

Analyzing the Make-to-Stock Queue in the Supply Chain and eBusiness Settings

by

René A. Caldentey

Civil Industrial Engineer, Univeristy of Chile (1995)

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Management

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2001

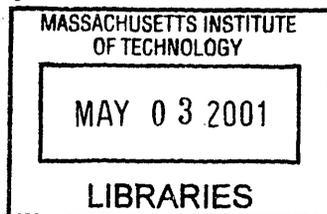
© René A. Caldentey, MMI. All rights reserved.

The author hereby grants to MIT permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part.

Author
Sloan School of Management
May 1, 2001

Certified by
Lawrence M. Wein
DEC Leaders for Manufacturing Professor of Management Science
Thesis Supervisor

Accepted by
Birger Wernerfelt
Professor of Management Science
Chair, Doctoral Program



ARCHIVES

Analyzing the Make-to-Stock Queue in the Supply Chain and eBusiness Settings

by

René A. Caldentey

Submitted to the Sloan School of Management
on May 1, 2001, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Management

Abstract

This thesis presents two applications of the prototypical make-to-stock queue model that are mainly motivated by supply chain management and e-commerce issues.

In the first part, we consider the decentralized version of the make-to-stock model. Two different agents that we call the supplier and the retailer control production and finish goods inventory level independently. The retailer carries finished goods inventory to service an exogenous demand and specifies a policy for replenishing his/her inventory from the upstream supplier. The supplier, on the other hand, chooses the capacity of his manufacturing facility. Demand is backlogged and both agents share the backorder cost. In addition, a linear inventory holding cost is charged to the retailer, and a linear cost for building production capacity is incurred by the supplier. The inventory level, demand rate and cost parameters are common knowledge to both agents. Under the continuous state approximation that the $M/M/1$ queue has an exponential rather than geometric steady-state distribution, we characterize the optimal centralized and Nash solutions, and show that a contract with linear transfer payments based on backorder, inventory and capacity levels coordinates the system in the absence of participation constraints. We also derive explicit formulas to assess the inefficiency of the Nash equilibrium, compare the agents' decision variables and the customer service level of the centralized versus Nash solutions, and identify conditions under which a coordinating contract is desirable for both agents.

In the second part, we return to the centralized version of the make-to-stock model and analyze the situation where the price that the end customers are willing to pay for the good changes dynamically and stochastically over time. We also assume that demand is fully backlogged and that holding and backordering costs are linearly incurred by the manufacturer. In this setting, we formulate the stochastic control problem faced by the manager. That is, at each moment of time and based on the current inventory position, the manager decides (*i*) whether or not to accept an incoming order and (*ii*) whether or not to idle the machine. We use the expected

long-term average criteria to compute profits. Under heavy traffic conditions, we approximate the problem by a dynamic diffusion control problem and derive optimality (Bellman) conditions. Given the mathematical complexity of the Bellman equations, numerical and approximated solutions are presented as well as a set of computational experiments showing the quality of the proposed policies.

Thesis Supervisor: Lawrence M. Wein

Title: DEC Leaders for Manufacturing Professor of Management Science

Acknowledgments

I am forever and deeply indebted to the many talented people who, during the course of my life, have given me not only knowledge, but more importantly taught me the secrets behind dedication, friendship, and generosity. I cannot possibly thank here by name the full complement of those to whom thanks are owed.

First and foremost, I am extremely grateful to my advisor Larry Wein who has provided inspiration, insight, and guidance without which this thesis would never have been completed. I will spend my career aspiring to match the quality of his research, teaching, and advising.

The energy and courage that I have had during these years to pursue this doctoral degree come from those that I have left behind in Chile. To my brothers in arms - René van Kilsdonk (Bob), Ivan Marinovic (Imán), Jorge Laval (Coco), Pedro Parra (Peter), and Rolando Cid (Rock&Rolo) with whom I grew up- I owe my curiosity for knowledge, my strong sense of friendship, and much much more. To my mentors -Don Leonardo Sanchez and Susana Mondschein- I owe my passion and love for academia.

Funding was obviously important enabling me to undertake and complete this research. Without the extreme generosity of the Department of Industrial Engineering of the University of Chile and the fellowship “Beca Presidente de la República” provided by Chilean government this work would have never been possible.

At MIT, one could not have wished to meet with a higher calibre group of persons. In particular, working with Professor Gabriel Bitran has been especially gratifying. He has not only supported my research these years with excellent advise, but more significantly he has always kept his door open to listen to my daily problems. I would like also to thank Professor Stephen Graves for his continuous and generous feedback.

Over my years at Cambridge I have made a number of friends who have always been source of comfort and strength. I thank the members of the Chilean brotherhood (“Cofradia”) -Fernando Ordoñez, Alejandro Neut (Misterioso), Mario Navarro, and Pablo García- for providing me the chance to periodically honor the Chilean cuisine always seasoned with friendship and unbalanced political discussion. I specially thank my classmates Juan Carlos Ferrer and Ariel Schilkrut for their unselfish academic support and for being great friends inside and outside the classroom. I deeply thank Hernan Alperin, Charles Krueger, Agneska Afelt, Machuca Ordoñez, and Ivan Muñoz for providing me all these years with the necessary escape from MIT, I will particularly miss those evenings of chess, bridge, and pool. I would like to

express my gratitude to Martin Haugh (Pl..) for promoting the healthy five o'clock breaks at ABP and for lifting my personal confidence when was needed. I also thank Nicolás Stier, Brian Tomlin, Paulo Gonçalves, Paulo Oliveira, Alp Muharremoglu, Ozie Ergun, John Bossert, Claudio Raddatz, Kevin Cowan, Manuel Amador, Patrio Mendez, Marina Zaretsky, Christian Broda, and Jérémie Gallien for making MIT a friendly and much better environment.

I was particularly fortunate to have Huy Nguyen as my roommate the first year at MIT. I am specially grateful to him for making my stay at Edgerton remarkably pleasant given our proximity to Necco. I also thank my current roommates, Armida Amador (Diva), Sapna Padte, Justin Flack, and Admar Verschoor for tolerating my mess these years and for making me feel always at home.

Finally I would like to thank my family: my parents Mateo and Mónica, whose love I have never doubted, to my brothers Esteban and Eugenio who have inspired me and always provided a model for me to live up to, and to my sisters Andrea y Margarita for their care and closed contact during all these years. My final and most important acknowledgment goes to my beloved Carola for all her support and love. Beyond all doubt, her optimism and sweet smile are the only responsible for making this moment unforgettable.

Dedicada a la memoria de la Pitita, Blanca Bonn Jarpa.

Contents

1	Introduction	17
2	The Make-to-Stock Queue Model	19
2.1	Fully Backlogged Demand	22
2.2	Lost Sales	25
2.3	Partially Backlogged Demand	29
3	A Decentralized Production-Inventory System	31
3.1	Introduction	31
3.2	The Model	33
3.3	The Centralized Solution	36
3.4	The Nash Solution	37
3.5	Comparison of Solutions	40
3.6	Contracts	45
3.7	The Stackelberg Games	50
3.8	Proofs	54
3.8.1	Proof of Proposition 3	54
3.8.2	Proof of Proposition 8	55
3.8.3	Proof of Proposition 9	56
4	Revenue Management of a Make-to-Stock Queue	59
4.1	Introduction	59
4.2	Literature Review	61
4.3	The Control Problem	63
4.4	The Diffusion Control Problem	65
4.5	Optimality Conditions	71
4.6	Numerical Solution	74

4.7	Approximations	80
4.7.1	Proposed Policy	80
4.7.2	Standard Form	81
4.7.3	Driftless Case	82
4.7.4	General Case	89
4.8	Computational Experiments	95
4.9	Extensions	99
4.9.1	Correlation Between Price and Demand	99
4.9.2	About the Error of the Proposed Policy	101
4.10	Proofs	104
4.10.1	Proof of Proposition (10)	104
4.10.2	Proof of Proposition (11)	105
4.10.3	Proof of Proposition (12)	106
4.10.4	Proof of Proposition (14)	108
4.10.5	Proof of Proposition (17)	109
4.10.6	Proof of Proposition (19)	109
5	Conclusions	111
A	Brownian Motion and Queueing Models	117
A.1	Introduction	117
A.2	Properties of Brownian Motions	120
A.2.1	Basic Properties	120
A.2.2	Wiener Measure and Donsker's Theorem	122
A.2.3	Reflection Principle	125
A.2.4	Some Extensions	126
A.3	Stochastic Calculus	129
A.3.1	Motivation	130
A.3.2	Stochastic Integration	131
A.3.3	Itô's Lemma	134
A.4	Regulated Brownian Motion	136
A.4.1	One-Sided Regulator	137
A.4.2	The Two-Sided Regulator	138
A.5	Queueing Models	139
A.5.1	Introduction	139

<i>CONTENTS</i>	11
A.5.2 Single-Stage Infinite Capacity Queueing Model	142
A.5.3 Single-Stage Finite Capacity Queueing Model	146
Bibliography	149

List of Figures

2.1	The Make-to-Stock Model	19
2.2	Exact versus Approximate Solution	27
2.3	Shape of the Function $H(Z)$	28
3.1	Agents' Profits as a Function of the Backorder Allocation Fraction . .	39
3.2	Nash Solution as a Function of the Backorder Allocation Fraction . .	44
4.1	Reflection Field at the Boundaries	75
4.2	Boundary Update	78
4.3	Numerical Solution to the Driftless Case	79
4.4	Myopic Approximation vs. Numerical Solution	87
4.5	Proposed Policy: Driftless Case	88
4.6	Normalized Expected Average Cost ($\tilde{g}(K)$) as a Function of K	89
4.7	Reflection Field	89
4.8	Proposed Policy: General Case	95
A.1	Donsker's Theorem	123
A.2	Simple Process	132
A.3	The One-Sided Regulator	138

List of Tables

3.1	The Competition Penalty	43
4.1	Performance as a Function of the Backordering Cost Rate	97
4.2	Holding/Backordering cost vs. Boundary Cost	98
4.3	Performance as a Function of the Demand and Production Variability	99

Chapter 1

Introduction

In this thesis we study the behavior of a make-to-stock queue model under two major variations of its traditional formulation.

First, and mainly motivated by the emerging literature on supply chain management, we consider in chapter 3 a decentralized version of the make-to-stock model. Two different agents that we call the supplier and the retailer control production and finish goods inventory level independently. More precisely, we model an isolated portion of a competitive supply chain as a $M/M/1$ make-to-stock queue. The retailer carries finished goods inventory to service a Poisson demand process, and specifies a policy for replenishing his inventory from an upstream supplier. The supplier chooses the service rate, i.e., capacity, of his manufacturing facility, which behaves as a single-server queue with exponential service times. Demand is backlogged and both agents share the backorder cost. In addition, a linear inventory holding cost is charged to the retailer, and a linear cost for building production capacity is incurred by the supplier. The inventory level, demand rate and cost parameters are common knowledge to both agents. Under the continuous state approximation that the $M/M/1$ queue has an exponential rather than geometric steady-state distribution, we characterize the optimal centralized and Nash solutions, and show that a contract with linear transfer payments based on backorder, inventory and capacity levels coordinates the system in the absence of participation constraints. We also derive explicit formulas to assess the inefficiency of the Nash equilibrium, compare the agents' decision variables and the customer service level of the centralized versus Nash solutions, and identify conditions under which a coordinating contract is desirable for both agents.

The second variation that we consider in this work is motivated by new business practices that e-commerce is creating on the B2B and B2C markets. In particular, we believe that manufacturers and suppliers that operate on the internet are beginning to experience higher variability on the price they receive for their production. For instance, Internet price search intermediaries (web aggregators) offer customers easy access to price lists and it is just a matter of time that most consumers' purchasing decisions will be based on this type of information. On the other hand, in the Business-to-Business setting the situation is not much different. The increasing popularity of online auctions is a good example showing how *spot markets* are winning ground over traditional long-term fixed price contracts. For these reasons, we studied in chapter 4 a make-to-stock model where the price that customers are willing to pay for the good changes dynamically and stochastically over time. We also assume that demand is fully backlogged and that holding and backordering costs are linearly incurred by the manufacturer. In this setting, we formulate the stochastic control problem faced by the manager. That is, at each moment of time and based on the current inventory position, the manager decides (i) whether or not to accept an incoming order and (ii) whether or not to idle the machine. We use the expected long-term average criteria to compute profits. Under heavy traffic conditions, we approximate the problem by a dynamic diffusion control problem and derive optimality (Bellman) conditions. Given the mathematical complexity of the Bellman equations, numerical and approximated solutions are presented as well as a set of computational experiments showing the quality of the proposed policies.

The rest of this thesis is organized as follows. In chapter 2 we present the basic make-to-stock model and analyze some of its main properties. In Chapter 3 we study the decentralized version of the system. We analyze the non-cooperative equilibrium between supplier and retailer and we present a mechanism that allows coordination. In chapter 4 we return to the centralized model and discuss the case where the selling price behaves as a geometric Brownian motion. The admission and production control problem is formulated and approximately solved. Chapter 5 contains our conclusions. Finally, because of the use of Brownian motion processes, heavy traffic analysis, and stochastic calculus in chapter 4, we conclude this thesis with an appendix that briefly detailed the main results of these research fields.

Chapter 2

The Make-to-Stock Queue Model

In this chapter we introduce the make-to-stock model and describe its main features. The goal is to present to the reader the basic elements and properties of this simple but powerful modelling device.

The basic make-to-stock queue model consists of three main components: (i) a demand process for a single durable good, (ii) a manufacturing facility that produces the good to satisfy the demand, and (iii) a buffer of finish goods inventory laying in between the end customer and the manufacturer. In the traditional make-to-stock model, the finish goods inventory is controlled by the manufacturer but it is also possible that a different entity owns this portion of the system. In this case, we will call (iii) the retailer's inventory. The analysis of different agents controlling the production process and the finish goods inventory is discussed in chapter 3.

The following picture schematically describe the basic element of the model. The

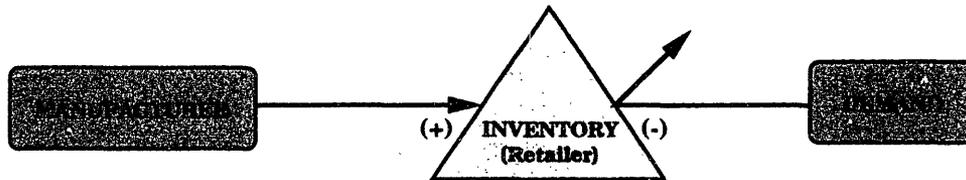


Figure 2.1: *The Make-to-Stock Model.*

plus sign in figure 2.1 indicates that the flow of finish goods coming out of the manufacturing facility contributes positively to the inventory position. On the other hand, the minus sign indicates that the flow of demand reduces the level of stock. We now give a more detailed description of the different components of the model.

Demand is generally modelled as a renewal process with known inter-arrival time distribution $f_D(t)$ and demand rate λ . We set $D(t)$ as the cumulative demand (number of orders) up to time t and $R(t)$ as the net revenue associated to selling one unit at time t (that is the difference between selling price and per unit production cost). In the traditional model that we will discuss in this chapter the selling price is constant over time, i.e., $R(t) = R$. Chapter 4 is devoted to removing this assumption considering a stochastic price process $R(t)$.

Production, on the other hand, is modelled as a single server queueing system. We set $S(t)$ as the cumulative production out of the manufacturer's facility up to time t . This is in general a stochastic process that depends on both (i) the service or production time of the machine and (ii) the production policy, i.e., the rule that decides when to have the machine producing or idle. Let $P(t)$ be the production policy, that is the cumulative amount of time up to time t that the machine has been working. $P(t)$ is a continuous and increasing process satisfying $0 \leq P(t) \leq t$. We also assume here that the machine has a service time distribution $f_S(t)$ and production rate μ .

Since production is capacitated, there is a natural delay between the moment an order is placed and the moment the final good corresponding to that particular order is finally produced. In order to avoid the customer to experience this delay, the manufacturer builds a buffer of finish goods inventory¹. Thus, each arriving customer simply takes one unit from the finish goods inventory leaving the system without further delay instead of placing an order to the manufacturer and waiting until that order is completed. Therefore, the type of product that we consider here can not be customized, that is every customer gets the same product that the manufacturer produces and stocks. We set X_0 as the initial level of inventory. Notice that since production and demand are both stochastic it is possible that an arriving customer sees no inventory on stock. This out-of-stock situation depends on the demand and service rate as well as in the initial inventory position and the production policy used. For instance, if the production rate is bigger than the demand rate ($\mu > \lambda$) and the manufacturer decides to have the machine working all the time ($P(t) = t$) then the probability of stock out is almost zero as time goes to infinity. However, in this case the inventory of finish goods will increase systematically with time making this approach very unattractive from an economic perspective.

¹For this reason we call these type of models *make-to-stock queues*.

More efficient policies can be constructed in order to balance the trade-off between the probability of stock-outs and the level of finish goods inventory. Here we will discuss one particular policy that we believe is commonly used in practice and has received most of the research attention. We call this policy *Re-order for each Item Sold* (RIS) which is the name used by Philip M. Morse in his seminal work *Queues, Inventories, and Maintenance*. The main idea behind this policy is that we place an order for another unit to the factory only when one unit is sold and that the machine is working only when there are orders waiting on queue. Let $X(t)$ and $Q(t)$ be the level of finish goods inventory and the number of orders at time t respectively. Suppose, moreover, that at time $t = 0$ there is no order at the manufacturer facility ($Q(0) = 0$) and that the initial inventory is $X(0) = X_0$. Then, under a RIS policy the following identity holds

$$X(t) + Q(t) = X_0 \quad \text{for all } t \geq 0.$$

This invariant property of the inventory position $X(t) + Q(t)$ characterizes the make-to-stock model that we consider in this work. As a first observation, we can notice that under a RIS policy the inventory of finish good $X(t) \leq X_0$ ensuring that the cost of holding inventory is bounded. On the other hand, stock-out will happen when $Q(t) \geq X_0$. The probability of this event depends on the particular admission policy that we consider. Three cases will be discussed here

- Case 1: Demand is fully backlogged. Every arriving order is accepted. If there is inventory on hand then the arriving customer gets immediately the product and one order for replenishment is place to the manufacturing facility. If there is no units on stock the arriving customer is backlogged until one unit of finish goods gets available for him/her. In addition, one order is place to the manufacturer. In this situation, the $X(t)$ can be arbitrarily negative in which case the absolute value represents the number of customers waiting for the good.
- Case 2: Lost sales. Customers are accepted only if there is inventory. In this case the customer gets the good and one order is place to the manufacturer. If there is no units on stock the arriving customer is lost and no order is place to the manufacturer. We notice that in this situation $X(t)$ is always nonnegative.
- Case 3: Demand is partially backlogged. This case is a combination of the previous two. Here some customers are backlogged but there is a maximum

number of customers that can be waiting for the product. We set $X_b \geq 0$ as the maximum number of backorders. Thus, in this case we have $X(t) \geq -X_b$. Case 1 is recovered making $X_b = \infty$ and case 2 is recovered making $X_b = 0$.

From a practical point of view which of these three cases is more appropriate to use will depend on the type of market that we consider. Case 1, for instance, seems to be appropriate for a monopoly type of market where customers do not have other alternative where to go to get the product. Case 2 on the other hand reflects the situation where consumers have many buying options, thus if one manufacturer does not have the product the customer will go somewhere else to get it. Finally, case three is an intermediate situation.

From a literature perspective, the $M/M/1$ make-to-stock queue was introduced by Morse (1958), but lay mysteriously dormant for the next three decades, perhaps because the multi-echelon version of it lacked the attractive decomposition property of the Clark-Scarf model and traditional (i.e., make-to-order) queueing networks, except under some restrictive inventory policies (Rubio and Wein 1996). Make-to-stock queueing systems have experienced a revival in the 1990s, including multi-product queues with (e.g., Federgruen and Katalan 1996, Markowitz *et al.* 1999) and without (e.g., Zheng and Zipkin 1990, Wein 1992) setups, and single-product, multi-stage systems in continuous time (e.g., Buzacott *et al.* 1992, Lee and Zipkin 1992) and discrete time (e.g., Glasserman and Tayur 1995 and Gavirneni *et al.* 1996, building on earlier work by Federgruen and Zipkin 1986).

In what follows, we will analyze in more detail the three cases mentioned above. In particular, we are interest on how to determine the optimal levels of X_0 (and X_b for case 3) as a function of the different system parameters. As a general comment, we would like to mention that we have selected the average profit (cost) criteria to compute the objective function to be maximize (minimize). Thus, steady-state average profit (cost) are considered.

2.1 Fully Backlogged Demand

When demand is fully backlogged the number of orders $Q(t)$ on the manufacturing facility is equivalent to the queue length of a single-server infinite-capacity queueing system. Certainly, the closed-form analysis of these systems is prohibited under general arrival and production processes. For this reason, we will focus on the Marko-

vian case, i.e., demand follows a Poisson process, production time is exponentially distributed, and both demand and production are independent.

In order to compute the optimal base-stock policy X_0 ² we have to balance the trade-off between holding inventory and having customers backlogged. We notice that the net revenue R associated to selling a unit does not play any significant role here since every order is eventually satisfied and the time average criteria is used.

On the other hand, we assume a simple linear holding/backordering cost function $c(X)$ such that

$$c(X) = \begin{cases} hX & \text{if } X \geq 0, \\ bX & \text{if } X < 0. \end{cases}$$

We can also write $c(X) = hX^+ + b(-X)^+$ where $X^+ = \max\{0, X\}$. Now, let Q be the number of orders in steady-state and let X be the steady-state finish goods inventory position. The existence of Q and X will be guaranteed under the Markovian assumption and the additional condition $\lambda < \mu$. Moreover, the invariant property of the make-to-stock model ensures that $X + Q = X_0$. The steady-state average cost as a function of X_0 is then given by

$$\bar{c}(X_0) = h E[X^+] + b E[(-X)^+] = h E[(X_0 - Q)^+] + b E[(Q - X_0)^+]. \quad (2.1)$$

Then, the problem that the manager of the production facility has to solve is to find a nonnegative integer X_0 that minimizes $\bar{c}(X_0)$ above. The solution of the problem is presented in the following proposition.

Proposition 1 *Let Q be the steady-state number of orders being processed by the manufacturer. Then, the optimal base-stock level X_0 satisfied*

$$X_0 = \min\{Z \geq 0 \text{ s.t. } b \leq (h + b) Pr(Q \leq Z)\}. \quad (2.2)$$

Proof: If X_0 is the optimal base-stock level then it must satisfied the condition

$$\bar{c}(X_0 - 1) \geq \bar{c}(X_0) \leq \bar{c}(X_0) + 1.$$

²Under a base-stock policy X_0 , the server is working as long as the finish goods inventory position is less than X_0 . When the inventory reaches the level X_0 the server turns to the idle or off position.

Combining these two inequality with the identities

$$\begin{aligned} E[(X+1-Q)^+] &= E[(X-Q)^+] + Pr(Q \leq X) \\ E[(Q-X)^+] &= E[(Q-X-1)^+] + Pr(Q \geq X+1) \quad \text{for all } X \geq 0 \end{aligned}$$

we get that X_0 must satisfy the optimality condition

$$(h+b)Pr(Q \leq X_0 - 1) - b \leq 0 \leq (h+b)Pr(Q \leq X_0) - b. \quad (2.3)$$

Since the expression $(h+b)Pr(Q \leq X_0) - b$ is nondecreasing in X_0 , condition (2.3) guarantees the uniqueness of X_0 . The existence is also guaranteed since $\lim_{X_0 \rightarrow \infty} Pr(Q \leq X_0) = 1$. Finally, condition (2.2) is simply a more concise version of (2.3). ■

If we forget for a moment that Q is a discrete random variable then X_0 is the solution to

$$Pr(Q \leq X_0) = \frac{b}{h+b}.$$

The solution is of a critical fractile type similar to those encountered on newsboy problems. Thus, we just need to compute the cumulative distribution $F_Q(x) = Pr(Q \leq x)$ for the number of orders on the production facility. The solution is given then by

$$X_0 = F_Q^{-1} \left(\frac{b}{h+b} \right).$$

For the moment, we have not make any particular assumption on Q rather than assuming that it has a steady-state distribution $F_Q(x)$. If we further impose the Markovian assumption we get that for the $M/M/1$ queue

$$Pr(Q \leq x) = 1 - \rho^{x+1} \quad \text{where } \rho = \frac{\lambda}{\mu}.$$

Thus, from condition (2.2) the optimal value of X_0 is given by

$$X_0 = \left\lceil \frac{\ln(h) - \ln(\rho(h+b))}{\ln(\rho)} \right\rceil. \quad (2.4)$$

The solution is consistent with the intuition. X_0 is increasing in b and ρ and decreasing

with h . We can also conclude that $X_0 = 0$ is optimal if

$$\rho \leq \frac{h}{h+b}.$$

In this case the make-to-stock system behaves like a make-to-order queue. We now turn to the analysis of Case 2.

2.2 Lost Sales

In this case we need to consider a different cost function. In particular, we do not have backordering but we do have rejections. We then assume that for each customer that is rejected because of stock-outs the system manager incurs a penalty cost f . We recall here that for each customer that is accepted the manager received a net revenue of R . Finally, the holding cost rate for keeping finish goods on stock is again h per unit. Also, since demand is not backlogged the number of orders in the system (Q) ranges from 0 to X_0 .

In this setting, the average revenue per unit time $\pi(X_0)$ as a function of X_0 is given by

$$\pi(X_0) = \lambda R(1 - Pr(Q = X_0)) - hE[(X_0 - Q)] - \lambda f Pr(Q = X_0).$$

The first term is the total expected revenue obtained by selling the good, we notice that $1 - Pr(Q = X_0)$ is the average fraction of customers that will effectively buy the good. The second term is the expected holding cost. Finally, the third term is the expected penalty associated to rejections, here $Pr(Q = X_0)$ is the stock-out probability.

The problem faced by the manager in this case is to find a nonnegative integer X_0 that maximizes $\pi(X_0)$ above. This problem was first analyzed by P. Morse (1958). In addition a quite similar but not identical problem was also studied by P. Naor (1969) in the context of socially optimal admission control to a queue. We notice that in this case, the distribution of Q does depend on the value of X_0 making hard the analysis of $\pi(X_0)$ without further assumption about the distribution of Q . Let us assume then that we are dealing with a Markovian system. In this case, Q is the steady-state number of order in an $M/M/1/X_0$ queueing system and it has the

following distribution:

$$\Pr(Q = k) = \frac{(1 - \rho) \rho^k}{1 - \rho^{X_0+1}} \quad \text{for all } k = 0, 1, \dots, X_0.$$

Proposition 2 *The optimal base-stock level X_0 for the lost sales case satisfies*

$$X_0 = \min\{Z \geq 0 \text{ s.t. } \frac{\lambda(R + f)}{h} \leq G(Z)\} \quad (2.5)$$

where the function $G(Z)$ is given by

$$G(Z) := \frac{1 - (Z + 2) \rho^{Z+1} + (Z + 1) \rho^{Z+2}}{\rho^Z (1 - \rho)^2}.$$

We notice that $G(Z)$ is a nondecreasing function with $G(0) = 1$ and $G(\infty) = \infty$. Thus, condition (2.5) is always well-defined. In addition, if $\lambda(R + f) \leq h$ then $X_0 = 0$ and the manager is better-off closing the production facility.

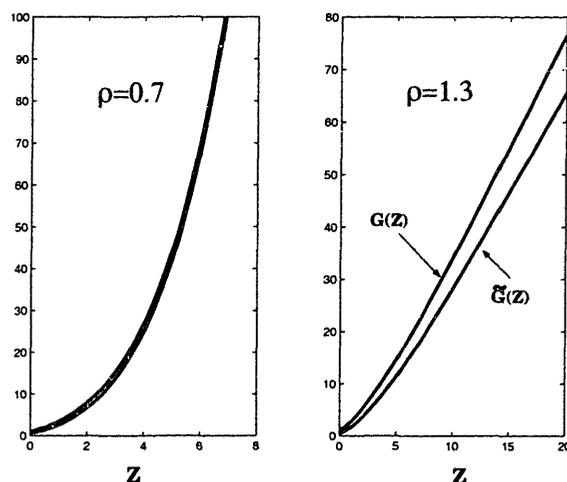
Proof: After some straightforward manipulations it can be shown that the optimality condition $\pi(X_0 - 1) \leq \pi(X_0) \geq \pi(X_0 + 1)$ is equivalent to

$$G(X_0 - 1) < \frac{\lambda(R + f)}{h} \leq G(X_0),$$

which turns out to be equivalent to (2.5). The existence of a solution is guaranteed because $\lim_{X_0 \rightarrow \infty} G(X_0) = \infty$. Finally, the uniqueness is also ensured since the function $G(X_0)$ is nondecreasing. ■

Unfortunately, and mainly because of the functional form of $G(Z)$, we can not get a closed form solution for the optimal level X_0 in this case.

One possible way to get good approximations for X_0 is to assume that Q is a continuous rather than discrete random variable. We can formally justify this transformation under *heavy traffic* arguments. The details of this transformation are presented in the Appendix A at the end. Here, we briefly mention that for system with traffic intensity closed to unity we can approximately model the behavior of the queue length $Q(t)$ by a regulated (one-sided or two-sided) Brownian motion. We do not pursue this approach here but we mention that the section §A.5 (and in particular §A.5.3) in the appendix contains all the elements required to complete this heavy

Figure 2.2: *Exact versus Approximate Solution.*

traffic analysis for this lost sales model.

Here, instead, we approximate X_0 in a much simpler and crude way. First, we notice that the function $G(Z)$ can be rewritten as follows

$$G(Z) = \sum_{k=0}^{Z+1} (Z+1-k) \rho^{-k}$$

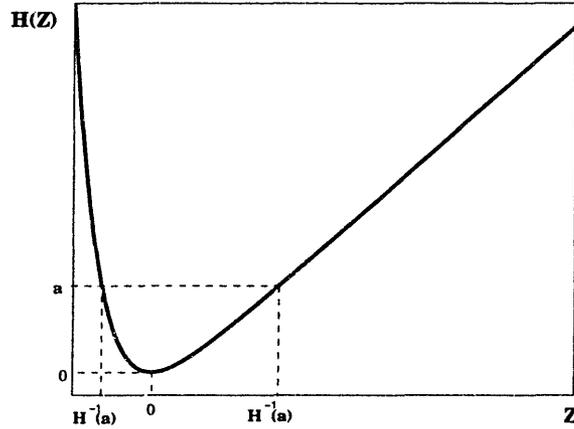
(The proof of this identity is direct and it is left to the reader). From this relation we can easily see that $G(Z)$ is nondecreasing in Z as it was claimed above. The key of our approximation is to replace the summation by an integral. That is,

$$G(Z) \approx \tilde{G}(Z) := \int_0^{Z+1} (Z+1-x) e^{-\theta x} dx \quad \text{where } \theta = \ln(\rho).$$

Solving the integral we get

$$\tilde{G}(Z) = \frac{e^{-\theta(Z+1)} + \theta(Z+1) - 1}{\theta^2}.$$

A quick comparison between $G(Z)$ and $\tilde{G}(Z)$ shows that in fact the approximation is quite accurate specially for values of $\rho < 1$. Figure 2.2 shows the behavior of $G(Z)$ and $\tilde{G}(Z)$ for two values of ρ . As we can see, for $\rho < 1$ the approximation is very precise. For the case $\rho > 1$ the approximation does not perform as good as in the previous case but still the prediction of X_0 obtained from it are good.

Figure 2.3: Shape of the function $H(Z)$,

The computation of the optimal level X_0 reduces to find the minimum Z satisfying

$$e^{-\theta(Z+1)} + \theta(Z+1) - 1 \geq \frac{\lambda\theta^2(R+f)}{h}.$$

Let define $H(Z) = \exp(-Z) + Z - 1$, then the approximation for X_0 is given by

$$X_0 = \left\lceil \frac{1}{\theta} H^{-1} \left(\frac{\lambda\theta^2(R+f)}{h} \right) - 1 \right\rceil. \quad (2.6)$$

The shape of the function $H(z)$ is plotted in figure 2.3. As we can see $H^{-1}(a)$ contains two different values for $a > 0$; one is positive and the other is negative. Which one of the two we will select depends on the sign of θ . If θ is negative we select the negative value of $H^{-1}(a)$. The opposite is true if $\theta > 0$. A second order Taylor expansion of $H(Z)$ suggest the following base stock policy.

$$X_0 \approx \left\lceil \sqrt{\frac{2\lambda(R+f)}{h}} - 1 \right\rceil.$$

Notice the similarities of this solution and the standard EOQ solution (e.g., Hadley and Whitin (1963)).

As a general comment about the solution in (2.6), we can see that X_0 is increasing in R and f and decreasing in h , as we should expect. It is important to notice that the impact of R and f in the optimal solution comes only through the sum $R + f$. Thus, low price high rejection penalty systems have the same optimal solution as high

price low penalty systems.

We now turn the last of the three cases which considers simultaneously backorders and rejections.

2.3 Partially Backlogged Demand

We consider in this case that the manager has the ability to backlog as well as to reject incoming orders. In this situation, the production and admission policy is characterized by two integers X_0 and X_b . X_0 is the base-stock policy that controls the production. That is, the machine will be working as long as the finish goods inventory level is under X_0 . On the other hand, X_b is the maximum level of backorders that the system is willing/able to hold. Thus, X_b controls the admission of orders to the system. Using the same notation that we have used in the previous two sections we can construct the expected utility per unit time for the system, $\pi(X_0, X_b)$, which is given by

$$\pi(X_0, X_b) = \lambda R - \lambda(R + f) Pr(Q = X_0 + X_b) - h E[(X_0 - Q)^+] - b E[(Q - X_0)^+].$$

In this case, Q ranges from $0, \dots, X_0 + X_b$. We notice that the optimization of $\pi(X_0, X_b)$ is more demanding in this case because of the two dimension. For this reason, we solve the optimization problem sequentially. First, we fix the sum $L := X_0 + X_b$ and we look for the optimal values of X_0 and X_b given L . Then, we select the optimal value of L .

If we fix L , then the distribution of Q is also fixed, independent of the particular values of X_0 and X_b . In this situation we get

$$Pr(Q = k) = \frac{(1 - \rho) \rho^k}{1 - \rho^{L+1}} \quad \text{for all } k = 0, 1, \dots, L.$$

Moreover, the first and last term in the utility function are constant and the problem reduces to minimize the holding/backordering cost

$$c(X_0, X_b) = h E[(X_0 - Q)^+] - b E[(Q - X_0)^+]$$

subject to $X_0 + X_b = L$. We can, in fact, rewrite the optimization problem only in terms of X_0 . It turns out that the optimality condition in this case with fixed L is

the same as (2.3), i.e.,

$$(h + b) Pr(Q \leq X_0 - 1) - b \leq 0 \leq (h + b) Pr(Q \leq X_0) - b.$$

The optimal solution X_0 is then characterized by the condition

$$Pr(Q \leq X_0 | X_0 + X_b = L) = \frac{b}{h + b}. \quad (2.7)$$

Under the Markovian assumption, we can rewrite (2.7) as

$$\frac{1 - \rho^{X_0+1}}{1 - \rho^{L+1}} = \frac{b}{h + b} \iff X_0 = \frac{\ln(h + b \rho^{L+1}) - \ln(h + b) - \ln(\rho)}{\ln(\rho)}.$$

The optimal solution is rigorously given by $\lceil X_0 \rceil$ and $X_b = L - \lceil X_0 \rceil$.

After some tedious algebra, the problem of finding the optimal value of L can be written as follows:

$$\begin{aligned} \min_{L \geq 0} & \left[\lambda(R + f) \frac{(1 - \rho)\rho^L}{(1 - \rho^L + 1)} \right. \\ & \left. + \frac{h(X_0 - (X_0 + 1)\rho + \rho^{X_0+1}) + b((L - X_0)\rho^{L+2} - (L - X_0 + 1)\rho^{L+1} + \rho^{X_0+1})}{(1 - \rho)(1 - \rho^{L+1})} \right] \\ \text{s.t.} & \\ X_0 & = \frac{\ln(h + b \rho^{L+1}) - \ln(h + b) - \ln(\rho)}{\ln(\rho)} \end{aligned}$$

Although messy, the optimization problem above is one-dimensional and it can be solved easily. We are, unfortunately, unable to compute the optimal solution in closed form. Again, we can try to use some type of approximations to estimate the optimal solutions (X_0, X_b) . The use of Brownian motion and heavy traffic conditions is one alternative. We postpone, however, its discussion to chapter 4 where we derive the optimal solution for this case of partially backlogged demand (see relations (4.67) and (4.68)).

Chapter 3

A Decentralized Production-Inventory System

3.1 Introduction

Within many supply chains, a devoted upstream agent, referred to here as the *supplier*, produces goods for a downstream agent, called the *retailer*, in a make-to-stock manner. Broadly speaking, the performance (e.g., service levels, cost to produce and hold items) of this isolated portion of the supply chain is dictated by three factors: (i) Retailer demand, which is largely exogeneous but can in some cases be manipulated via pricing and advertising, (ii) the effectiveness of the supplier's production process and the subsequent transportation of goods, and (iii) the inventory replenishment policy, by which retailer demand is mapped into orders placed with the supplier. If the supplier and retailer are under different ownership or are independent entities within the same firm, then their competing objectives can lead to severe coordination problems: The supplier typically wants to build as little capacity as possible and receive excellent demand forecasts and/or a steady stream of orders, while the retailer prefers to hold very little inventory and desires rapid response from the supplier. These tensions may deteriorate overall system performance.

The recent explosion in the academic supply chain management literature is aimed at this type of multi-agent problem. Almost without exception, the papers that incorporate stochastic demand employ variants of one of two prototypical operations management models: The newsvendor model or the Clark-Scarf (1960) multi-echelon inventory model. One-period and two-period versions of newsvendor supply chain

models have been studied intensively to address the three factors above; see Agrawal *et al.* (1999), Cachon (1999) and Lariviere (1999) for recent reviews. Although many valuable insights have been generated by this work, these models are primarily useful for style goods and products with very short life cycles. More complex (multi-period, and possibly multi-echelon and positive lead time) supply chain models have been used to analyze the case where a product experiences ongoing production and demand. Of the three factors in the last paragraph, these multi-period supply chain models successfully capture the replenishment policy and have addressed some aspects of retailer demand, e.g., information lead times in the Clark-Scarf model (Chen 1999), pricing in multi-echelon models with deterministic demand and ordering costs (Chen *et al.* 1999), and forecast updates (Anupindi and Bassok 1999 in a multi-period newsvendor model and Tsay and Lovejoy 1999 in a multi-stage model). However, the Clark-Scarf model, and indeed all of traditional inventory theory, takes a crude approach towards the supplier's production process, by assuming that lead times are independent of the ordering process, or equivalently, that the production process is an infinite-server queue.

In this chapter, we use an alternative prototypical model, an $M/M/1$ make-to-stock queue, to analyze a supply chain. Here, the supplier is modeled as a single-server queue, rather than an infinite-server queue, and the retailer's optimal inventory replenishment strategy is a base stock policy. Because the production system is explicitly incorporated, these make-to-stock queues are also referred to as *production-inventory systems*.

Much of the work done on make-to-stock queues (e.g., Wein 1992, Buzacott *et al.* 1992, Lee and Zipkin 1992, Rubio and Wein 1996, Federgruen and Katalan 1996, Markowitz *et al.* 1999, Glasserman and Tayur 1995 and Gavirneni *et al.* 1996) either undertake a performance analysis or consider a centralized decision maker (Gavirneni *et al.* analyze their system under various informational structures, but not in a game-theoretic setting), the make-to-stock queue is amenable to a competitive analysis because it explicitly captures the trade-off between the supplier's capacity choice and the retailer's choice of base stock level. However, the model treats the third key factor in a naive way, by assuming that retailer demand is an exogenous Poisson process. Moreover, we assume that the system state, the demand rate and the cost parameters are known by each agent. While this assumption is admittedly crude, we believe it is an appropriate starting point for exploring competitive make-to-stock queues. In the only other multi-agent production-inventory study that we are aware of, Plambeck

and Zenios (1999), contemporaneously to us, analyze a more complex dynamic system with information asymmetry.

In an attempt to isolate – and hence understand – the impact of incorporating capacity into a supply chain model, we intentionally mimic Cachon and Zipkin (1999). Their two-stage Clark-Scarf model is quite similar to our $M/M/1$ make-to-stock queue: Both models have two players, assume linear backorder and holding costs for retailer inventory (where the backorder costs are shared by both agents), employ steady-state analyses, and ignore fixed ordering costs. The key distinction between the two models is that the production stage is an infinite-server queue and the supplier controls his (local or echelon) inventory level in Cachon and Zipkin, whereas in our work the production stage is modeled as a single-server queue and the supplier controls the capacity level, which in turn dictates a steady-state lead time distribution. While Cachon and Zipkin’s supplier incurs a linear inventory holding cost, our supplier is subjected to linear capacity and production costs. Another deviation in the formulations is that Cachon and Zipkin’s agents minimize cost, while our agents maximize profit; this allows us to explicitly incorporate participation (i.e., nonnegative profits) constraints. A minor difference is that our queueing model is in continuous time, while Cachon and Zipkin’s inventory model is in discrete time. In fact, to make our results more transparent and to maintain a closer match of the two models, we use a continuous state approximation, essentially replacing the geometric steady-state distribution of the $M/M/1$ queue by an exponential distribution with the same mean.

The rest of this chapter is organized as follows. After defining the model in §3.2, we derive the centralized solution in §3.3, where a single decision maker optimizes system performance, and the Nash equilibrium in §3.4, where the supplier and retailer maximize their own profit. The two solutions are compared in §3.5. In §3.6, we describe the contract that *coordinates* the system; i.e., allows the decentralized system to achieve the same profit as the centralized system. In §3.7, we analyze the Stackelberg games, where one agent has all the bargaining power. Concluding remarks, on the other hand, are presented in chapter 5.

3.2 The Model

Our idealized supply chain consists of a supplier providing a single product to a retailer. Retailer demand is modeled as a homogeneous Poisson process with rate

λ . The retailer carries inventory to service this demand, and unsatisfied demand is backordered. The retailer uses a $(s - 1, s)$ base stock policy to replenish his inventory. That is, the inventory initially contains s units, and the retailer places an order for one unit with the supplier at each epoch of the Poisson demand process. Because we assume that there are no fixed ordering costs, the retailer's optimal replenishment policy is indeed characterized by the base stock level s .

The supplier's production facility is modeled as a single-server queue with service times that are exponentially distributed with rate μ . The supplier is responsible for choosing the parameter μ , which will also be referred to as the *capacity*. The server is only busy when retailer orders are present in the queue. The supplier's facility behaves as a $M/M/1$ queue because the demand process is Poisson and a base stock policy is used.

The selling price r that the retailer charges to the end customers and the wholesale price w that the retailer pays to the supplier are fixed. These conditions implicitly assume that the retailer and supplier operate in competitive markets. Each backordered unit generates a cost b per unit of time for the production-inventory system. As in Cachon and Zipkin, this backorder cost is split between the two agents, with a fraction $\alpha \in [0, 1]$ incurred by the retailer. The parameter α , which we refer to as the *backorder allocation fraction*, is exogenously specified in our model. Much of the academic literature assumes $\alpha = 1$; however, even if the supplier does not care about backorders *per se*, he would incur a cost in switching to a different retailer if he provided extremely poor service, which suggests that this extreme case is not very realistic.

In addition, the retailer incurs a holding cost h per unit of inventory per unit of time. The supplier pays the fixed cost of building production capacity and the variable production costs. The capacity cost parameter c is per unit of product, so that $c\mu$ represents the amortized cost per unit of time that the supplier incurs for having the capacity μ ; this fixed cost rate is independent of the demand level. The variable production cost per unit is denoted by p . We assume $r > w > c + p$, so that positive profits are not unattainable. To make our results more transparent, we normalize the expected profit per unit time by dividing it by the holding cost rate h . Towards this end, we normalize the cost parameters as follows:

$$\tilde{h} = \frac{h}{h} = 1, \quad \tilde{b} = \frac{b}{h}, \quad \tilde{c} = \frac{\lambda c}{h}, \quad \tilde{p} = \frac{\lambda p}{h}, \quad \tilde{r} = \frac{\lambda r}{h}, \quad \tilde{w} = \frac{\lambda w}{h}. \quad (3.1)$$

To ease the notation, we hereafter omit the tildes from these cost parameters.

Let N be the steady-state number of orders at the supplier's manufacturing facility. If we assume for now that $\mu > \lambda$ (this point is revisited later), then N is geometrically distributed with mean ν^{-1} , where

$$\nu = \frac{\mu - \lambda}{\lambda}. \quad (3.2)$$

This parameter, which represents the normalized excess capacity, is the supplier's decision variable in our analysis, and we often refer to it simply as capacity. To simplify our analysis, we assume that N is a continuous random variable, and replace the geometric distribution by an exponential distribution with parameter ν . This continuous state approximation can be justified by a heavy traffic approximation (e.g., §10 of Harrison 1988), and leads to slightly different quantitative results (the approximation tends to underestimate the optimal discrete base stock level). However, it has no effect on the qualitative system behavior, which is the object of our study.

Because the revenues for each agent are fixed, profit maximization and cost minimization lead to the same solution. We employ profit maximization to explicitly incorporate the agents' *participation constraints*, which take the form of nonnegative expected profits. However, we introduce some variable cost notation (C_R and C_S) in equations (3.3)-(3.4) for future reference when discussing the inefficiency of the Nash solution (§3.5) and contracts that coordinate the system (§3.6). In these equations, the quantities $r - w$ and $w - c - p$ are independent of the supply chain decisions (the total normalized capacity cost is $\frac{\mu}{\lambda} = c(1 + \nu)$, where c is the capacity cost if no excess capacity is built, and $c\nu$ is the cost of excess capacity) and represent fixed profits for the respective agents. The steady-state expected normalized profit per unit time for the risk-neutral retailer (Π_R) and supplier (Π_S) in terms of the two decision variables are given by

$$\begin{aligned} \Pi_R(s, \nu) &= r - w - C_R(s, \nu) \\ &= r - w - E[(s - N)^+] - \alpha b E[(N - s)^+] \\ &= r - w - s + \frac{1 - e^{-\nu s}}{\nu} - \alpha b \frac{e^{-\nu s}}{\nu}, \end{aligned} \quad (3.3)$$

and

$$\Pi_S(s, \nu) = w - p - c - C_S(s, \nu) \quad (3.4)$$

$$\begin{aligned}
&= w - p - c - (1 - \alpha) b E[(N - s)^+] - c\nu \\
&= w - p - c(1 + \nu) - (1 - \alpha) b \frac{e^{-\nu s}}{\nu}.
\end{aligned}$$

3.3 The Centralized Solution

As a reference point for the efficiency of the two-agent system, we start by finding the optimal solution to the centralized version of the problem, where there is a single decision maker that simultaneously optimizes the base stock level s and the normalized excess capacity ν . The steady-state expected normalized profit per unit time Π (defined in terms of the total variable cost $C = C_R + C_S$) for this decision maker is

$$\begin{aligned}
\Pi(s, \nu) &= \Pi_R(s, \nu) + \Pi_S(s, \nu) \\
&= r - p - c - C(s, \nu) \\
&= r - p - c(1 + \nu) - s + \frac{1 - (b + 1)e^{-\nu s}}{\nu}.
\end{aligned} \tag{3.5}$$

The centralized solution is given in Proposition 3; see the Appendix for the proof.

Proposition 3 *If $r - p - c \geq 2\sqrt{c \ln(1 + b)}$, then the optimal centralized solution is the unique solution to the first-order conditions*

$$\frac{\partial \Pi(s, \nu)}{\partial s} = 0 \iff \nu s = \ln(1 + b), \tag{3.6}$$

$$\frac{\partial \Pi(s, \nu)}{\partial \nu} = 0 \iff -(b + 1)(\nu s + 1) \frac{e^{-\nu s}}{\nu^2} + \frac{1}{\nu^2} + c = 0, \tag{3.7}$$

and is given by

$$\nu^* = \sqrt{\frac{\ln(1 + b)}{c}} \quad \text{and} \quad s^* = \sqrt{c \ln(1 + b)}. \tag{3.8}$$

The resulting profit is

$$\Pi(s^*, \nu^*) = r - p - c - 2\sqrt{c \ln(1 + b)}. \tag{3.9}$$

If $r - p - c < 2\sqrt{c \ln(1 + b)}$, then the system generates negative profits and the optimal centralized solution is to not operate the supply chain.

By relation (3.6), the ratio of the base stock level, s , to the supplier's mean queue length, ν^{-1} , satisfies $\nu s = \ln(1 + b)$ at optimality, which corresponds to a Pareto

frontier for the selection of s and ν . (The corresponding first-order conditions for the discrete inventory problem is $\ln(\nu + 1)s = \ln(1 + b)$, and so our continuous approximation can be viewed as using the Taylor series approximation $\ln(\nu + 1) \approx \nu$.) Although this ratio is independent of the capacity cost c , the optimal point on this Pareto frontier depends on c via $s = \nu c$ according to (3.8).

As expected, the optimal capacity level decreases with the capacity cost and increases with the backorder-to-holding cost ratio b . Similarly, because capacity and safety stock provide alternative means to avoid backorders, the optimal base stock level increases with the capacity cost and with the normalized backorder cost b . Finally, as expected, neither w nor α play any role in this single-agent optimization, because transfer payments between the retailer and the supplier do not affect the centralized profit.

3.4 The Nash Solution

Under the Nash equilibrium concept, the retailer chooses s to maximize $\Pi_R(s, \nu)$, assuming that the supplier chooses ν to maximize $\Pi_S(s, \nu)$; likewise, the supplier simultaneously chooses ν to maximize $\Pi_S(s, \nu)$ assuming the retailer chooses s to maximize $\Pi_R(s, \nu)$. Because each agent's strategy is a best response to the other's, neither agent is motivated to depart from this equilibrium.

Our results are most easily presented by deriving the Nash equilibrium in the absence of participation constraints, which is done in the next proposition, and then incorporating the participation constraints, $\Pi_R \geq 0$ and $\Pi_S \geq 0$. In anticipation of subsequent analysis, we express the Nash equilibrium in terms of the backorder allocation fraction α . Let us also define the auxiliary function

$$f_\alpha(b) = \sqrt{\frac{(1 - \alpha)b(\ln(1 + \alpha b) + 1)}{(1 + \alpha b)\ln(1 + b)}}, \quad (3.10)$$

which plays a prominent role in our analysis.

Proposition 4 *In the absence of participation constraints, the unique Nash equilibrium is*

$$\nu_\alpha^* = f_\alpha(b)\nu^*, \quad (3.11)$$

$$s_\alpha^* = \left(\frac{\ln(1 + \alpha b)}{\ln(1 + b)f_\alpha(b)} \right) s^*. \quad (3.12)$$

The resulting profits, $\Pi_R^*(\alpha)$ and $\Pi_S^*(\alpha)$, are

$$\Pi_R^*(\alpha) = \Pi_R(s_\alpha^*, \nu_\alpha^*) = r - w - s_\alpha^*, \quad (3.13)$$

$$\Pi_S^*(\alpha) = \Pi_S(s_\alpha^*, \nu_\alpha^*) = w - p - c - \left(\frac{\ln(1 + \alpha b) + 2}{\ln(1 + \alpha b) + 1} \right) c\nu_\alpha^*. \quad (3.14)$$

Proof: Let $s^*(\nu)$ be the retailer's reaction curve, i.e., the optimal base stock level given a capacity ν installed by the supplier. Because (3.3) is concave in s , $s^*(\nu)$ is characterized by the first-order condition

$$\nu s^*(\nu) = \ln(1 + \alpha b). \quad (3.15)$$

Using a similar argument, we find that the supplier's reaction curve $\nu^*(s)$ satisfies

$$e^{-\nu^*(s)s} \left(\frac{\nu^*(s)s + 1}{(\nu^*(s)s)^2} \right) = \frac{c}{(1 - \alpha)bs^2}. \quad (3.16)$$

The unique solution to (3.15)-(3.16) is (3.11)-(3.12), and substituting this solution into (3.3)-(3.4) yields (3.13)-(3.14). ■

Because $f_\alpha(b)$ is decreasing in α and $\ln(1 + \alpha b)$ is increasing in α for $b > 0$, it follows that as α increases, the retailer becomes more concerned with backorders and increases his base stock level, while the supplier cares less about backorders and builds less excess capacity.

As mentioned previously, we assume that the two agents do not participate in the game unless their expected normalized profits in (3.13)-(3.14) are nonnegative. Hence, if either of these profits are negative, the Nash equilibrium (in the presence of participation constraints) is an inoperative supply chain. The remainder of this section is devoted to an analysis of these profits as a function of α . The supplier's profit $\Pi_S^*(\alpha)$ is an increasing function of α that satisfies

$$\Pi_S^*(0) = w - p - c - 2\sqrt{bc}, \quad \Pi_S^*(1) = w - p - c,$$

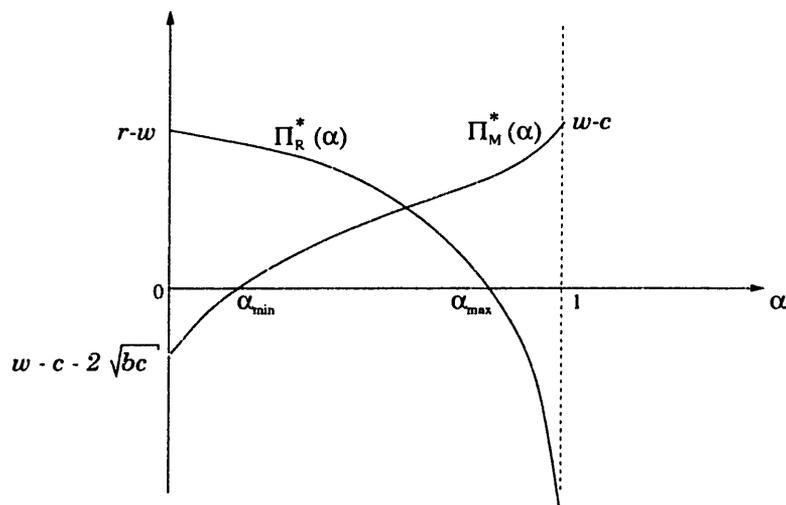


Figure 3.1: The retailer's (Π_R) and supplier's (Π_S) profits in the Nash equilibrium as a function of the backorder allocation fraction α .

and $\Pi_R^*(\alpha)$ is a decreasing function of α that satisfies

$$\Pi_R^*(0) = r - w, \quad \lim_{\alpha \rightarrow 1} \Pi_R^*(\alpha) \rightarrow -\infty, \quad (3.17)$$

as shown in Figure 3.1. (Many of the limits taken in this work, e.g., $\alpha \rightarrow 1$, are implicitly taken to be one-sided.)

To understand the unbounded retailer losses in (3.17), note that for the extreme case $\alpha = 1$, the supplier does not face any backorder cost and consequently has no incentive to build excess capacity, i.e., $\nu_1^* = 0$. This corresponds to the null recurrent case of a queueing system with an arrival rate equal to its service rate, and $s_1^* = \infty$: There is no base stock level that allows the retailer to maintain finite inventory (backorder plus holding) costs. Hence, this production-inventory system is unstable when $\alpha = 1$ and the Nash equilibrium is that the retailer does not participate, and the supply chain does not operate.

More generally, there exist α_{\min} and α_{\max} such that $\Pi_R^*(\alpha) \geq 0$ and $\Pi_S^*(\alpha) \geq 0$ if and only if $\alpha \in [\alpha_{\min}, \alpha_{\max}]$. That is, the Nash equilibrium is an inoperative supply chain when $\alpha < \alpha_{\min}$ or $\alpha > \alpha_{\max}$, because one of the agents is burdened with too much of the backorder cost. The threshold $\alpha_{\max} \in (0, 1)$, and solves

$$r - w = \left(\frac{\ln(1 + \alpha b)}{\ln(1 + b) f_\alpha(b)} \right) s^*. \quad (3.18)$$

If $w - p - c \geq 2\sqrt{bc}$, then $\alpha_{\min} = 0$. Otherwise, α_{\min} solves

$$w - p - c = \left(\frac{\ln(1 + \alpha b) + 2}{\ln(1 + \alpha b) + 1} \right) f_{\alpha}(b) s^*. \quad (3.19)$$

We have been unable to explicitly solve for α_{\max} and α_{\min} in (3.18)-(3.19). However, to increase our understanding of these two equations, we investigate the solution in two extreme cases: When backorders are much less costly than holding inventory ($b \ll 1$), we have

$$\alpha_{\max} \approx \frac{\sqrt{(r-w)^4 + 4(r-w)^2 bc} - (r-w)^2}{2bc}, \quad \alpha_{\min} \approx 1 - \frac{(w-p-c)^2}{4bc}. \quad (3.20)$$

When backorder costs are very large,

$$\alpha_{\max} \approx \frac{(r-w)^2}{\ln(1+b)c + (r-w)^2}, \quad \alpha_{\min} \approx \frac{\ln(1+b)c}{\ln(1+b)c + (w-p-c)^2}. \quad (3.21)$$

Even under the assumption $r > w > c + p$, it is possible that $\alpha_{\min} > \alpha_{\max}$ in (3.20)-(3.21). In this situation, even though each agent is willing to participate for some values of α , it is not possible for the retailer and supplier to simultaneously earn nonnegative profits.

3.5 Comparison of Solutions

In this section, we compare the centralized solution and the Nash equilibrium with respect to the total system profit, the agents' decisions, and the consumers of the product.

The Nash equilibrium is inefficient. As in §3.4, it is convenient to first quantify the inefficiency of the Nash equilibrium in the absence of participation constraints, and then to incorporate them later. In the absence of participation constraints, the centralized solution is not achievable as a Nash equilibrium. By equations (3.6) and (3.15), the first-order conditions are $\nu s^* = \ln(1+b)$ in the centralized solution and $\nu s^* = \ln(1+\alpha b)$ in the Nash solution. Hence, the two solutions are not equal when $\alpha < 1$, and the Nash equilibrium in the $\alpha = 1$ case is an unstable system, as discussed earlier.

The magnitude of the inefficiency of a Nash equilibrium is typically quantified by

comparing the profits under the centralized and Nash solutions. Because the profits $r - w$ and $w - p - c$ are fixed in (3.3)-(3.4), it is more natural to restrict ourselves to the variable costs. This assumption also allows us to follow Cachon and Zipkin and compute the *competition penalty*, which is defined as the percentage increase in variable cost of the Nash equilibrium over the centralized solution. By (3.5) and (3.8), the variable cost for the centralized solution is

$$C(s^*, \nu^*) = \left(s^* - \frac{1 - e^{-\nu^* s^*}}{\nu^*} \right) + b \frac{e^{-\nu^* s^*}}{\nu^*} + c\nu^* = 2\sqrt{c \ln(1 + b)},$$

and the variable cost C_α^* associated with the Nash equilibrium in the absence of participation constraints is, by (3.3)-(3.4) and Proposition 4,

$$C_\alpha^* = C_R(s_\alpha^*, \nu_\alpha^*) + C_S(s_\alpha^*, \nu_\alpha^*) = \left[f_\alpha(b) \left(\frac{\ln(1 + \alpha b) + 2}{\ln(1 + \alpha b) + 1} \right) + \frac{\ln(1 + \alpha b)}{\ln(1 + b) f_\alpha(b)} \right] s^*.$$

Hence, the competition penalty in the absence of participation constraints is $\frac{C_\alpha^* - C(s^*, \nu^*)}{C(s^*, \nu^*)} \times 100\%$, where

$$\frac{C_\alpha^* - C(s^*, \nu^*)}{C(s^*, \nu^*)} = \frac{1}{2} \left[f_\alpha(b) \left(\frac{\ln(1 + \alpha b) + 2}{\ln(1 + \alpha b) + 1} \right) + \frac{\ln(1 + \alpha b)}{\ln(1 + b) f_\alpha(b)} \right] - 1. \quad (3.22)$$

Surprisingly, the competition penalty in (3.22) is independent of the supplier's cost of capacity. This occurs because the centralized variable cost and the Nash variable cost are both proportional to \sqrt{c} at optimality, which is a consequence of the particular functional form arising from the make-to-stock formulation. However, this penalty is a function of α and b , and we can simplify equation (3.22) for the limiting values of these two parameters. The function $f_\alpha(b)$ is decreasing in α and $f_1(b) = 0$. Hence, the competition penalty goes to ∞ as $\alpha \rightarrow 1$. This inefficiency occurs because as the retailer bears more of the backorder cost, the supplier builds less excess capacity, and in the limit the lack of excess capacity causes instability of the $M/M/1$ system. At the other extreme, $f_\alpha(b) \rightarrow \sqrt{\frac{b}{\ln(1+b)}}$ as $\alpha \rightarrow 0$, and the competition penalty in this case is given by

$$\sqrt{\frac{b}{\ln(1+b)}} - 1 \quad \text{for } b > 0. \quad (3.23)$$

This function is increasing and concave in b , approaches zero as $b \rightarrow 0$ and grows

to ∞ as $b \rightarrow \infty$. Hence, when the supplier incurs most of the backorder cost, the retailer holds very little inventory and the competition penalty depends primarily on the backorder cost; if this cost is low then the supplier has little incentive to build excess capacity, which leads to a small competition penalty because the centralized planner holds neither safety stock nor excess capacity in this case. In contrast, if the backorder cost is very high, the supplier cannot overcome the retailer's lack of safety stock, and his backorders get out of control, leading to high inefficiency.

Turning to the backorder cost asymptotics, $f_\alpha(b) \rightarrow \sqrt{\frac{1-\alpha}{\alpha}}$ as $b \rightarrow \infty$, and the competition penalty approaches

$$\frac{1}{2\sqrt{\alpha(1-\alpha)}} - 1. \quad (3.24)$$

This quantity vanishes at $\alpha = 0.5$, is symmetric about $\alpha = 0.5$, is convex for $\alpha \in (0, 1)$, and approaches ∞ as $\alpha \rightarrow 0$ and $\alpha \rightarrow 1$. Thus, when backorders are very expensive, this cost component dominates both agents' objective functions when they care equally about backorders ($\alpha = 0.5$), and their cost functions – and hence decisions – coincide with the centralized solution. However, when there is a severe imbalance in the backorder allocation (α is near 0 or 1), one of the agents does not build enough of his buffer resource, and the other agent cannot prevent many costly backorders, which is highly inefficient from the viewpoint of the entire supply chain. Finally, for the case $b \rightarrow 0$, the competition penalty is given by

$$\frac{2-\alpha}{2\sqrt{1-\alpha}} - 1, \quad (3.25)$$

which is an increasing convex function of α . Consistent with the previous analyses, this penalty function vanishes as $\alpha \rightarrow 0$ and approaches ∞ as $\alpha \rightarrow 1$.

In summary, there are two regimes, $(\alpha = 0.5, b \rightarrow \infty)$ and $(\alpha \rightarrow 0, b \rightarrow 0)$, where the Nash equilibrium is asymptotically efficient, and two regimes, $\alpha \rightarrow 1$ and $(\alpha \rightarrow 0, b \rightarrow \infty)$, where the inefficiency of the Nash solution is arbitrarily large. However, because equation (3.22) does not consider the agents' participation constraints, some of the large inefficiencies in the latter regimes are not attainable by the supply chain.

To complement these asymptotic results, we compute in Table 3.1 the competition penalty in (3.22) for various values of α and b . Our asymptotic results agree with the numbers around the four edges of this table. Two new insights emerge from Table 3.1. First, the competition penalty is minimized by α near 0.5 (i.e., the backorder

b	α						
	0	0.1	0.3	0.5	0.7	0.9	1
10^{-3}	0.0%	0.2%	1.6%	6.1%	18.7%	73.9%	$\infty\%$
1	20.1%	14.8%	7.8%	5.9%	12.9%	59.5%	$\infty\%$
2	34.9%	23.8%	10.7%	5.8%	11.1%	56.1%	$\infty\%$
3	47.1%	30.1%	12.3%	5.6%	10.2%	54.8%	$\infty\%$
4	57.6%	34.9%	13.3%	5.4%	9.6%	54.1%	$\infty\%$
5	67.0%	38.7%	14.0%	5.3%	9.2%	53.8%	$\infty\%$
6	75.6%	41.7%	14.5%	5.2%	9.0%	53.6%	$\infty\%$
7	83.5%	44.3%	14.9%	5.1%	8.8%	53.5%	$\infty\%$
8	90.8%	46.4%	15.2%	5.0%	8.7%	53.5%	$\infty\%$
9	97.7%	48.3%	15.4%	4.9%	8.5%	53.5%	$\infty\%$
10	104.2%	49.9%	15.6%	4.8%	8.4%	53.5%	$\infty\%$
10^{200}	$4.6 \times 10^{100}\%$	66.7%	9.2%	0.0%	9.1%	66.5%	$\infty\%$

Table 3.1: The competition penalty in (3.22) in the absence of participation constraints for different values of the backorder-to-holding cost ratio b and the backorder allocation fraction α .

cost is split evenly) when $b \geq 1$. Second, the competition penalty appears to be an increasing function of b for $\alpha < 0.5$, and a U-shaped function of b for $\alpha > 0.5$.

Finally, we note that *double marginalization* (as introduced by Spengler 1950 or described by Cachon 1999) is not the source of the inefficiency in our model because demand is fully backlogged and therefore is independent of the agents' decisions. The negative externality in our model is due to the fact that the centralized planner balances the costs of backorders, safety stock and production capacity, whereas the agents in the decentralized model – by behaving selfishly – do not fully incorporate the impact of their decisions on the entire supply chain cost.

Comparison of decision variables. Figure 3.2 depicts the optimal Nash production capacity ν_α^* and the optimal Nash base stock level s_α^* as a function of α , and allows us to compare these functions to the centralized solutions, ν^* and s^* . Excess capacity and the base stock level are alternative ways for the supplier and retailer, respectively, to buffer against demand uncertainty, and Figure 3.2 shows that the inefficiency of the Nash solution does not necessarily imply that these agents have less buffer resources in the Nash solution than in the centralized solution. For both decision variables, there exist thresholds on the value of α , denoted by α_s and α_ν in

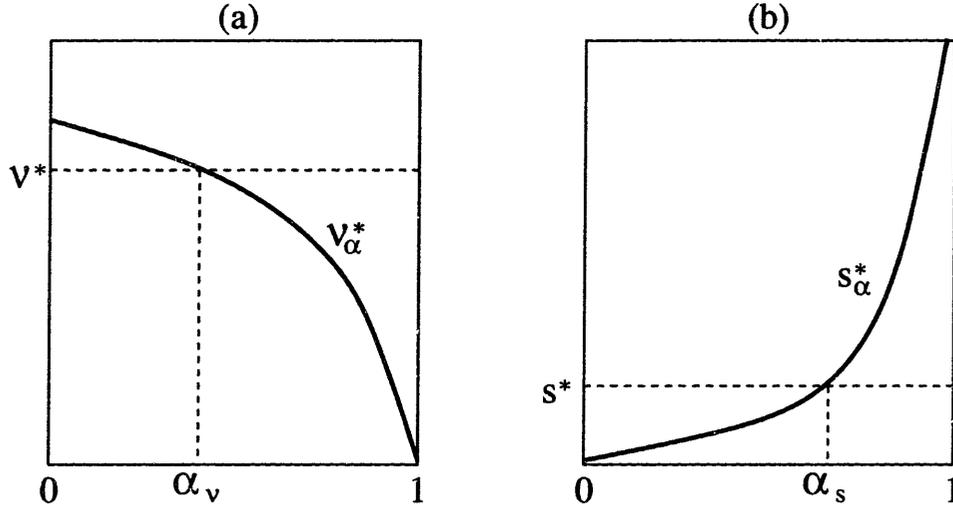


Figure 3.2: The optimal Nash production capacity (ν_α^*) and the optimal Nash base stock level (s_α^*) as a function of the backorder allocation fraction α . The centralized solutions are ν^* and s^* .

Figure 3.2, that divide the regions where the agents have more or less buffer resources than the optimal centralized solution. However, as shown in the next proposition, at least one agent in the Nash equilibrium possesses less of his buffer resource than the central planner.

Proposition 5 For α_s and α_ν defined in Figure 3.2, we have $\alpha_s > \alpha_\nu$.

Proof: By Figure 3.2, if $\alpha_s \leq \alpha_\nu$ then there exists $\hat{\alpha} \in [\alpha_s, \alpha_\nu]$ such that $\nu_{\hat{\alpha}}^* s_{\hat{\alpha}}^* \geq \nu^* s^*$. However, this inequality together with (3.6) and (3.15) implies that $f_{\hat{\alpha}}(b) = \ln(1+b)$, i.e., $\hat{\alpha} = 1$. But for $\alpha = 1$ the supply chain is unstable and does not operate. Hence, $\nu_\alpha^* s_\alpha^* < \nu^* s^*$ for $\alpha \in [0, 1)$, and consequently $\alpha_s > \alpha_\nu$. ■

We cannot solve for α_s and α_ν in closed form, except when b takes on a limiting value. By (3.12), α_s satisfies

$$\frac{\ln(1+\alpha b)}{\ln(1+b)f_\alpha(b)} = 1. \quad (3.26)$$

As $b \rightarrow 0$, we have $f_\alpha(b) \rightarrow \sqrt{1-\alpha}$ and $\frac{\ln(1+\alpha b)}{\ln(1+b)} \rightarrow \alpha$. Therefore, as $b \rightarrow 0$, α_s satisfies $\frac{\alpha}{\sqrt{1-\alpha}} = 1$, or $\alpha_s = \frac{\sqrt{5}-1}{2} \approx 0.618$, which is the inverse of the *golden-section number* that arises in a variety of disciplines (e.g., Vajda 1989). As $b \rightarrow \infty$, we have $f_\alpha(b) \rightarrow \sqrt{\frac{1-\alpha}{\alpha}}$, and $\frac{\ln(1+\alpha b)}{\ln(1+b)} \rightarrow 1$. In this case, α_s satisfies $\sqrt{\frac{\alpha}{1-\alpha}} = 1$, or $\alpha_s = 0.5$. Numerical computations reveal that α_s is unimodal in b , achieving a maximum of

0.627 at $b = 1.48$, and is rather insensitive to moderate values of b (e.g., $\alpha_s \geq 0.61$ for $b \in [1, 10]$). In other words, for moderate values of the backorder cost, the Nash retailer holds more inventory than optimal when his share of the backorder cost is more than about 61%.

By (3.11), α_ν solves $f_\alpha(b) = 1$; i.e., when the backorder cost is small, the Nash supplier holds a less-than-optimal level of capacity. As $b \rightarrow 0$, this condition becomes $\sqrt{1-\alpha} = 1$, which gives $\alpha_\nu = 0$. As $b \rightarrow \infty$, the condition becomes $\sqrt{\frac{1-\alpha}{\alpha}} = 1$, which is solved by $\alpha_\nu = 0.5$. Note that $\alpha_s = \alpha_\nu = 0.5$ as $b \rightarrow \infty$ is consistent with our previous claim that the Nash equilibrium is asymptotically efficient in the regime $(\alpha = 0.5, b \rightarrow \infty)$. A numerical study reveals that α_ν is more sensitive than α_s to the value of b . As b varies from 1 to 10, α_ν ranges from 0.28 to 0.49.

Customer service level. The exponential distribution of the queue length implies that the steady-state probability that a customer is forced to wait because of retailer shortages is equal to $Pr(N \geq s) = e^{-\nu s}$; consequently, we refer to $(1 - e^{-\nu s}) \times 100\%$ as the *service level*. By equations (3.6) and (3.15), the stockout probability $e^{-\nu s}$ equals $(1 + b)^{-1}$ in the centralized solution and $(1 + \alpha b)^{-1}$ in the Nash solution. Hence, customers receive better service in the centralized solution than in the Nash equilibrium. This is because (see Figure 2) the product of the two buffer resources (normalized excess capacity and base stock level) is always smaller in the Nash equilibrium than in the centralized solution, and the customers suffer from this less-than-optimal level of collective buffer resource; this degradation in customer service in decentralized systems also occurs in the traditional bilateral monopoly model (e.g., Tirole 1997), where double marginalization leads to a higher price charged to the customer and less goods sold. Finally, even though the system is not stable for $\alpha = 1$, customers generally desire a larger value of α ; i.e., they prefer that the penalty for shortages be absorbed primarily by the agent in direct contact with them.

3.6 Contracts

We showed in §3.5 that the Nash equilibrium is always inefficient when the supply chain operates. In this section, we construct a coordinating contract that specifies linear transfer payments based on retailer inventory and backorder levels, the capacity level and the cost parameters. As in our earlier analysis, this information is assumed to be common knowledge; although information sharing among entities in a supply chain is now common practice, we are not aware of any contracts in use that are similar

to the one proposed here. Cachon and Zipkin also use a linear transfer payment based on inventory levels to coordinate their supply chain, and readers are referred to §1.5 of Cachon for a survey of alternative types of contracts in the multi-echelon inventory setting. We do not impose an explicit constraint that forces either agent to build a predefined level of its buffer resource. Using Cachon and Lariviere's (1997) terminology, we assume a *voluntary compliance* regime, where both the retailer and the supplier choose their buffer resource levels to maximize their own profits.

Although we have used profit maximization thus far, because the revenues are fixed our presentation of the contract analysis is simpler – and perhaps more natural – in the setting of variable cost minimization. Consequently, we first present the contract in the absence of participation constraints, and at the end of this section we incorporate the revenues via the agents' participation constraints. Because the model contains three cost components, the most general linear transfer payment (without loss of generality, the payment is from the supplier to the retailer) contains cost coefficients for the holding cost, backorder cost and capacity cost, which are denoted by γ_h , γ_b and γ_c , respectively. The steady-state expected normalized (recall from (3.1) that $h = 1$) transfer payment per unit time is given by

$$T(s, \nu) = \gamma_h \left(E[(s - N)^+] \right) + \gamma_b \left(bE[(N - s)^+] \right) + \gamma_c c\nu. \quad (3.27)$$

This transfer payment modifies the profit functions in (3.3)-(3.4) for the retailer and supplier, respectively, to

$$\tilde{C}_R(s, \nu) = C_R(s, \nu) - T(s, \nu) = (1 - \gamma_h) \left(s - \frac{1 - e^{-\nu s}}{\nu} \right) + b(\alpha - \gamma_b) \frac{e^{-\nu s}}{\nu} - \gamma_c c\nu, \quad (3.28)$$

$$\tilde{C}_S(s, \nu) = C_S(s, \nu) + T(s, \nu) = \gamma_h \left(s - \frac{1 - e^{-\nu s}}{\nu} \right) + b(1 - \alpha + \gamma_b) \frac{e^{-\nu s}}{\nu} + (1 + \gamma_c) c\nu. \quad (3.29)$$

The following proposition provides a general result for the coordination of static games with additive utility structures, and may be applicable to other supply chain problems.

Proposition 6 *Consider a static game with n players, where player i has action space X_i for $i = 1, \dots, n$. For any action $x \in X_1 \times X_2 \times \dots \times X_n$, the utility of player i is given by $u_i(x)$. The utility function for the centralized planner problem is $\sum_{i=1}^n u_i(x)$, and let $x^* \in \operatorname{argmax}_x \sum_{i=1}^n u_i(x)$ be an optimal centralized solution. For*

$i \neq j$, let $T_{ij}(x)$ be a linear transfer payment that player i pays to player j if the action is x . If $T_{ij}(x) = \gamma_j u_i(x)$, where $\sum_{j=1}^n \gamma_j = 1$, then x^* is a Nash equilibrium of the modified static game in which player i 's utility function is given by

$$\tilde{u}_i(x) = u_i(x) - \sum_{j \neq i} T_{ij}(x) + \sum_{j \neq i} T_{ji}(x).$$

Proof: By our assumptions on $T_{ij}(x)$, the modified utility function for player i can be written as

$$\tilde{u}_i(x) = \left(1 - \sum_{j \neq i} \gamma_j\right) u_i(x) + \gamma_i \sum_{j \neq i} u_j(x) \quad (3.30)$$

$$= \gamma_i \sum_{j=1}^N u_j(x). \quad (3.31)$$

Hence, $\sum_{i=1}^n \tilde{u}_i(x) = \sum_{i=1}^n u_i(x)$, and x^* is a Nash equilibrium of the modified game.

■

Applying Proposition 6 to equations (3.28)-(3.29) shows that coordination in the absence of participation constraints can be achieved if

$$-\gamma_c = \alpha - \gamma_b = 1 - \gamma_h. \quad (3.32)$$

A comparison of equations (3.3)-(3.4) and (3.28)-(3.29) implies that the modified cost functions are given by

$$\tilde{C}_R(s, \nu) = (1 - \gamma_h)C(s, \nu), \quad \tilde{C}_S(s, \nu) = \gamma_h C(s, \nu). \quad (3.33)$$

This is a consequence of the more general result in (3.31), which shows that any split of the total profit is possible by selecting appropriate values of the $\{\gamma_i\}$ parameters. Note that γ_h need not be in the interval $[0, 1]$.

Although we appear to have a degree of freedom in splitting the profits via γ_h , two conditions must be met by γ_h to guarantee that both agents will enter into the contract. First, both agents must be better off under the Nash equilibrium with the transfer payments than under the Nash equilibrium without the transfer payments, i.e.,

$$C_R(s_\alpha^*, \nu_\alpha^*) \geq (1 - \gamma_h)C(s, \nu),$$

$$C_S(s_\alpha^*, \nu_\alpha^*) \geq \gamma_h C(s, \nu).$$

This condition can be rewritten as $\gamma_h \in [\underline{\gamma}_\alpha(b), \bar{\gamma}_\alpha(b)]$, where

$$\begin{aligned} \underline{\gamma}_\alpha(b) &= 1 - \frac{\ln(1 + \alpha b)}{2 \ln(1 + b) f_\alpha(b)}, \\ \bar{\gamma}_\alpha(b) &= \frac{f_\alpha(b)}{2} \left(\frac{\ln(1 + \alpha b) + 2}{\ln(1 + \alpha b) + 1} \right). \end{aligned}$$

The second condition on γ_h requires that both agents achieve a nonnegative profit. The resulting inequalities can be calculated using equations (3.3)-(3.4), (3.9) and (3.28)-(3.29). Combining these two conditions, we can characterize the range of coordinating contracts that are attractive to both agents as

$$\max \left\{ 1 - \frac{r - w}{2\sqrt{c \ln(1 + b)}}, \underline{\gamma}_\alpha(b) \right\} \leq \gamma_h \leq \min \left\{ \frac{w - p - c}{2\sqrt{c \ln(1 + b)}}, \bar{\gamma}_\alpha(b) \right\}. \quad (3.34)$$

If condition (3.34) is satisfied, we say that the *system can be coordinated* (by the contract), and the remainder of this section is devoted to an analysis of this condition. First we note that it is always possible to coordinate the system if both agents are willing to participate in the Nash equilibrium in §3.4. This conclusion stems from the fact that if both players are willing to participate in the Nash equilibrium, then the additional profits from the centralized solution can be split so that each agent is still willing to participate and is at least as well off as in the Nash solution. Hence, the nontrivial cases to analyze are $\alpha \in [0, \alpha_{\min})$ and $\alpha \in (\alpha_{\max}, 1]$; recall that these two thresholds characterize the participation constraints in the Nash equilibrium and are defined in (3.18)-(3.19).

We can analyze (3.34) when α approaches one of its extreme values. As $\alpha \rightarrow 0$, we have

$$\underline{\gamma}_\alpha(b) \rightarrow 1 \quad \text{and} \quad \bar{\gamma}_\alpha(b) \rightarrow \sqrt{\frac{b}{\ln(1 + b)}}. \quad (3.35)$$

Thus, because $\alpha \rightarrow 0$ implies that $\alpha \leq \frac{r-w}{\sqrt{bc}}$, equation (3.34) reduces to

$$1 \leq \gamma_h \leq \min \left\{ \frac{w - p - c}{2\sqrt{c \ln(1 + b)}}, \sqrt{\frac{b}{\ln(1 + b)}} \right\}. \quad (3.36)$$

If $w - p - c \geq 2\sqrt{bc}$, then the second term in the brackets in (3.36) achieves the minimum, and the interval in (3.36) is nonempty because $b \geq \ln(1 + b)$ for $b > 0$. Because $f_\alpha(b) \rightarrow \sqrt{\frac{b}{\ln(1+b)}}$ as $\alpha \rightarrow 0$, the inequality $w - p - c \geq 2\sqrt{bc}$ can be recognized as the supplier's participation constraint in (3.19) in the $\alpha \rightarrow 0$ case. If, on the other hand, $w - p - c < 2\sqrt{bc}$, then (3.36) reduces to

$$1 \leq \gamma_h \leq \frac{w - p - c}{2\sqrt{c \ln(1 + b)}}. \quad (3.37)$$

By (3.9), the interval in (3.37) is nonempty if and only if $\Pi(s^*, \nu^*) \geq r - w$.

To summarize, when the supplier absorbs almost all of the backorder cost (i.e., $\alpha \rightarrow 0$), a coordinating contract is always possible if the supplier is willing to participate in the absence of the contract (i.e., $w - p - c \geq 2\sqrt{bc}$). If the centralized system profit $\Pi(s^*, \nu^*)$ is less than the retailer's fixed profit, $r - w$, then the supplier will not enter into the contract. Most interesting is the intermediate case, $w - p - c \in (2\sqrt{c \ln(1 + b)}, 2\sqrt{bc})$, where the system profit is bigger than the retailer's fixed profit, but the supplier is unwilling to participate in the absence of a contract. Here, the excess system profit enables the contract to entice the supplier to participate in the supply chain. By (3.34), the contract in the $\alpha \rightarrow 0$ case has $\gamma_h \geq 1$, and the supplier subsidizes the retailer's entire operation; i.e., the only way to coordinate a system in which the retailer does not incur backorder costs is for the supplier to give the retailer the inventory on consignment.

Turning to the $\alpha \rightarrow 1$ case, we have $f_\alpha(b) \rightarrow 0$ and therefore $\underline{\gamma}_\alpha(b) \rightarrow -\infty$ and $\bar{\gamma}_\alpha(b) \rightarrow 0$. Thus, condition (3.34) reduces to

$$1 - \frac{r - w}{2\sqrt{c \ln(1 + b)}} \leq \gamma_h \leq 0, \quad (3.38)$$

Equation (3.9) shows that the interval in (3.38) is nonempty if and only if $\Pi(s^*, \nu^*) \geq w - p - c$. Hence, when the retailer incurs almost all of the backorder cost, the contract is attractive to both parties if and only if the centralized system profit exceeds the supplier's fixed profit, $w - p - c$. In this case, the contract coefficient $\gamma_h \leq 0$, and (3.33) implies that $\tilde{C}_S(s, \nu) \leq 0$; i.e., the retailer subsidizes the supplier's entire operation. Similar behavior, where where large manufacturers pay for their suppliers' capital equipment, has been observed in the automobile industry (e.g., Dyer and Ouchi 1993, Dyer et al. 1998).

The analysis is more difficult in the general case, $\alpha \in (0, \alpha_{\min}) \cup (\alpha_{\max}, 1)$. Here, we consider the extreme values of b . As $b \rightarrow \infty$, $\underline{\gamma}_\alpha(b) \rightarrow 1 - \frac{1}{2}\sqrt{\frac{\alpha}{1-\alpha}}$ and $\bar{\gamma}_\alpha(b) \rightarrow \frac{1}{2}\sqrt{\frac{1-\alpha}{\alpha}}$. Even though $\underline{\gamma}_\alpha(b) < \bar{\gamma}_\alpha(b)$ for $\alpha \in [0, 1]$ in this case (this can be derived with the change of variable $x = \sqrt{\frac{\alpha}{\alpha-1}}$), neither of these quantities are binding in (3.34), and coordination becomes impossible because the lower bound in (3.34) is at least 1 and the upper bound in (3.34) goes to 0. In contrast, as $b \rightarrow 0$, $\underline{\gamma}_\alpha(b)$ and $\bar{\gamma}_\alpha(b)$ are binding in (3.34), and $\underline{\gamma}_\alpha(b) \rightarrow 1 - \frac{\alpha}{\sqrt{1-\alpha}}$ and $\bar{\gamma}_\alpha(b) \rightarrow \sqrt{1-\alpha}$. Because $1 - \frac{\alpha}{\sqrt{1-\alpha}} \leq \sqrt{1-\alpha}$ for $\alpha \in [0, 1]$, coordination is always possible as $b \rightarrow 0$.

In conclusion, system coordination is most difficult when the backorder cost is incurred almost entirely by one agent (i.e., α is near 0 or 1). Furthermore, for a fixed extreme value of α , coordination is harder if backorders are costly. We can show (the proof is omitted, but the result follows from our analyses of the $\alpha \rightarrow 0$ and $\alpha \rightarrow 1$ cases) that a sufficient (but not necessary) condition for coordination is that the optimal centralized profit exceeds the fixed profits of both agents, $\Pi(s^*, \nu^*) \geq \max\{r - w, w - p - c\}$. Moreover, if we consider w as endogenous, then the likelihood of coordination is maximized by minimizing the right side of this inequality. This is achieved by $w = (r + p + c)/2$, which splits the fixed profits evenly. By (3.9), coordination is always possible in this case if $r - p - c \geq 4\sqrt{c \ln(1 + b)}$.

3.7 The Stackelberg Games

We conclude our study of this two-stage supply chain by considering the case where one agent dominates.

Supplier's Stackelberg game. When the supplier is the Stackelberg leader, he chooses ν to optimize $\Pi_S(s, \nu)$ in (3.4), given the retailer's best response, $s^*(\nu)$ in (3.15). This straightforward computation leads to the following proposition.

Proposition 7 *In the absence of participation constraints, the equilibrium in the supplier's Stackelberg game is*

$$\bar{s}_\alpha = s_\alpha^* \sqrt{1 + \ln(1 + \alpha b)}, \quad \bar{\nu}_\alpha = \frac{\nu_\alpha^*}{\sqrt{1 + \ln(1 + \alpha b)}}. \quad (3.39)$$

The agents' profits are

$$\Pi_S(\bar{s}_\alpha, \bar{\nu}_\alpha) = w - p - c - 2\sqrt{\frac{(1-\alpha)bc}{1+\alpha b}}, \quad \Pi_R(\bar{s}_\alpha, \bar{\nu}_\alpha) = r - w - \bar{s}_\alpha. \quad (3.40)$$

Equation (3.39) implies that $\bar{\nu}_\alpha \bar{s}_\alpha = \nu_\alpha^* s_\alpha^*$, and hence the customer service level is the same under the Stackelberg and Nash equilibria. Because the first-order conditions of the centralized problem dictate the service level, it also follows that the Stackelberg equilibrium is inefficient relative to the centralized solution. Not surprisingly, the supplier builds less capacity and the retailer holds more safety stock in (3.39) than in the Nash equilibrium. The discrepancy between the Stackelberg and Nash solutions increases as α and b increase.

Now we compare the profit of each agent and the entire system under the Nash and Stackelberg equilibria. By (3.14), the supplier's profit in the Nash equilibrium can be written as

$$\Pi_S(s_\alpha^*, \nu_\alpha^*) = w - p - c - 2\sqrt{\frac{(1-\alpha)bc}{1+\alpha b}} \left(\frac{\ln(1+\alpha b) + 2}{2\sqrt{\ln(1+\alpha b) + 1}} \right).$$

The function $\frac{x+2}{2\sqrt{x+1}}$ is strictly increasing in $[0, \infty)$, and is equal to 1 when $x = 0$. Thus, it is always the case that $\Pi_S(\bar{s}_\alpha, \bar{\nu}_\alpha) \geq \Pi_S(s_\alpha^*, \nu_\alpha^*)$; this is to be expected, because the supplier incorporates the retailer's best response when selecting his level of capacity. However, $\Pi_S(\bar{s}_\alpha, \bar{\nu}_\alpha) = \Pi_S(s_\alpha^*, \nu_\alpha^*)$ when $\alpha = 0$, $\alpha = 1$ or $b = 0$, and so the supplier does not benefit from being the leader in these extreme cases. When $\alpha = 1$ or $b = 0$, the supplier does not face any backorder costs and builds no excess capacity ($\nu = 0$). On the other hand, when $\alpha = 0$ the retailer – incurring no backorder costs – holds no safety stock ($s = 0$). Because these choices ($\nu = 0$ and $s = 0$) are independent of the bargaining power of the supplier in these cases, the Stackelberg and Nash equilibria provide the same utility to the supplier.

By (3.13) and (3.39), the difference in the retailer's profit between the Nash equilibrium and the Stackelberg equilibrium is

$$\Pi_R(s_\alpha^*, \nu_\alpha^*) - \Pi_R(\bar{s}_\alpha, \bar{\nu}_\alpha) = s_\alpha^* \left(\sqrt{1 + \ln(1 + \alpha b)} - 1 \right). \quad (3.41)$$

As expected, the retailer is worse off in the supplier's Stackelberg equilibrium than in the Nash equilibrium. By (3.8), (3.10), (3.12) and (3.41), the reduction in the retailer's profit from being the follower vanishes as $\alpha \rightarrow 0$ and $b \rightarrow 0$ (similar to the reasons given in the previous paragraph), and increases with α and b .

A comparison of the total system profit shows that $\Pi_R(s_\alpha^*, \nu_\alpha^*) + \Pi_S(s_\alpha^*, \nu_\alpha^*) \geq \Pi_R(\bar{s}_\alpha, \bar{\nu}_\alpha) + \Pi_S(\bar{s}_\alpha, \bar{\nu}_\alpha)$ if and only if

$$\frac{(1 + \alpha b) \ln(1 + \alpha b)}{(1 - \alpha)b} - \sqrt{1 + \ln(1 + \alpha b)} + 1 \geq 0. \quad (3.42)$$

Condition (3.42) holds for large values of α , but is not true in general. Because the left side of (3.42) equals zero when $\alpha = 0$, is increasing in $\alpha \geq \alpha_0$ if it is increasing in α at α_0 , and has a derivative with respect to α equal to $1 - \frac{b}{2}$ when $\alpha = 0$, we conclude that for $b \leq 2$ the Nash solution achieves a higher system profit than the Stackelberg equilibrium for any value of α . If $b > 2$, the Nash solution is more efficient if and only if $\alpha > \bar{\alpha}$, where $\bar{\alpha}$ is the unique positive value of α that solves (3.42) with equality. Hence, overall system performance suffers when the retailer incurs most of the expensive backorder costs, and – as the follower – has less power than in the Nash equilibrium to control these costs.

Now we turn to the participation constraints. We start with the follower (i.e., the retailer) because he performs the inner maximization in this game. By (3.8), (3.10), (3.12) and (3.39), we can express the optimal Stackelberg base stock level as

$$\bar{s}_\alpha = \ln(1 + \alpha b) \sqrt{\frac{(1 + \alpha b)c}{(1 - \alpha)b}}.$$

Hence, it follows from (3.40) that the retailer's participation constraint, $\Pi_R(\bar{s}_\alpha, \bar{\nu}_\alpha) \geq 0$, is equivalent to

$$(r - w) \sqrt{\frac{b}{c}} \geq \ln(1 + \alpha b) \sqrt{\frac{1 + \alpha b}{1 - \alpha}}. \quad (3.43)$$

The right side of (3.43) is increasing in $\alpha \geq 0$, and so there exists a threshold, call it $\bar{\alpha}_{\max}$, such that the participation constraint is satisfied if and only if $\alpha \leq \bar{\alpha}_{\max}$ (i.e., $\bar{\alpha}_{\max}$ solves (3.43) with equality). For future reference, let us also define the threshold

$$\bar{\alpha}_{\min} = \frac{4bc - (w - p - c)^2}{4bc + b(w - p - c)^2}. \quad (3.44)$$

The supplier's profit $\Pi_S(\bar{s}_\alpha, \bar{\nu}_\alpha)$ in (3.40) is nonnegative if and only if $\alpha \geq \bar{\alpha}_{\min}$.

There are three cases to examine: $\alpha > \bar{\alpha}_{\max}$, $\alpha \in [\bar{\alpha}_{\min}, \bar{\alpha}_{\max}]$ and $\alpha < \bar{\alpha}_{\min}$. If $\alpha > \bar{\alpha}_{\max}$ then the retailer incurs too much of the backorder costs and his participation constraint is violated. To avoid an inoperative supply chain, which would give zero

profits to both agents, the supplier – as the leader – has the luxury of selecting a capacity level that raises the retailer’s expected profit to zero. That is, by (3.15) the supplier chooses $\bar{v}_\alpha = \frac{\ln(1+\alpha b)}{r-w}$, and the retailer subsequently selects $\bar{s}_\alpha = r - w$. If the supplier’s resulting profit (see (3.4)) is nonnegative, i.e.,

$$w - p - c - \frac{(r-w)(1-\alpha)b}{(1+\alpha b)\ln(1+\alpha b)} - \frac{\ln(1+\alpha b)c}{r-w} \geq 0, \quad (3.45)$$

then this is the supplier’s Stackelberg equilibrium for the $\alpha > \bar{\alpha}_{\max}$. If (3.45) is violated, then the equilibrium is an inoperative supply chain.

In the second case, $\alpha \in [\bar{\alpha}_{\min}, \bar{\alpha}_{\max}]$, both participation constraints are satisfied and the Stackelberg equilibrium is given by (3.39). In the last case, $\alpha < \bar{\alpha}_{\min}$, the supplier’s high portion of the backorder costs prevents him from earning a nonnegative expected profit, and he decides not to participate. Finally, as in §3.4, it is possible that $\bar{\alpha}_{\min} > \bar{\alpha}_{\max}$; in this case, there is no value of α that simultaneously provides nonnegative profits for both agents. For brevity’s sake, we do not pursue asymptotics for $\bar{\alpha}_{\min}$ and $\bar{\alpha}_{\max}$.

Retailer’s Stackelberg game. The Stackelberg problem is less tractable when the retailer is the leader. However, the following proposition (see the Appendix for a proof) characterizes the solution.

Proposition 8 *Let $(\hat{v}_\alpha, \hat{s}_\alpha)$ be the equilibrium when the retailer is the Stackelberg leader in the absence of participation constraints. Define $\hat{\beta} \geq 0$ to be the unique nonnegative solution of*

$$\beta^2 + (\beta + 2) \left(1 - e^{-(\beta - \nu_\alpha^* s_\alpha^*)}\right) = 0. \quad (3.46)$$

Then the Stackelberg solution is

$$\hat{v}_\alpha = \sqrt{\frac{(1-\alpha)b(\hat{\beta}+1)e^{-\hat{\beta}}}{c}}, \quad \hat{s}_\alpha = \frac{\hat{\beta}}{\hat{v}_\alpha}. \quad (3.47)$$

Although we do not have a closed-form solution to the retailer’s Stackelberg game, the next proposition (see the Appendix for a proof) provides a comparison between this equilibrium and the Nash equilibrium.

Proposition 9 *The following five inequalities hold:*

$$\hat{\nu}_\alpha \hat{s}_\alpha \leq \nu_\alpha^* s_\alpha^*, \quad (3.48)$$

$$\hat{\nu}_\alpha \geq \nu_\alpha^*, \quad (3.49)$$

$$\hat{s}_\alpha \leq s_\alpha^*, \quad (3.50)$$

$$\Pi_R(\hat{s}_\alpha, \hat{\nu}_\alpha) \geq \Pi_R(s_\alpha^*, \nu_\alpha^*), \quad (3.51)$$

$$\Pi_S(\hat{s}_\alpha, \hat{\nu}_\alpha) \leq \Pi_S(s_\alpha^*, \nu_\alpha^*). \quad (3.52)$$

While inequalities (3.49)-(3.52) mirror our results for the supplier Stackelberg game, inequality (3.48) states that the customer service level, $(1 - e^{-\nu s}) \times 100\%$, is lower in the retailer's Stackelberg equilibrium than in the supplier's Stackelberg equilibrium (and the Nash equilibrium). Inequality (3.48) also implies that the retailer Stackelberg equilibrium is inefficient relative to the centralized solution. Analytical approximations (using $e^{-x} \approx 1 - x$ in (3.46)) and numerical computations reveal that when the service level is close to 0 or 100%, both Stackelberg games have asymptotically the same service level. The maximum difference is approximately 9.5%, and is achieved when the service level is 76.0% for the supplier's Stackelberg game and 66.5% for the retailer's Stackelberg game. In a more practical example, if the supplier's Stackelberg service level is 90.0% then the retailer's Stackelberg service level is approximately 82%. Hence, the deterioration in customer service is not trivial, and if there is a leader the customers prefer that it is the supplier.

Finally, numerical experiments under a wide range of values for α and b suggest that the retailer's Stackelberg game achieves a higher total profit than the supplier's Stackelberg game. However, we have been unable to provide a proof.

3.8 Proofs

3.8.1 Proof of Proposition 3

The function $\Pi(s, \nu)$ defined in (3.5) is continuously differentiable and bounded above by λr in $X = \{(s, \nu) \mid s \geq 0, \nu > 0\}$. Thus, a global maximum is either a local interior maximum that satisfies the first-order conditions or an element of the boundary of X ; alternatively, there could be no global maximum if the function increases as $s \rightarrow \infty$ or $\nu \rightarrow \infty$.

However, we have checked that $\lim_{s \rightarrow \infty} \Pi(s, \nu) \rightarrow -\infty$ for $\nu > 0$, and $\lim_{\nu \rightarrow \infty} \Pi(s, \nu) \rightarrow -\infty$ for $s \geq 0$, which implies that a global maximum exists. From the first-order conditions (3.6) and (3.7), the only interior point that is a candidate for the global maximum is (s^*, ν^*) . In addition, the Hessian of $\Pi(s, \nu)$ at (s^*, ν^*) is given by

$$H(s^*, \nu^*) = - \begin{pmatrix} \nu^* & s^* \\ s^* & \frac{c(\ln(1+b)+2)}{\nu^*} \end{pmatrix}.$$

Because $\ln(1+b) > 0$ for $b > 0$, the Hessian is negative definite and (s^*, ν^*) is the unique local maximum in the interior of X . The resulting profit is $\Pi(s^*, \nu^*) = r - p - c - 2\sqrt{c \ln(1+b)}$. Finally, $\lim_{\nu \rightarrow 0} \Pi(s, \nu) \rightarrow -\infty$ for $s \geq 0$, and

$$\Pi(0, \nu) = r - p - c - \frac{b}{\nu} + c\nu \leq r - p - c - 2\sqrt{cb} < \Pi(s^*, \nu^*) \quad \text{for } \nu > 0, b > 0.$$

Thus, (s^*, ν^*) is the unique global maximum for $\Pi(s, \nu)$. ■

3.8.2 Proof of Proposition 8

To derive the Stackelberg equilibrium, we find it convenient to define

$$\beta = \nu s, \tag{3.53}$$

and rewrite the supplier's reaction curve (3.16) as

$$e^{-\beta} \left(\frac{\beta + 1}{\beta^2} \right) = \frac{c}{(1 - \alpha)bs^2}. \tag{3.54}$$

The one-to-one correspondence between the base stock level s and the service level parameter β (recall that the service level is $e^{-\beta} \times 100\%$) allows the retailer in this Stackelberg game to choose β rather than s . By (3.3) and (3.53), the retailer's profit is

$$\Pi_R(\beta, \nu) = r - w - \left(\frac{\beta - 1 + (1 + \alpha b)e^{-\beta}}{\nu} \right). \tag{3.55}$$

Solving (3.54) for s and using (3.53) gives

$$\nu(\beta) = \sqrt{\frac{(1 - \alpha)b(\beta + 1)e^{-\beta}}{c}}. \tag{3.56}$$

Substituting (3.56) into (3.55) yields the retailer's profit as the following concave function of $\beta \geq 0$:

$$\Pi_R(\beta) = r - w - \sqrt{\frac{c}{(1-\alpha)b}} \left(\frac{\beta - 1 + (1 + \alpha b)e^{-\beta}}{\sqrt{(\beta + 1)e^{-\beta}}} \right). \quad (3.57)$$

Therefore, the first-order condition

$$\frac{\beta^2 + (\beta + 2)(1 - (1 + \alpha b)e^{-\beta})}{2(\beta + 1)^{\frac{3}{2}}e^{-\frac{\beta}{2}}} = 0 \quad (3.58)$$

is sufficient for optimality. Because $1 + \alpha b = e^{\nu_\alpha^* s_\alpha^*}$ and the denominator of (3.58) is always positive, condition (3.58) is equivalent to (3.46). Hence, by (3.46), (3.53) and (3.56), the Stackelberg equilibrium is given by (3.46)-(3.47). ■

3.8.3 Proof of Proposition 9

To prove (3.48), note that the left side of (3.46) is positive if $\beta > \nu_\alpha^* s_\alpha^*$. Thus, the root $\hat{\beta}$ of (3.46) must satisfy $\hat{\beta} \leq \nu_\alpha^* s_\alpha^*$, i.e., $\hat{\nu}_\alpha \hat{s}_\alpha \leq \nu_\alpha^* s_\alpha^*$.

To show that $\hat{\nu}_\alpha \geq \nu_\alpha^*$ and $\hat{s}_\alpha \leq s_\alpha^*$, we first observe that $\nu(s) = \arg \max_{\nu > 0} \{\Pi_R(s, \nu)\}$ and $\frac{\partial^2 \Pi_R(s, \nu)}{\partial s \partial \nu} = -(1 - \alpha)bse^{-\nu s} \leq 0$. Thus (e.g., Chapter 2 of Topkis 1998), $\Pi_R(s, \nu)$ satisfies the decreasing difference property,

$$\frac{d\nu(s)}{ds} \leq 0. \quad (3.59)$$

In addition, the function $e^{-\beta}(\beta + 1)/\beta^2$ is decreasing in $\beta > 0$. Hence, from (3.54) and inequality (3.48), we conclude that $\hat{s}_\alpha \leq s_\alpha^*$. Finally, (3.59) and $\hat{s}_\alpha \leq s_\alpha^*$ implies that $\hat{\nu}_\alpha \geq \nu_\alpha^*$.

The retailer's profit in (3.57) is a decreasing function of β for $\beta \geq \hat{\beta}$. Hence, inequality (3.51) follows from (3.48). To prove (3.52), i.e., $\Pi_S(\beta^*, \nu_\alpha^*) \geq \Pi_S(\hat{\beta}, \hat{\nu}_\alpha)$, we first use (3.4) to rewrite the supplier's profit as

$$\Pi_S(\beta, \nu) = w - p - c(\nu + 1) - \left(\frac{(1 - \alpha)be^{-\beta}}{\nu} \right).$$

The function $\Pi_S(\beta, \nu)$ is increasing in β for $\nu > 0$, and so inequality (3.48) implies that $\Pi_S(\beta^*, \hat{\nu}_\alpha) \geq \Pi_S(\hat{\beta}, \hat{\nu}_\alpha)$. Hence the proof of (3.52) will follow if we can show that $\Pi_S(\beta^*, \nu_\alpha^*) \geq \Pi_S(\beta^*, \hat{\nu}_\alpha)$. For any fixed nonnegative β , the function $\Pi_S(\beta, \nu)$ is concave in ν , achieves its only maximum at $\nu(\beta) = \sqrt{\frac{(1-\alpha)be^{-\beta}}{c}}$, and is decreasing for $\nu \in [\nu(\beta), \infty)$. In particular, we have $\nu(\beta^*) = \sqrt{\frac{(1-\alpha)b}{c(1+\alpha b)}} \leq \sqrt{\frac{(1-\alpha)b}{c(1+\alpha b)}} \sqrt{1 + \ln(1 + \alpha b)} = \nu_\alpha^*$. This inequality and (3.49) imply that $\Pi_S(\beta^*, \nu_\alpha^*) \geq \Pi_S(\beta^*, \hat{\nu}_\alpha)$, which completes the proof of (3.52). ■

Chapter 4

Revenue Management of a Make-to-Stock Queue

4.1 Introduction

In this chapter, we address the joint problem of admission and production/inventory control in a single-product manufacturing setting. Demand is stochastic in both the arrival pattern and the price that each customer is willing to pay. On the other hand, we embed the production strategy under a make-to-stock queue model framework. That is, production capacity is limited and stochastic and the manufacturer carries finish goods inventory to service demand. Our interest is to adequately balance the benefits from selling and the costs from providing the good to the end customers by rejecting some of the orders and controlling the levels of stock.

Certainly, admission and production control to a make-to-stock queue is not a new research topic. However, almost all of the analysis has been carried out under the assumption that demand prices are fixed or at most depend deterministically on the type of customer that is placing the order. That is, demand is partitioned into several independent classes parameterized by independent demand processes -usually Poisson- and prices (see the literature review below §4.2). A major drawback of these static price models is that they are unable to capture the natural variability and correlation of the price of successive orders. In this chapter, we present a mathematical model of the traditional make-to-stock queue that incorporates explicitly the fact that selling prices vary stochastically and continuously over time.

Price variations is a natural phenomenon of dynamic markets. The reason is that

price is one of the most effective variable that managers can manipulate to encourage or discourage demand in the short run. Price is not only important from a financial point of view but also as an operational tool that helps regulating workload and production pressures. Adding to this, the emerging growth of electronic commerce is facilitating price changes. The costs of physically relabelling the prices of goods and those associated to informing customers about these changes are being considerably reduced in this new channel (e.g., Brynjolfsson and Smith (1999) report a substantial increase in the number of price changes in the Internet with respect to conventional retailers). In addition, customers are getting more and better information about product variety. For instance, Internet price search intermediaries (web aggregators) offer customers easy access to price lists and it is just a matter of time that most consumers' purchasing decisions will be based on this type of information. On the other hand, in the Business-to-Business setting the situation is not much different. The increasing popularity of online auctions is a good example showing how *spot markets* are winning ground over traditional long-term fixed price contracts.

Is in this rapidly changing environment that managers need to be prepared to face continuous, and to some extent unpredictable, fluctuations in price and demand. For systems with limited production capacity, this situation implies that rejecting orders is possibly beneficial. During periods of congestion, low price orders might only exhaust dedicated capacity that will be needed to server future high profit orders. The challenge is to be able to serve the right customer at the right price at the right time. Thus, dynamic admission policies -together with production decisions- are crucial to balance workload and at the same time to maximize profitability of the business operation.

The rest of this chapter is organized as follows. Next section revises the relevant literature on optimal admission and inventory control to a queue. Section §4.3 introduces the stochastic control problem while §4.4 presents its diffusion approximation. Optimality conditions and numerical solutions are reported in §4.5 and §4.6 respectively. Section §4.7 is devoted to develop approximation policies and to analyze their performance. Finally, some concluding remarks are discussed in §8.

One of the key element of our formulation is the use of Brownian motion processes to characterize buffer sizes and price processes. Some elementary knowledge about these processes and stochastic calculus from the part of the reader will be required. For this reason, we have included at the end Appendix A describing the basic properties of the Brownian motion process, the key results on stochastic and Itô's calculus,

and some applications of these diffusion processes to queueing model.

4.2 Literature Review

The literature on admission control to a queueing system is extensive but by no means fully developed. In what follows, we describe some of the research that has been done and that closely relates to our work. This review is far from being complete and we recommend the survey papers by Shaler Stidham (1985, 1988) that are still good introductory references to the field.

Naor (1969) is one of the first to investigate the effects of customer rejection on the performance of a queue. The setting is a $M/M/1$ model with fixed reward R for order completion and fixed cost rate C per unit for queueing orders. In this case, the admission policy is purely based on the inventory position since prices are fixed. Naor shows that a simple threshold policy is optimal when maximizing the infinite horizon expected profit, i.e., accept a new job if the number of jobs already in the system is less than the critical value $n_0 := \min\{n : n + 1 - \rho - \rho^2 - \dots - \rho^{n+1} > R\mu(1 - \rho)C^{-1}\}$, where ρ is the traffic intensity before rejections and μ is the service rate. By the same time, Miller (1969) was looking at a different scenario. In his case, the setting is a $M/M/n/n$ system with m different classes of customers, which are differentiated only by the reward τ_k ($k = 1, \dots, m$) associated to the completion of a class- k job. The admission policy in this case is also of a threshold type but now not only depending on the inventory position but also on the class of the arriving order. Quantitative results about the control policy are only reported algorithmically in this case because of the underlying complexity of the infinite horizon continuous time Markov chain formulation. Although these two pioneer works are essentially different when combined they share much of the complexity of our current research. In particular, we share (i) the trade-off between collecting the reward from accepting orders and the holding cost incurred to queue them with Naor's paper and (ii) the trade-off between accepting a low price job now and eventually rejecting a better deal in the early future because of unavailable capacity with Miller's work.

Prototypical extensions to these earliest papers were presented during the 70's. Yechiali (1971, 1972) generalized Naor's formulation to the $GI/M/1$ and $GI/M/s$ cases respectively. Lippman (1975) improved Miller's model by introducing queue capacity ($M/M/n/k$, $k > n$) and allowing an infinite number of customer classes. Stidham (1978) enriches Yechiali's $GI/M/1$ formulation by considering random rewards and

allowing customers to pay a state dependent entrance fee instead of being rejected. Finally, Johansen and Stidham (1980) synthesize much of these results under a general semi-Markov decision process formulation.

After this first decade of research, the interest start shifting from the single-queue Markovian formulations towards more general and complicate settings. One of such extensions is the study of admission control policies for network of queues, e.g., Ephremides *et al.* (1980), Ghoneim and Stidham (1985), Veatch and Wein (1992). Helm and Waldmann (1984) look at something different. Using a general semi-Markovian formulation, the authors analyze optimal admission control policies to a single queue for which system parameters -such as prices, holding cost rates, and production capacity- are function of the *environment*. In particular, Helm and Waldmann consider the case where the state of the environment is a Markov process. The importance of this new modeling device is that the reward/price of successive orders are not longer independent as we should expect to see happening in real systems. Another important extension is to consider inventory decisions. All the articles that we have mentioned above consider make-to-order systems, i.e., the possibility of holding inventory is excluded. Even though inventory decisions can be seen as independent of admission policies, they are intimately related specially in models that assume *lost sales* (e.g., Li (1992), Ha (1997)). A recent paper in this line is Carr and Duenyas (2000). The authors consider a single machine production system with two classes of product: one class is served using a make-to-order policy while the other product is make-to-stock. Dynamic admission policies and scheduling decisions are studied for the two-class M/M/1 queue and the case or Erlang distributions.

So far, we have presented the research that uses Markov processes as the main modeling tool. In this setting, and except for very few exceptions, most of the results regarding admission policies are qualitative in the sense that only structural properties of the value function (such as monotonicity, concavity, modularity, etc) are reported. Explicit admission rules can only be obtained numerically solving some type of DP recursion. This is probably one of the reasons why a completely different stream of work has been running in parallel, namely, diffusion and *heavy traffic* models. These more crude but also more tractable approximations use an adequate time and space scaling to model inventory position as a (regulated) diffusion process (e.g., Harrison (1985)). The main advantage is that DP recursions or Bellman optimality equations are represented by systems of ordinary or partial differential equations whose solutions are the desired control policies. Although these systems of equations (specially PDE's)

are difficult to solve, they provide useful insight about the problem, insight that can be used to develop good approximations, as we do in this research.

One-dimensional diffusion control models (i.e., single queue systems) were studied in Harrison and Taylor (1978), Beneš *et al* (1980), Harrison and Taksar (1983), Menaldi and Robin (1984), Taksar (1985), Wein (1992), Krichagina *et al.* (1994), among others. In these papers, a one-dimensional diffusion process (e.g., inventory position, queue length) is regulated by adjusting an *input* process (e.g., demand) and/or an *output* process (e.g., production). Typically, a cost function depending on the regulated diffusion process is minimized. Finite/infinite horizon as well as discounted/average versions of this problem have been considered and threshold policies of bang-bang type have been shown to be optimal. That is, in the interior of a certain region G the process moves free of any control and it is only when it hits the boundary that maximum control is used to keep the process within G . It is important to mention that these diffusion control papers assume fixed (or deterministic in a few exceptions) objective function parameters such as holding cost rates, per unit production cost, and selling prices. Finally, we point out that uncertainty in these models is present only through the diffusion process used to represent the inventory position.

In view of the literature mentioned above, our work can be described as an extension to the one-dimensional diffusion control problem studied by Harrison and Taylor (1978), Harrison and Taksar (1983), and Taksar (1985). Our main contribution is to consider the price of incoming orders as stochastic rather than fixed. In that sense, our formulation shares much with Lippman (1975) and Stidham (1978). However a major distinction with these two papers is that we consider that the prices of successive orders are correlated. This feature makes our work to be closely related to Helm and Waldmann (1984) but using a continuous rather than discrete approach.

4.3 The Control Problem

Consider a single-server manufacturing system that serves according to a make-to-stock discipline an exogenous demand. Both production and demand are independent stochastic processes. We assume that demand is fully backlogged and that holding and backordering costs are linearly incurred by the manufacturer. We also consider that the price paid by an incoming customer is an exogenous stochastic process. The data and notation of the model are summarized as follows:

- $D(t)$: Cumulative demand process. Total number of orders up to time t . We set λ the average demand rate and c_d^2 the squared coefficient of variation (scv) for the interarrival time.
- $S(t)$: Cumulative production process. Total number of unit produced if the server has been continuously working during $[0, t]$. We set μ and c_s^2 the service rate and scv for the service time respectively.
- $R(t)$: Price of orders at time t . We model the price as a driftless geometric Brownian motion. Thus, $R(t) = R_0 e^{\delta W_R(t)}$, where R_0 is the initial price at time 0, δ is the diffusion parameter, and $W_R(t)$ is a Wiener process.
- $Z(t)$: Inventory process. Number of units (possibly negative) in inventory at time t .
- h : Holding cost per order per unit time.
- b : Backorder cost per order per unit time.
- $c(z)$: Cost rate if the inventory position is z (possibly negative).

$$c(z) = \begin{cases} hz & \text{if } z \geq 0, \\ bz & \text{if } z < 0. \end{cases}$$

Except for the selling price $R(t)$, the rest of the data are standard inputs for a make-to-stock manufacturing system. Thus, the novelty of our formulation is in the use of a stochastic process to model the selling price. The particular choice of a geometric Brownian motion (GBM) is mainly influenced by the Finance literature where GBM is the prototypical model used to represent price processes. Its simplicity and mathematical tractability make GBM an attractive modeling tool (e.g., Black-Scholes option-pricing formula). For more details on the use of GBM to model prices, readers are referred to the early works of Osborne (1964) and Samuelson (1965) and to Merton (1990).

Let $A(t)$ and $P(t)$ be the cumulative time that the system has been accepting orders and producing respectively during $[0, t]$. These are the decision variables that the manufacturer can manipulate in order to maximize the net profit. Both $A(t)$ and $P(t)$ are assumed to be nonnegative, nondecreasing, and non-anticipating with respect to the demand, production, and price processes.

We consider the expect long-term average criteria version of the problem. That is, the manufacturer is interested in the solution to the following stochastic problem

$$\max_{A,P} \inf \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T R(t) dD(A(t)) - \int_0^T c(Z(t)) dt \right], \quad (4.1)$$

s.t

$$Z(t) = Z(0) + S(P(t)) - D(A(t)). \quad (4.2)$$

In (4.1), $dD(A(t))$ represents the marginal demand that is accepted at time t . Thus, the first integral in (4.1) is the net selling revenue up to time T . The second integral, on the other hand, is the total holding and backordering cost up to time T . Finally, (4.2) describes the evolution of the inventory position, which is the difference between cumulative production and demand. $Z(0)$ is the initial inventory level that we will assume to be equal to 0 from here on.

Because (4.1)-(4.2) appears to be analytically intractable for the general case, we simplify the problem complexity by considering its diffusion version. That is, under a *heavy traffic scaling* transformation, we replace (4.1)-(4.2) by a sequence of stochastic control problems that weakly converge to a diffusion control problem. The details of this transformation are carried out in the following section.

4.4 The Diffusion Control Problem

The main tool underlying the heavy traffic approximation is a functional limit theorem, Donsker's Theorem. The basic elements of Donsker's theorem are described in section §A.2.2 on Appendix A. For a more detailed and technical description, we recommend Billingsley (1999) and the references therein. The key idea behind Donsker's result is that under an appropriate scaling of time and space (and other technical conditions that we omit here) a stochastic process converges weakly to a Brownian motion. To be more precise, let us first introduce the following notation.

Definition 1 Let $X(t)$ be a stochastic process with average rate $\bar{x} = \lim_{t \rightarrow \infty} \frac{X(t)}{t}$. We define the Brownian scaling, the fluid scaling, and the centered versions of X by

$$\tilde{X}(t) = \frac{X(nt)}{\sqrt{n}}, \quad \bar{X}(t) = \frac{X(nt)}{n}, \quad \text{and} \quad X^c(t) = X(t) - \bar{x}t$$

respectively, where n is a sufficiently large positive number (scaling factor).

What Donsker's theorem establishes is that $\tilde{X}(t)$ weakly converges to a Brownian motion as $n \rightarrow \infty$. Motivated by this result, we introduce the following sequence of transformation to our system.

From the definition above we have

$$\tilde{Z}(t) = \frac{S(P(nt))}{\sqrt{n}} - \frac{D(A(nt))}{\sqrt{n}} = \frac{S(n\frac{P(nt)}{n})}{\sqrt{n}} - \frac{D(n\frac{A(nt)}{n})}{\sqrt{n}} = \tilde{S}(\bar{P}(t)) - \tilde{D}(\bar{A}(t)).$$

Let us recall that the average rate for $S(t)$ and $D(t)$ are μ and λ respectively. Thus, $\tilde{Z}(t)$ becomes

$$\tilde{Z}(t) = \tilde{S}^c(\bar{P}(t)) - \tilde{D}^c(\bar{A}(t)) + \sqrt{n}(\mu\bar{P}(t) - \lambda\bar{A}(t)).$$

Finally, introducing the rejection and idleness processes $U(t) = t - A(t)$ and $I(t) = t - P(t)$ respectively, we get the following dynamics for the scaled inventory position

$$\tilde{Z}(t) = \chi(t) + \lambda\tilde{U}(t) - \mu\tilde{I}(t), \quad (4.3)$$

where $\chi(t) = \tilde{S}^c(\bar{P}(t)) - \tilde{D}^c(\bar{A}(t)) + \sqrt{n}(\mu - \lambda)t$.

The key of the Brownian approximation is to use Donsker's argument to replace $\chi(t)$ by a Brownian process. However, the presence of $\bar{P}(t)$ and $\bar{A}(t)$ as the argument of the production and demand processes makes the transformation less straightforward. In order to make the argument precise, we need to apply a *random time-change* result that replaces $\bar{P}(t)$ and $\bar{A}(t)$ by their time average limits (see proposition (26) on the Appendix A and/or Harrison (1988), section §5 and §11 for details). Both $\bar{P}(t)$ and $\bar{A}(t)$ are clearly two random time-change processes. However, it is not absolutely clear if they are weakly convergent to some limit. We simplify this problem by assuming that the admission process is in fact stationary from an average sense. That is

$$\lim_{t \rightarrow \infty} \frac{\bar{A}(t)}{t} = \beta. \quad (4.4)$$

This assumption obviously ensures weak convergence of the admission policy. On the other hand, the production policy can now be constrained using classical queueing arguments. In fact, given the previous assumption the average effective arrival rate is given by $\lambda\beta$. Thus, the server must be working on average a fraction $\rho\beta$ (where $\rho := \frac{\lambda}{\mu}$) in order to keep the inventory process under control. The approximation for

$\bar{T}(t)$ becomes now evident and is given by

$$\lim_{t \rightarrow \infty} \frac{\bar{T}(t)}{t} = \rho \beta. \quad (4.5)$$

We notice that both $\bar{A}(t)$ and $\bar{T}(t)$ are bounded below and above by 0 and t respectively. Thus, conditions (4.4) and (4.5) are well-defined only if $0 \leq \beta \leq \min\{1, \rho^{-1}\}$.

From (4.4) and (4.5) we approximate $\chi(t)$ as follows

$$\chi(t) \approx \tilde{S}(\rho \beta t) - \tilde{D}(\beta t) + \sqrt{n}(\mu - \lambda)t$$

and the Brownian approximation assumes that $\chi(t) \Rightarrow X(t)$, where $X(t)$ is a Brownian motion with drift $\sqrt{n}(\mu - \lambda)$ and variance $\lambda \beta (c_a^2 + c_s^2)$. The inventory dynamics is approximated by:

$$\tilde{Z}(t) = X(t) + \lambda \tilde{U}(t) - \mu \tilde{I}(t).$$

Thus, $\tilde{Z}(t)$ behaves as a two-sided controlled Brownian motion. Before analyzing the effects of the Brownian transformation on the objective function (4.1), let us discuss the implications of (4.4). Since the system manager controls both arrival and production rates, he/she has the flexibility to regulate the inventory position through different combinations of arrival and production rates. This “ambiguity” on the system operations avoids us to apply directly the convergence result. Therefore, by fixing the average arrival rate according to (4.4) we are also fixing the average production rate through (4.5). Although for real manufacturing systems we expect some kind of stationarity on the admission policy (i.e., (4.4) holding), it is not obvious a priori that a stationary admission policy is optimal. We will not dig into this issue here and we will concentrate our attention only on stationary admission policies. Finally, in order to find the optimal value of β we maximize the objective function (4.1) which by the additional constraint (4.4) has become a function of β .

Let us now turn to the objective function (4.1). First of all, given the piecewise linearity of the cost function $c(z)$, we have that

$$\int_0^T c(Z(t)) dt = \int_0^{\frac{T}{n}} c(Z(nt)) n dt = \sqrt{n^3} \int_0^{\frac{T}{n}} c(\tilde{Z}(t)) dt.$$

On the other hand, we have that

$$\int_0^T R(t) dD(A(t)) = \int_0^{\frac{T}{n}} R(nt) dD(n\bar{A}(t)) = \sqrt{n} \int_0^{\frac{T}{n}} R(nt) d\tilde{D}(\bar{A}(t)).$$

Moreover, given the Brownian scaling we have that $\tilde{D}(t)$ weakly converges to a BM with drift $\sqrt{n}\lambda$ and diffusion parameter λc_d^2 . Therefore,

$$d\tilde{D}(\bar{A}(t)) = \sqrt{n}\lambda d\bar{A}(t) + \sqrt{\lambda}c_d dW_D(\bar{A}(t)),$$

where $W_D(t)$ is the underlying Wiener process associated to \tilde{D} . Thus, after some algebra we get that the objective function is given by

$$\begin{aligned} \frac{1}{T} E \left[\underbrace{n \int_0^{\frac{T}{n}} R(nt) dt}_{(a)} + \underbrace{c_d \sqrt{\lambda n} \int_0^{\frac{T}{n}} R(nt) dW_D(\bar{A}(t))}_{(b)} \right. \\ \left. - \lambda \sqrt{n} \int_0^{\frac{T}{n}} R(nt) d\tilde{U}(t) - \sqrt{n^3} \int_0^{\frac{T}{n}} c(\tilde{Z}(t)) dt \right]. \end{aligned}$$

From Fubini's theorem and the fact that $E[R(t)] = E[R_0] e^{\frac{\delta^2 t}{2}}$ we have that

$$(a) = \left(\frac{2E[R_0]}{\delta^2} \right) \left(e^{\frac{\delta^2 T}{2}} - 1 \right).$$

We notice that (a) is independent of the controls, thus we can exclude it from the formulation. On the other hand, (b) is a stochastic integral in the Itô sense. Since $W_D(t)$ is a Wiener process, we expect (b) to have low impact in the objective in an average sense. This conclusion is true if $R(t)$ and $W_D(t)$ are independent. In this case the expected value of (b) is exactly equal to 0 (see proposition (10) below). However, in our setting we should expect demand and prices to be negatively correlated and (b) does not necessarily cancelled. To keep our formulation simple, we will assume for the moment that $R(t)$ and $W_D(t)$ are independent. We will remove this assumption later on in section §??. Under the independence assumption the following result holds.

Proposition 10 *If $R(t)$ is adapted and independent of $W_D(t)$ and if the control process $\bar{A}(t)$ is nonanticipating then,*

$$E \left[\int_0^T R(t) dW_D(\bar{A}(t)) \right] = 0.$$

Proof: See §4.10 at the end of this chapter.

We can now remove (a) and (b) from the objective function above. If in addition we make the change of variable $nT \leftarrow T$, the objective becomes

$$-\frac{\sqrt{n}}{T} E \left[\lambda \int_0^T \bar{R}(t) d\bar{U}(t) + \int_0^T c(\bar{Z}(t)) dt \right]. \quad (4.6)$$

The minus sign reveals that this objective is in fact a cost function. The first integral is the cost of rejecting orders while the second is the holding and backordering cost. A very important step in this heavy traffic scaling is to identify the orders of magnitude of the different parameters. In particular, we want the Brownian motion $X(t)$ to be well-defined and the two cost components in (4.6) to be in the same order of magnitude to avoid trivial solutions.

First, we recall that $X(t)$ is a (θ, σ^2) – Brownian motion where the drift $\theta := \sqrt{n} \mu(1 - \rho)$ and the diffusion $\sigma^2 := \lambda \beta(c_a^2 + c_s^2)$. Since our heavy traffic approximation works for large values of n , we require θ to be bounded as $n \rightarrow \infty$. This condition, which is traditionally known as the *Heavy traffic condition*, assumes then the existence of a bounded real $\bar{\theta}$ such that

$$\mu \lim_{n \rightarrow \infty} \sqrt{n}(1 - \rho) = \bar{\theta}.$$

The limiting drift $\bar{\theta}$ can be positive, negative, or zero depending on the relative value of the demand and production rates λ and μ . For instance, if $\mu > \lambda$ then $\bar{\theta}$ is positive.

The heavy traffic condition implies that the system works at a high level of utilization, in particular $1 - \rho$ has to be order $1/\sqrt{n}$. Therefore, we expect to observe very large queues in the unscaled system. If the penalty of rejecting orders ($R(t)$) is not large enough then the holding/backordering cost will dominate the objective in (4.6). To avoid this situation, we restrict ourselves to the cases that satisfy

$$O(\bar{R}(t)) \sim O(c(\bar{Z}(t))).$$

Given our scaling, we can write $\bar{R}(t)$ as follows

$$\bar{R}(t) = \frac{R_0}{n} e^{\sqrt{n}\delta \frac{W_R(nt)}{\sqrt{n}}}. \quad (4.7)$$

Thus, setting the holding/backordering cost to be order one, we get the following condition for the price parameters

$$R_0 \sim O(n) \quad \text{and} \quad \delta \sim O\left(\frac{1}{\sqrt{n}}\right).$$

The initial price has to be large enough (order n) to make the penalty of rejection on the same order of magnitude than the holding/backordering cost. In addition, we need a condition on the diffusion parameter. In particular, δ has to be small (order one over root n) to maintain the price bounded.

We summarize these heavy traffic conditions by assuming that there exist scalars \bar{R}_0 and $\bar{\delta}$, such that

$$\lim_{n \rightarrow \infty} \frac{R_0}{n} = \bar{R}_0 \quad \text{and} \quad \lim_{n \rightarrow \infty} \sqrt{n}\delta = \bar{\delta}.$$

As a representative example consider $n = 100$. If the holding and backordering cost rates are order 1, then the initial price R_0 has to be order 100, the diffusion δ has to be order 0.1, and the traffic intensity ρ has to be order 0.9.

We are now in condition to write the heavy traffic control problem, which is given by

$$\min_{\bar{U}, \bar{I}} \sup \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T \bar{R}(t) \lambda d\bar{U}(t) + \int_0^T c(\bar{Z}(t)) dt \right], \quad (4.8)$$

s.t.

$$\bar{Z}(t) = X(t) + \lambda \bar{U}(t) - \mu \bar{I}(t), \quad (4.9)$$

$$\lim_{T \rightarrow \infty} \frac{\bar{U}(T)}{T} = \bar{U}, \quad (4.10)$$

where $\bar{U} = \sqrt{n}(1 - \beta)$. This problem has almost the same structure than (4.1)-(4.2). The only structural difference is constraint (4.10) which appears as a direct consequence of (4.4). Moreover, this constraint is not easily handled in the derivation of optimality condition and we use a Lagrangian approach to put it in the objective. Let α be the Lagrangian multiplier associated to (4.10). The problem becomes

$$\min_{\bar{U}, \bar{I}} \sup \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T (\bar{R}(t) + \alpha) \lambda d\bar{U}(t) + \int_0^T c(\bar{Z}(t)) dt \right], \quad (4.11)$$

s.t.

$$\bar{Z}(t) = X(t) + \lambda \bar{U}(t) - \mu \bar{I}(t). \quad (4.12)$$

From weak duality, we know that if $(\tilde{U}^*, \tilde{I}^*)$ is a solution to (4.11)-(4.12) that satisfies (4.10), then it is also a solution to (4.8)-(4.10). This result is formalized in the next proposition.

Proposition 11 *Suppose that $(\tilde{U}^*, \tilde{I}^*)$ is a solution to the Lagrangian problem (4.11)-(4.12). Suppose, moreover, that*

$$\lim_{T \rightarrow \infty} \frac{\tilde{U}^*(T)}{T} = \bar{U}. \quad (4.13)$$

Then $(\tilde{U}^, \tilde{I}^*)$ is a solution to the constrained problem (4.8)-(4.10).*

For a proof see §4.10 at the end of this chapter. We can use the result of proposition (11) to find an optimal solution to the original problem as follows:

1. Solve the Lagrangian relaxation as a function of α -the Lagrangian multiplier. Let $(\tilde{U}_\alpha^*, \tilde{I}_\alpha^*)$ be the solution.
2. Find α^* such that $\tilde{U}_{\alpha^*}^*$ satisfies condition (4.13).

Naturally, we will first tackle step 1 and later we will address step 2. For this purpose, we start computing optimality conditions for the Lagrangian formulation.

4.5 Optimality Conditions

We now turn to the problem of finding the optimality equations that characterize the optimal policy (U^*, I^*) . We use an heuristic derivation following Taksar (1985). First, we look at the discounted version of the problem, that is

$$(P^\gamma) \quad \min_{U, I} \quad E \left[\int_0^\infty e^{-\gamma t} (R(t) + \alpha) dU(t) + \int_0^\infty e^{-\gamma t} c(Z(t)) dt \right], \quad (4.14)$$

s.t

$$Z(t) = X(t) + U(t) - I(t), \quad (4.15)$$

where $\gamma > 0$ is the discount rate. The attend reader should have noticed that we have simplified the notation in (4.14)-(4.15) by replacing $\tilde{U} \leftarrow \lambda \tilde{U}$ and $\tilde{I} \leftarrow \mu \tilde{I}$ and by removing the tildes and bars of the different processes. In what follows, it is important to remember that the two exogenous processes are (i) $R(t)$ a driftless geometric Brownian motion with diffusion δ starting at R_0 and (ii) $X(t)$ a Brownian

motion with drift θ and diffusion σ . For simplicity and without loss of generality in what follows we will assume $R_0 := 1$. This is equivalent to normalize holding and backordering cost rates (h and b respectively) by R_0 .

The approach is first to find the optimality equations for problem (P^γ) and then let $\gamma \rightarrow 0^+$ to derive the optimality equations for the average control problem by means of a Tauberian argument. We assume that $R(t) + \alpha > 0$ for otherwise the control problem would make no sense. Our approach for solving (4.14)-(4.15) mimics Harrison and Taksar (1983).

The essential step is heuristic. We assume the existence of a function $f^\gamma(x, r)$ twice continuously differentiable representing the optimal value in (4.14) starting at (x, r) . This assumption is known as *the principle of smoothness fit*. The main result obtained out of this assumption is summarized in the following proposition.

Proposition 12 *The Hamilton-Jacobi-Bellman equation for problem (P^γ) is*

$$[-f_x^\gamma(x, r)] \wedge [f_x^\gamma(x, r) + r + \alpha] \wedge [\tilde{\Gamma} f^\gamma(x, r) + c(x) - \gamma f^\gamma(x, r)] = 0, \quad (4.16)$$

where

$$\tilde{\Gamma} = \theta \frac{\partial}{\partial x} + \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} + \frac{\delta^2 r}{2} \frac{\partial}{\partial r} + \frac{(\delta r)^2}{2} \frac{\partial^2}{\partial r^2}$$

is the infinitesimal generator of $(X(t), R(t))$.

Proof: See the §4.10.

We now turn to the average cost case. First we put $V^\gamma(x, r) = f^\gamma(x, r) - f^\gamma(\hat{x}, \hat{r})$. The value function $V^\gamma(x, r)$ represents the relative discounted cost of starting in state (x, r) instead of starting in state (\hat{x}, \hat{r}) . We then let $\gamma \downarrow 0$ and assume that $g = \lim \gamma \cdot f^\gamma(\hat{x}, \hat{r})$ exists. Then, passing to the limit in (4.16) we get

$$[-V_x(x, r)] \wedge [V_x(x, r) + r + \alpha] \wedge [c(x) + \tilde{\Gamma} V(x, r) - g] = 0. \quad (4.17)$$

This is the Hamilton-Jacobi-Bellman equation for the average cost problem (4.11)-(4.12). Here g represents the minimal average cost of the problem.

Given the singular nature of the control processes, the first two terms in (4.17) are boundary conditions. For example and given the smoothness assumption, we can rewrite $V_x(x, r) + r + \alpha = 0$ as follows

$$V(x, r) - V(x + \epsilon, r) = (r + \alpha) \epsilon + o(\epsilon).$$

Given a price r , the left-hand side is the marginal cost of accepting an order of size ϵ and the right-hand side is the marginal cost of rejecting an order of the same size. Thus, the condition above defines the set of states (x, r) for which it is optimal to reject orders at the marginal cost $r + \alpha$. Similarly, $V_x(x, r) = 0$ characterizes those states for which it is optimal to stop producing.

Let us introduce two boundary curves $x = \eta(r)$ and $x = \xi(r)$ which will play a central role in our solution. We define these curves as follows

$$V_x(\eta(r), r) + r + \alpha = 0 \quad \text{and} \quad V_x(\xi(r), r) = 0. \quad (4.18)$$

According to our previous discussion, both $\eta(r)$ and $\xi(r)$ are switching curves for the accept/reject and the produce/idle decisions respectively and they are essentially the decision variables for the diffusion control problem. Combining this observation, the smoothness of fit assumption, and (4.17) we obtain the following partial differential equation (PDE) representation

$$\tilde{\Gamma}V(x, r) + c(x) - g = 0, \quad \forall(x, r) : \eta(r) \leq x \leq \xi(r), \quad r > 0, \quad (4.19)$$

$$V_x(\eta(r), r) + r + \alpha = 0, \quad V_x(\xi(r), r) = 0, \quad (4.20)$$

$$-(r + \alpha) \leq V_x(x, r) \leq 0, \quad (4.21)$$

$$V_{xx}(\eta(r), r) = V_{xx}(\xi(r), r) = 0. \quad (4.22)$$

This is a two-dimensional elliptic PDE problem with free boundary conditions. Closed form solutions for this type of problems are very rare in general. Unfortunately (4.19)-(4.22) is not an exception and we have not been able to solve it analytically. For this reason, we approach the solution through numerical and approximated methods.

Before jumping into these approximations, we simplify the system above by introducing the change of variable $r = \exp(y)$. With this change, the system becomes

$$\Gamma V(x, y) + c(x) - g = 0, \quad \forall(x, y) : \eta(y) \leq x \leq \xi(y), \quad (4.23)$$

$$V_x(\eta(y), y) + e^y + \alpha = 0, \quad V_x(\xi(y), y) = 0, \quad (4.24)$$

$$-(e^y + \alpha) \leq V_x(x, y) \leq 0, \quad (4.25)$$

$$V_{xx}(\eta(y), y) = V_{xx}(\xi(y), y) = 0, \quad (4.26)$$

where the infinitesimal generator Γ is given by

$$\Gamma = \theta \frac{\partial}{\partial x} + \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} + \frac{\delta^2}{2} \frac{\partial^2}{\partial y^2}.$$

The main advantage of this new formulation is that Γ is independent of the state, fact that will facilitate the analysis later. We omit here the proof of the validity of (4.23)-(4.26); it basically relies on the fact that under the change $R(t) = \exp(Y(t))$, $Y(t)$ is a driftless Brownian motion with diffusion δ .

4.6 Numerical Solution

As with any numerical method for solving PDE's, we start by discretizing the state space. We approach this problem using the *finite difference approximation* technique developed by H. Kushner (1977). This method relies on the fact that we can approximate a regulated diffusion process by a finite state Markov Chain. Readers are referred to Kushner and Dupuis (1992) for details.

Let us first define a bounded region Ω on the plane $\{(x, y) \in \mathfrak{R}^2\}$ where the process (X, Y) will lay on ($\partial\Omega$ denotes the boundary of Ω). We have to choose the set Ω large enough to ensure that (X, Y) will rarely reach the boundaries given the optimal control policies (U^*, I^*) . Next, we introduce \hbar as the *finite difference interval* which defines how finely the state space and time space are discretized. Then, the initial state space of the Markov chain associated to \hbar is given by the lattice

$$\Omega_{\hbar}^0 = \{(x, y) \in \Omega : x = n_x \hbar, y = n_y \hbar, |n_x| \leq N_x, |n_y| \leq N_y, n_x, n_y \text{ integers}\},$$

where N_x and N_y are positive integers that define the dimensions of Ω_{\hbar}^0 .

Starting with Ω_{\hbar}^0 , our solution method iteratively generates a sequence of regions $\Omega_{\hbar}^1, \Omega_{\hbar}^2, \dots$ such that $\Omega_{\hbar}^0 \supseteq \Omega_{\hbar}^1 \supseteq \Omega_{\hbar}^2 \supseteq \dots$ and that will approach the optimal region (Ω_{\hbar}^*) defined by the boundaries $\eta(y)$ and $\xi(y)$

$$\Omega_{\hbar}^* = \{(x, y) \in \Omega : x = n_x \hbar, y = n_y \hbar, \eta(y) \leq x \leq \xi(y), |n_y| \leq N_y, n_x, n_y \text{ integers}\}.$$

For this reason, we require N_x to be large enough to ensure $\Omega_{\hbar}^* \subseteq \Omega_{\hbar}^0$.

Suppose that we are in stage k and that the current region is Ω_{\hbar}^k . The remaining of this section is devoted to describing the details of the iteration that generates Ω_{\hbar}^{k+1} .

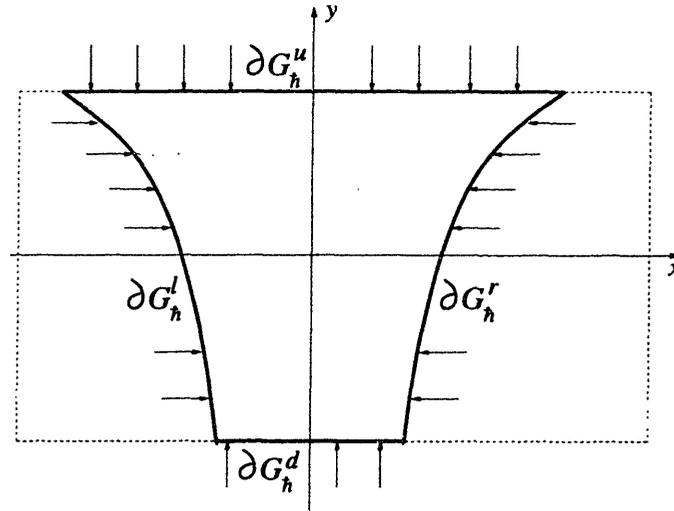


Figure 4.1: Reflection Field at the Boundaries.

For notational convenience, we omit the dependence on k .

The first step is to compute the transition probability distribution for the Markov chain associated to Ω_h . For this purpose, we assume that in the interior of Ω_h the manager does not exert any control over the evolution of the process but she does control it on the boundaries. Moreover, the boundaries of this region are reflecting according to figure 4.1.

The dashed line represents Ω_h^0 , the solid line is the current region Ω_h , and the arrows indicate the direction of the reflection.

Let $P_h((x, y), (x', y'))$ be the transition probability from state (x, y) to state (x', y') . Define (see Kushner and Dupuis (1992), section §5)

$$Q_h = \hbar |\theta| + \sigma^2 + \delta^2.$$

Then, the transition probability distribution is given by:

- If $(x, y) \in \Omega_h - \partial\Omega_h$ then

$$\begin{aligned} P_h((x, y), (x \pm \hbar, y); u) &= \frac{\sigma^2 + 2\hbar \theta^\pm}{2Q_h}, \\ P_h((x, y), (x, y \pm \hbar); u) &= \frac{\delta^2}{2Q_h}, \\ P_h(w, w'; u) &= 0 \text{ otherwise,} \end{aligned} \tag{4.27}$$

where $\theta^\pm = \frac{|\theta| \pm \theta}{2}$.

- If $(x, y) \in \partial\Omega_\hbar^l$ then $P_\hbar((x, y), (x + \hbar, y)) = 1$.
- If $(x, y) \in \partial\Omega_\hbar^r$ then $P_\hbar((x, y), (x - \hbar, y)) = 1$.
- If $(x, y) \in \partial G_\hbar^u$ then (4.27) still holds but replacing

$$P_\hbar((x, y), (x, y + \hbar))$$

by $P_\hbar((x, y), (x, y))$. This ensure that the process does not leave Ω_\hbar .

- Similarly, if $(x, y) \in \partial\Omega_\hbar^d$ then (4.27) still holds but replacing $P_\hbar((x, y), (x, y - \hbar))$ by $P^\hbar((x, y), (x, y))$.

The four regions $\partial\Omega_\hbar^u$, $\partial\Omega_\hbar^d$, $\partial\Omega_\hbar^l$, and $\partial\Omega_\hbar^r$ in which we have divided the boundary of Ω_\hbar are shown in Figure 4.1.

The *interpolation interval* Δt_\hbar is defined by

$$\Delta t_\hbar = \frac{\hbar^2}{Q_\hbar},$$

which, defined this way, ensures that the first two moments of the Markov chain are consistent with those of the diffusion process.

We use this Markov chain representation to approximate the value function $V(x, y)$ by $V_\hbar(x, y)$ and the expected average cost g by g_\hbar . We notice that $V_\hbar(x, y)$ and g_\hbar are approximations restricted to region Ω_\hbar . Let $\{\pi_\hbar(x, y) : (x, y) \in \Omega_\hbar\}$ be the steady state probability distribution for the Markov chain. Given the instantaneous (singular) nature of the control processes we have that

$$g_\hbar = \sum_{(x,y) \in \Omega_\hbar} k(x, y; 0) \pi_\hbar(x, y), \quad (4.28)$$

where

$$k(x, y; \epsilon) := \begin{cases} (c(x) - \epsilon)\Delta t_\hbar & \text{if } (x, y) \in \Omega_\hbar - \partial\Omega_\hbar^l - \partial\Omega_\hbar^r, \\ (e^y + \alpha)\hbar & \text{if } (x, y) \in \partial\Omega_\hbar^l, \\ 0 & \text{if } (x, y) \in \partial\Omega_\hbar^r. \end{cases}$$

On the other hand, the value function satisfies

$$V_{\hbar}(x, y) = k(x, y; g_{\hbar}) + \sum_{(x', y') \in \Omega_{\hbar}} P_{\hbar}((x, y), (x', y')) V_{\hbar}(x', y'). \quad (4.29)$$

We notice that the solution to(4.29) is unique up to an additive constant. Thus, we can arbitrarily fixed the value of $V(\hat{x}, \hat{y})$ for an arbitrary state $(\hat{x}, \hat{y}) \in \Omega_{\hbar}$. Given the sparse nature of the transition matrix P_{\hbar} , we can solve efficiently (4.28) and (4.29) using Gaussian elimination. However, the dimension of the system can be extremely large and value iteration methods are also recommended.

The final step for updating Ω_{\hbar} is based on a numerical method proposed by Kumar and Muthuraman (2000). Their method is based on the PDE representation of the problem and the approximation of the value function $V_{\hbar}(x, y)$. First, we observe that the Markov chain approximation is such that (4.23) is asymptotically satisfied as $\hbar \rightarrow 0$. This is, in fact, in the nature of Kushner's method and its *finite difference* structure. Then, only conditions (4.24)-(4.26) need to be checked. For this purpose, we approximate $V_x(x, y)$ for any $(x, y) \in \Omega_{\hbar}$ by

$$V_x(x, y) \approx \frac{V_{\hbar}(x + \hbar, y) - V_{\hbar}(x, y)}{\hbar}.$$

Given our current solution Ω_{\hbar} , the boundaries $\partial\Omega_{\hbar}^l$ and $\partial\Omega_{\hbar}^r$ are the natural candidates for $\eta(y)$ and $\xi(y)$ respectively. Moreover, it turns out that by (4.29) we have

$$V_{\hbar}(x, y) = (e^y + \alpha) \hbar + V_{\hbar}(x + \hbar, y), \quad \forall (x, y) \in \partial\Omega_{\hbar}^l$$

which is equivalent to

$$\frac{V_{\hbar}(x + \hbar, y) - V_{\hbar}(x, y)}{\hbar} + e^y + \alpha = 0 \quad \text{or} \quad V_x(x, y) + e^y + \alpha = 0, \quad \forall (x, y) \in \partial\Omega_{\hbar}^l.$$

In the same way (4.29) implies

$$V_{\hbar}(x, y) = V_{\hbar}(x - \hbar, y), \quad \forall (x, y) \in \partial\Omega_{\hbar}^r.$$

That is, $V_x(x, y) = 0, \forall (x, y) \in \partial\Omega_{\hbar}^r$. Therefore, given the approximation for $V_x(x, y)$ we have that (4.24) is satisfied with $\{(\eta(y), y)\} = \partial\Omega_{\hbar}^l$ and $\{(\xi(y), y)\} = \partial\Omega_{\hbar}^r$. It only remains to check (4.25)-(4.26). Unless Ω_{\hbar} is optimal at least one of these conditions is

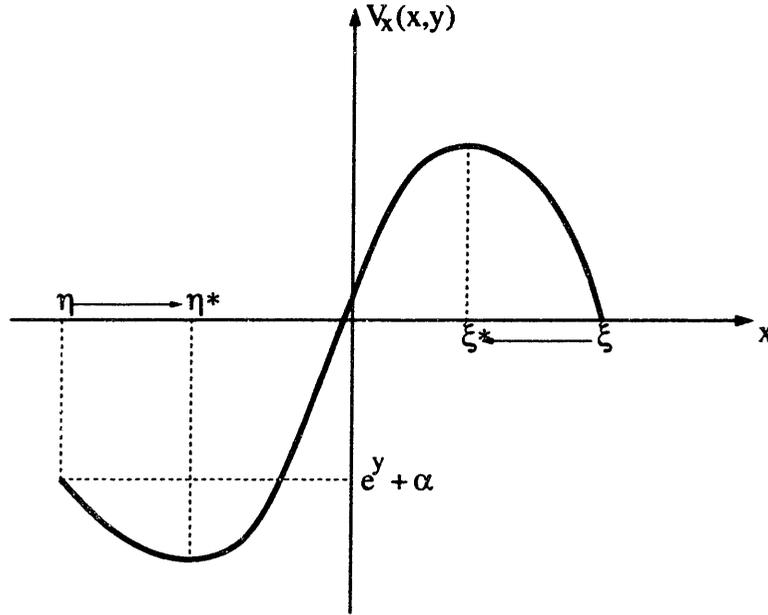


Figure 4.2: Boundary Update

violated and we need to update Ω_h . The following figure plots a prototypical example showing how conditions (4.25)-(4.26) are usually violated and how the update of Ω_h is done. Given the two-dimensional nature of the process, Figure 4.2 shows the update of Ω_h for a fixed value of y .

The solid line represents the value of $V_x(x, y)$. The abscissas η and ξ are the current values of the boundaries $\partial\Omega_h^l$ and $\partial\Omega_h^r$ at the level y . Since condition (4.25) is not satisfied in the example, we find the new values η^* and ξ^* looking for the minimum and maximum values of $V_x(x, y)$ ($V_{xx}(x, y) = 0$) in the range $[\eta, \xi]$. Thus, repeating this procedure for all values of y , we are able to update the values of $\partial\Omega_h^l$ and $\partial\Omega_h^r$.

We do not attempt here a proof of convergence of this iterative procedure. We only report that in all computational experiments that we have performed the method has found an approximated solution to (4.23)-(4.26) in few iterations. We conclude this section with an example that shows the convergence of the proposed method. In virtue of future discussion, we consider the following instance of the problem:

Example 1:

The data for this example are $\theta = 0$, $\sigma^2 = \delta^2 = 2$, $b = 2h = 2$, $\alpha = 3$, $N_x = 100$, $N_y = 40$, and $\hbar = 0.1$. In order to get an idea of the computational complexity of this instance, the solution to (4.28) and (4.29) requires the inversion of a $16,000 \times 16,000$

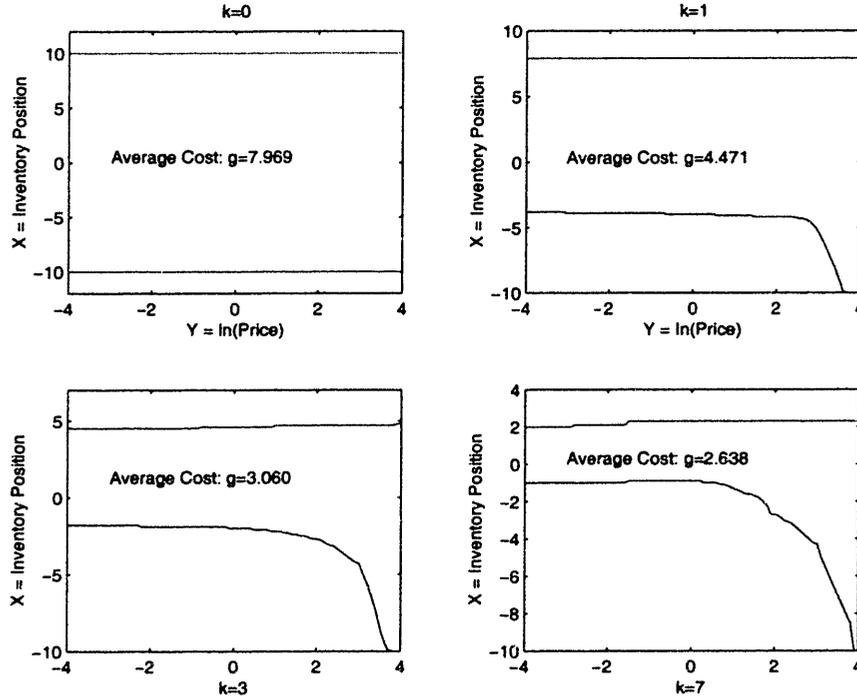


Figure 4.3: Numerical Solution to the Driftless Case.

matrix with 80,000 non-zero entries for the initial state space Ω_h^0 . Figure 4.3 plots the evolution of Ω_h^k . Notice that we have inverted the axis because the boundaries $\eta(y)$ and $\xi(y)$ are functions of y instead of x .

Initially ($k = 0$), Ω_h^0 is a rectangle defined by N_x , N_y , and h . As k increases, region Ω_h^k shrinks as we should expect given our previous discussion. The algorithm stops at $k = 7$ when the boundaries remain unchanged from the previous iteration. We observe that the upper boundary ($\xi(y)$) is approximately horizontal while the lower boundary ($\eta(y)$) has an exponential shape. If we consider that hitting $\xi(y)$ and $\eta(y)$ at level y has a cost of 0 and $\exp(y) + \alpha$ respectively, we can immediately imagine a few relations between the boundaries and the corresponding cost functions. We postpone, however, this discussion until we develop some additional insights about the problem. Finally, Figure 3 also shows the values of the expected average cost (g) for each region Ω_h^k . We observe that the improvements obtained during early iterations are significantly bigger than those obtained at the end.

4.7 Approximations

We devote this section to find approximate solutions to (4.23)-(4.26). Two types of approximations will be proposed. The first one -our *Proposed Policy*- uses as much as possible the optimality conditions but it modifies the reflection field on the boundary. With this transformation, we replace the PDE formulation by a simpler ODE system for which explicit solutions are presented. The second approximation -*Static Policy*- is computed assuming that the optimal policy is independent of the price. That is, production and admission decisions are based exclusively on the inventory position. We use this suboptimal policy because its simplicity makes it appealing from a practical point of view. Thus, it can be used as a benchmark for our *Proposed Policy*.

4.7.1 Proposed Policy

We turn now to our proposed solution. Let us rewrite here problem (4.23)-(4.26) in a more general form that will facilitate the exposition.

$$\theta V_x(x, y) + \frac{\sigma^2}{2} V_{xx}(x, y) + \frac{\delta^2}{2} V_{yy}(x, y) = g - C(x) \quad (4.30)$$

$$V_x(\xi(y), y) = H(y), \quad V_x(\eta(y), y) = -G(y), \quad (4.31)$$

where θ is the drift of the workload process, σ and δ are the diffusions of the workload and price processes respectively. The function $C(x)$ represents the holding/backordering cost rate while $H(y)$ and $G(y)$ are the boundary cost functions associated to stop production and to stop admission respectively. In this work, we consider the special instance $C(x) = h x^+ + b x^-$, $G(y) = \exp(y) + \alpha$, and $H(y) = 0$. These functions $C(x)$, $G(y)$, and $H(y)$ define the particular structure of the problem. We also notice that for the moment we are not imposing any continuity condition on $V(x, y)$ such as $V_{xx}(\xi(y), y) = V_{xx}(\eta(y), y) = 0$. Our objective is to find the switching curves $\xi(y)$ and $\eta(y)$ and the value of g such that (4.30) and (4.31) are satisfied and g is minimized.

Suppose that we fixed $\xi(y)$ and $\eta(y)$, then (4.30)-(4.31) becomes a Neumann problem¹ except for the fact that the boundary conditions do not involve *normal derivative* but *oblique derivative* on the direction of X . The key idea of our approximations is

¹For more details on the Neumann problem see John (1982).

to replace our particular boundary condition by the standard Neumann condition. That is, we will look for the solution to the following problem:

$$\theta V_x(x, y) + \frac{\sigma^2}{2} V_{xx}(x, y) + \frac{\delta^2}{2} V_{yy}(x, y) = g - C(x), \text{ for all } (x, y) \in \Omega \quad (4.32)$$

$$\frac{\partial V(x, y)}{\partial n} = F(y), \text{ for all } (x, y) \in \partial\Omega \quad (4.33)$$

where Ω , as it was introduced in section §(4.6), is a bounded domain given by

$$\Omega = \{(x, y) : \eta(y) \leq x \leq \xi(y), \quad y_1 \leq y \leq y_2\} \quad (4.34)$$

and n is the inward unit normal vector field on $\partial\Omega$. The boundary function $F(y)$ satisfies

$$F(y) = \begin{cases} H(y) & \text{if } x = \xi(y), \\ -G(y) & \text{if } x = \eta(y), \\ 0 & \text{otherwise.} \end{cases} \quad (4.35)$$

The rest of this section is organized as follows. First we normalized our PDE problem to what we call the *standard form*. Next, we present the approximate solution for the driftless case $\theta = 0$. Finally, we analyze the general case.

4.7.2 Standard Form

We normalized (4.30) and (4.31) by introducing new independent variables w and z through a real transformation:

$$w = w(x) = \frac{\sqrt{2}}{\sigma} x, \quad z = z(y) = \frac{\sqrt{2}}{\delta} y.$$

The problem in the new variables becomes

$$\vartheta V_w(w, z) + V_{ww}(w, z) + V_{zz}(w, z) = g - \hat{C}(w), \quad (4.36)$$

$$V_w(w, z) = \hat{F}(z), \quad (w, z) \in \partial\Omega, \quad (4.37)$$

where $\vartheta := \frac{\sqrt{2}\theta}{\sigma}$, $\hat{C}(w) := C\left(\frac{\sigma w}{\sqrt{2}}\right)$, and $\hat{F}(z) := F\left(\frac{\delta z}{\sqrt{2}}\right)$ (we also set $\hat{G}(z) := G\left(\frac{\delta z}{\sqrt{2}}\right)$ and $\hat{H}(z) := H\left(\frac{\delta z}{\sqrt{2}}\right)$).

Let Δ and ∇ be the Laplacian and gradient operators in \mathfrak{R}^2 . We set $\kappa := (\vartheta, 0) \in \mathfrak{R}^2$ the drift vector and $n(\beta) = (n_w(\beta), n_z(\beta))$ the inward unit vector field at $\beta \in \partial\Omega$.

We replace problem (4.36) and (4.37) by its standard form version:

$$(\Delta + \kappa \cdot \nabla)V(w, z) = g - \hat{C}(w), \quad (4.38)$$

$$\frac{\partial V(w, z)}{\partial n} \equiv (n(\beta) \cdot \nabla)V(w, z) = \hat{F}(z), \quad \text{for all } \beta \in \partial\Omega. \quad (4.39)$$

As we can see, the key approximation is to use at the boundary the reflection vector $n(\beta)$ instead of the true vector $(1, 0)$. The advantage of this approximation is that classic results such as Green's identity or the divergence theorem can be used directly to compute optimality conditions as we will see next.

4.7.3 Driftless Case

We present now the driftless case $\kappa = 0$. The main result is summarized in the following proposition.

Proposition 13 *For a fixed and bounded domain Ω with smooth boundary $\partial\Omega$ we have that the steady-state average cost satisfies:*

$$g = \pi \left[\iint_{\text{cl}(\Omega)} \hat{C}(w) dA + \oint_{\partial\Omega} \hat{F}(z) ds \right], \quad (4.40)$$

where $dA = dw dz$ is the element of surface in Ω , ds is the element of length in $\partial\Omega$, and π is the inverse of the area of the domain Ω (i.e. $\pi = \text{Area}(\Omega)^{-1}$).

Moreover, if Ω is given by (4.34) and $F(z)$ by (4.35) then

$$g = \pi \left[\int_{z_1}^{z_2} \int_{\eta(z)}^{\xi(z)} \hat{C}(w) dw dz + \int_{z_1}^{z_2} \left(\hat{H}(z) \sqrt{1 + \xi_z^2(z)} + \hat{G}(z) \sqrt{1 + \eta_z^2(z)} \right) dz \right]. \quad (4.41)$$

Proof:

For the driftless case the PDE problem becomes:

$$\begin{aligned} \Delta V(w, z) &= g - \hat{C}(w), \\ \frac{\partial V(w, z)}{\partial n} &= \hat{F}(z), \quad \text{for all } \beta \in \partial\Omega. \end{aligned}$$

In addition, for a bounded domain ω with smooth boundary $\partial\Omega$ and for two continuous functions $p(w, z)$ and $q(w, z)$ on $\text{cl}(\Omega)$, Green first identity states that

$$\oint_{\partial\Omega} p \frac{\partial q}{\partial n} ds = \int \int_{\text{cl}(\Omega)} (p \Delta q - \nabla p \cdot \nabla q) dA.$$

If we replace on Green's identity $p(w, z) = 1$ and $q(w, z) = V(w, z)$ we get

$$\oint_{\partial\Omega} \hat{F}(z) ds = \int \int_{\text{cl}(\Omega)} (g - \hat{C}(w)) dA.$$

From this relation (4.40) follows directly since $\int \int_{\text{cl}(\Omega)} g dA = g \text{Area}(\Omega)$.

Finally, (4.41) is a direct consequence of (4.40) and the representation of $\hat{F}(z)$ and Ω by (4.35) and (4.34) respectively. ■

The result of proposition (10) allows us to represent our original problem as a calculus of variation problem. In fact, the problem becomes now to identify the functions $\xi(z)$ and $\eta(z)$ that minimize g above. This is a variational problem that we can write in standard form as:

$$\min_{\xi, \eta} \pi \int_{z_1}^{z_2} \left[\underbrace{\int_{\eta(z)}^{\xi(z)} \hat{C}(w) dw + \hat{H}(z) \sqrt{1 + \xi_z^2(z)} + \hat{G}(z) \sqrt{1 + \eta_z^2(z)}}_M dz \right] \quad (4.42)$$

s.t

$$\int_{z_1}^{z_2} \underbrace{(\xi(z) - \eta(z))}_N dz = \frac{1}{\pi}. \quad (4.43)$$

Let $E = M - \gamma N$ be the corresponding Hamiltonian where γ is the Lagrangian multiplier for (4.43). Then, the Euler-Lagrange necessary conditions for optimality are (e.g., Gelfand (1963))

$$\frac{\partial E}{\partial \xi} - \frac{d}{dz} \left(\frac{\partial E}{\partial \xi_z} \right) = 0 \quad \text{and} \quad \frac{\partial E}{\partial \eta} - \frac{d}{dz} \left(\frac{\partial E}{\partial \eta_z} \right) = 0,$$

which in this case are equivalent to

$$\hat{C}(\xi(z)) = \gamma + \frac{d}{dz} \left(\frac{\hat{H}(z) \xi_z(z)}{\sqrt{1 + \xi_z^2(z)}} \right), \quad (4.44)$$

$$-\hat{C}(\eta(z)) = -\gamma + \frac{d}{dz} \left(\frac{\hat{G}(z) \eta_z(z)}{\sqrt{1 + \eta_z^2(z)}} \right). \quad (4.45)$$

We notice that (4.44)-(4.45) is a system of ordinary differential equations (ODE). That is, we have managed to replace a complex PDE system by a much standard and easier to solve ODE problem. Moreover, the system is diagonal in the sense that each equation has a single unknown and we can solve independently for ξ and η .

Further analysis of (4.44)-(4.45) requires more information about the function $\hat{C}(w)$, $\hat{H}(z)$, and $\hat{G}(z)$. Thus, we return now to our original formulation and replace those functions by their particular values. In this case system (4.44)-(4.45) becomes

$$\xi(z) = \frac{\sqrt{2} \gamma}{h \sigma} \quad (4.46)$$

$$\eta(z) = -\frac{\sqrt{2} \gamma}{b \sigma} + \frac{\sqrt{2}}{b \sigma} \frac{d}{dz} \left(\frac{(\alpha + e^{\frac{\delta z}{\sqrt{2}}}) \eta_z(z)}{\sqrt{1 + \eta_z^2(z)}} \right). \quad (4.47)$$

Equation (4.46) reveals that the production switching curve ($\xi(z)$) is a base-stock policy independent of the price. This result is particularly interesting because it is exactly the same observation that we have made based on the numerical solution in section §4.6. Thus, our approximation is able to capture the horizontal behavior of ξ .

On the other hand, $\eta(z)$ is the solution to the (4.47). This is a nonlinear second-order ordinary differential equation which we have not been able to solve in closed form. However, two general observations about (4.47) are important

1. $\eta(z) = -\frac{\sqrt{2} \gamma}{b \sigma}$ is a particular solution. Moreover, from the results of section §4.6 we can conjecture that this behavior is optimal for low values of z ($z \rightarrow -\infty$).
2. If $\eta_z^2(z) \gg 1$ then we can use the approximation

$$\eta(z) \approx -\frac{\sqrt{2} \gamma}{b \sigma} + \frac{\sqrt{2}}{b \sigma} \frac{d}{dz} \left((e^{\frac{\delta z}{\sqrt{2}}} + \alpha) \right) = -\frac{\sqrt{2} \gamma}{b \sigma} - \frac{\delta}{b \sigma} e^{\frac{\delta z}{\sqrt{2}}}. \quad (4.48)$$

We observe that for $z \rightarrow -\infty$, this approximation is consistent with our first observation $\eta(z) = \text{constant}$. In addition, for large value of z the assumption $\eta_z^2(z) \gg 1$ is satisfied. Thus, (4.48) asymptotically satisfies (4.47) when $z \rightarrow \pm\infty$.

It only remains to estimate the value of γ to complete the characterization of the *Proposed Policy*. For this purpose, we use a completely different approach. The idea here is to separate the influences of $\hat{C}(w)$, $\hat{G}(z)$, and $\hat{H}(z)$ from the general structure of our problem. For this purpose, we introduce the switching curves $\hat{\xi}(z)$ and $\hat{\eta}(z)$ together with the function $U(w, z)$ defined through

$$U_w(w, z) = \hat{H}(z) + \int_w^{\hat{\xi}(z)} (\hat{C}(w) - g - \lambda(z)) dt. \quad (4.49)$$

where $\lambda(z)$ is given by

$$\lambda(z) = \hat{C}(\hat{\xi}(z)) - g = \hat{C}(\hat{\eta}(z)) - g. \quad (4.50)$$

For the moment, all three functions $\lambda(z)$, $\hat{\xi}(z)$, and $\hat{\eta}(z)$ are unknown. Thus, (4.50) only partially characterizes $\lambda(y)$. We notice that (4.49) implies $U_{ww}(w, z) = (\lambda(y) - \hat{C}(w) + g)$. This relation together with (4.50) imply that $U_{ww}(\hat{\xi}(z), z) = U_{ww}(\hat{\eta}(z), z) = 0$. In addition, from (4.49) is clear that $U_w(\hat{\xi}(z), z) = \hat{H}(z)$. We can now select the value of $\lambda(y)$ solving $U_w(\hat{\eta}(z), z) = -\hat{G}(z)$, i.e.,

$$\int_{\hat{\eta}(z)}^{\hat{\xi}(z)} (\hat{C}(t) - g - \lambda(z)) dt + \hat{G}(z) + \hat{H}(z) = 0. \quad (4.51)$$

In summary, relations (4.49), (4.50), and (4.51) define the functions $U_w(w, z)$, $\hat{\eta}(z)$, $\hat{\xi}(z)$, and $\lambda(z)$ such that

$$\begin{aligned} U_w(\hat{\eta}(z), z) &= -\hat{G}(z), & U_w(\hat{\xi}(z), z) &= \hat{H}(z), \\ -\hat{G}(z) &\leq U_w(w, z) \leq \hat{H}(z), \\ U_{ww}(\hat{\eta}(z), z) &= U_{ww}(\hat{\xi}(z), z) = 0. \end{aligned}$$

That is, the function $U(w, z)$ satisfies the boundary conditions of our PDE formulation for the driftless case. Unfortunately, it can be shown that $U(w, z)$ does not satisfies the elliptic PDE (4.30). However, what makes this approximation appealing is the following result.

Proposition 14 *If $\delta = 0$ then the solutions $\hat{\eta}(z)$ and $\hat{\xi}(z)$ given by (4.50) and (4.51) are optimal for the driftless case.*

Proof: See the §4.10 at the end of this chapter.

This result reveals the myopic nature of this approximation. In fact, $\delta = 0$ implies that there is no variation on the price. Thus, (4.50)-(4.51) compute the switching curves $\hat{\eta}(z)$ and $\hat{\xi}(z)$ as if z were constant. It turns out that fixing z , our problem is equivalent to controlling a one-dimensional two-sided regulated Brownian motion (e.g., Harrison (1985) §5, Wein (1990)).

In our setting $\hat{C}(w) = h w^+ + b w^-$, $\hat{G}(z) = \exp(\delta z/\sqrt{2}) + \alpha$, and $H(z) = 0$, the solutions to (4.50) and (4.51) are

$$\lambda(z) + g = \sqrt{\frac{2hb(\hat{G}(z) + \hat{H}(z))}{h+b}}, \quad \hat{\xi}(z) = \frac{\lambda(z) + g}{h}, \quad \text{and} \quad \hat{\eta}(z) = -\frac{\lambda(z) + g}{b}. \quad (4.52)$$

Therefore,

$$\hat{\xi}(z) = \sqrt{\frac{2b(e^{\frac{\delta z}{\sqrt{2}}} + \alpha)}{h(h+b)}} \quad \text{and} \quad \hat{\eta}(z) = -\sqrt{\frac{2h(e^{\frac{\delta z}{\sqrt{2}}} + \alpha)}{b(h+b)}}.$$

A few observations are important with respect to these solutions. First, $\hat{\xi}(y)$ and $\hat{\eta}(y)$ are decreasing functions of h and b respectively as we might expect. For example, if h (the holding cost rate) increases, it becomes more expensive to hold inventory and the manager will stop producing at a lower level of inventory. Thus, $\hat{\xi}(y)$ decreases with h . A similar intuition applies to $\hat{\eta}(y)$ and b . Another observation is that the values of the boundaries depend on the sum of the boundary cost $\hat{G}(z)$ and $\hat{H}(y)$ and not on their distribution. This is in fact a direct consequence of the myopic nature of these switching curves. Figure 4.4 below plots the myopic solution $(\hat{\xi}(z), \hat{\eta}(z))$ for the Example 1 of the previous section. In Figure 4.4, the dashed lines correspond to the numerical solution. As we can see, for $z \rightarrow -\infty$ the approximation mimics the numerical solution. The reason is that for z small $\hat{G}(z) = \exp(\delta z/\sqrt{2}) + \alpha \approx \alpha$, i.e., it is approximately constant in this region. This is in fact the condition required by proposition (14). Thus, the solutions in (4.52) are asymptotically optimal as $z \rightarrow -\infty$. On the other hand, for large values of z the approximation is unable to represent the behavior of the optimal solution because of the high variability of $\hat{G}(z)$ in this range. However, we can use the asymptotic optimality when $z \rightarrow -\infty$ to compute the value of γ that is missing to completely describe the *Proposed Policy*.

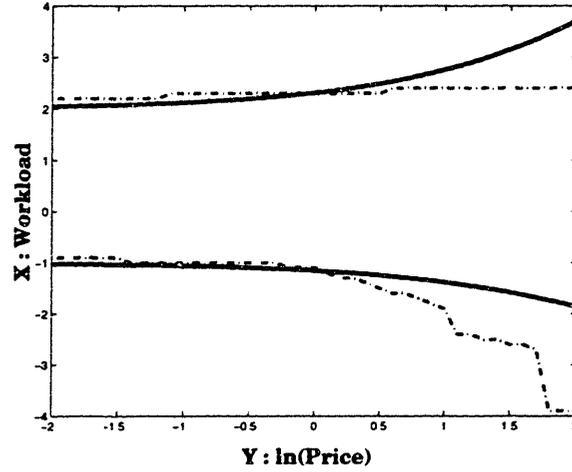


Figure 4.4: Myopic Approximation vs. Numerical Solution.

In fact, as $z \rightarrow -\infty$ the *Proposed Policy* is given by

$$\lim_{z \rightarrow -\infty} \xi(z) = \frac{\sqrt{2}\gamma}{h\sigma} \quad \text{and} \quad \lim_{z \rightarrow -\infty} \eta(z) = \frac{\sqrt{2}\gamma}{b\sigma}.$$

Similarly, as $z \rightarrow -\infty$ the myopic solution in (4.52) becomes

$$\lim_{z \rightarrow -\infty} \hat{\xi}(z) = \sqrt{\frac{2b\alpha}{h(h+b)}} \quad \text{and} \quad \lim_{z \rightarrow -\infty} \hat{\eta}(z) = \sqrt{\frac{2h\alpha}{b(h+b)}}.$$

From these two relations we conclude that

$$\gamma = \sigma \sqrt{\frac{hb\alpha}{h+b}}.$$

Finally, we can now give a full characterization of the *Proposed Policy* for the driftless case:

$$\xi(z) = \sqrt{\frac{2b\alpha}{h(h+b)}} \quad \text{and} \quad \eta(z) = -\sqrt{\frac{2h\alpha}{b(h+b)}} - \frac{\delta}{b\sigma} e^{\frac{\delta z}{\sqrt{2}}}.$$

Moreover, these boundary curves can be expressed in terms of our original variables (x, y) as follows

Proposed Policy

$$\boxed{\xi(y) = \sigma \sqrt{\frac{b\alpha}{h(h+b)}} \quad \text{and} \quad \eta(y) = -\sigma \sqrt{\frac{h\alpha}{b(h+b)}} - \frac{\delta}{\sqrt{2}b} e^y.} \quad (4.53)$$

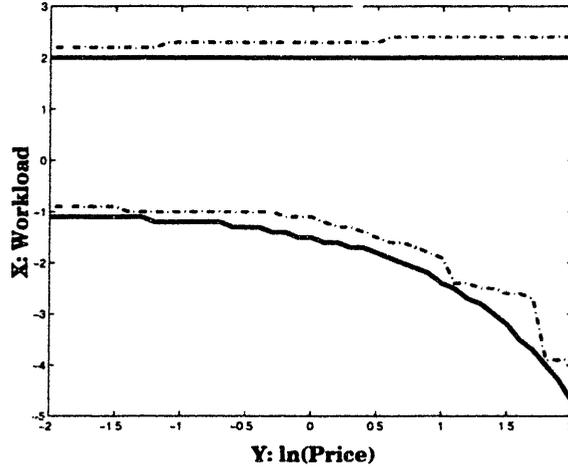


Figure 4.5: Proposed Policy: Driftless Case.

Figure 4.5 plots the boundaries in (4.53) and those obtained numerically. Two observations are important (i) the horizontal nature of $\xi(y)$ and (ii) the exponential behavior of $\eta(y)$. These results are consistent with those obtained numerically in section §4.6 as it is shown in the figure. Finally, let us comment on the approximation (4.48) above. This approximation, which is in fact asymptotically consistent with (4.45), is made because we cannot get closed form solutions for (4.45). Based on this observation, we can look for an improvement to (4.53) by considering a more general policy of the form

$$\xi(y) = \sigma \sqrt{\frac{b\alpha}{h(h+b)}} \quad \text{and} \quad \eta(y) = -\sigma \sqrt{\frac{h\alpha\sigma^2}{b(h+b)}} - K \frac{\delta}{\sqrt{2}b} e^y, \quad (4.54)$$

where $K > 0$ is a fixed parameter to be determined. In order to study the behavior of this modified policy we perform a numerical analysis. For a set of different scenarios (in terms of b , h , α , σ , and δ) we compute the expected average cost as a function of K . Let $g_i(K)$ be this expected average cost function for the i^{th} scenario. Figure 4.6 plots the results that we have obtained. In order to plot all the functions in the same graph, we normalize $g_i(K)$ and consider the functions

$$\tilde{g}_i(K) = \frac{g_i(K)}{\min_K \{g_i(K)\}},$$

which are those presented in the figure.

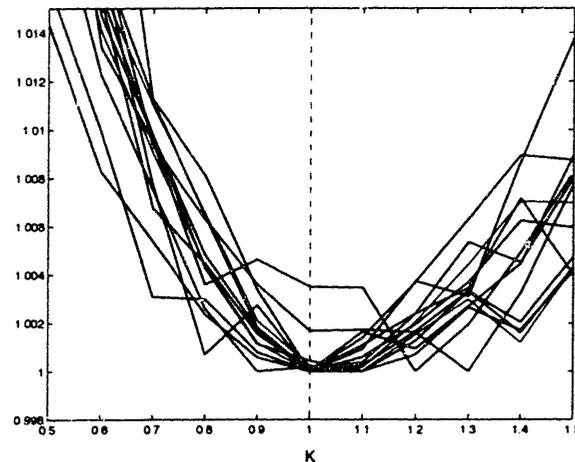


Figure 4.6: *Normalized Expected Average Cost ($\bar{g}(K)$) as a Function of K .*

We can see that the minimum value of $g_i(K)$ is attained in most of the cases at $K = 1$ or in the neighborhood of it. Moreover, $g_i(K)$ increases monotonically as $|K - 1|$ increases. This empirical result suggests that the boundaries given by (4.53) are the best within the family of policies in (4.54).

4.7.4 General Case

We now turn to the case of an arbitrary boundary vector κ . Much of our analysis is based on the paper by Harrison and Williams (1987). In this work, the authors study the steady-state distribution of a regulated Brownian motion that moves inside a bounded domain Ω . The special feature of their setting is the type of reflection on the boundary they assume. To be more precise, let us consider Figure 4.7. Here,

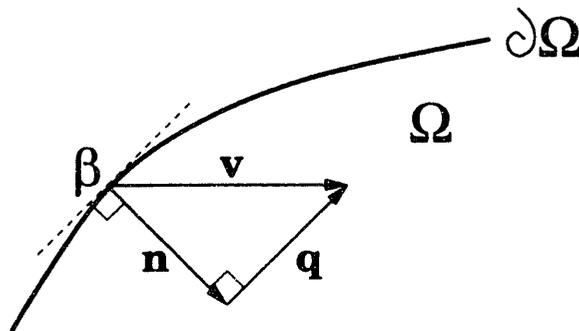


Figure 4.7: *Reflection Field*

n is the inward unit vector field and v is the reflection field on $\partial\Omega$. The tangential vector field q is given by $q = v - n$. The main result of Harrison and Williams is the following:

If the reflection field v satisfies the *skew symmetric* condition:

$$n(\beta) \cdot q(\beta^*) + q(\beta) \cdot n(\beta^*) = 0, \quad \text{for all } \beta, \beta^* \in \partial\Omega, \quad (4.55)$$

then the stationary distribution of the RBM with drift κ has a density of the *exponential form*

$$p(x) = K(\kappa) \exp\{\gamma(\kappa) \cdot x\}, \quad \text{for all } x \in \Omega.$$

The skew symmetric condition above is, in general, very restrictive. However, if we assume once again that we have normal reflection then $v = n$ and $q = 0$. In this situation (4.55) holds trivially. In particular, we have that in our setting

$$\gamma(\kappa) = \kappa \quad \text{and} \quad K(\kappa) = \left(\int \int_{\text{cl}(\Omega)} \exp\{\kappa \cdot x\} dA \right)^{-1}.$$

We now state the analogue of proposition (13) for the general case.

Proposition 15 *For a fixed and bounded domain Ω with smooth boundary $\partial\Omega$ we have that the steady-state average cost satisfies:*

$$g = K(\kappa) \left[\int \int_{\text{cl}(\Omega)} e^{\kappa \cdot x} \hat{C}(w) dA + \oint_{\partial\Omega} e^{\kappa \cdot x} \hat{F}(z) ds \right], \quad (4.56)$$

Moreover, if Ω is given by (4.34), $\hat{F}(z)$ by (4.35), and $\kappa = (\vartheta, 0)$ then

$$\begin{aligned} g &= K(\kappa) \left[\int_{z_1}^{z_2} \int_{\eta(z)}^{\xi(z)} e^{\vartheta w} \hat{C}(w) dw dz \right. \\ &\quad \left. + \int_{z_1}^{z_2} \left(e^{\vartheta \xi(z)} \hat{H}(z) \sqrt{1 + \xi_z^2(z)} + e^{\vartheta \eta(z)} \hat{G}(z) \sqrt{1 + \eta_z^2(z)} \right) dz \right]. \end{aligned} \quad (4.57)$$

Proof:

The proof in this case is almost identical to the one used for proposition (13). In this

case the PDE problem is:

$$\begin{aligned} (\Delta + \kappa \cdot \nabla)V(w, z) &= g - \hat{C}(w), \\ \frac{\partial V(w, z)}{\partial n} &= \hat{F}(z), \quad \text{for all } \beta \in \partial\Omega. \end{aligned}$$

Moreover, from Lemma 2.1 in Harrison and Williams (1987) we have that

$$\oint_{\partial\Omega} p(w, z) \frac{\partial f(w, z)}{\partial n} ds = \int \int_{\text{cl}(\Omega)} p(w, z) (\Delta + \kappa \cdot \nabla) f(w, z) dA, \quad \text{for all } f \in C^2(\text{cl}(\Omega)). \quad (4.58)$$

Here $p(w, z)$ is the steady-state distribution, i.e., the solution to the *adjoint problem*

$$\begin{aligned} (\Delta - \kappa \cdot \nabla)p(w, z) &= 0, \quad (w, z) \in \Omega, \\ \frac{\partial p(w, z)}{\partial n} &= 0, \quad (w, z) \in \partial\Omega. \end{aligned}$$

Relation (4.58) is in fact a direct consequence of Green's identity and the divergence theorem. If we replace $p(w, z) = K(\kappa) \exp(\kappa \cdot x)$ and $f(w, z) = V(w, z)$ in (4.58) we get (4.56). Again, (4.57) follows directly from (4.56) and the definitions of Ω and $\hat{F}(z)$. ■

Based on (4.57) we can write the following optimization problem:

$$\begin{aligned} \min_{\eta(z), \xi(z)} \quad & K(\kappa) \left[\int_{z_1}^{z_2} \int_{\eta(z)}^{\xi(z)} e^{\vartheta w} \hat{C}(w) dw dz \right. \\ & \left. + \int_{z_1}^{z_2} \left(e^{\vartheta \xi(z)} \hat{H}(z) \sqrt{1 + \xi_z^2(z)} + e^{\vartheta \eta(z)} \hat{G}(z) \sqrt{1 + \eta_z^2(z)} \right) dz \right], \end{aligned} \quad (4.59)$$

s.t.

$$\int_{z_1}^{z_2} \left(\frac{e^{\vartheta \xi(z)} - e^{\vartheta \eta(z)}}{\vartheta} \right) dz = \frac{1}{K(\kappa)}. \quad (4.60)$$

Since $\lim_{\vartheta \rightarrow 0} \vartheta^{-1} (\exp(\vartheta \xi(z)) - \exp(\vartheta \eta(z))) = \xi(z) - \eta(z)$, it is not hard to show that (4.42)-(4.43) is a special case of (4.59)-(4.60) with $\vartheta = 0$.

We look now at the solution to (4.59)-(4.60) for our particular case of interest.

Proposition 16 *The optimality Euler-Lagrange necessary conditions for (4.59)-(4.60)*

when $\hat{C}(w) = \frac{\sigma}{\sqrt{2}}(hw^+ + bw^-)$, $\hat{H}(z) = 0$ and $\hat{G}(z) = \exp(\delta z/\sqrt{2}) + \alpha$ are:

$$\frac{\sigma \xi(z)}{\sqrt{2}} = \frac{\gamma}{h}, \quad (4.61)$$

$$\begin{aligned} \frac{\sigma \eta(z)}{\sqrt{2}} &= -\frac{\gamma}{b} + \frac{\vartheta}{b} (e^{\frac{\delta z}{\sqrt{2}}} + \alpha) \sqrt{1 + \eta_z^2(z)} \\ &\quad - \frac{1}{b e^{\vartheta \eta(z)}} \frac{d}{dz} \left[\frac{e^{\vartheta \eta(z)} (e^{\frac{\delta z}{\sqrt{2}}} + \alpha) \eta_z(z)}{\sqrt{1 + \eta_z^2(z)}} \right]. \end{aligned} \quad (4.62)$$

where γ is the Lagrangian multiplier associated to (4.60).

The proof follows the same lines that we used for the driftless case and we omit it here. First, we notice that condition (4.61) reveals that the production switching curve is a base-stock policy independent of the price. Thus, $\xi(z)$ has the same behavior that in the driftless case. On the other hand, condition (4.62) is a highly nonlinear ODE that we have not been able to solve in closed form. However, we can study the asymptotic behavior of the solution for $z \rightarrow \pm\infty$.

In the first case, when $z \rightarrow -\infty$, the cost function $\hat{G}(z) \approx \alpha$. Thus, in this regime the price is constant and we can consider that both $\xi(z)$ and $\eta(z)$ are approximately constant. In particular, we can rewrite (4.59) assuming $\sqrt{1 + \xi_z^2(z)} \approx 1$ and $\sqrt{1 + \eta_z^2(z)} \approx 1$. Replacing this conditions in (4.59)-(4.60) and recomputing (4.61)-(4.62) we get

$$\frac{\sigma \xi(z)}{\sqrt{2}} = \frac{\gamma}{h}, \quad (4.63)$$

$$\frac{\sigma \eta(z)}{\sqrt{2}} = -\frac{\gamma + \vartheta \alpha}{b} - \frac{\vartheta e^{\frac{\delta z}{\sqrt{2}}}}{b}, \quad z \rightarrow -\infty. \quad (4.64)$$

Notice, moreover, that this solution satisfies our previous assumption $\sqrt{1 + \xi_z^2(z)} \approx 1$ and $\sqrt{1 + \eta_z^2(z)} \approx 1$ as $z \rightarrow -\infty$.

On the other hand, when $z \rightarrow \infty$ we can use the approximation $\eta_z^2(z) \gg 1$. In particular, we consider $\sqrt{1 + \eta_z^2(z)} \approx |\eta_z(z)| = -\eta_z(z)$ in (4.59). In this case and after some algebra, we can rewrite the optimality conditions (4.61)-(4.62) as follows

$$\frac{\sigma \xi(z)}{\sqrt{2}} = \frac{\gamma}{h}, \quad (4.65)$$

$$\frac{\sigma \eta(z)}{\sqrt{2}} = -\frac{\gamma}{b} - \frac{\delta}{\sqrt{2}b} e^{\frac{\delta z}{\sqrt{2}}}, \quad z \rightarrow \infty. \quad (4.66)$$

Again, this asymptotic solution is consistent with the assumption $\eta_z^2(z) \gg 1$ as $z \rightarrow \infty$. Although different, both asymptotic solutions (4.63)-(4.64) and (4.65)-(4.66) are quite similar. As $z \rightarrow -\infty$ the switching curves become price independent. On the other hand, as $z \rightarrow \infty$ the admission curve $\eta(z)$ exhibits an exponential decay.

We are once more left with the task of finding the optimal value of γ . We can again use the same technique that we have used for the driftless case. That is, we can look for the myopic solution obtained when we assume fixed prices. Since for $z \rightarrow -\infty$ the price function $\exp(z) + \alpha \approx \alpha$, we expect the fixed price assumption to be asymptotically optimal for $z \rightarrow -\infty$.

Proposition 17 *If the price is fixed and equal to α , then the optimal policy $(\xi_\alpha, \eta_\alpha)$ solves*

$$-\vartheta \alpha = h \xi_\alpha + b \eta_\alpha, \quad (4.67)$$

$$\begin{aligned} \exp\left(\frac{\vartheta^2 \alpha}{h+b}\right) &= \left(\frac{h}{h+b}\right) \exp\left(\frac{\vartheta b(\xi_\alpha - \eta_\alpha)}{h+b}\right) \\ &+ \left(\frac{b}{h+b}\right) \exp\left(\frac{-\vartheta h(\xi_\alpha - \eta_\alpha)}{h+b}\right). \end{aligned} \quad (4.68)$$

Proof: See §4.10 at the end of this chapter.

A closed-form solution for (4.67)-(4.68) is only available for the special case $h = b$ and it is given by

$$\xi_\alpha = \frac{\operatorname{arccosh}(S)}{|\vartheta|} - \frac{\ln(S)}{\vartheta} \quad \text{and} \quad \eta_\alpha = -\frac{\operatorname{arccosh}(S)}{|\vartheta|} - \frac{\ln(S)}{\vartheta}, \quad \text{where } S = \exp\left(\frac{\vartheta^2 \alpha}{2h}\right).$$

For the case ($h \neq b$) we can solve the system sequentially in two easy steps.

1. First, solve (4.68) for $I_\alpha = \xi_\alpha - \eta_\alpha$. A simple bisection algorithm can be used. For $\vartheta > 0$, we search on the interval $0 \leq I_\alpha \leq (h+b) \ln((h+b)h^{-1}S)(b\vartheta)^{-1}$. For $\vartheta < 0$ the search has to be conducted over the range $0 \leq I_\alpha \leq -(h+b) \ln((h+b)b^{-1}S)(h\vartheta)^{-1}$.

2. Once I_α has been found, the optimal solution is

$$\xi_\alpha = \frac{b I_\alpha - \vartheta \alpha}{h + b} \quad \text{and} \quad \eta_\alpha = -\frac{h I_\alpha + \vartheta \alpha}{h + b}.$$

As we discussed above, we expect this solution to be optimal as the price decreases ($z \rightarrow -\infty$). From this observation and (4.63)-(4.64) we get

$$\gamma = \frac{\sigma}{\sqrt{2}} \left(\frac{h(b I_\alpha - \vartheta \alpha)}{h + b} \right).$$

We can now present the *Proposed Policy* for the general case which in terms of the original variable (x, y) is given by

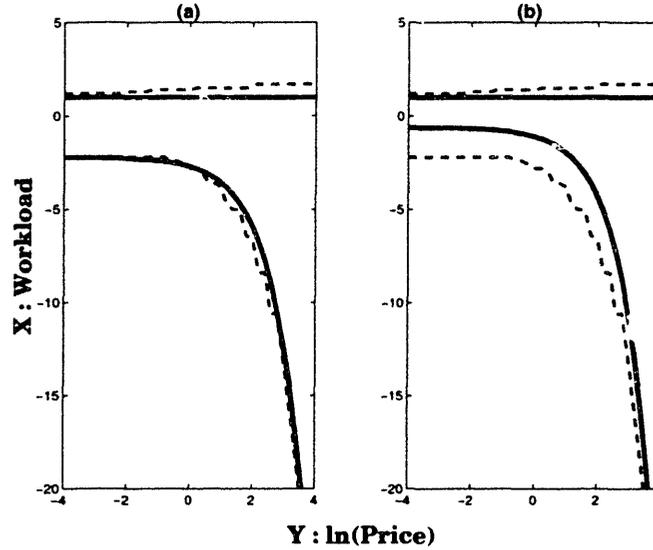
Proposed Policy for the General Case

$$\xi(y) = \frac{\sigma(b I_\alpha - \vartheta \alpha)}{\sqrt{2}(h + b)}, \quad \eta(y) = \begin{cases} -\left(\frac{\sigma(h I_\alpha + \vartheta \alpha)}{\sqrt{2}(h + b)} \right) - \frac{\vartheta}{b} \exp(y) & \text{if } z \rightarrow -\infty, \\ -\left(\frac{\sigma h(b I_\alpha - \vartheta \alpha)}{\sqrt{2} b(h + b)} \right) - \frac{\delta}{\sqrt{2} b} \exp(y) & \text{if } z \rightarrow \infty. \end{cases} \quad (4.69)$$

We notice that the driftless solution (4.53) is a particular case of (4.69) with

$$I_\alpha = \sigma \sqrt{\frac{\alpha(h + b)}{h b}} \quad \text{if } \vartheta = 0.$$

Figure 4.8 plots the boundaries $\xi(y)$ and $\eta(y)$ for the same data used in Example 1 in section §4.6 and a drift $\theta = 1$. The graph on the left (a) considers the approximation for $\eta(y)$ as $y \rightarrow -\infty$. On the other hand, (b) uses the result for $y \rightarrow \infty$. We can see that the upper boundary $\xi(y)$ mimics closely the behavior of the numerical solution (although it tends to be less accurate as y increases). On the other hand, the approximation for $\eta(y)$ in Figure 4.8(a) looks particularly good in this example for all ranges on y . The approximation in (b) does not fit well the numerical solution for low values of y but improves significantly as y increases (this result should be expected since (b) is the approximation obtained making $y \rightarrow \infty$).

Figure 4.8: *Proposed Policy: General Case*

4.8 Computational Experiments

In this section we study the performance of the policies presented in the previous section. In particular, we use the expected average cost (g) as the measure and the numerical solution of section §4.6 as our “optimal” solution. That is, for an arbitrary policy φ we define the *relative error* $\Delta g(\varphi)$ by

$$\Delta g(\varphi) = \frac{g(\varphi) - g(\text{numerical})}{g(\text{numerical})} \times 100,$$

where $g(\varphi)$ and $g(\text{numerical})$ are the expected average cost under policy φ and the numerical solution of section §4.6, respectively.

In order to get an idea of the performance of the proposed policies with respect to common practices in industry, we consider the simplest *static policy*. This policy is obtained by assuming that the admission and production decisions are independent of the price, that is, both $\eta(y)$ and $\xi(y)$ are constant. From a managerial perspective, this policy is easy to implement since it does not require any monitoring of the dynamics of the price. In addition, it also represents the solution that fixed priced models (such as those described in section §4.2) would suggest. It is not hard to show -specially after considering the solution in (4.52)- that the optimal *static policy* in the

interval $[y_1, y_2]$ is given by

$$\xi^{\text{static}}(y) = \sqrt{\frac{b(\bar{G})\sigma^2}{h(h+b)}}, \quad \eta^{\text{static}}(y) = -\sqrt{\frac{h(\bar{G})\sigma^2}{b(h+b)}} \quad (4.70)$$

for the driftless case $\theta = 0$ and by

$$\xi^{\text{static}}(y) = \frac{b\sigma I - \theta\bar{G}}{h+b}, \quad \eta^{\text{static}}(y) = -\frac{h\sigma I + \theta\bar{G}}{h+b} \quad (4.71)$$

for the general $\theta \neq 0$ case. \bar{G} is the average price

$$\bar{G} = \frac{\int_{y_1}^{y_2} (e^y + \alpha) dy}{y_2 - y_1} = \alpha + \frac{e^{y_2} - e^{y_1}}{y_2 - y_1}$$

and I is the corresponding amplitude² of the inventory position, which is the solution of the equation :

$$\exp\left(\frac{2\bar{G}\theta^2}{\sigma^2(h+b)}\right) = \left(\frac{h}{h+b}\right) \exp\left(\frac{2\theta bI}{\sigma(h+b)}\right) + \left(\frac{b}{h+b}\right) \exp\left(\frac{-2\theta hI}{\sigma(h+b)}\right).$$

The goal of the computational experiments that we have performed is to assess the quality of the proposed policy when used to solve the PDE problem

$$\begin{aligned} \theta V_x(x, y) + \frac{\sigma^2}{2} V_{xx}(x, y) + \frac{\delta^2}{2} V_{yy}(x, y) &= g - (hx^+ + b(-x)^+) \\ V_x(\xi(y), y) &= 0, \quad V_x(\eta(y), y) = -(R_0 e^y + \alpha). \end{aligned}$$

Equivalently, the real transformation presented in §4.7.2 allows us to rewrite the system in standard form as

$$\vartheta V_w(w, z) + V_{ww}(w, z) + V_{zz}(w, z) = g - (\hat{h} w^+ + \hat{b} (-w)^+), \quad (4.72)$$

$$V_w(\xi(z), z) = 0, \quad V_w(\eta(z), z) = -(\hat{R}_0 e^z + \hat{\alpha}), \quad (4.73)$$

(where $\hat{l} := \frac{\sigma l}{\sqrt{2}}$ and $\theta = \hat{\vartheta}$) thus reducing the number of parameters to consider.

In what follows we present a set of experiments measuring Δg for different values of the \hat{h} , \hat{b} , \hat{R}_0 , $\hat{\alpha}$ and $\hat{\vartheta}$. We will drop the hats from here on for ease of notation.

²We call *amplitude* the difference between the maximum level inventory and the maximum level of backorders that the system can ever have.

The first set of computational experiments studied the influence of the holding and backordering cost parameter on the performance of the proposed policies. Table 4.1 shows some of the results that we have obtained for the driftless case. Our first

b	$\Delta g(\text{proposed})$	$\Delta g(\text{static})$
1	3.47%	24.60%
2	2.78%	21.31%
3	2.52%	9.08%
4	2.26%	7.57%
5	2.49%	4.80%
6	2.14%	4.02%
7	2.05%	3.33%
8	2.10%	3.16%
9	1.94%	2.85%
10	1.91%	2.90%
Average	2.37%	5.94%

Table 4.1: Performance of the policies as a function of the backordering cost (b). The data used is $R_0 = 1$, $\alpha = 3$, $h = 1$, and $\sigma^2 = \delta^2 = 2$.

observation is that the proposed policy out performs the static policy. While the proposed policy has an average error of 2% the static policy has an average error above 10%. In addition, the quality of the approximations seems to improve as b increases.

The second analysis that we present aims to expose the relative effect of the two types of cost present, that is, the holding/backordering cost and the boundary cost associated to rejecting orders. The experiment that we conduct measures the performance of the proposed policy as a function of the relation between the holding/backordering cost $C(x) = hx^+ + b(-x)^+$ and the boundary cost $G(y) = (R_0 e^y + \alpha)$. In particular, we fix $G(y)$ and vary $C(x)$. The goal is to understand the effect of the relative weight of these two cost components on the quality of the proposed policy. We use the following simple scheme to vary $C(x)$

$$C(x) = k(hx^+ + b(-x)^+)$$

where the parameters h and b are kept fixed and $k > 0$ is changed. For low values of k the penalty of rejecting customers $G(y)$ is the main component of the cost function.

On the other hand, as k increases the manager cares more about holding inventory or backordering customers than rejecting orders. The following table presents the results that we have obtained. Notice that we have set the experiments for two value of the drift θ in order to studied simultaneously the impact of this parameter. The

	$\theta = 0$		$\theta = 1$	
k	$\Delta g(\text{proposed})$	$\Delta g(\text{static})$	$\Delta g(\text{proposed})$	$\Delta g(\text{static})$
0.5	1.04%	26.85%	0.06%	1.43%
1	2.61%	20.97%	0.23%	5.46%
2	3.88%	15.62%	2.22%	8.70%
4	4.72%	12.23%	2.91%	9.28%
8	5.20%	9.27%	3.21%	9.61%

Table 4.2: Performance as a function of the relative importance of the holding/backordering cost and boundary cost ($R_0 = 1$, $\alpha = 3$, $b = 2h = 2$, $\sigma^2 = \delta = 2$).

results of Table 4.2 shows that the proposed policy performs systematically better than the simplest static policy. However, as the holding/backordering cost increases with respect to the rejection penalty ($k \uparrow$) the difference between this two policy tends to diminish. The biggest difference are observed for the case for low values of k . It is also interesting to notice that in the present of a positive drift the quality of the proposed policy is better than in the driftless case.

Our next set of experiments are concerned with the quality of the solution as a function of the variability of the demand and production processes. We notice that since we have combined in a single process both demand and production, we can not isolate the variability of each of those processes. Rather, we have to look at both simultaneously. The parameter that characterizes this variability is σ the diffusion of the Brownian motion. The results in Table 4.1 show once again that the proposed policy is much better that the simplest static policy. Variability, on the other hand, seems to have a different impact on a driftless and a positive drift systems. On the driftless case, variability improves the performance of the proposed policy while in the case of a positive drift variance affects negatively the quality of the results. In general, the relative error of the proposed policy is about 2%, consistent with the results of the previous experiments.

σ^2	$\theta = 0$		$\theta = 1$	
	$\Delta g(\text{proposed})$	$\Delta g(\text{static})$	$g(\text{proposed})$	$\Delta g(\text{static})$
0.25	4.17%	77.31%	0.12%	2.38%
0.5	2.25%	42.00%	1.58%	5.99%
1	3.98%	24.99%	1.78%	9.99%
2	2.62%	20.95%	2.61%	13.41%
3	2.32%	25.65%	1.78%	14.92%
4	2.17%	32.12%	1.39%	16.34%

Table 4.3: Performance as a Function of the Demand and Production Variability ($R_0 = 1$, $\alpha = 3$, $b = 2h = 2$).

4.9 Extensions

We conclude this chapter looking at two additional results -that we have partially developed- concerning the general structure of the problem that we have analyzed so far. First, we studied in §4.9.1 the case that considers correlation between price and demand. Next, in §4.9.2 we look more closely at the specific type of approximations that we have used here. In particular, we quantify the error of replacing a general reflection field by the standard inward unit normal reflection.

4.9.1 Correlation Between Price and Demand

We consider in this section a natural extension to our basic model. In section §4.4 while building the diffusion model we made the assumption that $R(t)$ (the price process) and $D(t)$ (the demand process) were independent. Certainly, this is a major assumption since price and demand are tightly correlated in practice.

In this section, we relax this assumption and we extend our formulation to include correlation between $R(t)$ and $D(t)$. We tackle this problem, however, in a very simplistic way that we will explain shortly. First, we recall that according to the heavy traffic scaling we can represent the demand and price processes by diffusion processes. More precisely, what we have is that

$$D(t) \sim \sqrt{n} \lambda t + \sqrt{\lambda} c_a W_D(t) \quad \text{and} \quad R(t) \sim R_0 e^{\delta W_R(t)},$$

where $W_D(t)$ and $W_R(t)$ are two Wiener processes. Thus, the dependence between

$R(t)$ and $D(t)$ is reflected through the dependence between $W_D(t)$ and $W_R(t)$. Here, we assume the simplest type of dependence between this two processes, that is, we consider that the correlation between $W_D(t)$ and $W_R(t)$ is fixed and equal to $-\rho$, or more precisely we assume that the quadratic covariation $\langle W_D, W_R \rangle$ satisfies

$$\langle W_D, W_R \rangle = -\rho t. \quad (4.74)$$

Notice that $0 \leq \rho \leq 1$ does not have any relation with the traffic intensity of the system. In addition, the minus sign in (4.74) ensures that demand and price are in fact negatively correlated. In this new setting we have the following result.

Proposition 18 *The Hamilton-Jacobi-Bellman equation is*

$$[-V_x(x, y)] \wedge [V_x(x, y) + E^y + \alpha] \wedge [\Gamma V(x, y) + c(x) - g] = 0, \quad (4.75)$$

where

$$\Gamma = \theta \frac{\partial}{\partial x} + \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} + (\sigma \delta \rho) \frac{\partial^2}{\partial xy} + \frac{\delta^2}{2} \frac{\partial^2}{\partial y^2}$$

is the infinitesimal generator of $(X(t), Y(t))$, where $Y(t) = \ln(R(t))$.

We omit the proof of this result since it follows almost exactly the same lines of the proof of proposition (10). In PDE form, the problem becomes to find a value function $V(x, y)$, two boundary functions $\xi(y)$ and $\eta(y)$, and a the smallest scalar g that satisfy

$$\theta V_x(x, y) + \frac{\sigma^2}{2} V_{xx}(x, y) + (\sigma \delta \rho) V_{xy}(x, y) + \frac{\delta^2}{2} V_{yy}(x, y) = g - C(x) \quad (4.76)$$

$$V_x(x, y) = F(y), \quad \text{for all } (x, y) \in \partial\Omega. \quad (4.77)$$

This problem is almost equivalent to (4.30)-(4.31). In fact, for $\rho = 0$ those two problems are the same. The first step to solve this problem is to rewrite (4.76)-(4.77) on standard form (see section §4.7.2). We can achieve this goal if we find a real transformation $\{w = w(x, y); z = z(x, y)\}$ that satisfies the *Beltrami condition* (for details see for example John (1982) §2.3)

$$w_x = \frac{\sigma \rho z_x - \delta z_y}{\sigma \sqrt{1 - \rho^2}}, \quad w_y = \frac{\sigma z_x - \delta \rho z_y}{\delta \sqrt{1 - \rho^2}}.$$

It turns out that in this homogeneous case a simple rotation and scaling of the X and Y axes suffice to solve the Beltrami equation above. In fact, it can be shown that the

following transformation

$$w = \frac{1}{\sqrt{1-\rho}} \left(\frac{x}{\sigma} + \frac{y}{\delta} \right) \quad \text{and} \quad z = \frac{1}{\sqrt{1+\rho}} \left(\frac{x}{\sigma} - \frac{y}{\delta} \right).$$

reduces (4.76)-(4.77) to its standard form which is

$$\begin{aligned} (\Delta + \kappa \cdot \nabla)V(w, z) &= g - \hat{C}(w, z), \\ \frac{1}{\theta}(\kappa \cdot \nabla)V(w, z) &= \hat{F}(w, z) \quad \text{for all } (w, z) \in \partial\Omega, \end{aligned}$$

where drift (κ) is given by

$$\kappa = \frac{\theta}{\sigma} \left(\frac{1}{\sqrt{1-\rho}}, \frac{1}{\sqrt{1+\rho}} \right).$$

The cost function $\hat{C}(w, z)$ and the boundary function $\hat{F}(w, z)$ are computed using the transformation of coordinates above. Finally, we approximate this system by replacing the reflection field by the inward unit vector field n to get

$$(\Delta + \kappa \cdot \nabla)V(w, z) = g - \hat{C}(w, z), \tag{4.78}$$

$$(n \cdot \nabla)V(w, z) = \hat{F}(w, z) \quad \text{for all } (w, z) \in \partial\Omega, \tag{4.79}$$

We can now solve this problem using the same techniques that we used in section §4.7.4.

4.9.2 About the Error of the Proposed Policy

An important question that we have not formally addressed so far is related to the quality of our approximations. We have partially addressed this issue numerically comparing the results of our proposed policy and those of the numerical solution. In this section, we look more closely to this problem. We restrict, however, the scope of our analysis looking only at the driftless case which can be extended to the general case without much difficulty.

As a summary of our approximations, we know that the only structural difference between our original problem and the one we solve is the vector field that we use to reflect the diffusion process at the boundary. Suppose that we fix the domain Ω and its boundary $\partial\Omega$. Let \mathbf{u} be the original vector field and \mathbf{n} be the vector field that we

use to replace \mathbf{u} . In our particular case, we have that \mathbf{u} is the outward vector field point in the direction of X and \mathbf{n} is the outward normal vector field. Let $C(x)$ be the holding/backordering cost function and $G(y)$ be the boundary cost function. The generic problem for an arbitrary vector field \mathbf{w} is given by

$$\begin{aligned} & \text{Problem } P^{\mathbf{w}} \\ & \nabla^2 V^{\mathbf{w}}(x, y) = g^{\mathbf{w}} - C(x), \quad \text{in } \Omega, \\ & \mathbf{w} \cdot \nabla V^{\mathbf{w}}(x, y) = G(y), \quad \text{in } \partial\Omega, \end{aligned}$$

where $g^{\mathbf{w}}$ is the average cost under the reflection \mathbf{w} and $V^{\mathbf{w}}(x, y)$ is the corresponding value function.

We are interested in the solution to $P^{\mathbf{u}}$, that is, we look for a pair $(V^{\mathbf{u}}(x, y), g^{\mathbf{u}})$ that satisfies the system above when $\mathbf{w} = \mathbf{u}$. In particular, since we try to minimize g , we primarily care about the value of $g^{\mathbf{u}}$. The exact solution to this problem is in general hard to find. However, we can solve “easily” problem $P^{\mathbf{n}}$ and find $g^{\mathbf{n}}$ without much difficulty. Thus, if $g^{\mathbf{u}}$ and $g^{\mathbf{n}}$ are *similar* then we might view our approximation (i.e., replacing \mathbf{u} by \mathbf{n}) as appropriate. In order to be more precise about what *similar* means, we define the following criteria to compare the solutions. We define the relative error $\Delta(\mathbf{u}, \mathbf{n})$ as follows:

$$\Delta(\mathbf{u}, \mathbf{n}) = \frac{g^{\mathbf{u}} - g^{\mathbf{n}}}{g^{\mathbf{n}}}.$$

Our main result is the following.

Proposition 19

$$\Delta(\mathbf{u}, \mathbf{n}) = \frac{\int_{\partial\Omega} \mathbf{q} \cdot \nabla V^{\mathbf{u}}(x, y) d\sigma}{\int_{\Omega} C(x) ds + \int_{\partial\Omega} G(y) d\sigma} = \frac{\int_{\partial\Omega} (\mathbf{n} \cdot \nabla V^{\mathbf{u}}(x, y) - G(y)) d\sigma}{\int_{\Omega} C(x) ds + \int_{\partial\Omega} G(y) d\sigma}, \quad (4.80)$$

where $\mathbf{q} = \mathbf{n} - \mathbf{u}$, $d\sigma$, and ds are the element of length and surface on $\partial\Omega$ and Ω respectively.

Proof: See §4.10 at the end of the chapter. Before proving proposition (19) we observe the following properties of the solution.

- $\Delta(\mathbf{u}, \mathbf{n}) = 0$. This condition holds trivially if $\mathbf{q} = \mathbf{0}$. It can also hold for $\mathbf{q} \neq \mathbf{0}$ if for example $\mathbf{q} \cdot \nabla V^{\mathbf{u}}(x, y) = 0$ in $\partial\Omega$.

- More interesting, perhaps, is the fact that the nominator in (4.80) is a line integral which depends only on the value of the functions on the boundary and the length of the boundary. On the other hand, the denominator depends on both the boundary and the surface of Ω . Thus, if the boundary is well-behaved and the different functions are bounded we will expect the ratio in (4.80) to decrease as the area of Ω increases. In other, and more intuitive, words, the boundary conditions become negligible as the area of Ω increases.

In order to proceed further, we need to give more structure to the domain Ω and its boundary $\partial\Omega$. For this purpose and based on our particular problem, we consider the following region

$$\Omega = \{(x, y) \text{ s.t. } \eta(y) \leq x \leq \xi(y) \text{ and } y_1 \leq y \leq y_2\}.$$

We will, now, look at the behavior of $\Delta(\mathbf{u}, \mathbf{n})$ for a given pair of functions (η, ξ) when the length interval $[y_1, y_2]$ increases. In order to keep the same structure that we have in our problem, we will consider that the boundary function $G(y)$ is nonzero only in the lower boundary $(\eta(y), y)$. If we define α as the angle between \mathbf{u} and \mathbf{n} then we can rewrite the nominator in (4.80) as follows:

$$\int_{\partial\Omega} (\mathbf{n} \cdot \nabla V^{\mathbf{u}}(x, y) - G(y)) d\sigma = \int_{\partial\Omega} (\sin(\alpha)V_y^{\mathbf{u}}(x, y) - \cos(\alpha)V_x^{\mathbf{u}}(x, y) - G(y)) d\sigma.$$

Moreover, on the boundary we know that $V_x^{\mathbf{u}}(x, y) = -G(y)$. Thus,

$$\Delta(\mathbf{u}, \mathbf{n}) = \frac{\int_{\partial\Omega} (\sin(\alpha)V_y^{\mathbf{u}}(x, y) - (1 - \cos(\alpha))G(y)) d\sigma}{\int_{\Omega} C(x) ds + \int_{\partial\Omega} G(y) d\sigma}.$$

We can see from this last relation that that the holding/backordering cost $C(x)$ appears only in the denominator. Thus, we expect $\Delta(\mathbf{u}, \mathbf{n})$ to be decrease in $C(x)$, this result was in fact observed in the computational experiments performed in §4.8. Further analysis of the error requires the characterization of the partial derivative $V_y^{\mathbf{u}}(x, y)$ on $\partial\Omega$. We postpone this issue for a future research but we can mentioned that empirical observations suggest that the error $\Delta(\mathbf{u}, \mathbf{n})$ decreases with the magnitude of $\text{Area}(\Omega)$.

4.10 Proofs

4.10.1 Proof of Proposition (10)

The notation used in this proof is based on chapter 4 in Harrison (1985). We first check that $R(t) \in H^2$, that is

$$E \left[\int_0^t R^2(s) ds \right] < \infty, \quad \forall t \geq 0.$$

Let recall that $R(t) = R_0 e^{\delta W_R(t)}$ is a geometric Brownian motion. Now, from Fubini's theorem we have

$$E \left[\int_0^t R^2(s) ds \right] = \int_0^t E \left[e^{2\delta W_R(s)} \right] ds.$$

In addition, $W_R(s)$ is normally distributed with 0 mean and variance equal to s , thus $E[e^{2\delta W_R(s)}] = e^{\frac{\delta^2}{2}s}$ and

$$E \left[\int_0^t R^2(s) ds \right] = 2 \left(e^{\delta^2 t} - 1 \right) < \infty \quad \forall t \geq 0.$$

Since $R(t)$ is adapted and belongs to H^2 , we know that there exists a sequence of adapted processes $\{R^n(t)\} \subseteq S^2$ (*simple processes*) such that $R^n(t) \rightarrow R(t)$ (see §4 proposition (7) on Harrison (1985)). Given that $R^n(t)$ is a simple process, there exists a partition $\pi_n = \{t_0, t_1, \dots, t_n\}$ such that $R^n(s, \omega) = R^n(t_{k-1}, \omega)$ for all $s \in [t_{k-1}, t_k)$, $t_0 = 0$, $t_n = t$, and $mesh(\pi_n) \rightarrow 0$. Then,

$$I(R^n) \equiv \int_0^t R^n(s) dW_D(\bar{A}(s)) = \sum_{k=0}^{n-1} R^n(t_k) [W_D(\bar{A}(t_{k+1})) - W_D(\bar{A}(t_k))] \quad a.s.$$

Let $\{F_k = F(t_k)\}$, where $\mathcal{F} = \{F(t) : t \geq 0\}$ is the corresponding *filtration* on the probability space. Then,

$$E [I(R^n)] = \sum_{k=0}^{n-1} E \left\{ R^n(t_k) E \left[W_D(\bar{A}(t_{k+1})) - W_D(\bar{A}(t_k)) \middle| F_k \right] \right\}.$$

But $W_D(t)$ is a martingale and $\bar{A}(t)$ is random time change process (continuous and non-decreasing), therefore

$$E [W_D(\bar{A}(t_{k+1})) - W_D(\bar{A}(t_k)) | F_k] = 0$$

and $E [I(R^n)] = 0$ for all n . Finally, we apply §4 proposition (11) in Harrison (1985) to conclude that $E [I(R)] = 0$. ■

4.10.2 Proof of Proposition (11)

Let $\tilde{Z}^*(t)$ be the inventory position under the optimal control $(\tilde{U}^*, \tilde{I}^*)$, i.e., $\tilde{Z}^*(t) = X(t) + \lambda \tilde{U}^*(t) - \mu \tilde{I}^*(t)$. Then, for any policy (\tilde{U}, \tilde{I}) such that $\tilde{Z}(t) = X(t) + \lambda \tilde{U}(t) - \mu \tilde{I}(t)$ we have

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T (\bar{R}(t) + \alpha) \lambda d\tilde{U}^*(t) + \int_0^T c(\tilde{Z}^*(t)) dt \right] &\leq \\ \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T (\bar{R}(t) + \alpha) \lambda d\tilde{U}(t) + \int_0^T c(\tilde{Z}(t)) dt \right]. \end{aligned}$$

Since α is fixed and \tilde{U}^* satisfies (4.13), we get the following equivalent condition

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T \bar{R}(t) \lambda d\tilde{U}^*(t) + \int_0^T c(\tilde{Z}^*(t)) dt \right] + \alpha \lambda \bar{U} &\leq \\ \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T \bar{R}(t) \lambda d\tilde{U}(t) + \int_0^T c(\tilde{Z}(t)) dt + \alpha \lambda \tilde{U}(T) \right]. \end{aligned}$$

Moreover, if (\tilde{U}, \tilde{I}) is a feasible solution to the original problem (4.8)-(4.10) then \tilde{U} also satisfies (4.13) and

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T \bar{R}(t) \lambda d\tilde{U}^*(t) + \int_0^T c(\tilde{Z}^*(t)) dt \right] + \alpha \lambda \bar{U} &\leq \\ \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T \bar{R}(t) \lambda d\tilde{U}(t) + \int_0^T c(\tilde{Z}(t)) dt \right] + \alpha \lambda \bar{U}. \end{aligned}$$

After cancelling $\alpha \lambda \bar{U}$ in both side of the inequality above, we conclude that $(\tilde{U}^*, \tilde{I}^*)$ is an optimal solution to (4.8)-(4.10). ■

4.10.3 Proof of Proposition (12)

We consider a small neighborhood $[x - \epsilon, x + \epsilon]$, and the stopping time

$$T_\epsilon = \inf\{t \geq 0 : |X(t) - X(0)| = \epsilon\}.$$

The idea is that the scheduler when facing the state (x, r) can take three decision (i) he/she can increase I moving *instantaneously* to $(x - \epsilon, r)$, (ii) he/she can increase U moving to $(x + \epsilon, r)$, and (iii) he/she can do nothing until T_ϵ and re-evaluate what to do at that time.

In the first case the objective function is given by

$$f(x - \epsilon, r) = f(x, r) - f_x(x, r)\epsilon + o(\epsilon). \quad (4.81)$$

In the second case we have

$$r\epsilon + f(x + \epsilon, r) = f(x, r) + [f_x(x, r) + r]\epsilon + o(\epsilon). \quad (4.82)$$

Finally, in the third case

$$E_{(x,r)} \left[\int_0^{T_\epsilon} e^{-\gamma t} c(X(t)) dt + e^{-\gamma T_\epsilon} f(X(T_\epsilon), R(T_\epsilon)) \right]. \quad (4.83)$$

Now, using the following general result: For any continuous function $A(\cdot)$

$$\lim_{\epsilon \downarrow 0} \frac{E \left[\int_0^{T_\epsilon} A(X(t)) dt \right]}{E[T_\epsilon]} = A(X(0)), \quad \lim_{\epsilon \downarrow 0} \frac{\sigma^2 E[T_\epsilon]}{\epsilon^2} = 1$$

(where σ^2 is the variance of X). We can rewrite (4.83) when $\epsilon \downarrow 0$ as follows

$$f(x, r) + \left(c(x) + \frac{E_{(x,r)} \left[e^{-\gamma T_\epsilon} f(X(T_\epsilon), R(T_\epsilon)) - f(x, r) \right]}{E[T_\epsilon]} \right) \left(\frac{\epsilon}{\sigma} \right)^2 + o \left(E_{(x,r,l)}[T_\epsilon] \right).$$

Since f is twice continuously differentiable, $f(X, R)$ is an Itô process. Then applying integration by parts (see Harrison (1985) §4.9 proposition (2)) we have

$$e^{-\gamma t} f(X(t), R(t)) = f(x, r) + \int_0^t e^{-\gamma s} df(X(s), R(s)) - \gamma \int_0^t e^{-\gamma s} f(X(s), R(s)) ds.$$

In addition, from Itô's lemma

$$\begin{aligned} df(X, R) &= f_x dX + f_r dR \\ &+ \frac{1}{2} (f_{xx} dX + f_{xr} dR) dX \\ &+ \frac{1}{2} (f_{rx} dX + f_{rr} dR) dR. \end{aligned}$$

Moreover, we know that

$$dX(t) = \theta dt + \sigma dW_x(t), \quad dR(t) = \frac{\delta^2 R(t)}{2} dt + \delta R(t) dW_r(t),$$

where the processes $W_x(t)$ and $W_r(t)$ are two independent Wiener processes. Thus, using the above representation for X and R and the differential rule for quadratic variations³, we have

$$df(X, R) = \underbrace{\left[\theta f_x \frac{\delta^2 R(t)}{2} f_r + \frac{\sigma^2}{2} f_{xx} + \frac{(\delta R(t))^2}{2} f_{rr} \right]}_{\tilde{\Gamma} f(X, R)} dt + \sigma f_x dW_x + \delta R(t) f_r dW_r, \quad (4.84)$$

or equivalently

$$df(X, R) = \tilde{\Gamma} f(X, R) dt + \sigma f_x dW_x + \delta R(t) f_r dW_r.$$

Therefore,

$$\begin{aligned} e^{-\gamma t} f(X(t), R(t)) &= f(x, r, l) + \int_0^t e^{-\gamma s} (\tilde{\Gamma} f - \gamma f)(X(s), R(s)) ds \\ &+ \underbrace{\int_0^t e^{-\gamma s} \sigma f_x dW_x(s) + \int_0^t e^{-\gamma s} \delta R(s) f_r dW_r(s)}_{M(t)}. \end{aligned}$$

Now, taking expectation in the previous expression and noticing that the integrands on the two last integrals are bounded, we have that $E[M(t)] = 0$, therefore

$$\frac{E_{(x,r)} \left[e^{-\gamma T_\epsilon} f(X(T_\epsilon), R(T_\epsilon)) - f(x, r) \right]}{E[T_\epsilon]} = \frac{\int_0^t e^{-\gamma s} (\tilde{\Gamma} f - \gamma f)(X(s), R(s)) ds}{E[T_\epsilon]}.$$

³That is, $d \langle t, t \rangle = dt$, $d \langle t, W \rangle = 0$ and $d \langle W, W \rangle = dt$.

Combining the previous result, we can replace (4.83) by

$$f(x, r) + \left(c(x) + \tilde{\Gamma} f(x, r) - \gamma f(x, r) \right) \left(\frac{\epsilon}{\sigma} \right)^2 + o\left(E_{(x,r)}[T_\epsilon] \right). \quad (4.85)$$

The optimality of $f(x, r)$ together with (4.81), (4.82), and (4.85) imply that f satisfies

$$\begin{aligned} f(x, r) &= \min\{f(x, r) - f_x(x, r)\epsilon + o(\epsilon), \\ & f(x, r) + [f_x(x, r) + r]\epsilon + o(\epsilon), \\ & f(x, r) + \left(c(x) + \tilde{\Gamma} f(x, r) - \gamma f(x, r) \right) \left(\frac{\epsilon}{\sigma} \right)^2 + o\left(E_{(x,r)}[T_\epsilon] \right). \end{aligned}$$

Subtracting $f(x, r)$ from both side an letting $\epsilon \downarrow 0$, we conclude that $f(x, r)$ solves

$$0 = [f_x^\gamma(x, r)] \wedge [f_x^\gamma(x, r) + r] \wedge [c(x) + \tilde{\Gamma} f^\gamma(x, r) - \gamma f^\gamma(x, r)].$$

■

4.10.4 Proof of Proposition (14)

If $\delta = 0$, the price is not changing over time. Therefore, for each value of z there are optimal thresholds values $\eta(z)$ and $\xi(z)$ used to control the one-dimensional Brownian motion process X . From standard results (e.g., Harrison (1985), section §5.5), the steady state distribution of the driftless Brownian motion on the interval $[\eta(z), \xi(z)]$ is Uniform and the average amount of control per unit time used to keep the process within the range is

$$\frac{1}{(\xi(z) - \eta(z))},$$

in both extremes. Since $\hat{G}(z) = \exp(z) + \alpha$ and $\hat{H}(z) = 0$ are the cost of reflecting X to the right and to left respectively, we get the following expression for the average cost per unit time

$$g(z) = \frac{\int_{\eta(z)}^{\xi(z)} \hat{C}(w) dx + (\hat{G}(z) + \hat{H}(z))}{(\xi(z) - \eta(z))}.$$

Finally, in this case (z fixed) we maximize $g(z)$ solving

$$\frac{\partial g(z)}{\partial \xi(z)} = 0 \quad \text{and} \quad \frac{\partial g(z)}{\partial \eta(z)} = 0.$$

Solving this system gives the desired solution. ■

4.10.5 Proof of Proposition (17)

The proof uses the fact that the steady state distribution of a $(\vartheta, 1)$ Brownian motion in the interval $[a, a + I]$ is a truncated exponential

$$p(x) = \frac{\vartheta e^{\vartheta(x-a)}}{e^{\vartheta I} - 1} \text{ for } a \leq x \leq a + I.$$

For a proof of this result see §A.5.3 in the Appendix A or chapter 5 Harrison (1985). Moreover, the local times on the left ($x = a$) and right ($x = a + I$) boundaries are given by

$$LT(a) = \frac{\vartheta}{e^{\vartheta I} - 1} \text{ and } LT(I + a) = \frac{\vartheta}{1 - e^{-\vartheta I}},$$

respectively. Thus, in order to find the optimal values for a and I , the following optimization problem has to be solve

$$\min_{a, I \geq 0} \left\{ \int_0^{I+a} h x p(x) dx - \int_a^0 b x p(x) dx + \alpha \left(\frac{\vartheta}{e^{\vartheta I} - 1} \right) \right\}.$$

From the first order optimality conditions and after some algebra we recover conditions (4.67) and (4.68). ■

4.10.6 Proof of Proposition (19)

The proof is a straightforward application of Green's identity. In fact, Green's identity on problem P^n implies

$$\int_{\Omega} (g^n - C(x)) d\sigma = \int_{\partial\Omega} G(y) ds.$$

Letting π to be the inverse of the area of Ω , the condition above is equivalent to

$$g^n = \pi \left[\int_{\Omega} C(x) d\sigma + \int_{\partial\Omega} G(y) ds \right].$$

Similarly, problem P^u can be written in standard form as follows:

$$\begin{aligned} \nabla^2 V^u(x, y) &= g^u - C(x), & \text{in } \Omega, \\ \mathbf{n} \cdot \nabla V^u(x, y) &= G(y) + \mathbf{q} \cdot \nabla V^u(x, y) & \text{in } \partial\Omega. \end{aligned}$$

Thus, Green's identity implies

$$g^{\mathbf{u}} = \pi \left[\int_{\Omega} C(x) d\sigma + \int_{\partial\Omega} (G(y) + \mathbf{q} \cdot \nabla V^{\mathbf{u}}(x, y)) ds \right].$$

Combining the values of $g^{\mathbf{u}}$ and $g^{\mathbf{n}}$ we get the first equality in (4.80). The second equality is obtained replacing $\mathbf{q} = \mathbf{n} - \mathbf{u}$ and $\mathbf{u} \cdot \nabla V^{\mathbf{u}}(x, y) = G(y)$ in $\partial\Omega$. ■

Chapter 5

Conclusions

In this thesis we have presented two applications of the traditional make-to-stock queue model related to the fields of supply chain management and electronic commerce. As a general comment, it is our view that the make-to-stock queue is an attractive operations management model to embed into a game-theoretic framework or into an optimal admission control problem. The model is in most ways richer than the newsvendor model and is about as complex as – but considerably more tractable than – a two-stage Clark-Scarf model. It also allows us to capture the nonlinear effect of capacity and the impact of the retailer’s order process on the supplier’s lead times. Of course, none of these models attempt to mimic the complexities of an actual supply chain. Nevertheless, to the extent that queueing effects are present in manufacturers’ production facilities, the make-to-stock queue is a parsimonious and tractable model for deriving new insights into multi-agent models for supply chain management.

Regarding the decentralized make-to-stock model, the distinguishing feature of our simple supply chain model is that congestion at the supplier’s manufacturing facility is explicitly captured via a single-server queue. Each agent has a resource at his disposal (the supplier chooses the capacity level and the retailer chooses the base stock level) that buffers against expensive backorders of the retailer’s inventory. When the inventory backorder cost is incurred entirely by the retailer (i.e., the backorder allocation fraction $\alpha = 1$), the supplier has no incentive to build any excess capacity, which leads to system instability. When the supplier incurs some backorder cost ($\alpha \in [0, 1)$), there is a unique Nash equilibrium in the absence of participation constraints. The Nash equilibrium is always inefficient: The agents’ selfish behavior degrades

overall system performance. The Nash equilibrium is asymptotically efficient in two cases: (i) The backorder cost goes to zero and the supplier incurs all of the backorder cost, and (ii) the backorder cost goes to infinity and is split evenly between the two agents. In the absence of participation constraints, the Nash equilibrium has an arbitrarily high inefficiency in two cases: (i) The backorder cost goes to infinity and the supplier incurs all of the backorder cost, and (ii) the retailer incurs all of the backorder cost. Relative to the centralized solution, the agents in the Nash equilibrium have more buffer resources when they care sufficiently about backorders: The supplier builds more capacity than optimal when $\alpha < \alpha_\nu$ (and $\alpha_\nu \geq 0.28$ if backorders are more expensive than holding inventory) and the retailer has a larger than optimal base stock level when $\alpha > 0.63$ (and in some cases, an even smaller threshold). However, at least one of the agents in the Nash equilibrium holds a lower-than-optimal level of his buffer resource. Finally, customers receive better service in the centralized solution than in the Nash equilibrium, and customer service improves in the Nash setting when the retailer incurs most – but not all – of the backorder cost.

We assume that the agents only participate if their expected profits are nonnegative. In the Nash equilibrium, the retailer refuses to participate when his share of the backorder cost, α , is sufficiently large (the inventory costs become arbitrarily large when $\alpha = 1$), and the supplier may also (depending on the cost and revenue parameters) refuse to participate when α is too small. Hence, system instability due to a lack of excess capacity would not be expected to arise in practice in the $\alpha = 1$ case because it would be superceded by the participation constraint: The retailer would disengage from the relationship if the supplier offered such horrible service.

A simple linear transfer payment, which is based on actual inventory and backorder levels, the capacity level and the cost parameters, coordinates the system in the absence of participation constraints. We derive bounds on the backorder allocation fraction α for when the coordinating contract is attractive to both parties, in that each agent achieves a nonnegative profit that is no smaller than his Nash equilibrium profit. Interestingly, there are values of α for which the contract will lead to the operation of an otherwise inoperative supply chain; i.e., the extra system profit generated by the contract is sufficient to entice the nonparticipating agent into playing. Overall, we find that a contract is more likely to be entered into by both agents when the system is reasonably profitable (i.e., the optimal profit of the centralized system is large) and relatively well-balanced (i.e., α is near 0.5 and the wholesale price w is intermediate

between the retailer's selling price r and the supplier's manufacturing cost $c + p$). When α takes on an extreme value near 0 or 1, the coordinating contract dictates that one of the agents (the one who is incurring the backorder cost) subsidizes the other's entire operation.

Finally, when one of the agents is the Stackelberg leader, he builds less of his buffer resource and receives a higher profit than in the Nash equilibrium, and the other agent builds more of his buffer resource and receives a smaller profit. Customer service is the same in the Nash equilibrium as when the supplier is the Stackelberg leader, but customers fare worse when the retailer is the leader; this nonobvious result appears to stem from the asymmetric effects of capacity and safety stock on customer service. Taken together, these results provide operations managers with a comprehensive understanding of the competitive interactions in this system, and offer guidelines for when (i.e., for which sets of problem parameters) and how to negotiate contracts to induce participation and increase profits.

Recall that our model is quite similar to the two-stage inventory model of Cachon and Zipkin: The two main differences are the single-server vs. infinite-server model (i.e., queueing vs. inventory model) for the manufacturing process, and our inclusion of revenue, and hence participation constraints. Regarding the first difference, capacity in our model has a larger and more nonlinear impact on the service level than does upstream inventory in Cachon and Zipkin's model. As in single-agent models, the difference between queueing and inventory models in the multi-agent setting is magnified when the queueing system is heavily loaded, which occurs in our model when the capacity costs are large and/or the supplier is not concerned with backorders. In the extreme case when the supplier does not care about backorders ($\alpha = 1$), he builds no excess capacity in our queueing model, whereas he holds no inventory in Cachon and Zipkin's inventory model. The effect of the former is an unstable system, while the effect of the latter is to turn the supply chain into a stable – albeit ineffective – make-to-order system. In Cachon and Zipkin's echelon inventory game, the Nash solution is indeed highly inefficient when $\alpha = 1$, but in the local inventory game the median inefficiency in their computational study is only 1%. When $\alpha = 1$ in the local inventory game, the supplier's base stock level offers him little control over the system's cost, whereas the capacity level in our model affects the entire system in a more profound way. On the other hand, both models predict that the inefficiency is small when the backorder costs are shared equally. Another qualitative difference between the results in these two works is that Cachon and Zipkin's agents hold less inventory in

the Nash equilibrium than in the centralized solution, whereas our agents build/hold a higher-than-optimal level of their buffer resource when their share of the backorder cost is large (as suggested in the management literature by Buzzell and Ortmeyer 1995 and others). Finally, because our model incorporates participation constraints, we can go beyond the Cachon-Zipkin cost minimization analysis and predict that the high inefficiencies in most situations will be preempted by either participation constraints or coordinating contracts; for example, returning to the case where the retailer bears most of the backorder cost, our analysis suggests that the retailer either subsidizes the supplier's production capacity or – if this subsidization does not lead to any retailer profits – does not participate.

Regarding the second part of this thesis, we have formulated and approximately solved the admission/production control problem to a make-to-stock queue when selling prices are highly variable. The importance of such a problem and its solution is two fold.

First, we have the practical implications of this model. The increase in popularity of the Internet is certainly driven consumers, retailers, suppliers, and the whole supply chain to interact in a new way where fixed prices and long term contract agreements are becoming less common. Everybody realizes that the e-price of a product today can be completely different than the e-price of exactly the same product tomorrow. For instance, yield management practices on the airline industry are a good example of this phenomenon on the B2C world. Is in this rapidly changing environment that product managers must be alert to assess the right business opportunities that the Internet is offering. That is, in our setting they have to be able to understand if a particular order at a given price today is attractive or not to be served.

On the other hand, we have the academic interest of formulating and solving the stochastic control problem that comes out of this model. Brownian motions are very effective modelling devices. They are the natural extension to the more traditional Markov processes and they provide more tractable formulations and to some extent they are more insightful about the structure of the problem. In this perspective, the model that we have presented in chapter 4 generalized the standard two-sided regulated Brownian motion (RBM) formulation commonly used to model inventory position by considering in addition a geometric Brownian motion (frequently encountered in finance) to model the price process. The result is a two dimensional stochastic control problem for which approximate solutions are proposed.

In terms of the managerial insights that come out of our model and solution we have that (i) the production policy is characterized by a base-stock policy which is independent of the price and (ii) the admission control policy is a price sensitive function that depends linearly on the price. Thus, since price is rather unpredictable it is convenient to produce until we hit a target level in which case we should turn off production. On the other hand, if there are backorders on the system, let say B , we should reject a new order if the price offered by the incoming customer is below a threshold $R(B) = a_1 + a_2 B$, where the constant a_1 and a_2 are positives. The fact that we get a linear relations for the admission policy is due to our assumption that holding and backordering costs are linearly incur. However, the analysis that we have developed allows us to easily relax this assumption.

From a technical point of view, the analysis that we have performed reveals that the main complexity of this problem comes from the particular reflection field that we have to use on the boundaries. Moreover, our approximations are based on a simple rotation of this reflection field to obtain normal reflection. Under this transformation, we can compute explicitly the solutions to our problem. If we consider that the local time on the boundaries for these two dimensional regulated Brownian motions is 0, we can guess then that the quality of these approximations should be good. In fact, our computational experiments show that in average our proposed policy has an error around 2%. This result looks even better if we consider that a naïve policy such as the static policy produces average errors above 10%. In this sense, the next step to improve the quality of the results proposed here is to better understand the effect of the reflection field on the distribution of the RBM and to incorporate the non normality of these reflections on the structure of the solutions.

Appendix A

Brownian Motion and Queueing Models

This appendix has been written with the intention of exposing some elementary properties and features of Brownian motions that will help the development of chapter (4). Part of the analysis is based on the book by J.M. Harrison: *Brownian Motion and Stochastic Flow Systems*.

A.1 Introduction

The name *Brownian motion* comes from the studies done by the botanist Robert Brown in 1828 on the irregular movement (Brownian movement) of pollen suspended in water.

In order to define what is this object that we will call a Brownian motion, we start by introducing the notion of *stochastic process* and some basic concepts of measure theory.

Definition 2 A stochastic process is a collection of random variables $X = \{X_t : t \geq 0\}$ defined on a sample space (Ω, \mathcal{F}) .

For simplicity, we will assume on this notes that the index t above represents time. We can view a stochastic process as mapping that at each time $t \geq 0$ associated the occurrence of a random phenomenon represented by X_t .

The sample space (Ω, \mathcal{F}) consists on:

- Ω : An abstract space of points $\omega \in \Omega$.

• \mathcal{F} : a σ -field (or σ -algebra) on Ω . That is, a collection of subsets of Ω satisfying:

1. $\Omega \in \mathcal{F}$.
2. Let $A \subseteq \Omega$ such that $A \in \mathcal{F}$ then $A^c = \Omega - A \in \mathcal{F}$.
3. Let $A_1, A_2, A_3, \dots \in \mathcal{F}$ then $A_1 \cup A_2 \cup A_3 \cup \dots \in \mathcal{F}$.

The elements of \mathcal{F} are called *events*. Condition (1) above simply states that the space Ω is necessarily an event. Conditions (2) and (3) state that the collection of events is closed under the set operations of complement and countable union.

The space (Ω, \mathcal{F}) satisfying these conditions is called a *measurable space*. For a fixed point $\omega \in \Omega$, we called the *sample path* of the process X associated with ω the function $t \rightarrow X_t(\omega) : t \geq 0$.

An important consideration is related to the way we collect and use information over time. In particular, we would like to be able to isolate past and present from future. For example, let consider two events $A, B \in \mathcal{F}$ such that $X_s(\omega_1) = X_s(\omega_2); \forall \omega_1 \in A, \forall \omega_2 \in B, \forall s, 0 \leq s \leq t$. Then during the period $[0, t]$ the events A and B can not distinguished from the point of view of X . For this reason, we complement our sample space (Ω, \mathcal{F}) with a *filtration*, i.e., a nondecreasing family $\{\mathcal{F}_t : t \geq 0\}$ of sub- σ -fields of \mathcal{F} such that $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ for $0 \leq s \leq t$. The idea is that each \mathcal{F}_t contains the information available up to time t .

Another very important object that we will use during the analysis of Brownian motion is called *stopping time*. However before defining it, we have to introduce the notion of *measure* and *measurable function*.

Definition 3 Let (Ω, \mathcal{F}) be a measurable space. A set function μ on \mathcal{F} is called a measure if it satisfies the following conditions:

1. $\mu(\emptyset) = 0$,
2. $A \in \mathcal{F}$ implies $0 \leq \mu(A) \leq \infty$,
3. $A_1, A_2, \dots \in \mathcal{F}$ and $\{A_n\}$ are pairwise disjoint ($A_i \cap A_j = \emptyset$ for $i \neq j$), then

$$\mu \left(\bigcup_{n=1}^{\infty} A_n \right) = \sum_{n=1}^{\infty} \mu(A_n).$$

For example, the *Lebesgue measure* commonly denoted by λ and defined on the class of Borel sets \mathcal{R}^1 of the real line, is given by $\lambda(a, b] = b - a$.

We call a measure μ a *probability measure* if $\mu(\Omega) = 1$, in this case instead of μ , we use the notation P . If P is a probability measure then (Ω, \mathcal{F}, P) is called a *probability space*.

Definition 4 Let (Ω, \mathcal{F}) and (Ω', \mathcal{F}') two measurable spaces. A mapping $Y : \Omega \rightarrow \Omega'$ is $(\mathcal{F}, \mathcal{F}')$ -measurable if for each $A' \in \mathcal{F}'$

$$Y^{-1}(A') = \{\omega \in \Omega : Y(\omega) \in A'\} \in \mathcal{F}.$$

According to this definition, the mapping Y is measurable if for any well-defined event $A' \in \mathcal{F}'$ the pre-image of A' by Y (denoted by $Y^{-1}(A')$) is also a well-defined event in \mathcal{F} .

Definition 5 A stopping time is a measurable function T from (Ω, \mathcal{F}) to $[0, \infty)$ such that $\{\omega \in \Omega : T(\omega) \leq t\} \in \mathcal{F}_t$, for all $t \geq 0$.

We can now give a first definition of a Brownian motion.

Definition 6 A standard, one-dimensional Brownian motion (or Wiener process) is a continuous, adapted stochastic process $X = \{X_t, \mathcal{F}_t : 0 \leq t < \infty\}$, defined on some probability space (Ω, \mathcal{F}, P) with the properties that $X_0 = 0$ almost surely and for $0 \leq s < t$, the increment $X_t - X_s$ is independent of \mathcal{F}_s and is normally distributed with mean zero and variance $t - s$.

A few points about the previous definition:

- (a) We say that $X_0 = 0$ almost surely (a.s.) in (Ω, \mathcal{F}, P) if the set $A = \{\omega \in \Omega : X_0(\omega) \neq 0\}$ has probability 0, i.e., $P(A) = 0$.
- (b) The stochastic process X is adapted to the filtration $\{\mathcal{F}_t\}$ if for each $t \geq 0$, X_t is an \mathcal{F}_t -measurable random variable.

Two extremely important features that characterizes a Wiener process presented in the following property:

Property 1 *If X is a Wiener process then X has independent increments, that is, for any positive integer n and any sequences of times $0 \leq t_0 < t_1 < \dots < t_n < \infty$ the random variables $Y_i = X(t_i) - X(t_{i-1})$, $i = 1, 2, \dots, n$ are independent.*

In addition, X has stationary increments, that is, for any $0 \leq s < t < \infty$ the distribution of $X_t - X_s$ depends only on $t - s$.

These two properties are so attached to the Wiener process that they can be used as an alternative definition of a standard Brownian motion. We notice here that for discrete time stochastic processes the two properties above characterize the Poisson process. Thus, we might think for simplicity on the Wiener process as the continuous space extension of the Poisson process.

Once we have defined the Wiener process, we can extend its definition and define the general (μ, σ) Brownian motion process Y as follows:

$$Y(t) = Y(0) + \mu t + \sigma X(t)$$

where X is a Wiener process and $Y(0)$ (the initial value) is independent of X . We call μ the *drift* and σ^2 the *variance* or *diffusion* of Y . It follows directly from the definition of X that $Y(t+s) - Y(t)$ is normally distributed with mean μs and variance $\sigma^2 s$. Finally, we say that Z is a *geometric Brownian motion* if $Z_t = e^{Y_t}$, where Y is a Brownian motion.

A.2 Properties of Brownian Motions

In this section we present the main properties that make Brownian motions a very attractive modelling tool. However, we start ironically presenting some results showing the extremely erratic behavior of Brownian motion processes.

A.2.1 Basic Properties

Property 2 *Let X be a Brownian motion in (Ω, \mathcal{F}, P) . Then except for a set of probability 0, the sample path $X_t(\omega)$ is nowhere differentiable.*

Even though the variation of X over time is particular unstable, some measure of its variability can be computed. In fact, let define the random variable (quadratic variation) Q_t as follows:

$$Q_t \equiv \lim_{n \rightarrow \infty} \sum_{k=0}^{2^n-1} \left[X \left(\frac{(k+1)t}{2^n} \right) - X \left(\frac{kt}{2^n} \right) \right]^2. \quad (\text{A.1})$$

Then, we have the following result.

Property 3 For almost every $\omega \in \Omega$ we have $Q_t(\omega) = \sigma^2 t$ for all $t \geq 0$.

This last result implies that Brownian motion have infinite ordinary variation almost surely. In addition, as we will see later, property 3 contains the essence of the Itô's formula.

Property 4 If X is a (μ, σ) Brownian motion then:

- $E(X_t) = X_0 + \mu t$,
- $Var(X_t) = \sigma^2 t$,
- $Cov(X_t, X_s) = \sigma^2(t \wedge s) = \sigma^2 \min\{t, s\}$.

The following theorem is a very important result that reflect the memoryless property that characterizes Brownian motion processes.

Theorem 1 (Strong Markov Property)

Let X be a (μ, σ) Brownian motion and T be a finite stopping time. Then $Y_t = X_{T+t} - X_T$ is a (μ, σ) Brownian motion starting at 0 and it is independent of \mathcal{F}_T .

Property 5 (Brownian martingales)

Let X be a (μ, σ) Brownian motion then:

- (a) If $\mu = 0$ then X is a martingale, i.e., $E(X_t - X_s | \mathcal{F}_s) = 0$.
- (b) If $\mu = 0$ then $X_t^2 - \sigma^2 t$ is martingale.
- (c) Let $q(\beta) = \mu\beta + \frac{1}{2}\sigma^2\beta^2$ and $V_\beta(t) = e^{\beta X_t - q(\beta)t}$. Then, V_β is a martingale.

A.2.2 Wiener Measure and Donsker's Theorem

In this subsection we explore the nature of the Wiener process as a type of central limit theorem for stochastic processes. The notation and results are based on the textbook *Convergence of Probability Measures* by P. Billingsley (1999).

We start by introducing the Wiener measure, W , which is a probability measure on (C, \mathcal{C}) ¹ having two properties. First, each X_t is normally distributed under W with mean 0 and variance t , that is:

$$W[X_t \leq \alpha] = \frac{1}{\sqrt{2\pi t}} \int_{-\infty}^{\alpha} e^{-\frac{u^2}{2t}} du.$$

For $t = 0$, we have $W[X_0 = 0] = 1$. The second property is that the stochastic process X has independent increments under W .

In order to state the main result of this section (Donsker's theorem), we introduce the sequences $\{X^n : n = 0, 1, \dots\}$ of stochastic processes as follows. Let $\Xi = \{\xi_1, \xi_2, \dots\}$ be a sequence of IID random variables having mean 0 and finite variance σ^2 . Let $S_n = \xi_1 + \dots + \xi_n$ ($S_0 = 0$) be the partial sums of Ξ . We define X^n as follows:

$$X_t^n(\omega) = \frac{1}{\sigma\sqrt{n}} S_{[nt]}(\omega) + (nt - [nt]) \frac{1}{\sigma\sqrt{n}} \xi_{[nt]+1}(\omega) \quad (\text{A.2})$$

Figure A.1 shows the behavior of the process X_t^n for three values of n ². We can see that as n increases, the behavior of X^n resembles a Wiener process. This result is in fact Donsker's theorem. We can see the non differentiability of $X^n(t)$ as n increases.

Theorem 2 (Donsker's Theorem) *If ξ_1, ξ_2, \dots are independent and identically distributed random variables with mean 0 and variance σ^2 , and if X^n is the random process defined by (A.2), then $X^n \Rightarrow_n W$, a Wiener process.*

(Where the symbol \Rightarrow_n stands for *convergence in distribution* as $n \rightarrow \infty$.) This result can be understood as a generalization of the standard *central limit theorem* for random variables.

¹Where $C \equiv C[0, \infty)$ is the space of all continuous functions $x : [0, \infty) \rightarrow \mathcal{R}$ and \mathcal{C} is the Borel σ -algebra on C .

²In the construction of the three sample paths the ξ_n are uniformly distributed in $[-1, 1]$.

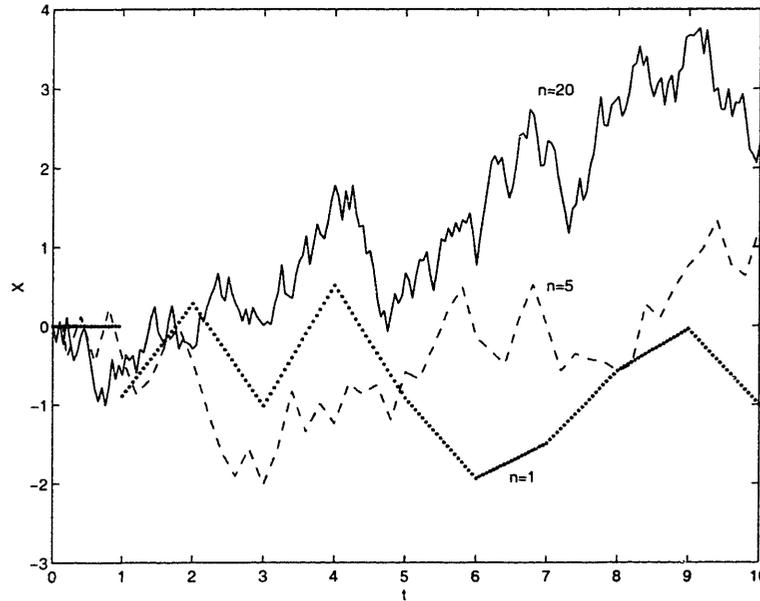


Figure A.1: Behavior of $\{X_t^n\}$ for $n = 1, 5,$ and 20 .

The previous result is intuitive in the sense that S_n -being the sum of IID random variables- converges in distribution to a $N(0, n\sigma^2)$. Another interesting property of the Wiener process, and more generally of any $(0, \sigma)$ Brownian motion is their *scale invariance* that can partially be observed in (A.2).

Property 6 (Scale Invariance) Let X be a $(0, \sigma)$ Brownian motion, then for any $c > 0$:

$$\left\{ \frac{X(ct)}{\sqrt{c}} : t \geq 0 \right\} \stackrel{D}{=} \{X(t) : t \geq 0\}. \quad (\text{A.3})$$

(Where $\stackrel{D}{=}$ stands for equality in distribution.)

This scaling property, that is of course related to the normal distribution, will be specially important (together with Donsker's theorem) later for our study of queueing systems and the use of heavy traffic approximations.

We can now use Donsker's theorem to find the distribution of $M \equiv \sup W$, however, we need before an additional result.

Theorem 3 (Mapping Theorem) Let $\{X_n\}$ be a sequence of processes such that $X_n \Rightarrow X$. Let h be a measurable function and let D_h be the set of its discontinuities.

If D_h has probability 0, then $h(X_n) \Rightarrow h(X)$.

Since $h(X) := \sup_t X_t$ is a continuous function on C , then from the mapping theorem and the fact that $X^n \Rightarrow W$, we have that:

$$\sup_t X_t^n \Rightarrow \sup_t W_t.$$

Let $M_n = \max_{0 \leq i \leq n} S_i$, then it is not hard to show that $\sup_t X_t^n = \frac{M_n}{\sigma\sqrt{n}}$. Thus,

$$\frac{M_n}{\sigma\sqrt{n}} \Rightarrow \sup_t W_t. \quad (\text{A.4})$$

Since we can peak any sequence $\{\xi_n\}$ such that $E(\xi_n) = 0$ and $E(\xi_n^2) < \infty$, let assume that ξ_n takes the values ± 1 with probability $\frac{1}{2}$. Therefore, S_0, S_1, \dots represents a symmetric random walk starting at 0. We first prove that

$$P(M_n \geq a, S_n < a) = P(M_n \geq a, S_n > a) \quad a \geq 0.$$

This should be clear from the fact that the behavior of the random walk is independent of its history and it is symmetric, thus if the random walk reach a at time $\hat{n} < n$ then the value of S_n is symmetric with respect to $S_{\hat{n}} = a$. In other words, for each path of the random walk (S_0, S_1, \dots, S_n) such that $M_n \geq a, S_n = a - k < a$ there exists another path such that $M_n \geq a, S_n = a + k > a$. This symmetry is an example of the *reflection principle*. Given this result, we have that:

$$\begin{aligned} P(M_n \geq a) &= P(M_n \geq a, S_n < a) + P(M_n \geq a, S_n = a) + P(M_n \geq a, S_n > a) \\ &= 2P(M_n \geq a, S_n > a) + P(M_n \geq a, S_n = a) \\ &= 2P(S_n > a) + P(S_n = a) \end{aligned}$$

By the central limit theorem $P(S_n > a\sqrt{n}) \rightarrow P(N > a)$ and $P(S_n = a\sqrt{n}) \rightarrow 0$, where N is a standard $(0, 1)$ normally distributed random variable. In addition $2P(N > a) = P(|N| > a)$. Thus, combining this results we have that $M = \sup_t W_t$ has the same distribution of $|N|$ and

$$P(M \leq a) = \frac{2}{\sqrt{2\pi}} \int_0^a e^{-\frac{u^2}{2}} du. \quad (\text{A.5})$$

A.2.3 Reflection Principle

In this subsection, we look with more detail at the distribution of $M_t = \sup_{0 \leq s \leq t} X_s$, where X is a general (μ, σ) Brownian motion. We first start the analysis for the special case of $\mu = 0, \sigma = 1$. In this case, we can apply a similar argument that the one used in the previous subsection based on the *reflection principle* to show that

$$P(M_t \geq x) = 2P(X_t \geq x) = P(|X_t| \geq x).$$

We can also compute the joint distribution for (X_t, M_t) , that is

$$F_t(x, y) = P(X_t \leq x, M_t \leq y).$$

Since $X_0 = 0$ and $M_t \geq X_t$ w.p.1, we can focus our attention to the case $x \leq y$ and $y \geq 0$. First of all, we notice that

$$\begin{aligned} F_t(x, y) &= P(X_t \leq x) - P(X_t \leq x, M_t > y) \\ &= \Phi(xt^{-\frac{1}{2}}) - P(X_t \leq x, M_t > y), \end{aligned}$$

where $\Phi(\cdot)$ is the $N(0, 1)$ distribution function. From the reflection principle $P(X_t \leq x, M_t > y) = P(X_t \geq 2y - x) = P(X_t \leq x - 2y)$. Thus, we have the following result.

Property 7 *If $\mu = 0$ and $\sigma = 1$, then*

$$P(X_t \leq x, M_t \leq y) = \Phi(xt^{-\frac{1}{2}}) - \Phi((x - 2y)t^{-\frac{1}{2}}).$$

The previous result depends heavily on the assumption $\mu = 0$ or in other words on the reflection principle. In order to extend the result to general Brownian motion, it is required first to understand how making a change of measure can lead to a change of drift.

Let P and Q be two probability measures on the same probability space (Ω, \mathcal{F}) with the important property that P is *dominated* by Q . That is, $Q(A) = 0 \implies P(A) = 0$. Then, there exists a non-negative random variable ξ (also denoted by $\frac{dP}{dQ}$) such that

$$P(A) = \int_A \xi dQ, \quad \forall A \in \mathcal{F}.$$

An important implication of the above relation is that if Y is a random variable and $E_Q(|\xi Y|) < \infty$ then $E_P(Y)$ exists and $E_P(Y) = E_Q(\xi Y)$. The random variable ξ

is usually called the *density* or Radon-Nikodym derivative (or likelihood ratio) of P with respect to Q . In order to find the density ξ associated to two Brownian motion measures of different drift, we use an heuristical approach.

Let consider now a (μ, σ) Brownian motion and a sequence of instants $0 = t_0 < t_1 < \dots < t_n = t$ such that $t_i - t_{i-1} = \delta$, $i = 1, \dots, n$. The density associated to that particular sequence of instance is given by:

$$\frac{1}{(\sigma\sqrt{2\pi\delta})^n} \prod_{i=1}^n e^{-\frac{(X_{t_i} - X_{t_{i-1}} - \mu\delta)^2}{2\sigma^2\delta}}.$$

If the drift were instead $\mu + \theta$ then density is obtained replacing μ by $\mu + \theta$ above. Thus, the density is given by:

$$e^{\frac{1}{2\sigma^2\delta} \sum_{i=1}^n (X_{t_i} - X_{t_{i-1}} - \mu\delta)^2 - (X_{t_i} - X_{t_{i-1}} - (\mu + \theta)\delta)^2}.$$

After some algebra, we have that the density is given by:

$$\begin{aligned} \xi(t) &= e^{\frac{\theta}{\sigma^2}(X_t - \mu t - \frac{\theta t}{2})} \\ &= V_{\frac{\theta}{\sigma^2}}(t). \end{aligned} \tag{A.6}$$

Where $V_{\beta}(t)$ is Wald martingale defined in property (5). We can compute the distribution of M_t for the case of $\mu \neq 0$ as follows (we start with $\sigma = 1$):

$$\begin{aligned} P_{\mu}(M_t \geq x) &= E_0(V_{\mu}(t); M_t \geq x) \\ &= 1 - \Phi\left(\frac{x - \mu t}{\sqrt{t}}\right) + e^{2\mu x} \Phi\left(\frac{-x - \mu t}{\sqrt{t}}\right). \end{aligned}$$

Finally, for the general case (μ, σ) , we can rescale the probability measure to obtain:

$$P(M_t \leq x) = \Phi\left(\frac{x - \mu t}{\sigma\sqrt{t}}\right) - e^{\frac{2\mu x}{\sigma^2}} \Phi\left(\frac{-x - \mu t}{\sigma\sqrt{t}}\right), \tag{A.7}$$

which is called the *inverse Gaussian distribution*.

A.2.4 Some Extensions

In this subsection, we present some additional results concerning Brownian motions. The first important extensions is related to the initial condition X_0 . Before, we have

impose the restriction that $X_0 = 0$ w.p.1. We now turn to the general case $X_0 = x$ w.p.1, where x is any real number. In order to make explicit this new value of the initial state, we introduce the notation P_x to refer to the probability measure that satisfies $P_x(X_0 = x) = 1$ (the same is valid for E_x , the expected value operator under P_x).

A first important result is related to the way we represent Brownian motions (BM). Of course, we have already given a concrete definition of a BM, however, let look at an alternative representation. We know that $X_{t+s} - X_t$ has a $N(\mu s, \sigma^2 s)$ distribution. Thus, the transition density

$$p(t, x, y) dy \equiv P_x(X_t \in dy) = \frac{1}{\sigma\sqrt{t}} \phi\left(\frac{y - x - \mu t}{\sigma\sqrt{t}}\right) dy$$

satisfies the following differential equation:

$$\frac{\partial}{\partial t} p(t, x, y) = \left(\frac{1}{2} \sigma^2 \frac{\partial^2}{\partial x^2} + \mu \frac{\partial}{\partial x} \right) p(t, x, y),$$

with initial condition

$$p(0, x, y) = \delta(x - y) = \begin{cases} 1 & \text{if } y = x, \\ 0 & \text{otherwise.} \end{cases}$$

The differential equation above characterizes BM and is called *Kolmogorov's backward equation*. If instead of differentiating with respect to the initial state x , we differentiate with respect y , the final state, we get the *Kolmogorov's forward equation*

$$\frac{\partial}{\partial t} p(t, x, y) = \left(\frac{1}{2} \sigma^2 \frac{\partial^2}{\partial y^2} - \mu \frac{\partial}{\partial y} \right) p(t, x, y).$$

In the special case when $\mu = 0$, the previous equation reduces to the traditional *heat equation* (or diffusion equation), for this reason Brownian motion are usually called *diffusion processes*.

Another important extension is the *Hitting Time* problem. That is, the problem of determining the first time when the process reaches a predefined state. Let define $T(y) = \inf\{t \geq 0 : X_t = y\}$, i.e., the first time at which X reaches the value y . Suppose that the process start at $x \geq 0$ and let $0 < x \leq b$. Then, we are interested in finding the distribution of $T \equiv T(0) \wedge T(b)$. A first step is the following result:

Property 8

$$E_x(T) < \infty, \quad 0 \leq x \leq b.$$

The proof is based on the *martingale stopping* theorem, that is

Theorem 4 Martingale Stopping Theorem

Let T be a stopping time and X a martingale (with right-continuous sample paths) on certain filtered probability space. Then the stopped process $\{X(t \wedge T), t \geq 0\}$ is also a martingale.

Thus, if we apply this result to $M_t = X_t - \mu t$, which is clearly a martingale, we have that:

$$E_x(M(T \wedge t)) = E_x(M(0)) = x.$$

But $E_x(M(T \wedge t)) = E_x(X(T \wedge t)) - \mu E_x(T \wedge t)$. Thus, for $\mu \neq 0$, the result in (8) follows directly. For the case, $\mu = 0$, we have to apply the martingale stopping theorem to the martingale $X_t^2 - \sigma^2 t$.

Let us now recall *Wald martingale* introduced in property (5), that is, $V_\beta(t) = e^{\beta X_t - q(\beta)t}$ where the function $q(\cdot)$ is given by $q(\beta) = \mu\beta + \frac{\sigma^2\beta^2}{2}$. Now, it can be shown that

$$E_x(V_\beta(T)) = E_x(V_\beta(0)) = e^{\beta x}, \quad 0 \leq x \leq b.$$

Therefore, we have the following decomposition:

$$\begin{aligned} e^{\beta x} &= E_x(V_\beta(T); X_T = 0) + E_x(V_\beta(T); X_T = b), \\ &= \psi_*(x|q(\beta)) + e^{\beta b}\psi^*(x|q(\beta)), \end{aligned} \tag{A.8}$$

where $\psi_*(x|\lambda) \equiv E_x(e^{-\lambda T}; X_T = 0)$ and $\psi^*(x|\lambda) \equiv E_x(e^{-\lambda T}; X_T = b)$. Solving the equation $q(\beta) = \lambda$, we get:

$$\beta_*(\lambda) = \frac{\mu + \sqrt{\mu^2 + 2\sigma^2\lambda}}{\sigma^2} > 0; \quad \beta^*(\lambda) = \frac{\mu - \sqrt{\mu^2 + 2\sigma^2\lambda}}{\sigma^2} < 0.$$

Thus, combining this result and (A.8), we get the following system of equation:

$$\begin{aligned} e^{-\beta_*(\lambda)x} &= \psi_*(x|\lambda) + e^{-\beta_*(\lambda)b}\psi^*(x|\lambda), \\ e^{-\beta^*(\lambda)x} &= \psi_*(x|\lambda) + e^{-\beta^*(\lambda)b}\psi^*(x|\lambda). \end{aligned}$$

The solution of this system gives the following result.

Property 9 Let $\lambda > 0$ be fixed. For $0 \leq x \leq b$,

$$\begin{aligned}\psi^*(x|\lambda) &= \frac{\theta^*(x, \lambda) - \theta_*(x, \lambda)\theta^*(0, \lambda)}{1 - \theta_*(b, \lambda)\theta^*(0, \lambda)}, \\ \psi_*(x|\lambda) &= \frac{\theta_*(x, \lambda) - \theta^*(x, \lambda)\theta_*(b, \lambda)}{1 - \theta_*(b, \lambda)\theta^*(0, \lambda)}, \\ \theta_*(x, \lambda) &= e^{-\beta_*(\lambda)x}, \\ \theta^*(x, \lambda) &= e^{\beta^*(\lambda)(b-x)}.\end{aligned}$$

Finally, from the previous result, we can obtain the distribution (or more precisely the Laplace transform) of T .

Proposition 20 Let θ^* and θ_* be defined as above. Then,

$$E_x(e^{-\lambda T(0)}; T(0) < \infty) = \theta_*(x, \lambda) ; E_x(e^{-\lambda T(b)}; T(b) < \infty) = \theta^*(x, \lambda), \quad 0 \leq x \leq b.$$

In addition, if $\mu = 0$ then $P_x(X_T = b) = \frac{x}{b}$. Otherwise,

$$P_x(X_T = b) = \frac{1 - \xi(x)}{1 - \xi(b)} ; \quad \xi(z) \equiv e^{\frac{-2\mu z}{\sigma^2}}, \quad 0 \leq x \leq b.$$

So far, we have presented the basic properties associated to Brownian motions. We would like next to present the applications of BM to queueing theory. However, in order to extend these results to queueing systems, we require some additional background, in particular, the notion of *Regulated Brownian Motion* (RBM) is fundamental. It turns out that the analysis of RBM is much easier if we use results from Stochastic Calculus. For this reason, we postpone the study of RBM and we present in next section the basic elements of Stochastic Calculus.

A.3 Stochastic Calculus

The main goals of this section is to present Itô's lemma and the use of stochastic differential equations as important tools for modelling stochastic processes. We start the analysis introducing heuristically *Itô's Stochastic differential equation*.

A.3.1 Motivation

It is a common practice when modelling physical systems to express the dynamics of the system, i.e., its evolution over time through a difference or differential equation. For example, when describing the position ($y(t)$) of a certain object at time t , we might use the relation

$$\frac{dy(t)}{dt} = v(t),$$

where $v(t)$ is the instantaneous velocity of the object at time t . In general, differential equations have been used extensively in science, and probably one of their biggest advantages is that they are able to capture the essence of the physical system without incorporating the natural and necessary difficulties that are imposed by border conditions.

Let us now look at the “general” (deterministic) differential equation:

$$\frac{dx(t)}{dt} = f(t, x(t)).$$

Solving this equation is an old problem in mathematics, and it is not the purpose of this note to go into the details of how to solve it. We would like, however, to introduce some type of uncertainty into the model. One easy way of doing this is to use the traditional trick used by econometricians, that is, to simply add an stochastic term to the above relation. In order to do that, we proceed as follows. We first approximate the dynamics by:

$$x(t + \Delta t) - x(t) = f(t, x(t))\Delta t + o(\Delta t),$$

where $o(t)$ is function such that $t^{-1}o(t) \rightarrow 0$ as $t \rightarrow 0$. If we assume now that uncertainty can be model by an stochastic process $v(t)$ that we simply add in to the dynamics of the system, we have:

$$x(t + \Delta t) - x(t) = f(t, x(t))\Delta t + v(t + \Delta) - v(t) + o(\Delta t).$$

In particular, we might think on this uncertainty as being the sum of independent and small perturbations. Thus, a reasonable model is to suppose that

$$v(t + \Delta) - v(t) = \sigma(t, x(t))(z(t + \Delta t) - z(t)),$$

where $z(t)$ is Wiener process and σ accounts for the variance of v . We can then

rewrite the dynamics of the system as follows:

$$dx = f(t, x)dt + \sigma(t, x)dz,$$

which is called *Itô's stochastic differential equation*. Notice that we can not divide by dt above since z is nowhere differentiable.

A.3.2 Stochastic Integration

Since the Wiener process is nowhere differentiable, Itô's differential equation does not have a clear meaning per-se. In this subsection, we will give it one, which is based on the notion of stochastic integral. The idea is the following, we use the notion:

$$dx = f(t, x)dt + \sigma(t, x)dz,$$

as a shorthand for

$$x(t) = x(0) + \int_0^t f(s, x(s))ds + \int_0^t \sigma(s, x(s))dz(s).$$

The first integral in the right-hand side is understood in the usual Riemann sense. The second integral, however, does not have a clear meaning for the reasons that we have already mentioned. Let us then focus in the following stochastic process:

$$I_t(X) = \int_0^t X_s dW_s, \quad t \geq 0, \quad (\text{A.9})$$

that we called the *stochastic integral*. Here X is any stochastic process and W is a Wiener process. Stochastic integration was first presented by Itô (1944) and extended later by Doob (1953). Here, we will not go into the formal details behind the theory of stochastic integration. We will rather give a more simpler and intuitive analysis.

From traditional calculus, we know that if a function is relatively well-behaved in the interval $[0, t]$ (i.e., it is integrable), then we can approximate the value of

$$I = \int_0^t f(s) ds,$$

as follows. We first introduce a sequence of partitions $\{P_n : n \geq 1\}$ where $P_n = \{t_i : 0 \leq i \leq n\}$ is a partition of the interval $[0, t]$, i.e., $0 = t_0 < t_1 < \dots < t_n = t$. We denote by $\|P_n\| = \max\{t_i - t_{i-1} : 1 \leq i \leq n\}$. Then, if $\lim_{n \rightarrow \infty} \|P_n\| = 0$, we have

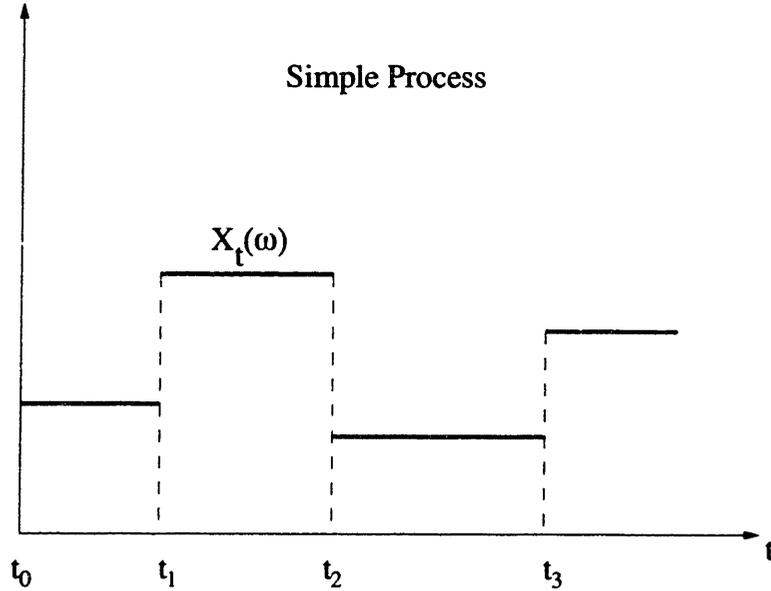


Figure A.2: Simple process.

that

$$I = \lim_{n \rightarrow \infty} \sum_{t_i \in P_n} f(\xi_i)(t_{i+1} - t_i),$$

where $\xi_i \in [t_i, t_{i+1}]$. In particular, $\xi_i = t_i$ or $\xi_i = t_{i+1}$ does not make much difference for deterministic real-valued function. This analysis is exactly the one that we will apply to compute stochastic integral, however, the analysis requires some extra attention.

We first introduce a special class of stochastic processes that we call *simple*. A process X is *simple* if there exist a sequence of times $\{t_k\}$ such that

$$0 = t_0 < t_1 < \cdots < t_k \rightarrow \infty$$

and

$$X(t, \omega) = X(t_{k-1}, \omega), \quad \forall t \in [t_{k-1}, t_k) \quad k = 1, 2, \dots$$

We notice that the sequence $\{t_k\}$ is independent of ω . Given the special form of simple processes, it is possible to give a clear definition of the stochastic integral in this case. In fact, if X is simple, then

$$I(X) = \int_0^t X dW = \sum_{k=0}^{n-1} X(t_k)[W(t_{k+1}) - W(t_k)],$$

where $t_0 = 0$ and $t_n = t$. The importance of simple processes is not only that we are able to compute easily their integrals but also that they can be use to approximate other more complex processes. That is,

Property 10 *Let consider the stochastic process $X \in H^2$, i.e.,*

$$E \left[\int_0^t X^2(s) ds \right] < \infty, \quad \forall t \geq 0.$$

Then, there exists a sequence of simple processes $\{X_n\} \in H^2$ such that

$$X_n \Rightarrow_n X.$$

Moreover, the value of $I_t(X)$ can be obtained from the fact that

$$I_t(X_n) \Rightarrow_n I_t(X).$$

Let define the norm $\|\cdot\|$ in H^2 as follows:

$$\|X\| = E \left[\int_0^t X^2(s) ds \right]^{\frac{1}{2}}.$$

Then we have the following result

Proposition 21 *Let $X \in H^2$, then $E[I_t(X)] = 0$ and $\|I_t(X)\| = \|X\|$.*

Example:

Let consider the case when $X = W$, it is not hard to show that $W \in H^2$, moreover, $\|W\| = \frac{t}{\sqrt{2}}$. Now, in order to compute $I(W)$ we introduce the simple processes

$$X_n(s) = W \left(\frac{kt}{2^n} \right), \quad s \in \left[\frac{kt}{2^n}, \frac{(k+1)t}{2^n} \right).$$

Let define $t_k = \frac{kt}{2^n}$. Then, for these simple processes we have:

$$\begin{aligned} I_t(X_n) &= \sum_{k=0}^{2^n-1} W(t_k)[W(t_{k+1}) - W(t_k)] \\ &= \frac{1}{2} \sum_{k=0}^{2^n-1} [W^2(t_{k+1}) - W^2(t_k)] - \frac{1}{2} \sum_{k=0}^{2^n-1} [W(t_{k+1}) - W(t_k)]^2 \\ &= \frac{1}{2} W^2(t) - \frac{1}{2} \sum_{k=0}^{2^n-1} [W(t_{k+1}) - W(t_k)]^2. \end{aligned}$$

But in equation (A.1), we saw that the summation above converges to t . Therefore, we conclude that:

$$I_t(W) \equiv \int_0^t W dW = \frac{1}{2}W^2(t) - \frac{t}{2}. \quad (\text{A.10})$$

A.3.3 Itô's Lemma

In this section, we define the notion of *stochastic differential* and state and prove Itô's lemma which is the fundamental rule for computing stochastic differentials.

We start by defining some notation. As usual, we consider a probability space (Ω, \mathcal{F}, P) , a Wiener process $W(t, \omega)$, a process $Y(t, \omega)$ that is jointly measurable in t and ω with respect to \mathcal{F}_t , is adapted and satisfies $\int_0^T |Y(t, \omega)| dt < \infty$ w.p.1. We also consider a process X that is non-anticipating on $[0, T]$. We say that Z is an *Itô process* if it has the following functional form:

$$Z(t, \omega) = Z(0, \omega) + \int_0^t X(s, \omega) dW(s, \omega) + \int_0^t Y(s, \omega) ds. \quad (\text{A.11})$$

The first integral in the right-hand side is called the *Brownian component* of Z and it has to be computed according to the analysis that we did in the previous section. The second integral is called the *drift component* (or *VF component*) of Z and it is evaluated in the usual Reimann sense. Instead of using (A.11) to represent Z , we say that Z has an *Itô differential* (or *stochastic differential*) dZ given by:

$$dZ = X dW + Y dt.$$

Proposition 22 (Itô's Lemma)

Let $u(t, x)$ be a continuous non-random function with continuous partial derivatives u_t , u_x , and u_{xx} . Suppose that Z is a process with stochastic differential $dZ = X dW + Y dt$. Let define the process $V(t) = u(t, Z(t))$, then V has a stochastic differential given by:

$$dV = \left[\frac{\partial}{\partial t} u(t, Z) + \frac{\partial}{\partial Z} u(t, Z) Y + \frac{1}{2} \frac{\partial^2}{\partial Z^2} u(t, Z) X^2 \right] dt + \frac{\partial}{\partial Z} u(t, Z) X dW. \quad (\text{A.12})$$

(The proof of the Lemma uses a second order Taylor expansion of the function $u(t, x)$.)

Let us take a look at Itô's lemma in a particular case. Let suppose that $u(t, Z) =$

$f(Z)$, for some twice continuously differentiable function f . Then (A.12) implies:

$$df(Z) = \left[f'(Z)Y + \frac{1}{2}f''(Z)X^2 \right] dt + f'(Z)XdW.$$

Rearranging terms we get:

$$\begin{aligned} df(Z) &= f'(Z)[Ydt + XdW] + \frac{1}{2}f''(Z)X^2dt \\ &= f'(Z)dZ + \frac{1}{2}f''(Z)(dZ)^2. \end{aligned} \tag{A.13}$$

Relation (A.13) is a simplify way of expressing Itô's lemma and uses the convention $(dZ)^2 = (Ydt + XdW)^2 = X^2dt$. The idea is that in differential terms only $(dW)^2 \neq 0$, this is consistent with our finding in (A.1) about the quadratic variation of Wiener processes. Let notice that for ordinary differentials $df(Z) = f'(Z)dZ$, thus the second term in (A.13) is the main difference for stochastic differential that, as we have already mentioned, reflects the infinite variation of Brownian paths.

Example:

Suppose that $V(t) = e^{X(t)}$ where X satisfies

$$dX = -\alpha dt + \beta dW,$$

and $X(0) = 0$. Applying Itô's lemma we have:

$$\begin{aligned} dV &= \left[-\alpha V + \frac{1}{2}V \right] dt + \beta V dW, \\ &= V \left(\left[\frac{1}{2} - \alpha \right] dt + \beta dW \right). \end{aligned}$$

That is, V is a *geometric Brownian motion*. In order to solve the stochastic differential above, we introduce the following change of variable:

$$A = \ln(V).$$

Then, using Itô's lemma we get:

$$dA = \underbrace{\frac{1}{V}dV}_{[1]} - \frac{1}{2V^2} \underbrace{(dV)^2}_{[2]}.$$

Thus, replacing [1] by $[\frac{1}{2} - \alpha]dt + \beta dW$ and [2] by $\beta^2 V^2 dt$ we get:

$$\begin{aligned} dA &= \left(\frac{1}{2} - \alpha\right)dt + \beta dW - \frac{\beta^2}{2}dt, \\ &= \left(\frac{1}{2} - \alpha - \frac{\beta^2}{2}\right)dt + \beta dW. \end{aligned}$$

This linear differential implies $A(t) = (\frac{1}{2} - \alpha - \frac{\beta^2}{2})t + \beta W(t)$. Finally, combining this result and the transformation $A = \ln(V)$ we conclude:

$$V(t) = e^{(\frac{1}{2} - \alpha - \frac{\beta^2}{2})t + \beta W(t)}.$$

We finish this section with an important result about the existence of solutions for stochastic differential.

Theorem 5 *Let consider the Itô process Z defined through the following stochastic differential:*

$$dZ(t) = f(t, Z(t))dt + g(t, Z(t))dW(t),$$

with initial condition $Z(0, \omega) = c(\omega) = c$. If

1. f and g are both measurable with respect to all their arguments,
2. There exists a constant $K > 0$ such that

$$|f(t, x) - f(t, y)| + |g(t, x) - g(t, y)| \leq K|x - y|,$$

$$|f(t, x)|^2 + |g(t, x)|^2 \leq K^2(1 + |x|^2),$$

3. The initial condition $Z(0, \omega)$ does not depend on $W(t)$ and $E[Z(0, \omega)^2] < \infty$.

Then, there exists a solution $Z(t)$ satisfying the initial condition which is unique w.p.1, has continuous paths and $\sup_t E[Z(t)^2] < \infty$.

A.4 Regulated Brownian Motion

So far we have study Brownian motion and other stochastic processes trying to understand how they behave and which are they main properties. In this section, we will develop the concept of *Regulated Brownian Motion* (RBM) that is a stochastic

process build from an Brownian motion and a set of constraints that we will impose to ensure certain desired behavior. In particular, we will be looking at two types of constraints (i) we would like to impose that the process never goes below 0 (one-sided regulator), and (ii) we would like that the process evolves only inside the interval $[0, b]$ for some $b > 0$ (two-sided regulator).

A.4.1 One-Sided Regulator

Let $x \in C[0, \infty)$ be a continuous function on $[0, \infty)$. Suppose that we are interested in constructing a new function $y \in C[0, \infty)$ such that y is as similar as x as possible but y satisfies the condition $y(t) \geq 0, \forall t \in [0, \infty)$. In order to do this, we introduce two mappings $\phi, \psi : C \rightarrow C$ defined as follows

$$\psi_t(x) = \sup_{0 \leq s \leq t} \{x^-(s)\}, \quad (\text{A.14})$$

$$\phi_t(x) = x(t) + \psi_t(x), \quad (\text{A.15})$$

where $z^- = \max(0, -x)$. The pair (ψ, ϕ) is called *one-sided regulator* with lower barrier at zero. The function y that we were looking for is in fact given by $y = \phi(x)$, and the next proposition makes this fact explicit.

Proposition 23 *Suppose that $x \in C$ and $x(0) \geq 0$. Then $\psi(x)$ is the unique function l such that:*

1. l is continuous and increasing with $l(0) = 0$,
2. $y(t) = x(t) + l(t) \geq 0$ for all $t \geq 0$, and
3. l increases only when $y = 0$.

Figure A.3 plots the behavior of the one-sided regulator when applied to an arbitrary function $x(t)$. We might think of $\psi(t) = l(t)$ (the step function in the figure) as the cumulative amount of *control* used by an observer of the sample path of x up to time t . The observer wants to increase l fast enough to keep $y = x + l$ positive but using as little control as possible.

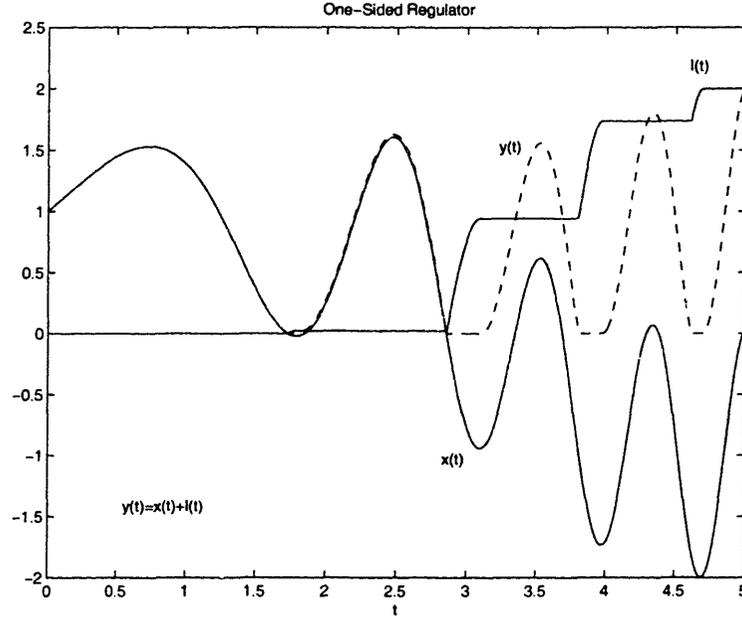


Figure A.3: The one-sided regulator.

A.4.2 The Two-Sided Regulator

Let us now consider the case when we wish to keep a process within the interval $[0, b]$ for some $b > 0$. We can view the previous case (one-sided regulator) as a special case with $b = \infty$.

Let $x \in C[0, \infty)$ be a continuous function such that $x(0) \in [0, b]$. If we mimic the approach that we took in the case of the one-sided regulator, we are interested in a pair of function (l, u) such that

1. l and u are continuous, increasing and $l(0) = u(0) = 0$,
2. $y(t) = x(t) + l(t) - u(t) \in [0, b]$ for all $t \geq 0$, and
3. l and u increase only when $y = 0$ and $y = b$ respectively.

Using the same reasoning than before we have that

$$l(t) = \psi_t(x - u) = \sup_{0 \leq s \leq t} (x(s) - u(s))^- . \quad (\text{A.16})$$

Similarly, we can compute u as follows:

$$u(t) = \psi_t(b - x - l) = \sup_{0 \leq s \leq t} (b - x(s) - l(s))^- . \quad (\text{A.17})$$

Proposition 24 *For each $x \in C$ with $x(0) \in [0, b]$, there is a unique pair of continuous functions (l, u) satisfying (A.16) and (A.17), and this same pair uniquely satisfies points 1, 2, and 3 above.*

Let consider now a general (μ, σ) Brownian motion and

$$W_t = \frac{1}{\sigma}(X_t - X_0 - \mu t),$$

a standard Brownian motion defined by X . Let define L and U as the two-sided regulator for X and $Z = X + L - U$ as the regulated Brownian motion. Let consider a function f twice continuously differentiable. Then, applying Itô's lemma, we have

Proposition 25 *$f(Z)$ is an Itô process with differential*

$$df(Z) = \sigma f'(Z)dW + [\Gamma f(Z)dt + f'(0)dL - f'(b)dU],$$

where Γ is a differential operator such that $\Gamma f \equiv \frac{1}{2}\sigma^2 f'' + \mu f'$.

The proof of the proposition is based on a direct application of Itô's lemma and the fact that L and U increase only when Z reaches one of the boundaries. Thus,

$$\int_0^t f'(Z)dL = \int_0^t f'(0)dL = f'(0)L_t.$$

A.5 Queueing Models

In this section we present the applications that Brownian motions have on Queueing Theory. We start by motivating the use of diffusion models to represent queues, part of this analysis and notation are based on the Newell (1971). After this introduction, we will present some concrete examples and formulations.

A.5.1 Introduction

Let consider a generic queue, and let $Q(t)$ be the queue size at time t . If $Q(t)$ is large enough, then after a *short* period of time we do not expect that queue size has changed by more than a few customers, or in other words, we do not expect to observe an empty system. Thus, we might try to analyze locally the behavior of $Q(t)$ without worrying about the border condition $Q \geq 0$. Let define $f(x, t)$ as the density of $Q(t)$

at x . Of course, for a real system $Q(t)$ is always integer, however, we will assume here that $Q(t)$ is a continuous variable for obvious reasons.

Let us now introduce, in some sense, the dynamics of the system. That is, let define $f(x, t + \tau | x_0, t)$ as the probability density for $Q(t + \tau)$ at x given that $Q(t) = x_0$. This conditional distribution completely captures the evolution of the system since

$$f(x, t + \tau) = \int f(x, t + \tau | x_0, t) f(x_0, t) dx_0. \quad (\text{A.18})$$

The previous relations captures the essence of the dynamics of the system and it will be the base to obtain a differential equation that characterizes the queue. In order to do this, we need first to work a little more on the conditional distribution $f(x, t + \tau | x_0, t)$. Let $A(t)$ be the cumulative number of arrivals to the queue up to time t and let $D(t)$ be the cumulative number of departures up to time t . Then we have

$$Q(t + \tau) - Q(t) = [A(t + \tau) - A(t)] - [D(t + \tau) - D(t)].$$

We can argue that for τ large enough many arrivals and departures will take place between t and $t + \tau$, then $Q(t + \tau) - Q(t)$ is approximately normally distributed almost independently of the arrival or departure processes.

It is important to mention here that in one side we would like to have τ small so that the border conditions can be omitted and on the other side we would like to have τ large so that we can approximate the conditional distribution by a normal. This situation is not necessarily infeasible. The fact is that if we choose a sufficiently coarse scale to measure $Q(t)$ then it should be possible to choose τ such that the queue does not change much during this time. Thus, what we really required is that τ is negligible compared with the time required to change the queue by a significant fraction of its normal values (a heavy traffic condition). On the other hand, we need τ to be sufficiently large so that the change in the queue is large compared with the integer scale of counting single customers.

Let $\lambda(t)$ and $\mu(t)$ the arrival and service rates. Then, we can use the following approximation:

$$E[Q(t + \tau) - Q(t)] \approx [\lambda(t) - \mu(t)]\tau.$$

Similarly, we can approximate the variance of $Q(t + \tau) - Q(t)$ by some function $\delta(t)\tau$. Thus, assuming normality we can find the distribution of $Q(t + \tau)$ given $Q(t) = x_0$

by

$$f(x, t + \tau | x_0, t) \approx \frac{1}{\sqrt{2\pi\delta(t)\tau}} e^{-\frac{(x-x_0-(\lambda(t)-\mu(t))\tau)^2}{2\delta(t)\tau}}.$$

Given this particular form for the conditional distribution of the queue, we can try to solve (A.18). In order to do that we first expand $f(x_0, t)$ in a power series. Thus,

$$f(x_0, t) = f(x, t) + (x_0 - x) \frac{\partial f(x, t)}{\partial x} + \frac{(x_0 - x)^2}{2} \frac{\partial^2 f(x, t)}{\partial x^2} + \dots$$

If we substitute this relation in (A.18) and integrate we get:

$$\frac{f(x, t + \tau) - f(x, t)}{\tau} = -[\lambda(t) - \mu(t)] \frac{\partial f(x, t)}{\partial x} + [\delta(t) + (\lambda(t) - \mu(t))^2 \tau] \frac{\partial^2 f(x, t)}{2\partial x^2} + \dots$$

Thus, if we take limit as $\tau \rightarrow 0$, we obtain the following differential equation:

$$\frac{\partial f(x, t)}{\partial t} = -[\lambda(t) - \mu(t)] \frac{\partial f(x, t)}{\partial x} + \frac{\delta(t)}{2} \frac{\partial^2 f(x, t)}{\partial x^2}. \quad (\text{A.19})$$

This is exactly the *Kolmogorov's forward equation* (or diffusion equation) that characterizes Brownian motions. Thus, for heavily loaded system, we can approximate the behavior of the queue when it far from 0 as a BM. The drift is the difference between the arrival and service rates, and the variance is the difference between the variance of the arrival and service process.

Before moving to some concrete models of queueing system, we would like to introduce a scaling transformation for (A.19). This transformation applies in the stationary case, that is when the arrival and service process are not time dependent. In this case (A.19) becomes

$$\frac{\partial f(x, t)}{\partial t} = -[\lambda - \mu] \frac{\partial f(x, t)}{\partial x} + \frac{\delta}{2} \frac{\partial^2 f(x, t)}{\partial x^2}.$$

Let introduce the following transformation of units

$$\bar{x} = \frac{x}{L}, \quad \bar{t} = \frac{t}{T}.$$

That is, queue length is measure in units of L and time in units of T . With this

transformation, the diffusion equation becomes:

$$\frac{\partial f(\bar{x}L, \bar{t}T)}{T\partial\bar{t}} = -[\lambda - \mu] \frac{\partial f(\bar{x}L, \bar{t}T)}{L\partial\bar{x}} + \frac{\delta}{2L^2} \frac{\partial^2 f(\bar{x}L, \bar{t}T)}{\partial\bar{x}^2}.$$

If we write $\bar{f}(\bar{x}, \bar{t}) \equiv f(\bar{x}L, \bar{t}T)$ which is the distribution of $\bar{Q} = \frac{Q}{L}$, then

$$\frac{\partial\bar{f}(\bar{x}, \bar{t})}{\partial\bar{t}} = -\frac{[\lambda - \mu]T}{L} \frac{\partial\bar{f}(\bar{x}, \bar{t})}{\partial\bar{x}} + \frac{\delta T}{2L^2} \frac{\partial^2\bar{f}(\bar{x}, \bar{t})}{\partial\bar{x}^2}.$$

Finally, if we choose L and T such that

$$\frac{(\mu - \lambda)T}{L} = 1, \quad \frac{\delta T}{2L^2} = 1,$$

we get

$$L = \frac{\delta}{\mu - \lambda}, \quad T = \frac{\delta}{(\mu - \lambda)^2},$$

and the diffusion equation is

$$\frac{\partial\bar{f}(\bar{x}, \bar{t})}{\partial\bar{t}} = \frac{\partial\bar{f}(\bar{x}, \bar{t})}{\partial\bar{x}} + \frac{1}{2} \frac{\partial^2\bar{f}(\bar{x}, \bar{t})}{\partial\bar{x}^2}.$$

A somewhat different approach can also be used to justify the usage of Brownian motions as a good approximation for modelling queue lengths. This approach is commonly known as the *heavy traffic* approximation for reasons that will become clear shortly. The details of this approximation are explained in the next section by mean of an example.

A.5.2 Single-Stage Infinite Capacity Queueing Model

Let us now look at the simplest queueing model, that is a single-stage system with infinite capacity. In order to model this system, we make use of two stochastic processes $A = \{A_t : t \geq 0\}$ (arrival process) and $D = \{D_t : t \geq 0\}$ (departure process). The value A_t represents the cumulative number of arrivals (or input) up to time t and B_t represents the cumulative *potential* number of departures (or output) up to time t . That is, B_t is the cumulative number of departures if we assume that the system is never empty. For simplicity we assume that $A_0 = B_0 = 0$. Since it is possible for the system to get empty eventually, B_t does not properly model the output process. In order to model the real output process, we need to consider an

additional process $L = \{L_t : t \geq 0\}$, where L_t represents the amount of output lost up to time t because of system emptiness. Thus, the actual departure process is $B - L$.

We define the auxiliary process

$$X_t = X_0 + A_t - B_t,$$

where X_0 is the initial queue size. By the discussion of the previous subsection we can approximate the behavior of X by a Brownian motion. We can now define the queue size (or inventory) at time t as follows:

$$Z_t = X_0 + A_t - (B_t - L_t) = X_t + L_t.$$

We can notice, at this point, that L_t is the one-sided regulator that keeps $Z_t \geq 0$, thus

$$L_t = \sup_{0 \leq s \leq t} X_s^-.$$

Now in order to find the distribution of Z , we use the following identity:

$$\begin{aligned} Z_t &= X_t + \sup_{0 \leq s \leq t} X_s^- \\ &= \sup_{0 \leq s \leq t} (X_t - X_s) \\ &= \sup_{0 \leq s \leq t} \hat{X}_s, \end{aligned}$$

where $\hat{X}_s = X_t - X_s$, $0 \leq s \leq t$. But since X is a BM, we have that \hat{X} is also a BM with the same distribution of X . Thus,

$$Z_t \sim M_t \equiv \sup_{0 \leq s \leq t} X_s.$$

In section (A.2.3) we have already determined the distribution of M_t , which is the inverse Gaussian distribution. Therefore, if X is a (μ, σ) Brownian motion we have:

$$P(Z_t \leq x) = \Phi\left(\frac{x - \mu t}{\sigma\sqrt{t}}\right) - e^{\frac{2\mu x}{\sigma^2}} \Phi\left(\frac{-x - \mu t}{\sigma\sqrt{t}}\right).$$

Finally, we can get the limiting distribution of Z by taking limit as $t \rightarrow \infty$. Let Z_∞ the steady state value of the queue, then

$$P(Z_\infty \leq x) = \begin{cases} 1 - e^{-\frac{2\mu x}{\sigma^2}} & \text{if } \mu < 0; \\ 0 & \text{if } \mu \geq 0. \end{cases}$$

That is, under the natural condition $\mu < 0$ (i.e., the service rate is bigger than the arrival rate) the limiting distribution of the one-sided regulated Brownian motion is exponential.

The previous analysis was based on the assumption that we can approximate $X_t = X_0 + A_t - B_t$ by a (μ, σ) Brownian motion. This assumption is consistent the arguments that we develop in the previous subsection (A.5.1), where we argue that the BM approximation is appropriate for systems that are heavy loaded. In what follows, we will give additional support to the use of Brownian motion under heavy traffic condition.

We start by redefining our primitive data which are two renewal processes (arrival and service or departure):

$$\begin{aligned} \{A_t : t \geq 0\} & \text{ rate } \lambda, \text{ variance } \alpha^2, \\ \{D_t : t \geq 0\} & \text{ rate } \mu, \text{ variance } \beta^2. \end{aligned}$$

Notice that we now use μ for the service rate, which is standard notation in queueing models. As before, A_t is the cumulative numbers of arrivals up to time t , while D_t is the cumulative number of service completions if the system were always busy. Let Z_t be the queue length at time t . We can define the cumulative busy-time process in $[0, t]$ (i.e., the total amount of time that the server was serving) as follows:

$$B(t) = \int_0^t 1_{\{Z(s) > 0\}} ds,$$

and then the inventory process is given by

$$Z(t) = A(t) - S(B(t)).$$

In order to simplify the analysis, it is convenient to redefine the arrival and departure process defining their centered version. We use the same notation to avoid excessive notation.

$$A(t) \leftarrow A(t) - \lambda t \quad , \quad D(t) \leftarrow D(t) - \mu t.$$

With this new definitions we have that

$$Z(t) = \underbrace{A(t) - D(B(t))}_{X(t)} + (\lambda - \mu)t + \underbrace{\mu(t - B(t))}_{L(t)},$$

The next step is related to a *functional central limit* result. We have already presented Donsker's theorem that asserts that

$$\frac{S_{\lfloor nt \rfloor} - (nt)m}{\nu\sqrt{n}} \Rightarrow W,$$

where $S_n = \sum_{i=1}^n \xi_i$ and $\{\xi_i\}$ is a sequence of i.i.d. random variable with mean m and variance ν^2 . Thus, we rescale our processes as follows,

$$Z_n(t) = \frac{Z(nt)}{\sqrt{n}}, \quad A_n(t) = \frac{A(nt)}{\sqrt{n}}, \quad D_n(t) = \frac{D(nt)}{\sqrt{n}}, \quad X_n(t) = \frac{X(nt)}{\sqrt{n}}, \quad L_n(t) = \frac{L(nt)}{\sqrt{n}},$$

and we also set

$$B_n(t) = \frac{B(nt)}{n}.$$

We use a fluid rather than a Brownian scaling for B for reasons that will become clear shortly.

Given this transformation, we have that

$$X_n(t) = A_n(t) - D_n(B_n(t)) + \sqrt{n}(\lambda - \mu)t.$$

That is, X_n has almost the required form to apply Donsker's theorem. The only problem is that the argument of D_n is the busy process. In order to solve this problem, we apply the following result:

Proposition 26 (Random Time-Change Theorem)

Suppose that $\{T_n\}$ is a sequence of random time changes (i.e., T_n is a process with increasing, real-value sample paths) converging to T . If $V_n \Rightarrow V$, then $V_n \circ T_n \Rightarrow V \circ T$.

$B_n(t)$ is clearly a random time change. However, it is not absolutely clear if the sequence $\{B_n(t)\}$ is weakly convergent to some limit. If we use classical queueing results, we can argue that in order to satisfy demand and keep inventory under control, the server must be working on average a fraction $\rho = \frac{\lambda}{\mu}$ of the time. Thus, we use

the following approximation

$$B_n(t) \Rightarrow \rho t.$$

We are now in condition to apply Donsker's Theorem to $X_n(t)$ and argue that

$$X_n(t) \Rightarrow X^*,$$

where X^* is a (θ, σ) Brownian motion with drift $\theta = -\mu\sqrt{n}(1 - \rho)$ ³ and variance $\sigma^2 = \lambda(\nu_a^2 + \nu_s^2)$ (where ν_a^2 and ν_s^2 are the scv for the arrival and service time). Finally, and by construction $Z_n \Rightarrow Z^*$ a (θ, σ) reflected Brownian motion.

We conclude this section with the following more general result:

Proposition 27 *Consider a queueing system with J arrival streams and K servers. Suppose that the j^{th} arrival stream has parameters λ_j and ν_{aj}^2 and the k^{th} server has parameter μ_k and ν_{sk}^2 . Then the total number of customers in the system is approximated by a reflected Brownian motion with drift μ and variance σ^2 given by*

$$\mu = \sqrt{n} \left(\sum_{j=1}^J \lambda_j - \sum_{k=1}^K \mu_k \right), \quad \sigma^2 = \sum_{j=1}^J \lambda_j \nu_{aj}^2 + \rho \sum_{k=1}^K \mu_k \nu_{sk}^2,$$

where $\rho = \frac{\sum_j \lambda_j}{\sum_k \mu_k}$.

A.5.3 Single-Stage Finite Capacity Queueing Model

If the queueing system has a maximum capacity C then the Brownian motion approximations has to consider the two-sided regulator. In this case, we say that $Z(t)$ (the queue size) behaves like a two-sided regulated Brownian motion in the interval $[0, C]$. Thus, it can be written as

$$Z_t = X_t + L_t - U_t$$

where X_t is a uncontrolled (μ, σ^2) Brownian motion and L_t and U_t are the lower and upper controls respectively. This process are the minimum amount of control

³Since n is large to ensure the convergence of $X_n(t)$, we need $(1 - \rho)$ to be order $1/\sqrt{n}$ to ensure a finite value of the drift *theta*. This condition is known as the *heavy traffic* condition and also gives the name to these type of approximations.

necessary to use to keep Z_t within $[0, C]$. If we call

$$\Gamma := \frac{\sigma^2}{2} \frac{d^2}{dx^2} + \mu \frac{d}{dx}$$

the infinitesimal generator of X_t then a major result is

Proposition 28 *If $f : [0, C] \rightarrow \mathcal{R}$ is a twice continuously differentiable function then $f(Z(t))$ is an Itô process with differential*

$$df(Z) = \sigma f'(Z)dW + \Gamma f(Z) dt + f'(0) dL - f'(C) dU.$$

The proof can be found in chapter 5 in Harrison (1985). This result allows us to use the whole set of tools of Stochastic Calculus. In particular, the main result (that we present without proof) for our queueing system is the following.

Proposition 29 *The steady state distribution of $Z(t)$ is given by the following truncated exponential distribution*

$$p(z) = \frac{\nu e^{\nu z}}{e^{\nu C} - 1} \text{ for all } 0 \leq z \leq C,$$

where $\nu = \frac{2\mu}{\sigma^2}$. In addition, the average amount of control used (or the local time on the boundaries) is

$$\lim_{t \rightarrow \infty} \frac{E[R(t)]}{t} = \frac{\nu}{e^{\nu C} - 1} \quad \text{and} \quad \lim_{t \rightarrow \infty} \frac{E[R(t)]}{t} = \frac{\nu}{1 - e^{-\nu C}}.$$

Bibliography

- [1] Agrawal, N., S. Nahmias, A. Tsay. 1999. Modeling supply chain contracts: A review. Chapter 10 in Tayur S., M. Magazine, R. Ganeshan (Eds.), *Quantitative models for supply chain management*. Kluwer Academic Publishers.
- [2] Anupindi, R., Y. Bassok. 1999. Supply contracts with quantity commitments and stochastic demand. Chapter 7 in Tayur S., M. Magazine, R. Ganeshan (Eds.), *Quantitative models for supply chain management*. Kluwer Academic Publishers.
- [3] Beneš, V., Shepp L.A., Witsenhausen H.S. (1980). Some Solvable Stochastic Control Problems. *Stochastics* 4, 39-83.
- [4] Billingsley, P. (1999). *Convergence of Probability Measures*. Second Edition, Wiley.
- [5] Brynjolfsson, E., Smith M. (1999). Frictionless Commerce? A Comparison of Internet and Conventional Retailers. *Mgmt. Sci.* 46, No. 4, 563-585.
- [6] Buzacott, J. A., S. Price, J. G. Shanthikumar. 1992. Service levels in multi-stage MRP and base stock controlled production systems. In Fandel, G., T. Gullledge, A. Jones (Eds.), *New directions for operations research in manufacturing*. Springer, Berlin, 445-463.
- [7] Buzzell, R., G. Ortmeyer. 1995. Channel partnerships streamline distribution. *Sloan Management Review* 36, 85-96.
- [8] Cachon, G. 1999. Competitive supply chain inventory management. Chapter 5 in Tayur S., M. Magazine, R. Ganeshan (Eds.), *Quantitative models for supply chain management*. Kluwer Academic Publishers.

- [9] Cachon, G., M. Lariviere. 1997. Contracting to assure supply or what did the supplier know and when did he know it? Working paper, Fuqua School of Business, Duke University.
- [10] Cachon, G., P. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Science* **45**, 936-953.
- [11] Chen, F. 1999. Decentralized supply chains subject to information delays. To appear in *Management Science*.
- [12] Chen, F., A. Federgruen, Y. Zheng. 1999. Coordination mechanisms for a distribution system with one supplier and multiple retailers. Graduate School of Business, Columbia University, New York, NY.
- [13] Clark, A., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Science* **6**, 475-490.
- [14] Dyer, J., W. Ouchi. 1993. Japanese-style partnerships: Giving companies a competitive edge. *Sloan Management Review* **34**, 51-63.
- [15] Dyer J., D. S. Cho, W. Chu. 1998. Strategic supplier segmentation: The next "best practice" in supply chain management. *California Management Review* **40**, 57-77.
- [16] Ephremides A., Varaiya P., Walrand J. (1980). A simple Dynamic Routing Problem. *IEEE Trans. Auto. Control* **AC-25**, 690-693.
- [17] Federgruen, A., Z. Katalan. 1996. The stochastic economic lot scheduling problem: Cyclical base stock policies with idle times. *Management Science* **42**, 783-796.
- [18] Federgruen, A., P. Zipkin. 1986. An inventory model with limited production capacity and uncertain demands I: The average-cost criterion. *Math. Operations Research* **11**, 193-207.
- [19] Gavirneni, S., R. Kapuscinski, S. Tayur. 1996. Value of information in capacitated supply chains. Working paper, GSIA, Carnegie Mellon University, Pittsburgh, PA.
- [20] Gelfand, I.M., Fomin S.V. (1963). *Calculus of Variations*. Prentice-Hall, Inc.

- [21] Ghoneim, H. Stidham S. (1985). Control of Arrivals to Two Queues in Series. *Euro. J. Opns. Res.* **21**, 399-409.
- [22] Glasserman P., S. Tayur. 1995. Sensitivity analysis for base stock levels in multi-echelon production-inventory systems. *Management Science* **41**, 263-281.
- [23] Ha, A. (1997). Inventory Rationing in a Make-to-Stock Production System with Several Demand Classes and Lost Sales. *Magment Sci.* **43**, No. 8, 1093-1103.
- [24] Hadley, G., Whitin T.M. (1963). *Analysis of Inventory Systems*. Prentice-Hall, Inc.
- [25] Harrison, J.M., Taylor A.M. (1978). Optimal Control of Brownian Storage Systems. *Stochastic Processes and teir Applications* **6**, 179-194.
- [26] Harrison, J.M., Taksar M. (1983). Instantaneous Control of Brownian Motion. *Math. Opers. Res.* **8**, 439-453.
- [27] Harrison, J.M. (1985) *Brownian Motion and Stochastic Flow Systems*. John Wiley & Sons, New York.
- [28] Harrison, J.M., Landau H.J., Shepp L.A. (1985). The stationary Distribution of Reflected Brownian Motion in a Planar Region. *The Annals of Probability*, vol. **13**, No. 3, 744-757.
- [29] Harrison M., Williams R. (1987). Multidimensional Reflected Brownian Motions Having Exponential Stationary Distributions. *The Annals of Probability*, Vol. **15**, No. 1, 115-137.
- [30] Harrison, J. M. 1988. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic differential systems, stochastic control theory and applications*, W. Fleming, P. L. Lions (eds.). IMA Volume **10**, Springer-Verlag, New York, 147-186.
- [31] Johansen, S.G., Stidham S. (1980). Control of Arrivals to a Stochastic Input-Output System. *Adv. Appl. Prob.* **12**, 972-999.
- [32] John, F. (1982). *Partial Differential Equations*. Fourth Edition, Springer.

- [33] Krichagina, E.V., Lou S.X., Taksar M.I. (1994). Double Band Policy for Stochastic Manufacturing Systems in Heavy Traffic. *Math. Opns. Res.* **19**, No. 3, 560-596.
- [34] Kumar, S., Muthuraman M. (2000). A Numerical Method for Solving Singular Brownian Control Problems. Working paper, Stanford University.
- [35] Kushner, H.J. (1977). *Probability Methods for Approximations in Stochastic Control and for Elliptic Equations*. Academic Press, New York.
- [36] Kushner, H.J., Dupuis P.G. (1992). *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer, New York.
- [37] Lariviere, M. 1999. Supply chain contracting and coordination with stochastic demand. Chapter 8 in Tayur S., M. Magazine, R. Ganeshan (Eds.), *Quantitative models for supply chain management*. Kluwer Academic Publishers.
- [38] Lee, Y. J., P. Zipkin. 1992. Tandem queues with planned inventories. *Operations Research* **40**, 936-947.
- [39] Li, L. (1992). The Role of Inventory in Delivery-Time Competition. *Magment. Sci.* **38**, No. 2, 182-197.
- [40] Markowitz, D. M., M. I. Reiman, L. M. Wein. 1999. The stochastic economic lot scheduling problem: Heavy traffic analysis of dynamic cyclic policies. To appear in *Operations Research*.
- [41] Menaldi, J.L., Robin M. (1984). Some Singular Problem with Long term Average Criterion. Proceedings of the Eleventh IFIP Conference on System Modelling and Optimization, Copenhagen, Denmark. In *Lectures Notes in Control and Information Sciences*, Ed. P. Thopt-Christensen, **59**, 424-432.
- [42] Merton, R. 1990. *Continuous-Time Finance*. Basil Blackwell Inc.
- [43] Miller, B. (1969). A Queueing Reward System with Several Customer Classes. *Magment. Sci.* **16**, No. 3, 234-245.
- [44] Morse, P. 1958. *Queues, inventories and maintenance*. Wiley, New York.
- [45] Naor, P. (1969). The regulation of Queues Size by Levying Tolls. *Econometrica* **37**, 15-24.

- [46] Osborne, M. 1964. Brownian Motion in the Stock Market. Chapter 4 in Cootner P. (Ed.), *The Random Character of Stock Market Prices*. MIT.
- [47] Plambeck, E. L., S. A. Zenios. 1999. Incentive efficient control of a make-to-stock production system. Working Paper, Graduate School of Business, Stanford University.
- [48] Rubio, R., L. M. Wein. 1996. Setting base stock levels using product-form queueing networks. *Management Science* **42**, 259-268.
- [49] Samuelson, P. (1965). Rational Theory of Warrant Pricing. *Industrial Management Review*, 6, 13-31.
- [50] Sato, K., Tanaka H. (1962). Local Times on the Boundary for Multi-Dimensional Reflecting Diffusion.
- [51] Stidham, S. (1978). Socially and Individually Optimal Control of Arrivals to a GI/M/1 Queue. *Mgmt. Sci.* **24**, No. 15, 1598-1610.
- [52] Stidham, S. (1985). Optimal Control of Admission to a Queueing System. *IEEE Trans. Automat. Contr.* **AC-30**, No. 8, 705-713.
- [53] Stidham, S. (1988). Scheduling, Routing, and Flow Control of Stochastic Network. In *Stochastic Differential Systems, Stochastic Control Theory and Applications*. IMA Volume 10, W. Fleming and P.L. Lions (eds.). Springer-Verlag, New York, 529-554.
- [54] Taksar, M. (1985). Average Optimal Singular Control and a Related Stopping Problem. *Math. Oper. Res.* **10** 63-81.
- [55] Tirole, J. 1997. *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.
- [56] Topkis, D. M. 1998. *Supermodularity and complementarity*. Princeton University Press, Princeton, New Jersey.
- [57] Tsay, A. A., W. S. Lovejoy. 1999. Quantity flexibility contracts and supply chain performance. To appear in *Manufacturing & Service Operations Management*.
- [58] Vajda, S. 1989. *Fibonacci & Lucas numbers, and the golden section: Theory and applications*. Ellis Horwood Limited, England.

- [59] Veatch, M., Wein L. (1992). Monotone Control of Queueing Networks. *Queueing Systems* **12**, 391-408.
- [60] Wein, L.M. (1990). Optimal Control of a Two-Station Brownian Network. *Math. Opns. Res.* **15**, 215-242.
- [61] Wein, L.M. (1992). Dynamic Scheduling of a Multiclass Make-To-Stock Queue. *Opns. Res.* **40**, No. 4, 724-735.
- [62] Wein, L. M. 1992. Dynamic scheduling of a multiclass make-to-stock queue. *Operations Research* **40**, 724-735.
- [63] Yechiali, U. (1971). On Optimal Balking Rules and Toll Charges in the GI/M/1 Queueing Process. *Opns. Res.* **19**, 349-370.
- [64] Yechiali, U. (1972). Customers' Optimal Joining Rules for the GI/M/s Queue. *Mgmt. Sci.* **18**, No. 7, 434-443.
- [65] Zheng, Y., P. Zipkin. 1990. A queueing model to analyze the value of centralized inventory information. *Operations Research* **38**, 296-307.

THESIS PROCESSING SLIP

FIXED FIELD: ill. _____ name _____

index

biblio

► COPIES: Archives Aero Dewey Barker Hum
Lindgren Music Rotch Science Sche-Plough

TITLE VARIES: ► _____

NAME VARIES: ► Antonio

IMPRINT: (COPYRIGHT) _____

► COLLATION: _____

► ADD: DEGREE: _____ ► DEPT.: _____

► ADD: DEGREE: _____ ► DEPT.: _____

SUPERVISORS: _____

NOTES:

cat'r:

date:

► DEPT: mat

page:	<u>5212</u>
-------	-------------

► YEAR: 2001 ► DEGREE: Ph.D.

► NAME: CALDENTEY, Rene' A.