory base level developed by Shang and Song [6] in two ways. First, the average of the bounds is considered as an approximation of the optimal echelon inventory base level. Moreover, we modify the dual-balancing policy with the constrain of these two bounds. The accuracy of all the balancing policies is evaluated in terms of the error against the approximatly optimal policy and the efficiency of the balancing policies is evaluated in terms of the computational time for each decision.

The dual-balancing has a worst-case guarantee of 2 but the numerical results show surprisingly good performance. The modification of different parameters has great impact on the performance of the balancing policies. The errors of the balancing policies without bounds are increasing with the per-unit backorder cost but decreasing with the demand parameter. As the per-unit echelon holding cost or the lead time increases at some stage, the errors decreases accordingly depending on that stage. On the other hand, the performance of the bounded balancing policies are more consistent under different parameters. In terms of the efficiency, the computational time is monotone increasing with the per-unit backorder cost and the demand parameter. The constrain of bounds on the dual-balancing policy makes large improvement in the accuracy without extending the computational time. In some cases, the combination of adding bounds and modifying balancing ratio shows the least cost among all the policies. We also observe several cases when some of the balancing policies obviously dominate the approximately optimal policy.

Bibliography

[1] Restef, L., R. O. Roundy and D. B. Shmoys, V. A. Truong: Approximation algorithms for capacitated stochastic inventory control models, *Math. of Oper. Res.*, Vol. 32-2, p284-302, 2007.

[2] Restef, L., R. O. Roundy and V. A. Truong: Provably near-optimal balancing policies for multi-echelon stochastic inventory control models. Submitted, 2005.

[3] Gavin H., Peter J., Restef L., R. O. Roundy and D. B. Shmoys: New policies for stochastic inventory control models: theoretical and computational results. Submitted, 2006.

[4] Axsäter: Simple solution procedures for a class of two-echelon inventory problems, *Oper. Res.*, Vol. 38, p64-69, 1990.

[5] Joneja, D.: The joint replenishment problem: New heuristics and worst-case performance bounds, *Oper. Res.*, Vol. 38, p723-771, 1990.

[6] Shang, H. K. and Song J.: Newsvendor bounds and heuristics for optimal policies in a serial supply-chain, *Management Science*, Vol 49, p618-638, 2003.

[7] A. Clark and H.Scarf: Optimal Policies for a multiechelon inventory problem, *Management Science*, Vol. 6, pg 475-490, 1960.

[8] F. Chen and J. Song: Optimal policies for multi-echelon inventory problems with Markov-modulated demand, *Operations Research*, Vol. 49, pg 226-234, 2001.

[9] A. and Zipkin, P.: Evaluation of the one-for-one replenishment policies for multiechelon inventory systems, *Management Science*, Vol. 37, pg 68-83, 1991.

[10] Zipkin, P.: Evaluation of base-stock policies in multiechelon inventory systems with compound-Poisson demands, *Naval Research Logistics*, Vol. 38, pg 397-412, 1991.

[11] A. Muharremoglu and J. N. Tsitsiklis: A single-unit decomposition approach to multi-echelon inventory systems, *Operations Research*, Vol. 56, pg 1089-1103, 2008.

[12] F. Chen and Y. S. Zheng: Lower Bounds for multi-echelon stochastic inventory systems, *Management Sci*, Vol. 6, pg 475-490, 1994.

[13] Erkip, N., W. H. Hausman, S. Nahmias. Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands, *Management Sci.* 36 381-392, 1990.

[14] Lee, H. L., K. C. So, C. S. Tang. The value of information sharing in two-level supply chains, *Management Sci.* 46:626-643, 1999.

[15] Zipkin, P. H. 2000. Foundations of Inventory Management. McGraw-Hill, New York.

[16] L. Dong and H. L. Lee. Optimal policies and approximations for a serial multi-echelon inventory system with time-correlated demand. *Operations Research*, 51, 2003.

[17] J. Song and P. Zipkin. Inventory control in a fluctuating demand environment. *Operations Research*, 41:351-370, 1993.