Optimizing Beer Distribution Game Order Policy Using Numerical Simulations

by

Qinwen Xiao

B.Eng, Zhejiang University (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

September 2009

Author .....
School of Engineering
August 7, 2009

Certified by .....
James B. Orlin
Professor of Sloan School of Management
Thesis Supervisor

Certified by .....
David Simchi-Levi
Professor of Civil and Environmental Engineering
Thesis Supervisor

Accepted by .....
Jaime Peraire
Professor of Aeronautics and Astronautics
Director, Computation for Design and Optimization Program

Optimizing Beer Distribution Game Order Policy Using Numerical Simulations

by

Qinwen Xiao

Submitted to the School of Engineering
on August 7, 2009, in partial fulfillment of the
requirements for the degree of
Master of Science in Computation for Design and Optimization

Abstract

One of the major challenges in supply chain management is the level of information availability. It is very hard yet important to coordinate each stage in the supply chain when the information is not centralized and the demand is uncertain.

In this thesis, I analyzed the bullwhip effect in supply chain management using the MIT Beer Distribution Game. I also proposed heuristics and models to optimize the MIT Beer Distribution Game order policy when the customer's demand is both known and unknown. The proposed model provides each player with an order policy based on how many weeks of inventory the player needs to keep ahead to minimize the global cost of the supply chain. The optimized order policy is robust, practical, and generated by numerical simulations. The model is applied in a number of experiments involving deterministic and random demand and lead time. The simulation results of my work are compared with two other artificial agent algorithms, and the improvements brought by my results are presented and analyzed.

Thesis Supervisor: James B. Orlin
Title: Professor of Sloan School of Management

Thesis Supervisor: David Simchi-Levi
Title: Professor of Civil and Environmental Engineering

Acknowledgments

I would like to thank my supervisor, Professor James B. Orlin and Professor David Simchi-Levi, for coming up with this meaningful topic and their guidance throughout my research. Without their invaluable advice and assistance, it would not have been possible for me to complete this thesis. During my stay at MIT, I learned how to consider a problem and conduct research, and gained much other knowledge, which, I think, will benefit my future study, career, and my whole life. I really thank them for giving me this opportunity to study with them.

I would like to thank all my friends at MIT: Jan, Höskuldur, Michael, Jonathan, Alex, Ben, Constantin, etc. With all you guys, I have spent a wonderful time during my study at MIT. I appreciate all your encouragement, companionship, and help.

Also, I want to thank all the members in the program of Computation for Design and Optimization and the Singapore-MIT Alliance. I would like to thank Laura Koller and Jocelyn Sales for their help with my academic affairs.

Contents

1 Introduction **13**

1.1 Bullwhip Effect 14

1.2 MIT Beer Distribution Game 16

1.3 Demand Forecast 18

1.4 Literature Review 20

1.4.1 Studies of the Bullwhip Effect 20

1.4.2 Study of the Beer Distribution Game 22

1.5 Thesis Objectives and Organization 23

2 Modeling the Beer Distribution Game Order Policy **25**

2.1 Assumptions 25

2.2 Heuristics 27

2.2.1 Notation 27

2.2.2 Q-model Description 28

2.2.3 Expected Demand 29

2.2.4 Expected Inventory 32

2.2.5 Order Quantity 32

3 Simulating the Beer Distribution Game **35**

3.1 Simulation Process 35

3.2 Current Simulation 37

3.3 Revised Simulation 38

4 Simulation Results and Analysis **43**

4.1 Simulation Results 43

4.1.1 Testing Scenarios 43

4.1.2 Order Policies 44

4.1.3 Testing Results 45

4.1.4 Analysis of Results 45

4.2 Smoothing Technique 49

4.2.1 Testing Results 50

4.2.2 Analysis of Results 51

4.3 General Scenario Simulations 53

4.3.1 Randomly Distributed Demand 54

4.3.2 Normally Distributed Demand 55

4.3.3 Initial Inventory 56

4.3.4 Seasonal Demand 58

4.3.5 Robust Analysis 59

5 Conclusions and Future Work **61**

List of Figures

1-1 Goods and Information Flows in a Supply Chain 14

1-2 Order Flow Demonstration in the Beer Distribution Game with One-week Information Delay 18

1-3 Multi-retailer Structure in a Supply Chain 21

2-1 Beer Game Supply Chain with Constant Lead Time 27

3-1 Beer Game Supply Chain with Stochastic Lead Time 36

3-2 Beer Game Supply Chain Simulation Order and Delivery Flowchart . 40

4-1 Inventory Level under the Genetic Algorithm Order Policy in a Scenario with Constant Lead Time and Stochastic Demand 47

4-2 Demand forecast under the Q-model Order Policy with Stochastic Demand and Lead Time 48

4-3 Inventory Level under the Q-model Order Policy $Q = [0, 2.75, 4.5, 4]]$ in a Scenario with Constant Lead Time and Stochastic Demand . . . 49

4-4 Inventory Level under the adjusted Q-model Order Policy $Q = [0, 2.75, 4.5, 4]$ in Comparison to Figure 4-3 50

4-5 Orders Placed by Every Player under Optimized Q-model Order Policy $Q = [0, 2.75, 4.5, 4]$ in a Scenario with Constant Lead Time and Stochastic Demand 51

4-6 Orders Placed by Every Player under Adjusted Q-model Order Policy $Q = [1, 2.75, 4.5, 4]$ in a Scenario with Constant Lead Time and Stochastic Demand 52

4-7 Inventory Level with Order Policy under the smoothed Q-model $Q = [1.5, 2.5, 4.5, 3]$ in a Scenario with Stochastic Lead Time and Demand 53

4-8 Order Placed by Each Player under the smoothed Q-model $Q = [1.5, 2.5, 4.5, 3]$ in a Scenario with Stochastic Lead Time and Demand 54

4-9 Comparison of Normally Distributed Demand (12,2) and Random Distributed Demand between 0-15 with Different Initial Inventory 58

4-10 Robust Analysis of Order Policy Suggested by the Smoothed Q-model 60

List of Tables

3.1 Simulation Results of the Retailer and the Wholesaler Using Simulation Approach and RLOM Order Policy Proposed by Chaharsooghi et al. . 37

3.2 Revised Simulation Data of the Retailer and the Wholesaler under Stochastic Lead Time in Comparison to Table 3.1 41

4.1 Basic Testing Scenarios of Constant and Stochastic Demand and Lead Time 44

4.2 Optimized Order Policies Recommended by GA, RLOM, and Q-model in the Basic Constant and Stochastic Scenarios 45

4.3 Comparison of Results of the GA, RLOM, and Q-model in the Basic Constant and Stochastic Scenarios 45

4.4 Optimized Order Policies under the Smoothed Q-model with Comparison to the GA, RLOM, and the Original Q-model 51

4.5 Inventory Costs under the Smoothed Q-model with Comparison to the GA, RLOM, and the Original Q-model 52

4.6 Order Policy Comparison under 50 Sample Sets of Random Distributed Demand between 0 and 15 55

4.7 Comparison of Results under 50 Sets of Stochastic Demand in 0-15 . 55

4.8 Order Policy and Cost under Normally Distributed Demand 55

4.9 Comparison of Order Policies and Costs under Normally Distributed Demand with Initial Inventory 12 and 20 Cases of Beer 56

4.10 Comparison of Order Policies and Costs under Normally Distributed Demand (12, 2) with Different Initial Inventory 57

4.11 Comparison of Order Policies and Costs under Stochastically Distributed Demand between (0,15) with Different Initial Inventory 57
4.12 Seasonal Demand Distribution 58
4.13 Optimized Order Policy under the Smoothed Q-model and Its Costs in the Scenario of Seasonal Demand 59
4.14 Robust Analysis of Order Policy Generated by Smoothed Q-model . . 59

Chapter 1

Introduction

Globalization has brought companies great opportunities for business while it also leads companies to fierce global competition and more complex operations. Over the past 20 years, looking for ready sources of components, lower-priced labor, and talented designers and engineers, companies, especially high-tech ones, have ranged throughout the world. In a highly competitive and fast-moving marketplace, they have sought to maximize their strengths and flexibility as products change rapidly and prices continue to fall. However, the sprawl and complexity of such networks have made it harder to manage end-to-end operations smoothly.

Supply chain management is one of the major challenges in the operations of a company. A well-planned supply chain will increase companies' profits significantly; for example, Procter & Gamble has integrated business plans between manufacturers and suppliers, which saved retailers USD 65 million in an 18-month supply chain [19]. In supply chain management, the value of information cannot be ignored. The increase in demand variability as one moves upstream in the supply chain causes inefficient use of resources and huge supply chain costs. These costs are due to excessive inventory cost and backorder cost because of difficulties in acquiring the appropriate customer information. Researchers have been studying the importance of information in supply chains and the bullwhip effect. The MIT Beer Distribution Game is one of the approaches that can explain and aid in understanding the problems.

1.1 Bullwhip Effect

An important observation in supply chain management, known as the *bullwhip effect*, suggests that demand variability increases as one moves up a supply chain. The more stages there are from the company to the end customer, the larger the variation is. Consider a supply chain consisting of supplier, manufacturer, wholesaler, and retailer; the goods and information flows of this supply chain can be shown in Figure 1-1. In such a typical supply chain, suppliers provide raw materials to manufacturers, who process the raw materials into finished goods and then provide the finished goods to wholesalers who combine products from a number of manufacturers for sale to retailers. The retailers then sell the products to the customers. In addition to the physical goods flow, there is another flow of demand information that proceeds upstream. The retailers have direct contact with the ultimate customers who are at the end of the supply chain. The wholesalers predict the demand from the orders received from the retailer since they have no access to the customers, and so on for the upstream entities of the supply chain.

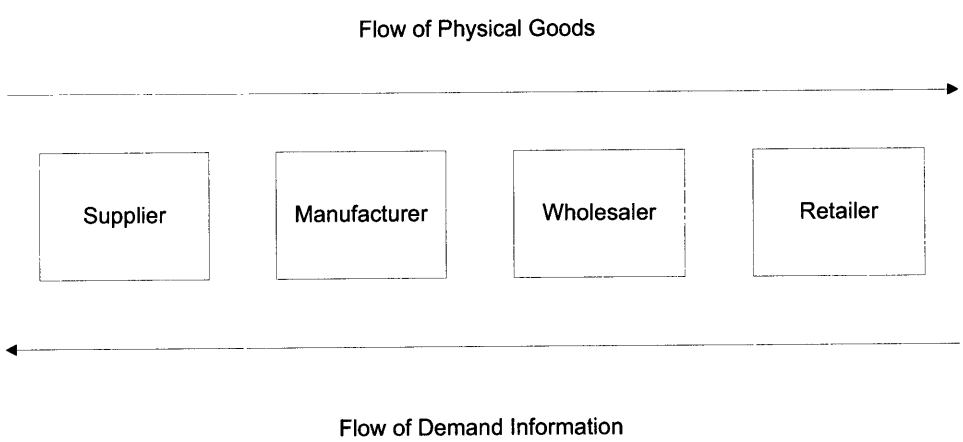

Figure 1-1: Goods and Information Flows in a Supply Chain

To understand the impact of the increase in variability in a supply chain, we consider the supply chain shown in Figure 1-1. Take the wholesaler as an example: the wholesaler here does not have direct access to customers' orders. To determine the order quantity, the wholesaler can only forecast the demand from the orders received from the retailer. However, the orders placed by the retailer already have a higher variability than the orders from the ultimate customer. Therefore, the wholesaler is forced to carry more safety stock than the retailer or else to maintain higher capacity than the retailer in order to meet the same service level as the retailer. This analysis can be carried over to the factory and the supplier, which results in even higher capacity and cost in the upstream.

The bullwhip effect leads to inefficiencies in supply chains, since it increases the cost for logistics and makes the companies less competitive in the market. In particular, the phenomenon affects a supply chain in the following three respects [17]:

1. **Capacities**: The variation of the customer demand increases when it moves up the supply chain. How much inventory to keep becomes a dilemma for companies. It is hard to meet the demand when a huge order comes, yet maintaining the maximum capacity leads to poor use of inventory.

2. **Inventory Level**: The variation of the customer demand leads to variation of inventory level. The farther away the company is from the end customer, the higher the variation of inventory level it faces. A high level of inventory costs a lot and makes the company less profitable, while a low level of inventory leaves the company unreliable.

3. **Safety Stock**: Higher variation of demand requires much more safety stock to maintain a certain service level. Thus, the stronger the bullwhip effect is in a supply chain, the more safety stock should be kept to meet the same service level.

Thus, it is important to identify techniques and tools that will allow us to control the effect – in other words, to control the variability in the supply chain. For this

purpose, we need to understand the main factors contributing to the increase in variability in the supply chain [19].

1. **Demand forecasting**. If the demand forecast of a company is based on orders from its downstream instead of the end customer, the variation of demand is amplified up the supply chain. The amplified variation leads to a higher level of safety stock, which will lead to an inaccurate forecast compared to the end demand.

2. **Lead Time**. With longer lead times, a small change in the estimate of demand variability implies a significant change in safety stock. Thus, it is easy to see that the increase in variability is magnified with increasing lead time.

3. **Batch ordering**. In order to reduce set-up costs and order-fixed costs, retailers tend to order in batches. This leads to distorted information for the wholesalers, who see a distorted and highly variable pattern of orders.

4. **Price fluctuation**. Wholesalers vary the price of their products and sometimes offer a lower price temporarily to end customers and retailers for marketing purposes. If the prices fluctuate, the retailers often attempt to stock up when prices are lower. This action increases the variation of the demand.

5. **Inflated orders**. When a huge order comes, retailers tend to order more than actual demand due to the shortage they see in the inventory. As soon as the huge order gets fulfilled, the retailers will face excessive inventory caused by the inflated orders.

1.2 MIT Beer Distribution Game

The Beer Distribution Game ("beer game"), developed at the Sloan School of Management in the early 1960s [10], is a popular simulation used in the classroom to help students understand a number of key principles of supply chain management

[19]. The game was originally played with a board and cards, and it was later developed as a web-based game [15], to demonstrate the value of integrated supply chain management. The game is played by teams of at least four players, often in heated competition, and takes from one to one-and-a-half hours to complete. In this game, participants (students, managers, analysts, and so on) play the roles of retailer, wholesaler, distributor, and factory of a popular brand of beer.

The purpose of the game is to meet customer demand for cases of beer, through a multi-stage supply chain with minimal expenditure on back orders and inventory. The backorder cost is usually twice as much as the inventory cost. No communication is allowed during the game. Each player only has access to the order information of its downstream who places orders to him/her. Most of the players feel frustrated because they are not getting the results they want. Players wonder whether someone in their team did not understand the game, or assume customer demand is following a very erratic pattern as backlogs mount and/or massive inventories accumulate. However, in this game these feelings are common, and the huge cost of back orders and inventory in this multi-stage supply chain is caused by the bullwhip effect.

The game's rules are as follows. Each player starts the game with an initial inventory of beer, say, 12 cases. There are usually eight cases of beer in the pipeline for each player for a two-week supply (four cases for each week). Every player places orders on a weekly basis and there is no communication among the players throughout the game. At the beginning of each week during the game, the retailer receives the current week's order from the customer. Due to the one-week delay of information, the wholesaler, the distributor, and the factory receive the order placed by their downstream in the previous week, shown in Figure 1-2.

After receiving the orders, the players fulfill the orders with their available inventory, and the beer shipped out is delivered in two weeks due to the shipping delay. All the players then place an order in the stack, and the order will reach the upstream in

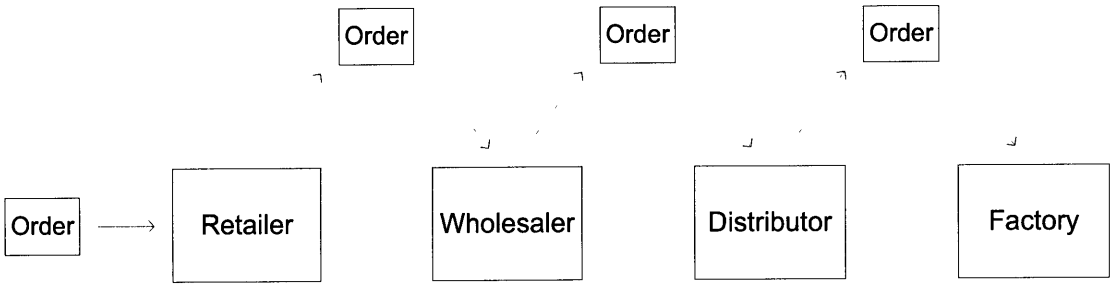

Figure 1-2: Order Flow Demonstration in the Beer Distribution Game with One-week Information Delay

the next week. For example, a retailer places an order in Week Two which is received by the wholesaler in Week Three; if the wholesaler has sufficient beer to satisfy the retailer's order, the retailer will receive the beer in Week Five; otherwise, the retailer must wait until the wholesaler receives additional beer from the distributor. The factory faces a constant delay in acquiring inventory but no capacity limit for production.

The object of the game is to minimize the total cost of all the players throughout the game. The cost is cumulatively calculated on a weekly basis with \$1 per case inventory cost and \$2 per case backorder cost. Players are not informed of the demand before the game starts, and the total inventory and backorder costs are usually very high when played by human players without communication, such as students and supply chain managers. Therefore, a manageable order policy is needed to help each part of the supply chain to reduce its cost and improve its performance in the game.

1.3 Demand Forecast

Given the importance of the demand forecast, a practical way of calculating it is needed to help players in the Beer Distribution Game to achieve better performance.

The simplest way of calculating demand forecast is to calculate the average of all the previous received orders x_i:

$$s_t = \frac{1}{n}\sum_{i=1}^{n} x_i.$$

The sales of some products, such as beverages and jerseys, are highly dependent on the seasons and special periods. For these products (take football jerseys as an example), if we calculate the forecast by simply taking the average of past sales, there will be a huge shortage during the sport's season, and excessive inventory during the rest of the year. Therefore, we introduce a moving average into the forecast. In the simple moving average, the past observations are weighted equally [1]:

$$s_t = \frac{1}{k}\sum_{n=0}^{k-1} x_{t-n} = \frac{x_t + x_{t-1} + x_{t-2} + ... + x_{t-k+1}}{k} = s_{t-1} + \frac{x_t - x_{t-k}}{k}.$$

By introducing weight to each of the previous orders, the weighted moving average is:

$$\{w_1, w_2, ..., w_k\} \text{ such that } \sum_{n=1}^{k} w_n = 1,$$

$$s_t = \sum_{n=1}^{k} w_n x_{t+1-n} = w_1 x_t + w_2 x_{t-1} + w_3 x_{t-2} + ... + w_k x_{t-k+1}.$$

In statistics, exponential smoothing is a technique that can be applied to time series data, either to produce smoothed data for presentation or to make forecasts. The exponential smoothing assigns exponentially decreasing weights as the observations get older. The simple exponential smoothing (exponentially weighted moving average) formulas are expressed as the following [2]:

$$s_0 = x_0,$$

$$s_t = \alpha x_t + (1-\alpha)s_{t-1} = s_{t-1} + \alpha(x_t - s_{t-1}).$$

Thus, the current smoothed value is an interpolation between the previous smoothed value and the current observation, where α controls the closeness of the interpolated value to the most recent observation. Notice that if $\alpha = 1$, the model is equivalent to a random walk model (without growth). If $\alpha = 0$, the model is equivalent to the mean model, which assumes the first smoothed value is set to the mean.

1.4 Literature Review

1.4.1 Studies of the Bullwhip Effect

As we have seen, the bullwhip effect refers to the tendency of orders to increase in variation as one moves up a supply chain. Distorted information from one end of a supply chain to the other can lead to tremendous inefficiencies: excessive inventory and unnecessary waste of inventory investment, poor customer service and loss of revenues, misguided capacity plans and inefficient production schedules, etc. Forrester was the first to point out this effect and its possible causes [9]. In the past few years, industry managers and researchers have focused attention on the operational causes of the bullwhip effect.

Lee et al. analyzed the bullwhip effect in a supply chain [13]. They identified four major causes of the bullwhip effect: demand forecast updating, order batching, price fluctuation, and rationing and shortage gaming. They devised a set of simple models to illustrate how each of these factors can lead to the amplification of order variance as one moves up a production/distribution supply chain [14]. They also suggested a few approaches to counteract the bullwhip effect, such as 1) avoiding multiple demand forecast updating, 2) breaking up order batches and devising strategies that lead to smaller batches or more frequent resupply, 3) stabilizing prices, and 4) eliminating gaming in shortage situations – that is, when a supplier faces a shortage, instead of allocating products based on orders, it can allocate in proportion to past sales records.

Metters established an empirical lower bound on the profitability impact of the bullwhip effect [16]. He quantified the impact of the effect on the profitability of a company and studied the problem with a linear programming approach. He reported that elimination of the bullwhip effect can result in a 5% profitability increase using managerially relevant parameter settings and that this saving can be even higher in a capacity-limited supply chain.

Baganha and Cohen presented a hierarchical model framework for the analysis of the stabilizing effect of inventories in multi-echelon manufacturing/distribution supply chains [3]. In their model, they have N retailers. All of the retailers are supplied by one distribution center which reviews its inventory position at the beginning of each period, shown in Figure 1-3. They solved the overall optimization problem by a multi-echelon decomposition procedure assuming that each location only has local information. That is, each location will determine its optimal policy only after all those locations it supplies have determined their inventory policies and their associated ordering policies.

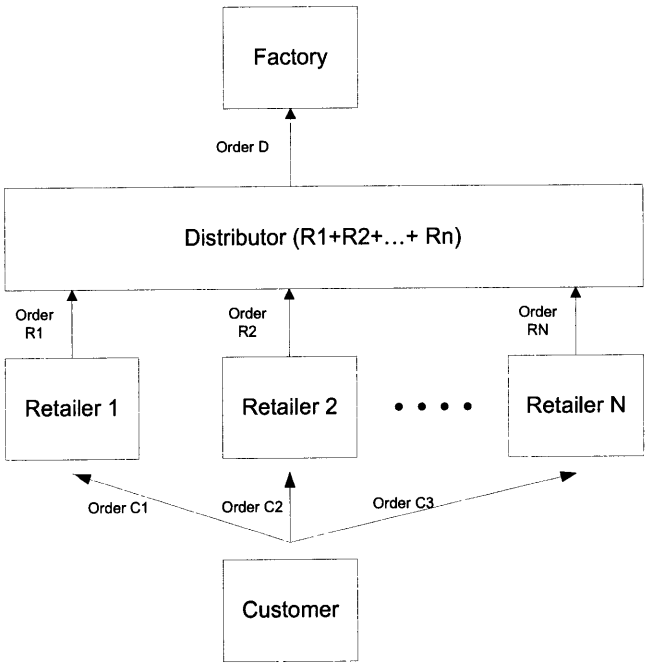

Figure 1-3: Multi-retailer Structure in a Supply Chain

Chen et al. have introduced ways to alleviate these operational problems by introducing quantitative models to improve demand forecasting [5]. They considered a simple two-echelon supply chain comprising a single manufacturer and a single retailer and quantified the impact of exponential smoothing forecasts on order-rate fluctuations within order-up-to policies [6]. The authors provided important managerial insights in their work, such as the fact that the bullwhip effect is caused by the need to forecast, and that the smoother the demand forecasts, the smaller the

bullwhip effect.

1.4.2 Study of the Beer Distribution Game

The Beer Distribution Game has been a useful approach for the study of the bullwhip effect. Sterman was the first to use the Beer Distribution Game to rigorously test the existence of the bullwhip effect in an experimental context [20]. Croson and Donohue have conducted two experiments on different sets of participants. They have proved that the bullwhip effect still exists even when the normal operational causes (e.g., batching, price fluctuations, demand estimation, etc.) are removed [8].

The beer game is also considered an important tool for teaching supply chain management. Kaminsky and Simchi-Levi have used the Beer Distribution Game as a tool for teaching the value of integrated supply chain management [11]. Chen and Samroengraja later developed a variant of the popular Beer Distribution Game called the *stationary beer game* for classroom use [7]. This new game models the material and information flows in a production/distribution channel serving a stationary market, and the players in this new game all know the demand distribution while they manage the different stages of the supply chain.

A number of researchers have discussed the order policy of the beer game under the Genetic Algorithm. O'Donnell et al. proposed the Genetic Algorithm to reduce the bullwhip effect of sales promotions in supply chains [18]. In 2006, Kimbrough et al. proposed an order policy optimization based on the Genetic Algorithm [12]. Later, Chaharsooghi et al. proposed a reinforcement learning model to suggest an order policy for players on a weekly basis [4]. However, some of the methods are based on the condition that the customer's demand is known, and they both simulate the optimized order policy under the customer's demand. In this paper, we suggest an easily implemented robust order policy that can work with both known and unknown customer demand by using numerical simulations.

1.5 Thesis Objectives and Organization

This thesis is organized as follows. In the next chapter, I will introduce the Q-model developed by my supervisor and me to suggest an easily learned order policy for each echelon in the beer supply chain. In Chapter 3, I will analyze the simulation approach developed by other researchers and also introduce my own simulation. In Chapter 4, a number of experiments involving deterministic and stochastic demand and lead times will be examined and the results of some other recommended order policies will be compared with the Q-model order policy. I will apply the Q-model to more general scenarios and analyze its performance.

Chapter 2

Modeling the Beer Distribution Game Order Policy

Chapter 1 introduced the bullwhip effect in supply chain management, the MIT Beer Distribution Game, and studies related to them. This chapter proposes the Q-model to suggest a practical order policy for each of the players in the Beer Distribution Game to reduce the bullwhip effect and the total logistic cost in the beer supply chain. The model will be simulated in Chapter 4, and the results will be compared to the other prevailing models which are based on the Genetic Algorithm.

2.1 Assumptions

The Beer Distribution Game mimics the mechanics of a decentralized periodic review inventory system with four serial echelons. It has been played for teaching and research purposes for years. However, different scholars simulate the game in different ways. For instance, the initial inventory, the delay period of information, and the delivery lead time vary in different cases and simulations. Therefore, before introducing the Q-model, certain necessary assumptions and details in the model proposed in this chapter should be stated:

1. **Initial Inventory**

   All the players start with 12 cases of beer in their inventory at the beginning

of the game. There are eight cases of beer in the pipeline, of which four cases will be delivered in Week One and four cases will be delivered in Week Two to every player.

2. **Game Duration**

   The duration of the game is 35 weeks; the game is played on a weekly basis by four players, who are the retailer, the wholesaler, the distributor, and the factory. Each week, every player receives and places one order.

3. **Order Placement**

   Each week, the retailer receives an order from customers, fulfills the order, and places an order to the wholesaler. The wholesaler, the distributor, and the factory each know only the orders of their own downstream. There is a one-week delay for passing along orders (shown in Figure 2.1). They fulfill the received orders and place orders to their upstream.

4. **Lead Time**

   The delivery delay is either a fixed one or two weeks or a random time between zero and four weeks. Take a fixed two-week delivery lead time as an example. If the retailer places an order in the current week, and the wholesaler has sufficient beer to fulfill the retailer's order, the retailer will receive the beer in three weeks (one-week delay for the order, two-week delay for the delivery); otherwise, the retailer must wait until the wholesaler receives additional beer from the distributor, which is the upstream of the wholesaler. The factory, which is the producer, faces a three-week delay in acquiring inventory but no capacity limit for production. The supply chain of the Beer Distribution Game is shown in Figure 2-1.

5. **Cost**

   The inventory cost is $1 per case every week; the backorder cost is $2 per case every week. The objective is to minimize the total cumulative cost of all the players throughout the game.

6. **Process**

   In every period of the game the following occurs: (1) each player receives orders placed by downstream in the previous week (the retailer receives the current week's order from the ultimate customer at the end of the supply chain); (2) each player receives delivery from its upstream; (3) the received order is fulfilled from the available inventory; and (4) each player places an order to its upstream.

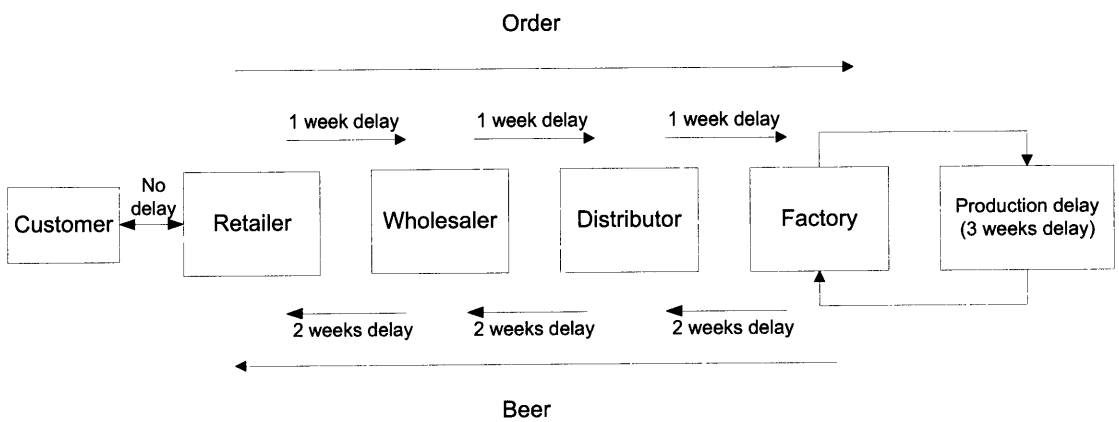

Figure 2-1: Beer Game Supply Chain with Constant Lead Time

2.2 Heuristics

This section presents the Q-model to suggest an order policy for each of the players based on expected demand and previous orders.

2.2.1 Notation

Here, we introduce some necessary notation:

- R, W, D, F, X: Retailer, Wholesaler, Distributor, Factory, $X = \{R, W, D, F\}$;
- $D_X(t)$: The order received in period t by X;
- $O_X(t)$: The order placed in period t by X;

- $\hat{D}_X(t)$: The demand forecast in period t of X;
- $I_X(t)$: Inventory of X in period t, initial inventory ($t = 0$);
- $\hat{I}_X(t)$: The expected inventory at the beginning of period $t+1$ of X;
- $RB_X(t)$: The beer received in period t by X;
- $BB_X(t)$: The back order of beer of X in period t;
- Q_X: Ratio of expected inventory to expected demand of X in all periods;
- C_X: Total cost of all the periods of X.

2.2.2 Q-model Description

In the heuristics of the Q-model, we introduce four factors $Q_X(X \in \{R, W, D, F\})$ to each of the players. The factor Q_X indicates how many weeks' inventory the players should maintain. Every week, all the players update their demand forecasts based on the orders they have received, and multiply the estimated demand by Q_X to determine how many weeks' inventory they should keep. Then they place orders, based on the expected inventory and on the orders they have placed, according to the following relation:

In Week T, for the retailer,

$$\begin{aligned}\textit{Total beer ordered by } R \textit{ in Week } 1 \textit{ to } T &+ \textit{Initial inventory and pipeline supply of } R \\ &- \textit{Total demand from Week } 1 \textit{ to } T \\ &= Q_R \times \textit{Expected demand.}\end{aligned}$$

In Week T, for the wholesaler, distributor, and factory,

$$\begin{aligned}\textit{Total beer ordered by } X \textit{ in Week } 1 \textit{ to } T &+ \textit{Initial inventory and pipeline supply of } X \\ &- \textit{Initial pipeline supply} - \textit{Total demand from Week } 2 \textit{ to } T \\ &= Q_X \times \textit{Expected demand.}\end{aligned}$$

The above heuristic ensures that all the players in the supply chain always order up to Q_X weeks of inventory ahead throughout the game. Each week, their order quantity is based on the updated forecast, their initial inventory, and their previous orders. The advantage of this heuristic is that it prevents the players from continuously placing huge orders when the demand increases or the back order increases. Therefore, it is very important to make accurate demand forecasts to ensure the players make appropriate ordering decisions in the game without mutual communication.

2.2.3 Expected Demand

In Chapter 1, we introduced a few approaches to calculate the demand forecast using moving average and exponential smoothing. For different scenarios, different methods should be used to determine the demand forecast. In order to satisfy more general scenarios in the Beer Distribution Game, such as seasonal demand, simply calculating the average of all the orders received is a poor approach. With the exponential smoothing, it is hard to determine the weight for the last order when the demand has a huge fluctuation. Therefore, the moving average introduced in Chapter 1 is more suitable, and we have the equations of the moving average for the demand forecast in the beer game as follows:

$$\hat{D}_X(t) = \begin{cases} \dfrac{\sum_{i=t-p+1}^{t} D_R(i)}{p}, & X = R, t > 1+p \\ \dfrac{\sum_{i=t-p}^{t-1} D_{X-1}(i)}{p}, & X = W, D, F, t > 1+p. \end{cases} \tag{2.1}$$

In the above equation, the demand forecast of the retailer for the first p periods is the average of all the t orders received. For W, D, and F, the demand forecast of the first p periods is derived from the average of the orders received and the initial inventory of their downstream. The forecast of the $(p+1)$th period and after is the average of the previous p orders received. The reason for considering the downstream's initial

inventory in the first p periods is that simply calculating the average of the t orders received will lead to inaccurate forecasts for the first few weeks. Every player makes a forecast demand in the first few weeks to determine how much inventory should be kept. Each orders the amount of beer using the total amount of beer needed in inventory, including safety stock, minus the initial inventory. Therefore, the expected demand should consider both the previous orders received and the initial inventory of the downstream.

Meanwhile, at the beginning of the game, each of the players places huge orders to acquire safety stock according to their demand forecast. If the upstream considers those huge orders simply as orders based on weekly demand, it will lead to distorted demand forecasts for the first few weeks. Therefore, we have to consider the factor Q_X in the demand forecast to make the demand more accurate at every stage of the supply chain.

By considering the initial inventory and the factor Q_X in the demand forecast, we derive the demand forecast for the first p periods as follows. For the retailer:

$$\hat{D}_R(t) = \frac{\sum_{i=1}^{t} D_R(i)}{t},$$

and for the wholesaler:

$$\hat{D}_W(t) = \frac{\sum_{i=1}^{t-1} O_R(i) + I_R(0)}{t-1+Q_R}.$$

We could calculate the first p weeks' demand forecast for the distributor and the factory in the same way as for the wholesaler. However, the performance of the forecast is not as good as expected. This is because the order decision of the wholesaler has already been affected by the initial inventory of the retailer and the factor Q_R, so the distributor and the factory can only make better demand forecasts by considering these two factors. Therefore, we have all the demand forecasts of all

the players as follows:

for the retailer:

$$\hat{D}_R(t) = \begin{cases} \dfrac{\sum_{i=1}^{t} D_R(i)}{t}, & t \leq p \\ \dfrac{\sum_{i=t-p+1}^{t} D_R(i)}{p}, & t \geq 1+p \end{cases} \tag{2.2}$$

for the wholesaler:

$$\hat{D}_W(t) = \begin{cases} 4, & t = 1 \\ \dfrac{\sum_{i=1}^{t-1} O_R(i) + I_R(0)}{t-1+Q_R}, & t = 2, ..., 1+p \\ \dfrac{\sum_{i=t-p}^{t-1} O_R(i)}{p}, & t > 1+p \end{cases} \tag{2.3}$$

for the distributor:

$$\hat{D}_D(t) = \begin{cases} 4, & t = 1 \\ \dfrac{\sum_{i=1}^{t-1} O_W(i) + I_W(0) + I_R(0)}{t-1+Q_R+Q_W}, & t = 2, ..., 1+p \\ \dfrac{\sum_{i=t-p}^{t-1} O_W(i)}{p}, & t > 1+p \end{cases} \tag{2.4}$$

and for the factory:

$$\hat{D}_F(t) = \begin{cases} 4, & t = 1 \\ \dfrac{\sum_{i=1}^{t-1} O_D(i) + I_D(0) + I_W(0) + I_R(0)}{t-1+Q_R+Q_W+Q_D}, & t = 2, ..., 1+p \\ \dfrac{\sum_{i=t-p}^{t-1} O_D(i)}{p}, & t > 1+p. \end{cases} \tag{2.5}$$

The first period's demand forecast for the wholesaler, the distributor, and the

factory is four cases of beer. The reason for this is that in the first week of the game, none of the three players receives an order from its downstream due to the one-week information delay. Meanwhile, each of the above three players has to ship out four cases of beer during the first two weeks of the game for the pipeline supply. Therefore, the demand forecast for the first week is four cases.

2.2.4 Expected Inventory

Each player calculates the expected inventory by multiplying expected demand by Q_X to store Q_X weeks of expected future demand as follows:

$$\hat{I}_X(t) = Q_X \cdot \hat{D}_X(t), \tag{2.6}$$

$$X = \{R, W, D, F\}.$$

The Q_X set is determined by a computer simulation that will be discussed in the next chapter. At the beginning of the game, each player is assigned a unique Q to determine how many weeks' inventory they should keep. Our objective is to find an appropriate set of Q_X to minimize the total cost for all the players in the beer supply chain and make the Q_X robust in different scenarios.

2.2.5 Order Quantity

The previous sections introduced how to determine the demand forecast and the expected inventory. The last step in the Q-model is to decide how much to order for each player. We have learned that it is more reasonable to make decisions based on the previous orders and demand instead of on each player's inventory level. In this way, the players can avoid ordering too much. The order quantities of each player are as follows:

In Week T, for the retailer:

$$O_R(t) = \begin{cases} max\{0, D_R(1) - I_R(0) - 4 + \hat{I}_R(t)\}, & t = 1 \\ max\{0, \sum_{i=1}^{t} D_R(i) - I_R(0) - 8 - \sum_{i=1}^{t-1} O_R(i) + \hat{I}_R(t)\}, & t = 2, ..., T. \end{cases} \tag{2.7}$$

In Week T, for the wholesaler, the distributor, and the factory:

$$O_X(t) = \begin{cases} D_X(1) - I_X(0) - 4 + \hat{I}_X(1), & t = 1 \\ max\{0, \sum_{i=1}^{t-1} O_{X-1}(i) + 4 - I_X(0) - 8 - \sum_{i=1}^{t-1} O_X(i) + \hat{I}_X(t)\}, & t = 2, ..., T. \end{cases} \tag{2.8}$$

As we have explained, by using this manageable Q-model order policy, players can avoid amassing excessive inventory. The next chapter introduces the simulation process for evaluating the models.

Chapter 3

Simulating the Beer Distribution Game

Chapter 2 introduced a manageable order policy for each of the players in the Beer Distribution Game. This chapter introduces the simulation process of the game in order to evaluate the performance of different order policies and compare the results.

3.1 Simulation Process

In the MIT Beer Distribution Game, each player (except the retailer) has a one-week delay in receiving its downstream's orders as an information delay. Meanwhile, there is a lead time for delivery, which is a constant, caused by the shipping and production delay. In the web-based MIT Beer Distribution Game, players are able to choose to play the game with either a one-week or two-week delay of the delivery. However, due to the uncertainty of delivery time and production time, we consider the delay as both constant and stochastic (randomly distributed between 0 and 4) in the simulation. The process of the stochastic lead time is explained as follows.

Each week, there is a delivery lead time $t_i \in \{0, 1, 2, 3, 4\}$ assigned to all the players. The order placed by each player during the week will be ready in inventory in $1 + t_i$ weeks if the upstream has enough inventory on hand. If the upstream does not

have enough products to ship out, it will generate a back order and the back order will be fulfilled along with a later order once the upstream replenishes its inventory. Figure 3-1 details the Beer Distribution Game supply chain.

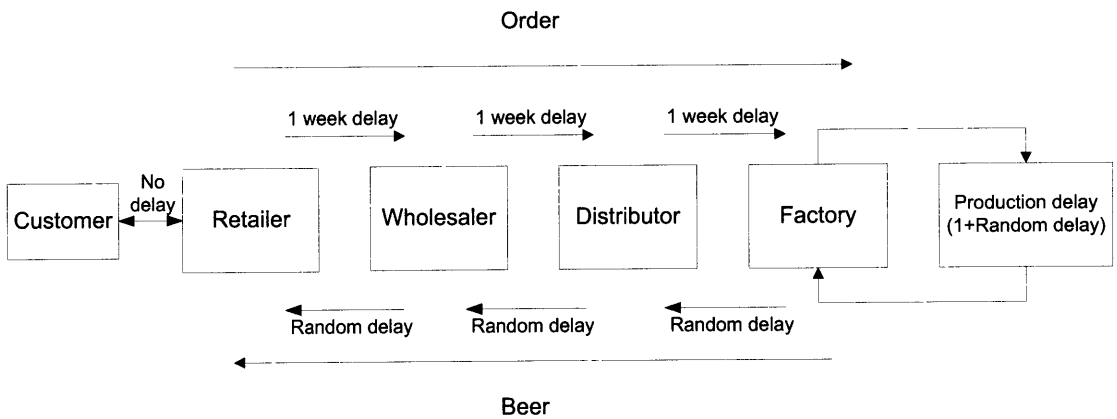

Figure 3-1: Beer Game Supply Chain with Stochastic Lead Time

Given $t_1 = 2, t_2 = 0$, let us take the retailer and the wholesaler as an example. In Week One, the retailer orders 17 cases of beer from the wholesaler, which is the upstream of the retailer. The order reaches the wholesaler in Week Two due to the one-week delay of information. However, the wholesaler has only 16 cases of beer in Week Two, so 16 cases of beer are shipped out from the wholesaler, and 1 case of beer back order is generated due to the shortage of inventory. These 16 cases of beer are delivered to the retailer in two weeks and are ready for inventory in Week Four. In Week Three, the wholesaler receives another order of 13 cases of beer that the retailer placed in Week Two, while it also receives 18 cases of beer from the distributor. Therefore, in Week Three, the wholesaler ships out 14 cases of beer, which fulfills the order it received in Week Three and the back order generated in Week Two. These 14 cases of beer are delivered to the retailer within Week Three since the delivery lead time of the order in Week Two is 0 weeks.

3.2 Current Simulation

Many researchers have studied the Beer Distribution Game and have performed simulations to evaluate their order policy models [4][12][18]. However, some of the current available simulations ignore the delivery of the back-ordered beer. Therefore, it is difficult to evaluate the performance of the order policy models using those simulations.

Table 3.1 fully details the first 16 weeks' order and inventory information from the order policy suggested by Chaharsooghi et al. [4]. The inventory data is from their simulation results in a scenario of stochastic demand and lead time under the Reinforcement Learning Ordering Mechanism introduced in their article [4]. I have retrieved part of the data, shown in Table 3.1. In the table, the second and third columns are the stochastic demand and lead time, respectively.

Table 3.1: Simulation Results of the Retailer and the Wholesaler Using Simulation Approach and RLOM Order Policy Proposed by Chaharsooghi et al.

Week	**D**	**LT**	**Retailer**					**Wholesaler**				
			Inv.	**New**	**Out**	**Order**	**Cost**	**Inv.**	**New**	**Out**	**Order**	**Cost**
0			12					12				
1	15		1	4	15	17	1	**12**	4	4	6	12
2	10	2	-5	4	10	**13**	10	**-1**	**4**	**17**	18	2
3	8	0	**0**	**13**	8	11	0	4	18	**13**	13	4
4	14	2	2	**16**	14	15	2	-1	6	11	11	2
5	9	4	-7	0	9	10	14	-16	0	15	15	32
6	3	4	0	10	3	5	0	-19	7	10	11	38
7	13	4	-13	0	13	13	26	-24	0	5	6	48
8	2	0	3	18	2	4	3	-19	18	13	14	38
9	13	2	-10	0	13	16	20	-23	0	4	5	46
10	11	4	-14	7	11	13	28	-24	15	16	17	48
11	3	1	-17	0	3	5	34	-37	0	13	14	74
12	4	1	-21	0	4	7	42	-20	22	5	8	40
13	6	0	2	29	6	9	2	6	33	7	10	6
14	11	0	15	24	11	12	15	16	19	9	10	16
15	15	1	0	0	15	15	0	4	0	12	12	4
16	12	1	0	12	12	12	0	-1	10	15	15	2

Source: Table 3 in [4]

In the above table, some of the data listed appear to be inaccurate. For instance, in Week Two, the wholesaler ships out 17 cases of beer even though there are only 16 cases (Inv.(12) + New(4) in Week One) on hand. According to the received beer data from the retailer, the retailer receives 16 cases of beer in Week Four, and 13 cases in Week Three (see the bold data in the table). However, for these two orders the retailer actually has ordered $(13 + 17)$ cases of beer; but only $(16 + 13)$ cases of beer are actually shipped out and delivered to the retailer, which means the back-ordered beer has never been delivered, even though the data of the beer shipped out by the wholesaler shows all the beer ordered by the retailer has been shipped out.

In addition, we can also examine the delivery of beer by setting the last five weeks' orders from the retailer to be zero, so that all the beer ordered by the retailer could be delivered from the wholesaler before the game ends. I have found the total beer shipped out from the wholesaler to be 336 cases, and the beer received by the retailer to be only 289 cases: there are 47 cases of beer missing in the simulation. The reason for this is because the original simulation assumes that the upstream will ship out the back-ordered beer while actually the back order has never been fulfilled during the simulation. Therefore, the original simulation should be revised and a more appropriate simulation should be developed to evaluate the models.

3.3 Revised Simulation

I have developed a new simulation and recalculated the inventory level of each player in the Beer Distribution Game. The new simulation keeps track of the back-ordered beer as negative inventory. In the simulation, each week every player ships out the beer according to their received orders and their previous inventory level. For example, if the player's inventory of the previous week is -3, then there is a 3-case back order. During the current week, the player receives another order of 10 cases of beer and a delivery of 11 cases of beer from its upstream; the player will then ship out 11 cases of beer, which will satisfy the current order and part of the back order, to its

downstream and have -2 as the current inventory level, which means the player still has a 2-case back order to be fulfilled. The flowchart of the simulation is shown in Figure 3-2.

The above flowchart explains how the order goes from downstream to upstream, and how the simulation determines the inventory and the delivered beer quantity. In the simulation, at the beginning of every week, each player updates the player's inventory with the received beer from the upstream $RB(t)$, received order from the downstream $O(t-1)$, which has been delayed for one week, and the player's own previous week's inventory status $I(t-1)$. The simulation will then determine the quantity of the delivered beer $DB(t)$ (shown in Figure 3-2), and then update the downstream's inventory by the delivered beer and delay time as follows:

$$RB_{downstream}(t + t_{delay}(t-1)) = RB_{downstream}(t + t_{delay}(t-1)) + DB(t).$$

With the simulation introduced above, I have re-simulated the scenario simulated in Table 3.1, and Table 3.2 shows the results.

In Table 3.2, the beer shipped out from upstream equals the beer received by the downstream. Meanwhile, the beer shipped out is related to the inventory on hand, the order received in the current period, and the back order from the previous week. The beer received by all the players is based on the delivery lead time and the quantity shipped out by their upstream.

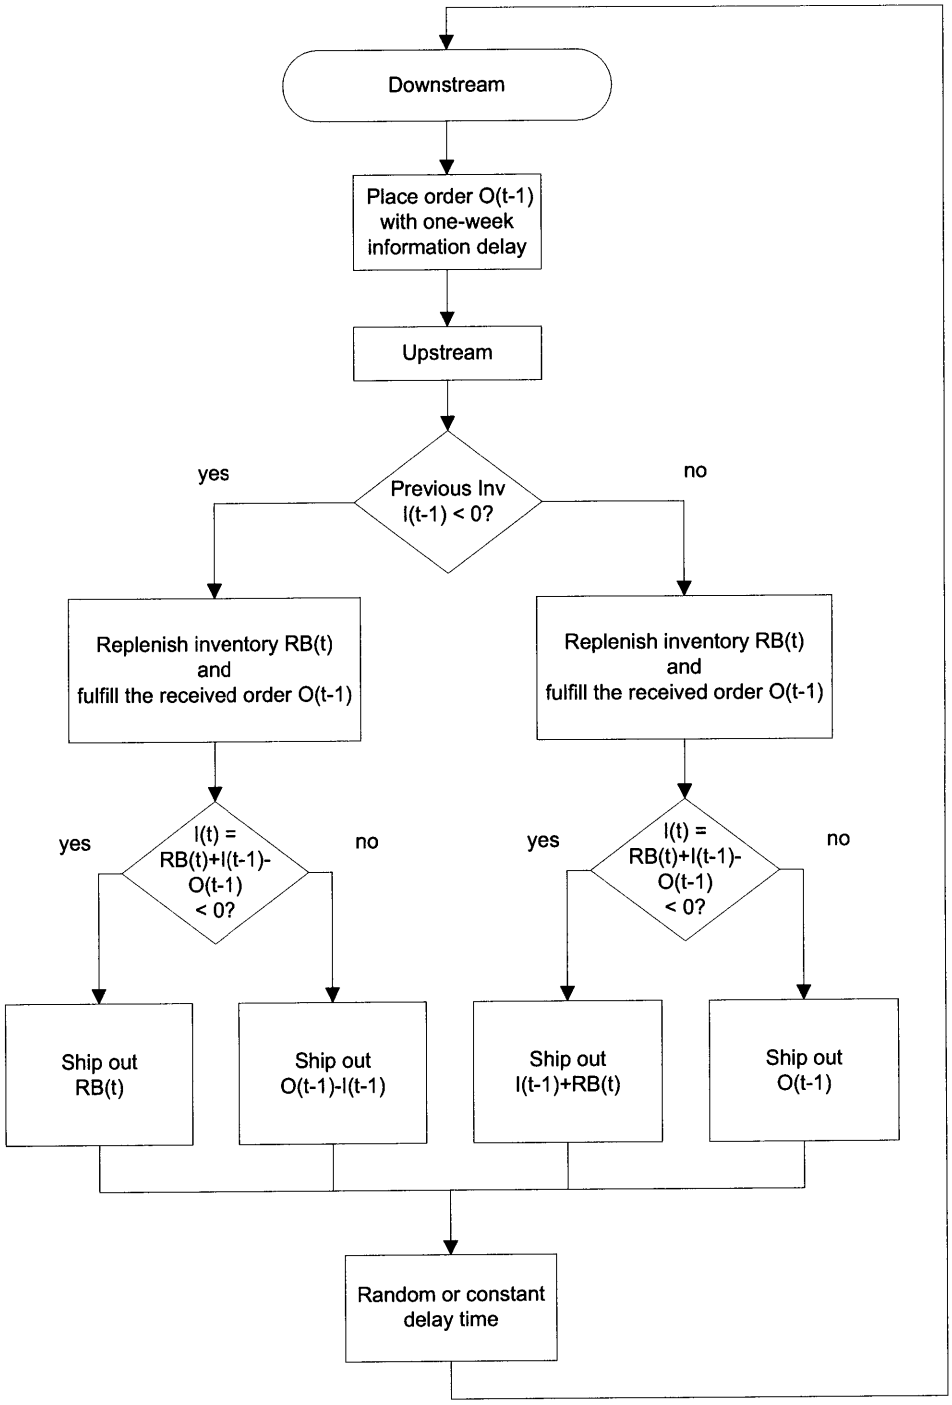

Figure 3-2: Beer Game Supply Chain Simulation Order and Delivery Flowchart

Table 3.2: Revised Simulation Data of the Retailer and the Wholesaler under Stochastic Lead Time in Comparison to Table 3.1

Week	D	LT	Retailer					Wholesaler				
			Inv.	New	Out	Order	Cost	Inv.	New	Out	Order	Cost
0			12					12				
1	15		1	4	15	17	1	12	4	4	6	12
2	10	2	-5	4	5	13	10	-1	4	16	18	2
3	8	0	1	14	13	11	1	4	18	14	13	4
4	14	2	3	16	14	15	3	-1	6	10	11	2
5	9	4	-6	0	3	10	12	-16	0	0	15	32
6	3	4	1	10	9	5	1	-19	7	7	11	38
7	13	4	-12	0	1	13	24	-24	0	0	6	48
8	2	0	4	18	14	4	4	-19	18	18	14	38
9	13	2	-9	0	4	16	18	-23	0	0	5	46
10	11	4	-13	7	7	13	26	-24	15	15	17	48
11	3	1	-16	0	0	5	32	-37	0	0	14	74
12	4	1	-20	0	0	7	40	-20	22	22	8	40
13	6	0	23	49	26	9	23	16	43	27	10	16
14	11	0	36	24	11	12	36	26	19	9	10	26
15	15	1	21	0	15	15	21	14	0	12	12	24
16	12	1	21	12	12	12	21	9	10	15	15	9

Chapter 4

Simulation Results and Analysis

Chapter 2 introduced the Q-model and presented the formulas for the order policy under the Q-model. Chapter 3 analyzed the simulations that appeared in other available sources and also proposed a more appropriate simulation process to evaluate the performance of the models. This chapter simulates different scenarios and examines the performance of the Q-model by comparing it to the other models.

4.1 Simulation Results

This section simulates the total costs of the game using the order policies recommended by the Q-model and the Genetic Algorithm (GA) in various scenarios. The simulations consider both demand and lead time in constant and stochastic scenarios. The testing scenarios adopted are as follows.

4.1.1 Testing Scenarios

I have tested scenarios with various demands and both constant and stochastic lead times. Table 4.1 shows the constant and stochastic demands and lead times tested in the simulations. In the table, the constant demand is given by the MIT Beer Distribution Game; the stochastic demand is randomly distributed between 0 and 15, and the stochastic lead time is randomly distributed between 0 and 4. I have adopted

the stochastic demand and lead time scenario data from the papers using the Genetic Algorithm [12][4] for ease of comparison.

Table 4.1: Basic Testing Scenarios of Constant and Stochastic Demand and Lead Time

Type	**Value**
Stochastic demand	15, 10, 8, 14, 9, 3, 13, 2, 13, 11, 3, 4, 6, 11, 15, 12, 15, 4, 12, 3, 13, 10, 15, 15, 3, 11, 1, 13, 10, 10, 0, 0, 8, 0, 14
4-8 demand	4 cases of beer for the first 4 weeks; 8 cases of beer for the rest of weeks.
Stochastic lead time	2, 0, 2, 4, 4, 4, 0, 2, 4, 1, 1, 0, 0, 1, 1, 0, 1, 1, 2, 1, 1, 1, 4, 2, 2, 1, 4, 3, 4, 1, 4, 0, 3, 3, 4
Constant lead time	2 weeks

4.1.2 Order Policies

The order policy recommended by Kimbrough et al. [12], using the Genetic Algorithm, suggests that each player orders $[x_t + y]$ each week. Here, the x_t refers to the order every player received in Week T, and y is an integer generated by the algorithm for each player to minimize the total cost. The order policy recommended by Chaharsooghi et al. [4], using the Reinforcement Learning Ordering Mechanism, suggests that every player orders $[x_t + y_t]$ each week. Here, y_t is an integer generated by the algorithm for each player for Week T to minimize the total cost. As discussed before, the Q-model generates a set of $[Q_R, Q_W, Q_D, Q_F]$ for the players to estimate how much inventory they should obtain and then make the order decision. Table 4.2 shows the optimized order policies of the above three models for the basic constant and stochastic demand and lead time scenarios (shown in Table 4.1).

In Table 4.2, for the stochastic demand and lead time case, the Genetic Algorithm dictates that the retailer, the wholesaler, and the factory order one more case of beer along with the order they received from their downstream player each week, and that

Table 4.2: Optimized Order Policies Recommended by GA, RLOM, and Q-model in the Basic Constant and Stochastic Scenarios

Demand	Lead Time	GA	RLOM	Q-model
Stochastic	2	$[x+1, x+1, x+2, x+1]$	Table 3 [3]	$[0, 2.75, 4.5, 4]$
Stochastic	Stochastic	$[x+1, x+1, x+2, x+1]$		$[0, 2.5, 4.5, 3.25]$
4-8	2	$[x+0, x+0, x+0, x+0]$		$[0.25, 4, 4.75, 3.75]$
4-8	Stochastic	$[x+0, x+0, x+0, x+0]$		$[0.75, 4, 4, 3.25]$

the distributor orders two more cases of beer along with the order the distributor receives from the wholesaler each week. The optimized Q-model order policy for the stochastic demand and lead time case suggests that the retailer, the wholesaler, the distributor, and the factory order up to keep $[0, 2.75, 4.5, 4]$ weeks of expected future demand, respectively.

4.1.3 Testing Results

With the order policies generated by different models shown in Table 4.2, we can calculate the total costs with the simulation. The simulation results are shown in Table 4.3.

Table 4.3: Comparison of Results of the GA, RLOM, and Q-model in the Basic Constant and Stochastic Scenarios

Demand	Lead Time	GA	RLOM	Q-model
Stochastic	2	2954	3155	3438
Stochastic	Stochastic	2977	3249	3429
4-8	2	2224		2633
4-8	Stochastic	2480		2217

4.1.4 Analysis of Results

In Table 4.3, the results from the GA perform the best. However, the major part of the cost in the results of the GA order policy is the extra beer kept in the inventory. The reason for this extra cost is that for the GA model, the players always order more than the actual demand. The longer the game lasts, the more the cost will

be. For example, if the order policy for the retailer is $x + 2$, and the mean value of the actual customer's demand is 10 cases of beer in the 35 weeks of the game; the average beer ordered by the retailer will be 12 cases of beer. Therefore, there will be approximately 70 cases in total left in the inventory of the retailer at the end of the game. These 70 cases are accumulated throughout the game and cause a huge inventory cost. This excessive order of beer applies all the way up to the factory.

Figure 4-1 shows the inventory level of all the players throughout the game under the order policy of $[x+1, x+1, x+2, x+1]$ when the demand is stochastic and lead time is constant. We can see that at the end of the game, the retailer, the wholesaler, and the factory all have about $30 - 35$ cases of beer kept in their inventory, and the distributor has about 70 cases of beer in stock. These cases are continuously being accumulated from Week 15 until the end of the game. The excessive inventory cost takes more than 60% of the total cost of \$2954 as shown in Table 4.3.

The Q-model is a manageable approach and it captures the expected demand very well. Figure 4-2 shows the actual customer demand and the forecasts of each player in the game using the Q-model. From the figure we can see that every player can make a very similar demand forecast without access to the end customer's demand. We can also see that the order policy can capture the seasonality of the demand as well, since the Q-model generates the demand forecast using the moving average.

However, from Table 4.3, we can see that the order policy recommended by the Q-model in this thesis is not as good as the policies suggested by GA and RLOM. This is because the optimized order policy suggested by the Q-model recommends that the retailer keep 0 weeks of inventory ahead, which means whenever an order comes to the retailer, it will become a back order and be fulfilled later. Figure 4-3 shows the inventory level of all the players under the Q-model order policy. We can see that the retailer is almost always in the status of having back orders.

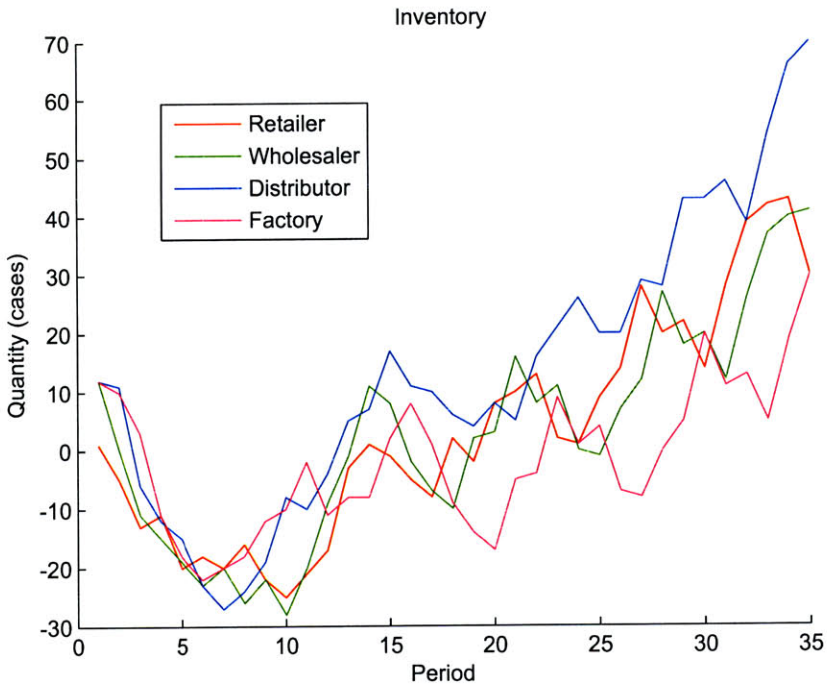

Figure 4-1: Inventory Level under the Genetic Algorithm Order Policy in a Scenario with Constant Lead Time and Stochastic Demand

Noticing the backorder cost is twice as much as the inventory cost, it seems inappropriate that the optimized order policy suggest that the retailer keep back orders throughout the game. Let us manually adjust Q_R from 0 to 1 to let the retailer keep one week's inventory ahead on hand whenever an order comes. By making this change in the order policy, we can see that the inventory of the retailer has been improved, and the backorder cost of the retailer has been slightly reduced, shown in Figure 4-4.

The adjusted order policy reduces the backorder cost of the retailer. However, it increases the backorder cost of the wholesaler and the inventory cost of the distributor and the factory. This is because at the beginning of the game, each player orders a huge amount of beer to acquire safety stock, which is determined by the expected demand and the Q. For example, if the order policy is $[1, 2.75, 4.5, 4]$, the retailer orders one week's inventory ahead at the beginning of the game, while the distributor orders 2.75 weeks' inventory for itself and 1 week's inventory for the retailer. This

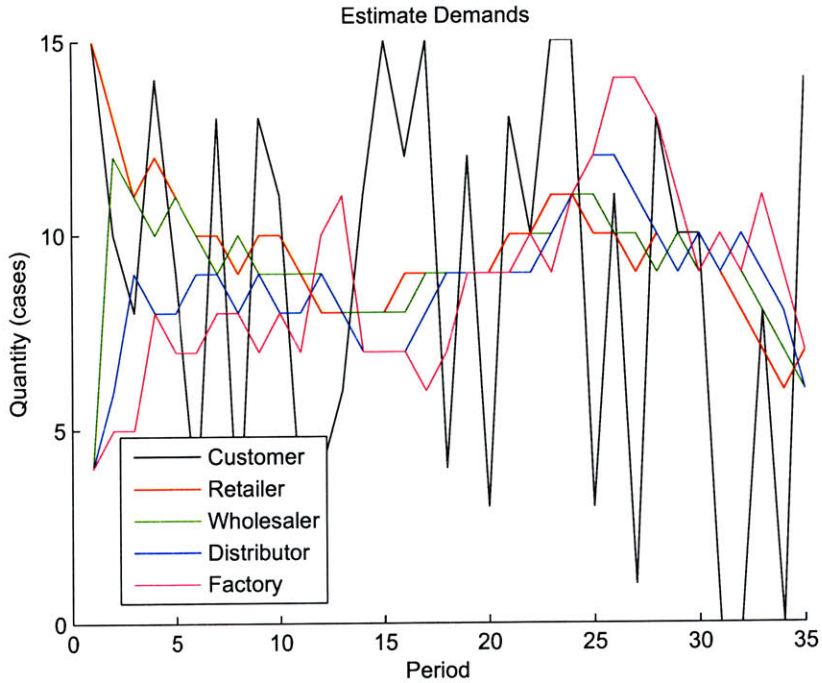

Figure 4-2: Demand forecast under the Q-model Order Policy with Stochastic Demand and Lead Time

goes all the way up to the factory and leads the factory to place a huge order at the beginning of the game, shown in Figure 4-5. In the simulation, we found that the factory orders 43 cases in Week Four when order policy is $Q = [0, 2.75, 4.5, 4]$. If we manually change Q_R from 0 to 1, even though the inventory situation of the retailer has improved (shown in Figure 4-4), the factory will place an order of 76 cases of beer in Week Four (shown in Figure 4-6), which is far more than the actual demand.

In this section, we can see that all the players in the beer supply chain are able to forecast the demand well using the Q-model. However, even though they have good demand forecast, they are still not able to achieve a low logistic cost due to the huge orders for acquiring safety stock at the beginning of the game. The next section introduces a solution to this problem.

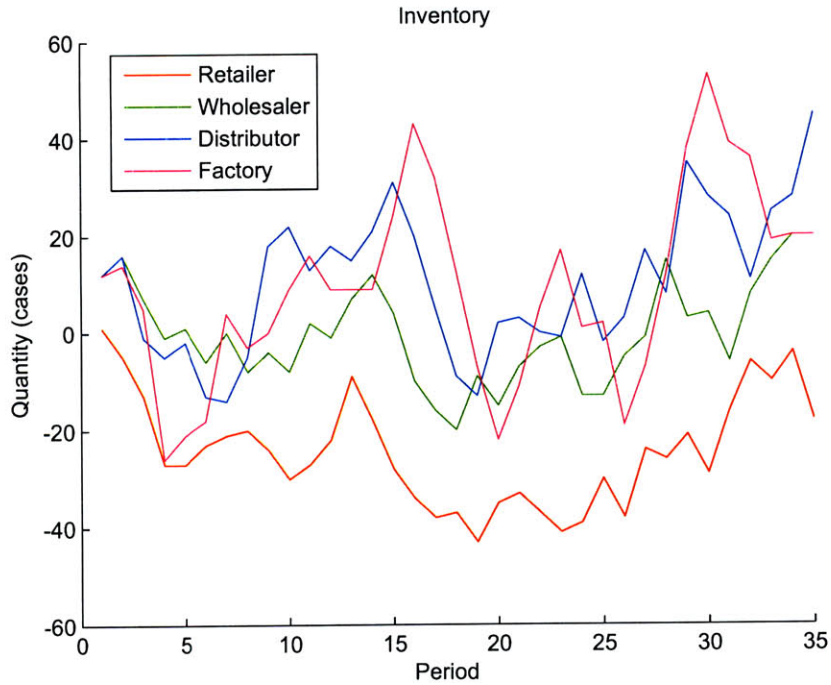

Figure 4-3: Inventory Level under the Q-model Order Policy $Q = [0, 2.75, 4.5, 4]]$ in a Scenario with Constant Lead Time and Stochastic Demand

4.2 Smoothing Technique

The previous section simulated the Q-model and compared the total cost of the Q-model with the order policies recommended by the Genetic Algorithm and the Reinforcement Learning Order Mechanism. However, the results from the order policy recommended by the Q-model are not as good as the results from the other two. This is mainly because of the huge orders for acquiring safety stock at the beginning of the game, which causes high inventory cost during the game, and continuous backorder cost for the retailer. In this section, we introduce a smoothing technique to increase the optimized value of Q_R and slow down the process of acquiring safety stock for all the players. The heuristics are as follows:

1. Choose H as the smoothing period length;

2. Given a set of Q' value as $[Q'_R, Q'_W, Q'_D, Q'_F]$;

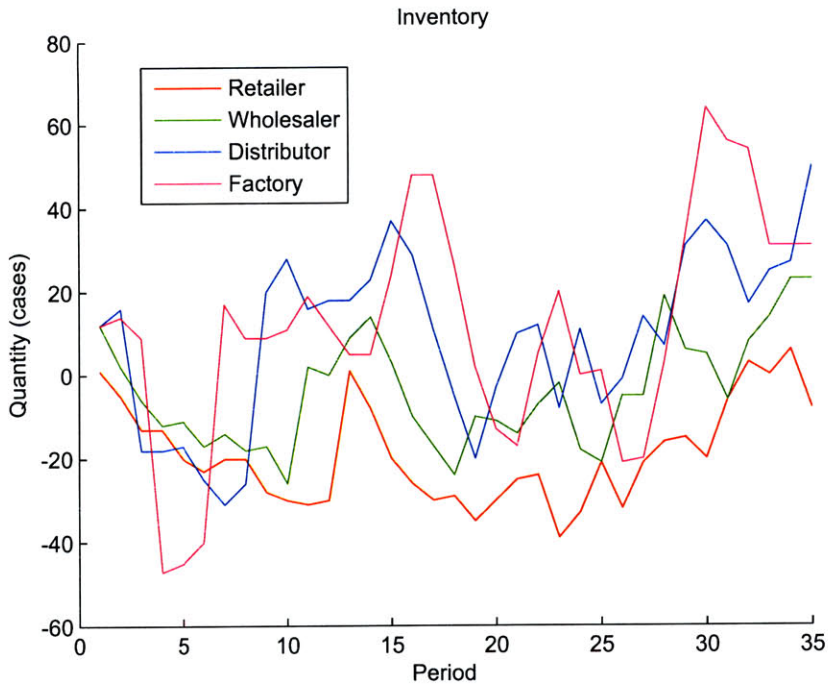

Figure 4-4: Inventory Level under the adjusted Q-model Order Policy $Q = [0, 2.75, 4.5, 4]$ in Comparison to Figure 4-3

3. During Week 1 to Week H, gradually change Q_R from 0 to Q'_R while Q'_W, Q'_D, Q'_F remain the same;

4. Calculate the cost of inventory under given Q', change the set of Q' until the optimum $Q^* = [Q^*_R, Q^*_W, Q^*_D, Q^*_F]$ has been found.

By adding the smoothing heuristics introduced above into the Q-model, we can slow down the safety stock acquisition process duration for the retailer from one week to H weeks. In this way, the retailer will not always keep back orders since the optimized Q^*_R is no longer 0. Meanwhile, the orders placed by the upstream players at the beginning of the game will not cause excessive inventory cost.

4.2.1 Testing Results

Let us call the new model with the smoothing technique implemented as the smoothed Q-model. The optimized order policies generated by the smoothed Q-model are shown

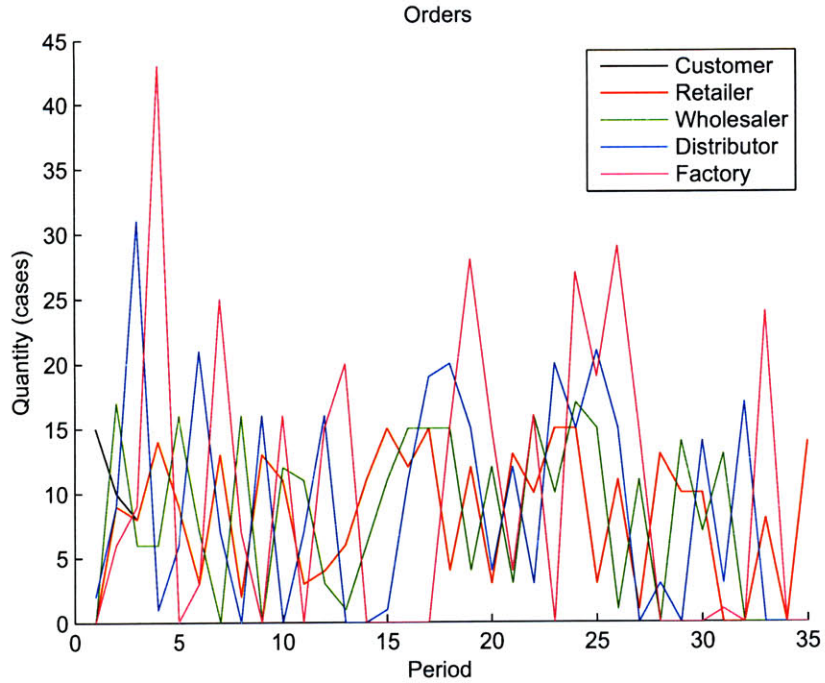

Figure 4-5: Orders Placed by Every Player under Optimized Q-model Order Policy $Q = [0, 2.75, 4.5, 4]$ in a Scenario with Constant Lead Time and Stochastic Demand

in Table 4.4 with comparison to the other models.

Table 4.4: Optimized Order Policies under the Smoothed Q-model with Comparison to the GA, RLOM, and the Original Q-model

Demand	**Lead Time**	**GA**	**Q-model**	**Smoothed Q-model**
Stochastic	2	$[x+1, x+1, x+2, x+1]$	$[0, 2.75, 4.5, 4]$	$[1, 3, 5, 3.5]$
Stochastic	Stochastic	$[x+1, x+1, x+2, x+1]$	$[0, 2.5, 4.5, 3.25]$	$[1.5, 2.5, 4.5, 3]$
4-8	2	$[x+0, x+0, x+0, x+0]$	$[0.25, 4, 4.75, 3.75]$	$[1, 4.25, 4.5, 4.5]$
4-8	Stochastic	$[x+0, x+0, x+0, x+0]$	$[0.75, 4, 4, 3.25]$	$[1, 3.5, 3.5, 2.75]$

The total accumulated inventory cost of all the players throughout the game under the smoothed Q-model are shown in Table 4.5 with comparison to the other models.

4.2.2 Analysis of Results

In Table 4.4, the optimized order policy generated by the Q-model is no longer 0 for the retailer after we apply the smoothing technique. This change has significantly

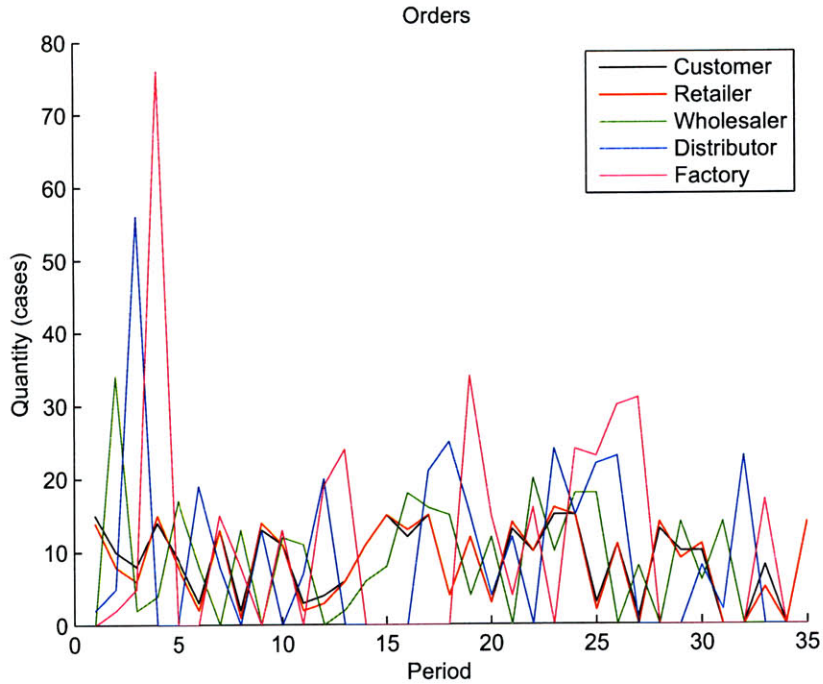

Figure 4-6: Orders Placed by Every Player under Adjusted Q-model Order Policy $Q = [1, 2.75, 4.5, 4]$ in a Scenario with Constant Lead Time and Stochastic Demand

Table 4.5: Inventory Costs under the Smoothed Q-model with Comparison to the GA, RLOM, and the Original Q-model

Demand	**Lead Time**	**GA**	**RLOM**	**Q-model**	**Smoothed Q-model**
Stochastic	2	2954	3155	3438	3310
Stochastic	Stochastic	2977	3249	3429	2686
4-8	2	2224		2633	2429
4-8	Stochastic	2480		2217	1913

reduced the backorder cost at the retailer. Figure 4-7 shows that the retailer no longer always keeps back orders, and all the players' inventory level throughout the game fluctuates mainly between -10 to 10 cases.

Table 4.5 shows the smoothed Q-model results, which have greatly improved compared to the results from the original Q-model, especially when the demands and lead times are both stochastic. Section 4.1.4 analyzed the disadvantage of the Q-model,

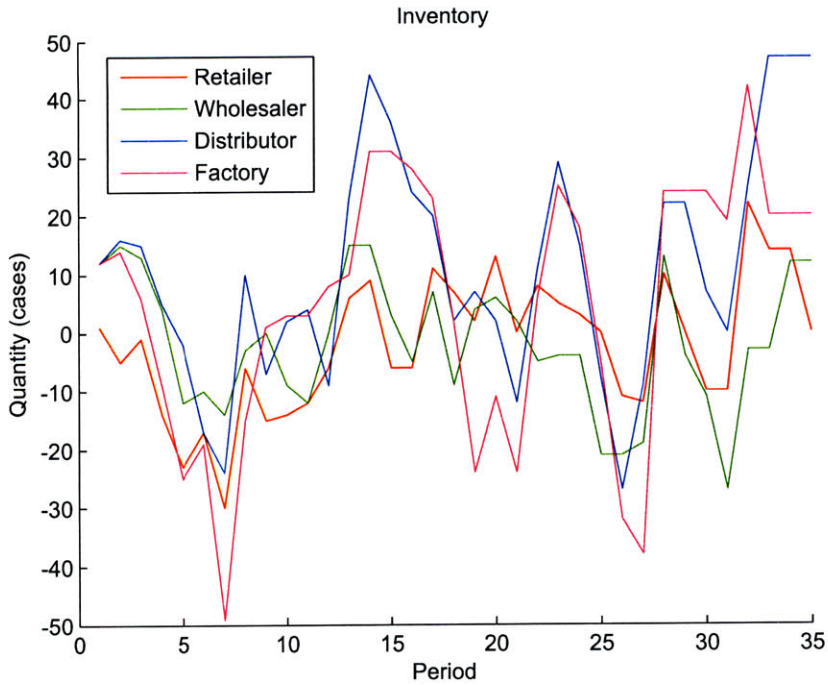

Figure 4-7: Inventory Level with Order Policy under the smoothed Q-model $Q = [1.5, 2.5, 4.5, 3]$ in a Scenario with Stochastic Lead Time and Demand

which is that either the retailer keeps continuous back orders or the factory and the distributor keep large excessive inventory. By implementing the smoothing technique, we can see that the back orders have been reduced at the retailer. Figure 4-8 shows the orders placed by all the players under the smoothed Q-model order policy. In this figure, the factory places orders more appropriately compared to the orders from the original Q-model shown in Figure 4-6.

4.3 General Scenario Simulations

In the Beer Distribution Game, the customer demand is unknown. It could follow a random distribution or a normal distribution or even any other unknown distribution. Section 4.1 and 4.2 simulated the scenarios with one set of samples in either stochastic or constant demand and lead time. This section conducts more general experiments to compare the performance of the smoothed Q-model with the GA order policies.

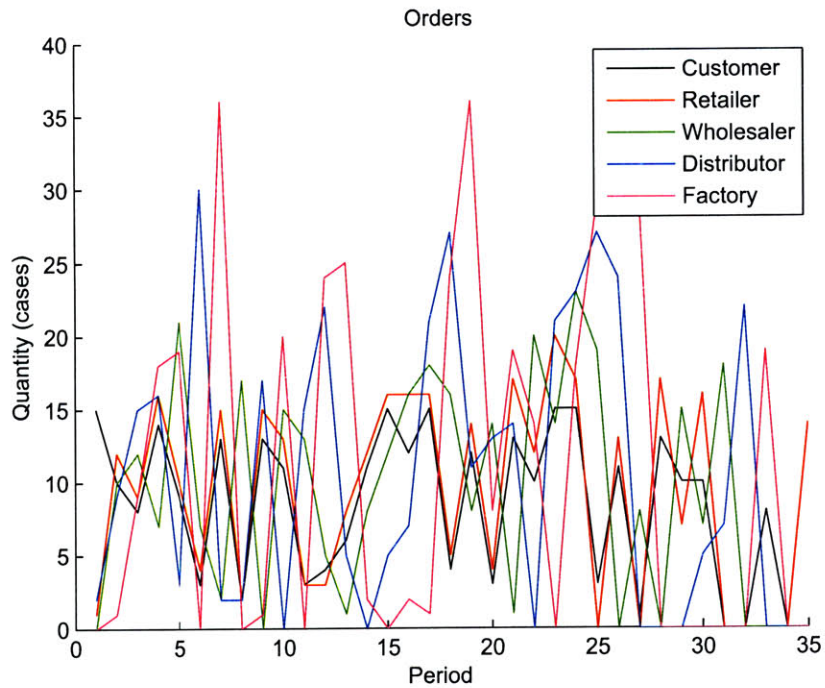

Figure 4-8: Order Placed by Each Player under the smoothed Q-model $Q = [1.5, 2.5, 4.5, 3]$ in a Scenario with Stochastic Lead Time and Demand

The purpose of performing the experiments is to see if the smoothed Q-model is robust and if the model can perform well in different scenarios.

4.3.1 Randomly Distributed Demand

Fifty sets of demands stochastically distributed in $0 - 15$ are generated as the input of the smoothed Q-model to simulate the optimized order policy, shown in Table 4.6. In the table, we can see that the optimized order policies of the stochastic lead time and the constant lead time are almost the same except the stochastic lead time policy is slightly smaller. This is because there are periods when the lead time is 0 or 1, so that the upstream can ship out the beer for the back orders or fulfill the current order in a very short time.

The total logistic costs of the general stochastic demand are shown in Table 4.7. The results from the smoothed Q-model are a little higher than the results from the

Table 4.6: Order Policy Comparison under 50 Sample Sets of Random Distributed Demand between 0 and 15

Demand	Lead Time	Test Sets	GA	Smoothed Q-model
0-15 Stochastic	2	50	[x+1,x+1,x+2,x+1]	[1,3.5,4.5,4.75]
0-15 Stochastic	0-4 Stochastic	50	[x+1,x+1,x+2,x+1]	[1,3,3.75,4.5]

GA and RLOM. We can see that the GA performs better when the demand has a huge fluctuation.

Table 4.7: Comparison of Results under 50 Sets of Stochastic Demand in 0-15

Demand	Lead Time	Test Sets	GA	RLOM	Smoothed Q-model
0-15 Stochastic	0-4 Stochastic	50	159493	154141	159199
0-15 Stochastic	2	50	144082	135496	155082

4.3.2 Normally Distributed Demand

For many products in their mature stage, the demand is usually very stable and follows a normal distribution. This section simulates demand in normal distributions, shown in Table 4.8.

Table 4.8: Order Policy and Cost under Normally Distributed Demand

Mean	Dev.	Lead Time	Sets	Q_R	Q_W	Q_D	Q_F	Total Cost
6	2	2	50	1	3.5	3.75	4.5	89923
6	2	0-4 Stochastic	50	1.25	3	3.75	3.75	97034
8	3	2	50	1	3.5	4.5	4.75	129060
8	3	0-4 Stochastic	50	1.25	3	3.75	4	143849
10	2	2	50	1	3	4.25	5.25	139710
10	2	0-4 Stochastic	50	1.25	3	3.75	4.5	161141
12	2	2	50	1	3	3.75	5.5	188943
12	2	0-4 Stochastic	50	1.25	2.75	4	4.5	218246

From the above table, we can see that by increasing the mean value of the distribution, the order policy does not change much. However, the total cost has changed dramatically. Initial inventory is one of the main reasons that contribute to the change of the cost.

4.3.3 Initial Inventory

In the previous section, we found that the cost increases significantly when the mean value of the demand increases under the smoothed Q-model order policy. This section examines the effect of the initial inventory on the cost. Table 4.9 shows the order policies and costs when changing the initial inventory from 12 cases of beer to 20. With this change, the cost has been greatly reduced for the normally distributed demand with mean value of 12.

Table 4.9: Comparison of Order Policies and Costs under Normally Distributed Demand with Initial Inventory 12 and 20 Cases of Beer

Initial Inventory	Demand		Lead Time	Sets	Q_R	Q_W	Q_D	Q_F	Total Cost
	Mean	Deviation							
12	6	2	2	50	1	3.5	3.75	4.5	89923
12	6	2	Stochastic	50	1.25	3	3.75	3.75	97034
20	6	2	2	50	1.5	3.25	4	3.5	116868
20	6	2	Stochastic	50	2	3	3.75	3.25	118671
12	10	2	2	50	1	3	4.25	5.25	139710
12	10	2	Stochastic	50	1.25	3	3.75	4.5	161141
20	10	2	2	50	1.25	3.5	4	4	119971
20	10	2	Stochastic	50	1.25	3	3.75	4	141160
12	12	2	2	50	1	3	3.75	5.5	188943
12	12	2	Stochastic	50	1.25	2.75	4	4.5	218246
20	12	2	2	50	1	3.25	3.75	4.5	132616
20	12	2	Stochastic	50	1.5	2.75	3.75	4	169558

Table 4.9 shows that when changing the initial inventory from 12 cases to 20 cases, the logistic cost of the scenario, which has a mean value of 6 cases, has increased by 20%. However, the logistic cost of the scenario, which has a mean value of 12, has decreased by 30%. Therefore, the initial inventory does have a huge impact on cost under the smoothed Q-model order policy. Now, let us examine the relationship between the initial inventory and the normally distributed demand. Table 4.10 shows the optimized order policies and costs when changing the initial inventory of every player. Table 4.11 shows the change of the order policies and the costs of the demand stochastically distributed between 0 and 15 when changing the initial inventory.

Table 4.10: Comparison of Order Policies and Costs under Normally Distributed Demand (12, 2) with Different Initial Inventory

Init. Inv.	Mean	Dev.	Lead Time	Sets	Q_R	Q_W	Q_D	Q_F	Total Cost
10	12	2	2	50	0.75	3.25	4	5.25	212028
12	12	2	2	50	1	3	3.75	5.5	188943
14	12	2	2	50	1	3	4	5.25	165418
16	12	2	2	50	1	3.25	4.25	4.75	150088
18	12	2	2	50	1	3.25	3.75	4.5	137391
20	12	2	2	50	1	3.25	3.75	4.5	132616
22	12	2	2	50	1.25	3.5	3.75	4.5	133301
24	12	2	2	50	1.25	3.5	3.75	4.5	136264
26	12	2	2	50	1.5	3.5	3.75	4.5	143733
28	12	2	2	50	1.75	3.5	3.75	4.5	153495
30	12	2	2	50	1.75	3.5	3.5	4.5	162450

Table 4.11: Comparison of Order Policies and Costs under Stochastically Distributed Demand between (0,15) with Different Initial Inventory

Initial Inventory	Demand	Lead Time	Sets	Q_R	Q_W	Q_D	Q_F	Total Cost
8	Stochastic	2	50	0.5	3	4.75	4.5	178729
12	Stochastic	2	50	1	3.5	4.5	4.75	155082
15	Stochastic	2	50	1	3	3.75	4.5	153928
20	Stochastic	2	50	1.5	3	3.75	4	154954
25	Stochastic	2	50	2	3	3.5	4.5	172672
30	Stochastic	2	50	2	2.5	3.5	4.5	194946

From Table 4.10 and Table 4.11, we can see that the change of initial inventory does not change the order policies much; however, it changes the costs a lot. Figure 4-9 shows that the best initial inventory should be about 20 or 21 cases of beer if the demand is normally distributed with a mean value of 12 and deviation of 2, while it also shows that the cost does not change much if the initial inventory is between 12 and 20 for the demand stochastically distributed between 0 and 15. Therefore, the smoothed Q-model is more sensitive to the initial inventory if the demand follows a normal distribution than if the demand follows a stochastic distribution.

Figure 4-9 shows that when the initial inventory is less than 10 cases of beer or more than 25 cases of beer, the cost changes dramatically if the initial inventory decreases or increases more. The reason is that when the initial inventory falls to

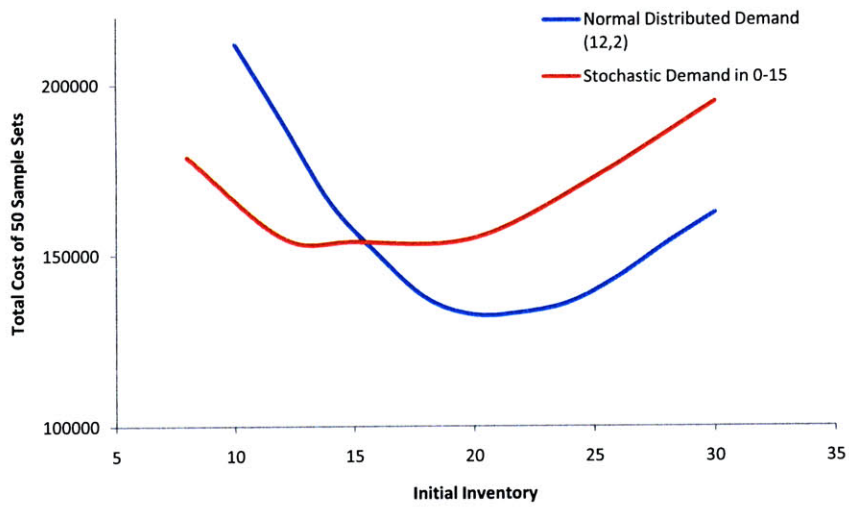

Figure 4-9: Comparison of Normally Distributed Demand (12,2) and Random Distributed Demand between 0-15 with Different Initial Inventory

less than 10, in the first few weeks, all the players stack a lot of back orders which continuously cause a huge backorder cost. However, when the initial inventory is more than 25 cases of beer, it will contribute unnecessary storage cost to the total cost. The best initial inventory is approximately 1.6 times the average demand from our experiments.

4.3.4 Seasonal Demand

Beer is a seasonal product, the demand is usually higher in the summer than in the winter. Therefore, it would also be interesting to simulate a scenario with seasonal demand. The demand is designed as Table 4.12, and the results are shown in Table 4.13.

Table 4.12: Seasonal Demand Distribution

Week	0 - 15	16 - 25	26 - 35
Demand	(0,15) stochastic	Normal distribution (12, 3)	(0,15) stochastic

Table 4.13: Optimized Order Policy under the Smoothed Q-model and Its Costs in the Scenario of Seasonal Demand

Case	Demand	Lead time	Scenarios	Q_X	Cost
1	Seasonal	2	100	[0.75,3.5,4.75,4.75]	382380
2	Seasonal	Stochastic	100	[1.25,3,3.75,4.5]	335061

Table 4.13 shows the optimized order policies generated by the smoothed Q-model for the seasonal demand scenarios and the costs. The order policies for the seasonal demand are very similar to the order policies for the other general scenarios. The next section will discuss the robustness of the order policy generated by the smoothed Q-model.

4.3.5 Robust Analysis

Let us take a look at the order policies suggested by the smoothed Q-model in various demand and lead time scenarios in the previous sections. Table 4.14 shows that the order policies are fairly the same and very robust. The order policy changes a little when the testing sample changes from 1 set to 50 sets. However, it is very robust when it comes to changing lead time from constant to stochastic, or varying the distribution of the demand.

Table 4.14: Robust Analysis of Order Policy Generated by Smoothed Q-model

Case	Demand	Lead Time	Sets	Q_R	Q_W	Q_D	Q_F
1	Stochastic	2	1	1	3	5	3.5
2	4-8	2	1	1	3	5	3.5
3	Stochastic	2	50	1	3.5	4.5	4.75
4	(10,2)	2	50	1.25	3.5	4	4
5	(6,2)	2	50	1	3.5	3.75	4.5
6	Seasonal	2	100	0.75	3.5	4.75	4.75
7	Stochastic	Stochastic	1	1.5	2.5	4.5	3
8	4-8	Stochastic	1	1	3.5	3.5	2.75
9	Stochastic	Stochastic	50	1	3	3.75	4.5
10	(10,2)	Stochastic	50	1.25	3	3.75	4
11	(6,2)	Stochastic	50	1.25	3	3.75	3.75
12	Seasonal	Stochastic	100	1.25	3	3.75	4.5

Figure 4-10 shows the order policies optimized by the smoothed Q-model for scenarios that have different demands and lead times. In this figure, we can see that the optimized order policy for the retailer is about 1, 3.25 for the wholesaler, 4 for the distributor, and around 4 for the factory. This suggests that the smoothed Q-model is able to generate a very robust order policy no matter whether the lead time and the demand are stochastic or constant.

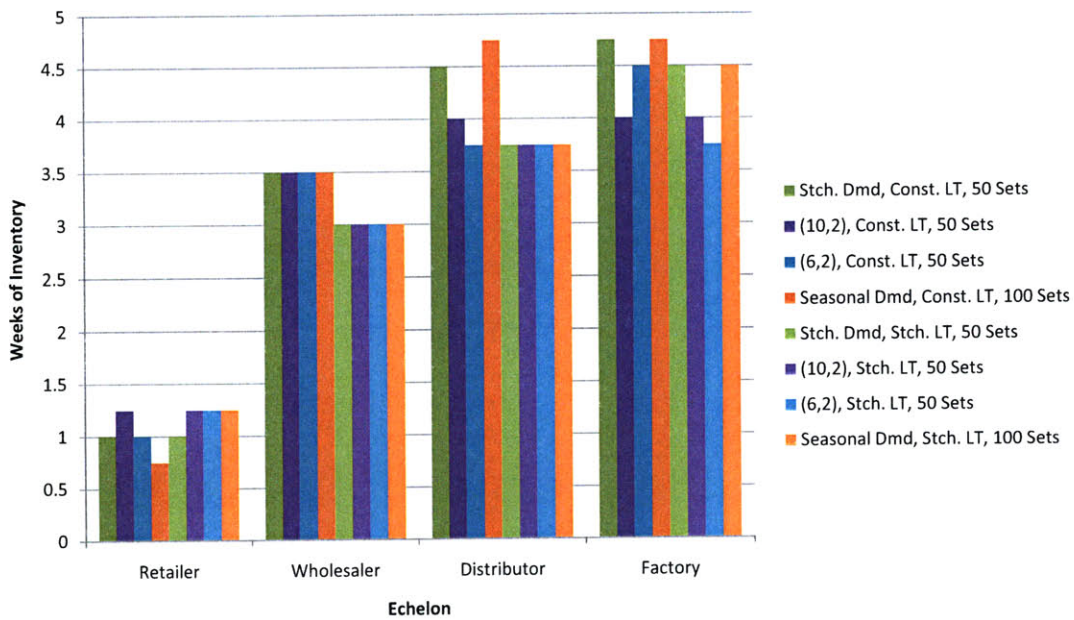

Figure 4-10: Robust Analysis of Order Policy Suggested by the Smoothed Q-model

Chapter 5

Conclusions and Future Work

In this thesis, we studied the bullwhip effect in supply chain management and the MIT Beer Distribution Game. In order to help the players perform better in the Beer Distribution Game and reduce the bullwhip effect in the beer supply chain, we proposed the Q-model to suggest an order policy for each of the players to decide how many cases of beer to order throughout the game. The policy provides all the players a Q set $[Q_R, Q_W, Q_D, Q_F]$ to suggest how many weeks' inventory each player should keep ahead. In order to evaluate the performance of the Q-model and compare it to the other models published, we have examined the current available simulations for both constant and stochastic lead times. Some of the simulations ignore the delivery of the back-ordered beer under the stochastic lead time scenarios, which makes it difficult to evaluate the performance of different order policies. I have developed another simulation to simulate the Beer Distribution Game under both constant and stochastic lead times with various demand scenarios.

Chapter 4 simulates the Beer Distribution Game using the order policy suggested by the Q-model proposed in Chapter 2. By introducing the smoothing technique into the Q-model, the results from the model have been improved. The simulation results show that the smoothed Q-model has almost the same or sometimes better results than the Genetic Algorithm. Moreover, from the experiments we have found that the smoothed Q-model generates a very robust order policy for all the players in various

scenarios. The optimized policy does not change much when changing the lead time and the demand distribution. We also examined the effect of the initial inventory to the total cost.

The following work will be focusing on 1) making better forecasts and eliminating the dependence on the downstream's Q value in the forecasts; 2) reducing the total cost by applying the smoothing technique to all the players in the supply chain; 3) improving the model to perform better when the demand has a huge fluctuation.

Bibliography

[1] *Averaging and Exponential Smoothing Models.* http://www.duke.edu/ rnau/411avg.htm.

[2] *NIST/SEMATECH e-Handbook of Statistical Methods.* http://www.itl.nist.gov/div898/handbook/, 2006.

[3] Manuel P. Baganha and Morris A. Cohen. The stabilizing effect of inventory in supply chains. *Operations Research*, 46(3):S27–S83, 1998.

[4] S. K. Chaharsooghi, J. Heydari, and S. H. Zegordi. A reinforcement learning model for supply chain ordering management: An application to the beer game. *Decision Support Systems*, 45:949–959, 2008.

[5] F. Chen, Z. Drezner, and J. Ryan amd D. Simchi-Levi. Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Science*, 46(3):436–443, 2000.

[6] F. Chen, J. Ryan, and D. Simchi-Levi. The impact of exponential smoothing forecasts on the bullwhip effect. *Naval Research Logistics*, 47(4):271–286, 2000.

[7] F. Chen and R. Samroengraja. The stationary beer game. *Production and Operations Management*, 9:19–30, 2000.

[8] R. Croson and K. Donohue. Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Science*, 00(0):1–14, 2005.

[9] J. Forrester. Industrial dynamics: A major breakthrough for decision makers. *Harvard Business Review*, 36:37–66, 1958.

[10] W. E. Jarmain. *Problems in Industrial Dynamics.* MIT Press, 1963.

[11] P. Kaminsky and D. Simchi-Levi. A new computerized beer distribution game: A tool for teaching the value of integrated supply chain management. *Supply Chain and Technology Management. Hau Lee and Shu Ming Ng, eds.*, 1:The Production and Operations Management Society, 216–225, 1998.

[12] S. O. Kimbrough, D. J. Wu, and F. Zhong. Computers play the beer game: can artificial agents manage supply chains? *Decision Support Systems*, 33:323–333, 2002.

[13] Hau L. Lee, V. Padmanabhan, and Seungjin Whang. The bullwhip effect in supply chain. *Sloan Management Review*, 38(4):93–102, 1997.

[14] Hau L. Lee, V. Padmanabhan, and Seungjin Whang. Information distortion in a supply chain: The bullwhip effect. *Management Science*, 43(4):546–558, 1997.

[15] Michael Li and David Simchi-Levi. *The web based beer game.* MIT Forum for Supply Chian Innovation, http://beergame.mit.edu/, 2002.

[16] Richard Metters. Quantifying the bullwhip effect in supply chains. *Journal of Operations Management*, 15:89–100, 1997.

[17] Joerg Nienhaus, Arne Ziegenbein, and Christoph Duijts. How human behaviour amplifies the bullwhip effect - a study based on the beer distribution game online. *Supply Chian World Europe*, 2002.

[18] T. O'Donnell, P. Humphreys, R. Mclvor, and L. Maguire. Reducing the negative effects of sales promotions in supply chain using genetic algorithms. *Expert Systems with Applications*, 36:7827–7837, 2009.

[19] D. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi. *Designing and Managing the Supply Chain, Concepts, Strategies and Case Studies.* McGraw-Hill Irwin, second edition, 2002.

[20] J. D. Sterman. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science*, 35:321–339, 1989.